

Notes for Computational Linear Algebra

Math and Programming
at the Scale of Life

Author

Jessy Grizzle, Director

Contributors

Maani Ghaffari, Kira Biener, Tribhi Kathuria, Madhav Achar,
Fangtong (Miley) Liu, Shaoxiong Yao, Eva Mungai, Bruce JK Huang,
Grant A. Gibson, Oluwami Dosunmu-Ogunbi, Lu Gan, Ray Zhang

Inspiration

Chad Jenkins, Associate Director of Undergraduate Programs



Cover design by Dan Newman, Head of Communications, Michigan Robotics

© 2022, 2021, 2020 Jessy Grizzle, former Director of Robotics, University of Michigan
JERRY W. AND CAROL L. LEVIN PROFESSOR OF ENGINEERING
ELMER G. GILBERT DISTINGUISHED UNIVERSITY PROFESSOR
PROFESSOR OF EECS AND ME

Composed between April 2020 and August 2020 for use in ROB 101, Computational Linear Algebra, a first-semester first-year course. Revised Fall 2020 when put to the test with our Pilot Class. Revised Summer and Fall 2021. Extensive proof reading provided by Muhsinun Chowdhury in Fall 2022.

First release, August 31, 2020. Second release, August 30, 2021. Third release August 29, 2022.

Contents

Preface	9
Philosophy of the Course	11
List of Algorithms and Methods or Things You Need Code to Do Well	13
1 Introduction to Systems of Linear Equations	15
1.1 For Review on Your Own: You Have Done Algebra Before	16
1.2 Linear Systems of Equations: What can happen?	19
1.3 Naming Variables	21
1.4 A 3 x 3 Example	21
1.5 Greek Alphabet	22
1.6 (Optional Read) Where is Computational Linear Algebra Used?	22
1.7 Looking Ahead	23
2 Vectors, Matrices, and Determinants	25
2.1 Scalars and Arrays	26
2.2 Row Vectors and Column Vectors	27
2.3 Remark on Brackets	28
2.4 Matrices, Rectangular and Square, and the Matrix Diagonal	29
2.5 Expressing a System of Linear Equations in terms of Vectors and Matrices	30
2.6 The Matrix Determinant	33
2.6.1 First Contact with the Matrix Determinant	34
2.6.2 Examples of Using the Determinant	34
2.7 (Optional Read): Other Facts on the Matrix Determinant Covered in “Standard” Courses	35
2.8 (Optional Read): A Few “Practical” Examples Using Linear Algebra	36
2.9 (Optional Read) Yet Another Determinant Example	38
2.10 Looking Ahead	39
3 Triangular Systems of Equations: Forward and Back Substitution	41
3.1 Background	42
3.2 A Warm-up with Diagonal Systems of Linear Equations	42
3.3 Lower Triangular Systems of Linear Equations	43
3.4 Determinant of a Lower Triangular Matrix	45
3.5 Lower Triangular Systems and Forward Substitution	45
3.6 Upper Triangular Systems, Upper Triangular Matrices, Determinants, and Back Substitution	46
3.7 General Cases	47
3.8 A Simple Trick with Systems of Equations: Re-arranging their Order	50
3.9 Looking Ahead	50
4 Matrix Multiplication	51
4.1 Multiplying a Row Vector by a Column Vector	52
4.2 Examples of Row and Column Partitions	53
4.3 General Case of Partitions	54
4.4 Standard Matrix Multiplication: It’s All About Rows and Columns	54
4.4.1 Examples	55

4.4.2	Optional Read: General case: what is happening inside Julia	57
4.5	Multiplication by Summing over Columns and Rows	57
4.6	(Optional Read:) Why the Second Method of Matrix Multiplication Works	59
4.7	Permutation Matrices: The Matrix View of Swapping the Order of Equations	59
4.8	(Optional Read): Useful Fact on Matrix Multiplication	61
4.9	(Optional Read): Some Geometric Aspects of Matrix-Vector Multiplication	62
4.10	Looking Ahead	64
5	LU (Lower-Upper) Factorization	65
5.1	Recalling Forward and Back Substitution	66
5.2	Recalling Matrix Multiplication in the Form of Columns Times Rows	66
5.3	The Secret Sauce of LU (Lower-Upper) Factorization: Peeling the Onion	67
5.4	(Optional Read:) Peeling the Onion in Pictures	71
5.5	LU (Lower-Upper) Factorization (without row permutations)	72
5.6	LU Factorization for Solving Linear Equations	78
5.7	(Optional Read): Toward LU Factorization with Row Permutations	81
5.8	(Optional Read): An Algorithm for LU Factorization with Row Permutations	85
5.9	Solving $Ax=b$ via LU with Row Permutations	91
5.10	(Optional Read): Using the LU command native to Julia	91
5.11	(Optional Read): Large Scale Example: Computing the Forces in a Truss Bridge	96
5.12	(Optional Read): An Algorithm for Rectangular LU Factorization with Row Permutations	100
5.13	Looking Ahead	101
6	Determinant of a Matrix Product, Matrix Inverses, Matrix Transposes, and Permutation Matrices	103
6.1	A very Useful Fact Regarding the Matrix Determinant	104
6.2	Identity Matrix and Matrix Inverse	105
6.3	Utility of the Matrix Inverse and its Computation	107
6.4	Matrix Transpose and Symmetric Matrices	110
6.5	Revisiting Permutation Matrices	112
6.6	Matrix Determinants, Matrix Inverses, and the Matrix Transpose	114
6.7	Looking Ahead	115
7	The Vector Space \mathbb{R}^n: Part 1	117
7.1	Motivation	118
7.2	Vectors in \mathbb{R}^n and Some Basic Algebra	118
7.3	The Notion of Linear Combinations	122
7.3.1	Linear combinations through the lens of $Ax=b$	123
7.3.2	Linear Combinations in \mathbb{R}^n	124
7.4	Existence of Solutions to $Ax=b$	126
7.5	Linear Independence of a Set of Vectors	127
7.5.1	Preamble	127
7.5.2	Linear Independence through the Lens of $Ax=0$	127
7.5.3	Linear Independence in \mathbb{R}^n (and why theory matters)	128
7.5.4	A Pro Tip for Checking Linear Independence	130
7.5.5	(Optional Read): Why the Pro Tip Works	132
7.6	LDLT and the Number of Linearly Independent Vectors in a Set	134
7.7	Attractive Test for Linear Combinations	140
7.8	Existence and Uniqueness of Solutions to $Ax=b$	141
7.9	When are the Solutions to $Ax = b$ the same as the Solutions to $M \cdot Ax = M \cdot b$?	143
7.10	(Optional Read): Why LDLT and not Simply LU?	143
7.11	(Optional Read): LU Factorization for Symmetric Matrices	144
7.12	Looking Ahead	148

8	Euclidean Norm, Least Squared Error Solutions to Linear Equations, and Linear Regression	149
8.1	Norm or “Length” of a Vector	150
8.2	Least Squared Error Solutions to Linear Equations	151
8.3	Linear Regression or Fitting Functions to Data	154
8.4	(Optional Read): How to Derive the Main Regression Formula	160
8.4.1	Completing the Square for the Quadratic Formula:	160
8.4.2	Completing the Square for Least Squares Solutions to Systems of Linear Equations:	161
8.5	Looking Ahead	162
9	The Vector Space \mathbb{R}^n: Part 2	163
9.1	Motivation	164
9.2	\mathbb{R}^n as a Vector Space	164
9.3	Subspaces	166
9.3.1	Subspaces of \mathbb{R}^n	166
9.3.2	(Optional Read:) A Broader View of Subspaces	168
9.4	Three Sources of Subspaces in \mathbb{R}^n : Matrix Null Space, Span of a set of Vectors, and Column Span of a Matrix	169
9.5	Dot Product and Orthonormal Vectors	178
9.6	Orthogonal Matrices or Why Orthonormal Vectors are Super Useful	181
9.7	Constructing Orthonormal Vectors: the Gram-Schmidt Process	182
9.7.1	Building Orthogonal Vectors	182
9.7.2	Gram-Schmidt Process (aka the Gram-Schmidt Algorithm) for Building Orthonormal Vectors	184
9.8	QR Factorization and Solutions of Linear Equations	187
9.9	Underdetermined Equations or What to do When $Ax=b$ has an Infinite Number of Solutions?	191
9.10	Steering a Mobile Robot as a Practical Example of an Underdetermined System of Linear Equations	194
9.11	(Optional Read): In the QR Factorization, Why R is Upper Triangular and How to Efficiently Obtain its Coefficients from the Gram-Schmidt Process	198
9.12	(Optional Read): Modified Gram-Schmidt Algorithm	199
9.13	(Optional Read) Source of the Definition of Orthogonal Vectors	201
9.14	Looking Ahead	202
10	The Vector Space \mathbb{R}^n: Part 3	203
10.1	Motivation	204
10.2	Basis Vectors, Coordinates, and Dimension	204
10.3	Eigenvalues and Eigenvectors	210
10.4	Range, Column Span, and Null Space	217
10.5	Rank and Nullity	223
10.6	Finding a Basis for the Null Space without Using Orthogonality	224
10.7	(Optional Read): Proofs of the Rank and Nullity Relations	225
10.8	Looking Ahead	228
11	Changing Gears: Solutions of Nonlinear Equations	229
11.1	Motivation and Simple Ideas	230
11.2	Bisection	231
11.3	The Concept of a Derivative and Its Numerical Approximation	236
11.4	Newton’s Method for Scalar Problems	239
11.5	Vector Valued Functions: Linear Approximations, Partial Derivatives, Jacobians, and the Gradient	243
11.5.1	Linear Approximation about a Point: Take 1	244
11.5.2	Partial Derivatives	245
11.5.3	The Jacobian and Linear Approximation of a Function about a Point	247
11.5.4	The Gradient and Linear Approximation of a Function about a Point	249
11.5.5	Summary on Partial Derivatives	251
11.6	Newton-Raphson for Vector Functions	252
11.7	(Optional Read): From the Gradient or Jacobian of a Function to its Linear Approximation	254
11.7.1	The Gradient	254
11.7.2	Expanding on Vector Notation	256
11.7.3	The Jacobian	258
11.7.4	Linear Approximation of Vector Valued Functions	259

11.8	Looking Ahead	260
12	Changing Gears Again: Basic Ideas of Optimization	261
12.1	Motivation and Basic Ideas	262
12.2	Contour Plots and the Gradient of the Cost	265
12.3	Gradient Descent	267
12.4	Extrinsic Calibration Using Gradient Descent	269
12.4.1	Introduction	269
12.4.2	Problem Setup and Initial Solution	270
12.5	Optimization as a Root Finding Problem: the Hessian	273
12.5.1	Second Derivatives	273
12.5.2	The Hessian is the Jacobian of the Transpose of the Gradient	274
12.5.3	Use of the Second Derivative and Hessian in Optimization	275
12.6	Local vs Global	285
12.7	Maximization as Minimizing the Negative of a Function	287
12.8	(Optional Read): Quadratic Programs: Our first Encounter with Constraints	287
12.9	(Optional Read): QP Solver in Julia	293
12.10	(Optional Read): Optimization Packages: The Sky is the Limit	295
12.11	Looking Ahead	296
13	Background for Classification and Machine Learning	297
13.1	Separating Hyperplanes	298
13.1.1	Lines in \mathbb{R}^2 as Separating Hyperplanes	298
13.1.2	Hyper Subspaces	300
13.1.3	Translations of Sets, Hyper Subspaces, and Hyperplanes	302
13.2	Signed Distance to a Hyperplane	304
13.3	Max-margin Classifier	306
13.4	Remarks on Soft Margin Classifiers	311
13.5	Orthogonal Projection	318
13.5.1	Orthogonal Projection for Subspaces	319
13.5.2	Orthogonal Projection onto Linear Varieties (translations of subspaces)	324
A	To Learn on Your Own (if you want to): Cool and Important Things We Omitted From our Linear Algebra Introduction	327
A.1	Complex Numbers and Complex Vectors	328
A.1.1	Arithmetic of Complex Numbers: Enough to Get You By	329
A.1.2	Angles of Complex Numbers and Euler's Formula: More Advanced Aspects	330
A.1.3	Iterating with Complex Numbers: Background for Eigenvalues	332
A.1.4	\mathbb{C}^n , the Space of Complex Vectors	334
A.1.5	Iterating with Matrices: The Case for Eigenvalues and Eigenvectors	335
A.2	Eigenvalues and Eigenvectors	337
A.2.1	General Square Matrices	338
A.2.2	Real Symmetric Matrices	340
A.3	Positive Definite Matrices	342
A.4	Singular Value Decomposition or SVD	345
A.4.1	Motivation	345
A.4.2	Definition and Main Theorem	346
A.4.3	Numerical Linear Independence	348
A.5	Linear Transformations and Matrix Representations	350
A.6	Affine Transformations	353
B	What is an Ordinary Differential Equation?	355
B.1	Preliminaries: Expanding our Concept of an Equation	356
B.2	Time in a Digital Computer is Discrete	356
B.3	Digital Time may be Discrete, but We Can Make the Time Increment δt Quite Small	357
B.4	Equations with Derivatives in them are called Differential Equations	359
B.5	Discretization of ODEs of Higher Dimension	360

B.6	Julia Code for Generating the Figures	363
B.6.1	For Figures B.1 and B.2	363
B.6.2	Code for Figure B.3	363
B.6.3	Code for Figure B.4	364
B.6.4	Code for Figure B.5	364
B.6.5	Code for Figure B.6	364
C	Camera and LiDAR Models for Students of Robotics	367
C.1	Pinhole Camera Model	368
C.2	Preliminaries: Geometrical Transformations	368
C.2.1	Homogeneous Coordinates	368
C.2.2	2D Translation	368
C.2.3	2D Scaling	369
C.2.4	2D Rotation	369
C.2.5	Other 2D Geometrical Transformations in Homogeneous Coordinates	370
C.3	Pinhole Model	370
C.4	Geometric and Mathematical Relations	373
C.4.1	Intrinsic Matrix K	373
C.4.2	Projection Matrix	373
C.4.3	Extrinsic Matrix	373
C.4.4	Full Projection Matrix Workspace to Camera	373
C.5	Intrinsic Matrix Calibration	374
C.6	Nonlinear Phenomena in Calibration	374
C.7	Projection Map from LiDAR to Camera	375
C.8	Projection Map	375

Preface

This collection of course notes is dedicated to the group of students who were brave enough to take the pilot offering of ROB 101, Computational Linear Algebra in Fall 2020.

ROB 101 was conceived by Prof. Chad Jenkins as one part of a complete undergraduate curriculum in Robotics. Chad and I both thank Associate Dean for Undergraduate Education, Prof. Joanna Mirecki Millunchick, for her support in the offering of this course as ROB 101.

The following remarks are adapted from an Education Proposal led by Prof. Jenkins, Dr. Mark Guzdial, Ella Atkins, and myself.

A challenge for the current undergraduate curricular structure at the University of Michigan (and elsewhere) is that the best Robotics major is a quadruple major in ME, EE, CS, and Math with a minor in Psychology. The Robotics Institute at Michigan is poised to address this challenge in revolutionary ways as it evolves toward a department. The Robotics faculty are using the opportunity of a clean slate in the area of undergraduate robotics education—and the absolute necessity to integrate learning across a wide range of traditional disciplines—to design a new curricular system where computation, hardware, and mathematics are on an equal footing in preparing a future innovation workforce.

By integrating linear algebra, optimization, and computation in the first semester, students will experience mathematics as a means of making predictions and reasoning about experimental outcomes and robot designs. They will be prepared to encounter physics as a means to build mathematical models that describe the movement of objects, those with a palpable mass as well as electrons and waves, while grounding all of this in design. And computation? They will find it is the engine that drives discovery and the means of extracting information from data, allowing their machines to make decisions in real-time.

In addition to **ROB 101 (*Computational Linear Algebra*)** in the first year, we are planning **ROB 102 (*Graph Search for Robotics and AI*)**, which will show how computers can reason autonomously through graph search algorithms. The objective of the course is for students to implement a path planner for autonomous navigation with a given mobile robot at a known location in a known environment. The course will build towards providing a broader conceptual foundation for modeling problems as graphs and inferring solutions through search.

In **ROB 103 (*Robotic Mechanisms*)**, students will experience hands-on robotic systems design, build, and operation. The objective of the course is for students, in teams, to build an omni-drive mobile robot that can be teleoperated, as a step towards the gateway course: ROB 204. Students will learn to safely and efficiently operate modern shop tools such as 3D printers and laser cutters as well as traditional machining tools such as lathes, mills, and drill presses. Students will learn basic electronic and power systems principles including safe battery management, wiring harness design and assembly, and signal measurement for test and debugging. Students will design and build their real-world mobile robot based on application requirements from conception with CAD software through manufacturing, assembly, and test. Each student team will be given a “kit” of servos, sensors, and low-cost embedded processing components to use in their design.

The new first-year curriculum allows for major innovation in the second and third year curriculum, and the advancing of many graduate topics to much earlier places in Robotics education. We hope that you will follow our efforts in bringing this curriculum to the University of Michigan.

Jessy Grizzle
Fall Term, 2020

Philosophy of the Course

The Robotics faculty want out of the Sputnik era of educating engineers, where we are forced to pound our students with four semesters of Calculus before we are able to engage them in any interesting engineering. We seek to prepare students for the era of Information, AI, Data, and of course, Robotics. While we believe our ideas for this Linear Algebra course will work for most engineering departments, we understand that the College of Engineering needs a skunkworks to test some of our more revolutionary suggestions before turning them loose on everyone. We are proud to serve that role.

ROB 101 assumes a High School course in Algebra and no background in programming. With these entry requirements, we seek to open up mathematics and programming to everyone with the drive and skills to arrive at the University of Michigan's College of Engineering. From its inception in December 2019, the plan has always been to teach the course in a hybrid mode. We are being very intentional to design the course for inclusivity with a focus on making sure that one's zip code is not the best predictor of success. To do that, we are re-imagining the way mathematics is introduced to first-semester Y1 undergrads. We want to break the Sputnik era 4-semester calculus chain where AP credits are a huge predictor of success.

We will begin mathematics with Linear Algebra, the workhorse of modern autonomous systems, machine learning, and computer vision. We are integrating it with computation, and to make sure that students without access to high-end tools can be successful, all the programming will be run in a web browser through cloud services that are hidden from the student. If you can google at home or the library, then you can complete our programming assignments. Setting this up is a challenge, but we feel it is very important to have a platform that has equity in mind. Our plans for a hybrid mode of teaching are motivated by a desire to have students from Minority Serving Institutions join the course this fall.

The material in the course leaves out many traditional Linear Algebra topics, such as eigenvalues, eigenvectors, or how to compute determinants for matrices larger than 2×2 . Sounds crazy! Good. The course focuses on solving systems of linear equations at scale, meaning hundreds or thousands of variables, and all the mathematics in the book should be implementable in HW sets in Julia. With that as a premise, we chose to focus on a few things in Linear Algebra that work well at scale, such as triangular systems of equations. This led to the conviction that LU Factorization should be introduced early and in an understandable manner, which was made possible by a magical video by Prof. Gilbert Strang (MIT). Once you understand that matrix multiplication $C = A \cdot B$ can be done by multiplying the columns of A by the rows of B and summing them up, the LU Factorization becomes completely transparent, allowing triangular matrices to rule the day. And once you focus on triangular structures, even linear independence and linear combinations become much more approachable. Of course, some basic geometry is good, and thus Gram-Schmidt is featured along with the QR Factorization. In between, least squared error solutions to linear equations and regression are treated along with other useful tools.

The book wraps up with a treatment of root finding for both scalar and vector nonlinear equations, and a users' view of optimization that highlights the role that the abstract function "arg min" is playing in modern engineering.

Readers of the book will not find many applications of the juicy computational tools as they are treated in the HW sets, via jupyter notebooks, and in three amazing Projects. The first project leads students through the process of building a map for robot navigation from LiDAR data collected on the UofM North Campus Grove. The main Linear Algebra concept being explored is the transformation of points and collections of points in \mathbb{R}^3 under rigid body transformations. The second project is built around regression and will give students insight into the power of Machine Learning. The third project will focus on the control of a planar Segway using a simplified version of Model Predictive Control. Students will experience the excitement of balancing a challenging ubiquitous mobile platform, and while doing so, will learn about ODEs and their relation to iterative discrete-time models.

Finally, we thank Prof. Steven Boyd for making his Y1 Applied Linear Algebra course material open source (<http://vmls-book.stanford.edu/>). We have done the same (<https://tinyurl.com/sxk8n4u9>).

Jessy Grizzle and Maani Ghaffari Ann Arbor, Michigan USA

List of Algorithms and Methods or Things You Need Code to Do Well

The book covers the following algorithms. In HW, you will subsequently code most of them in the Julia Programming Language and apply them “at the scale of life”! The class projects will bring Linear Algebra to life.

- Forward substitution; see Chap. 3.5
- Back substitution; see Chap. 3.6
- Standard matrix multiplication; see Chap. 4.4
- A second way to do matrix multiplication; see Chap. 4.5
- LU Factorization without row permutations; see Chap. 5.5
- Using LU to solve $Ax = b$; see Chap. 5.6 and Chap. 5.9
- (Optional Read): LU Factorization with row permutations; see Chap. 5.7
- Using LU to compute $\det(A)$; see Chap. 6.1
- Building a row permutation matrix; see Chap. 6.5
- Checking linear independence; see Chap. 7.5.4
- (Optional Read): LDLT Factorization (aka Cholesky Factorization); see Chap. 7.6
- Counting number of linearly independent columns of a matrix; see Chap. 7.6
- Checking linear combinations; see Chap. 7.7
- Linear regression for overdetermined equations; see Chap. 8.2
- Gram-Schmidt; see Chap. 9.7.2 and (Optional Read): Modified Gram-Schmidt; see Chap. 9.12
- QR Factorization; see Chap. 9.8
- Linear regression for underdetermined equations; see Chap. 9.9
- Null space of a matrix; see Chap. 10.4
- Bisection Algorithm; see Chap. 11.2
- Numerical derivatives; see Chap. 11.3 and Chap. 11.5
- Newton’s Method; see Chap. 11.4 and Newton-Raphson Algorithm; see Chap. 11.6
- Gradient Descent; see Chap. 12.3
- Second-order optimization using the Hessian; see Chap. 12.5.3
- (Optional Read): Quadratic Program (QP) and max-margin classifier; see Chap. 12.8 and Chapter 13.3
- (Optional Read): Solving simple ODEs; see Appendix B

Chapter 1

Introduction to Systems of Linear Equations

Learning Objectives

- Get you going on Algebra, just in case Calculus has erased it from your mind
- Review on your own the quadratic equation.
- Set the stage for cool things to come.

Outcomes

- Refresher on the quadratic equation.
- Examples of systems of linear equations with two unknowns. Show that three things are possible when seeking a solution:
 - there is one and only one solution (one says there is a unique solution, which is shorthand for “there is a solution and it is unique”);
 - there is no solution (at all); and
 - there are an infinite number of solutions
- Notation that will allow us to have as many unknowns as we’ll need.
- Remarks that you can mostly ignore on why some numbers are called counting numbers, rational numbers, irrational numbers, or complex numbers
- Remarks on your first project.
- There is a programming manual associated with this book. Please see https://www.dropbox.com/s/ev6v8veutdjhkuk/ROB_101_Julia_Programming_Guide.pdf?dl=0.

1.1 For Review on Your Own: You Have Done Algebra Before

Quadratic Equation

$ax^2 + bx + c = 0$, where x is an *unknown variable* and typically a , b , and c are fixed real numbers, called *constants*. If $a \neq 0$, the solutions to this *nonlinear algebraic equation* are

$$x = \frac{-b \pm \sqrt{b^2 - 4ac}}{2a},$$

where the symbol \pm (read, plus minus) indicates that there is one solution x corresponding to the plus sign and a second solution x corresponding to the minus sign.

The *discriminant* of a quadratic equation is $\Delta := b^2 - 4ac$. If $\Delta \geq 0$, the two solutions are *real numbers*, while if $\Delta < 0$, the two solutions are *complex numbers*.

Complex Numbers

If you have not learned complex numbers (also known as (aka) *imaginary numbers*) or are fuzzy on the details, let your GSI know. We may not use complex numbers at all in the graded portion of the course. We're not sure yet! We will need complex numbers when we study eigenvalues of matrices, which could happen at the end of the term, or it could be that we do not get that far. Piloting a Linear Algebra course during a pandemic has never been done before!

Example 1.1 (*two distinct real solutions*) Find the roots of $2x^2 + 8x - 10 = 0$.

Solution:

$$\begin{aligned} a &= 2, b = 8, c = -10 \\ b^2 - 4ac &= 144 > 0 \\ x &= \frac{-8 \pm \sqrt{144}}{4} \\ &= \frac{-8 \pm 12}{4} \\ &= -2 \pm 3 \end{aligned}$$

The two solutions are $x = 1$ and $x = -5$. ■




Figure 1.1: The two roots are at the intersection of the quadratic with the line $y = 0$. Why? Because, by definition, the roots are values of x where $ax^2 + bx + c$ equals zero!

Example 1.2 (two repeated real solutions) Find the roots of $2x^2 + 8x + 8 = 0$.

Solution:

$$\begin{aligned} a &= 2, b = 8, c = 8 \\ b^2 - 4ac &= 0 \\ x &= \frac{-8 \pm \sqrt{0}}{4} \\ &= \frac{-8 \pm 0}{4} \\ &= -2 \pm 0 \end{aligned}$$

The two solutions are $x = -2$ and $x = -2$. ■




Figure 1.2: Note that there is now only a single intersection with $y = 0$ at $x = -2$. Because quadratic equations have two solutions, we say **the root at $x = -2$ is repeated**.

Example 1.3 (two distinct complex solutions) Find the roots of $2x^2 + 8x + 10 = 0$.

Solution:

$$\begin{aligned} a &= 2, b = 8, c = 10 \\ b^2 - 4ac &= -16 \\ x &= \frac{-8 \pm \sqrt{-16}}{4} \\ &= \frac{-8 \pm \sqrt{16}\sqrt{-1}}{4} \\ &= \frac{-8 \pm 4\sqrt{-1}}{4} \\ &= -2 \pm \sqrt{-1} \\ &= -2 \pm i \end{aligned}$$

The two solutions are $x = -2 + i$ and $x = -2 - i$, where $i := \sqrt{-1}$ is an *imaginary number*. If you have not already learned complex numbers, do not sweat it. If we need them at all in ROB 101, it will be at the end of the term and we will teach complex numbers before we use them! ■




Figure 1.3: Note that there are no intersections with $y = 0$ and hence there are not any real solutions.

Remark: As in many subjects, the vocabulary in mathematics reflects the many twists and turns of history. If you are expecting mathematics to be a *pure science* in the sense that it is 100% logical, then you will often be disappointed! The Greeks and Egyptians called $1, 2, 3, \dots$ *natural numbers* or *counting numbers* because they arose “naturally” in “counting objects”: one cow, two sheep, three coins, four sacks of flour, etc. They were also comfortable with *fractions*, $\frac{m}{n}$, where both m and n were counting numbers, and hence, n could never be zero. Fractions were called *rational numbers* based on the word *ratio*. The square root of two, $\sqrt{2}$, was known to the Babylonians, Indians and Greeks. Our notion of $\sqrt{2}$ is probably thanks to Pythagoras, the Ionian Greek philosopher known to you for developing the Pythagorean Theorem relating the sides of right triangles, yes, the famous formula $a^2 + b^2 = c^2$, which gives $1^2 + 1^2 = (\sqrt{2})^2$, where the symbol $\sqrt{2}$ stands for the “quantity that when squared gives the counting number 2”. It took a long time for the Greeks to accept that $\sqrt{2}$ could not be written as a *ratio of two counting numbers*. Eventually, it was proved to be *irrational*, that is, *not rational*, which means precisely that it cannot be expressed as the ratio of two counting numbers. In case you are interested, while Euclid was not the first to prove $\sqrt{2}$ was irrational, the beautiful reasoning he invented for the proof of $\sqrt{2}$ not being a rational number is still taught today. It is called *proof by contradiction*¹.

So far, so good with regards to vocabulary. But why call things involving $\sqrt{-1}$ complex numbers? Initially, mathematicians could not justify the existence of *quantities involving square roots of negative numbers*. Using them in calculations was a sign of “careless” mathematics and such numbers were treated as being *figments of one’s imagination*, literally, *imaginary numbers*. Nevertheless, some (brave) mathematicians found them convenient for solving equations and others even for describing physical phenomena. Eventually, formal algebra caught up with the notion of imaginary numbers and their rules of use were rigorously justified. The name imaginary numbers stuck, however! Here is a link to a slightly simplified explanation of how mathematicians “codify” the existence of complex numbers: <http://www.math.toronto.edu/mathnet/answers/imagexist.html> Your instructors are unsure if they could have followed the reasoning in this document when they were at your stage of college, so do not sweat it.

Just for fun, you might want to learn more about numbers on Wikipedia <https://en.wikipedia.org/wiki/Number>. Were negative numbers always accepted? What about the notion of zero? **None of these remarks on numbers are required reading.** You will not be tested on the history of numbers!

To Know

- What are the counting numbers (also known as the natural numbers)?
- What are rational numbers?
- Be able to give an example of an irrational number, but of course, you are not expected to prove it is irrational.
- Later in the course (though it is not sure): what are imaginary numbers?

¹It may seem remarkable to you that the two words “proof” and “contradiction” can coexist in logical statements. As you advance in your mathematical training, you may come across the term again. “Proof by contradiction” is not for amateurs...think about it as semi-professional-grade logic and math.

In ROB 101, the first ten Chapters focus on linear equations, hence, equations that do not have quadratic terms x^2 , cubic terms x^3 , nor terms with higher powers; they also do not have $\sin(x)$, $\cos(x)$, \sqrt{x} , or e^x or anything like that. It is perhaps hard to believe that equations with only order-one terms and constants could be interesting or important, but they are both interesting and important!

1.2 Linear Systems of Equations: What can happen?

We begin with several examples to illustrate what can happen.

Same number of equations as unknowns, and there is a unique solution:

$$\begin{aligned}x + y &= 4 \\2x - y &= -1\end{aligned}\tag{1.1}$$

One way to compute a solution is to solve for x in terms of y in the first equation and then substitute that into the second equation,

$$\begin{aligned}x + y = 4 &\implies x = 4 - y \\2x - y = -1 &\implies 2(4 - y) - y = -1 \\&\implies -3y = -9 \\&\implies y = 3 \\&\text{going back to the top} \\x = 4 - y &\implies x = 1\end{aligned}$$

You can try on your own solving for y in terms of x and repeating the above steps. You will obtain the same answer, namely $x = 1, y = 3$.

Another “trick” you can try, just to see if we can generate a different answer, is to add the first equation to the second, which will eliminate y ,

$$\begin{array}{r}x + y = 4 \\+ 2x - y = -1 \\ \hline3x + 0y = 3 \\ \implies x = 1\end{array}$$

Going back to the top and using either of the two equations

$$\begin{aligned}x + y = 4 &\implies y = 4 - x \\&\implies y = 3 \\&\text{or} \\2x - y = -1 &\implies -y = -2x - 1 \\&\implies -y = -3 \\&\implies y = 3\end{aligned}$$

gives the same answer as before, namely, $x = 1, y = 3$.

In fact, the set of equations (1.1) has one, and only one, solution. In math-speak, one says the set of equations (1.1) has a *unique solution*. Often, we will stack x and y together and write the solution as

$$\begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} 1 \\ 3 \end{bmatrix}.$$

Same number of equations as unknowns, and there is no solution:

$$\begin{aligned}x - y &= 1 \\2x - 2y &= -1\end{aligned}\tag{1.2}$$

Because there are only two equations with a very nice set of numbers you might notice almost immediately that the left-hand side of the second equation is twice the left-hand side of the first equation, namely, $2x - 2y = 2(x - y)$, but when we look to the right-hand sides, $-1 \neq 2 \cdot 1$, and hence the equations (1.2) are *inconsistent*.

While the above is one way to analyze the problem, let's try to find a solution just as we did for the first set of equations,

$$\begin{aligned}x - y = 1 &\implies x = y + 1 \\2x - 2y = -1 &\implies 2(y + 1) - 2y = -1 \\&\implies 2 + 2y - 2y = -1 \\&\implies 2 = -1.\end{aligned}$$

Hence, trying to solve the equations has led us to a *contradiction*, namely $2 = -1$. Perhaps if we had solved for y in terms of x , we could have found a solution? Let's try

$$\begin{aligned}x - y = 1 &\implies -y = -x + 1 \\&\implies y = x - 1 \\2x - 2y = -1 &\implies 2x - 2(x - 1) = -1 \\&\implies 2x - 2x + 2 = -1 \\&\implies 2 = -1\end{aligned}$$

again! No matter how you manipulate the equations (1.2), while obeying “the rules of algebra” (which we have **not yet learned** for systems of linear equations), you cannot extract a sensible answer from these equations. They really are *inconsistent*.

Same number of equations as unknowns, and there are an infinite number of solutions:

$$\begin{aligned}x - y &= 1 \\2x - 2y &= 2\end{aligned}\tag{1.3}$$

Because there are only two equations with a very nice set of numbers, you might notice that the left-hand side of the second equation is twice the left-hand side of the first equation, namely, $2x - 2y = 2(x - y)$, and this time, when we look to the right-hand sides, $2 = 2 \cdot 1$, and hence the two equations are actually the “same” in the sense that one equation can be obtained from the other equation by multiplying both sides of it by a non-zero constant. We could elaborate a bit, but it's not worth it at this point. Instead, let's approach the solution just as we did in the first case.

We solve for x in terms of y in the first equation and then substitute that into the second equation,

$$\begin{aligned}x - y = 1 &\implies x = y + 1 \\2x - 2y = 2 &\implies 2(y + 1) - 2y = 2 \\&\implies 2y + 2 - 2y = 2 \\&\implies 2 = 2.\end{aligned}$$

The conclusion $2 = 2$, is perfectly correct, but tells us nothing about y . In fact, we can view the value of y as an arbitrary *free parameter* and hence the solution to (1.3) is

$$x = y + 1, \quad -\infty < y < \infty.$$

The solution can also be expressed as

$$y = x - 1, \quad -\infty < x < \infty,$$

which perhaps inspires you to plot the solution as a line in \mathbb{R}^2 , with slope $m = 1$ and y -intercept $b = -1$.

Summary So Far

Consider a set of two equations with two unknowns x and y

$$\begin{aligned}a_{11}x + a_{12}y &= b_1 \\a_{21}x + a_{22}y &= b_2,\end{aligned}\tag{1.4}$$

constants $a_{11}, a_{12}, a_{21}, a_{22}$ and b_1, b_2 . Depending on the values of the constants, the linear equations (1.4) can have a unique solution, no solution, or an infinity of solutions. The equations cannot have two and only two distinct solutions.

More equations than unknowns typically means that there are no solutions: The system of equations

$$\begin{aligned}x &= 1 \\y &= 2 \\x + y &= a\end{aligned}\tag{1.5}$$

will only have a solution when $a = 3$. For all other values of a , it will not have a solution. We will learn to recognize later when the set of equations are consistent. At this point in the course, we do not have adequate mathematical tools to address the issue.

Limit of Hand Solutions

When there are only two equations and two unknowns, determining *by hand* if the equations have one solution, no solution, or an infinity of solutions is quite do-able. With sufficient motivation and “nice numbers”, three equations and three unknowns is also not too bad. At four, it becomes super tedious and errors begin to sprout like weeds. Hence, what about 100 equations and 100 unknowns? **Our four-week goal is for you to handle systems of linear equations with hundreds of variables with confidence and ease** As you can imagine, this is where the “computational” part of ROB 101’s name comes into play!

1.3 Naming Variables

If we have two variables (also called unknowns), it is natural to call them x and y . If we have three variables, naming them x , y , and z works fine. If we have 26 variables, would we start with a and go all the way to z ? What if we have more variables? Well, there are 24 Greek letters, do we add them to the list? Throw in Mandarin Characters to get us to several thousand variables? Clearly, this is not a practical way to go. The only real possibility is to add counting numbers, because there are as many of them as we need, and computers understand numbers!

Welcome to $x_1, x_2, x_3, \dots, x_n$, or y_1, y_2, \dots, y_m etc.

1.4 A 3 x 3 Example

Here is a *system of linear equations* with variables x_1, x_2, x_3 .

$$\begin{aligned}x_1 + x_2 + 2x_3 &= 7 \\2x_1 - x_2 + x_3 &= 0.5 \\x_1 + 4x_3 &= 7\end{aligned}\tag{1.6}$$

The strategy for seeking a solution is the same as with two equations and two unknowns: solve for one of the variables in terms of the other variables, substitute into the remaining equations, simplify them, and repeat. We can start anywhere, so let’s solve for x_1 in the bottom equation and plug that back into the two equations above it. Doing so gives us,

$$x_1 = 7 - 4x_3\tag{1.7}$$

and then

$$\begin{aligned}\underbrace{(7 - 4x_3)}_{x_1} + x_2 + 2x_3 &= 7 \implies x_2 - 2x_3 = 0 \\2\underbrace{(7 - 4x_3)}_{2x_1} - x_2 + x_3 &= 0.5 \implies -x_2 - 7x_3 = -13.5 \\x_2 - 2x_3 &= 0 \implies x_2 = 2x_3 \\-x_2 - 7x_3 &= -13.5 \implies \underbrace{-(2x_3)}_{-x_2} - 7x_3 = -13.5 \\&\implies -9x_3 = -13.5 \\&\implies x_3 = 1.5\end{aligned}\tag{1.8}$$

Plugging “ $x_3 = 1.5$ ” into (1.7) gives

$$x_1 = 7 - 4 \cdot (1.5) = 7 - 6 = 1.$$

And then from (1.7), we have that

$$x_2 = 2x_3 = 2 \cdot (1.5) = 3.$$

Hence, the solution is

$$\begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} 1.0 \\ 3.0 \\ 1.5 \end{bmatrix}. \quad (1.9)$$

Tedium and Pain \implies Motivation!

The hope is that you found the 3×3 example super tedious and something you’d prefer to avoid! The more tedious and painful you found it, the more motivation you will have to learn some new math that will make solving systems of linear equations a snap.

1.5 Greek Alphabet

It’s hard to believe that the 26 Latin letters we know and love, a, b, c, \dots, x, y, z , especially after being augmented with numerical subscripts, do not provide enough symbols! You might as well get used to it, **engineers and mathematicians often use the 24 Greek letters as well**: $\alpha, \beta, \gamma, \dots, \chi, \psi, \omega$. Here is a link to the Greek alphabet with names of the letters and their pronunciations: <https://web.mit.edu/jmorzins/www/greek-alphabet.html>. 99.9% of all engineers, including your author, mispronounce them, so do not feel intimidated by that. Just gleefully join the club!

There are several accepted modern pronunciations of the Greek alphabet, it seems, giving you even less worry about your own pronunciation. The links below give a few of them:

1. Greek Alphabet - Pronunciation of each letter: <https://www.youtube.com/watch?v=1FyEWbwBarQ>
2. Greek Alphabet Rap Song: <https://www.youtube.com/watch?v=w3D5ERMOpMk>
3. Greek Alphabet Song (Nursery-rhyme Style): <https://www.youtube.com/watch?v=3gaeIUsPJ-Y>
4. Learn the Greek Alphabet in Less Than 10 Minutes: <https://www.youtube.com/watch?v=BQVoz-HX2cA>

If you lookup the origin of the word “AlphaBet”, you will learn that it comes from Alpha Beta, the first two letters of the Greek Alphabet.

1.6 (Optional Read) Where is Computational Linear Algebra Used?

Figure 1.5 shows numerous areas where Computational Linear Algebra is used. In ROB 101, we have built projects around three different themes to give you a chance to really dive into a subject and feel that you learned something beyond the math itself:

- For **Project 1: Map Building from LiDAR Data** you will be given data collected on the Cassie Blue bipedal robot during an experiment on the North Campus Grove and are asked to transform the data for building a map and visualizing it using Julia. The underlying work you will do is very similar to what is done in real-time² on Cassie in order to build a map for autonomous navigation. You may wish to view the videos <https://www.youtube.com/watch?v=pNyXsZ5zVZk> and <https://youtu.be/gE3Y-2Q3gco> to see mapping and navigation done in real-time. Yes, this is very similar to what is done by AVs (Autonomous Vehicles). One difference is that an AV has 200 Kg of electronics in its trunk to process all the data, while Cassie’s entire autonomy package weighs in at 9 Kg and runs off a hobbyist LiPo battery.

²Real-time means the computations are done quickly enough on the robot that they can be used almost immediately. This is opposed to “off-line” where you collect data on the robot and then do the processing on a desktop in the lab. You are clearly doing offline data processing, which is easier, because there are no computation time requirements.

- For **Project 2: Precipitation Data in Alaska (A True Story)**, you will use linear regression methods taught in Chap. 8 and apply them to a much larger dataset from the U.S. National Oceanic and Atmospheric Administration (NOAA). The project will give you insight into Machine Learning, one facet of the burgeoning area of Artificial Intelligence or AI for short. You will learn an important new function called the radial basis function and apply it towards building a mathematical model of a surface.
- For **Project 3: Feedback Control of a Segway**, you will learn how to balance and “drive” an inherently unstable mobile robot! If you have not seen a Segway, here is a video of the bideal robot Cassie Blue riding a Segway <https://youtu.be/0gauVSUJzd0>. Your project will not be this awesome, but the project will put you on the road to awesome things!

Now, in case just reading about these projects intimidates you, we want to assure you that their difficulty has been thoughtfully scaled to a first-semester Y1 experience, while also giving you insight into how the real things are done. This is possible because we will teach you math and programming in an integrated manner. The math reinforces the programming and the programming reinforces the math. At each step of the way, you will increase your confidence in Engineering, Mathematics, or the Sciences as a career choice. In High School, you were mostly taught math as being disconnected from real applications. There was a huge emphasis on memorizing formulas or theorems, without ever really using the concepts to do something fun. In ROB 101, we hope to show you how empowering it is to do mathematics at the scale of life.

1.7 Looking Ahead

We need a straightforward way to check whether a set of equations falls into the nice case of having a unique solution or is it one of the “problematic cases” where there may be no solution at all or an infinity of solutions. To get there, we need to learn:

- what are *matrices* and *vectors*;
- how to express a system of linear equations in terms of *matrices* and *vectors*;
- what it means for a set of linear equations to be *triangular* or *square*; and
- what is the *determinant* of a matrix.






Figure 1.4: Cassie Blue with her perception package: a 32-Beam Velodyne LiDAR and an Intel RealSense Camera. In your first project, you will learn what is a LiDAR sensor and how to process the LiDAR data to build a “map” that a robot can use for navigation. Each colored dot in Fig. 1.5a is a LiDAR measurement. How many are there? Ten thousand, maybe? The image is the result of combining multiple LiDAR scans into a single image. You will learn that this involves applying matrices to vectors. You will learn about coordinate frames. Right now, this seems like magic. In a few weeks, you will be comfortable enough with the Julia programming language to manipulate a dataset with 10,000 points or more. **Remember, one of our goals is mathematics and programming at the scale of life!**




(a)




(b)




(c)



(d)



(e)



(f)

Figure 1.5: Just a few of the almost infinitely many application areas of linear algebra and programming. (a) Cassie Blue building a map on the UofM North Campus Grove. You will study this in Project 1. (b) Fitting a function to data allows one to model new situations that no one has seen before; image courtesy of Wikimedia Commons. You will study this in Project 2. (c) A Segway-like structure that requires a feedback loop to balance it; image courtesy of www.soarboards.com. You will study this in Project 3. (d) Gradient descent is used to find minima of functions, image courtesy of <https://www.analyticsvidhya.com>. You will study this in Chap. 12. (e) Checkerboard or other regular patterns are used to calibrate cameras so that they can measure the position of objects in a scene; image courtesy of <https://www.imatest.com/> (f) An electric circuit with a non-inverting amplifier that converts a voltage signal to a desired current, such as one might use to drive a motor or an electromagnet; image courtesy of <https://wiki.analog.com/university/courses/electronics/text/chapter-4>. Circuits are studied in EECS 215.

Chapter 2

Vectors, Matrices, and Determinants

Learning Objectives

- Begin to understand the vocabulary of mathematics and programming
- Answer the question: why make a distinction between integers and decimal numbers.
- An introduction to the most important tools of linear algebra: vectors and matrices.
- Find out an easy way to determine when a set of linear equations has a unique answer.

Outcomes

- Scalars vs Array
- Row vectors and column vectors
- Rectangular matrices and square matrices
- Learn new mathematical notation $x := y$, which means that x is by definition equal to y . You may be more used to $x \triangleq y$.
- Using matrices and vectors to express systems of linear equations
- Determinant of a square matrix and its relation to uniqueness of solutions of systems of linear equations

2.1 Scalars and Arrays

Scalars are simply numbers such as the ones you have been using for a long time: 25.77763 , $\sqrt{17}$, 10 , -4 , π . In the Julia programming language, you will soon learn that for computational and storage efficiency, Julia differentiates between scalars that require decimal points and those that do not. Get ready for that! In ROB 101, when we do math, scalars are just numbers. When we do programming, scalars that do not require decimal points are called INTEGERS and those that do require decimal points are called FLOATING POINT NUMBERS because, with the same number of zeros and ones¹, Julia has to represent very large numbers such as 5.972×10^{24} , the mass of the earth in kilograms, and very small numbers, such as the mass of one atom of lead in kilograms $3.4406366 \times 10^{-22}$. To do that, where the decimal point appears in a list of zeros and ones has to “float”, that is, it varies from number to number.

In ROB 101, we will not need to learn how computers represent numbers in *binary*. We will simply accept that computers use a different representation for numbers than we humans use. The more zeros and ones we allow in the representation of number, the more space it takes and the longer it takes to add them or multiply them, etc. The **computer** that provided navigational assistance for the moon landing on 20 July 1969 had a 16 bit word length, meaning its computations were based on groups of 16 binary digits (zeros and ones), called *bits*. In Julia, we’ll typically use 64 zeros and ones to represent numbers:

- Int64
- Float64

Remark (can skip): In case you are wondering, ∞ (infinity) is a concept: it is NOT a number. Because it is not a number, ∞ is neither an integer nor a decimal number. The symbol ∞ is used to indicate that there is no positive upper bound for how big something can grow (i.e., it can get bigger and bigger and bigger . . .), while $-\infty$ is similar to indicate that something can get more and more negative without end. In ROB 101, you will only encounter ∞ when you try to divide something by zero. Julia has been programmed to rock your world when you do that! (Just kidding).

	A	B	C	D	E	F	G
1	36	0.157822	114.2628	2536	10675	4808	14122
2	36	0.157822	114.2628	2536	10672	4808	14122
3	36	0.157822	114.2628	2536	10672	4808	14122
4	36	0.157822	114.2628	2536	10672	4808	14089
5	35	0.153377	111.0451	2535	10672	4808	14061
6	35	0.153377	111.0451	2535	10672	4808	14001
7	34	0.148936	107.8298	2534	10672	4808	13982
8	33	0.144499	104.617	2533	10672	4808	13911
9	32	0.140065	101.4068	2532	10672	4808	13823
10	30	0.131207	94.99374	2530	10672	4808	13756
11	30	0.131207	94.99374	2530	10672	4808	13683
12	29	0.126783	91.79099	2529	10672	4808	13609
13	24	0.104717	75.81477	2524	10672	4808	13494
14	24	0.104717	75.81477	2524	10672	4808	13470
15	20	0.087125	63.07885	2520	10672	4808	13290
16	6	0.025992	18.81852	2506	10673	4808	13209
17	-1	-0.00432	-3.12766	2499	10670	4807	13155
18	-18	-0.07723	-55.9149	2482	10666	4806	13119
19	-27	-0.11543	-83.5682	2473	10666	4806	13059
20	-56	-0.23659	-171.294	2444	10666	4806	12965

Figure 2.1: Here is an array of numbers from a spreadsheet. The rows are numbered while the columns are labeled with letters.

¹Zeros and ones are sometimes called the binary language of computers. If you are interested in binary numbers and binary arithmetic, you can find many sources on the web.

Arrays are scalars that have been organized somehow into lists. The lists could be rectangular or not, they could be made of columns of numbers, or rows of numbers. If you’ve ever used a spreadsheet, then you have seen an array of numbers. In Fig. 2.1, columns A, D, E, F, and G have only integer numbers in the them, while columns B and C use decimal numbers. Each row has both integer numbers and decimal numbers.

2.2 Row Vectors and Column Vectors

For us, a *vector* is a finite ordered list of numbers or of unknowns. In Julia, a vector is a special case of an *array*. For example

$$v = \begin{bmatrix} 1.1 \\ -3 \\ 44.7 \end{bmatrix} \quad (2.1)$$

is a vector of *length* three; sometimes we may call it a *3-vector*. By *ordered* we mean the list has a first element $v_1 = 1.1$, a second element $v_2 = -3$, and a third element $v_3 = 44.7$, and we also mean that if we change the order in which we list the elements, we (usually) obtain a different vector. For example, the vector

$$w = \begin{bmatrix} 1.1 \\ 44.7 \\ -3 \end{bmatrix} \quad (2.2)$$

is not equal to the vector v , even though they are made of the same set of three numbers.

Vectors written from “top” to “bottom” as in (2.1) and (2.2) are called *column* vectors. It is also useful to write vectors from “left” to “right” as in

$$v^{\text{row}} = [1.1 \quad -3 \quad 44.7] \quad (2.3)$$

and we call them *row* vectors. What we said about the word “ordered” applies equally well to row vectors in that the row vector v has a first element $v_1^{\text{row}} = 1.1$, a second element $v_2^{\text{row}} = -3.0$, and a third element $v_3^{\text{row}} = 44.7$. Moreover, if we change the order in which we list the elements, we (usually) obtain a different row vector. Here, we were super careful and added the superscript ^{row} to v^{row} to clearly distinguish the symbol v being used for the column vector (2.1) from the row vector (2.3). Normally, we will not be that fussy with the notation, BUT, you must be aware that column vectors and row vectors are different “animals” except in the case that they have length one, as in $v = [v_1]$ is both a row vector and a column vector.

A general length n column vector is written like this,

$$v = \begin{bmatrix} v_1 \\ \vdots \\ v_n \end{bmatrix}, \quad (2.4)$$

while a general length n row vector is written like this

$$v = [v_1 \quad \cdots \quad v_n]. \quad (2.5)$$

This general notation allows for the case that $n = 1$, yielding

$$v = [v_1],$$

which as we noted above, is both a row vector and a column vector.

Here are some examples in Julia, thanks to Prof. Maani Ghaffari:

```
1 # we define an array of numbers.
2 # a is a 1x5 array
3 a = [1 -2 4 8.1 2^0.5]
```

Output

```
1x5 Matrix{Float64}:
 1.0 -2.0 4.0 8.1 1.41421
```

```
1 # b is a 5x1 array
2 b = [1, -2, 4, 8.1, 2^0.5]
```

Output

```
5-element Vector{Float64}:
 1.0
-2.0
 4.0
 8.1
1.4142135623730951
```

```
1 # b is a 5x1 array
2 b = [1, -2, 4, 8.1, 2^0.5]
```

Output

```
5-element Vector{Float64}:
 1.0
-2.0
 4.0
 8.1
1.4142135623730951
```

```
1 # or
2 c = [1; -2; 4; 8.1; 2^0.5]
```

Output

```
5-element Vector{Float64}:
 1.0
-2.0
 4.0
 8.1
1.4142135623730951
```

2.3 Remark on Brackets

Usually, in this course and in most books, *square brackets* [] are used to enclose vectors, but that is mainly *a matter of taste* as you can also use *parentheses* () as in

$$v = \begin{pmatrix} v_1 \\ \vdots \\ v_n \end{pmatrix} \quad (2.6)$$

and

$$v = (v_1 \ \cdots \ v_n). \quad (2.7)$$

In most programming languages, and Julia is no exception, you can only use square brackets! The choice of symbols, words, and their allowed arrangements in a language is called *syntax*. In a programming language, syntax is typically much more restrictive than in a written version of a spoken language. At some point, you may appreciate that throwing errors for bad syntax helps us to reduce *ambiguity* and *bugs* when we program. *You have likely experienced that ambiguities in spoken language can sometimes lead to bad outcomes.*

2.4 Matrices, Rectangular and Square, and the Matrix Diagonal

Matrices are generalizations of vectors that allow multiple columns and rows, where each row must have the same number of elements². Here is a 3×2 matrix,

$$A = \begin{bmatrix} 1 & 2 \\ 3 & 4 \\ 5 & 6 \end{bmatrix}, \quad (2.8)$$

meaning it has three rows and two columns, while here is a 2×3 matrix,

$$A = \begin{bmatrix} 1.2 & -2.6 & 11.7 \\ 3.1 & \frac{11}{7} & 0.0 \end{bmatrix}, \quad (2.9)$$

meaning it has two rows and three columns. It is customary to call a $1 \times n$ matrix a row vector and an $n \times 1$ matrix a column vector, but it is perfectly fine to call them matrices too!

A general $n \times m$ matrix

$$A = \begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1m} \\ a_{21} & a_{22} & \cdots & a_{2m} \\ \vdots & \vdots & \ddots & \vdots \\ a_{n1} & a_{n2} & \cdots & a_{nm} \end{bmatrix}, \quad (2.10)$$

is said to be *rectangular of size $n \times m$* (one reads this as “ n by m ”), and when $n = m$, we naturally say that the $n \times n$ matrix

$$A = \begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1n} \\ a_{21} & a_{22} & \cdots & a_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ a_{n1} & a_{n2} & \cdots & a_{nn} \end{bmatrix} \quad (2.11)$$

is *square*. We note that a_{ij} , the ij -element of A , lies on the intersection of the i -th row and the j -th column.

Definition: The **diagonal** of the square matrix A in (2.11) is

$$\text{diag}(A) = [a_{11} \ a_{22} \ \cdots \ a_{nn}] \quad (2.12)$$

What we are calling the diagonal is sometimes called the *main diagonal of a matrix*. We now highlight the diagonal in red to help those with a “visual memory”

$$A = \begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1n} \\ a_{21} & a_{22} & \cdots & a_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ a_{n1} & a_{n2} & \cdots & a_{nn} \end{bmatrix} \iff \text{diag}(A) = [a_{11} \ a_{22} \ \cdots \ a_{nn}].$$

While it is possible to define the diagonal of general rectangular matrices, we will not do so at this time. Here are some examples in Julia, once again thanks to Prof. Maani Ghaffari

```
1 # Let's define a 3x2 matrix
2 A = [1 2; 3 4; 5 6]
```

Output

```
3x2 Array{Int64,2}:
 1  2
 3  4
 5  6
```

²Equivalently, each column has the same number of elements.

```

1 # Let's define a 2x3 matrix
2 A = [1.2 -2.6 11.7; 3.1 11/7 0]

```

Output

```

2x3 Array{Float64,2}:
 1.2  -2.6   11.7
 3.1  1.57143  0.0

```

```

1 # Here is a big matrix
2 using Random
3 A = randn(6,6)

```

Output

```

6x6 Matrix{Float64}:
-0.518479  0.693952 -0.137698  0.200556  2.03276 -1.09174
 0.575811  0.999473 -0.0427593 -0.903915  2.14338  0.240728
 0.570255  0.477523 -0.503201 -0.864054 -0.661544  0.0821051
-0.681514 -1.02591 -0.418878  0.248959 -0.776872  0.466698
 0.395184  1.72782  0.976437  1.10196  0.892258 -1.48822
 1.15146 -0.161273 -0.691775 -1.07168  0.909486 -1.02277

```

```

1 # compute its diagonal
2 using LinearAlgebra
3 diagA=diag(A)

```

Output

```

6-element Vector{Float64}:
-0.5184785168661076
 0.9994728076152639
-0.5032006177084569
 0.24895902626040853
 0.8922580664116776
-1.022769405730864

```

2.5 Expressing a System of Linear Equations in terms of Vectors and Matrices

Before we even talk about “matrix-vector multiplication”, we can address the task of writing a system of linear equations in “matrix-vector form”. In the beginning, we will diligently follow the notation $Ax = b$, so let’s see how we can identify a matrix A , a column vector x , and a column vector b . We’ll do a few examples before giving a “general method”.

Example 2.1 *Express the System of Linear Equations in Matrix Form:*

$$\begin{aligned}
 x_1 + x_2 &= 4 \\
 2x_1 - x_2 &= -1.
 \end{aligned}
 \tag{2.13}$$

Solution: When your instructors look at this equation, they see two unknowns x_1 and x_2 and *coefficients*³ associated with them on the left-hand and right-hand sides of the two equations.

³Coefficients in this case are the numbers multiplying x_1 and x_2 . An equation typically has variables (unknowns) and coefficients (numbers) that multiply the unknowns or that determine what the equation is supposed to equal, as in $3x_1 + 4x_2 = 7$.

We can use the unknowns to define a column vector of length two, the four coefficients multiplying the unknowns to build a 2×2 matrix A , and the two coefficients on the right-hand side to define another column vector of length two,

$$\begin{aligned} x_1 + x_2 = 4 \\ 2x_1 - x_2 = -1 \end{aligned} \iff \underbrace{\begin{bmatrix} 1 & 1 \\ 2 & -1 \end{bmatrix}}_A \underbrace{\begin{bmatrix} x_1 \\ x_2 \end{bmatrix}}_x = \underbrace{\begin{bmatrix} 4 \\ -1 \end{bmatrix}}_b. \quad (2.14)$$

■

Example 2.2 Express the System of Linear Equations in Matrix Form:

$$\begin{aligned} x_1 - x_2 = 1 \\ 2x_1 - 2x_2 = -1. \end{aligned} \quad (2.15)$$

Solution: We now jump straight into it. We form the 2×1 vector x of unknowns as before and place the coefficients by rows into the 2×2 matrix A ,

$$\begin{aligned} x_1 - x_2 = 1 \\ 2x_1 - 2x_2 = -1 \end{aligned} \iff \underbrace{\begin{bmatrix} 1 & -1 \\ 2 & -2 \end{bmatrix}}_A \underbrace{\begin{bmatrix} x_1 \\ x_2 \end{bmatrix}}_x = \underbrace{\begin{bmatrix} 1 \\ -1 \end{bmatrix}}_b. \quad (2.16)$$

■

Example 2.3 Express the System of Linear Equations in Matrix Form:

$$\begin{aligned} 3x_1 + x_2 + 2x_3 = 7 \\ 2x_1 - x_2 + 4x_3 = 4 \end{aligned} \quad (2.17)$$

Solution: We can have the number of equations different from the number of unknown variables. In this case, we have more unknowns than equations.

$$\begin{aligned} 3x_1 + x_2 + 2x_3 = 7 \\ 2x_1 - x_2 + 4x_3 = 4 \end{aligned} \iff \underbrace{\begin{bmatrix} 3 & 1 & 2 \\ 2 & -1 & 4 \end{bmatrix}}_A \underbrace{\begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}}_x = \underbrace{\begin{bmatrix} 7 \\ 4 \end{bmatrix}}_b \quad (2.18)$$

■

From Multiple Linear Equations to Matrices and Vectors

We will work an example with three equations and four variables. The same method works for n equations in m variables.

Goal: Transform a system of linear equations into matrix form $Ax = b$

$$\begin{aligned} a_{11}x_1 + a_{12}x_2 + a_{13}x_3 + a_{14}x_4 &= b_1 \\ a_{21}x_1 + a_{22}x_2 + a_{23}x_3 + a_{24}x_4 &= b_2 \\ a_{31}x_1 + a_{32}x_2 + a_{33}x_3 + a_{34}x_4 &= b_3. \end{aligned} \tag{2.19}$$

Two parts are really easy. We define two column vectors

$$x := \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{bmatrix} \text{ and } b := \begin{bmatrix} b_1 \\ b_2 \\ b_3 \end{bmatrix}. \tag{2.20}$$

The vector x has four rows because there are $m = 4$ unknowns, x_1, x_2, x_3 and x_4 . The vector b has three rows because there are $n = 3$ equations.

Now, what about the matrix A ? First of all, it's size is 3×4 because it has one row for each equation and one column for each unknown. You can think about it as follows. We place each term $a_{ij}x_j$ in the entry corresponding to the i -th row and j -th column. For example $a_{23}x_3$ goes in the second row and third column; we highlight it below

$$\begin{array}{c} \text{row}_1 \\ \text{row}_2 \\ \text{row}_3 \end{array} \underbrace{\begin{bmatrix} a_{11}x_1 & a_{12}x_2 & a_{13}x_3 & a_{14}x_4 \\ a_{21}x_1 & a_{22}x_2 & a_{23}x_3 & a_{24}x_4 \\ a_{31}x_1 & a_{32}x_2 & a_{33}x_3 & a_{34}x_4 \end{bmatrix}}_{\begin{array}{cccc} \text{col}_1 & \text{col}_2 & \text{col}_3 & \text{col}_4 \end{array}} = \begin{bmatrix} b_1 \\ b_2 \\ b_3 \end{bmatrix}. \tag{2.21}$$

We note that row_1 of the matrix has the terms for the first equation, row_2 has the terms for the second equation, etc, while col_1 is for the x_1 terms of each of the equations, col_2 is for the x_2 terms, etc. If we add up all of the entries in a given row, we obtain the corresponding equation.

Next, we move the x_i 's out of the matrix to the right-hand side while leaving the coefficients where they are, like so

$$\begin{array}{c} \text{row}_1 \\ \text{row}_2 \\ \text{row}_3 \end{array} \underbrace{\begin{bmatrix} a_{11}x_1 & a_{12}x_2 & a_{13}x_3 & a_{14}x_4 \\ a_{21}x_1 & a_{22}x_2 & a_{23}x_3 & a_{24}x_4 \\ a_{31}x_1 & a_{32}x_2 & a_{33}x_3 & a_{34}x_4 \end{bmatrix}}_{\begin{array}{cccc} \text{col}_1 & \text{col}_2 & \text{col}_3 & \text{col}_4 \end{array}} = \begin{bmatrix} b_1 \\ b_2 \\ b_3 \end{bmatrix} \iff \begin{array}{c} Eq_1 \\ Eq_2 \\ Eq_3 \end{array} \underbrace{\begin{bmatrix} a_{11} & a_{12} & a_{13} & a_{14} \\ a_{21} & a_{22} & a_{23} & a_{24} \\ a_{31} & a_{32} & a_{33} & a_{34} \end{bmatrix}}_{\begin{array}{cccc} \text{var}_1 & \text{var}_2 & \text{var}_3 & \text{var}_4 \end{array}} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{bmatrix} = \begin{bmatrix} b_1 \\ b_2 \\ b_3 \end{bmatrix} \tag{2.22}$$

to obtain

$$\begin{aligned} a_{11}x_1 + a_{12}x_2 + a_{13}x_3 + a_{14}x_4 &= b_1 \\ a_{21}x_1 + a_{22}x_2 + a_{23}x_3 + a_{24}x_4 &= b_2 \\ a_{31}x_1 + a_{32}x_2 + a_{33}x_3 + a_{34}x_4 &= b_3. \end{aligned} \iff \underbrace{\begin{bmatrix} a_{11} & a_{12} & a_{13} & a_{14} \\ a_{21} & a_{22} & a_{23} & a_{24} \\ a_{31} & a_{32} & a_{33} & a_{34} \end{bmatrix}}_A \underbrace{\begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{bmatrix}}_x = \underbrace{\begin{bmatrix} b_1 \\ b_2 \\ b_3 \end{bmatrix}}_b \tag{2.23}$$

Remark: After a bit of practice, you will go straight from (2.19) to (2.22) (in other words, from the left-hand side of (2.23) to the right-hand side of (2.23)), without passing through the intermediate step given in (2.21).

Example 2.4 Express the System of Linear Equations in Matrix Form:

$$\begin{aligned} x_1 + x_2 + 2x_3 &= 7 \\ 2x_1 - x_2 + x_3 &= 0.5 \\ x_1 + 4x_3 &= 7 \end{aligned} \tag{2.24}$$

Solution: Note that we treat “any missing coefficients” (as in the “missing x_2 ” in the third equation) as being zeros! This is super

important to remember.

$$\begin{aligned} x_1 + x_2 + 2x_3 &= 7 \\ 2x_1 - x_2 + x_3 &= 0.5 \\ x_1 + 4x_3 &= 7 \end{aligned} \iff \underbrace{\begin{bmatrix} 1 & 1 & 2 \\ 2 & -1 & 1 \\ 1 & \boxed{0} & 4 \end{bmatrix}}_A \underbrace{\begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}}_x = \underbrace{\begin{bmatrix} 7 \\ 0.5 \\ 7 \end{bmatrix}}_b \quad (2.25)$$

Example 2.5 We redo Example 2.4 with the equations in a different order

$$\begin{aligned} x_1 + 4x_3 &= 7 \\ x_1 + x_2 + 2x_3 &= 7 \\ 2x_1 - x_2 + x_3 &= 0.5 \end{aligned} \quad (2.26)$$

Solution: Once again, we treat “any missing coefficients” (as in the “missing x_2 ” in the first equation) as being zeros! This is super important to remember.

$$\begin{aligned} x_1 + \boxed{0x_2} + 4x_3 &= 7 \\ x_1 + x_2 + 2x_3 &= 7 \\ 2x_1 - x_2 + x_3 &= 0.5 \end{aligned} \iff \underbrace{\begin{bmatrix} 1 & \boxed{0} & 4 \\ 1 & 1 & 2 \\ 2 & -1 & 1 \end{bmatrix}}_A \underbrace{\begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}}_x = \underbrace{\begin{bmatrix} 7 \\ 7 \\ 0.5 \end{bmatrix}}_b \quad (2.27)$$

Example 2.6 (Optional Read) We redo Example 2.4 with the equations and variables in a different order

$$\begin{aligned} x_1 + x_2 + 2x_3 &= 7 \\ x_1 + 4x_3 &= 7 \\ 2x_1 - x_2 + x_3 &= 0.5 \end{aligned} \quad (2.28)$$

Solution: This time, just to drive home the point that you can place the variables in any order that you wish, we will define x by

$$x_{\text{odd}} = \begin{bmatrix} x_2 \\ x_3 \\ x_1 \end{bmatrix}.$$

With the variables x_1 , x_2 , and x_3 arranged in this deliberately strange order, the matrices become

$$\begin{aligned} x_1 + x_2 + 2x_3 &= 7 \\ x_1 + \boxed{0x_2} + 4x_3 &= 7 \\ 2x_1 - x_2 + x_3 &= 0.5 \end{aligned} \iff \underbrace{\begin{bmatrix} 1 & 2 & 1 \\ \boxed{0} & 4 & 1 \\ -1 & 1 & 2 \end{bmatrix}}_{A_{\text{odd}}} \underbrace{\begin{bmatrix} x_2 \\ x_3 \\ x_1 \end{bmatrix}}_{x_{\text{odd}}} = \underbrace{\begin{bmatrix} 7 \\ 7 \\ 0.5 \end{bmatrix}}_b \quad (2.29)$$

Re-ordering the components of x re-orders the columns of A .

How to Handle “Missing” Variables or Coefficients

A zero in row i and column j of a matrix corresponds to the variable x_j being absent from the i -th equation

$$\underbrace{\begin{bmatrix} 0 & 0 & 2 \\ 2 & 0 & 0 \\ 0 & -1 & 4 \end{bmatrix}}_A \underbrace{\begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}}_x = \underbrace{\begin{bmatrix} 7 \\ 0.5 \\ 7 \end{bmatrix}}_b \iff \begin{aligned} 2x_3 &= 7 \\ 2x_1 &= 0.5 \\ -x_2 + 4x_3 &= 7 \end{aligned} \iff \begin{aligned} 0x_1 + 0x_2 + 2x_3 &= 7 \\ 2x_1 + 0x_2 + 0x_3 &= 0.5 \\ \underbrace{0x_1 - x_2 + 4x_3}_{\text{we'd never write it like this}} &= 7 \end{aligned} \quad (2.30)$$

2.6 The Matrix Determinant

One of the more difficult topics in *Linear Algebra* is the notion of the determinant of a matrix. Most people who claim to understand it are wrong. Instead of teaching you right away a bunch of details of the matrix determinant that are rather useless in the actual practice of Linear Algebra, we are going to cut directly to the “good stuff”. Moreover, we’ll only give you part of the good stuff here, and save some “really excellent stuff” for Chapter 6.

2.6.1 First Contact with the Matrix Determinant

What you want and need to know about the determinant function: the first five points are extra important.

Enough Facts about the Determinant to Get Us Going

Fact 1 The determinant of a square matrix A is a real number, denoted $\det(A)$.

Fact 2 A square system of linear equations (i.e., n equations and n unknowns), $Ax = b$, has a unique solution x for any $n \times 1$ vector b if, and only if, $\det(A) \neq 0$.

Fact 3 When $\det(A) = 0$, the set of linear equations $Ax = b$ may have either no solution or an infinite number of solutions. To determine which case applies (and to be clear, only one case can apply), depends on how “ b is related to A ”, which is another way of saying that we will have to address this later in the course.

Fact 4 The determinant of the 1×1 matrix $A = [a]$ is $\det(A) := a$.

Fact 5 The determinant of the 2×2 square matrix $A = \begin{bmatrix} a & b \\ c & d \end{bmatrix}$ is $\det(A) := ad - bc$.

Fact 6 The determinant is only defined for square matrices; see Fact 1.

Fact 7 We will learn later how to compute by hand the determinant of $n \times n$ matrices with special structure, but for general matrices, you only need to know how to *compute by hand the determinant of a 2×2 matrix*.

Remark: In case you are “sad” that you can only compute determinants for 2×2 matrices, do not worry. *Very soon, you will also be able to do it for a 100×100 matrix and understand what you are doing.* In this book, we take a “structural approach” to Computational Linear Algebra in general, and to the matrix determinant in particular. In Chap. 3, we will first focus on matrices with a special “triangular” structure and then we’ll see how to “factor” or “decompose” a non-triangular matrix into (the product of) two triangular matrices. This structural approach has many advantages when it comes to scaling up computations to problems with hundreds or even thousands of variables. In particular, it will provide a practical way of understanding and computing the matrix determinant for really big matrices.

Remark: A “clear but advanced” derivation of the classical formula for the determinant of an $n \times n$ matrix can be found in a blog post by Akihiro Matsukawa: <https://mtskw.com/posts/determinant/>. In case the blog post gets taken down, here is a link to a PDF <https://tinyurl.com/5n7pm4x> of it.

2.6.2 Examples of Using the Determinant

We reuse some previous examples and add in a few more to illustrate Fact 2 of Sec. 2.6.1

Example 2.7 Check for existence and uniqueness of solutions:

$$\begin{aligned} x_1 + x_2 &= 4 \\ 2x_1 - x_2 &= -1 \end{aligned} \iff \underbrace{\begin{bmatrix} 1 & 1 \\ 2 & -1 \end{bmatrix}}_A \underbrace{\begin{bmatrix} x_1 \\ x_2 \end{bmatrix}}_x = \underbrace{\begin{bmatrix} 4 \\ -1 \end{bmatrix}}_b. \quad (2.31)$$

Solution: We compute $\det(A) = (1) \cdot (-1) - (1) \cdot (2) = -3 \neq 0$ and conclude that (2.31) has a unique solution. ■

Example 2.8 Check for existence and uniqueness of solutions:

$$\begin{aligned} x_1 - x_2 &= 1 \\ 2x_1 - 2x_2 &= -1 \end{aligned} \iff \underbrace{\begin{bmatrix} 1 & -1 \\ 2 & -2 \end{bmatrix}}_A \underbrace{\begin{bmatrix} x_1 \\ x_2 \end{bmatrix}}_x = \underbrace{\begin{bmatrix} 1 \\ -1 \end{bmatrix}}_b \quad (2.32)$$

Solution: We compute $\det(A) = (1) \cdot (-2) - (-1) \cdot (2) = 0$. We therefore conclude that (2.32) does not have a unique solution. In fact, it may have either no solution at all or it may have an infinite number of solutions. At this point in the course, we do not have the tools to determine which case applies here without grinding through the equations. Referring back to (1.2), where we did grind through the equations, we know that this system of linear equations has no solution. ■

Example 2.9 Check for existence and uniqueness of solutions:

$$\begin{aligned} x_1 - x_2 &= 1 \\ 2x_1 - 2x_2 &= 2 \end{aligned} \iff \underbrace{\begin{bmatrix} 1 & -1 \\ 2 & -2 \end{bmatrix}}_A \underbrace{\begin{bmatrix} x_1 \\ x_2 \end{bmatrix}}_x = \underbrace{\begin{bmatrix} 1 \\ 2 \end{bmatrix}}_b \quad (2.33)$$

Solution: We compute $\det(A) = (1) \cdot (-2) - (1) \cdot (2) = 0$. We therefore conclude that (2.33) does not have a unique solution. In fact, it may have either no solution at all or it may have an infinite number of solutions. At this point in the course, we do not have the tools to determine which case applies here without grinding through the equations. Referring back to (1.3), we know that this system of linear equations has an infinite number of solutions. ■

Example 2.10 Check for existence and uniqueness of solutions:

$$\begin{aligned} x_1 + x_2 + 2x_3 &= 7 \\ 2x_1 - x_2 + x_3 &= 0.5 \\ x_1 + 4x_3 &= 7 \end{aligned} \iff \underbrace{\begin{bmatrix} 1 & 1 & 2 \\ 2 & -1 & 1 \\ 1 & 0 & 4 \end{bmatrix}}_A \underbrace{\begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}}_x = \underbrace{\begin{bmatrix} 7 \\ 0.5 \\ 7 \end{bmatrix}}_b \quad (2.34)$$

Solution: Using Julia, we compute $\det(A) = -9 \neq 0$. We therefore conclude that (2.34) has a unique solution for x .

```
1 # using LinearAlgebra
2 A=[1 1 2; 2 -1 1; 1 0 4]
3 det(A)
```

Output

-9.0

Example 2.11 Check for existence and uniqueness of solutions in an example where the unknowns are not in the “correct” order and we have a bunch of “missing coefficients”:

$$\begin{aligned} x_1 + x_2 + 2x_3 &= 7 \\ -x_2 + x_3 + 2x_1 &= 0.5 \\ x_1 + 4x_3 &= 7 \\ x_4 + 2x_3 - 5x_5 - 11 &= 0 \\ -4x_2 + 12x_4 &= 0 \end{aligned} \iff \underbrace{\begin{bmatrix} 1 & 1 & 2 & 0 & 0 \\ 2 & -1 & 1 & 0 & 0 \\ 1 & 0 & 4 & 0 & 0 \\ 0 & 0 & 2 & 1 & -5 \\ 0 & -4 & 0 & 12 & 0 \end{bmatrix}}_A \underbrace{\begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \end{bmatrix}}_x = \underbrace{\begin{bmatrix} 7 \\ 0.5 \\ 7 \\ 11 \\ 0 \end{bmatrix}}_b \quad (2.35)$$

Solution: Using Julia, one computes $\det(A) = -540 \neq 0$. We therefore conclude that (2.35) has a unique solution. Grinding through the equations would have been no fun at all!


```
1 A=[1 1 2 0 0; 2 -1 1 0 0; 1 0 4 0 0; 0 0 2 1 -5; 0 -4 0 12 0]
2 det(A)
```

Output


-540.0

2.7 (Optional Read): Other Facts on the Matrix Determinant Covered in “Standard” Courses

Here are some other facts that are sometimes useful but are not required in ROB 101.



(a)



(b)

Figure 2.2: (a) Image borrowed from the Khan Academy. A simple electric circuit as you would learn to analyze in the first month of EECS 215 using Kirchhoff’s Current and Voltage Laws, which can be thought of as the electrical versions of Newton’s equations for balancing forces. (b) A quadratic function that you need to find based upon someone having measured for you the three values given in red! Can you do it? Does it matter that the data were selected from only one side of the parabola?

- In written mathematics, but not in programming, you will also encounter the notation $\det(A) = |A|$, so that

$$\begin{vmatrix} a & b \\ c & d \end{vmatrix} := ad - bc.$$

- In Julia, once you state that you are using the linear algebra package, that is `using LinearAlgebra`, the command is `det(A)`, if A is already defined, or, for example, `det([1 2 3; 4 5 6; 7 8 9])`
- The determinant of a 3×3 matrix is

$$\det \left(\begin{bmatrix} a & b & c \\ d & e & f \\ g & h & i \end{bmatrix} \right) := a \cdot \det \left(\begin{bmatrix} e & f \\ h & i \end{bmatrix} \right) - b \cdot \det \left(\begin{bmatrix} d & f \\ g & i \end{bmatrix} \right) + c \cdot \det \left(\begin{bmatrix} d & e \\ g & h \end{bmatrix} \right). \quad (2.36)$$

In ROB 101, you do not need to use, much less memorize, the formula (2.36).

Optional, and for sure, skip on your first read: Khan Academy

Here are the (Tiny URLs) of two videos by the Khan Academy that provide a “classical” introduction to the matrix determinant.

- <https://tinyurl.com/gqt9313> works a 3×3 example based on (2.36).
- <https://tinyurl.com/zhlwb5v> shows an alternative way to compute the determinant of a 3×3 matrix.
- **In ROB 101**, you do not need to use either of these methods.

2.8 (Optional Read): A Few “Practical” Examples Using Linear Algebra

You are not responsible for any of these examples. They are given here to provide a sense of how vectors and matrices are used in engineering. Because we are so early in the course, the examples we can work are rather boring and are not on par with what you will do in the projects.

Example 2.12 [Circuit Equations] Write the equations satisfied by the currents I_1 , I_2 , and I_3 in Fig. 2.2a and then solve them.

Solution: Once you have taken EECS 215, you would quickly write down the following equations which are based on three facts called Kirchhoff’s Current and Voltage Laws:

- the sum of the voltages around any loop in the circuit is zero;
- current I flowing through resistor R creates a voltage $V = IR$; and
- when two loops touch, as for the $25\ \Omega$ resistor, the currents add with their sign determined by their directions.

At the $25\ \Omega$ resistor, for example, the current is $I_1 - I_2$ because they flow in opposite directions. In short, after taking EECS 215 and applying Kirchhoff's Current and Voltage Laws, you would obtain

$$\begin{aligned} -10 + I_1 + 25(I_1 - I_2) + 50(I_1 - I_3) &= 0 \\ 25(I_2 - I_1) + 30I_2 + (I_2 - I_3) &= 0 \\ 55I_3 + 50(I_3 - I_1) + (I_3 - I_2) &= 0 \end{aligned} \iff \underbrace{\begin{bmatrix} 76 & -25 & -50 \\ -25 & 56 & -1 \\ -50 & -1 & 106 \end{bmatrix}}_A \underbrace{\begin{bmatrix} I_1 \\ I_2 \\ I_3 \end{bmatrix}}_x = \underbrace{\begin{bmatrix} 10 \\ 0 \\ 0 \end{bmatrix}}_b \quad (2.37)$$

While it is not important to you, the first equation is obtained by adding up the voltages in the loop with I_1 , proceeding in the direction of I_1 , the second equation is obtained by adding up the voltages in the loop with I_2 , proceeding in the direction of I_2 , and similarly for the last equation. In Julia, we compute that $\det(A) = 242,310$, kind of a wild number, but it's non-zero. Solving the equation results in

$$\begin{bmatrix} I_1 \\ I_2 \\ I_3 \end{bmatrix} = \begin{bmatrix} 0.2449 \\ 0.1114 \\ 0.1166 \end{bmatrix},$$

in units of Amperes in case you are curious! ■

Example 2.13 [Fitting a Function to Data] Figure 2.2b shows a quadratic function of the form

$$y = a_2x^2 + a_1x + a_0. \quad (2.38)$$

You are given the data in Table 2.1 and need to find the coefficients a_2 , a_1 , and a_0 that define the quadratic.

Table 2.1: Data for a Quadratic Function.

x	y
1	6
2	11
3	18

Solution: Our unknowns are the coefficients a_2 , a_1 , and a_0 . Hence, we rewrite the quadratic as

$$y = a_2x^2 + a_1x + a_0 = x^2a_2 + xa_1 + a_0,$$

and note that this equation is linear in the unknowns. From Table 2.1, we have the following equations

$$\begin{aligned} 6 &= a_2 + a_1 + a_0 \\ 11 &= 4a_2 + 2a_1 + a_0 \\ 18 &= 9a_2 + 3a_1 + a_0 \end{aligned} \iff \underbrace{\begin{bmatrix} 1 & 1 & 1 \\ 4 & 2 & 1 \\ 9 & 3 & 1 \end{bmatrix}}_A \underbrace{\begin{bmatrix} a_2 \\ a_1 \\ a_0 \end{bmatrix}}_x = \underbrace{\begin{bmatrix} 6 \\ 11 \\ 18 \end{bmatrix}}_b \quad (2.39)$$

Using Julia, we compute $\det(A) = -2 \neq 0$ and hence a solution to (2.39) exists and is unique. Using our current methods, we compute that the answer is

$$\begin{bmatrix} a_2 \\ a_1 \\ a_0 \end{bmatrix} = \begin{bmatrix} 1.0 \\ 2.0 \\ 3.0 \end{bmatrix}.$$

In other words, the data are compatible with the function $y = x^2 + 2x + 3$. ■

Example 2.14 [Fitting a Function to Data: Take 2] We re-visit the quadratic function in Fig. 2.2b with unknown coefficients a_2 , a_1 , and a_0 , namely $y = a_2x^2 + a_1x + a_0$. This time, however, we use the data in Table 2.2

Table 2.2: Data for a Quadratic Function.

x	y
-3	6
1	6
2	11
3	18

Solution: Our unknowns are once again the coefficients a_2 , a_1 , and a_0 . We rewrite the quadratic as

$$y = a_2x^2 + a_1x + a_0 = x^2a_2 + xa_1 + a_0,$$

and note again that this equation is linear in the unknowns. From Table 2.2, we have the following equations

$$\begin{aligned} 6 &= 9a_2 - 3a_1 + a_0 \\ 6 &= a_2 + a_1 + a_0 \\ 11 &= 4a_2 + 2a_1 + a_0 \\ 18 &= 9a_2 + 3a_1 + a_0 \end{aligned} \iff \underbrace{\begin{bmatrix} 9 & -3 & 1 \\ 1 & 1 & 1 \\ 4 & 2 & 1 \\ 9 & 3 & 1 \end{bmatrix}}_A \underbrace{\begin{bmatrix} a_2 \\ a_1 \\ a_0 \end{bmatrix}}_x = \underbrace{\begin{bmatrix} 6 \\ 6 \\ 11 \\ 18 \end{bmatrix}}_b \quad (2.40)$$

In this case, we have more equations than unknowns, which translates into A being non-square. Because A is non-square, we cannot compute its determinant and then use the answer to determine whether solutions exist or are unique, which is kind of a bummer, because it kind of says that more data is not necessarily better! Later in the course, we will be able to confront this issue head on and learn that, yes, in most cases, more data is better. For now, we punt! ■

2.9 (Optional Read) Yet Another Determinant Example

In the following examples, the determinant is computed using the method of (2.36). **You are not required to know this method because you would never want to use it on a 10×10 matrix, for example**, whereas we will learn a method that scales easily to matrices that are 100×100 . **The method in (2.36) is only illustrated here to prove a point: it's painful.**

Example 2.15 Compute the determinant of the matrix below using the method that is normally taught in Linear Algebra, namely, (2.36).

$$A = \begin{bmatrix} 4 & -1 & 1 \\ 4 & 5 & 3 \\ -2 & 0 & 0 \end{bmatrix}$$

Solution:

$$\begin{aligned} \det(A) &= \begin{vmatrix} 4 & -1 & 1 \\ 4 & 5 & 3 \\ -2 & 0 & 0 \end{vmatrix} \\ &= 4 \cdot \begin{vmatrix} 5 & 3 \\ 0 & 0 \end{vmatrix} - (-1) \cdot \begin{vmatrix} 4 & 3 \\ -2 & 0 \end{vmatrix} + 1 \cdot \begin{vmatrix} 4 & 5 \\ -2 & 0 \end{vmatrix} \\ &\quad \text{(after applying our determinant formula to each of the } 2 \times 2 \text{ matrices above)} \\ &= 4 \cdot (0) + 1 \cdot (6) + 1 \cdot 10 \\ &= 16. \end{aligned}$$

For 3×3 it's not so bad, but already at 4×4 , it becomes tedious. ■

Example 2.16 Compute the determinant of the matrix below using the method that is normally taught in Linear Algebra.

$$A = \begin{bmatrix} 4 & -1 & 1 & 2 \\ 4 & 5 & 3 & 7 \\ -2 & 0 & 0 & 4 \\ -2 & 8 & 1 & 4 \end{bmatrix}$$

Solution:

$$\begin{aligned} \det(A) &= \begin{vmatrix} 4 & -1 & 1 & 2 \\ 4 & 5 & 3 & 7 \\ -2 & 0 & 0 & 4 \\ -2 & 8 & 1 & 4 \end{vmatrix} \\ &= 4 \cdot \begin{vmatrix} 5 & 3 & 7 \\ 0 & 0 & 4 \\ 8 & 1 & 4 \end{vmatrix} - (-1) \begin{vmatrix} 4 & 3 & 7 \\ -2 & 0 & 4 \\ -2 & 1 & 4 \end{vmatrix} + (1) \begin{vmatrix} 4 & 5 & 7 \\ -2 & 0 & 4 \\ -2 & 8 & 4 \end{vmatrix} - (2) \begin{vmatrix} 4 & 5 & 3 \\ -2 & 0 & 0 \\ -2 & 8 & 1 \end{vmatrix} \\ &= \text{(after applying our determinant formula to each of the } 3 \times 3 \text{ matrices above)} \\ &= 4(76) + 1(-30) + 1(-240) - 2(-38) \\ &= 110. \end{aligned}$$

Once again, the point is not to be understand why this method works or how to do it. The point is how unwieldy it is. Since you have read this far, we'll let you in on a secret. In Chapter 5, we'll learn how to "factor" or "decompose" the matrix A above into two matrices called

$$L = \begin{bmatrix} 1.0000 & 0.0000 & 0.0000 & 0.0000 \\ -0.5000 & 1.0000 & 0.0000 & 0.0000 \\ 1.0000 & 0.8000 & 1.0000 & 0.0000 \\ -0.5000 & -0.0667 & 0.7500 & 1.0000 \end{bmatrix} \text{ and } U = \begin{bmatrix} 4.0 & -1.0 & 1.0 & 2.0 \\ 0.0 & 7.5 & 1.5 & 5.0 \\ 0.0 & 0.0 & 0.8 & 1.0 \\ 0.0 & 0.0 & 0.0 & 55/12 \end{bmatrix}.$$

In Chapter 6, we'll learn that all of the information concerning the determinant is contained in U and that $\det(A) = \det(U) = 4 \times 7.5 \times 0.8 \times 55/12 = 110.0$, the product of the terms on the diagonal of U . How mind blowing is that?

2.10 Looking Ahead

We want to solve very large sets of linear equations, say more than 100 variables. We could teach you the matrix inverse command in Julia, but then you'd understand nothing. In the next chapter, we will begin exploring how to solve problems with special structure. Soon after that, we'll show how to transform all linear systems of n -equations in n -variables to a form where they are solvable with "special structure". To get there we need to understand:

- what are *square and triangular matrices*;
- what is *forward and back substitution*;
- what does it mean to multiply two matrices; and
- how to *factor* a square matrix as the product of two triangular matrices.

Help! Help! How am I supposed to remember all of this?

You probably can't. In any case, we don't want you to memorize the ROB 101 material. Instead, open up a `google doc` or `google sheet` and make notes! You need an organized method for keeping track of stuff. In High School, you may have been able to remember all the new notation without any special effort. In College, it's a bit different.

Chapter 3

Triangular Systems of Equations: Forward and Back Substitution

Learning Objectives

- Equations that have special structure are often much easier to solve
- Some examples to show this.

Outcomes

- Recognize triangular systems of linear equations and distinguish those with a unique answer.
- Learn that the determinant of a square triangular matrix is the product of the terms on its diagonal.
- How to use forward and back substitution.
- Swapping rows of equations and permutation matrices.

3.1 Background

- A system of linear equations is **square** if it has the same number of equations as unknowns.
- If the system of linear equations is expressed in the form $Ax = b$, then it is **square** if, and only if, A has the same number of rows as it has columns; that is, A is $n \times n$.
- A square system of linear equations $Ax = b$, has a **unique solution** x for any $n \times 1$ vector b if, and only if, the $n \times n$ matrix A satisfies $\det(A) \neq 0$.
- Computing the **determinant** of a general square matrix is tedious.
- The **diagonal** of the square matrix

$$A = \begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1n} \\ a_{21} & a_{22} & \cdots & a_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ a_{n1} & a_{n2} & \cdots & a_{nn} \end{bmatrix}$$

is

$$\text{diag}(A) = [a_{11} \quad a_{22} \quad \dots \quad a_{nn}].$$

The diagonal consists of all elements a_{ij} of A such that $j = i$.

- The elements of A that are **not on the diagonal** are called the **off-diagonal elements**.

3.2 A Warm-up with Diagonal Systems of Linear Equations

As a warm-up to the main topic, we want to illustrate that some systems of linear equations are easy to solve, independently of the number of unknowns. We'll start small, and then go big!

This is an example of a square system of linear equations that is *Diagonal*,

$$\begin{aligned} 4x_1 &= 6 \\ -x_2 &= 7 \\ 3x_3 &= 2. \end{aligned} \tag{3.1}$$

Computing the solution of the set of diagonal equations is trivial because we just have to divide each row of the equation by the element multiplying the unknown, namely

$$\begin{aligned} 4x_1 &= 6 & x_1 &= \frac{6}{4} = \frac{3}{2} \\ -x_2 &= 7 & \iff x_2 &= \frac{7}{-1} = -7 \\ 3x_3 &= 2. & x_3 &= \frac{2}{3}. \end{aligned} \tag{3.2}$$

In this example, all of the constants multiplying the unknowns were non-zero. If one of them had been zero, for example, if the second row were $0x_2 = 7$, then the equations would not have a solution; if the second row were $0x_2 = 0$, then the equations would have an infinite number of solutions because x_2 could take on any value.

When we write the system as $Ax = b$, we have

$$\begin{aligned} 4x_1 &= 6 \\ -x_2 &= 7 \\ 3x_3 &= 2. \end{aligned} \iff \underbrace{\begin{bmatrix} 4 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 3 \end{bmatrix}}_A \underbrace{\begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}}_x = \underbrace{\begin{bmatrix} 6 \\ 7 \\ 2 \end{bmatrix}}_b. \tag{3.3}$$

where we note that all of the **off-diagonal terms** of the matrix A are zero. More precisely, the condition is $a_{ij} = 0$ for all $i \neq j$. Such matrices are called *Diagonal Matrices*.

The corresponding system of linear equations is easy to solve because each of the coefficients $a_{11} = 4$, $a_{22} = -1$, and $a_{33} = 3$ multiplying the unknowns x_1 , x_2 , and x_3 is non-zero. For later use, we note that n real numbers $a_{11}, a_{22}, \dots, a_{nn}$ are all non-zero if, and only if, their product is non-zero, that is,

$$a_{11} \neq 0, a_{22} \neq 0, \dots, a_{nn} \neq 0 \iff a_{11}a_{22} \cdots a_{nn} \neq 0.$$

This result is true because the product of any two real numbers is zero if, and only if, at least one of the numbers is zero.

More generally, a square $n \times n$ matrix A where all of the **off-diagonal** elements are zero is said to be a **diagonal matrix**,

$$A = \begin{bmatrix} a_{11} & 0 & \cdots & 0 \\ 0 & a_{22} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & a_{nn} \end{bmatrix};$$

in other words, if $a_{ij} = 0$ for all $i \neq j$, then A is diagonal. Computing the **determinant** of a diagonal matrix is very easy,

$$\det(A) = \begin{vmatrix} a_{11} & 0 & \cdots & 0 \\ 0 & a_{22} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & a_{nn} \end{vmatrix} = a_{11}a_{22} \cdots a_{nn},$$

the product of the terms on its diagonal.

A diagonal system of n linear equations

$$\begin{aligned} a_{11}x_1 &= b_1 \\ a_{22}x_2 &= b_2 \\ \vdots &= \vdots \\ a_{nn}x_n &= b_n \end{aligned} \iff \underbrace{\begin{bmatrix} a_{11} & 0 & \cdots & 0 \\ 0 & a_{22} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & a_{nn} \end{bmatrix}}_A \underbrace{\begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix}}_x = \underbrace{\begin{bmatrix} b_1 \\ b_2 \\ \vdots \\ b_n \end{bmatrix}}_b \quad (3.4)$$

is straightforward to solve, no matter its size, because, whenever $a_{ii} \neq 0$, we have

$$x_i = \frac{b_i}{a_{ii}}. \quad (3.5)$$

If one of the coefficients a_{ii} on the diagonal is zero, then, if the corresponding element b_i is non-zero, we deduce that the equations do not have a solution, and if the corresponding element b_i is zero, we deduce that the equations can have an infinite number of solutions. This is identical to the discussion below (3.2).

For diagonal matrices, it follows that if $\det(A) \neq 0$, then $Ax = b$ has a unique solution; moreover the solution is very easy to compute via (3.5). In the next Section, we explore other types of matrices that make computing a solution very fast and easy.

3.3 Lower Triangular Systems of Linear Equations

This is an example of a square system of linear equations that is *Lower Triangular*

$$\begin{aligned} 3x_1 &= 6 \\ 2x_1 - x_2 &= -2 \\ x_1 - 2x_2 + 3x_3 &= 2. \end{aligned} \quad (3.6)$$

When we write the system as $Ax = b$, in the lower triangular case we have

$$\begin{aligned} 3x_1 &= 6 \\ 2x_1 - x_2 &= -2 \\ x_1 - 2x_2 + 3x_3 &= 2 \end{aligned} \iff \underbrace{\begin{bmatrix} 3 & 0 & 0 \\ 2 & -1 & 0 \\ 1 & -2 & 3 \end{bmatrix}}_A \underbrace{\begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}}_x = \underbrace{\begin{bmatrix} 6 \\ -2 \\ 2 \end{bmatrix}}_b. \quad (3.7)$$

where we note that all terms “above” the diagonal of the matrix A are zero. More precisely, the condition is $a_{ij} = 0$ for all $j > i$. Such matrices are called *lower-triangular*.

Here are two more examples of square lower triangular systems

$$\begin{aligned} 3x_1 &= 6 \\ 2x_1 - x_2 &= 0 \end{aligned} \iff \underbrace{\begin{bmatrix} 3 & 0 \\ 2 & -1 \end{bmatrix}}_A \underbrace{\begin{bmatrix} x_1 \\ x_2 \end{bmatrix}}_x = \underbrace{\begin{bmatrix} 6 \\ 0 \end{bmatrix}}_b \quad (3.8)$$

$$\begin{aligned} 3x_1 &= 6 \\ 2x_2 &= -2 \\ x_1 - 2x_2 &= 2 \\ x_1 - x_3 + 2x_4 &= 10 \end{aligned} \iff \underbrace{\begin{bmatrix} 3 & 0 & 0 & 0 \\ 0 & 2 & 0 & 0 \\ 1 & -2 & 0 & 0 \\ 1 & 0 & -1 & 2 \end{bmatrix}}_A \underbrace{\begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{bmatrix}}_x = \underbrace{\begin{bmatrix} 6 \\ -2 \\ 2 \\ 10 \end{bmatrix}}_b \quad (3.9)$$

The following systems of linear equations are square, but they are not lower triangular. The “offending term” is boxed,

$$\begin{aligned} 3x_1 &= 6 \\ 2x_1 - x_2 - \boxed{x_3} &= -2 \\ x_1 - 2x_2 + 3x_3 &= 2 \end{aligned} \iff \underbrace{\begin{bmatrix} 3 & 0 & 0 \\ 2 & -1 & \boxed{-1} \\ 1 & -2 & 3 \end{bmatrix}}_A \underbrace{\begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}}_x = \underbrace{\begin{bmatrix} 6 \\ -2 \\ 2 \end{bmatrix}}_b \quad (3.10)$$

and

$$\begin{aligned} 3x_1 + \boxed{3x_2} &= 6 \\ 2x_1 - 2x_2 &= 0 \end{aligned} \iff \underbrace{\begin{bmatrix} 3 & \boxed{3} \\ 2 & -1 \end{bmatrix}}_A \underbrace{\begin{bmatrix} x_1 \\ x_2 \end{bmatrix}}_x = \underbrace{\begin{bmatrix} 6 \\ 0 \end{bmatrix}}_b \quad (3.11)$$

Triangular Matrices

The key ingredients of a square lower triangular system are

- The unknowns are ordered, as in $(x_1 \ x_2 \ \dots \ x_n)$ or $(u \ v \ w \ x \ y \ z)$
- The first equation only involves the first unknown.
- The second equation involves only the first two unknowns
- More generally, the i -th equation involves only the first i unknowns.
- The coefficients a_{ij} on or below the diagonal can be zero or non-zero.

From a matrix point of view, the condition is, all terms above the diagonal are zero. What does this mean? It means that A looks like this,

$$A = \begin{bmatrix} \mathbf{a_{11}} & 0 & 0 & 0 & \cdots & 0 & 0 \\ a_{21} & \mathbf{a_{22}} & 0 & 0 & \cdots & 0 & 0 \\ a_{31} & a_{32} & \mathbf{a_{33}} & 0 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \ddots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & 0 & 0 \\ a_{(n-1)1} & a_{(n-1)2} & a_{(n-1)3} & a_{(n-1)4} & \cdots & \mathbf{a_{(n-1)(n-1)}} & 0 \\ a_{n1} & a_{n2} & a_{n3} & a_{n4} & \cdots & a_{n(n-1)} & \mathbf{a_{nn}} \end{bmatrix}, \quad (3.12)$$

namely, $a_{ij} = 0$ for $j > i$. Note that the diagonal is in bold and the terms above the diagonal are in red. The matrices in (3.7) through (3.9) are *lower triangular*, while the matrices in (3.10) and (3.11) are not lower triangular.

3.4 Determinant of a Lower Triangular Matrix

Highly Useful Fact: The matrix determinant of a square lower triangular matrix is equal to the product of the elements on the diagonal. For the matrix A in (3.12), its determinant is

$$\det(A) = a_{11} \cdot a_{22} \cdot \dots \cdot a_{nn}. \quad (3.13)$$

Hence, for lower triangular matrices of size 10×10 or so, the determinant can be computed by inspection! Let's do a few:

$$\det \left(\begin{bmatrix} 3 & 0 & 0 \\ 2 & -1 & 0 \\ 1 & -2 & 3 \end{bmatrix} \right) = 3 \cdot (-1) \cdot 3 = -9 \neq 0.$$

Hence, the system of equations (3.7) has a unique solution.

Let's now use the alternative notation for the determinant of a matrix,

$$\begin{vmatrix} 3 & 0 & 0 & 0 \\ 0 & 2 & 0 & 0 \\ 1 & -2 & 0 & 0 \\ 1 & 0 & -1 & 2 \end{vmatrix} = 3 \cdot 2 \cdot 0 \cdot 2 = 0.$$

Hence, the system of equations (3.9) is one of the problem cases: it may have no solution or it may have an infinite number of solutions.

Remark: For a square lower triangular matrix, the matrix determinant is nonzero if, and only if, all of the elements on the diagonal of the matrix are non-zero. Equivalently, for a square lower triangular matrix, the matrix determinant is zero if, and only if, at least one of the elements on the diagonal of the matrix is zero. In other words, we do not really need to multiply them out to check for $\det(A) \neq 0$, the condition for uniqueness of solutions of square systems of linear equations.

3.5 Lower Triangular Systems and Forward Substitution

We develop a solution to a lower triangular system by starting at the top and working our way to the bottom via a method called **forward substitution**. As an example, we use (3.7), which for convenience, we repeat here:

$$\begin{aligned} 3x_1 &= 6 \\ 2x_1 - x_2 &= -2 \\ x_1 - 2x_2 + 3x_3 &= 2 \end{aligned} \iff \underbrace{\begin{bmatrix} 3 & 0 & 0 \\ 2 & -1 & 0 \\ 1 & -2 & 3 \end{bmatrix}}_A \underbrace{\begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}}_x = \underbrace{\begin{bmatrix} 6 \\ -2 \\ 2 \end{bmatrix}}_b. \quad (3.14)$$

Because we have ordered the variables as $(x_1 \ x_2 \ x_3)$, we isolate x_1 in the first equation, x_2 in the second equation, and x_3 in the third equation by moving the other variables to the right-hand side,

$$\begin{aligned} 3x_1 &= 6 \\ -x_2 &= -2 - 2x_1 \\ 3x_3 &= 2 - x_1 + 2x_2. \end{aligned} \quad (3.15)$$

You can see how the first equation is trivial, and so is the second one, once the first is solved, etc. Next, we can make the coefficients of the leading variables all equal to one by dividing through by the coefficients that multiply them, yielding

$$\begin{aligned} x_1 &= \frac{1}{3}6 = 2 \\ x_2 &= -[-2 - 2x_1] = 2 + 2x_1 \\ x_3 &= \frac{1}{3}[2 - x_1 + 2x_2]. \end{aligned} \quad (3.16)$$

Next, we substitute in, working from top to bottom:

$$\begin{aligned}x_1 &= \frac{1}{3}6 = 2 \\x_2 &= 2 + 2x_1 = 6 \\x_3 &= \frac{1}{3}[2 - x_1 + 2x_2] = \frac{1}{3}[12] = 4,\end{aligned}\tag{3.17}$$

Forward Substitution

The method is called **forward substitution** because once we have solved x_1 , we take its value forward to the next equation where we solve for x_2 , and once we have solved for x_1 and x_2 , we take their values forward to the next equation where we solve for x_3 . I am guessing that you see the pattern.

When can forward substitution go wrong? Well first of all, we only use it for lower triangular systems of linear equations. If the diagonal of the matrix corresponding to the system of linear equations has a zero on it, then the matrix determinant is zero (\implies no solution or an infinite number of solutions) and forward substitution would lead us to **divide by zero, which we know is a major error in mathematics**.

As an example, we take (3.9), which we repeat here

$$\begin{aligned}3x_1 &= 6 \\2x_2 &= -2 \\x_1 - 2x_2 &= 2 \\x_1 - x_3 + 2x_4 &= 10.\end{aligned}\iff \underbrace{\begin{bmatrix} 3 & 0 & 0 & 0 \\ 0 & 2 & 0 & 0 \\ 1 & -2 & 0 & 0 \\ 1 & 0 & -1 & 2 \end{bmatrix}}_A \underbrace{\begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{bmatrix}}_x = \underbrace{\begin{bmatrix} 6 \\ -2 \\ 2 \\ 10 \end{bmatrix}}_b\tag{3.18}$$

When we try to isolate x_3 in the third equation, we have a problem. To make it more obvious, we make x_3 explicit with its corresponding coefficient of zero

$$\left. \begin{aligned}3x_1 &= 6 \\2x_2 &= -2 \\x_1 - 2x_2 + 0x_3 &= 2 \\x_1 - x_3 + 2x_4 &= 10\end{aligned} \right\} \implies \left. \begin{aligned}3x_1 &= 6 \\2x_2 &= -2 \\0x_3 &= 2 - (x_1 - 2x_2) \\2x_4 &= 10 - (x_1 - x_3)\end{aligned} \right\} \begin{aligned}x_1 &= 2 \\x_2 &= -1 \\x_3 &= \frac{1}{0} \cdot (-2) \text{ ??? Divide by zero: not allowed!} \\x_4 &= \frac{1}{4}(10 - 2 + \frac{-2}{0}) \text{ !?!? Wrong!}\end{aligned}\tag{3.19}$$

3.6 Upper Triangular Systems, Upper Triangular Matrices, Determinants, and Back Substitution

This is an example of a square **upper triangular system of equations** and its corresponding **upper triangular matrix**

$$\begin{aligned}x_1 + 3x_2 + 2x_3 &= 6 \\2x_2 + x_3 &= -2 \\3x_3 &= 4,\end{aligned}\iff \underbrace{\begin{bmatrix} 1 & 3 & 2 \\ 0 & 2 & 1 \\ 0 & 0 & 3 \end{bmatrix}}_A \underbrace{\begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}}_x = \underbrace{\begin{bmatrix} 6 \\ -2 \\ 4 \end{bmatrix}}_b.\tag{3.20}$$

We note that all of the coefficients of A below the diagonal are zero; that is $a_{ij} = 0$ for $i > j$.

Highly Useful Fact: The matrix determinant of a square upper triangular matrix is equal to the product of the elements on the diagonal. For the matrix A in (3.20), its determinant is

$$\det(A) = 1 \cdot 2 \cdot 3 = 6.\tag{3.21}$$

Hence, we know the system of equations has a unique solution.

Back Substitution

We develop a solution to an **upper triangular system of linear equations** by starting at the bottom and work our way to the top via **back substitution**. We first solve for the leading variables, which here we do in one step,

$$\begin{aligned}x_1 &= 6 - (3x_2 + 2x_3) \\x_2 &= \frac{1}{2}(-2 - x_3) \\x_3 &= \frac{4}{3},\end{aligned}\tag{3.22}$$

but of course, you can do it in two steps as we did for lower triangular systems. Next, we do back substitution, from bottom to top, sequentially plugging in numbers from the previous equations

$$\begin{aligned}x_1 &= 6 - (3x_2 + 2x_3) = 6 - (3 \cdot (-\frac{5}{3}) + 2 \cdot \frac{4}{3}) = \frac{18}{3} + \frac{7}{3} = \frac{25}{3} = 8\frac{1}{3} \\x_2 &= \frac{1}{2} \cdot (-2 - x_3) = \frac{1}{2} \cdot (-2 - \frac{4}{3}) = \frac{1}{2} \cdot (-\frac{10}{3}) = -\frac{5}{3} = -1\frac{2}{3} \\x_3 &= \frac{4}{3}.\end{aligned}\tag{3.23}$$

3.7 General Cases

The general form of a lower triangular system with a non-zero determinant is

$$\begin{aligned}a_{11}x_1 &= b_1 \quad (a_{11} \neq 0) \\a_{21}x_1 + a_{22}x_2 &= b_2 \quad (a_{22} \neq 0) \\&\vdots \\a_{n1}x_1 + a_{n2}x_2 + a_{n3}x_3 + \cdots + a_{nn}x_n &= b_n \quad (a_{nn} \neq 0)\end{aligned}\tag{3.24}$$

and the solution proceeds from top to bottom, like this

$$\begin{aligned}x_1 &= \frac{b_1}{a_{11}} \quad (a_{11} \neq 0) \\x_2 &= \frac{b_2 - a_{21}x_1}{a_{22}} \quad (a_{22} \neq 0) \\&\vdots \\x_n &= \frac{b_n - a_{n1}x_1 - a_{n2}x_2 - \cdots - a_{n,n-1}x_{n-1}}{a_{nn}} \quad (a_{nn} \neq 0).\end{aligned}\tag{3.25}$$

The general form of an upper triangular system with a non-zero determinant is

$$\begin{aligned}a_{11}x_1 + a_{12}x_2 + a_{13}x_3 + \cdots + a_{1n}x_n &= b_1 \quad (a_{11} \neq 0) \\a_{22}x_2 + a_{23}x_3 + \cdots + a_{2n}x_n &= b_2 \quad (a_{22} \neq 0) \\a_{33}x_3 + \cdots + a_{3n}x_n &= b_3 \quad (a_{33} \neq 0) \\&\vdots \\a_{nn}x_n &= b_n \quad (a_{nn} \neq 0),\end{aligned}\tag{3.26}$$

and the solution proceeds from the **bottom** of (3.26) to its **top**, like this,

$$\begin{aligned}
 x_n &= \frac{b_n}{a_{nn}} && (a_{nn} \neq 0) \\
 x_{n-1} &= \frac{b_{n-1} - a_{n-1,n}x_n}{a_{n-1,n-1}} && (a_{n-1,n-1} \neq 0) \\
 &\vdots && \vdots \\
 x_2 &= \frac{b_2 - a_{23}x_3 - \cdots - a_{2n}x_n}{a_{22}} && (a_{22} \neq 0) \\
 x_1 &= \frac{b_1 - a_{12}x_2 - \cdots - a_{1n}x_n}{a_{11}} && (a_{11} \neq 0)
 \end{aligned} \tag{3.27}$$

Remark: For $1 \leq i < n$, we need to form the product

$$a_{i,i+1}x_{i+1} + a_{i,i+2}x_{i+2} + \cdots + a_{i,n}x_n.$$

In Julia, there are two ways to do it

```

1 method1 = A[i, i+1 : n]' * x[i+1 : n]
2 method2 = (A[i : i, i+1 : n] * x[i+1 : n])[1]

```

In method one, $A[i, i+1 : n]$ is a column vector while $A[i, i+1 : n]'$ is a row vector. When it is multiplied by the column vector, $x[i+1 : n]$, it produces a scalar. In method two, $A[i : i, i+1 : n]$ is a $1 \times (n-i)$ matrix, which in Julia, is different than a row vector. When it is multiplied by the column $x[i+1 : n]$, it produces a 1×1 matrix. You then have to extract its value, which is done via $(A[i : i, i+1 : n] * x[i+1 : n])[1]$. See the “tinyMatrix” example in Lab #1.

For those of you with a “visual memory”, here is a graphical representation for upper and lower triangular matrices

$$\text{(lower triangular)} \begin{bmatrix} \times & & & & \\ \times & \times & & & \\ \times & \times & \times & & \\ \times & \times & \times & \times & \\ \times & \times & \times & \times & \times \end{bmatrix} \begin{matrix} \\ \\ \mathbf{0} \\ \\ \end{matrix} \qquad \begin{bmatrix} \times & \times & \times & \times & \times \\ & \times & \times & \times & \times \\ & & \times & \times & \times \\ & & & \times & \times \\ \mathbf{0} & & & & \times \end{bmatrix} \text{(upper triangular)}. \tag{3.28}$$

Lower: everything above the diagonal is zero. **Upper:** everything below the diagonal is zero. For us to be able to solve the equations for arbitrary values b on the right-hand side of $Ax = b$, we need the elements on the diagonal to be non-zero.

Example 3.1 Use back substitution to solve the upper triangular system of linear equations $Ux = b$, where

$$U = \begin{bmatrix} \mathbf{0.9555} & -0.8218 & -1.2433 & -0.5536 & 0.9102 & 1.2047 \\ 0.0000 & \mathbf{-0.2728} & 0.3770 & 2.0805 & -1.1050 & 1.0576 \\ 0.0000 & 0.0000 & \mathbf{0.2126} & 1.0730 & -1.3323 & 2.3487 \\ 0.0000 & 0.0000 & 0.0000 & \mathbf{-0.2295} & 0.9807 & 0.3360 \\ 0.0000 & 0.0000 & 0.0000 & 0.0000 & \mathbf{-1.2425} & -1.5521 \\ 0.0000 & 0.0000 & 0.0000 & 0.0000 & 0.0000 & \mathbf{-0.7935} \end{bmatrix} \text{ and } b = \begin{bmatrix} 0.8399 \\ -0.8898 \\ 0.0069 \\ -1.1286 \\ -0.0115 \\ -1.1136 \end{bmatrix}.$$

The diagonal of U is in bold font. We note that U is indeed upper triangular because all of its elements below the diagonal are zero. Moreover, $\det(U) \neq 0$ because all of the elements on the diagonal are non-zero.

Solution: We first program a function in Julia that implements back substitution. In HW you will develop Julia code to solve triangular systems of equations. Your data will be given to you sometimes as systems of equations written out as formulas and sometimes directly as matrices. The function below checks that there are no “tiny” entries on the diagonal of U , but it does not check that U is really upper triangular. In HW, we’ll have you do that check.

```

1 function backwardsub(U, b)
2     # U a square upper triangular matrix

```



```

3 # b has same number of rows as U
4 #
5 # Assert no entries on the diagonal of U
6 # are 0 (or very close to 0)
7 if minimum(abs.(diag(U))) < 1e-6
8     return false
9 end
10 n = length(b)
11 x = Vector{Float64}(undef, n)
12 # Start from the bottom and work our way to the top
13 x[n] = b[n] / U[n,n]
14 for i = n-1:-1:1
15     #x[i]=(b[i] - U[i,(i+1):n]' * x[(i+1):n]) / U[i,i]
16     x[i]=( b[i] - (U[i:i, (i+1):n] * x[(i+1):n])[1] ) / U[i,i]
17 end
18 # The for loop works with either line 15 or line 16. They differ in how
19 # a row is extracted from a matrix in Julia
20 return x
21 end
22
23 U=[ 0.955467  -0.821842  -1.24331  -0.553594  0.910181  1.20471
24     0.0        -0.272776  0.376981  2.08047  -1.10505  1.05765
25     0.0         0.0        0.212559  1.07301  -1.33234  2.3487
26    -0.0         0.0         0.0       -0.229487  0.980719  0.336002
27    -0.0         0.0        -0.0       0.0       -1.24249  -1.55205
28    -0.0         0.0         0.0       0.0       -0.0     -0.793501]
29
30 b=[0.8398952455773964
31    -0.8897505302659705
32     0.006884706336738545
33    -1.1285718398040936
34    -0.011546427596053652
35    -1.1135689635657877]
36
37 x=backwardsub(U, b)

```

```

6-element Vector{Float64}:
-48.810775767214956
-21.035398066599555
-23.98466204387721
-0.47925601967730014
-1.7437077113383561
 1.40336197583662

```

```
1 U*x-b
```

```

6-element Vector{Float64}:
-7.771561172376096e-16
-8.881784197001252e-16
 2.211772431870429e-16
 0.0
 3.122502256758253e-17
 0.0

```

It looks like the computed x is a pretty good solution of $Ux = b$! ■

3.8 A Simple Trick with Systems of Equations: Re-arranging their Order

This system of equations is neither upper triangular nor lower triangular

$$\begin{aligned}3x_1 &= 6 \\x_1 - 2x_2 + 3x_3 &= 2 \\2x_1 - x_2 &= -2.\end{aligned}\tag{3.29}$$

We can simply re-arrange the order of the equations to arrive at

$$\begin{aligned}3x_1 &= 6 \\2x_1 - x_2 &= -2 \\x_1 - 2x_2 + 3x_3 &= 2,\end{aligned}\tag{3.30}$$

which happens to be lower triangular. We still have the same equations and hence their solution has not changed. The equations have just been re-arranged to make the process of computing their solution fall into a nice pattern that we already know how to handle.

This is a really useful trick in mathematics: re-arranging a set of equations without changing the answer. Right here, it may seem kind of trivial, but when we look at this in terms of matrices in the next Chapter, we get a cool piece of insight.

3.9 Looking Ahead

Once again, our goal is to solve very large sets of linear equations, say more than 100 variables. In this chapter, we saw how to solve problems with triangular structure. Our next major way point is to reduce all linear systems of n -equations in n -variables to being solvable with forward substitution and back substitution. To get there we need to understand:

- the standard way to multiply two matrices
- a little known way to multiply two matrices
- how to *factor* a square matrix as the product of two triangular matrices

Help! Help! How am I supposed to remember all of this?

You probably can't. In any case, we don't want you to memorize the ROB 101 material. Instead, open up a `google doc` or `google sheet` and make notes! You need an organized method for keeping track of stuff. In High School, you may have been able to remember all the new notation without any special effort. In College, it's a bit different.

Chapter 4

Matrix Multiplication

Learning Objectives

- How to partition matrices into rows and columns
- How to multiply two matrices
- How to swap rows of a matrix

Outcomes

- Multiplying a row vector by a column vector.
- Recognizing the rows and columns of a matrix
- Standard definition of matrix multiplication $A \cdot B$ using the rows of A and the columns of B
- Size restrictions when multiplying matrices
- Examples that work and those that don't because the sizes are wrong
- Introduce a second way of computing the product of two matrices using the columns of A and the rows of B . This will later be used to compute the LU decomposition in a very simple way.
- Permutation matrices

4.1 Multiplying a Row Vector by a Column Vector

The summation symbol

If you have not used it before, then the best way to learn it is via examples.

- $1 + 2 = \sum_{k=1}^2 k$. Here, k is called an index and \sum is the symbol for **sum or summation**. $\sum_{k=1}$ gives the initial value of the index in the sum, which in this case is 1, and \sum^2 gives the final value of the index in the sum, which in this case, is 2.
- $1 + 2 + 3 = \sum_{k=1}^3 k$
- $1 + 2 + \dots + n = \sum_{k=1}^n k$. We note that $\sum_{k=1}$ defines the initial value of the index in the sum to be 1, and \sum^n defines the final value of the index in the sum to be n .
- Changing the name of the index from k to i , for example, does not change anything

$$\sum_{i=1}^n i = 1 + 2 + \dots + n = \sum_{k=1}^n k$$

- $a_1 + a_2 + a_3 = \sum_{i=1}^3 a_i$
- $a_7 + a_8 + a_9 = \sum_{i=7}^9 a_i$. We note that the index is i , the initial value of the index is 7, and final value of the index is 9.
- $a_1 + a_2 + \dots + a_n = \sum_{i=1}^n a_i$

Let $a^{\text{row}} = [a_1 \ a_2 \ \dots \ a_k]$ be a row vector with k elements and let $b^{\text{col}} = \begin{bmatrix} b_1 \\ b_2 \\ \vdots \\ b_k \end{bmatrix}$ be a column vector with the same number of elements as a^{row} . The **product** of a^{row} and b^{col} is defined¹ as

$$a^{\text{row}} \cdot b^{\text{col}} := \sum_{i=1}^k a_i b_i := a_1 b_1 + a_2 b_2 + \dots + a_k b_k. \quad (4.1)$$

For many, the following visual representation is more understandable,

$$[a_1 \ a_2 \ \dots \ a_k] \cdot \begin{bmatrix} b_1 \\ b_2 \\ \vdots \\ b_k \end{bmatrix} := a_1 b_1 + a_2 b_2 + \dots + a_k b_k. \quad (4.2)$$

Remarks on Multiplying a Row Vector Times a Column Vector

- It is very important to observe that the two vectors **MUST** have the same number of elements for their product to be defined.
- While (4.1) and (4.2) are equivalent (meaning they represent the same mathematical information), the formula in (4.1) is more convenient when writing a program and the visual representation in (4.2) is more convenient for hand computations.
- At this point, a column vector “times” a row vector is not defined. When we do make sense of it, it will not be equal to $\sum_{i=1}^k a_i b_i$.

¹If you have not seen the “summation” symbol before, here are some examples: $\sum_{i=1}^3 i := 1 + 2 + 3$, $\sum_{i=5}^8 i := 5 + 6 + 7 + 8$, and $\sum_{k=1}^n k^2 := 1 + (2)^2 + (3)^2 + \dots + (n)^2$.

Example 4.1 Let $a^{\text{row}} = [1 \ 2 \ 3]$ be a row vector with $k = 3$ elements and let $b^{\text{col}} = \begin{bmatrix} 4 \\ 5 \\ 6 \end{bmatrix}$ be a column vector with $k = 3$ elements. Perform their multiplication if it makes sense.

Solution: Because they have the same number of elements, we can form their product and we compute

$$a^{\text{row}} \cdot b^{\text{col}} := \sum_{i=1}^3 a_i b_i = (1)(4) + (2)(5) + (3)(6) = 32,$$

or we can look at it in the form

$$[1 \ 2 \ 3] \cdot \begin{bmatrix} 4 \\ 5 \\ 6 \end{bmatrix} = (1)(4) + (2)(5) + (3)(6) = 32. \quad \blacksquare$$

Example 4.2 Let $a^{\text{row}} = [1 \ 2 \ 3]$ be a row vector with $k = 3$ elements and let $b^{\text{col}} = \begin{bmatrix} 4 \\ 5 \end{bmatrix}$ be a column vector with $k = 2$ elements. Perform their multiplication if it makes sense.

Solution: Because they do NOT have the same number of elements, we cannot form their product.

$$a^{\text{row}} \cdot b^{\text{col}} := \text{undefined},$$

or

$$[1 \ 2 \ 3] \cdot \begin{bmatrix} 4 \\ 5 \end{bmatrix} = \text{error!} \quad \blacksquare$$

Example 4.3 Let $a^{\text{row}} = [2 \ -3 \ -1 \ 11]$ be a row vector with $k = 4$ elements and let $b^{\text{col}} = \begin{bmatrix} 3 \\ 5 \\ -1 \\ -2 \end{bmatrix}$ be a column vector with $k = 4$ elements. Perform their multiplication if it makes sense.

Solution: Because they have the same number of elements, we can form their product and we compute

$$a^{\text{row}} \cdot b^{\text{col}} := \sum_{i=1}^4 a_i b_i = (2)(3) + (-3)(5) + (-1)(-1) + (11)(-2) = -30,$$

or, equivalently, we write it like this

$$[2 \ -3 \ -1 \ 11] \cdot \begin{bmatrix} 3 \\ 5 \\ -1 \\ -2 \end{bmatrix} = (2)(3) + (-3)(5) + (-1)(-1) + (11)(-2) = -30. \quad \blacksquare$$

4.2 Examples of Row and Column Partitions

Let $A = \begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{bmatrix}$ be a 2×3 matrix. Then a **partition** of A **into rows** is

$$\begin{bmatrix} a_1^{\text{row}} \\ a_2^{\text{row}} \end{bmatrix} = \begin{bmatrix} \boxed{1 \ 2 \ 3} \\ \boxed{4 \ 5 \ 6} \end{bmatrix}, \quad \text{that is, } \begin{matrix} a_1^{\text{row}} = [1 \ 2 \ 3] \\ a_2^{\text{row}} = [4 \ 5 \ 6]. \end{matrix}$$

We note that a_1^{row} and a_2^{row} are row vectors of size 1×3 ; they have the same number of entries as A has columns.

A partition of $A = \begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{bmatrix}$ into columns is

$$[a_1^{\text{col}} \quad a_2^{\text{col}} \quad a_3^{\text{col}}] = \left[\begin{array}{|c|} \hline 1 \\ \hline 4 \\ \hline \end{array} \right] \left[\begin{array}{|c|} \hline 2 \\ \hline 5 \\ \hline \end{array} \right] \left[\begin{array}{|c|} \hline 3 \\ \hline 6 \\ \hline \end{array} \right], \text{ that is, } a_1^{\text{col}} = \begin{bmatrix} 1 \\ 4 \end{bmatrix}, a_2^{\text{col}} = \begin{bmatrix} 2 \\ 5 \end{bmatrix}, a_3^{\text{col}} = \begin{bmatrix} 3 \\ 6 \end{bmatrix}$$

We note that a_1^{col} , a_2^{col} , and a_3^{col} are column vectors of size 2×1 ; they have the same number of entries as A has rows.

4.3 General Case of Partitions

Let A be an $n \times m$ matrix. We recall that n is the number of rows in A , m is the number of columns, and a_{ij} is the notation for the ij element of A , that is, its value on the i -th row and j -th column. A partition of A into rows is

$$\begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1m} \\ a_{21} & a_{22} & \cdots & a_{2m} \\ \vdots & \vdots & \ddots & \vdots \\ a_{n1} & a_{n2} & \cdots & a_{nm} \end{bmatrix} = \begin{bmatrix} a_1^{\text{row}} \\ a_2^{\text{row}} \\ \vdots \\ a_n^{\text{row}} \end{bmatrix} = \begin{bmatrix} \boxed{a_{11} \ a_{12} \ \cdots \ a_{1m}} \\ \boxed{a_{21} \ a_{22} \ \cdots \ a_{2m}} \\ \vdots \\ \boxed{a_{n1} \ a_{n2} \ \cdots \ a_{nm}} \end{bmatrix}.$$

That is, the i -th row is the $1 \times m$ row vector

$$a_i^{\text{row}} = [a_{i1} \ a_{i2} \ \cdots \ a_{im}],$$

where i varies from 1 to n .

A partition of A into columns is

$$\begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1m} \\ a_{21} & a_{22} & \cdots & a_{2m} \\ \vdots & \vdots & \ddots & \vdots \\ a_{n1} & a_{n2} & \cdots & a_{nm} \end{bmatrix} = [a_1^{\text{col}} \quad a_2^{\text{col}} \quad \cdots \quad a_m^{\text{col}}] = \left[\begin{array}{|c|} \hline a_{11} \\ a_{21} \\ \vdots \\ a_{n1} \\ \hline \end{array} \right] \left[\begin{array}{|c|} \hline a_{12} \\ a_{22} \\ \vdots \\ a_{n2} \\ \hline \end{array} \right] \cdots \left[\begin{array}{|c|} \hline a_{1m} \\ a_{2m} \\ \vdots \\ a_{nm} \\ \hline \end{array} \right].$$

That is, the j -th column is the $n \times 1$ column vector

$$a_j^{\text{col}} = \begin{bmatrix} a_{1j} \\ a_{2j} \\ \vdots \\ a_{nj} \end{bmatrix},$$

where j varies from 1 to m .

4.4 Standard Matrix Multiplication: It's All About Rows and Columns

Let A be an $n \times k$ matrix, meaning it has n rows and k columns, and let B be a $k \times m$ matrix, so that it has k rows and m columns. When the number of columns of the first matrix A equals the number of rows of the second matrix B , the **matrix product** of A and B is defined and results in an $n \times m$ matrix:

$$[n \times k \text{ matrix}] \cdot [k \times m \text{ matrix}] = [n \times m \text{ matrix}].$$

The values of n , m , and k can be any integers greater than or equal to one. In particular, they can all be different numbers.

- $[4 \times 2 \text{ matrix}] \cdot [2 \times 5 \text{ matrix}] = [4 \times 5 \text{ matrix}]$.
- $[1 \times 11 \text{ matrix}] \cdot [11 \times 1 \text{ matrix}] = [1 \times 1 \text{ matrix}]$, which in Julia is different than a scalar.
- $[4 \times 4 \text{ matrix}] \cdot [4 \times 4 \text{ matrix}] = [4 \times 4 \text{ matrix}]$.
- $[4 \times 3 \text{ matrix}] \cdot [4 \times 4 \text{ matrix}] = \text{undefined}$.

Matrix multiplication using rows of A and columns of B

The standard way of doing **matrix multiplication** $A \cdot B$ involves multiplying the rows of A with the columns of B . We do some small examples before giving the general formula. **We will observe that the order in which we multiply two matrices is very important: in general, $A \cdot B \neq B \cdot A$ even when A and B are square matrices of the same size.**

Source of the Definition of Matrix Multiplication

Here is an optional video by Michael Penn explaining WHY matrix multiplication needs to work this way: <https://youtu.be/cc1ivD1Z71U>.

4.4.1 Examples

Example 4.4 We consider a 2×2 matrix A and a 2×1 matrix B , where we partition A by rows and B by columns. We note that the number of columns of A matches the number of rows of B so their product $A \cdot B$ is supposed to be defined. Let's do it!

Solution:

$$A = \begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix} = \begin{bmatrix} a_1^{\text{row}} \\ a_2^{\text{row}} \end{bmatrix} = \begin{bmatrix} \boxed{1 \ 2} \\ \boxed{3 \ 4} \end{bmatrix} \quad \text{and} \quad B = \begin{bmatrix} 5 \\ 6 \end{bmatrix} = \begin{bmatrix} b_1^{\text{col}} \end{bmatrix} = \begin{bmatrix} \boxed{5} \\ \boxed{6} \end{bmatrix}. \quad (4.3)$$

The matrix product of A and B is

$$A \cdot B = \begin{bmatrix} a_1^{\text{row}} \\ a_2^{\text{row}} \end{bmatrix} \cdot \begin{bmatrix} b_1^{\text{col}} \end{bmatrix} := \begin{bmatrix} a_1^{\text{row}} \cdot b_1^{\text{col}} \\ a_2^{\text{row}} \cdot b_1^{\text{col}} \end{bmatrix} = \begin{bmatrix} 17 \\ 39 \end{bmatrix} \quad (4.4)$$

because

$$\begin{aligned} a_1^{\text{row}} \cdot b_1^{\text{col}} &= [1 \ 2] \cdot \begin{bmatrix} 5 \\ 6 \end{bmatrix} = 5 + 12 = 17 \\ a_2^{\text{row}} \cdot b_1^{\text{col}} &= [3 \ 4] \cdot \begin{bmatrix} 5 \\ 6 \end{bmatrix} = 15 + 24 = 39. \end{aligned}$$

A more visual way to do the multiplication is like this,

$$A \cdot B = \begin{bmatrix} \boxed{1 \ 2} \\ \boxed{3 \ 4} \end{bmatrix} \cdot \begin{bmatrix} \boxed{5} \\ \boxed{6} \end{bmatrix} = \begin{bmatrix} (1)(5) + (2)(6) \\ (3)(5) + (4)(6) \end{bmatrix} = \begin{bmatrix} 17 \\ 39 \end{bmatrix} \quad (4.5)$$

■

Remarks

- A 2×2 matrix times a 2×1 matrix yields a 2×1 matrix.
- Let's observe that the operations for computing the matrix product boil down to performing the products of row vectors and column vectors of equal lengths! This is a general fact, as we will see.
- Let's also recall that we do not know how to form the product of a row vector with a column vector when their lengths are not equal.
- The number elements in a_i^{row} is equal to the number of columns of A , while the number of elements in b_j^{col} is equal to the number of rows of B . This is why the number of columns of A must equal the number of rows of B for their matrix product $A \cdot B$ to be defined.

Example 4.5 We reuse A and B above and ask if we can form the matrix product in the order $B \cdot A$.

Solution: We have that the first matrix B is 2×1 and the second matrix A is 2×2 . The number of columns of the first matrix does not match the number of rows of the second matrix, and hence the product cannot be defined in this direction. ■

Example 4.6 We consider a 2×2 matrix A and a 2×2 matrix B . We note that the number of columns of A matches the number of rows of B so their product $A \cdot B$ is defined. We also note that the number columns of B matches the number of rows of A , so the product $B \cdot A$ is also defined. This is a general property of square matrices A and B of the same size: their matrix product can be performed in either order. **We will note, however, that $A \cdot B \neq B \cdot A$. Hence, when multiplying matrices, the order matters!**

Solution:

$$A = \begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix} = \begin{bmatrix} a_1^{\text{row}} \\ a_2^{\text{row}} \end{bmatrix} = \begin{bmatrix} \boxed{1 \ 2} \\ \boxed{3 \ 4} \end{bmatrix} \quad \text{and} \quad B = \begin{bmatrix} 5 & -2 \\ 6 & 1 \end{bmatrix} = \begin{bmatrix} b_1^{\text{col}} & b_2^{\text{col}} \end{bmatrix} = \begin{bmatrix} \boxed{5} & \boxed{-2} \\ \boxed{6} & \boxed{1} \end{bmatrix}. \quad (4.6)$$

The matrix product of A and B is

$$A \cdot B = \begin{bmatrix} \boxed{1 \ 2} \\ \boxed{3 \ 4} \end{bmatrix} \cdot \begin{bmatrix} \boxed{5} & \boxed{-2} \\ \boxed{6} & \boxed{1} \end{bmatrix} = \begin{bmatrix} (1)(5) + (2)(6) & (1)(-2) + (2)(1) \\ (3)(5) + (4)(6) & (3)(-2) + (4)(1) \end{bmatrix} = \begin{bmatrix} 17 & 0 \\ 39 & -2 \end{bmatrix}. \quad (4.7)$$

The matrix product of B and A is

$$B \cdot A = \begin{bmatrix} \boxed{5 \ -2} \\ \boxed{6 \ 1} \end{bmatrix} \cdot \begin{bmatrix} \boxed{1} & \boxed{2} \\ \boxed{3} & \boxed{4} \end{bmatrix} = \begin{bmatrix} (5)(1) + (-2)(3) & (5)(2) + (-2)(4) \\ (6)(1) + (1)(3) & (6)(2) + (1)(4) \end{bmatrix} = \begin{bmatrix} -1 & 2 \\ 9 & 16 \end{bmatrix}. \quad (4.8)$$

Order Matters When Multiplying Matrices

- We note that $\begin{bmatrix} 17 & 0 \\ 39 & -2 \end{bmatrix} = A \cdot B \neq B \cdot A = \begin{bmatrix} -1 & 2 \\ 9 & 16 \end{bmatrix}$. The order in which matrices are multiplied matters. This is very different from the order of two real numbers, such as π and $\sqrt{2}$, which you can multiply in either order and you always get the same answer!
- $A \cdot B := \begin{bmatrix} a_1^{\text{row}} \cdot b_1^{\text{col}} & a_1^{\text{row}} \cdot b_2^{\text{col}} \\ a_2^{\text{row}} \cdot b_1^{\text{col}} & a_2^{\text{row}} \cdot b_2^{\text{col}} \end{bmatrix}$, where you can find a_i^{row} and b_j^{col} highlighted in (4.6).
- $B \cdot A := \begin{bmatrix} b_1^{\text{row}} \cdot a_1^{\text{col}} & b_1^{\text{row}} \cdot a_2^{\text{col}} \\ b_2^{\text{row}} \cdot a_1^{\text{col}} & b_2^{\text{row}} \cdot a_2^{\text{col}} \end{bmatrix}$, where you can find b_i^{row} and a_j^{col} highlighted in (4.8).

Example 4.7 For the given 3×2 matrix A and 2×2 matrix B , compute $A \cdot B$ and $B \cdot A$ if the given multiplications make sense.

$$A = \begin{bmatrix} 1 & 2 \\ 3 & 4 \\ 5 & 6 \end{bmatrix} \quad \text{and} \quad B = \begin{bmatrix} 3 & 4 \\ 2 & 1 \end{bmatrix}.$$

Solution: We note that the number of columns of A equals the number of rows of B so their product $A \cdot B$ is defined. We also note that the number columns of B does not equal the number of rows of A , so the product $B \cdot A$ is not defined.

$$A = \begin{bmatrix} 1 & 2 \\ 3 & 4 \\ 5 & 6 \end{bmatrix} = \begin{bmatrix} a_1^{\text{row}} \\ a_2^{\text{row}} \\ a_3^{\text{row}} \end{bmatrix} = \begin{bmatrix} \boxed{1 \ 2} \\ \boxed{3 \ 4} \\ \boxed{5 \ 6} \end{bmatrix} \quad \text{and} \quad B = \begin{bmatrix} 3 & 4 \\ 2 & 1 \end{bmatrix} = \begin{bmatrix} b_1^{\text{col}} & b_2^{\text{col}} \end{bmatrix} = \begin{bmatrix} \boxed{3} & \boxed{4} \\ \boxed{2} & \boxed{1} \end{bmatrix}$$

By now, you may have a favorite method, and using it, you should compute that

$$A \cdot B = \begin{bmatrix} 7 & 6 \\ 17 & 16 \\ 27 & 26 \end{bmatrix}$$

Remark: When doing calculations by hand, 3×3 matrices are about as big as you ever really want to do. In Julia, “the sky is the limit”. The following section is to help us understand what Julia is doing when it multiplies two matrices with the command `A * B`.

4.4.2 Optional Read: General case: what is happening inside Julia

We partition the $n \times k$ matrix A into rows and the $k \times m$ matrix B into columns, as in

$$A = \begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1k} \\ a_{21} & a_{22} & \cdots & a_{2k} \\ \vdots & \vdots & \ddots & \vdots \\ a_{n1} & a_{n2} & \cdots & a_{nk} \end{bmatrix} = \begin{bmatrix} a_1^{\text{row}} \\ a_2^{\text{row}} \\ \vdots \\ a_n^{\text{row}} \end{bmatrix} = \begin{bmatrix} \boxed{a_{11} \ a_{12} \ \cdots \ a_{1k}} \\ \boxed{a_{21} \ a_{22} \ \cdots \ a_{2k}} \\ \vdots \\ \boxed{a_{n1} \ a_{n2} \ \cdots \ a_{nk}} \end{bmatrix} \quad (4.9)$$

and

$$B = \begin{bmatrix} b_{11} & b_{12} & \cdots & b_{1m} \\ b_{21} & b_{22} & \cdots & b_{2m} \\ \vdots & \vdots & \ddots & \vdots \\ b_{k1} & b_{k2} & \cdots & b_{km} \end{bmatrix} = [b_1^{\text{col}} \ b_2^{\text{col}} \ \cdots \ b_m^{\text{col}}] = \begin{bmatrix} \boxed{b_{11}} & \boxed{b_{12}} & \cdots & \boxed{b_{1m}} \\ \boxed{b_{21}} & \boxed{b_{22}} & \cdots & \boxed{b_{2m}} \\ \vdots & \vdots & \cdots & \vdots \\ \boxed{b_{k1}} & \boxed{b_{k2}} & \cdots & \boxed{b_{km}} \end{bmatrix}, \quad (4.10)$$

then

$$A \cdot B := \begin{bmatrix} a_1^{\text{row}} \cdot b_1^{\text{col}} & a_1^{\text{row}} \cdot b_2^{\text{col}} & \cdots & a_1^{\text{row}} \cdot b_m^{\text{col}} \\ a_2^{\text{row}} \cdot b_1^{\text{col}} & a_2^{\text{row}} \cdot b_2^{\text{col}} & \cdots & a_2^{\text{row}} \cdot b_m^{\text{col}} \\ \vdots & \vdots & \ddots & \vdots \\ a_n^{\text{row}} \cdot b_1^{\text{col}} & a_n^{\text{row}} \cdot b_2^{\text{col}} & \cdots & a_n^{\text{row}} \cdot b_m^{\text{col}} \end{bmatrix}. \quad (4.11)$$

Another way to see the pattern is like this

$$A \cdot B := \begin{bmatrix} \boxed{a_{11} \ a_{12} \ \cdots \ a_{1k}} \\ \boxed{a_{21} \ a_{22} \ \cdots \ a_{2k}} \\ \vdots \\ \boxed{a_{n1} \ a_{n2} \ \cdots \ a_{nk}} \end{bmatrix} \cdot \begin{bmatrix} \boxed{b_{11}} & \boxed{b_{12}} & \cdots & \boxed{b_{1m}} \\ \boxed{b_{21}} & \boxed{b_{22}} & \cdots & \boxed{b_{2m}} \\ \vdots & \vdots & \cdots & \vdots \\ \boxed{b_{k1}} & \boxed{b_{k2}} & \cdots & \boxed{b_{km}} \end{bmatrix} = \begin{bmatrix} \sum_{i=1}^k a_{1i} b_{i1} & \sum_{i=1}^k a_{1i} b_{i2} & \cdots & \sum_{i=1}^k a_{1i} b_{im} \\ \sum_{i=1}^k a_{2i} b_{i1} & \sum_{i=1}^k a_{2i} b_{i2} & \cdots & \sum_{i=1}^k a_{2i} b_{im} \\ \vdots & \vdots & \ddots & \vdots \\ \sum_{i=1}^k a_{ni} b_{i1} & \sum_{i=1}^k a_{ni} b_{i2} & \cdots & \sum_{i=1}^k a_{ni} b_{im} \end{bmatrix}. \quad (4.12)$$

Bottom Line on Standard Multiplication

In the standard way of defining matrix multiplication, the ij -entry of $A \cdot B$ is obtained by multiplying the i -th row of A by the j -th column of B (whenever the multiplication makes sense, meaning that the number of columns of A equals the number of rows of B). We will not use the following notation on a regular basis, but some of you may like it; if we let $[A \cdot B]_{ij}$ denote the ij -element of the matrix, $A \cdot B$, then

$$[A \cdot B]_{ij} := a_i^{\text{row}} \cdot b_j^{\text{col}}.$$

4.5 Multiplication by Summing over Columns and Rows

Rock your World: an Alternative Formula for Matrix Multiplication

- We now introduce an alternative way to compute the product of two matrices. It gives the same answer as the “standard method”. **Very few people use or even know this definition because, for hand calculations, it takes more time to write out the individual steps.** ROB 101, however, is about computational linear algebra, and hence we ignore such trivial concerns as what is best for doing hand calculations!
- We will see shortly that understanding this alternative way of matrix multiplication will allow us to solve almost any system of linear equations via a combination of forward and back substitution.
- Instead of solving one hard system of linear equations, we will find the solution by solving two triangular systems of linear equations, one upper triangular and one lower triangular!

We reconsider two matrices A and B from Example 4.7, but this time we partition A into columns and B into rows

$$A = \begin{bmatrix} 1 & 2 \\ 3 & 4 \\ 5 & 6 \end{bmatrix} = [a_1^{\text{col}} \quad a_2^{\text{col}}] = \left[\begin{array}{|c|} \hline 1 \\ \hline 3 \\ \hline 5 \\ \hline \end{array} \quad \begin{array}{|c|} \hline 2 \\ \hline 4 \\ \hline 6 \\ \hline \end{array} \right] \quad \text{and} \quad B = \begin{bmatrix} 3 & 4 \\ 2 & 1 \end{bmatrix} = \begin{bmatrix} b_1^{\text{row}} \\ b_2^{\text{row}} \end{bmatrix} = \left[\begin{array}{|c|c|} \hline 3 & 4 \\ \hline 2 & 1 \\ \hline \end{array} \right].$$

What happens when we form the sum of the columns of A times the rows of B ,

$$a_1^{\text{col}} b_1^{\text{row}} + a_2^{\text{col}} b_2^{\text{row}} = ?$$

We first observe that the columns of A are 3×1 while the rows of B are 1×2 . Hence, their product with the usual rules of matrix multiplication will be 3×2 , the same size as $A \cdot B$. That seems interesting. We do the two multiplications and obtain

$$a_1^{\text{col}} \cdot b_1^{\text{row}} = \begin{bmatrix} 1 \\ 3 \\ 5 \end{bmatrix} \cdot [3 \quad 4] = \begin{bmatrix} (1) \cdot (3) & (1) \cdot (4) \\ (3) \cdot (3) & (3) \cdot (4) \\ (5) \cdot (3) & (5) \cdot (4) \end{bmatrix} = \begin{bmatrix} 3 & 4 \\ 9 & 12 \\ 15 & 20 \end{bmatrix}$$

$$a_2^{\text{col}} \cdot b_2^{\text{row}} = \begin{bmatrix} 2 \\ 4 \\ 6 \end{bmatrix} \cdot [2 \quad 1] = \begin{bmatrix} (2) \cdot (2) & (2) \cdot (1) \\ (4) \cdot (2) & (4) \cdot (1) \\ (6) \cdot (2) & (6) \cdot (1) \end{bmatrix} = \begin{bmatrix} 4 & 2 \\ 8 & 4 \\ 12 & 6 \end{bmatrix}$$

and hence their sum is

$$a_1^{\text{col}} b_1^{\text{row}} + a_2^{\text{col}} b_2^{\text{row}} = \begin{bmatrix} 3 & 4 \\ 9 & 12 \\ 15 & 20 \end{bmatrix} + \begin{bmatrix} 4 & 2 \\ 8 & 4 \\ 12 & 6 \end{bmatrix} = \begin{bmatrix} 7 & 6 \\ 17 & 16 \\ 27 & 26 \end{bmatrix} = A \cdot B$$

So while the bookkeeping is a bit different, the individual computations are easier than in the standard way of performing matrix multiplication, and perhaps they are also easier to get right. You're free to do your matrix multiplications as you wish. This particular method has a theoretical benefit that we will uncover shortly. Keep in mind that not so many people think of matrix multiplication as "sums of columns times rows" and you might confuse friends and other instructors when you talk about it in this manner. In fact, you may be told that you are wrong and that no one ever does it that way, even if it is correct.

General Case of Matrix Multiplication using Columns of A and Rows of B

Suppose that A is $n \times k$ and B is $k \times m$ so that the two matrices are compatible for matrix multiplication. Then

$$A \cdot B = \sum_{i=1}^k a_i^{\text{col}} \cdot b_i^{\text{row}},$$

the "sum of the columns of A multiplied by the rows of B ". A more precise way to say it would be "the sum over i of the i -th column of A times the i -th row of B ." In HW, we'll play with this idea enough that it will become ingrained into your subconscious!

Alpha Go Takes on Matrix Multiplication

Who would have guessed that there is still much innovation to be done in this space! You may enjoy these videos:

- <https://youtu.be/8IILk4Wjo5rc> AlphaTensor by DEEPMIND finds new Algorithms for Matrix Multiplication
- <https://youtu.be/CoFvu-n-fFs> AI Just Solved a 53-Year-Old Problem! | AlphaTensor, Explained

4.6 (Optional Read:) Why the Second Method of Matrix Multiplication Works

ROB 101 does not focus on proofs. Yet, sometimes it is genuinely fun to understand how or why things are as they are! We consider two 2×2 matrices A and B , where

$$A = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} \text{ and } B = \begin{bmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{bmatrix}.$$

Then, using the standard rows of A times the columns of B formulation of matrix multiplication yields

$$\begin{aligned} A \cdot B &:= \begin{bmatrix} a_1^{\text{row}} \cdot b_1^{\text{col}} & a_1^{\text{row}} \cdot b_2^{\text{col}} \\ a_2^{\text{row}} \cdot b_1^{\text{col}} & a_2^{\text{row}} \cdot b_2^{\text{col}} \end{bmatrix} \\ &= \begin{bmatrix} a_{11}b_{11} + a_{12}b_{21} & a_{11}b_{12} + a_{12}b_{22} \\ a_{21}b_{11} + a_{22}b_{21} & a_{21}b_{12} + a_{22}b_{22} \end{bmatrix} \text{ (we take the sum outside the matrix)} \\ &= \begin{bmatrix} a_{11}b_{11} & a_{11}b_{12} \\ a_{21}b_{11} & a_{21}b_{12} \end{bmatrix} + \begin{bmatrix} a_{12}b_{21} & a_{12}b_{22} \\ a_{22}b_{21} & a_{22}b_{22} \end{bmatrix} \text{ (we recognize each term)} \\ &= \begin{bmatrix} a_{11} \\ a_{21} \end{bmatrix} \cdot \begin{bmatrix} b_{11} & b_{12} \end{bmatrix} + \begin{bmatrix} a_{12} \\ a_{22} \end{bmatrix} \cdot \begin{bmatrix} b_{21} & b_{22} \end{bmatrix} \\ &= a_1^{\text{col}} \cdot b_1^{\text{row}} + a_2^{\text{col}} \cdot b_2^{\text{row}} \end{aligned}$$

In a similar manner, we can treat the general case. For those who wish to brave a blaze of indices, from (4.12), we have

$$\begin{aligned} A \cdot B &:= \begin{bmatrix} \boxed{a_{11} \ a_{12} \ \cdots \ a_{1k}} \\ \boxed{a_{21} \ a_{22} \ \cdots \ a_{2k}} \\ \vdots \\ \boxed{a_{n1} \ a_{n2} \ \cdots \ a_{nk}} \end{bmatrix} \cdot \begin{bmatrix} \boxed{b_{11}} & \boxed{b_{12}} & \cdots & \boxed{b_{1m}} \\ \boxed{b_{21}} & \boxed{b_{22}} & & \boxed{b_{2m}} \\ \vdots & \vdots & & \vdots \\ \boxed{b_{k1}} & \boxed{b_{k2}} & & \boxed{b_{km}} \end{bmatrix} \\ &= \begin{bmatrix} \sum_{i=1}^k a_{1i}b_{i1} & \sum_{i=1}^k a_{1i}b_{i2} & \cdots & \sum_{i=1}^k a_{1i}b_{im} \\ \sum_{i=1}^k a_{2i}b_{i1} & \sum_{i=1}^k a_{2i}b_{i2} & \cdots & \sum_{i=1}^k a_{2i}b_{im} \\ \vdots & \vdots & \ddots & \vdots \\ \sum_{i=1}^k a_{ni}b_{i1} & \sum_{i=1}^k a_{ni}b_{i2} & \cdots & \sum_{i=1}^k a_{ni}b_{im} \end{bmatrix} \text{ (we pull the sum outside the matrix)} \\ &= \sum_{i=1}^k \begin{bmatrix} a_{1i}b_{i1} & a_{1i}b_{i2} & \cdots & a_{1i}b_{im} \\ a_{2i}b_{i1} & a_{2i}b_{i2} & \cdots & a_{2i}b_{im} \\ \vdots & \vdots & \ddots & \vdots \\ a_{ni}b_{i1} & a_{ni}b_{i2} & \cdots & a_{ni}b_{im} \end{bmatrix} \text{ (we recognize what this is)} \\ &= \sum_{i=1}^k a_i^{\text{col}} \cdot b_i^{\text{row}} \\ &= \begin{bmatrix} \boxed{a_{11}} & \boxed{a_{12}} & \cdots & \boxed{a_{1k}} \\ \boxed{a_{21}} & \boxed{a_{22}} & & \boxed{a_{2k}} \\ \vdots & \vdots & & \vdots \\ \boxed{a_{n1}} & \boxed{a_{n2}} & & \boxed{a_{nk}} \end{bmatrix} \cdot \begin{bmatrix} \boxed{b_{11} \ b_{12} \ \cdots \ b_{1m}} \\ \boxed{b_{21} \ b_{22} \ \cdots \ b_{2m}} \\ \vdots \\ \boxed{b_{k1} \ b_{k2} \ \cdots \ b_{km}} \end{bmatrix}. \end{aligned} \tag{4.13}$$

4.7 Permutation Matrices: The Matrix View of Swapping the Order of Equations

Back in Sec. 3.8, we made a brief remark on swapping rows of equations to put them into a nicer form. For simplicity, we recall the equations again here.

This system of equations is neither upper triangular nor lower triangular

$$\begin{aligned} 3x_1 &= 6 \\ x_1 - 2x_2 + 3x_3 &= 2 \\ 2x_1 - x_2 &= -2, \end{aligned} \tag{4.14}$$

but if we simply re-arrange the order of the equations, we arrive at the lower triangular equations

$$\begin{aligned} 3x_1 &= 6 \\ 2x_1 - x_2 &= -2 \\ x_1 - 2x_2 + 3x_3 &= 2. \end{aligned} \tag{4.15}$$

As we noted before, we still have the same equations and hence their solution has not changed. The equations have just been re-arranged to make the process of computing their solution fall into a nice pattern that we already know how to handle.

We now write out the matrix equations for the “unfortunately ordered” equations (4.14) and then the “nicely” re-arranged system of equations (4.15)

$$\underbrace{\begin{bmatrix} 3 & 0 & 0 \\ 1 & -2 & 3 \\ 2 & -1 & 0 \end{bmatrix}}_{A_O} \underbrace{\begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}}_x = \underbrace{\begin{bmatrix} 6 \\ 2 \\ -2 \end{bmatrix}}_{b_O} \quad \underbrace{\begin{bmatrix} 3 & 0 & 0 \\ 2 & -1 & 0 \\ 1 & -2 & 3 \end{bmatrix}}_{A_L} \underbrace{\begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}}_x = \underbrace{\begin{bmatrix} 6 \\ -2 \\ 2 \end{bmatrix}}_{b_L}.$$

We see the second and third rows are swapped when we compare A_O to A_L and b_O to b_L . The swapping of rows can be accomplished by multiplying A_O and b_O on the left by

$$\mathbf{P} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix}.$$

Indeed, we check that

$$P \cdot A_O = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} 3 & 0 & 0 \\ 1 & -2 & 3 \\ 2 & -1 & 0 \end{bmatrix} = \begin{bmatrix} 3 & 0 & 0 \\ 2 & -1 & 0 \\ 1 & -2 & 3 \end{bmatrix} = A_L$$

and

$$P \cdot b_O = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} 6 \\ 2 \\ -2 \end{bmatrix} = \begin{bmatrix} 6 \\ -2 \\ 2 \end{bmatrix} = b_L.$$

P is called a permutation matrix.

Key Insight on Swapping Rows

We note that the permutation matrix P is constructed from the 3×3 identity matrix I by swapping its second and third rows, exactly the rows we wanted to swap in A_O and b_O . This observation works in general.

For example, suppose we want to do the following rearrangement

$$\begin{bmatrix} 1 \\ 2 \\ 3 \\ 4 \\ 5 \end{bmatrix} \leftrightarrow \begin{bmatrix} 4 \\ 2 \\ 5 \\ 1 \\ 3 \end{bmatrix}, \tag{4.16}$$

where rows one and four are swapped and three and five are swapped. We have shown the arrow as going back and forth (\leftrightarrow) because applying the swap twice results in the original arrangement. To see the resulting structure at a matrix level, we put the 5×5 identity matrix on the left and the permutation matrix P on the right

$$I = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix} \leftrightarrow P = \begin{bmatrix} 0 & 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \end{bmatrix}$$

so that it is apparent that P is just a re-ordering of the rows of I . Moreover, if we want to go from

$$\begin{bmatrix} 4 \\ 2 \\ 5 \\ 1 \\ 3 \end{bmatrix} \leftrightarrow \begin{bmatrix} 1 \\ 2 \\ 3 \\ 4 \\ 5 \end{bmatrix},$$

we can apply the matrix, P , once again. Indeed, you are invited to check in Julia that $P \cdot P = I$.

The kind of permutation matrix we have introduced here, where permuting the rows twice results in the original ordering of the rows, is a special case of a more general kind of permutation. While we do not need the more general matrices in ROB 101, we feel you need to know that they exist.

Heads Up! The above is only half the story on permutation matrices.

(Optional Read): A more general rearrangement of rows would be something like this

$$\begin{bmatrix} 1 \\ 2 \\ 3 \\ 4 \\ 5 \end{bmatrix} \rightarrow \begin{bmatrix} 4 \\ 2 \\ 1 \\ 5 \\ 3 \end{bmatrix}, \tag{4.17}$$

where applying the rearrangement twice does not result in the original order.

To see how this looks at the matrix level, we put the 5×5 identity matrix on the left and the permutation matrix P on the right

$$I = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix} \leftrightarrow P = \begin{bmatrix} 0 & 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 & 0 \end{bmatrix}.$$

P is still just a re-ordering of the rows of I , but it is no longer true that $P \cdot P = I$.

Permutation Matrices, the Full Story

In general, matrices that consist of all ones and zeros, with each row and column having a single one, are called **permutation matrices**.

In Chapter 6.4, we'll cover the matrix transpose, P^\top . Indeed, you are invited to look ahead and see what that is about. If you do so, you can check in Julia that $P \neq P^\top$ (we'll say that the matrix is not symmetric) and $P \cdot P \neq I$, but that $P^\top \cdot P = I$.

In Chapter 6.2, we introduce the matrix inverse, denoted P^{-1} . In general, inverses of matrices do not exist and they are hard to compute. For permutation matrices, $P^{-1} = P^\top$. They automatically satisfy $P^\top \cdot P = P \cdot P^\top = I$.

4.8 (Optional Read): Useful Fact on Matrix Multiplication

This fact is used in Project 1 to speed up the calculations. Consider an $n \times k$ matrix A and a $k \times m$ matrix B . Then

$$A \cdot B = A \cdot \underbrace{\begin{bmatrix} b_1^{\text{col}} & \dots & b_j^{\text{col}} & \dots & b_m^{\text{col}} \end{bmatrix}}_B = \begin{bmatrix} Ab_1^{\text{col}} & \dots & Ab_j^{\text{col}} & \dots & Ab_m^{\text{col}} \end{bmatrix};$$

that is, the j -th column of $A \cdot B$ is A times the j -th column of B . To see why this is true, we know that the ij -element of $C := A \cdot B$ is given by

$$c_{ij} := a_i^{\text{row}} \cdot b_j^{\text{col}},$$

and hence

$$c_j^{\text{col}} = \begin{bmatrix} c_{1j} \\ c_{2j} \\ \vdots \\ c_{nj} \end{bmatrix} = \begin{bmatrix} a_1^{\text{row}} \cdot b_j^{\text{col}} \\ a_2^{\text{row}} \cdot b_j^{\text{col}} \\ \vdots \\ a_n^{\text{row}} \cdot b_j^{\text{col}} \end{bmatrix} = \underbrace{\begin{bmatrix} a_1^{\text{row}} \\ a_2^{\text{row}} \\ \vdots \\ a_n^{\text{row}} \end{bmatrix}}_A \cdot b_j^{\text{col}} = A \cdot b_j^{\text{col}}.$$




Figure 4.1: Sixty LiDAR images (scans) are combined in this image of the Ford Motor Company Robotics Building. Image courtesy of Bruce Huang. (a) Wide view. (b) Zoom on Cassie’s position so you can see the road and the sidewalk. In Project 1, the images are collected while Cassie is walking. You will use matrix-vector multiplication to compensate for Cassie’s motion so that the images all appear to be taken from the same position and direction in 3D space. This is called “image registration”.

4.9 (Optional Read): Some Geometric Aspects of Matrix-Vector Multiplication

Our initial efforts in Computational Linear Algebra focus on the “algebraic” part of Linear Algebra, in other words, on equation solving. There are also deep connections with geometry that are used in Robotics. Your Project 1 on Robot Map Building focuses on geometry. We live in a three-dimensional (3D) world and we also find it useful to navigate with two-dimensional (2D) representations of the world, such as Google Maps.

How do vectors, matrices, and geometry all come together? Figure 4.1 shows a LiDAR image of the Ford Motor Company Robotics Building. The image consists of sixty individual LiDAR images (scans) that have been combined into one single image. The 3D geometry of the building, trees, and even the road passing in front of the building can be clearly seen. That doesn’t seem so special, right? We’re used to camera images that show even more detail than that! Yes, but while an image today contains millions of pixels, data from a LiDAR are “sparse”, in the sense that each scan may have only 10,000 points in it, or less than 0.1% of the number of pixels in today’s camera images. Each return from one of its lasers gives a 3D point in space, the (x, y, z) coordinates of the point in space off which the laser beam reflected! In addition, the intensity of each laser’s return is measured, so we really have a 4-vector (x, y, z, I) . While we use the intensity to color the image, we’ll ignore it in what follows.

The final image has roughly $60 \times 10,000$ laser returns, that is, points of light, in it, and each point of light in the image is a 3-vector. So, we have roughly 600,000 3-vectors in the image. For each of the scans, Cassie’s “head” was pointing in a different direction. In addition, the body sways a bit. Hence, each of the 10,000 vectors in an individual scan is being collected from a different position in 3D space, and looking toward a different direction or angle in 3D space, with respect to the other 59 scans. If you were to overlay the 60 images without adjusting for the offsets in position and direction, you would obtain a blurry mess. To see what we mean, just point your cell phone at the night sky sometime with the camera at a long exposure interval; the jittering of your hand will cause a blurry night sky.

How to remove the blur? Well, you can declare the position and angle of Cassie in the first scan as a reference and then compensate the vectors in each of the other scans to remove variations in its relative position and angle with respect to the first scan. This compensation process is called “image registration” and it can be accomplished by multiplying the (x, y, z) coordinates, which form a 3-vector², by an appropriate matrix. **Bottom line, building a map involves multiplying vectors by matrices, where the matrices**

²In Project 1, you’ll learn that it is better to work with 4-vectors, but the details would be confusing here, so we omit them.

somehow encode the relative changes in position and direction of a scan!

Who knew that Cassie had to be a Linear Algebra whiz to build a map! We'll now simplify the problem of imagining how matrices transform 3-vectors by reducing the dimension by one. Moreover, instead of 10,000 2-vectors, we take a single 2-vector $v = \begin{bmatrix} 2.0 \\ 0.5 \end{bmatrix}$ shown in blue in Fig. 4.2 and multiply it by four matrices,

$$(a) A_1 := \begin{bmatrix} \cos(\pi/10) & -\sin(\pi/10) \\ \sin(\pi/10) & \cos(\pi/10) \end{bmatrix} \implies A_1 v = \begin{bmatrix} 1.7476 \\ 1.0936 \end{bmatrix}$$

$$(b) A_2 := \begin{bmatrix} \cos(3\pi/4) & -\sin(3\pi/4) \\ \sin(3\pi/4) & \cos(3\pi/4) \end{bmatrix} \implies A_2 v = \begin{bmatrix} -1.7678 \\ 1.0607 \end{bmatrix}$$

$$(c) A_3 := \begin{bmatrix} 0.5 & 0.0 \\ 0.0 & 0.5 \end{bmatrix} \implies A_3 v = \begin{bmatrix} 1.0000 \\ 0.2500 \end{bmatrix}$$

$$(d) A_4 := \begin{bmatrix} -0.6000 & 0.6000 \\ -0.3360 & -0.0840 \end{bmatrix} \implies A_4 v = \begin{bmatrix} -0.9000 \\ -0.7140 \end{bmatrix}$$

The results are given by the red vectors in Fig. 4.2. In the context of Cassie and the LiDAR data, Fig. 4.2-(a) and -(b) are the most relevant as they show a vector being rotated by a fixed angle. Fig. 4.2-(c) shows a vector being scaled, while Fig. 4.2-(d) shows a vector that appears to be rotated and scaled.

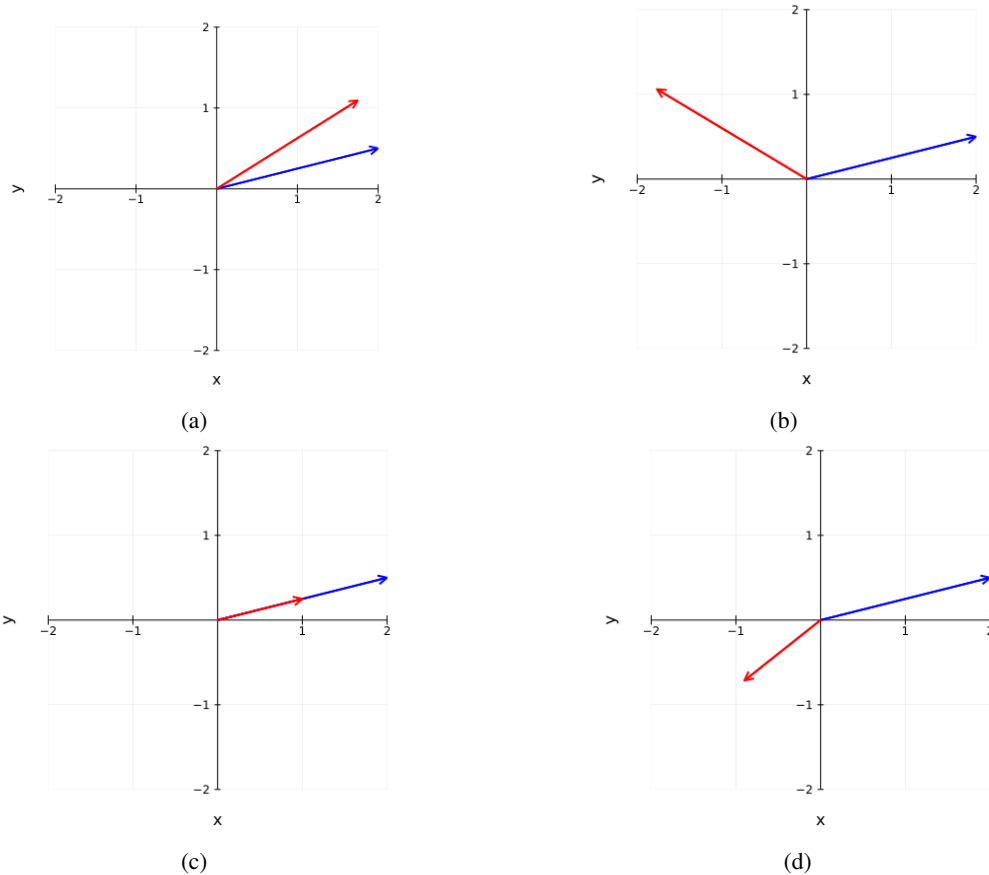


Figure 4.2: The vector v is in blue while the transformed vector, Av , is shown in red. (a) The matrix rotates the blue vector clockwise by an angle of $\pi/10$ or 18° . (b) The matrix rotates the blue vector clockwise by an angle of $3\pi/4$ or 135° . (c) The matrix scales the blue vector by 0.5. And (d), the matrix acts on the blue vector in a general fashion that we will be able to “decode” when we study eigenvalues and eigenvectors in Chapter 10.

4.10 Looking Ahead

Once again, our first major goal in the course is to solve very large sets of linear equations, say more than 100 variables. In the previous chapter, we saw how to solve problems with triangular structure. In this Chapter, we learned how to multiply two matrices.

Let's use our knowledge and see what happens when we multiply a lower triangular matrix times an upper triangular matrix. We define two matrices by L and U for lower and upper triangular, respectively,

$$L = \begin{bmatrix} 1 & 0 & 0 \\ -2 & 1 & 0 \\ 3 & -2 & 1 \end{bmatrix} \text{ and } U = \begin{bmatrix} 3 & 3 & 2 \\ 0 & 6 & 1 \\ 0 & 0 & -3 \end{bmatrix} \implies L \cdot U = \begin{bmatrix} 3 & 3 & 2 \\ -6 & 0 & -3 \\ 9 & -3 & 1 \end{bmatrix} =: A.$$

Hence the product of a lower triangular matrix and an upper triangular matrix seems to be a “general matrix” A , meaning it has no particular structure.

Question: Is it possible to go backwards? That is, starting with a general matrix A , is it possible to write it as the product of a lower triangular matrix and an upper triangular matrix? And if it is possible, is it useful? What's the trick?

Answers: Yes. Very! Our alternative way of doing matrix multiplication.

Chapter 5

LU (Lower-Upper) Factorization

Learning Objectives

- How to reduce a hard problem to two much easier problems
- The concept of “factoring” a matrix into a product of two simpler matrices that are in turn useful for solving systems of linear equations.

Outcomes

- Our first encounter in lecture with an explicit algorithm
- Learn how to do a *special case* of the LU factorization, where L is a lower triangular matrix and U is an upper triangular matrix.
- Use the LU factorization to solve linear equations
- More advanced: what we missed in our first pass at LU factorization: a (row) permutation matrix.

5.1 Recalling Forward and Back Substitution

For a lower triangular system of equations with non-zero leading coefficients, such as

$$\begin{aligned} 3x_1 &= 6 \\ 2x_1 - x_2 &= -2 \\ x_1 - 2x_2 + 3x_3 &= 2, \end{aligned} \iff \begin{bmatrix} 3 & 0 & 0 \\ 2 & -1 & 0 \\ 1 & -2 & 3 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} 6 \\ -2 \\ 2 \end{bmatrix} \quad (5.1)$$

we can find a solution via forward substitution,

$$\begin{aligned} x_1 &= \frac{1}{3}6 = 2 \\ x_2 &= -[-2 - 2x_1] \\ x_3 &= \frac{1}{3}[2 - x_1 + 2x_2], \end{aligned} \implies \begin{cases} x_1 = \frac{1}{3}6 = 2 \\ x_2 = 2 + 2x_1 = 2 + 2(2) = 6 \\ x_3 = \frac{1}{3}[2 - x_1 + 2x_2] = \frac{1}{3}[2 - (2) + 2(6)] = \frac{12}{3} = 4. \end{cases} \quad (5.2)$$

On the other hand, for an upper triangular system of equations with non-zero leading coefficients, such as

$$\begin{aligned} x_1 + 3x_2 + 2x_3 &= 6 \\ 2x_2 + x_3 &= -2 \\ 3x_3 &= 4, \end{aligned} \iff \begin{bmatrix} 1 & 3 & 2 \\ 0 & 2 & 1 \\ 0 & 0 & 3 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} 6 \\ -2 \\ 4 \end{bmatrix} \quad (5.3)$$

we can find a solution by back substitution,

$$\begin{aligned} x_1 &= 6 - (3x_2 + 2x_3) \\ x_2 &= \frac{1}{2}(-2 - x_3) \\ x_3 &= \frac{4}{3}, \end{aligned} \implies \begin{cases} x_1 = 6 - (3x_2 + 2x_3) = 6 - \left(\frac{3}{2} + \frac{8}{3}\right) = \frac{18}{3} - \frac{11}{3} = 2\frac{1}{3} \\ x_2 = \frac{1}{2}(-2 - x_3) = \frac{1}{2}\left(-2 - \frac{4}{3}\right) = \frac{1}{2}\left(-\frac{10}{3}\right) = -\frac{5}{3} \\ x_3 = \frac{4}{3}. \end{cases} \quad (5.4)$$

5.2 Recalling Matrix Multiplication in the Form of Columns Times Rows

Suppose that A is $n \times k$ and B is $k \times m$ so that the two matrices are compatible for matrix multiplication. Then

$$A \cdot B = \sum_{i=1}^k a_i^{\text{col}} \cdot b_i^{\text{row}},$$

the “sum of the columns of A multiplied by the rows of B ”.

Example 5.1 Form the matrix product of $A = \begin{bmatrix} 1 & 0 \\ 3 & 4 \end{bmatrix}$ and $B = \begin{bmatrix} 5 & 2 \\ 0 & -1 \end{bmatrix}$.

Solution

$$\begin{aligned} A &= \begin{bmatrix} 1 & 0 \\ 3 & 4 \end{bmatrix} = [a_1^{\text{col}} \quad a_2^{\text{col}}] = \left[\begin{bmatrix} 1 \\ 3 \end{bmatrix} \quad \begin{bmatrix} 0 \\ 4 \end{bmatrix} \right] \quad \text{and} \quad B = \begin{bmatrix} 5 & 2 \\ 0 & -1 \end{bmatrix} = \begin{bmatrix} b_1^{\text{row}} \\ b_2^{\text{row}} \end{bmatrix} = \left[\begin{bmatrix} 5 & 2 \end{bmatrix} \\ \begin{bmatrix} 0 & -1 \end{bmatrix} \right] \\ a_1^{\text{col}} b_1^{\text{row}} &= \begin{bmatrix} 1 \\ 3 \end{bmatrix} \cdot [5 \quad 2] = \begin{bmatrix} 5 & 2 \\ 15 & 6 \end{bmatrix} \\ a_2^{\text{col}} b_2^{\text{row}} &= \begin{bmatrix} 0 \\ 4 \end{bmatrix} \cdot [0 \quad -1] = \begin{bmatrix} 0 & 0 \\ 0 & -4 \end{bmatrix} \end{aligned}$$

and the matrix product is

$$A \cdot B = a_1^{\text{col}} b_1^{\text{row}} + a_2^{\text{col}} b_2^{\text{row}} = \begin{bmatrix} 5 & 2 \\ 15 & 6 \end{bmatrix} + \begin{bmatrix} 0 & 0 \\ 0 & -4 \end{bmatrix} = \begin{bmatrix} 5 & 2 \\ 15 & 2 \end{bmatrix}$$

■

5.3 The Secret Sauce of LU (Lower-Upper) Factorization: Peeling the Onion

As a lead in to our main topic, note that in Example 5.1, A is a lower triangular matrix, B is an upper triangular matrix, while their product $A \cdot B$ is neither. Can this process be reversed? That is, given a generic square matrix, can we factor it as the product of a lower-triangular matrix and an upper-triangular matrix? And even if we can do such a factorization, would it be helpful?

We'll delay an explicit answer to the question of utility because you have a sense already that triangular matrices make your life easier. Our goal here is to show you the secret sauce the underlies a very nice method for constructing the required triangular matrices. We call it **peeling the onion: working from the top left corner and working down the diagonal, it successively eliminates columns and rows from a matrix!** Watch, it's cool, it's kind of fun, and it relies heavily on the column times row form of matrix multiplication.

Peeling the Onion: Consider the square matrix

$$M = \begin{bmatrix} 1 & 4 & 5 \\ 2 & 9 & 17 \\ 3 & 18 & 58 \end{bmatrix}.$$

Our goal is to find a column vector C_1 and a row vector R_1 such that

$$M - C_1 \cdot R_1 = \begin{bmatrix} 0 & 0 & 0 \\ 0 & * & * \\ 0 & * & * \end{bmatrix},$$

where $*$ denotes “don't care” in the sense that we do not care about their particular values. We want to zero out the first column and the first row of M . That means, C_1 and R_1 are chosen so that the first column and first row of their matrix product $C_1 \cdot R_1$ match the first column and first row of M . How can you do that?

Here's a **special case of the method that works when the top left entry equals 1.0**. We define C_1 and R_1 to be the first column of M and the first row of M , respectively, that is

$$C_1 = \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix} \text{ and } R_1 = [1 \quad 4 \quad 5].$$

Then

$$C_1 \cdot R_1 = \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix} \cdot [1 \quad 4 \quad 5] = \begin{bmatrix} 1 & 4 & 5 \\ 2 & 8 & 10 \\ 3 & 12 & 15 \end{bmatrix},$$

and voilà,

$$M = \begin{bmatrix} 1 & 4 & 5 \\ 2 & 9 & 17 \\ 3 & 18 & 58 \end{bmatrix} \text{ and } C_1 \cdot R_1 = \begin{bmatrix} 1 & 4 & 5 \\ 2 & 8 & 10 \\ 3 & 12 & 15 \end{bmatrix}.$$

Consequently,

$$\begin{aligned} M - C_1 \cdot R_1 &= \begin{bmatrix} 1 & 4 & 5 \\ 2 & 9 & 17 \\ 3 & 18 & 58 \end{bmatrix} - \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix} \cdot [1 \quad 4 \quad 5] \\ &= \begin{bmatrix} 1 & 4 & 5 \\ 2 & 9 & 17 \\ 3 & 18 & 58 \end{bmatrix} - \begin{bmatrix} 1 & 4 & 5 \\ 2 & 8 & 10 \\ 3 & 12 & 15 \end{bmatrix} \\ &= \begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 7 \\ 0 & 6 & 43 \end{bmatrix}. \end{aligned}$$

Oh! We have taken a 3×3 matrix and essentially made it into a 2×2 matrix!! Can we do this again? Let's try. We define C_2 and R_2 to be the second column and second row of $M - C_1 \cdot R_1$, that is

$$C_2 = \begin{bmatrix} 0 \\ 1 \\ 6 \end{bmatrix} \text{ and } R_2 = [0 \quad 1 \quad 7].$$

Then we compute that

$$\begin{bmatrix} 0 \\ 1 \\ 6 \end{bmatrix} \cdot [0 \ 1 \ 7] = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 7 \\ 0 & 6 & 42 \end{bmatrix},$$

and we obtain

$$(M - C_1 \cdot R_1) = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 7 \\ 0 & 6 & 43 \end{bmatrix} \text{ and } C_2 \cdot R_2 = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 7 \\ 0 & 6 & 42 \end{bmatrix}.$$

Consequently,

$$\begin{aligned} (M - C_1 \cdot R_1) - C_2 \cdot R_2 &= \begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 7 \\ 0 & 6 & 43 \end{bmatrix} - \begin{bmatrix} 0 \\ 1 \\ 6 \end{bmatrix} \cdot [0 \ 1 \ 7] \\ &= \begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 7 \\ 0 & 6 & 43 \end{bmatrix} - \begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 7 \\ 0 & 6 & 42 \end{bmatrix} \\ &= \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix} \end{aligned}$$

Oh! Now we are essentially down to a 1×1 matrix!! You might be seeing the pattern! We very quickly note that if we define C_3 and R_3 to be the third column and third row of $M - C_1 \cdot R_1 - C_2 \cdot R_2$,

$$C_3 = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \text{ and } R_3 = [0 \ 0 \ 1],$$

then

$$C_3 \cdot R_3 = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix},$$

and hence, $M - C_1 \cdot R_1 - C_2 \cdot R_2 - C_3 \cdot R_3 = 0_{3 \times 3}$. We prefer to write this as

$$M = C_1 \cdot R_1 + C_2 \cdot R_2 + C_3 \cdot R_3 = \underbrace{\begin{bmatrix} C_1 & C_2 & C_3 \end{bmatrix}}_L \cdot \underbrace{\begin{bmatrix} R_1 \\ R_2 \\ R_3 \end{bmatrix}}_U.$$

Moreover,

- $L := [C_1 \ C_2 \ C_3] = \begin{bmatrix} 1 & 0 & 0 \\ 2 & 1 & 0 \\ 3 & 6 & 1 \end{bmatrix}$ is **lower triangular**,

- $U := \begin{bmatrix} R_1 \\ R_2 \\ R_3 \end{bmatrix} = \begin{bmatrix} 1 & 4 & 5 \\ 0 & 1 & 7 \\ 0 & 0 & 1 \end{bmatrix}$ is **upper triangular**, and

- $M = L \cdot U$, the product of a lower triangular matrix and an upper triangular matrix! We did it!

The Need for Pivots

Is that all there is to LU Factorization? No, there's a bit more to it. The real algorithm has a "normalization step". To motivate it we take a matrix M that has something other than a one in its a_{11} -entry and naively form " C_1 " as the first column of the matrix and " R_1 " as the first row of the matrix. Then

$$\underbrace{\begin{bmatrix} 2 & 3 \\ 4 & 5 \end{bmatrix}}_M - \underbrace{\begin{bmatrix} 2 \\ 4 \end{bmatrix}}_{C_1} \cdot \underbrace{\begin{bmatrix} 2 & 3 \end{bmatrix}}_{R_1} = \underbrace{\begin{bmatrix} 2 & 3 \\ 4 & 5 \end{bmatrix}}_M - \underbrace{\begin{bmatrix} 4 & 6 \\ 8 & 12 \end{bmatrix}}_{C_1 \cdot R_1} = \begin{bmatrix} -2 & -3 \\ -4 & -7 \end{bmatrix},$$

which does NOT result in a matrix with zeros in its leading column and row! However, if we keep R_1 as the first row of M but this time, we form C_1 from the first column of M normalized by its first entry, then we can make it work. Watch!

$$\underbrace{\begin{bmatrix} 2 & 3 \\ 4 & 5 \end{bmatrix}}_M - \underbrace{\begin{bmatrix} 2/2 \\ 4/2 \end{bmatrix}}_{C_1} \cdot \underbrace{\begin{bmatrix} 2 & 3 \end{bmatrix}}_{R_1} = \underbrace{\begin{bmatrix} 2 & 3 \\ 4 & 5 \end{bmatrix}}_M - \underbrace{\begin{bmatrix} 1 \\ 2 \end{bmatrix}}_{C_1} \cdot \underbrace{\begin{bmatrix} 2 & 3 \end{bmatrix}}_{R_1} = \underbrace{\begin{bmatrix} 2 & 3 \\ 4 & 5 \end{bmatrix}}_M - \underbrace{\begin{bmatrix} 2 & 3 \\ 4 & 6 \end{bmatrix}}_{C_1 \cdot R_1} = \begin{bmatrix} 0 & 0 \\ 0 & -1 \end{bmatrix}.$$

The a_{11} entry of M is called a **pivot**.

Warning: When we do the normalization and divide by the pivot, we could get into trouble with a divide by zero! We will worry about this later. First things first!

We'll illustrate a more general case of **peeling the onion** and then do it in Julia code. **More generally**, if we assume $a_{11} \neq 0$ and use it as the pivot

$$\begin{bmatrix} \mathbf{a}_{11} & \mathbf{a}_{12} & \cdots & \mathbf{a}_{1n} \\ \mathbf{a}_{21} & a_{22} & \cdots & a_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{a}_{n1} & a_{n2} & \cdots & a_{nn} \end{bmatrix} - \begin{bmatrix} \mathbf{a}_{11}/\mathbf{a}_{11} \\ \mathbf{a}_{21}/\mathbf{a}_{11} \\ \vdots \\ \mathbf{a}_{n1}/\mathbf{a}_{11} \end{bmatrix} \cdot \begin{bmatrix} \mathbf{a}_{11} & \mathbf{a}_{12} & \cdots & \mathbf{a}_{1n} \end{bmatrix} = \begin{bmatrix} \mathbf{a}_{11} & \mathbf{a}_{12} & \cdots & \mathbf{a}_{1n} \\ \mathbf{a}_{21} & a_{22} & \cdots & a_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{a}_{n1} & a_{n2} & \cdots & a_{nn} \end{bmatrix} - \begin{bmatrix} \mathbf{a}_{11} & \mathbf{a}_{12} & \cdots & \mathbf{a}_{1n} \\ \mathbf{a}_{21} & * & \cdots & * \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{a}_{n1} & * & \cdots & * \end{bmatrix} \\ = \begin{bmatrix} \mathbf{0} & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & * & \cdots & * \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & * & \cdots & * \end{bmatrix}$$

```
1 # Run me don't change me
2 using Random
3 using LinearAlgebra
4 Random.seed!(09182021)
5 A=randn(3,3)
6 Temp=copy(A)
```

Output because $Temp$ is the last thing in the cell, only it is printed

```
3×3 Matrix{Float64}:
 2.85947  0.0572321 -1.12303
 0.139854 -0.904663  1.57054
-0.59195  0.43911  0.108526
```

```
1 # Run me don't change me
2 # Note that we have defined pivot as the one-one-entry of Temp
3 pivot = Temp[1,1]
4 C1=Temp[:,1]/pivot
5 R1=Temp[1:1,: ]
6 @show C1*R1
```

```
7 Temp=Temp-C1*R1
```

Output You can compare the first row and first column of **Temp** with $C1 * R1$. The last thing printed is the new version of **Temp**.

```
C1 * R1 = [2.8594704859391835 0.05723214703413922 -1.1230256584654308;  
0.13985399718460348 0.0027991701853682135 -0.054926122878193;  
-0.591949779941339 -0.011847842811814082 0.23248178103812325]
```

```
3x3 Matrix{Float64}:  
 0.0  0.0  0.0  
 0.0 -0.907462  1.62546  
 0.0  0.450958 -0.123956
```

We now look at the second row and column of **Temp** and call **Temp[2,2]** the **pivot** (we note it is non-zero).

```
1 pivot = Temp[2,2]  
2 C2 = Temp[:,2]/pivot  
3 R2 = Temp[2:2,:]  
4 @show (C2*R2)  
5 Temp = Temp - (C2*R2)
```

Output

```
C2 * R2 = [0.0 0.0 0.0; 0.0 -0.9074620727541318 1.6254630955038092;  
0.0 0.45095794284898383 -0.8077643305803952]
```

```
3x3 Matrix{Float64}:  
 0.0  0.0  0.0  
 0.0  0.0  0.0  
 0.0  0.0  0.683809
```

We next look at the third row and column of **Temp** and call **Temp[3,3]** the **pivot** (we note it is non-zero).

```
1 pivot = Temp[3,3]  
2 C3 = Temp[:,3]/pivot  
3 R3 = Temp[3:3,:]  
4 Temp = Temp - (C3*R3)
```

Output

```
C3 * R3 = [0.0 0.0 0.0; 0.0 0.0 0.0; 0.0 0.0 0.6838085188740579]
```

```
3x3 Matrix{Float64}:  
 0.0  0.0  0.0  
 0.0  0.0  0.0  
 0.0  0.0  0.0
```

We now put all of this in a **for loop** to show how we can successively zero the columns and rows of a matrix.

```
1 Temp=copy(A)  
2 nRows, nCols = size(Temp)  
3 for k = 1:nRows  
4     @show Temp  
5     pivot=Temp[k,k]  
6     # boldly assume we never divide by zero  
7     C=Temp[:,k]/pivot  
8     R=Temp[k:k,:]  
9     Temp=Temp-C*R  
10 end  
11 @show Temp
```

Output

```
Temp = [2.8594704859391835 0.05723214703413922 -1.1230256584654308;  
0.13985399718460348 -0.9046629025687636 1.5705369726256162;  
-0.591949779941339 0.43911010003716977 0.108525969331786]
```

```
Temp = [0.0 0.0 0.0;  
0.0 -0.9074620727541318 1.6254630955038092;  
0.0 0.45095794284898383 -0.12395581170633725]
```

```
Temp = [0.0 0.0 0.0;  
0.0 0.0 0.0;  
0.0 0.0 0.6838085188740579]
```

```
Temp = [0.0 0.0 0.0; 0.0 0.0 0.0; 0.0 0.0 0.0]
```

5.4 (Optional Read:) Peeling the Onion in Pictures

Summary: Peeling the Onion

$$M_k = \begin{bmatrix} 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & \bar{a}_{kk} & \cdots & \bar{a}_{kn} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \bar{a}_{nk} & & \bar{a}_{nn} \end{bmatrix}$$

(k-1) rows of zeros

(k-1) cols of zeros

If $\bar{a}_{kk} \neq 0$, then $p_k = \bar{a}_{kk}$, $C_k = \frac{1}{p_k} \begin{bmatrix} 0 \\ \vdots \\ 0 \\ \bar{a}_{kk} \\ \vdots \\ \bar{a}_{nk} \end{bmatrix} = \begin{bmatrix} 0 \\ \vdots \\ 0 \\ 1 \\ \vdots \\ \frac{\bar{a}_{nk}}{\bar{a}_{kk}} \end{bmatrix}$

$R_k = [0 \cdots 0 \bar{a}_{kk} \cdots \bar{a}_{kn}]$ results in

$$M_k - C_k R_k = \begin{bmatrix} 0 & 0 & 0 & \dots & 0 \\ 0 & 0 & 0 & \dots & 0 \\ 0 & 0 & a_{kk} & \dots & a_{kn} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & a_{nk} & \dots & a_{nn} \end{bmatrix} - \begin{bmatrix} 0 & 0 & 0 & \dots & 0 \\ 0 & 0 & 0 & \dots & 0 \\ 0 & 0 & a_{kk} & \dots & a_{kn} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & a_{nk} & \dots & a_{nn} \end{bmatrix}$$

$$= \begin{bmatrix} 0 & 0 & 0 & \dots & 0 \\ 0 & 0 & 0 & \dots & 0 \\ 0 & 0 & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & 0 \end{bmatrix} \left. \vphantom{\begin{bmatrix} 0 & 0 & 0 & \dots & 0 \\ 0 & 0 & 0 & \dots & 0 \\ 0 & 0 & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & 0 \end{bmatrix}} \right\} k \text{ rows}$$

$$= \begin{bmatrix} 0 & 0 & 0 & \dots & 0 \\ 0 & 0 & 0 & \dots & 0 \\ 0 & 0 & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & 0 \end{bmatrix} \begin{bmatrix} \bar{a}_{(k+1)k} & \dots & \bar{a}_{(k+1)n} \\ \vdots & \ddots & \vdots \\ \bar{a}_{nk} & \dots & \bar{a}_{nn} \end{bmatrix}$$

$$A - C_1 R_1 - C_2 R_2 - \dots - C_n R_n = O_{n \times n}$$

∴

$$A = C_1 R_1 + C_2 R_2 + \dots + C_n R_n$$

$$= \begin{bmatrix} 1 & 0 & \dots & 0 \\ l_{21} & 1 & & \\ \vdots & \vdots & \ddots & \\ l_{n1} & l_{n2} & \dots & 1 \end{bmatrix} \cdot \begin{bmatrix} u_{11} & u_{12} & \dots & u_{1n} \\ 0 & u_{22} & \dots & u_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & u_{nn} \end{bmatrix}$$

Lab 4 should clear up any lingering issues.

5.5 LU (Lower-Upper) Factorization (without row permutations)

We state our first algorithm

LU Factorization (without permutations)

Algorithm 1: LU-Factorization (without permutations)

Result: For M an $n \times n$ matrix, with $n \geq 2$, find L and U such that $M = LU$

```
# initialization:
Temp=copy(M);
L = Array{Float64,2}(undef, n, 0) # L=[] Empty matrix
U = Array{Float64,2}(undef, 0, n) # U=[] Empty matrix
# end initialization:

for k=1:n
    pivot = Temp[k,k]
    if ! isapprox(pivot, 0, atol = 1E-8) # check for zero
        C=Temp[:,k] ./ pivot # normalize so that k-th entry is equal to 1.0
        R=Temp[k:k,:]
        # Alternative also works in Julia for k-th row
        # R=Temp[k,:]' k-th row, which in Julia, requires a transpose
        Temp=Temp-C*R;
        L=[L C] # Build the lower-triangular matrix by columns
        U=[U;R] # Build the upper-triangular matrix by rows
    else
        # pivot equals zero and we do not want to divide by zero
        println("Matrix requires row permutations to avoid divide by zero")
        println("Step where algorithm failed is k= $k")
        break # Jump out of the for loop and terminate the algorithm
    end
end

return L, U
```

The algorithm is easy to do by hand for very small matrices. In Julia, you can write code that will do the LU factorization of a 1000×1000 matrix in a few seconds. The `lu` function in the `LinearAlgebra` package of Julia can do it in a few milliseconds! The exact command is

```
1 using LinearAlgebra
2 L, U = lu(M, Val(false))
```

to return L and U **without permutations**. If you leave off “`Val(false)`”, then the algorithm uses row permutations and you should call it as

```
1 using LinearAlgebra
2 F = lu(M)
3 L=F.L
4 U=F.U
5 P=F.P
```

Example 5.2 (2×2 Matrix) Perform the LU Factorization of $M = \begin{bmatrix} 5 & 2 \\ 15 & 2 \end{bmatrix}$, a 2×2 square matrix.

Solution: We'll do every step, just as you would program it up.

We *initialize* our algorithm by

$$\text{Temp} := M = \begin{bmatrix} 5 & 2 \\ 15 & 2 \end{bmatrix}, L := [\text{empty matrix}], \text{ and } U := [\text{empty matrix}]$$

k=1: We start with the pivot is equal to $\text{Temp}[1, 1]$ and C as the first column of Temp scaled by the pivot, and thus¹

$$\begin{aligned}\text{pivot} &= \text{Temp}[1, 1] = 5 \\ C &= \text{Temp}[:, 1]/\text{pivot} \\ &= \begin{bmatrix} 5 \\ 15 \end{bmatrix} \cdot \frac{1}{5} \\ &= \begin{bmatrix} 1 \\ 3 \end{bmatrix},\end{aligned}$$

while R is the first row of Temp

$$R = \text{Temp}[1 : 1, :] = \begin{bmatrix} 5 & 2 \end{bmatrix}.$$

We finish up the first step by defining

$$\begin{aligned}\text{Temp} &:= \text{Temp} - C \cdot R \\ &= \begin{bmatrix} 5 & 2 \\ 15 & 2 \end{bmatrix} - \begin{bmatrix} 1 \\ 3 \end{bmatrix} \cdot \begin{bmatrix} 5 & 2 \end{bmatrix} \\ &= \begin{bmatrix} 5 & 2 \\ 15 & 2 \end{bmatrix} - \begin{bmatrix} 5 & 2 \\ 15 & 6 \end{bmatrix} \\ &= \begin{bmatrix} 0 & 0 \\ 0 & -4 \end{bmatrix}\end{aligned}$$

$$\begin{aligned}L &= [L, C] = \begin{bmatrix} 1 \\ 3 \end{bmatrix} \\ U &= [U; R] = \begin{bmatrix} 5 & 2 \end{bmatrix}\end{aligned}$$

Why LU Factorization Works: Peeling the Onion

We observe that L and U were defined so that their product exactly matched the first column and row of M . Hence, when we subtract them from M , we are left with a problem that is effectively one dimension lower, meaning, in this case, the non-zero part is 1×1 .

While you may be able to finish the factorization by “inspection”, we will do the complete algorithm just like you would do in a program.

k=2: The pivot is the second entry on the diagonal of Temp while C is the second column of Temp scaled by the pivot

$$\begin{aligned}\text{pivot} &= \text{Temp}[2, 2] = -4 \\ C &= \text{Temp}[:, 2]/\text{pivot} \\ &= \begin{bmatrix} 0 \\ -4 \end{bmatrix} \cdot \frac{-1}{4} \\ &= \begin{bmatrix} 0 \\ 1 \end{bmatrix},\end{aligned}$$

while R is the second row of Temp ,

$$R = \text{Temp}[2 : 2, :] = \begin{bmatrix} 0 & -4 \end{bmatrix}.$$

¹In HW solutions or exams, it is fine to skip a step and write down L directly, however, it is then more difficult to give you a lot of partial credit.

$$\begin{aligned}
\text{Temp} &:= \text{Temp} - C \cdot R \\
&= \begin{bmatrix} 0 & 0 \\ 0 & -4 \end{bmatrix} - \begin{bmatrix} 0 \\ 1 \end{bmatrix} \cdot \begin{bmatrix} 0 & -4 \end{bmatrix} \\
&= \begin{bmatrix} 0 & 0 \\ 0 & -4 \end{bmatrix} - \begin{bmatrix} 0 & 0 \\ 0 & -4 \end{bmatrix} \\
&= \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}
\end{aligned}$$

$$L = [L; C] = \begin{bmatrix} 1 & 0 \\ 3 & 1 \end{bmatrix} = \text{lower triangular matrix}$$

$$U = [U; R] = \begin{bmatrix} 5 & 2 \\ 0 & -4 \end{bmatrix} = \text{upper triangular matrix}$$

Peeling the Onion:

The matrix product $L \cdot U$ is equal to the sum of the columns of L times the rows of U . The columns and rows were iteratively designed to remove columns and rows from M . Indeed,

$$\begin{aligned}
\text{Temp} &:= M - L \cdot U \\
&= \begin{bmatrix} 5 & 2 \\ 15 & 2 \end{bmatrix} - \begin{bmatrix} 1 & 0 \\ 3 & 1 \end{bmatrix} \cdot \begin{bmatrix} 5 & 2 \\ 0 & -4 \end{bmatrix} \\
&= \begin{bmatrix} 5 & 2 \\ 15 & 2 \end{bmatrix} - \left(\begin{bmatrix} 1 \\ 3 \end{bmatrix} \cdot \begin{bmatrix} 5 & 2 \end{bmatrix} + \begin{bmatrix} 0 \\ 1 \end{bmatrix} \cdot \begin{bmatrix} 0 & -4 \end{bmatrix} \right) \\
&= \begin{bmatrix} 5 & 2 \\ 15 & 2 \end{bmatrix} - \left(\begin{bmatrix} 5 & 2 \\ 15 & 6 \end{bmatrix} + \begin{bmatrix} 0 & 0 \\ 0 & -4 \end{bmatrix} \right) \\
&= \begin{bmatrix} 5 & 2 \\ 15 & 2 \end{bmatrix} - \begin{bmatrix} 5 & 2 \\ 15 & 6 \end{bmatrix} - \begin{bmatrix} 0 & 0 \\ 0 & -4 \end{bmatrix} \\
&= \left(\begin{bmatrix} 5 & 2 \\ 15 & 2 \end{bmatrix} - \begin{bmatrix} 5 & 2 \\ 15 & 6 \end{bmatrix} \right) - \begin{bmatrix} 0 & 0 \\ 0 & -4 \end{bmatrix} \\
&= \left(\begin{bmatrix} 0 & 0 \\ 0 & -4 \end{bmatrix} \right) - \begin{bmatrix} 0 & 0 \\ 0 & -4 \end{bmatrix} \\
&= \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}
\end{aligned}$$

Hence,

$$\underbrace{\begin{bmatrix} 5 & 2 \\ 15 & 2 \end{bmatrix}}_M = \underbrace{\begin{bmatrix} 1 & 0 \\ 3 & 1 \end{bmatrix}}_L \cdot \underbrace{\begin{bmatrix} 5 & 2 \\ 0 & -4 \end{bmatrix}}_U$$

■

Example 5.3 (3×3 Matrix) Perform the LU Factorization of $M = \begin{bmatrix} -2 & -4 & -6 \\ -2 & 1 & -4 \\ -2 & 11 & -4 \end{bmatrix}$, a 3×3 square matrix.

Solution: We initialize our algorithm by

$$\text{Temp} := M = \begin{bmatrix} -2 & -4 & -6 \\ -2 & 1 & -4 \\ -2 & 11 & -4 \end{bmatrix}, L := [\text{empty matrix}], \text{ and } U := [\text{empty matrix}]$$

k=1: We start with $\text{pivot} = T[1, 1]$ and C as first column of Temp normalized by the pivot so that its first entry is one

$$\begin{aligned}\text{pivot} &= \text{Temp}[1, 1] = -2 \\ C &= \text{Temp}[:, 1] / \text{pivot} \\ &= \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix},\end{aligned}$$

while R is the first row of Temp

$$R = \text{Temp}[1 : 1, :] = \begin{bmatrix} -2 & -4 & -6 \end{bmatrix}.$$

We finish up the first step by defining

$$\begin{aligned}\text{Temp} &:= \text{Temp} - C \cdot R \\ &= \begin{bmatrix} -2 & -4 & -6 \\ -2 & 1 & -4 \\ -2 & 11 & -4 \end{bmatrix} - \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} \cdot \begin{bmatrix} -2 & -4 & -6 \end{bmatrix} \\ &= \begin{bmatrix} -2 & -4 & -6 \\ -2 & 1 & -4 \\ -2 & 11 & -4 \end{bmatrix} - \begin{bmatrix} -2 & -4 & -6 \\ -2 & -4 & -6 \\ -2 & -4 & -6 \end{bmatrix} \\ &= \begin{bmatrix} 0 & 0 & 0 \\ 0 & 5 & 2 \\ 0 & 15 & 2 \end{bmatrix}\end{aligned}$$

$$\begin{aligned}L &= [L, C] = \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} \\ U &= [U; R] = \begin{bmatrix} -2 & -4 & -6 \end{bmatrix}\end{aligned}$$

Reminder of Why LU Factorization Works: Peeling the Onion

We observe that L and U have been constructed so that their product exactly matches the first column and first row of M . Hence, when we subtract $L \cdot U$ from M , we are left with a problem that is effectively one dimension lower, meaning, in this case, the non-zero part is 2×2 .

Your wily instructor has arranged for the 2×2 matrix that remains to be an example that we have already worked! Nevertheless, we will complete the algorithmic steps just like you would do in Julia.

k=2: C is the second column of Temp scaled by its $k = 2$ entry,

$$\begin{aligned}\text{pivot} &= \text{Temp}[2, 2] = 5 \\ C &= \begin{bmatrix} 0 \\ 5 \\ 15 \end{bmatrix} / \text{pivot} \\ &= \begin{bmatrix} 0 \\ 5 \\ 15 \end{bmatrix} \cdot \frac{1}{5} \\ &= \begin{bmatrix} 0 \\ 1 \\ 3 \end{bmatrix},\end{aligned}$$

while R is the second row of Temp ,

$$R = \text{Temp}[2 : 2, 2] = \begin{bmatrix} 0 & 5 & 2 \end{bmatrix}.$$

$$\begin{aligned}
\text{Temp} &:= \text{Temp} - C \cdot R \\
&= \begin{bmatrix} 0 & 0 & 0 \\ 0 & 5 & 2 \\ 0 & 15 & 2 \end{bmatrix} - \begin{bmatrix} 0 \\ 1 \\ 3 \end{bmatrix} \cdot [0 \ 5 \ 2] \\
&= \begin{bmatrix} 0 & 0 & 0 \\ 0 & 5 & 2 \\ 0 & 15 & 2 \end{bmatrix} - \begin{bmatrix} 0 & 0 & 0 \\ 0 & 5 & 2 \\ 0 & 15 & 6 \end{bmatrix} \\
&= \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & -4 \end{bmatrix}
\end{aligned}$$

$$\begin{aligned}
L = [L, C] &= \begin{bmatrix} 1 & 0 \\ 1 & 1 \\ 1 & 3 \end{bmatrix} \\
U = [U; R] &= \begin{bmatrix} -2 & -4 & -6 \\ 0 & 5 & 2 \end{bmatrix}
\end{aligned}$$

As before, we do the complete algorithm just like you would do in Julia.

k=3: C is the third column of Temp scaled by the (3,3) entry,

$$\begin{aligned}
\text{pivot} &= \text{Temp}[3, 3] = -4 \\
C &= \begin{bmatrix} 0 \\ 0 \\ -4 \end{bmatrix} / \text{pivot} \\
&= \begin{bmatrix} 0 \\ 0 \\ -4 \end{bmatrix} \cdot \frac{1}{-4} \\
&= \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix},
\end{aligned}$$

while R is the third row of Temp,

$$R = [0 \ 0 \ -4].$$

$$\begin{aligned}
\text{Temp} &:= \text{Temp} - C \cdot R \\
&= \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & -4 \end{bmatrix} - \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \cdot [0 \ 0 \ -4] \\
&= \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & -4 \end{bmatrix} - \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & -4 \end{bmatrix} \\
&= \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}
\end{aligned}$$

$$L = [L; C] = \begin{bmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ 1 & 3 & 1 \end{bmatrix} = \text{lower triangular matrix}$$

$$U = [U; R] = \begin{bmatrix} -2 & -4 & -6 \\ 0 & 5 & 2 \\ 0 & 0 & -4 \end{bmatrix} = \text{upper triangular matrix}$$

Hence,

$$\underbrace{\begin{bmatrix} -2 & -4 & -6 \\ -2 & 1 & -4 \\ -2 & 11 & -4 \end{bmatrix}}_M = \underbrace{\begin{bmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ 1 & 3 & 1 \end{bmatrix}}_L \cdot \underbrace{\begin{bmatrix} -2 & -4 & -6 \\ 0 & 5 & 2 \\ 0 & 0 & -4 \end{bmatrix}}_U$$

■

5.6 LU Factorization for Solving Linear Equations

Solving $Ax = b$ via LU without Row Permutations

We seek to solve the system of linear equations $Ax = b$. Suppose that we can factor $A = L \cdot U$, where U is upper triangular and L is lower triangular. Hence, we are solving the equation

$$L \cdot Ux = b. \tag{5.5}$$

If we define $Ux = y$, then (5.5) becomes two equations

$$Ly = b \tag{5.6}$$

$$Ux = y. \tag{5.7}$$

Our solution strategy is to solve (5.6) by forward substitution, and then, once we have y in hand, we solve (5.7) by back substitution to find x , the solution to (5.5).

Major Important Fact

This process may seem a bit long and tedious when doing it by hand. On a computer, it is a snap. Once you have your Julia functions created, you'll see that this scales to big problems in a very nice way.

Example 5.4 Use LU Factorization to solve the system of linear equations

$$\underbrace{\begin{bmatrix} -2 & -4 & -6 \\ -2 & 1 & -4 \\ -2 & 11 & -4 \end{bmatrix}}_A \underbrace{\begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}}_x = \underbrace{\begin{bmatrix} 2 \\ 3 \\ -7 \end{bmatrix}}_b. \quad (5.8)$$

Solution: From our hard work above, we know that

$$A = \underbrace{\begin{bmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ 1 & 3 & 1 \end{bmatrix}}_L \cdot \underbrace{\begin{bmatrix} -2 & -4 & -6 \\ 0 & 5 & 2 \\ 0 & 0 & -4 \end{bmatrix}}_U \quad (5.9)$$

We first solve, using forward substitution,

$$\underbrace{\begin{bmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ 1 & 3 & 1 \end{bmatrix}}_L \underbrace{\begin{bmatrix} y_1 \\ y_2 \\ y_3 \end{bmatrix}}_y = \underbrace{\begin{bmatrix} 2 \\ 3 \\ -7 \end{bmatrix}}_b \implies \begin{bmatrix} y_1 \\ y_2 \\ y_3 \end{bmatrix} = \begin{bmatrix} 2 \\ 1 \\ -12 \end{bmatrix}.$$

And then we use this result to solve, using back substitution,

$$\underbrace{\begin{bmatrix} -2 & -4 & -6 \\ 0 & 5 & 2 \\ 0 & 0 & -4 \end{bmatrix}}_U \underbrace{\begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}}_x = \underbrace{\begin{bmatrix} 2 \\ 1 \\ -12 \end{bmatrix}}_y \implies \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} -8 \\ -1 \\ 3 \end{bmatrix}.$$

■

Do we need row permutations at all? The next section develops the LU Factorization with row Permutations. Why is that needed? **Some matrices do not have an LU Factorization without row permutations.** Suppose we wanted to compute an LU Factorization for

$$A = \begin{bmatrix} 0 & 1 \\ 2 & 3 \end{bmatrix}$$

with our current algorithm. We'd define

$$C = \begin{bmatrix} 0 \\ 2 \end{bmatrix} \cdot \frac{1}{0} = \text{undefined}, \quad R = \begin{bmatrix} 0 & 1 \end{bmatrix},$$

so we are stuck. If we permute rows one and two however, we'd have

$$P = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \\ P \cdot A = \begin{bmatrix} 2 & 3 \\ 0 & 1 \end{bmatrix}$$

which is already factored because we can define

$$L = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \\ U = \begin{bmatrix} 2 & 3 \\ 0 & 1 \end{bmatrix},$$

giving us $P \cdot A = L \cdot U$.

Solving $Ax = b$ via LU with Row Permutations

We once again seek to solve the system of linear equations $Ax = b$. Suppose this time we can factor $P \cdot A = L \cdot U$, where P is a permutation matrix, U is upper triangular and L is lower triangular. Would that even be helpful for solving linear equations?

A beautiful property of permutation matrices is that $\det(P) = \pm 1$, which means they are always invertible. This property leads to

$$Ax = b \iff P \cdot Ax = P \cdot b \iff L \cdot Ux = P \cdot b.$$

Hence, you can think of our previous result as the special case of P being the identity matrix, I_n . When row permutations are involved in the LU Factorization, we solve the slightly modified equation

$$L \cdot Ux = P \cdot b. \quad (5.10)$$

If we define $Ux = y$, then (5.10) again becomes two equations

$$Ly = P \cdot b \quad (5.11)$$

$$Ux = y. \quad (5.12)$$

Our solution strategy is to solve (5.11) by forward substitution, and then, once we have y in hand, we solve (5.12) by back substitution to find x , the solution to (5.10).

Example 5.5 Use LU Factorization with permutations to solve the system of linear equations

$$\underbrace{\begin{bmatrix} -2 & -4 & -6 \\ -2 & 1 & -4 \\ -2 & 11 & -4 \end{bmatrix}}_A \underbrace{\begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}}_x = \underbrace{\begin{bmatrix} 2 \\ 3 \\ -7 \end{bmatrix}}_b. \quad (5.13)$$

Solution: This time we use the native LU function in Julia to compute $P \cdot A = L \cdot U$, with

$$\begin{aligned} P &= \begin{bmatrix} 1.0 & 0.0 & 0.0 \\ 0.0 & 0.0 & 1.0 \\ 0.0 & 1.0 & 0.0 \end{bmatrix} \\ L &= \begin{bmatrix} 1.000 & 0.000 & 0.000 \\ 1.000 & 1.000 & 0.000 \\ 1.000 & 0.333 & 1.000 \end{bmatrix} \\ U &= \begin{bmatrix} -2.000 & -4.000 & -6.000 \\ 0.000 & 15.000 & 2.000 \\ 0.000 & 0.000 & 1.333 \end{bmatrix}. \end{aligned} \quad (5.14)$$

Even though A admits an LU Factorization without row permutations, Julia inserts a permutation matrix. This is to improve the numerical accuracy on large problems. On our small problem, it's not really needed. Nevertheless, we'll use it to show that we obtain the same answer with essentially the same amount of work.

We first compute

$$Pb = \begin{bmatrix} 1.0 & 0.0 & 0.0 \\ 0.0 & 0.0 & 1.0 \\ 0.0 & 1.0 & 0.0 \end{bmatrix} \begin{bmatrix} 2.0 \\ 3.0 \\ -7.0 \end{bmatrix} = \begin{bmatrix} 2.0 \\ -7.0 \\ 3.0 \end{bmatrix}.$$

We then solve $Ly = Pb$ for the intermediate variable y , using forward substitution,

$$\underbrace{\begin{bmatrix} 1.000 & 0.000 & 0.000 \\ 1.000 & 1.000 & 0.000 \\ 1.000 & 0.333 & 1.000 \end{bmatrix}}_L \underbrace{\begin{bmatrix} y_1 \\ y_2 \\ y_3 \end{bmatrix}}_y = \underbrace{\begin{bmatrix} 2.0 \\ -7.0 \\ 3.0 \end{bmatrix}}_{Pb} \implies \begin{bmatrix} y_1 \\ y_2 \\ y_3 \end{bmatrix} = \begin{bmatrix} 2.0 \\ -9.0 \\ 4.0 \end{bmatrix}.$$

And finally, we use this result to solve $Ux = y$ for x , using back substitution,

$$\underbrace{\begin{bmatrix} -2.000 & -4.000 & -6.000 \\ 0.000 & 15.000 & 2.000 \\ 0.000 & 0.000 & 1.333 \end{bmatrix}}_U \underbrace{\begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}}_x = \underbrace{\begin{bmatrix} 2.0 \\ -9.0 \\ 4.0 \end{bmatrix}}_y \implies \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} -8.0 \\ -1.0 \\ 3.0 \end{bmatrix}.$$

■

5.7 (Optional Read): Toward LU Factorization with Row Permutations

Are there Cases where our Simplified LU Factorization Process Fails?

Yes. They are not that hard to handle, but as your first introduction to LU factorization, you do not have to learn all of the nitty-gritty details. You've done well to get this far. You may wish to use the following in a project, in some other course, or just to understand what the Julia `lu` function is doing.

Case 1 (easiest): C becomes an entire column of zeros Suppose we have $M = \begin{bmatrix} 2 & -1 & 2 \\ 6 & -3 & 9 \\ 2 & -1 & 6 \end{bmatrix}$. To start off the process, we define

$$C = \begin{bmatrix} 2 \\ 6 \\ 2 \end{bmatrix} \cdot \frac{1}{2} = \begin{bmatrix} 1 \\ 3 \\ 1 \end{bmatrix}, \quad R = [2 \quad -1 \quad 2],$$

and arrive at

$$\begin{aligned} Temp = M - C \cdot R &= \begin{bmatrix} 2 & -1 & 2 \\ 6 & -3 & 9 \\ 2 & -1 & 6 \end{bmatrix} - \begin{bmatrix} 2 & -1 & 2 \\ 6 & -3 & 6 \\ 2 & -1 & 2 \end{bmatrix} \\ &= \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 3 \\ 0 & 0 & 4 \end{bmatrix} \end{aligned}$$

The next step would be

$$C = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix} \cdot \frac{1}{0} = \begin{bmatrix} \text{undefined} \\ \text{undefined} \\ \text{undefined} \end{bmatrix}, \quad R = [0 \quad 0 \quad 3],$$

leading us to attempt a divide by zero. The solution is to define

$$C = \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}, \quad \text{while maintaining } R = [0 \quad 0 \quad 3].$$

With this definition,

$$\begin{aligned} Temp - R \cdot C &= \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 3 \\ 0 & 0 & 4 \end{bmatrix} - \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix} \cdot [0 \quad 0 \quad 3] \\ &= \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 3 \\ 0 & 0 & 4 \end{bmatrix} - \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 3 \\ 0 & 0 & 0 \end{bmatrix} \\ &= \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 4 \end{bmatrix}. \end{aligned}$$

We've seen this last step before and note that we have

$$C = \begin{bmatrix} 0 \\ 0 \\ 4 \end{bmatrix} \cdot \frac{1}{4} = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}, \quad R = [0 \ 0 \ 4].$$

Assembling all the pieces, we have

$$\underbrace{\begin{bmatrix} 2 & -1 & 2 \\ 6 & -3 & 9 \\ 2 & -1 & 6 \end{bmatrix}}_M = \underbrace{\begin{bmatrix} 1 & 0 & 0 \\ 3 & 1 & 0 \\ 1 & 0 & 1 \end{bmatrix}}_L \cdot \underbrace{\begin{bmatrix} 2 & 1 & -2 \\ 0 & 0 & 3 \\ 0 & 0 & 4 \end{bmatrix}}_U,$$

and we note that U has a zero on its diagonal².

In general, at the k -th step, if C becomes a column of all zeros, set its k -th entry to one and define R as usual.

Case 2 (requires row permutation): C is not all zeros, but the pivot value is zero. In this case, we need to swap rows in the matrix $Temp$ so that the new C has a non-zero pivot value. This introduces a permutation matrix, which you can review in Chapter 4.7. Keeping track of the row permutations can be a bit tricky, which is why we left this to the end.

Suppose we have $M = \begin{bmatrix} 2 & -1 & 2 \\ 6 & -3 & 9 \\ 2 & -3 & 6 \end{bmatrix}$. We initialize the process by $L = []$, $U = []$, and $Temp = M$.

Step $k = 1$ We have

$$C = \begin{bmatrix} 2 \\ 6 \\ 2 \end{bmatrix} \cdot \frac{1}{2} = \begin{bmatrix} 1 \\ 3 \\ 1 \end{bmatrix}, \quad R = [2 \ -1 \ 2],$$

and

$$\begin{aligned} Temp &= Temp - CR \\ &= \begin{bmatrix} 2 & -1 & 2 \\ 6 & -3 & 9 \\ 2 & -3 & 6 \end{bmatrix} - \begin{bmatrix} 2 & -1 & 2 \\ 6 & -3 & 6 \\ 2 & -1 & 2 \end{bmatrix} \\ &= \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 3 \\ 0 & -2 & 4 \end{bmatrix}. \end{aligned}$$

We then update L and U by

$$\begin{aligned} L &= [L, C] = \begin{bmatrix} 1 \\ 3 \\ 1 \end{bmatrix} \\ U &= [U; R] = [2 \ -1 \ 2] \end{aligned}$$

Step $k = 2$ The problem appears when we attempt to define $\mathbf{pivot} = Temp[2, 2] = 0$

$$C = \begin{bmatrix} 0 \\ 0 \\ -2 \end{bmatrix} \cdot \frac{1}{\mathbf{pivot}} = \begin{bmatrix} 0 \\ 0 \\ -2 \end{bmatrix} \cdot \frac{1}{\mathbf{0}},$$

²Hence, $\det(U) = 0$, and if we are trying to solve $Ux = y$, we are then in the problematic case of having either no solution or an infinite number of solutions.

which sets off the divide by zero warning because we have $\text{pivot} = 0$. The solution is to permute the second and third rows of M so that at step $k = 2$, Temp will become

$$\text{Temp}_{\text{new}} := \begin{bmatrix} 0 & 0 & 0 \\ 0 & -2 & 4 \\ 0 & 0 & 3 \end{bmatrix}.$$

We could do this by going all the way back to the very beginning and starting the factorization over after doing the row permutation on M , but that is not very efficient. **It turns out that all we really have to do is permute the two rows in the current value of Temp AND the corresponding rows in L .** Nothing has to be done to U .

We'll first refresh our memory on permutation matrices. Then we'll write down the full LU factorization algorithm with row permutations and do an example that has both of these new steps in them.

Row Permutation Matrices Consider

$$M_a = \begin{bmatrix} 2 & -1 & 2 \\ 6 & -3 & 9 \\ 2 & -3 & 6 \end{bmatrix}.$$

And suppose we want to swap rows one and three to obtain

$$M_b = \begin{bmatrix} 2 & -3 & 6 \\ 6 & -3 & 9 \\ 2 & -1 & 2 \end{bmatrix}$$

Then this is the result of pre-multiplying M_a by the permutation matrix

$$I := \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \rightarrow \text{swap first and third rows} \rightarrow P_a := \begin{bmatrix} 0 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \end{bmatrix},$$

which we can verify by

$$P_a \cdot M_a = \begin{bmatrix} 0 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \end{bmatrix} \cdot \begin{bmatrix} 2 & -1 & 2 \\ 6 & -3 & 9 \\ 2 & -3 & 6 \end{bmatrix} = \begin{bmatrix} 2 & -3 & 6 \\ 6 & -3 & 9 \\ 2 & -1 & 2 \end{bmatrix} = M_b$$

In fact, with our "second way" of doing matrix multiplication as the sum over columns times rows, the above product can be written out as

$$\begin{aligned} P_a \cdot M_a &= \begin{bmatrix} 0 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \end{bmatrix} \cdot \begin{bmatrix} 2 & -1 & 2 \\ 6 & -3 & 9 \\ 2 & -3 & 6 \end{bmatrix} \\ &= \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \cdot [2 \quad -1 \quad 2] + \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix} \cdot [6 \quad -3 \quad 9] + \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} \cdot [2 \quad -3 \quad 6] \\ &= \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 2 & -1 & 2 \end{bmatrix} + \begin{bmatrix} 0 & 0 & 0 \\ 6 & -3 & 9 \\ 0 & 0 & 0 \end{bmatrix} + \begin{bmatrix} 2 & -3 & 6 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \\ &= \begin{bmatrix} 2 & -3 & 6 \\ 6 & -3 & 9 \\ 2 & -1 & 2 \end{bmatrix} \\ &= M_b \end{aligned}$$

It is fine to stop here, but if you want more practice with permutation matrices, then continue reading.

Example 5.6 Let's now swap rows one and two of M_b to obtain

$$M_c = \begin{bmatrix} 6 & -3 & 9 \\ 2 & -3 & 6 \\ 2 & -1 & 2 \end{bmatrix}.$$

Find the permutation matrix P_b such that $M_c = P_b \cdot M_b$ and the permutation matrix P_c such that $M_c = P_c \cdot M_a$. Verify that $P_c = P_b \cdot P_a$. This will help you to understand matrix multiplication.

Solution: To find the permutation matrix P_b

$$I := \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \rightarrow \text{swap first and second rows} \rightarrow P_b = \begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix},$$

which we can verify by

$$P_b \cdot M_b = \begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix} \cdot \begin{bmatrix} 2 & -3 & 6 \\ 6 & -3 & 9 \\ 2 & -1 & 2 \end{bmatrix} = \begin{bmatrix} 6 & -3 & 9 \\ 2 & -3 & 6 \\ 2 & -1 & 2 \end{bmatrix} = M_c$$

To find the permutation matrix P_c It is observed that M_c can be obtained by swapping the first with the last two rows of M_a .

$$I := \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \rightarrow \text{swap first and last two rows} \rightarrow P_c = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{bmatrix},$$

which we can verify by

$$P_c \cdot M_a = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{bmatrix} \cdot \begin{bmatrix} 2 & -1 & 2 \\ 6 & -3 & 9 \\ 2 & -3 & 6 \end{bmatrix} = \begin{bmatrix} 6 & -3 & 9 \\ 2 & -3 & 6 \\ 2 & -1 & 2 \end{bmatrix} = M_c$$

We verify $P_c = P_b \cdot P_a$ by

$$P_b \cdot P_a = \begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix} \cdot \begin{bmatrix} 0 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \end{bmatrix} = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{bmatrix} = P_c$$

■

How General Can LU Go?

All of our examples have been done for M a square matrix. LU even works for M rectangular; see Chapter 5.12.

5.8 (Optional Read): An Algorithm for LU Factorization with Row Permutations

LU Factorization of Square Matrices

Algorithm 2: LU-Factorization (square matrix with permutations)

Result: For M an $n \times n$ square matrix, with $n \geq 2$, find L , U , and P such that $PM = LU$

initialization:

Temp=copy(M);

L = Array{Float64,2}(undef, n, 0) # L=[] Empty matrix

U = Array{Float64,2}(undef, 0, n) # U=[] Empty matrix

n,m = size(M)

P= zeros(n,n)+I # I_n $n \times n$ identity matrix

eps= 1e-16 # estimate of machine epsilon

Kappa = 10 #How small is too small for you?

end initialization:

for k=1:n

C=Temp[:,k]; # k-th column

if max(abs(C)) <= Kappa*eps # (check for C all zeros)

C=0*C # Set tiny values to zero

C[k]=1.0

R=Temp[k:k,:]; # k-th row

Temp=Temp-C*R

L=[L,C] # Build the lower-triangular matrix by columns

U=[U;R] # Build the upper-triangular matrix by rows

else

We know C has at least one non-zero quantity

We'll bring its largest entry to the top;

while this is overkill, it helps with numerical aspects of the algorithm

Bringing biggest to top will always avoid divide by zero

It's enough to bring ANY non-zero value to the top

#

nrow=argmax(abs(C)) # find the index where C is "biggest"

P[[k,nrow],:]=P[[nrow,k],:] # permute (i.e., swap) rows of P

Temp[[k,nrow],:]=Temp[[nrow,k],:] # do same for Temp

if L is non-empty, also swap its rows

if k > 1 L[[k,nrow],:]=L[[nrow,k],:] **end**

C=Temp[:,k]; # k-th column

pivot = C[k]

C=C/pivot #divide all elements of C by the pivot

R=Temp[k:k,:]; # k-th row

Temp=Temp-C*R

L=[L C] # Build the lower-triangular matrix by columns

U=[U;R] # Build the upper-triangular matrix by rows

end

end

return L, U, P

Example 5.7 We do it by hand so that you can see all the steps. $M = \begin{bmatrix} 6 & -2 & 4 & 4 \\ 3 & -1 & 8 & 10 \\ 3 & -1 & 2 & 8 \\ 3 & -1 & -1 & 20 \end{bmatrix}$, a 4×4 square matrix.

Solution: We initialize our algorithm by

$$\text{Temp} := M = \begin{bmatrix} 6 & -2 & 4 & 4 \\ 3 & -1 & 8 & 10 \\ 3 & -1 & 2 & 8 \\ 3 & -1 & -1 & 20 \end{bmatrix}, P := I_4 = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, L := [\text{empty matrix}], \text{ and } U := [\text{empty matrix}]$$

k=1: We start with C as first column of Temp normalized by $C[1]$ so that its first entry is one

$$C = \begin{bmatrix} 6 \\ 3 \\ 3 \\ 3 \end{bmatrix}$$

$$\text{nrow} = \text{argmax}(\text{abs}(C)) = 1$$

$$P = P$$

$$\text{Temp} = \text{Temp}$$

$$C = \text{Temp}[:, 1] = C$$

$$\text{pivot} = C[1] = 6$$

$$C = C./\text{pivot}$$

$$= \begin{bmatrix} 1 \\ 0.5 \\ 0.5 \\ 0.5 \end{bmatrix},$$

while R is the first row of Temp

$$R = [6 \quad -2 \quad 4 \quad 4].$$

We finish up the first step by defining

$$\begin{aligned} \text{Temp} &:= \text{Temp} - C \cdot R \\ &= \begin{bmatrix} 6 & -2 & 4 & 4 \\ 3 & -1 & 8 & 10 \\ 3 & -1 & 2 & 8 \\ 3 & -1 & -1 & 20 \end{bmatrix} - \begin{bmatrix} 1 \\ 0.5 \\ 0.5 \\ 0.5 \end{bmatrix} \cdot [6 \quad -2 \quad 4 \quad 4] \\ &= \begin{bmatrix} 6 & -2 & 4 & 4 \\ 3 & -1 & 8 & 10 \\ 3 & -1 & 2 & 8 \\ 3 & -1 & -1 & 20 \end{bmatrix} - \begin{bmatrix} 6 & -2 & 4 & 4 \\ 3 & -1 & 2 & 2 \\ 3 & -1 & 2 & 2 \\ 3 & -1 & 2 & 2 \end{bmatrix} \\ &= \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 6 & 8 \\ 0 & 0 & 0 & 6 \\ 0 & 0 & -3 & 18 \end{bmatrix} \end{aligned}$$

$$L = [L, C] = \begin{bmatrix} & & & & 1 \\ & & & & 0.5 \\ & & & & 0.5 \\ & & & & 0.5 \end{bmatrix}$$

$$U = [U; R] = [6 \quad -2 \quad 4 \quad 4]$$

$$P = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

We observe that L and U have been constructed so that their product exactly matches the first two columns of M . Hence, when we subtract $L \cdot U$ from M , we are left with a problem that the second column of Temp becomes an entire column of zeros, meaning, in this case, the non-zero part is 3×2 .

k=2: C is the second column of Temp and it is all zeros,

$$C[2] = 1$$

$$C = \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \end{bmatrix},$$

while R is the second row of Temp ,

$$R = [0 \ 0 \ 6 \ 8].$$

$$\text{Temp} := \text{Temp} - C \cdot R$$

$$= \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 6 & 8 \\ 0 & 0 & 0 & 6 \\ 0 & 0 & -3 & 18 \end{bmatrix} - \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \end{bmatrix} \cdot [0 \ 0 \ 6 \ 8]$$

$$= \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 6 & 8 \\ 0 & 0 & 0 & 6 \\ 0 & 0 & -3 & 18 \end{bmatrix} - \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 6 & 8 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

$$= \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 6 \\ 0 & 0 & -3 & 18 \end{bmatrix}$$

$$L = [L, C] = \begin{bmatrix} 1 & 0 \\ 0.5 & 1 \\ 0.5 & 0 \\ 0.5 & 0 \end{bmatrix}$$

$$U = [U; R] = \begin{bmatrix} 6 & -2 & 4 & 4 \\ 0 & 0 & 6 & 8 \end{bmatrix}$$

$$P = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

When we subtract $L \cdot U$ from temp , we are left with a problem that is effectively two dimensions lower, meaning, in this case, the non-zero part is 2×2 and we are making good progress on the LU factorization! Moreover, we observe that the third column of Temp is not all zeros but the pivot value is zero. Hence, row permutation is required.

k=3: C is the third column of Temp and has 0 pivot,

$$C = \begin{bmatrix} 0 \\ 0 \\ 0 \\ -3 \end{bmatrix}$$

$$\text{nrow} = \text{argmax}(\text{abs}(C)) = 4$$

$$P[[3, 4], :] = P[[4, 3], :] = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix}$$

$$\text{Temp}[[3, 4], :] = \text{Temp}[[4, 3], :] = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & -3 & 18 \\ 0 & 0 & 0 & 6 \end{bmatrix}$$

$$L[[3, 4], :] = L[[4, 3], :] = \begin{bmatrix} 1 & 0 \\ 0.5 & 1 \\ 0.5 & 0 \\ 0.5 & 0 \end{bmatrix}$$

$$C = \text{Temp}[:, 3] = \begin{bmatrix} 0 \\ 0 \\ -3 \\ 0 \end{bmatrix}$$

$$\text{pivot} = C[3] = -3$$

$$C = C./\text{pivot}$$

$$= \begin{bmatrix} 0 \\ 0 \\ 1 \\ 0 \end{bmatrix},$$

while R is the third row of Temp ,

$$R = [0 \ 0 \ -3 \ 18].$$

$$\text{Temp} := \text{Temp} - C \cdot R$$

$$= \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & -3 & 18 \\ 0 & 0 & 0 & 6 \end{bmatrix} - \begin{bmatrix} 0 \\ 0 \\ 1 \\ 0 \end{bmatrix} \cdot [0 \ 0 \ -3 \ 18]$$

$$= \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & -3 & 18 \\ 0 & 0 & 0 & 6 \end{bmatrix} - \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & -3 & 18 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

$$= \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 6 \end{bmatrix}$$

$$L = [L, C] = \begin{bmatrix} 1 & 0 & 0 \\ 0.5 & 1 & 0 \\ 0.5 & 0 & 1 \\ 0.5 & 0 & 0 \end{bmatrix}$$

$$U = [U; R] = \begin{bmatrix} 6 & -2 & 4 & 4 \\ 0 & 0 & 6 & 8 \\ 0 & 0 & -3 & 18 \end{bmatrix}$$

$$P = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix}$$

k=4: C is the fourth column of Temp scaled by its $k = 4$ entry,

$$C = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 6 \end{bmatrix}$$

$$\text{nrow} = \text{argmax}(\text{abs}(C)) = 4$$

$$P = P$$

$$\text{Temp} = \text{Temp}$$

$$L = L$$

$$C = \text{Temp}[:, 4] = C$$

$$\text{pivot} = C[4] = 6$$

$$C = C./\text{pivot}$$

$$= \begin{bmatrix} 0 \\ 0 \\ 0 \\ 1 \end{bmatrix},$$

while R is the fourth row of Temp,

$$R = [0 \ 0 \ 0 \ 6].$$

$$\text{Temp} := \text{Temp} - C \cdot R$$

$$= \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 6 \end{bmatrix} - \begin{bmatrix} 0 \\ 0 \\ 0 \\ 1 \end{bmatrix} \cdot [0 \ 0 \ 0 \ 6]$$

$$= \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 6 \end{bmatrix} - \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 6 \end{bmatrix}$$

$$= \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

$$L = [L; C] = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0.5 & 1 & 0 & 0 \\ 0.5 & 0 & 1 & 0 \\ 0.5 & 0 & 0 & 1 \end{bmatrix} = \text{lower triangular matrix}$$

$$U = [U; R] = \begin{bmatrix} 6 & -2 & 4 & 4 \\ 0 & 0 & 6 & 8 \\ 0 & 0 & -3 & 18 \\ 0 & 0 & 0 & 6 \end{bmatrix} = \text{upper triangular matrix}$$

$$P = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix}$$

Let's remind ourselves why this works: As before,

$$\text{Temp} := P \cdot M - L \cdot U$$

$$= \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix} \cdot \begin{bmatrix} 6 & -2 & 4 & 4 \\ 3 & -1 & 8 & 10 \\ 3 & -1 & 2 & 8 \\ 3 & -1 & -1 & 20 \end{bmatrix} - \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0.5 & 1 & 0 & 0 \\ 0.5 & 0 & 1 & 0 \\ 0.5 & 0 & 0 & 1 \end{bmatrix} \cdot \begin{bmatrix} 6 & -2 & 4 & 4 \\ 0 & 0 & 6 & 8 \\ 0 & 0 & -3 & 18 \\ 0 & 0 & 0 & 6 \end{bmatrix}$$

$$= \begin{bmatrix} 6 & -2 & 4 & 4 \\ 3 & -1 & 8 & 10 \\ 3 & -1 & -1 & 20 \\ 3 & -1 & 2 & 8 \end{bmatrix} - \left(\begin{bmatrix} 1 \\ 0.5 \\ 0.5 \\ 0.5 \end{bmatrix} \cdot [6 \ -2 \ 4 \ 4] + \dots + \begin{bmatrix} 0 \\ 0 \\ 0 \\ 1 \end{bmatrix} \cdot [0 \ 0 \ 0 \ 6] \right)$$

$$= \begin{bmatrix} 6 & -2 & 4 & 4 \\ 3 & -1 & 8 & 10 \\ 3 & -1 & -1 & 20 \\ 3 & -1 & 2 & 8 \end{bmatrix} - \left(\begin{bmatrix} 6 & -2 & 4 & 4 \\ 3 & -1 & 2 & 2 \\ 3 & -1 & 2 & 2 \\ 3 & -1 & 2 & 2 \end{bmatrix} + \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 6 & 8 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} + \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & -3 & 18 \\ 0 & 0 & 0 & 0 \end{bmatrix} + \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 6 \end{bmatrix} \right)$$

$$= \left(\begin{bmatrix} 6 & -2 & 4 & 4 \\ 3 & -1 & 8 & 10 \\ 3 & -1 & -1 & 20 \\ 3 & -1 & 2 & 8 \end{bmatrix} - \begin{bmatrix} 6 & -2 & 4 & 4 \\ 3 & -1 & 2 & 2 \\ 3 & -1 & 2 & 2 \\ 3 & -1 & 2 & 2 \end{bmatrix} \right) - \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 6 & 8 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} - \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & -3 & 18 \\ 0 & 0 & 0 & 0 \end{bmatrix} - \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 6 \end{bmatrix}$$

$$= \left(\begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 6 & 8 \\ 0 & 0 & -3 & 18 \\ 0 & 0 & 0 & 6 \end{bmatrix} - \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 6 & 8 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \right) - \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & -3 & 18 \\ 0 & 0 & 0 & 0 \end{bmatrix} - \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 6 \end{bmatrix}$$

$$= \left(\begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & -3 & 18 \\ 0 & 0 & 0 & 6 \end{bmatrix} - \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & -3 & 18 \\ 0 & 0 & 0 & 0 \end{bmatrix} \right) - \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 6 \end{bmatrix}$$

$$= \left(\begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 6 \end{bmatrix} - \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 6 \end{bmatrix} \right)$$

$$= \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

Hence,

$$\underbrace{\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix}}_P \cdot \underbrace{\begin{bmatrix} 6 & -2 & 4 & 4 \\ 3 & -1 & 8 & 10 \\ 3 & -1 & 2 & 8 \\ 3 & -1 & -1 & 20 \end{bmatrix}}_M = \underbrace{\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0.5 & 1 & 0 & 0 \\ 0.5 & 0 & 1 & 0 \\ 0.5 & 0 & 0 & 1 \end{bmatrix}}_L \cdot \underbrace{\begin{bmatrix} 6 & -2 & 4 & 4 \\ 0 & 0 & 6 & 8 \\ 0 & 0 & -3 & 18 \\ 0 & 0 & 0 & 6 \end{bmatrix}}_U. \quad (5.15)$$

■

Disclaimer

This is your book's author speaking: We will never do this by hand again. Don't even ask. The good thing for you, I will not ask you do to something that I would not do myself! (Thank you, Miley.)

5.9 Solving $Ax=b$ via LU with Row Permutations

Solving $Ax = b$ when A is square and $P \cdot A = L \cdot U$

So, how does the permutation matrix change how one solves $Ax = b$? Because **permutation matrices are invertible**, the main logic is

$$Ax = b \iff P \cdot Ax = P \cdot b \iff L \cdot Ux = P \cdot b, \quad (5.16)$$

where we have used $P \cdot A = L \cdot U$. Hence, (5.6) and (5.7) are modified as follows: if we define $Ux = y$ as before, then (5.16) becomes two equations

$$Ly = P \cdot b \quad (5.17)$$

$$Ux = y. \quad (5.18)$$

The solution strategy is to solve (5.17) by forward substitution, and then, once we have y in hand, we solve (5.18) by back substitution to find x , the solution to (5.16).

5.10 (Optional Read): Using the LU command native to Julia

Julia has a built-in command for LU Factorizations. It has a few quirks that are highlighted here.

- **Quirk 1:** The base command assumes row permutations. You can force it to return an LU factorization without row permutations

```
1 using LinearAlgebra
2 using Random
3 Random.seed!(12345678)
4
5 A=randn(4,4)
6 b=rand(4,1)
```

Output because b is the last thing in the cell, only it is printed

```
4×1 Matrix{Float64}:
 0.9298795160436026
 0.9240212608175158
 0.5693478013892368
 0.6619748541328005
```

```
1 # The option Val(false) turns off row permutations
2 F=lu(A, Val(false))
```

Output

```
LU{Float64, Matrix{Float64}}
L factor:
4×4 Matrix{Float64}:
 1.0      0.0      0.0      0.0
-0.77706  1.0      0.0      0.0
-0.943967 0.53299  1.0      0.0
-0.223178 0.12119  0.548212 1.0
U factor:
4×4 Matrix{Float64}:
-0.889751 -1.11357 -0.256081 -0.272776
 0.0      -1.99388  0.63173  -0.3698
 0.0      0.0      1.49897  -0.882234
 0.0      0.0      0.0      -0.277314
```

We now have L and U and thus we can efficiently solve for x

```
1 y=forwardsub(L, b)
2 x=backwardsub(U, y)
```

Output

```
4-element Vector{Float64}:
 1.1147933490640047
-0.6334718017433187
-1.1343956246238571
-2.4120580258922018
```

- **Quirk 2:** You can get L , U , and “the permutation” in **two different ways that do not return the same information!**

```
1 using LinearAlgebra
2 using Random
3
4 Random.seed!(12345678)
5 A=randn(4, 4)
6 b=rand(4, 1)
```

Output

```
4×1 Matrix{Float64}:
 0.9298795160436026
 0.9240212608175158
 0.5693478013892368
 0.6619748541328005
```

Method one for doing an LU Factorization in Julia: The thing to note here is that we deliberately used a lowercase p because it is not the permutation matrix itself, but the vector of indices defining the permutation. What? Read on.

```
1 L, U, p = lu(A)
2 # Note that we are using a lowercase p anticipating it will NOT be a matrix
```

Output

```
LU{Float64, Matrix{Float64}}
L factor:
4×4 Matrix{Float64}:
```

```

1.0      0.0      0.0      0.0
-0.77706 1.0      0.0      0.0
-0.943967 0.53299 1.0      0.0
-0.223178 0.12119 0.548212 1.0
U factor:
4×4 Matrix{Float64}:
-0.889751 -1.11357 -0.256081 -0.272776
 0.0      -1.99388  0.63173  -0.3698
 0.0      0.0      1.49897  -0.882234
 0.0      0.0      0.0      -0.277314

```

```
1 p
```

Output

```

4-element Vector{Int64}:
 4
 2
 3
 1

```

We see that lowercase p is indeed not a matrix. It is column vector of indices that will bring the fourth row to the top, then keep the second row where it is, also keep the third row where it is, and finally, the first row is moved to the bottom. Moreover, look at how easily we can use the indices to permute the rows of b by forming $b[p]$,

```
1 [b b[p]]
```

Output shows that the rows of b have been permuted, which is what we need

```

4×2 Matrix{Float64}:
 0.92988  0.661975
 0.924021 0.924021
 0.569348 0.569348
 0.661975 0.92988

```

This is very important so we say it again. The above shows that the rows of b have indeed been permuted.

```

1 y=forwardsub(L, b[p])
2 x=backwardsub(U, y)

```

Output

```

4-element Vector{Float64}:
 1.1147933490640065
-0.633471801743319
-1.1343956246238587
-2.4120580258922035

```

Method two for doing an LU Factorization in Julia: The thing to note is that when we call the function using the structure notation, then we can obtain the permutation matrix instead of only the list of indices.

```

1 # Structure method of calling a native Julia function
2 F=lu(A)
3 L=F.L
4 U=F.U
5 P=F.P

```

F is called a structure, or `struct` for short. It contains the elements **L**, **U**, and **P** as fields. You extract **L**, **U**, and **P** as shown above with the “dot” notation, such as `L=F.L`

Output Because P is the last thing in the code cell, the output is P , which we see is a matrix, namely

```
4×4 Matrix{Float64}:
 0.0  0.0  0.0  1.0
 0.0  1.0  0.0  0.0
 0.0  0.0  1.0  0.0
 1.0  0.0  0.0  0.0
```

We see that P will bring the fourth row to the top, then keep the second row where it is, also keep the third row where it is, and finally, the first row is moved to the bottom.

```
1 [b P*b]
```

Output

```
4×2 Matrix{Float64}:
 0.92988  0.661975
 0.924021 0.924021
 0.569348 0.569348
 0.661975 0.92988
```

We see that when P is a matrix, we permute the rows of b by multiplying it by P , namely Pb , and when p is a column vector of indices, we permute the rows by $b[p]$. **Either way, we get the same answer for $Ax = b$ because $b[p] = P * b$**

```
1 [b P*b b[p]]
```

Output two equivalent ways to permute the elements of b

```
4×3 Matrix{Float64}:
 0.92988  0.661975  0.661975
 0.924021 0.924021  0.924021
 0.569348 0.569348  0.569348
 0.661975 0.92988  0.92988
```

```
1 y1=forwardsub(L, P*b)
2 x1=backwardsub(U, y1)
3 y2=forwardsub(L, b[p])
4 x2=backwardsub(U, y2)
5 [x1 x2]
```

Output we obtain the same solution x in each case

```
4×2 Matrix{Float64}:
 1.11479  1.11479
-0.633472 -0.633472
-1.1344   -1.1344
-2.41206  -2.41206
```

As you gain more experience, you may find that you like the structure method better. As a final note, you can obtain the list of indices used in the permutation from the structure via the field **F.p**

```

1 Random.seed!(12345678)
2 A=randn(4,4)
3 b=rand(4,1)
4 F=lu(A)
5 y=forwardsub(F.L, b[F.p])
6 x=backwardsub(F.U, y)

```

Output we obtain the same x as before

```

4-element Vector{Float64}:
 1.1147933490640065
-0.633471801743319
-1.1343956246238587
-2.4120580258922035

```

Once again, we note that “**F.p** ← lowercase p” provides the list of permutation indices, while “**F.P** ← uppercase P” provides the full permutation matrix. When using **F.P**, the correct code is

```

1 F=lu(A)
2 y=forwardsub(F.L, F.P*b)
3 x=backwardsub(F.U, y)

```

• Getting help

```

1 using LinearAlgebra
2
3 ? lu

```

Output It’s a lot to parse and understand. In the grey box, we highlight some very important information on the struct **F**.

```
search: lu lu! Luv LU LuvA blue flush ALuv values include include_string
```

```
lu(A, pivot=Val(true); check = true) -> F::LU
Compute the LU factorization of A.
```

When `check = true`, an error is thrown if the decomposition fails. When `check = false`, respo

In most cases, if `A` is a subtype `S` of `AbstractMatrix{T}` with an element type `T` supporting +

The individual components of the factorization `F` can be accessed via `getproperty`:

```

Component Description
F.L L (lower triangular) part of LU
F.U U (upper triangular) part of LU
F.p (right) permutation Vector
F.P (right) permutation Matrix
Iterating the factorization produces the components F.L, F.U, and F.p.

```

The relationship between `F` and `A` is

```
F.L*F.U == A[F.p, :]
```

Left unsaid is:

$$F.L * F.U == F.P * A$$

The following is where the help information gives you the elements of the structure, F:

Component Description
F.L L (lower triangular) part of LU
F.U U (upper triangular) part of LU
F.p (right) permutation Vector
F.P (right) permutation Matrix

We note the distinction between **Vector** and **Matrix**.

Help! Help! How am I supposed to remember all of this?

You probably can't. In any case, we don't want you to memorize the ROB 101 material. Instead, open up a google doc or google sheet and make notes! You need an organized method for keeping track of stuff. In High School, you may have been able to remember all the new notation without any special effort. In College, it's a bit different.

5.11 (Optional Read): Large Scale Example: Computing the Forces in a Truss Bridge



(a)

(b)

Figure 5.1: Trusses are ubiquitous in construction. Here we show two uses in bridges and roofs. The techniques used to analyze the forces in each member of a truss are taught in ME 211. They result in a large set of linear equations. We'll use a bridge to exemplify the main ideas. You can learn more at https://www.youtube.com/watch?v=Hn_iozUo9m4.

Figure 5.1 shows two truss structures, one for a railroad bridge and the other for a roof in a house or small office building. We'll analyze the simpler planar truss bridge in Fig. 5.2, which has joints labeled 1 to 13. In case you are interested, the basic method used

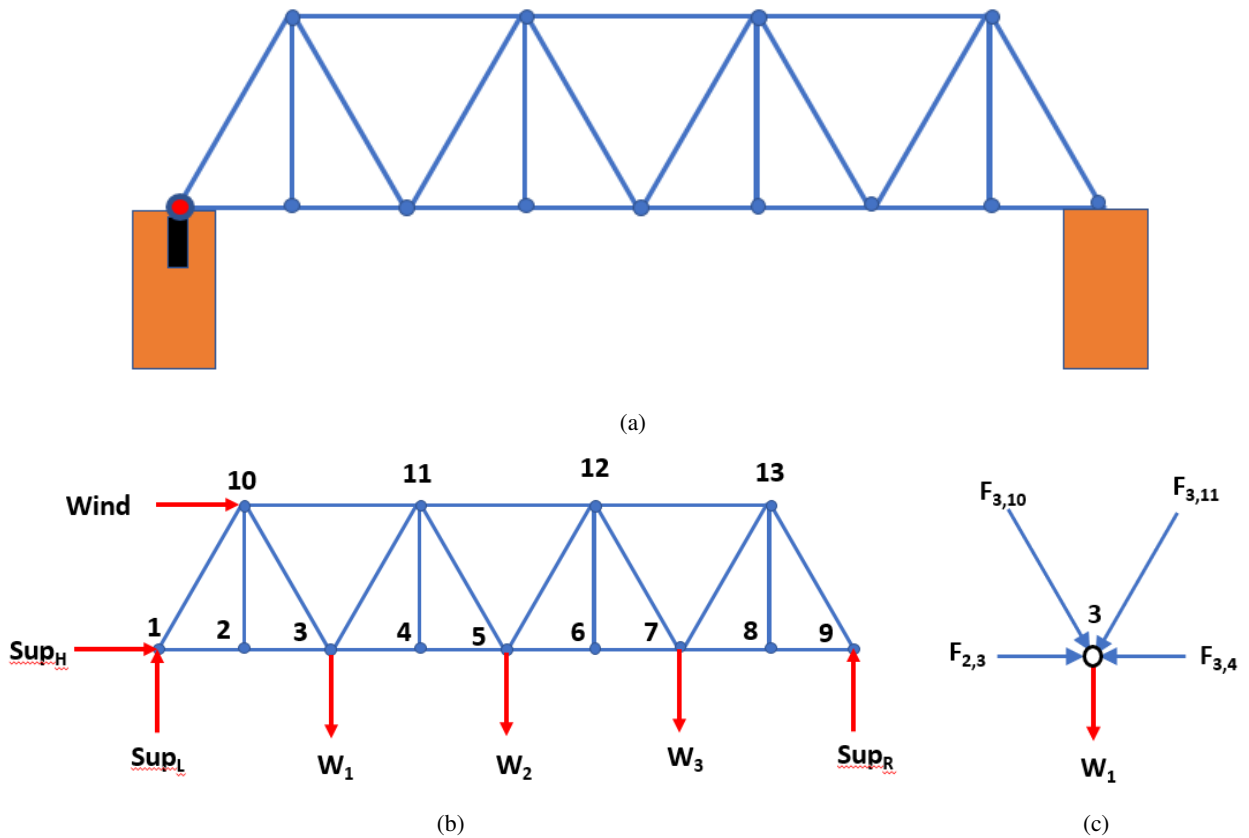


Figure 5.2: A planar truss bridge over an open span, say a road or body of water. As shown in (a), the bridge is pinned on the left and “floating” on the right. In (b), the joints of the truss structure have been labeled along with some of the forces, such as the horizontal and vertical forces at the left support point, the vertical force at the right support point, wind loading on the structure, and a simplified representation of the structure’s weight pulling down on it. A free body diagram for joint 3 is shown in (c), where it is assumed *arbitrarily* that each member is in compression. If a member is in *tension* instead of compression, the force will be a negative number. To build the linear model for the truss structure, force balances must be computed at each joint in the horizontal and vertical directions. Because there are 13 joints, the model will have 26 variables in it. The real-life structures in Fig. 5.1 would have several hundreds of variables.

to develop a mathematical model of the structure is taught in ME 211 Introduction to Dynamics and Vibrations. At each joint, a free body diagram is constructed as shown in Fig. 5.2-(c), where it has been arbitrarily assumed that all of the members in the truss are in *compression*, and thus the forces transmitted by them are directed *into the joint*. With this convention, when we eventually solve for the indicated forces, any member that is in *tension* will have a negative force associated to it.

The angled members in the diagram form an angle of $\frac{\pi}{3}$, or 60° , with respect to the horizontal axis. The force balance at joint 3 yields two equations,

$$\begin{aligned} F_{2,3} - F_{3,4} + \cos(\pi/3)F_{3,10} - \cos(\pi/3)F_{3,11} &= 0.0 & \text{horizontal components} \\ -F_{w1} - \sin(\pi/3)F_{3,10} - \sin(\pi/3)F_{3,11} &= 0.0 & \text{vertical components,} \end{aligned} \tag{5.19}$$

the first for the horizontal components of the forces and the second for the vertical components of the forces. A similar process at joint 10 yields

$$\begin{aligned} F_{\text{wind}} - F_{10,11} + \cos(\pi/3)F_{1,10} - \cos(\pi/3)F_{3,10} &= 0.0 & \text{horizontal components} \\ F_{2,10} + \sin(\pi/3)F_{1,10} + \sin(\pi/3)F_{3,10} &= 0.0 & \text{vertical components.} \end{aligned} \tag{5.20}$$

Repeating this for each of the 13 joints yields a total of 26 equations for the 26 unknown forces. **The resulting linear system of equations $Ax = b$ has an A matrix that is 26×26 , which means it has 676 entries. As you can imagine, building the matrix by hand would be very tedious! What to do about it? Write some code! The Julia code below is a bit advanced for where we are right now in ROB 101. It is included for the sole purpose of illustrating that programming can be used to build models as**

well as to solve equations.

The Julia code contains each of the 26 force balance equations, then from the force balance equations it creates the model $Ax = b$, and finally it solves for x by LU factorization with row permutations. You are not expected to work through and understand the code. Copying it into a jupyter notebook and playing with it may be fun. You can check that $\det(A) \approx 1.46$.

```
1 using LinearAlgebra
2
3 function Model(x, u)
4 # Units are kilo Newtons, or the force required to support 220 pounds
5 # Bridge model variables
6 # External forces
7 Fw1=u[1]; Fw2=u[2]; Fw3=u[3]; Fwind=u[4]
8 # Internal Forces, where "c" stands for "comma"
9 F1c2=x[1]
10 F1c10=x[2]
11 F2c3=x[3]
12 F2c10=x[4]
13 F3c4=x[5]
14 F3c10=x[6]
15 F3c11=x[7]
16 F4c5=x[8]
17 F4c11=x[9]
18 F5c6=x[10]
19 F5c11=x[11]
20 F5c12=x[12]
21 F6c7=x[13]
22 F6c12=x[14]
23 F7c8=x[15]
24 F7c12=x[16]
25 F7c13=x[17]
26 F8c9=x[18]
27 F8c13=x[19]
28 F9c13=x[20]
29 F10c11=x[21]
30 F11c12=x[22]
31 F12c13=x[23]
32 FhorizLeft=x[24]
33 FvertLeft=x[25]
34 FvertRight=x[26]
35 #Bridge force balance equations
36 # Note cos(pi/3)=0.5 and sin(pi/3)=sqrt(3)/2
37 #joint 1, x component first then y
38 Eq1x=FhorizLeft-F1c2-F1c10*cos(pi/3)
39 Eq1y=FvertLeft-F1c10*sin(pi/3)
40 #joint 2, x component first then y
41 Eq2x=F1c2-F2c3
42 Eq2y=F2c10
43 #joint 3, x component first then y
44 Eq3x=F2c3-F3c4+cos(pi/3)*F3c10-cos(pi/3)*F3c11
45 Eq3y=-Fw1-sin(pi/3)*F3c10-sin(pi/3)*F3c11
46 #joint 4, x component first then y
47 Eq4x=F3c4-F4c5
48 Eq4y=-F4c11
49 #joint 5, x component first then y
50 Eq5x=F4c5-F5c6+cos(pi/3)*F5c11-cos(pi/3)*F5c12
51 Eq5y=-Fw2-sin(pi/3)*F5c11-sin(pi/3)*F5c12
```

```

52 #joint 6, x component first then y
53 Eq6x=F5c6-F6c7
54 Eq6y=-F6c12
55 #joint 7, x component first then y
56 Eq7x=F6c7-F7c8+cos(pi/3)*F7c12-cos(pi/3)*F7c13
57 Eq7y=-Fw3-sin(pi/3)*F7c12-sin(pi/3)*F7c13
58 #joint 8, x component first then y
59 Eq8x=F7c8-F8c9
60 Eq8y=-F8c13
61 #joint 9, x component first then y
62 Eq9x=F8c9+F9c13*cos(pi/3)
63 Eq9y=FvertRight-F9c13*sin(pi/3)
64 #joint 10, x component first then y
65 Eq10x=Fwind-F10c11+cos(pi/3)*F1c10-cos(pi/3)*F3c10
66 Eq10y=F2c10+sin(pi/3)*F1c10+sin(pi/3)*F3c10
67 #joint 11, x component first then y
68 Eq11x=F10c11-F11c12+cos(pi/3)*F3c11-cos(pi/3)*F5c11
69 Eq11y=F4c11+sin(pi/3)*F3c11+sin(pi/3)*F5c11
70 #joint 12, x component first then y
71 Eq12x=F11c12-F12c13+cos(pi/3)*F5c12-cos(pi/3)*F7c12
72 Eq12y=F6c12+sin(pi/3)*F5c12+sin(pi/3)*F7c12
73 #joint 13, x component first then y
74 Eq13x=F12c13+cos(pi/3)*F7c13-cos(pi/3)*F9c13
75 Eq13y=F8c13+sin(pi/3)*F7c13+sin(pi/3)*F9c13
76 # Place all of the equations in a vector
77 y=[Eq1x; Eq1y; Eq2x; Eq2y; Eq3x; Eq3y; Eq4x; Eq4y; Eq5x; Eq5y; Eq6x; Eq6y; Eq7x; Eq7y; Eq8x; Eq8y; Eq9x;
Eq9y; Eq10x; Eq10y; Eq11x; Eq11y; Eq12x; Eq12y; Eq13x; Eq13y]
78 return y
79 end
80 # Query the equations to build the model in the form Ax=b
81 n=26
82 m=4
83 A=Array{Float64, 2}(undef, n, 0)
84 Id=zeros(n, n)+I
85 #
86 for k =1 : n
87     ek=Id[:, k]
88     ak=Model(ek, zeros(m, 1))
89     A=[A ak]
90     @show ak
91 end
92 @show det(A)
93 @show size(A)
94 Fw1=2000 #kilo Newtons
95 Fw2=3000
96 Fw3=2000
97 Fwind=500.
98 u=[Fw1; Fw2; Fw3; Fwind]
99 b=-Model(zeros(n, 1), u)
100 # Solve the equations
101 # Uncomment the next line to see if A can be factored without permutations
102 #F=lu(A, Val(false))
103 F=lu(A)
104 L=F.L
105 U=F.U
106 P=F.P
107 # the functions forwardsub and backwardsub are NOT given. You learned how to

```

```

108 # write them in previous chapters.
109 y=forwardsub (L, P*b)
110 x=backwardsub (U, y)

```

Here is the solution to $Ax = b$,

$$x = \begin{bmatrix} -2458.2 \\ 3916.5 \\ -2458.2 \\ 0.0 \\ -5220.0 \\ -3916.5 \\ 1607.1 \\ -5220.0 \\ 0.0 \\ -5095.0 \\ -1607.1 \\ -1857.1 \\ -5095.0 \\ 0.0 \\ -2083.2 \\ 1857.1 \\ -4166.5 \\ -2083.2 \\ 0.0 \\ 4166.5 \\ 4416.5 \\ 6023.5 \\ 4166.5 \\ -500.0 \\ 3391.7 \\ 3608.3 \end{bmatrix}_{26 \times 1} \quad (5.21)$$

The variable names are in the code. We see that there are four “zero force members” highlighted in red, the beams between joints 2 and 10, 4 and 11, 6 and 12, and 8 and 13. While they are not load bearing elements in the nominal structure, they provide important rigidity to the structure and they serve as safety elements in case one of the angled beams should fail. Finally, from the minus signs, we see that under the given loading conditions, more than half of the beams are in tension. We emphasize that this does not invalidate our analysis as tension vs compression is merely a direction convention, just like us assuming that forces that point up or to the right are acting in a positive direction.

5.12 (Optional Read): An Algorithm for Rectangular LU Factorization with Row Permutations

In Chapter 5.6, you were given an algorithm which generates the LU factorization of square matrices. However, it turns out that we do not need much modification to also handle rectangular matrices!

```

1 using LinearAlgebra
2 function luJWG (M: :Array{<:Number, 2})
3     a, b = size (M)
4     n=min (a, b)
5     Temp = deepcopy (M)
6     L = Array{Float64, 2} (undef, a, 0)
7     U = Array{Float64, 2} (undef, 0, b)
8     epsilon=1e-16
9     P=zeros (a, a) + I
10    K=100;
11    for k = 1:n

```

```

12     C = Temp[:,k] # k-th column
13     if maximum(abs.(C)) <= K*epsilon #column of zeros
14         C=0.0*C; # set tiny entries to zero
15         C[k]=1.0;
16         R = Temp[k:k, :] # k-th row
17         Temp=Temp.-C*R;
18         L=[L C];
19         U=[U;R];
20     else # put the biggest entry to the top
21         ii=argmax(abs.(C));
22         nrow=ii[1]
23         P[[k, nrow], :]=P[[nrow, k], :];
24         Temp[[k, nrow], :]=Temp[[nrow, k], :];
25         if k>1
26             L[[k, nrow], :]= L[[nrow, k], :];
27         end
28         pivot = Temp[k,k] # k-th column
29         C=Temp[:,k]/pivot #normalize all entires by pivot
30         R = Temp[k:k, :] # k-th row
31         Temp=Temp-C*R;
32         L=[L C];
33         U=[U;R];
34     end
35 end
36 return L, U, P
37 end

```

5.13 Looking Ahead

Once you have programmed in Julia

- (a) algorithms for forward and back substitution (methods for solving triangular systems of equations), and
- (b) the LU Factorization Algorithm,

you will be well placed for solving systems of linear equations of very large dimension. In the next chapter, we will fill in some of the more standard results that students learn in linear algebra, where the focus is on solving “drill problems” and not “real engineering problems”. In your later courses, most instructors will simply assume you know the more standard material and that you have no clue about the more nuanced approach to problem solving that we have developed in ROB 101. We’ll do our best to throw in some cool computational aspects as we go.

Chapter 6

Determinant of a Matrix Product, Matrix Inverses, Matrix Transposes, and Permutation Matrices

Learning Objectives

- Fill in some gaps that we left during our sprint to an effective means for solving large systems of linear equations.

Outcomes

- Whenever two square matrices A and B can be multiplied, it is true that $\det(A \cdot B) = \det(A) \cdot \det(B)$
- You will learn what it means to “invert a matrix,” and you will understand that you rarely want to actually compute a matrix inverse!
- If $ad - bc \neq 0$, then $\begin{bmatrix} a & b \\ c & d \end{bmatrix}^{-1} = \frac{1}{ad-bc} \begin{bmatrix} d & -b \\ -c & a \end{bmatrix}$.
- Moreover, this may be the only matrix inverse you really want to compute explicitly, unless a matrix has special structure.
- Matrix transpose takes columns of one matrix into the rows of another.

6.1 A very Useful Fact Regarding the Matrix Determinant

Let A and B be two $n \times n$ matrices. Then,

$$\det(A \cdot B) = \det(A) \cdot \det(B).$$

We note that A and B must be square and have the same size for the above useful fact to be true.

Matrix Determinant via LU Factorization

Now suppose that we have done the LU factorization of a square matrix A . Then, using the fact that the determinant of a product of two square matrices is the product their respective determinants, we have that

$$\det(A) = \det(L \cdot U) = \det(L) \cdot \det(U).$$

Because L and U are triangular matrices, each of their determinants is given by the product of the diagonal elements. Hence, we have a way of computing the determinant for square matrices of arbitrary size.

Example 6.1 Compute the matrix determinant of

$$\begin{bmatrix} -2 & -4 & -6 \\ -2 & 1 & -4 \\ -2 & 11 & -4 \end{bmatrix}.$$

Solution Going back to Chap. 5, specifically, (5.8) and (5.9), we have that

$$\underbrace{\begin{bmatrix} -2 & -4 & -6 \\ -2 & 1 & -4 \\ -2 & 11 & -4 \end{bmatrix}}_A = \underbrace{\begin{bmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ 1 & 3 & 1 \end{bmatrix}}_L \cdot \underbrace{\begin{bmatrix} -2 & -4 & -6 \\ 0 & 5 & 2 \\ 0 & 0 & -4 \end{bmatrix}}_U$$

Hence,

$$\det(A) = \underbrace{(1) \cdot (1) \cdot (1)}_{\det(L)} \cdot \underbrace{(-2) \cdot (5) \cdot (-4)}_{\det(U)} = 40.$$

■

Once we know that $\det(A \cdot B) = \det(A) \cdot \det(B)$, we immediately obtain from it

$$\boxed{\det(A \cdot B \cdot C) = \det(A) \cdot \det(B) \cdot \det(C)}$$

because $\det(A \cdot B \cdot C) = \det((A \cdot B) \cdot C) = \det(A \cdot B) \cdot \det(C) = \det(A) \cdot \det(B) \cdot \det(C)$. This formula extends to any product of $n \times n$ matrices.

6.2 Identity Matrix and Matrix Inverse

There is a special square matrix denoted I , or sometimes I_n to emphasize that it is an $n \times n$ matrix, which has ones on its diagonal and zeros everywhere else,

$$I_1 = [1], I_2 = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, I_3 = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}, I_4 = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \text{ etc.}$$

Because I_n is diagonal, we see immediately that $\det(I_n) = 1$.

Multiplication by the Identity Matrix

Suppose that A is an $n \times m$ matrix. Then

$$A = I_n \cdot A = A \cdot I_m.$$

In other words, as long as the matrix product is defined, multiplying a matrix (on either side) by an identity matrix does not change it.

Example 6.2 Suppose that A is 2×3 . Show that $I_2 \cdot A = A$ and $A \cdot I_3 = A$.

Solution: We apply our “second definition” of matrix multiplication,

$$\begin{aligned} \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \cdot \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \end{bmatrix} &= \begin{bmatrix} 1 \\ 0 \end{bmatrix} \cdot \begin{bmatrix} a_{11} & a_{12} & a_{13} \end{bmatrix} + \begin{bmatrix} 0 \\ 1 \end{bmatrix} \cdot \begin{bmatrix} a_{21} & a_{22} & a_{23} \end{bmatrix} \\ &= \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ 0 & 0 & 0 \end{bmatrix} + \begin{bmatrix} 0 & 0 & 0 \\ a_{21} & a_{22} & a_{23} \end{bmatrix} \\ &= \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \end{bmatrix} \end{aligned}$$

and

$$\begin{aligned} \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \end{bmatrix} \cdot \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} &= \begin{bmatrix} a_{11} \\ a_{21} \end{bmatrix} \cdot \begin{bmatrix} 1 & 0 & 0 \end{bmatrix} + \begin{bmatrix} a_{12} \\ a_{22} \end{bmatrix} \cdot \begin{bmatrix} 0 & 1 & 0 \end{bmatrix} + \begin{bmatrix} a_{13} \\ a_{23} \end{bmatrix} \cdot \begin{bmatrix} 0 & 0 & 1 \end{bmatrix} \\ &= \begin{bmatrix} a_{11} & 0 & 0 \\ a_{21} & 0 & 0 \end{bmatrix} + \begin{bmatrix} 0 & a_{12} & 0 \\ 0 & a_{22} & 0 \end{bmatrix} + \begin{bmatrix} 0 & 0 & a_{13} \\ 0 & 0 & a_{23} \end{bmatrix} \\ &= \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \end{bmatrix}. \end{aligned}$$

So, if multiplying a matrix by an identity matrix does not change it, for what is it used? The answer is similar to the number 1.0 in the reals: we know that

$$1.0 \times x = x \times 1.0 = x$$

for all $x \in \mathbb{R}$, which is handy to know, but it’s real importance is that we use it to define the multiplicative inverse of x .

Recall: For a given number $x \in \mathbb{R}$, a second number $y \in \mathbb{R}$ is called **the multiplicative inverse of x** if

$$x \cdot y = y \cdot x = 1.$$

We can prove that if a multiplicative inverse exists, it is unique. And, we can prove that a multiplicative inverse exists if, and only if, $x \neq 0$. Finally, when a multiplicative inverse exists, we typically denote it by $\frac{1}{x}$, though sometimes we might use x^{-1} .

In a similar, manner, the identity matrix is very useful for defining the inverse of a matrix. **In ROB 101, we will only define inverses for square matrices.** It is possible to define inverses for non-square matrices, which are often called Moore-Penrose Inverses, after their inventors. For a first introduction to matrix inverses, handling the topic for square matrices is really enough!

```

1 using LinearAlgebra
2 # Creating an identity matrix in Julia
3
4 n=5
5 myI=zeros(n, n) +I

```

Output

```

5×5 Matrix{Float64}:
 1.0  0.0  0.0  0.0  0.0
 0.0  1.0  0.0  0.0  0.0
 0.0  0.0  1.0  0.0  0.0
 0.0  0.0  0.0  1.0  0.0
 0.0  0.0  0.0  0.0  1.0

```

The Inverse of a Square Matrix

Let A be an $n \times n$ matrix. A second $n \times n$ matrix B is a multiplicative inverse of A if

$$A \cdot B = B \cdot A = I_n.$$

- When such a matrix exists, it can be shown to be unique.
- We say that A is invertible and we denote the inverse by A^{-1} and we simply call it A inverse or the inverse of A .
- It is a major *faux pas* (means, no-no, in French) to write $1/A$ in place of A^{-1} .
- **Very Useful Fact:** A is invertible if, and only if, $\det(A) \neq 0$. It is so useful we will state it a second time below!

Example 6.3 Consider $A = \begin{bmatrix} 4 & 2 \\ 5 & 3 \end{bmatrix}$. Let's check if $B = \frac{1}{2} \begin{bmatrix} 3 & -2 \\ -5 & 4 \end{bmatrix}$ is in fact an inverse for A .

Solution: According to the definition, we need to check that $A \cdot B = B \cdot A = I_2$. Well,

$$A \cdot B = \begin{bmatrix} 4 & 2 \\ 5 & 3 \end{bmatrix} \cdot \left(\frac{1}{2} \begin{bmatrix} 3 & -2 \\ -5 & 4 \end{bmatrix} \right) = \frac{1}{2} \begin{bmatrix} (4 \cdot 3 - 2 \cdot 5) & (4 \cdot (-2) + 2 \cdot 4) \\ (5 \cdot 3 - 3 \cdot 5) & (5 \cdot (-2) + 3 \cdot 4) \end{bmatrix} = \frac{1}{2} \begin{bmatrix} 2 & 0 \\ 0 & 2 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$$

and, because we passed the first part of the test, we do the second part

$$B \cdot A = \left(\frac{1}{2} \begin{bmatrix} 3 & -2 \\ -5 & 4 \end{bmatrix} \right) \cdot \begin{bmatrix} 4 & 2 \\ 5 & 3 \end{bmatrix} = \frac{1}{2} \begin{bmatrix} (3 \cdot 4 - 2 \cdot 5) & (3 \cdot 2 - 2 \cdot 3) \\ ((-5) \cdot 4 + 4 \cdot 5) & ((-5) \cdot 2 + 4 \cdot 3) \end{bmatrix} = \frac{1}{2} \begin{bmatrix} 2 & 0 \\ 0 & 2 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$$

and hence we conclude that $B = A^{-1}$. ■

Most Important Matrix Inverse for ROB 101

Consider $A = \begin{bmatrix} a & b \\ c & d \end{bmatrix}$ and suppose that $\det(A) = a \cdot d - b \cdot c \neq 0$. Then,

$$A^{-1} = \frac{1}{a \cdot d - b \cdot c} \begin{bmatrix} d & -b \\ -c & a \end{bmatrix} = \frac{1}{\det(A)} \begin{bmatrix} d & -b \\ -c & a \end{bmatrix}$$

Applying the above formula for the inverse of a 2×2 matrix immediately gives that

$$\begin{bmatrix} 4 & 2 \\ 5 & 3 \end{bmatrix}^{-1} = \frac{1}{2} \begin{bmatrix} 3 & -2 \\ -5 & 4 \end{bmatrix}.$$

Remarks: We've already stated that if a matrix has an inverse, it is unique (that is, there is only one of them). The definition of the inverse we used was, if A and B are both $n \times n$, then

$$(A \cdot B = B \cdot A = I_n) \iff B = A^{-1} \iff A = B^{-1}.$$

For square matrices with **real (or complex) elements**, you do not have to check both $A \cdot B = I$ and $B \cdot A = I$ to conclude that $B = A^{-1}$. **It is enough to check one of them because**

$$(A \cdot B = I) \iff (B \cdot A = I) \iff B = A^{-1}.$$

For more general notions of numbers, one does have to check both sides.

The Matrix Inverse and the Matrix Determinant:

- Suppose that A is $n \times n$ and invertible. Because the “determinant of a product is the product of the determinants”, we have that

$$1 = \det(I_n) = \det(A \cdot A^{-1}) = \det(A) \cdot \det(A^{-1}).$$

- It follows that if A has an inverse, then $\det(A) \neq 0$ and $\det(A^{-1}) = \frac{1}{\det(A)}$
- The other way around is also true: if $\det(A) \neq 0$, then A has an inverse (one also says that A^{-1} exists). Putting these facts together gives the next result.

Major Important Fact

An $n \times n$ matrix A is invertible if, and only if, $\det(A) \neq 0$.

Another useful fact about matrix inverses is that if A and B are both $n \times n$ and invertible, then their product is also invertible and

$$(A \cdot B)^{-1} = B^{-1} \cdot A^{-1}.$$

Note that the order is swapped when you compute the inverse. To see why this is true, we note that

$$(A \cdot B) \cdot (B^{-1} \cdot A^{-1}) = A \cdot (B \cdot B^{-1}) \cdot A^{-1} = A \cdot (I) \cdot A^{-1} = A \cdot A^{-1} = I.$$

Hence, $(A \cdot B)^{-1} = B^{-1} \cdot A^{-1}$ **and NOT** $A^{-1} \cdot B^{-1}$.

LU and Matrix Inverses

A consequence of this is that if A is invertible and $A = L \cdot U$ is the LU factorization of A , then

$$A^{-1} = U^{-1} \cdot L^{-1}.$$

While we will not spend time on it in ROB 101, it is relatively simple to compute the inverse of a triangular matrix whose determinant is non-zero.

6.3 Utility of the Matrix Inverse and its Computation

The primary use of the matrix inverse is that it provides a **closed-form solution** to linear systems of equations. Suppose that A is square and invertible, then

$$Ax = b \iff x = A^{-1} \cdot b. \tag{6.1}$$

While it is a beautiful thing to write down the closed-form solution given in (6.1) as the answer to $Ax = b$, one should rarely use it in computations. It is much better to solve $Ax = b$ by factoring $A = L \cdot U$ and using back and forward substitution, than to first compute A^{-1} and then multiply A^{-1} and b . Later in ROB 101, we'll learn another good method called the QR factorization.

We (your instructors) know the above sounds bizarre to you! You've been told that $Ax = b$ has a unique solution for x if, and only if, $\det(A) \neq 0$, and you know that A^{-1} exists if, and only if, $\det(A) \neq 0$. Hence, logic tells you know that $x = A^{-1}b$ is the unique solution to $Ax = b$ if, and only if, $\det(A) \neq 0$. So what gives?

(Optional Read) Theory vs Reality: I

If what you really want to solve is $Ax = b$, then it is numerically inefficient to first compute the matrix inverse, A^{-1} , and then to do the matrix multiplication $A^{-1}b$. This is because each column of A^{-1} is the solution to a system of linear equations, namely,

$$A \cdot A^{-1} = I_n \iff A^{-1} = [x^{\text{sol } 1} \quad x^{\text{sol } 2} \quad \dots \quad x^{\text{sol } n}], \text{ where, } Ax^{\text{sol } i} = e_i, 1 \leq i \leq n, \quad (6.2)$$

and, in Julia notation,

$$e_i[j] = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{otherwise} \end{cases}.$$

In other symbols, just to further drive home the point, solving for A^{-1} is equivalent to solving n different linear systems of equations,

$$Ax^{\text{sol } i} = e_i,$$

and then using the solutions as the columns of A^{-1} .

Do you really want to solve n systems of linear equations just to solve one? (That's a rhetorical question, and the answer is NO!)

Example 6.4 Use the method in (6.2) to compute A^{-1} and then apply the formula $x = A^{-1}b$ to solve

$$\begin{bmatrix} 0.9737 & 0.4123 & 1.3861 \\ 0.7551 & 0.6366 & 1.3918 \\ 0.6529 & 0.1277 & 0.7807 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} 0.5568 \\ 0.4081 \\ 0.5018 \end{bmatrix} \quad (6.3)$$

Solution: We compute A^{-1} using the method indicated in (6.2), namely, solving $Ax = e_i$ three times.

$$\begin{bmatrix} 0.9737 & 0.4123 & 1.3861 \\ 0.7551 & 0.6366 & 1.3918 \\ 0.6529 & 0.1277 & 0.7807 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} 1.0 \\ 0.0 \\ 0.0 \end{bmatrix} \implies \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}^{\text{sol } 1} = \begin{bmatrix} 19151.4036 \\ 19150.4446 \\ -19149.6564 \end{bmatrix}$$

$$\begin{bmatrix} 0.9737 & 0.4123 & 1.3861 \\ 0.7551 & 0.6366 & 1.3918 \\ 0.6529 & 0.1277 & 0.7807 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} 0.0 \\ 1.0 \\ 0.0 \end{bmatrix} \implies \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}^{\text{sol } 2} = \begin{bmatrix} -8691.1842 \\ -8688.3033 \\ 8689.9911 \end{bmatrix}$$

$$\begin{bmatrix} 0.9737 & 0.4123 & 1.3861 \\ 0.7551 & 0.6366 & 1.3918 \\ 0.6529 & 0.1277 & 0.7807 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} 0.0 \\ 0.0 \\ 1.0 \end{bmatrix} \implies \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}^{\text{sol } 3} = \begin{bmatrix} -18507.8638 \\ -18511.2973 \\ 18508.1696 \end{bmatrix}$$

Putting these together gives

$$A^{-1} = \begin{bmatrix} 19151.4036 & -8691.1842 & -18507.8638 \\ 19150.4446 & -8688.3033 & -18511.2973 \\ -19149.6564 & 8689.9911 & 18508.1696 \end{bmatrix} \implies x = A^{-1}b = \begin{bmatrix} -8860.1245 \\ -8860.7769 \\ 8860.1837 \end{bmatrix}$$

Remark: Why are there relatively large numbers in the solution of $Ax = b$ when A and b have modestly sized entries? In this case, we could have computed $\det(A) = 1.67e - 5$ and seen ahead of time that A is close to singular. Unfortunately, in the real world, it is easy to have examples where the determinant of a matrix is “relatively far from zero” and yet the matrix is “barely invertible”. ■

(Optional Read) Theory vs Reality: II

While in the world of perfect arithmetic, “ A^{-1} exists if, and only if, $\det(A) \neq 0$ ”, in the approximate arithmetic done by a computer, or by a hand calculator, for that matter, the determinant being nonzero is not a reliable indicator of “numerical invertibility”. The problem is that the determinant of a matrix can be “very nice”, meaning its absolute value is near 1.0, while, from a numerical point of view, the matrix is “barely invertible”.

Example 6.5 Consider $A = \begin{bmatrix} 1 & 10^{-15} \\ 10^{10} & 1 \end{bmatrix}$. For this small example, we can see that A has a huge number and a tiny number in it, but imagine that A is the result of an intermediate computation in your algorithm and hence you’d never look at it, or the matrix is so large, say 50×50 , you would not notice such numbers. If you want to check whether A is invertible or not, you find that $\det(A) = 1 - 10^{-5} = 0.99999$, which is very close to 1.0, and thus the determinant has given us no hint that A has crazy numbers of vastly different sizes in it. ■

(Optional Read) Theory vs Reality: III

The value of $|\det(A)|$ (magnitude of the determinant of A) is a poor predictor of whether or not A^{-1} has very large and very small elements in it, and hence poses numerical challenges for its computation. For typical HW “drill” problems, you rarely have to worry about this. However, for “real” engineering problems, where a typical dimension (size) of A may be 50 or more, then **please please please avoid the computation of the matrix inverse whenever you can!**

The LU factorization is a more reliable predictor of numerical problems that may be encountered when computing A^{-1} . But once you have the LU Factorization, you must ask yourself, do you even need A^{-1} ? In the majority of cases, the answer is NO! Once you have the LU Factorization, solving for $Ax = b$ is cake.

Example 6.6 Consider a 3×3 matrix

$$A = \begin{bmatrix} 100.0000 & 90.0000 & -49.0000 \\ 90.0000 & 81.0010 & 5.4900 \\ 100.0000 & 90.0010 & 59.0100 \end{bmatrix}. \quad (6.4)$$

We compute the determinant and check that it is not close to zero. Indeed, $\det(A) = 0.90100$ (correct to more than ten decimal places¹), and then bravely, we use Julia to compute the inverse, yielding

$$A^{-1} = \begin{bmatrix} -178.9000 & -10,789.0666 & 9,889.0666 \\ 198.7791 & 11,987.7913 & -10,987.981 \\ -110.9878 & -0.1110 & 0.1110 \end{bmatrix},$$

which we see has some really large numbers, such as 11,987. As a contrast, we compute $A = L \cdot U$, the LU factorization (without permutation), yielding

$$L = \begin{bmatrix} 1.0 & 0.0 & 0.0 \\ 0.9 & 1.0 & 0.0 \\ 1.0 & 1.0 & 1.0 \end{bmatrix} \quad U = \begin{bmatrix} 100.000 & 90.000 & -49.000 \\ 0.000 & -0.001 & 99.000 \\ 0.000 & 0.000 & 9.010 \end{bmatrix}.$$

We see that U has a small number on the diagonal, and hence, if we do back substitution to solve $Ux = y$, for example, as part of solving

$$Ax = b \iff L \cdot Ux = b \iff (Ly = b \text{ and } Ux = y)$$

¹ $\det(A) - 0.9010 = -1.7 \times 10^{-11}$.

we know in advance that we'll be dividing by -0.001 , which could easily yield a big number.

Moreover, we see the diagonal of L is $[1.0, 1.0, 1.0]$, and hence $\det(L) = 1$. The diagonal of U is $[100.0, -0.001, 9.01]$, and hence $\det(U) = 0.901$, and we realize that we ended with a number close to 1.0 in magnitude by multiplying a large number and a small number. ■

```

1 # Using the inverse command in Julia
2 using LinearAlgebra
3 using Random
4 Random.seed! (12345678)
5
6 A=randn(6, 6)
7 @show det(A)
8 inv(A)

```

Output

det(A) = 49.345449482491595

```

6×6 Matrix{Float64}:
-0.0313214  0.407173  -0.00884134  -0.550231  0.420031  -0.17968
 0.0915914 -0.0232623  0.268212   -0.0834243  0.400871  0.174108
-0.303567  0.103855   -0.240916   -0.325997  0.345296  0.406821
 0.346315  0.24365    -0.162719   -0.347234  0.346664  0.00357876
-0.121959  0.284931   0.135561    0.305418  -0.0151208 0.140373
 0.0575381  0.248593   0.467928    0.0124006 -0.0322815 0.440638

```

6.4 Matrix Transpose and Symmetric Matrices

Transpose of a Matrix

Let A be an $n \times m$ matrix with entries $[A]_{ij} = a_{ij}$, for $1 \leq i \leq n$ and $1 \leq j \leq m$. Then A^\top , the transpose of A , is an $m \times n$ matrix with its ij -entry equal to the ji entry of A ,

$$[A^\top]_{ij} = [A]_{ji}, 1 \leq i \leq m, 1 \leq j \leq n.$$

Most of us remember the definition more “graphically” as the transpose takes the rows of A and turns them into columns to form the matrix transpose,

$$A^\top := \begin{bmatrix} \boxed{a_{11} \ a_{12} \ \cdots \ a_{1m}} \\ \boxed{a_{21} \ a_{22} \ \cdots \ a_{2m}} \\ \vdots \\ \boxed{a_{n1} \ a_{n2} \ \cdots \ a_{nm}} \end{bmatrix}^\top := \begin{bmatrix} \boxed{a_{11}} & \boxed{a_{21}} & \cdots & \boxed{a_{n1}} \\ \boxed{a_{12}} & \boxed{a_{22}} & \cdots & \boxed{a_{n2}} \\ \vdots & \vdots & \cdots & \vdots \\ \boxed{a_{1m}} & \boxed{a_{2m}} & \cdots & \boxed{a_{nm}} \end{bmatrix}.$$

Equivalently, you can view the matrix transpose as taking each column of one matrix and laying the elements out as rows in the transposed matrix

$$A^\top := \begin{bmatrix} \boxed{a_{11}} & \boxed{a_{12}} & \cdots & \boxed{a_{1m}} \\ \boxed{a_{21}} & \boxed{a_{22}} & \cdots & \boxed{a_{2m}} \\ \vdots & \vdots & \cdots & \vdots \\ \boxed{a_{n1}} & \boxed{a_{n2}} & \cdots & \boxed{a_{nm}} \end{bmatrix}^\top := \begin{bmatrix} \boxed{a_{11} \ a_{21} \ \cdots \ a_{n1}} \\ \boxed{a_{12} \ a_{22} \ \cdots \ a_{n2}} \\ \vdots \\ \boxed{a_{1m} \ a_{2m} \ \cdots \ a_{nm}} \end{bmatrix}.$$

Example 6.7 Compute the transpose of

$$A = \begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{bmatrix} \quad (6.5)$$

Solution:

$$A^T = \begin{bmatrix} 1 & 4 \\ 2 & 5 \\ 3 & 6 \end{bmatrix}. \quad (6.6)$$

Did we just turn the rows into columns or the columns into rows? Trick question! We get the same result either way. :-) Also, is it clear that transposing the matrix in (6.6) gives back the original matrix in (6.5)? We'll let you try this on your own. ■

Appearances can be Deceiving

At first blush, the matrix transpose seems kind of pointless. We will learn, however, special types of matrices whose inverses are equal to their transpose:

- Permutation matrices, which we revisit in the next section.
- Orthogonal matrices, which we study in Chap. 9.5.

Clearly the transpose is super easy to compute. It's awesome that there are useful matrices with inverses that can be computed by simply transposing the matrix!

Properties of the Transpose Operation

For any matrix A ,

$$(A^T)^T = A \text{ and if } A \text{ is square, } \det(A^T) = \det(A).$$

Suppose that A is $n \times m$ and B is $m \times p$, so that $A \cdot B$ makes sense. Then

$$(A \cdot B)^T = B^T \cdot A^T,$$

and just as with the matrix inverse, the order of the matrices is reversed.

Note to self: For non-square matrices, it's often easy to detect your mistake when you attempt to form the product in the wrong order, namely $A^T \cdot B^T$ because, in general, you cannot multiply an $m \times n$ matrix with a $p \times m$ matrix [consider $n = 7$, $m = 5$ and $p = 3$, for example]. **Of course, when A and B are square, Julia will not complain and you will have simply introduced an error in your code! The same can happen when $n = p \neq m$.**

Symmetric and Skew-symmetric Matrices

An $n \times n$ matrix is **symmetric** if $A^T = A$ (the transpose of A is equal to A itself). An $n \times n$ matrix is **skew-symmetric** if $A^T = -A$ (the transpose of A is equal to minus A).

Example 6.8 Determine which matrices are symmetric, skew-symmetric, or neither.

$$A_1 = \begin{bmatrix} 1 & -2 & 5 \\ 2 & 0 & 6 \\ -5 & -6 & 7 \end{bmatrix}, \quad A_2 = \begin{bmatrix} 1 & 3 & 7 \\ 3 & 0 & -6 \\ 7 & -6 & 7 \end{bmatrix}, \quad A_3 = \begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{bmatrix}, \quad \text{and } A_4 = \begin{bmatrix} 0 & -2 & 5 \\ 2 & 0 & 6 \\ -5 & -6 & 0 \end{bmatrix}.$$

Solution: We compute each of the transposes and then compare them to either the original matrix or its negative, to check for the matrix being symmetric, skew-symmetric, or neither.

$$A_1^T = \begin{bmatrix} 1 & 2 & -5 \\ -2 & 0 & -6 \\ 5 & 6 & 7 \end{bmatrix}, \quad A_2^T = \begin{bmatrix} 1 & 3 & 7 \\ 3 & 0 & -6 \\ 7 & -6 & 7 \end{bmatrix}, \quad A_3^T = \begin{bmatrix} 1 & 4 \\ 2 & 5 \\ 3 & 6 \end{bmatrix}, \quad \text{and } A_4^T = \begin{bmatrix} 0 & 2 & -5 \\ -2 & 0 & -6 \\ 5 & 6 & 0 \end{bmatrix}.$$

Because A_1^\top is neither equal to A_1 nor $-A_1$, it is neither symmetric nor skew-symmetric. Because A_2^\top equals A_2 it is symmetric. Because A_3 is not square, it cannot be either symmetric or skew-symmetric. Because A_4^\top equals $-A_4$, it is skew-symmetric. ■

(Optional Read) Remarks: For an $n \times n$ matrix to be skew-symmetric, can you see that its main diagonal must be zero (because it must equal the negative of itself)? Also, that its super-diagonal and sub-diagonal must be negatives of one another, and the same for each successive pair of diagonals? This is illustrated below where the main diagonal is in red font, the super- and sub-diagonals in blue font, and the next diagonals are in bold black font,

$$A_4 = \begin{bmatrix} 0 & -2 & \mathbf{5} \\ 2 & 0 & \mathbf{6} \\ -5 & -6 & 0 \end{bmatrix}.$$

Similarly, for an $n \times n$ matrix to be symmetric, can you see that its main diagonal can be arbitrary (because it just has to be equal to itself)? Also, that its super-diagonal and sub-diagonal must be equal to one another, and the same for each successive pair of diagonals? This is illustrated below where the main diagonal is in red font, the super- and sub-diagonals in blue font, and the next diagonals are in bold black font,

$$A_2 = \begin{bmatrix} 1 & 3 & \mathbf{7} \\ 3 & 0 & -6 \\ \mathbf{7} & -6 & 7 \end{bmatrix}. \quad (6.7)$$

A Key Source of Symmetric Matrices

Let A be an $n \times m$ real matrix. Then $A^\top \cdot A$ is an $m \times m$ symmetric matrix and $A \cdot A^\top$ is an $n \times n$ symmetric matrix.

Remark: Why is the above true? Because $(A^\top \cdot A)^\top = (A)^\top \cdot (A^\top)^\top = A^\top \cdot A$, where we have used the property $(A \cdot B)^\top = B^\top \cdot A^\top$. Similar reasoning shows that $(A \cdot A^\top)^\top = A \cdot A^\top$.

6.5 Revisiting Permutation Matrices

In Chap. 4.7, we jumped ahead and introduced permutation matrices before properly treating the matrix transpose and the matrix inverse. We'll set things right here. Let's recall the definition first.

Permutation Matrices

Matrices that consist of all ones and zeros, with each row **and** each column having a single one, are called permutation matrices. In fact, the requirement that each and every column and row have exactly one 1 and all other entries zero implies that a permutation matrix has to be square. Hence, an $n \times n$ permutation matrix must have exactly n ones and $n(n-1)$ zeros.

Example 6.9 Determine which matrices are permutation matrices.

$$A_1 = \begin{bmatrix} 1 & -1 & 0 \\ 0 & 1 & 2 \\ 2 & 0 & 1 \end{bmatrix}, \quad A_2 = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{bmatrix}, \quad A_3 = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix}, \quad A_4 = \begin{bmatrix} 1 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix}, \quad \text{and} \quad A_5 = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{bmatrix}.$$

Solution: A_1 is not a permutation matrix because it has entries that are neither zero nor one, such $a_{12} = -1$ and $a_{23} = a_{31} = 2$. A_2 is a permutation matrix because we can check that **each and every row and column** has precisely one 1 and all other entries are 0's. In passing, we note that A_2 is not symmetric, and thus permutation matrices are not required to be symmetric. A_3 is not a permutation matrix because its third column is all zeros; one could also say that it fails because it is not square. A_4 is 3×3 and has exactly 3 entries equal to one and all the rest zero; however, it is not a permutation matrix because its first column has two ones. A_5 is not a permutation matrix because its second row has two ones; it also fails because its third column has two ones. ■

Inverting a Permutation Matrix is Cake

If P is a permutation matrix, then $P^\top \cdot P = P \cdot P^\top = I$. In other words,

$$P^{-1} = P^\top.$$

Here is one of the permutations we considered before where applying the permutation twice does not undo the swapping of rows,

$$\begin{bmatrix} 1 \\ 2 \\ 3 \\ 4 \\ 5 \end{bmatrix} \rightarrow \begin{bmatrix} 4 \\ 2 \\ 1 \\ 5 \\ 3 \end{bmatrix}; \quad (6.8)$$

see (5.15). As we did before, we put the 5×5 identity matrix on the left and the corresponding permutation matrix P on the right

$$I = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix} \leftrightarrow P = \begin{bmatrix} 0 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \end{bmatrix} \leftrightarrow \begin{bmatrix} 3 \rightarrow 1 \\ 2 \rightarrow 2 \\ 5 \rightarrow 3 \\ 1 \rightarrow 4 \\ 4 \rightarrow 5 \end{bmatrix}. \quad (6.9)$$

The matrix P is still just a re-ordering of the rows of I . We can check that P is not a symmetric matrix by either applying the “diagonal test” indicated in (6.7), or by computing the transpose

$$P^\top = \begin{bmatrix} 0 & 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 & 0 \end{bmatrix}$$

and noting that $P^\top \neq P$. Direct, albeit tedious multiplication gives that

$$P \cdot P^\top = e_4 \cdot e_4^\top + e_2 \cdot e_2^\top + e_1 \cdot e_1^\top + e_5 \cdot e_5^\top + e_3 \cdot e_3^\top = I_5 \quad \text{and} \\ P^\top \cdot P = e_3 \cdot e_3^\top + e_2 \cdot e_2^\top + e_5 \cdot e_5^\top + e_1 \cdot e_1^\top + e_4 \cdot e_4^\top = I_5,$$

where e_i is the i -th column of the identity matrix. Hence, computing the inverse of a permutation matrix is a snap!

Below is a snippet of Julia code that shows how to take a desired permutation of rows and build the corresponding permutation matrix from the identity matrix.

Objective: Bring the fourth row to the top, leave the second row where it is, move the first row to the third row, move the fifth row to the fourth row, and move the third row to the last position. In other symbols,

$$\begin{bmatrix} 4 \rightarrow 1 \\ 2 \rightarrow 2 \\ 1 \rightarrow 3 \\ 5 \rightarrow 4 \\ 3 \rightarrow 5 \end{bmatrix}$$

```
1 # Building a permutation matrix
2 using LinearAlgebra
3 p=[4, 2, 1, 5, 3]
4 myI=zeros(5,5)+I
5 P=myI[p, :]
```

Output we obtain the permutation matrix, P ,

```
5×5 Matrix{Float64}:
 0.0  0.0  0.0  1.0  0.0
 0.0  1.0  0.0  0.0  0.0
 1.0  0.0  0.0  0.0  0.0
 0.0  0.0  0.0  0.0  1.0
 0.0  0.0  1.0  0.0  0.0
```

We illustrate two ways to permute the elements of a column vector, similar to what we saw in Chapter 5.10,

```
1 b=[1 2 3 4 5]'  
2 [b b[p] P*b]
```

Output yielding,

```
5x3 Matrix{Float64}:  
 1.0  4.0  4.0  
 2.0  2.0  2.0  
 3.0  1.0  1.0  
 4.0  5.0  5.0  
 5.0  3.0  3.0
```

so that we see once again that $b[p] = P * b$.

6.6 Matrix Determinants, Matrix Inverses, and the Matrix Transpose

Our final useful facts of the Chapter deal with the transpose of a square matrix.

Determinant and Matrix Inverse of A^T

Some useful facts:

- For any $n \times n$ square matrix A , $\det(A^T) = \det(A)$.
- Hence, if A is invertible, then so is its transpose.
- Moreover, when $\det(A) \neq 0$, the inverse of A^T is the transpose of the inverse of A , that is

$$(A^T)^{-1} = (A^{-1})^T.$$

Remark: Because the order does not matter when inverting and transposing a matrix, one often writes A^{-T} for $(A^T)^{-1}$.

(Optional Read:) Even though proofs are not a big thing in ROB 101, we'll sketch out the reasoning for these two facts.

We recall that by the definition of a matrix inverse, it must satisfy $A \cdot A^{-1} = A^{-1} \cdot A = I_n$. By the product rule for the matrix transpose, we have that

$$(A^{-1})^T \cdot A^T = (A \cdot A^{-1})^T = (A^{-1} \cdot A)^T = A^T \cdot (A^{-1})^T = I_n^T = I_n.$$

Because

$$(A^{-1})^T \cdot A^T = A^T \cdot (A^{-1})^T = I_n,$$

we deduce that $(A^T)^{-1} = (A^{-1})^T$.

For the next part, we will need to use one fact that we have not covered in the book, namely, if P is a permutation matrix, then $\det(P) = \det(P^T) = \pm 1$. Suppose that A is square and perform its LU Factorization, $P \cdot A = L \cdot U$. Then, multiplying both sides by P^T and using $P^T \cdot P = I_n$, we have $A = P^T \cdot L \cdot U$. Hence, by the product rule of determinants, we have that

$$\det(A) = \det(P^T) \cdot \det(L) \cdot \det(U) = \det(P^T) \cdot \det(U) = \det(P) \cdot \det(U),$$

where we used $\det(L) = 1$ (because L is uni-lower-triangular) and $\det(P^T) = \det(P)$. Next, we note that $A^T = U^T \cdot L^T \cdot P$, using the product rule for the matrix transpose. Hence,

$$\det(A^T) = \det(U^T) \cdot \det(L^T) \cdot \det(P).$$

The transpose operation takes U to a lower-triangular matrix with $\text{diag}(U) = \text{diag}(U^\top)$ and hence $\det(U) = \det(U^\top)$. Similarly, the transpose operation takes L to a uni-upper-triangular matrix with $\text{diag}(L) = \text{diag}(L^\top)$ and hence $\det(L) = \det(L^\top) = 1$. We conclude that

$$\begin{aligned}\det(A) &= \det(P) \cdot \det(U) \\ \det(A^\top) &= \det(U) \cdot \det(P),\end{aligned}$$

and hence, $\det(A^\top) = \det(A) = \pm \det(U)$.

6.7 Looking Ahead

You have now seen the most important material for solving systems of linear equations. You have been cautioned about the **“theory vs practice gap”** when it comes to computing matrix inverses. In fact, your instructors often say, with a mix of seriousness and joking, that “we don’t hang out with people who compute matrix inverses.” The purpose of the statement is to drive home the fact that computing a matrix inverse is a numerically delicate operation and is often a sloppy way to solve a real problem with many variables. for “toy problems” typically seen in lower-level courses, sure, you can use the matrix inverse and not suffer any adverse consequences. **But, did you come to Michigan to learn how to solve toy problems?**

Our next task is to develop a deeper understanding of vectors and special sets built out of vectors, called “vector spaces.” We’ll learn about linear independence, basis vectors, dimension, as well as how to measure the size of a vector and the distance between two vectors. **This knowledge will allow us to analyze systems of linear equations $Ax = b$ that do not have an exact solution!** Right now, this might seem bizarre: if you already know there is not an exact solution, what could you possibly be analyzing? We will be seeking an approximate answer that minimizes the error in the solution, where the error is defined as

$$e := Ax - b.$$

We seek a value of x that makes e as “small as possible in magnitude”. **In Robotics as in life, most interesting problems do not have exact answers. The goal is then to find approximate answers that are good enough!**

Chapter 7

The Vector Space \mathbb{R}^n : Part 1

Learning Objectives

- Instead of working with individual vectors, we will work with a collection of vectors.
- Our first encounter with some of the essential concepts in Linear Algebra that go beyond systems of equations.

Outcomes

- Vectors as n -tuples of real numbers
- \mathbb{R}^n as the collection of all n -tuples of real numbers
- Linear combinations of vectors
- Linear independence of vectors
- Relation of these concepts to the existence and uniqueness of solutions to $Ax = b$.
- How the LU Factorization makes it very straightforward to check the linear independence of a set of vectors, and how a small modification to the LU factorization makes it equally straightforward to check if one vector is a linear combination of a set of vectors.

Warning! Theory Ahead!! Additional Study Time May Be Required

For many students, the concepts in Chapters 7, 9 and 10 are significantly more challenging than the material in any of the other chapters of the book. Why? Well, the reasons vary from person to person, but the biggest reason seems to be the level of abstraction. All of our previous work has pretty much dealt with solving equations and that is something most students of ROB 101 can wrap their heads around.

To be considered knowledgeable in Linear Algebra, you need to understand

- **Linear Combinations and Linear Independence of Vectors**
- **Subspaces**
- **Span of a set of vectors**
- **Basis vectors and coordinates**
- **Dimension of a subspace**
- **Relations of the above to Matrices**

7.1 Motivation

At this point, if we have a matrix equation, $Ax = b$, we understand that \bar{x} is a solution to the equation if, and only if, $A\bar{x} - b = 0$; in other words, plugging \bar{x} into the equation results in there being “no error.” In this chapter, we are setting ourselves up for deeper questions, such as what is the set of all solutions to an equation of the form $Ax = b$? What does it even mean to talk about “all solutions” to an equation?

To make this last question a bit more concrete, let’s suppose that we have found two solutions to the equation $Ax = 0_{n \times 1}$ (yes, we took $b = 0_{n \times 1}$), and we call them \bar{x} and $\bar{\bar{x}}$. We now let $\alpha \in \mathbb{R}$ and $\beta \in \mathbb{R}$ be two real numbers and we define a third potential solution by

$$\tilde{x} := \alpha\bar{x} + \beta\bar{\bar{x}}. \quad (7.1)$$

Is there any chance it is really a solution? In other words, just by doing simple algebraic operations on existing solutions, can we find other solutions? Let’s find out by plugging \tilde{x} into the equation,

$$\begin{aligned} A\tilde{x} &= A(\alpha\bar{x} + \beta\bar{\bar{x}}) \\ &= A(\alpha\bar{x}) + A(\beta\bar{\bar{x}}) \\ &= \alpha A\bar{x} + \beta A\bar{\bar{x}} \\ &= 0 + 0 \\ &= 0. \end{aligned}$$

Bingo, \tilde{x} is also a solution. Since there can be more than one solution, we are motivated to define

$$S_{\text{all}} := \{x \mid Ax = 0_{n \times 1}\}, \quad (7.2)$$

the set of all possible solutions. Assuming we could compute such a set, what would we even do with it? Later in the course, we’ll define what it means to say that one vector $\bar{x} \in S_{\text{all}}$ is “better than” another vector $\bar{\bar{x}} \in S_{\text{all}}$, in other words, we’ll define what it means for one solution (to an engineering problem) to be better than another solution (to an engineering problem). If you can compare one solution to another, then you can potentially quantify what is a “best” solution to a problem! For example, we could potentially pose problems about how to go from point A to point B using the least amount of energy, or how to manufacture something at lowest cost.

We’re straying a bit from Linear Algebra here, but hopefully, the teaser was kind of fun. We summarize:

- It makes sense to perform algebraic operations on solutions of equations, as in (7.1).
- It is at least plausible that understanding the “structure” of abstract sets like (7.2) will make it easier to sort through the set to find a “best element” in the set, when later in the course, we are able to define what we mean by “best”!

Need Intuition?

The YouTube Channel **3Blue1Brown** has a video series on the “essence” of Linear Algebra. Here is one you may like which talks about vectors, points, and coordinates https://youtu.be/fNk_zzaMoSs; there are many others. The series has wonderful animations. **If you find yourself struggling with the abstract concepts we are about to encounter, then this video series may be for you!** Others of you will not need the animations and the fact that the animations have to be done in dimensions two and three may actually inhibit you from mastering the material. We all learn in different ways.

7.2 Vectors in \mathbb{R}^n and Some Basic Algebra

An n -tuple is a fancy name for an ordered list of n numbers, (x_1, x_2, \dots, x_n) . If this sounds a lot like a vector, then good, because **we will systematically identify**

$$\mathbb{R}^n := \{(x_1, x_2, \dots, x_n) \mid x_i \in \mathbb{R}, 1 \leq i \leq n\} \quad (7.3)$$

with column vectors. In other words, for us,

$$\mathbb{R}^n := \{(x_1, x_2, \dots, x_n) \mid x_i \in \mathbb{R}, 1 \leq i \leq n\} \iff \left\{ \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix} \mid x_i \in \mathbb{R}, 1 \leq i \leq n \right\} =: \mathbb{R}^n. \quad (7.4)$$

The choice of identifying n -tuples of numbers with column vectors instead of row vectors is completely arbitrary, and yet, it is very common.

Columns of Matrices are Vectors and Vice Versa

Suppose that A is an $n \times m$ matrix, then its columns are vectors in \mathbb{R}^n and conversely, given vectors in \mathbb{R}^n , we can stack them together and form a matrix:

$$A = \begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1m} \\ a_{21} & a_{22} & \cdots & a_{2m} \\ \vdots & \vdots & \ddots & \vdots \\ a_{n1} & a_{n2} & \cdots & a_{nm} \end{bmatrix} = [a_1^{\text{col}} \quad a_2^{\text{col}} \quad \cdots \quad a_m^{\text{col}}] \iff a_j^{\text{col}} := \begin{bmatrix} a_{1j} \\ a_{2j} \\ \vdots \\ a_{nj} \end{bmatrix} \in \mathbb{R}^n, 1 \leq j \leq m \quad (7.5)$$

Example 7.1 The following are all vectors in \mathbb{R}^4

$$u = \begin{bmatrix} 1 \\ -2 \\ \pi \\ \sqrt{17} \end{bmatrix}, \quad v = \begin{bmatrix} 4.1 \\ -1.1 \\ 0.0 \\ 0.0 \end{bmatrix}, \quad w = \begin{bmatrix} 10^3 \\ 0 \\ 0.7 \\ 1.0 \end{bmatrix}.$$

Use them to make a 4×3 matrix.

Solution:

$$A = \begin{bmatrix} 1 & 4.1 & 10^3 \\ -2 & -1.1 & 0.0 \\ \pi & 0.0 & 0.7 \\ \sqrt{17} & 0.0 & 1.0 \end{bmatrix}.$$

■

Example 7.2 The matrix A is a 2×4 . Extract its columns to form vectors in \mathbb{R}^2 .

$$A = \begin{bmatrix} 1 & 2 & 3 & 4 \\ 5 & 6 & 7 & 8 \end{bmatrix}.$$

Solution:

$$a_1^{\text{col}} = \begin{bmatrix} 1 \\ 5 \end{bmatrix} \in \mathbb{R}^2, \quad a_2^{\text{col}} = \begin{bmatrix} 2 \\ 6 \end{bmatrix} \in \mathbb{R}^2, \quad a_3^{\text{col}} = \begin{bmatrix} 3 \\ 7 \end{bmatrix} \in \mathbb{R}^2, \quad a_4^{\text{col}} = \begin{bmatrix} 4 \\ 8 \end{bmatrix} \in \mathbb{R}^2.$$

■

For many applications, we go back and forth between columns of matrices and vectors. In such cases, we are really treating vectors as *ordered lists of numbers*, where *ordered* simply means there is a first element, a second element, etc., just as we did way back in Chapter 2. In other applications, we think of vectors as points in space. This is especially common with \mathbb{R}^2 , the two-dimensional plane, or \mathbb{R}^3 , the three-dimensional world in which we live, but if you've heard of Einstein's "space-time continuum", the space in

which his theory of general relativity lives, then you know about four-dimensional space!

Some Books use Different Notation for Vectors in \mathbb{R}^n

For $\mathbb{R}^n := \{(x_1, x_2, \dots, x_n) \mid x_j \in \mathbb{R}, j = 1, 2, \dots, n\}$, the current way to denote a vector is

$$x = \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix}.$$

The real numbers x_j are called the **components** of the **vector** x or the **entries** of the **vector** x .

You will find many textbooks that place arrows over vectors, as in

$$\vec{x} = \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix}$$

The “arrow notation” for vectors is falling out of favor. **When your instructors write out vectors, they will not use the “arrow” notation.** This actually helps to make your mathematics look closer to your programming, which is ALWAYS a good thing, but it is not the primary motivation.

Definitions: Consider two vectors $x \in \mathbb{R}^n$ and $y \in \mathbb{R}^n$. We define their vector sum by

$$x + y = \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix} + \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_n \end{bmatrix} := \begin{bmatrix} x_1 + y_1 \\ x_2 + y_2 \\ \vdots \\ x_n + y_n \end{bmatrix},$$

that is, we sum their respective components or entries. Let α be a real number. Then we define

$$\alpha x := \begin{bmatrix} \alpha x_1 \\ \alpha x_2 \\ \vdots \\ \alpha x_n \end{bmatrix},$$

that is, to multiply a vector by a real number, we multiply each of the components of the vector by the same real number. ■

With these definitions, two vectors x and y are **equal** if, and only if, they have the same components,

$$x := \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix} = \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_n \end{bmatrix} =: y \iff x_j = y_j \text{ for all } 1 \leq j \leq n.$$

Said another way,

$$x = y \iff x - y = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix} = 0_{n \times 1},$$

the zero vector in \mathbb{R}^n .

Remark: Normally, the zero vector in \mathbb{R}^n is simply denoted by 0. From time to time, to help our learning, we'll be very explicit and write it as $0_{n \times 1}$ in order to emphasize its dimension!




Figure 7.1: Images of the xy -plane with Cartesian coordinates x and y , courtesy of the WikiMedia Commons. (a) shows a vector drawn to the **point** $(2, 3)$ and labeled suggestively as a vector from the origin O to point A . (b) shows a **point** P represented with coordinates (x, y) , also with a vector drawn to it. These are both fine interpretations of vectors. In ROB 101, points and vectors are the same thing: an ordered list of numbers.

Figure 7.1 shows two geometric depictions of vectors in 2D, one emphasizing a vector as a directed line segment from the origin O to a point A located at $(2, 3)$, while the second is a directed line segment from the origin O to a point P located at (x, y) . Both of these are fine interpretations of vectors. In \mathbb{R}^2 , graphical depictions of vectors are relatively straightforward to understand. Figure 7.2a illustrates so-called unit vectors along the xyz -axes in \mathbb{R}^3 , while Fig. 7.2b shows a point in \mathbb{R}^3 with Cartesian coordinates (x, y, z) . These depictions of vectors and points in \mathbb{R}^3 are fairly simple, but it is not hard to give other examples in \mathbb{R}^3 that are harder to “picture” as a 2D-image, which is all we can show on a sheet of paper. So yes, it is easy to get lost in 3D-space and if you have a hard time “imagining” (that is, representing visually in your head for the purpose of building intuition) objects in 3D, you are not alone. Anyone want to try \mathbb{R}^{27} , that is, 27D-space?

In ROB 101, points in \mathbb{R}^n are vectors. We treat points and vectors as being different ways of representing the same mathematical object: an ordered list of numbers. Don't get hung up on the geometric aspect of vectors. In Physics, Dynamics, and Electromagnetics courses, you'll be challenged to “imagine” vectors in 3D. Deal with it there. For now, it is OK to think of vectors as lists of numbers that obey some basic algebraic relationships and nothing more. Besides, when dealing with a vector that has hundreds of components, what else can you do?

3 Special Vectors in \mathbb{R}^3

Figure 7.2a defines “unit” vectors \hat{i} , \hat{j} , and \hat{k} . You will see these a lot in Physics. The more accepted mathematical notation in Linear Algebra is

$$e_1 := \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, e_2 := \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}, \text{ and } e_3 := \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}.$$




Figure 7.2: Images of 3D space with Cartesian coordinates x , y , and z using two different color schemes. (a), courtesy of Wikimedia Commons, emphasizes that the vector $a = a_x + a_y + a_z$ is the sum of three vectors lying along the x -axis, y -axis, and z -axis, respectively. (b), courtesy of Prof. Maani Ghaffari, emphasizes that the point $P = (x, y, z)$ has an x -component P_x , a y -component P_y , and z -component P_z . In ROB 101, points and vectors are the same thing. Moreover, we do not rely on your ability to imagine vectors in n -dimensional space for $n > 2$.

Properties of Vector Addition and Scalar Times Vector Multiplication

1. Addition is commutative: For any $x \in \mathbb{R}^n$ and $y \in \mathbb{R}^n$,

$$x + y = y + x.$$

2. Addition is associative: For any $x \in \mathbb{R}^n$, $y \in \mathbb{R}^n$ and $z \in \mathbb{R}^n$,

$$(x + y) + z = x + (y + z).$$

3. Scalar multiplication is associative: For any $\alpha \in \mathbb{R}$, $\beta \in \mathbb{R}$, and any $x \in \mathbb{R}^n$ in

$$\alpha(\beta x) = (\alpha\beta)x.$$

4. Scalar multiplication is distributive: For any $\alpha \in \mathbb{R}$, $\beta \in \mathbb{R}$ and for any $x \in \mathbb{R}^n$, $y \in \mathbb{R}^n$,

$$(\alpha + \beta)x = \alpha x + \beta x,$$

and

$$\alpha(x + y) = \alpha x + \alpha y.$$

All of these properties follow from the corresponding properties for real numbers. Hence, there is no real reason to make a big deal about them. You've been adding vectors and multiplying them by constants in Julia for several weeks now!

7.3 The Notion of Linear Combinations

Sums of vectors times scalars are called linear combinations. We'll introduce the concept in two ways: first by studying what Ax really means in terms of the columns of A and the components of x , and in a second pass, a more "pure" or "abstract" definition in

\mathbb{R}^n . If you enjoyed the previous 3Blue1Brown video, then I suggest you check out this one as well <https://www.youtube.com/watch?v=k7RM-ot2NWY>.

7.3.1 Linear combinations through the lens of $Ax=b$

Let's consider $Ax = b$ and see if we truly understand the product of an $n \times m$ matrix A and an $m \times 1$ column vector x , and then the equation $Ax = b$! We write

$$A = \begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1m} \\ a_{21} & a_{22} & \cdots & a_{2m} \\ \vdots & \vdots & \ddots & \vdots \\ a_{n1} & a_{n2} & \cdots & a_{nm} \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_m \end{bmatrix}. \quad (7.6)$$

Then, using our column times row method for matrix multiplication, we have

$$Ax = \begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1m} \\ a_{21} & a_{22} & \cdots & a_{2m} \\ \vdots & \vdots & \ddots & \vdots \\ a_{n1} & a_{n2} & \cdots & a_{nm} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_m \end{bmatrix} = \begin{bmatrix} a_{11} \\ a_{21} \\ \vdots \\ a_{n1} \end{bmatrix} x_1 + \begin{bmatrix} a_{12} \\ a_{22} \\ \vdots \\ a_{n2} \end{bmatrix} x_2 + \cdots + \begin{bmatrix} a_{1m} \\ a_{2m} \\ \vdots \\ a_{nm} \end{bmatrix} x_m. \quad (7.7)$$

If the above is not clicking for you, go back and look at our second method for matrix multiplication in Chapter 4. Note that the rows of the vector x are simply its components; another way to look at it, because x is $m \times 1$, its rows are 1×1 .

The next step is to move the x_i in front of the column vectors of A , as in

$$Ax = x_1 \begin{bmatrix} a_{11} \\ a_{21} \\ \vdots \\ a_{n1} \end{bmatrix} + x_2 \begin{bmatrix} a_{12} \\ a_{22} \\ \vdots \\ a_{n2} \end{bmatrix} + \cdots + x_m \begin{bmatrix} a_{1m} \\ a_{2m} \\ \vdots \\ a_{nm} \end{bmatrix}.$$

Because we are used to thinking of the x_i as variables or unknowns, let's substitute in a "numerical" value, such as α_i . Of course, this has not changed anything, but it might look different to you,

$$A \begin{bmatrix} \alpha_1 \\ \alpha_2 \\ \vdots \\ \alpha_m \end{bmatrix} = \alpha_1 \begin{bmatrix} a_{11} \\ a_{21} \\ \vdots \\ a_{n1} \end{bmatrix} + \alpha_2 \begin{bmatrix} a_{12} \\ a_{22} \\ \vdots \\ a_{n2} \end{bmatrix} + \cdots + \alpha_m \begin{bmatrix} a_{1m} \\ a_{2m} \\ \vdots \\ a_{nm} \end{bmatrix}.$$

Linear Combination of the Columns of A

Definition The following sum of scalars times vectors,

$$\alpha_1 \begin{bmatrix} a_{11} \\ a_{21} \\ \vdots \\ a_{n1} \end{bmatrix} + \alpha_2 \begin{bmatrix} a_{12} \\ a_{22} \\ \vdots \\ a_{n2} \end{bmatrix} + \cdots + \alpha_m \begin{bmatrix} a_{1m} \\ a_{2m} \\ \vdots \\ a_{nm} \end{bmatrix}, \quad (7.8)$$

is called a **linear combination** of the columns of A . When we set $A\alpha = b$, and turn it around as $b = A\alpha$, we arrive at an important **Fact**: a vector $\alpha \in \mathbb{R}^m$ is a solution to $Ax = b$ (that is, $A\alpha = b$) if, and only if

$$b = \alpha_1 \begin{bmatrix} a_{11} \\ a_{21} \\ \vdots \\ a_{n1} \end{bmatrix} + \alpha_2 \begin{bmatrix} a_{12} \\ a_{22} \\ \vdots \\ a_{n2} \end{bmatrix} + \cdots + \alpha_m \begin{bmatrix} a_{1m} \\ a_{2m} \\ \vdots \\ a_{nm} \end{bmatrix}, \quad (7.9)$$

that is, b can be expressed as a linear combination of the columns of A . We will revisit this later, but (7.9) is a primary motivation for introducing the notion of linear combinations.

7.3.2 Linear Combinations in \mathbb{R}^n

Linear Combination

A vector $v \in \mathbb{R}^n$ is a **linear combination** of $\{u_1, u_2, \dots, u_m\} \subset \mathbb{R}^n$ if there exist real numbers $\alpha_1, \alpha_2, \dots, \alpha_m$ such that

$$v = \alpha_1 u_1 + \alpha_2 u_2 + \dots + \alpha_m u_m. \quad (7.10)$$

Example 7.3 *Because*

$$\underbrace{\begin{bmatrix} -3 \\ -5 \\ -7 \end{bmatrix}}_v = 2 \underbrace{\begin{bmatrix} 3 \\ 2 \\ 1 \end{bmatrix}}_{u_1} - 9 \underbrace{\begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}}_{u_2},$$

we have that v is a linear combination of $\{u_1, u_2\}$.

When you are given the coefficients, it is easy to observe that a given vector is a linear combination of other vectors. But when you have to check if such coefficients exist or do not exist, then it's just a wee bit more challenging!

Example 7.4 *Is the vector $v = \begin{bmatrix} 4 \\ 4 \\ 4 \end{bmatrix}$ a linear combination of*

$$u_1 = \begin{bmatrix} 3 \\ 1 \\ -1 \end{bmatrix} \quad \text{and} \quad u_2 = \begin{bmatrix} 2 \\ -2 \\ 1 \end{bmatrix} ?$$

Solution: The vector v is a linear combination of the vectors u_1 and u_2 if, and only if, there exist real numbers α_1 and α_2 such that

$$\alpha_1 u_1 + \alpha_2 u_2 = v. \quad (7.11)$$

What we have done here is apply¹ **the definition of a linear combination**. So, we write down the linear equations corresponding to (7.11), namely

$$\begin{aligned} \alpha_1 \begin{bmatrix} 3 \\ 1 \\ -1 \end{bmatrix} + \alpha_2 \begin{bmatrix} 2 \\ -2 \\ 1 \end{bmatrix} &= \begin{bmatrix} 4 \\ 4 \\ 4 \end{bmatrix} \\ &\Downarrow \\ \begin{bmatrix} 3 & 2 \\ 1 & -2 \\ -1 & 1 \end{bmatrix} \begin{bmatrix} \alpha_1 \\ \alpha_2 \end{bmatrix} &= \begin{bmatrix} 4 \\ 4 \\ 4 \end{bmatrix} \end{aligned} \quad (7.12)$$

and see if we can find a solution for the **unknowns** α_1 and α_2 ! What makes this look hard is that we have a non-square (rectangular) system of linear equations, namely, three equations and two unknowns.

What are the three equations? They may look more familiar to you if we write them out like this,

$$\begin{aligned} 3\alpha_1 + 2\alpha_2 &= 4 \\ \alpha_1 - 2\alpha_2 &= 4 \\ -\alpha_1 + \alpha_2 &= 4. \end{aligned}$$

A systematic way to approach such problems is to observe that if α_1 and α_2 are to simultaneously satisfy all three equations, then they must necessarily satisfy any two of them. Moreover, if you find that the solution to the resulting square system of two equations and two unknowns is unique, then two things are possible:

¹In the beginning, many of us are confused about how to go about solving a problem. Sound advice is to start with the definitions at your disposal. Then try to turn the definition into a set of equations to be solved.

- the solution of the smaller square system of equations also satisfies the equation(s) you did not use, giving you a solution to the full set of equations, or
- the solution of the smaller square system of equations does not satisfy the equation(s) you did not use, telling you that the full set of equations is inconsistent and does not have a solution.

In our case, if we remove the last equation, we have

$$\begin{bmatrix} 3 & 2 \\ 1 & -2 \end{bmatrix} \begin{bmatrix} \alpha_1 \\ \alpha_2 \end{bmatrix} = \begin{bmatrix} 4 \\ 4 \end{bmatrix}, \quad (7.13)$$

and we observe that the matrix multiplying the vector of unknowns has determinant -8 , and thus (7.13) has a unique solution. A bit of matrix magic with our formula for a 2×2 inverse gives

$$\alpha_1 = 2 \text{ and } \alpha_2 = -1.$$

Do these values satisfy the equation we did not use, namely, the last row of (7.13),

$$\begin{bmatrix} -1 & 1 \end{bmatrix} \begin{bmatrix} \alpha_1 \\ \alpha_2 \end{bmatrix} = \begin{bmatrix} 4 \end{bmatrix}?$$

Clearly not, because $-1\alpha_1 + \alpha_2 = -3 \neq 4$, and **hence, v is NOT a linear combination of u_1 and u_2 .**

The key point is that α_1 and α_2 must satisfy all three equations in (7.12). It does not matter in what order we use the equations when seeking a solution. To illustrate that, we'll first solve the last two equations,

$$\begin{aligned} \alpha_1 - 2\alpha_2 &= 4 \\ -\alpha_1 + \alpha_2 &= 4, \end{aligned}$$

which yields, $\alpha_1 = -12$ and $\alpha_2 = -8$. We then check if these values satisfy the first equation

$$3\alpha_1 + 2\alpha_2 = 4 \iff -36 - 16 = 4 \iff -52 = 4,$$

which is false, and therefore, v is NOT a linear combination of u_1 and u_2 .

Rework the problem: Let's keep u_1 and u_2 as given and change v to

$$\tilde{v} = \begin{bmatrix} 0 \\ -8 \\ 5 \end{bmatrix}.$$

Applying the above strategy to the equations

$$\begin{aligned} 3\alpha_1 + 2\alpha_2 &= 0 \\ \alpha_1 - 2\alpha_2 &= -8 \\ -\alpha_1 + \alpha_2 &= 5 \end{aligned}$$

yields $\alpha_1 = -2$ and $\alpha_2 = 3$, as you can verify by direct substitution. **Hence, \tilde{v} is a linear combination of u_1 and u_2 .** ■

Remark: Checking whether a given vector can be written as a linear combination of other vectors always comes down to solving a system of linear equations. If that system is square with a non-zero determinant, then solving the system of equations is cake. Otherwise, solving the equations by hand is mostly painful. **We will need to develop a better way to check if a vector is, or is not, a linear combination of other vectors!**

7.4 Existence of Solutions to $Ax=b$

Because it is so important, we now revisit the relation of linear combinations to the existence of solutions to systems of linear equations. We want to make sure you did NOT miss something important when reading Chapter 7.3.1. If you are confident of your knowledge, then feel free to move on to the next section. Otherwise, please read on!

Let A be an $n \times m$ matrix, not necessarily square, as in (7.5). Recall that columns of A and vectors in \mathbb{R}^n are the same thing. We write out $Ax = b$ in all its glory,

$$\underbrace{\begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1m} \\ a_{21} & a_{22} & \cdots & a_{2m} \\ \vdots & \vdots & \ddots & \vdots \\ a_{n1} & a_{n2} & \cdots & a_{nm} \end{bmatrix}}_A \underbrace{\begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_m \end{bmatrix}}_x = \underbrace{\begin{bmatrix} b_1 \\ b_2 \\ \vdots \\ b_n \end{bmatrix}}_b. \quad (7.14)$$

Existence of Solutions

The equation $Ax = b$ has a solution if, and only if, b can be written as a linear combination of the columns of A .

The following is more or less a proof: Suppose that \bar{x} satisfies $A\bar{x} = b$, which is the same thing as saying \bar{x} is a solution of $Ax = b$. Then, doing the indicated multiplication of $A\bar{x}$ via our now favorite method, the columns of A times the rows of the vector \bar{x} , which are its scalar entries \bar{x}_i , yields

$$\begin{bmatrix} a_{11} \\ a_{21} \\ \vdots \\ a_{n1} \end{bmatrix} \bar{x}_1 + \begin{bmatrix} a_{12} \\ a_{22} \\ \vdots \\ a_{n2} \end{bmatrix} \bar{x}_2 + \cdots + \begin{bmatrix} a_{1m} \\ a_{2m} \\ \vdots \\ a_{nm} \end{bmatrix} \bar{x}_m = \begin{bmatrix} b_1 \\ b_2 \\ \vdots \\ b_m \end{bmatrix} \quad (7.15)$$

Exchanging the two sides of the equal sign and moving the scalars \bar{x}_i to the front of the vectors give

$$\begin{bmatrix} b_1 \\ b_2 \\ \vdots \\ b_m \end{bmatrix} = \bar{x}_1 \begin{bmatrix} a_{11} \\ a_{21} \\ \vdots \\ a_{n1} \end{bmatrix} + \bar{x}_2 \begin{bmatrix} a_{12} \\ a_{22} \\ \vdots \\ a_{n2} \end{bmatrix} + \cdots + \bar{x}_m \begin{bmatrix} a_{1m} \\ a_{2m} \\ \vdots \\ a_{nm} \end{bmatrix}; \quad (7.16)$$

in other words, b is a linear combination of the columns of A . The other way around works as well: if we manage to write

$$b = c_1 a_1^{\text{col}} + c_2 a_2^{\text{col}} + \cdots + c_m a_m^{\text{col}}$$

for some real numbers $c_i \in \mathbb{R}$, then $\bar{x} = [c_1 \ c_2 \ \cdots \ c_m]^\top$ satisfies $A\bar{x} = b$, and hence it is a solution to $Ax = b$.

Recall that $a_j^{\text{col}} := [a_{1j} \ a_{2j} \ \cdots \ a_{nj}]^\top$, $1 \leq j \leq m$.

A weakness of the above result is that we do not yet have an effective means for checking whether or not a given vector is a linear combination of a specified set of vectors. We will solve this problem too, a little bit later in the Chapter. All we have done so far is indicate that the concept of a linear combination is related to the existence of a solution to a corresponding system of linear equations. **The result is mainly conceptual in that we have not yet provided a practical way to check this in code. That will come, though!**
Long Live the LU Factorization!

7.5 Linear Independence of a Set of Vectors

7.5.1 Preamble

The concept of *linear independence*, or its logical opposite, the concept of *linear dependence*, is one of the most important ideas in all of linear algebra. In the beginning, most of us struggle with two things: (1) why is the property of linear independence so important; and (2), how do you test for it? In this section, we highlight the importance of linear independence of vectors by relating it to uniqueness of solutions to $\mathbf{Ax} = \mathbf{0}$ and eventually, to uniqueness of solutions to $\mathbf{Ax} = \mathbf{b}$. We'll also show that our friend, the LU Factorization, makes testing for linear independence very straightforward.

7.5.2 Linear Independence through the Lens of $\mathbf{Ax}=\mathbf{0}$

Let A be an $n \times m$ matrix. We say that $x \in \mathbb{R}^m$ is a **nontrivial solution** to $Ax = 0_{n \times 1}$ if

- $Ax = 0_{n \times 1}$ (x is a solution), and
- $x \neq 0_{m \times 1}$ (x is not the zero vector in \mathbb{R}^m).

Because $x = 0$ is always a solution to $Ax = 0$, for it to be the **unique solution**, there must not be a non-trivial solution to $Ax = 0$.

Why are we making such a big deal about this? Because uniqueness is a desirable property and we want to learn how to check for it even in the case of rectangular systems of equations². Our immediate goal is to see what this means in terms of the columns of A , just as we did when introducing the notion of a linear combination.

To see what it means to have “non-trivial solutions” to $Ax = 0_{n \times 1}$, once again, we replace $x \in \mathbb{R}^m$ by a vector $\alpha \in \mathbb{R}^m$ and write

$$A = \begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1m} \\ a_{21} & a_{22} & \cdots & a_{2m} \\ \vdots & \vdots & \ddots & \vdots \\ a_{n1} & a_{n2} & \cdots & a_{nm} \end{bmatrix} \quad \text{and} \quad \alpha = \begin{bmatrix} \alpha_1 \\ \alpha_2 \\ \vdots \\ \alpha_m \end{bmatrix}. \quad (7.17)$$

Then, using our column times row method for matrix multiplication, and re-arranging terms a bit as we did in Chapter 7.3.1,

$$\begin{aligned} A\alpha &= \begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1m} \\ a_{21} & a_{22} & \cdots & a_{2m} \\ \vdots & \vdots & \ddots & \vdots \\ a_{n1} & a_{n2} & \cdots & a_{nm} \end{bmatrix} \begin{bmatrix} \alpha_1 \\ \alpha_2 \\ \vdots \\ \alpha_m \end{bmatrix} \\ &= \alpha_1 \begin{bmatrix} a_{11} \\ a_{21} \\ \vdots \\ a_{n1} \end{bmatrix} + \alpha_2 \begin{bmatrix} a_{12} \\ a_{22} \\ \vdots \\ a_{n2} \end{bmatrix} + \cdots + \alpha_m \begin{bmatrix} a_{1m} \\ a_{2m} \\ \vdots \\ a_{nm} \end{bmatrix}. \end{aligned} \quad (7.18)$$

Hence, $A\alpha = 0_{n \times 1}$ if, and only if,

$$\alpha_1 \begin{bmatrix} a_{11} \\ a_{21} \\ \vdots \\ a_{n1} \end{bmatrix} + \alpha_2 \begin{bmatrix} a_{12} \\ a_{22} \\ \vdots \\ a_{n2} \end{bmatrix} + \cdots + \alpha_m \begin{bmatrix} a_{1m} \\ a_{2m} \\ \vdots \\ a_{nm} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}_{n \times 1}. \quad (7.19)$$

Our takeaway is, $\alpha = 0_{m \times 1}$ is the unique solution to $A\alpha = 0$ if, and only if, the only way we can add up the columns of A and obtain the zero vector in \mathbb{R}^n is with

$$\alpha_1 = 0, \alpha_2 = 0, \dots, \alpha_m = 0.$$

Or equivalently, $A\alpha = 0$ has a non-trivial solution $\alpha \neq 0_{m \times 1}$, if, and only if, there exists real numbers $\alpha_1, \alpha_2, \dots, \alpha_m$, **NOT ALL ZERO**, such that (7.19) is satisfied.

²Recall, when A is square, uniqueness is equivalent to $\det(A) \neq 0$.

7.5.3 Linear Independence in \mathbb{R}^n (and why theory matters)

The definition of linear independence of an “abstract” set of vectors in \mathbb{R}^n is 100% motivated by the study above on the uniqueness of solutions to $Ax = 0$. Of course, vectors in \mathbb{R}^n are columns of $n \times m$ matrices, so who’s to say what is abstract and what is not!

Linear Independence of a Set of Vectors

The set of vectors $\{v_1, v_2, \dots, v_m\} \subset \mathbb{R}^n$ is **linearly dependent** if there exist real numbers $\alpha_1, \alpha_2, \dots, \alpha_m$ **NOT ALL ZERO** yielding a linear combination of vectors that adds up to the zero vector,

$$\alpha_1 v_1 + \alpha_2 v_2 + \dots + \alpha_m v_m = 0_{n \times 1}. \quad (7.20)$$

On the other hand, the vectors $\{v_1, v_2, \dots, v_m\}$ are **linearly independent** if the **only** real numbers $\alpha_1, \alpha_2, \dots, \alpha_m$ yielding a linear combination of vectors that adds up to the zero vector,

$$\alpha_1 v_1 + \alpha_2 v_2 + \dots + \alpha_m v_m = 0_{n \times 1}, \quad (7.21)$$

are $\alpha_1 = 0, \alpha_2 = 0, \dots, \alpha_m = 0$.

Concise Definition of Linear Independence:

$$\alpha_1 v_1 + \alpha_2 v_2 + \dots + \alpha_m v_m = 0_{n \times 1} \iff \begin{bmatrix} \alpha_1 \\ \alpha_2 \\ \vdots \\ \alpha_m \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix} = 0_{m \times 1}.$$

Example 7.5 By applying the definition, determine if the set of vectors

$$v_1 = \begin{bmatrix} \sqrt{2} \\ 0 \\ 0 \end{bmatrix}, v_2 = \begin{bmatrix} 4 \\ 7 \\ 0 \end{bmatrix}, v_3 = \begin{bmatrix} 3 \\ 1 \\ -1 \end{bmatrix}$$

is linearly independent or dependent.

Solution: We form the linear combination and do the indicated multiplications and additions

$$\alpha_1 v_1 + \alpha_2 v_2 + \alpha_3 v_3 = \alpha_1 \begin{bmatrix} \sqrt{2} \\ 0 \\ 0 \end{bmatrix} + \alpha_2 \begin{bmatrix} 4 \\ 7 \\ 0 \end{bmatrix} + \alpha_3 \begin{bmatrix} 3 \\ 1 \\ -1 \end{bmatrix} = \begin{bmatrix} \sqrt{2} \alpha_1 + 4 \alpha_2 + 3 \alpha_3 \\ 7 \alpha_2 + \alpha_3 \\ -1 \alpha_3 \end{bmatrix}. \quad (7.22)$$

Setting the right hand side of (7.22) to the zero vector yields

$$\begin{bmatrix} \sqrt{2} \alpha_1 + 4 \alpha_2 + 3 \alpha_3 \\ 7 \alpha_2 + \alpha_3 \\ -1 \alpha_3 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}. \quad (7.23)$$

This is one of our friendly triangular systems of linear equations which we can solve via back substitution. We see immediately that the only solution to the bottom equation is $\alpha_3 = 0$, the only solution to the middle equation is then $\alpha_2 = 0$, and finally, the only solution to the top equation is $\alpha_1 = 0$. Hence, the only solution to

$$\alpha_1 v_1 + \alpha_2 v_2 + \alpha_3 v_3 = 0$$

is $\alpha_1 = 0, \alpha_2 = 0, \alpha_3 = 0$, and hence the set of vectors $\{v_1, v_2, v_3\}$ is linearly independent.

The above analysis becomes even more clear when we write (7.23) as

$$\begin{bmatrix} \sqrt{2} & 4 & 3 \\ 0 & 7 & 1 \\ 0 & 0 & -1 \end{bmatrix} \begin{bmatrix} \alpha_1 \\ \alpha_2 \\ \alpha_3 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}. \quad (7.24)$$

We see that the matrix is square with a non-zero determinant, and hence we know it has a unique solution. ■

Example 7.6 By applying the definition, determine if the set of vectors

$$v_1 = \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix}, v_2 = \begin{bmatrix} 1 \\ -2 \\ -4 \end{bmatrix}$$

is linearly independent or dependent.

Solution: We seek to determine if there are non-zero coefficients α_1 and α_2 resulting in a linear combination that forms the zero vector in \mathbb{R}^3 ,

$$\alpha_1 v_1 + \alpha_2 v_2 = \alpha_1 \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix} + \alpha_2 \begin{bmatrix} 1 \\ -2 \\ -4 \end{bmatrix} = \begin{bmatrix} \alpha_1 + \alpha_2 \\ 2\alpha_1 - 2\alpha_2 \\ 3\alpha_1 - 4\alpha_2 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}. \quad (7.25)$$

We observe that we have three equations in two unknowns and no obvious triangular structure to help us! What do we do? Well, if we can take any two of the three equations and show that the only solution is the trivial solution, $\alpha_1 = 0$ and $\alpha_2 = 0$, then we are done, because the trivial solution will always satisfy the remaining equation. Please check this reasoning out for yourself.

We arbitrarily group the equations into the first two equations and then the last equation, as follows

$$\alpha_1 \begin{bmatrix} 1 \\ 2 \end{bmatrix} + \alpha_2 \begin{bmatrix} 1 \\ -2 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix} \quad (7.26)$$

$$\alpha_1 \begin{bmatrix} 3 \end{bmatrix} + \alpha_2 \begin{bmatrix} -4 \end{bmatrix} = 0. \quad (7.27)$$

We can rewrite (7.26) in the form $A\alpha = b$

$$\underbrace{\begin{bmatrix} 1 & 1 \\ 2 & -2 \end{bmatrix}}_A \underbrace{\begin{bmatrix} \alpha_1 \\ \alpha_2 \end{bmatrix}}_\alpha = \underbrace{\begin{bmatrix} 0 \\ 0 \end{bmatrix}}_b. \quad (7.28)$$

We note that $\det(A) = (1)(-2) - (1)(2) = -4 \neq 0$, and hence (7.28) has a unique solution. Because $\alpha_1 = 0$ and $\alpha_2 = 0$ is a solution and the solution is unique, we know that there cannot exist a different set of non-zero α_1 and α_2 that also solve the equation. We, therefore, conclude that the vectors $\{v_1, v_2\}$ are linearly independent.

Remark: Let's step back and see what is going on here. We have three equations and two unknowns. The trivial solution is always a solution to the full set of equations. We wonder if there is any other solution. If we make a choice of two of the equations, so that we have a more manageable square system of equations, and then find that those two equations constrain the solution to being the trivial solution, we are done! (Why, because the trivial solution will automatically satisfy the remaining equation and it is then the only solution that satisfies all three equations.) ■

That was a lot of work! It does not seem like it will scale to bigger sets of vectors. We'll do one more example to convince you that we are in dire need of a **Pro Tip!**

Example 7.7 By applying the definition, determine if the vectors

$$v_1 = \begin{bmatrix} 1 \\ 2 \\ 3 \\ 1 \end{bmatrix}, v_2 = \begin{bmatrix} 0 \\ -2 \\ 4 \\ 5 \end{bmatrix}, v_3 = \begin{bmatrix} 2 \\ 6 \\ 2 \\ -3 \end{bmatrix}$$

are linearly independent or dependent.

Solution: We form the linear combination and set it equal to zero,

$$\alpha_1 v_1 + \alpha_2 v_2 + \alpha_3 v_3 = \alpha_1 \begin{bmatrix} 1 \\ 2 \\ 3 \\ 1 \end{bmatrix} + \alpha_2 \begin{bmatrix} 0 \\ -2 \\ 4 \\ 5 \end{bmatrix} + \alpha_3 \begin{bmatrix} 2 \\ 6 \\ 2 \\ -3 \end{bmatrix} = \begin{bmatrix} \alpha_1 + 2\alpha_3 \\ 2\alpha_1 - 2\alpha_2 + 6\alpha_3 \\ 3\alpha_1 + 4\alpha_2 + 2\alpha_3 \\ \alpha_1 + 5\alpha_2 - 3\alpha_3 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}. \quad (7.29)$$

We must check whether or not there are *non-trivial* solutions to (7.29), where non-trivial means at least one of the coefficients α_1 , α_2 , or α_3 is non-zero.

We observe that we have four equations in three unknowns. Similar to Example 7.6, we could select any three of the four equations, solve those three, and then see if that solution is compatible with the remaining equation. Once again, this will be a lot of work. We'll save ourselves the effort and let you verify that $\alpha_1 = 2$, $\alpha_2 = -1$, $\alpha_3 = -1$ is a non-trivial solution to (7.29) and hence the vectors v_1, v_2, v_3 are linearly dependent. **Where is that Pro Tip?** ■

7.5.4 A Pro Tip for Checking Linear Independence

Pro-tip! Linear Independence in a Nutshell

Consider the vectors in \mathbb{R}^n ,

$$\left\{ v_1 = \begin{bmatrix} a_{11} \\ a_{21} \\ \vdots \\ a_{n1} \end{bmatrix}, v_2 = \begin{bmatrix} a_{12} \\ a_{22} \\ \vdots \\ a_{n2} \end{bmatrix}, \dots, v_m = \begin{bmatrix} a_{1m} \\ a_{2m} \\ \vdots \\ a_{nm} \end{bmatrix} \right\},$$

and use them as the columns of a matrix that we call A ,

$$A = \begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1m} \\ a_{21} & a_{22} & \cdots & a_{2m} \\ \vdots & \vdots & \ddots & \vdots \\ a_{n1} & a_{n2} & \cdots & a_{nm} \end{bmatrix}. \quad (7.30)$$

The following statements are equivalent:

- The set of vectors $\{v_1, v_2, \dots, v_m\}$ is linearly independent.
- The $m \times m$ matrix $A^\top \cdot A$ is invertible.
- $\det(A^\top \cdot A) \neq 0$.
- For any LU Factorization $P \cdot (A^\top \cdot A) = L \cdot U$ of $A^\top A$, the $m \times m$ upper triangular matrix U has no zeros on its diagonal.

We'll prove the Pro Tip shortly. For now, we'll focus on how it makes our lives so much easier in terms of checking linear independence.

Example 7.8 We apply the Pro Tip to Example 7.5

Solution: We use the vectors $v_1 = \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix}$ and $v_2 = \begin{bmatrix} 1 \\ -2 \\ -4 \end{bmatrix}$ as the columns of a matrix $A := [v_1 \ v_2]$, so that

$$A = \begin{bmatrix} 1 & 1 \\ 2 & -2 \\ 3 & -4 \end{bmatrix}.$$

We compute that

$$A^\top \cdot A = \begin{bmatrix} 14.0 & -15.0 \\ -15.0 & 21.0 \end{bmatrix}.$$

Because it is 2×2 , we can compute its determinant easily and obtain

$$\det(A^\top \cdot A) = 69.0 \neq 0,$$

and hence the vectors $\{v_1, v_2\}$ are linearly independent. ■

Example 7.9 We apply the Pro Tip to Example 7.7.

Solution: We use the vectors

$$v_1 = \begin{bmatrix} 1 \\ 2 \\ 3 \\ 1 \end{bmatrix}, v_2 = \begin{bmatrix} 0 \\ -2 \\ 4 \\ 5 \end{bmatrix}, v_3 = \begin{bmatrix} 2 \\ 6 \\ 2 \\ -3 \end{bmatrix}$$

and form the matrix

$$A := \begin{bmatrix} 1 & 0 & 2 \\ 2 & -2 & 6 \\ 3 & 4 & 2 \\ 1 & 5 & -3 \end{bmatrix}.$$

We go to Julia and compute that

$$A^\top \cdot A = \begin{bmatrix} 15.0 & 13.0 & 17.0 \\ 13.0 & 45.0 & -19.0 \\ 17.0 & -19.0 & 53.0 \end{bmatrix},$$

and that its LU Factorization is $P \cdot (A^\top \cdot A) = L \cdot U$, where

$$P = \begin{bmatrix} 0.0 & 0.0 & 1.0 \\ 0.0 & 1.0 & 0.0 \\ 1.0 & 0.0 & 0.0 \end{bmatrix}, L = \begin{bmatrix} 1.0 & 0.0 & 0.0 \\ 0.8 & 1.0 & 0.0 \\ 0.9 & 0.5 & 1.0 \end{bmatrix}, \text{ and } U = \begin{bmatrix} 17.0 & -19.0 & 53.0 \\ 0.0 & 59.5 & -59.5 \\ 0.0 & 0.0 & \boxed{0.0} \end{bmatrix}.$$

We observe that U has a zero on its diagonal and hence the set $\{v_1, v_2, v_3\}$ is linearly dependent. ■

Yeah, that Pro Tip on Linear Independence is certainly worth using! Can something like it be used for testing whether a given vector is (or is not) a linear combination of another set of vectors? The answer is YES! And we'll get to that after we show why our current Pro Tip works.

What if you really want to know a set of non-trivial coefficients such that $A\alpha = 0$?

Instead of solving $A\alpha = 0$, you can solve the triangular system of equations, $U\alpha = 0$. It is emphasized that we only do the additional work of solving $U \cdot \alpha = 0$ if we want to find the explicit coefficients $\alpha_1, \alpha_2, \dots, \alpha_m$ such that

$$\alpha_1 v_1 + \alpha_2 v_2 + \dots + \alpha_m v_m = 0.$$

Many times we do not really need the coefficients. In HW and Quizzes, we'll tell you if you need to find the coefficients or whether a YES vs NO answer on linear independence is acceptable.

Example 7.10 Find a specific set of non-trivial coefficients for Example 7.7 that results in the linear combination equaling zero. For ease of the reader, this is the same as finding $A\alpha = 0_{4 \times 1}$ for

$$A := \begin{bmatrix} 1 & 0 & 2 \\ 2 & -2 & 6 \\ 3 & 4 & 2 \\ 1 & 5 & -3 \end{bmatrix}.$$

Solution: From the LU Factorization of A computed in Example 7.9, we have that

$$U = \begin{bmatrix} 17.0 & -19.0 & 53.0 \\ 0.0 & 59.5 & -59.5 \\ 0.0 & 0.0 & 0.0 \end{bmatrix}.$$

If we solve

$$\begin{bmatrix} 17.0 & -19.0 \\ 0.0 & 59.5 \end{bmatrix} \begin{bmatrix} \alpha_1 \\ \alpha_2 \end{bmatrix} = \begin{bmatrix} -53.0 \\ 59.5 \end{bmatrix} \iff \begin{bmatrix} 17.0 & -19.0 \\ 0.0 & 59.5 \end{bmatrix} \begin{bmatrix} \alpha_1 \\ \alpha_2 \end{bmatrix} - \begin{bmatrix} -53.0 \\ 59.5 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix},$$

we obtain

$$\begin{bmatrix} \alpha_1 \\ \alpha_2 \end{bmatrix} = \begin{bmatrix} -2.0 \\ 1.0 \end{bmatrix}.$$

It follows that

$$A \begin{bmatrix} \alpha_1 \\ \alpha_2 \\ 1 \end{bmatrix} = L \cdot U \begin{bmatrix} \alpha_1 \\ \alpha_2 \\ 1 \end{bmatrix} = L \cdot \left(\begin{bmatrix} 17.0 & -19.0 \\ 0.0 & 59.5 \\ 0.0 & 0.0 \end{bmatrix} \begin{bmatrix} \alpha_1 \\ \alpha_2 \end{bmatrix} + \begin{bmatrix} 53.0 \\ -59.5 \\ 0.0 \end{bmatrix} \right) = L \cdot \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}.$$

An alternative perspective: Let's write U in what is called “block form”, namely

$$U = \begin{bmatrix} 17.0 & -19.0 & 53.0 \\ 0.0 & 59.5 & -59.5 \\ 0.0 & 0.0 & 0.0 \end{bmatrix} =: \begin{bmatrix} A & b \\ 0_{1 \times 2} & 0 \end{bmatrix},$$

where

$$A := \begin{bmatrix} 17.0 & -19.0 \\ 0.0 & 59.5 \end{bmatrix} \text{ and } b := \begin{bmatrix} 53.0 \\ -59.5 \end{bmatrix}.$$

We seek a non-zero vector such that $Uv = 0_{3 \times 1}$. Because of the last row of U being all zeros, we are motivated to take v having a special form, namely

$$\underbrace{\begin{bmatrix} 17.0 & -19.0 & 53.0 \\ 0.0 & 59.5 & -59.5 \\ 0.0 & 0.0 & 0.0 \end{bmatrix}}_U \underbrace{\begin{bmatrix} \alpha_1 \\ \alpha_2 \\ 1 \end{bmatrix}}_v = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix} \iff \begin{bmatrix} A & b \\ 0_{1 \times 2} & 0 \end{bmatrix} \begin{bmatrix} \alpha_1 \\ \alpha_2 \\ 1 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix} \iff A \begin{bmatrix} \alpha_1 \\ \alpha_2 \end{bmatrix} + b = \begin{bmatrix} 0 \\ 0 \end{bmatrix},$$

where in the last step we used the fact that the last row of U is all zeros. From this we obtain that

$$Uv = 0_{3 \times 1} \iff A \begin{bmatrix} \alpha_1 \\ \alpha_2 \end{bmatrix} = -b.$$

However, A is upper triangular with non-zero elements on its diagonal, and thus

$$\begin{bmatrix} \alpha_1 \\ \alpha_2 \end{bmatrix} = -A^{-1}b = \begin{bmatrix} -2.0 \\ 1.0 \end{bmatrix}.$$

We therefore have

$$v = \begin{bmatrix} -2.0 \\ 1.0 \\ 1.0 \end{bmatrix}.$$

■

Remark: In (7.37), we will introduce a refinement of the LU Factorization that makes it a snap to find solutions to $Ax = 0$. Right now, it is still a bit clunky. Example 10.13 illustrates an algorithmic method that works on large matrices.

7.5.5 (Optional Read): Why the Pro Tip Works

We first consider a vector in \mathbb{R}^n that we denote as

$$y = \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_n \end{bmatrix}.$$

We next note that

$$y^\top y = [y_1 \ y_2 \ \cdots \ y_n] \cdot \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_n \end{bmatrix} = (y_1)^2 + (y_2)^2 + \cdots + (y_n)^2.$$

From this, we deduce the useful fact that

$$y = 0_{n \times 1} = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix} \iff y^\top \cdot y = 0.$$

To determine linear independence or dependence, we are looking to exclude (or find) non-trivial solutions to $A\alpha = 0$, where

$$\alpha = \begin{bmatrix} \alpha_1 \\ \alpha_2 \\ \vdots \\ \alpha_m \end{bmatrix},$$

and $A = [v_1 \ v_2 \ \cdots \ v_m]$, the matrix formed from our set of vectors $\{v_1, v_2, \dots, v_m\}$. Motivated by this, we let

$$y = A\alpha.$$

We then note the following chain of implications³

$$(A\alpha = 0) \implies (A^\top \cdot A\alpha = 0) \implies (\alpha^\top A^\top \cdot A\alpha = 0) \implies ((A\alpha)^\top \cdot (A\alpha) = 0) \implies (A\alpha = 0),$$

where the last implication follows from $y = 0_{n \times 1} \iff y^\top y = 0$.

From logic, we know that when we have

$$(a) \implies (b) \implies (c) \implies (d) \implies (a),$$

a chain of implications that begins and ends with the same proposition, then we deduce that

$$(a) \iff (b) \iff (c) \iff (d).$$

In our case, we are only interested in $(a) \iff (b)$, that is,

$$\boxed{A\alpha = 0 \iff (A^\top A)\alpha = 0.} \tag{7.31}$$

We next note that the matrix $A^\top \cdot A$ is $m \times m$, because it is the product of $m \times n$ and $n \times m$ matrices, A^\top and A , respectively. Hence, the equation

$$(A^\top \cdot A)\alpha = 0 \tag{7.32}$$

has a unique solution if, and only if, $\det(A^\top \cdot A) \neq 0$.

Now, why are we done? If $\alpha = 0_{m \times 1}$ is the ONLY solution to (7.32), then it is also the only solution to $A\bar{\alpha} = 0$, and we deduce that the columns of A are linearly independent. If $\alpha = 0_{m \times 1}$ is not a unique solution to (7.32), then there exists a non-zero vector $\bar{\alpha} \in \mathbb{R}^m$ that is also a solution to (7.32), meaning that $(A^\top A)\bar{\alpha} = 0$. But we know from (7.31) that this also means that $\bar{\alpha} \neq 0$ is a solution of $A\bar{\alpha} = 0$, and hence the columns of A are linearly dependent. ■

Don't miss seeing the forest for the trees!

Don't let the last details of the proof distract you too much. The main steps of the Pro Tip are

- [rectangular system of equations] $A\alpha = 0 \iff A^\top \cdot A\alpha = 0$ [square system of equations].
- The square system of equations $A^\top \cdot A\alpha = 0$ has a unique solution of $\alpha = 0$, the 0-vector in \mathbb{R}^m , if, and only if, $\det(A^\top \cdot A) \neq 0$.
- Hence, $A\alpha = 0$ has a unique solution of $\alpha = 0$, the 0-vector in \mathbb{R}^m , if, and only if, $\det(A^\top \cdot A) \neq 0$.
- Our final result is, $A\alpha = 0$ has a unique solution of $\alpha = 0$, the 0-vector in \mathbb{R}^m , if, and only if, the columns of A are linearly independent, where the last implication uses the **definition of linear independence**.

³The second implication follows from the first by multiplying on the left by A^\top . The third implication follows from the second by multiplying on the left by α^\top . The fourth implication follows from the third by recognizing that $\alpha^\top A^\top = (A\alpha)^\top$. The last implication uses our fact that $y = 0 \iff y^\top \cdot y = 0$.

7.6 LDLT and the Number of Linearly Independent Vectors in a Set

Consider a finite set of vectors in \mathbb{R}^n , such as $\{v_1, \dots, v_m\}$. Is there an intelligent way to talk about the largest set of linearly independent vectors that we could build from a given set of vectors? That is, the number of vectors that would remain if we discarded the fewest vectors so that the resulting set of vectors is linearly independent?




Figure 7.3: Checking linear independence from left to right. You could also start from the right and go to the left, or you could start in the middle and proceed to the two ends. You just need to do an organized search of the vectors!

In fact there is. As illustrated in Fig. 7.3, we can start from left to right and ask, is the set $\{v_1\}$ linearly independent? If it is, keep v_1 and if not, discard it (meaning, in this case, v_1 was the zero vector). For the sake of argument, let's suppose that $v_1 \neq 0$ and hence we keep it. Next, we ask, is the set $\{v_1, v_2\}$ linearly independent? If not, then v_2 is a linear combination of v_1 and we discard it, otherwise, we keep it. For the sake of argument, let's suppose that v_2 is a linear combination of v_1 and hence we discard it. We next ask, is the set $\{v_1, v_3\}$ linearly independent? Let's say it is, and then we would ask if the set $\{v_1, v_3, v_4\}$ is linearly independent, etc. In the end, we have built the largest set of linearly independent vectors from the given set and we can ask, how many elements does

it contain?

Number of Linearly Independent Vectors in a Finite Set

The following statements are equivalent:

- One can divide the set of vectors in \mathbb{R}^n

$$v_1 = \begin{bmatrix} a_{11} \\ a_{21} \\ \vdots \\ a_{n1} \end{bmatrix}, v_2 = \begin{bmatrix} a_{12} \\ a_{22} \\ \vdots \\ a_{n2} \end{bmatrix}, \dots, v_m = \begin{bmatrix} a_{1m} \\ a_{2m} \\ \vdots \\ a_{nm} \end{bmatrix} \quad (7.33)$$

into k linearly independent vectors and $m - k$ vectors that are linearly dependent on them.

- The matrix

$$A = \begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1m} \\ a_{21} & a_{22} & \cdots & a_{2m} \\ \vdots & \vdots & \ddots & \vdots \\ a_{n1} & a_{n2} & \cdots & a_{nm} \end{bmatrix} \quad (7.34)$$

has k linearly independent columns and $m - k$ columns that are linearly dependent on them.

- Let $P \cdot (A^\top A) = L \cdot U$ be an LU Factorization of $A^\top A$. Then U has k linearly independent columns and $m - k$ dependent columns. Because U is triangular, as in Example 7.5, checking linear independence is much easier than for the original matrix A .

Example 7.11 How many columns of the matrix

$$A = \begin{bmatrix} -0.2 & -0.2 & -0.4 & 0.3 & 0.3 \\ 0.3 & 1.0 & -0.1 & -1.1 & -1.7 \\ 0.7 & -1.9 & 1.5 & -0.0 & -3.0 \\ 0.9 & -1.0 & -0.7 & 0.6 & -1.8 \\ -0.5 & 0.8 & -1.1 & -0.5 & -0.5 \\ -2.0 & -0.9 & -0.5 & 0.2 & 0.3 \\ -1.0 & 0.6 & 0.7 & -0.9 & 0.2 \end{bmatrix}_{7 \times 5} \quad (7.35)$$

are linearly independent? Doing this as shown in Fig. 7.3 would be painful.

Solution: We turn to Julia and perform the LU Factorization: $\mathbf{P} \cdot (\mathbf{A}^\top \cdot \mathbf{A}) = \mathbf{L} \cdot \mathbf{U}$, for which we only report the upper triangular matrix

$$U = \begin{bmatrix} 6.680 & -1.090 & 1.320 & 0.900 & -4.840 \\ 0.000 & 7.282 & -1.965 & -2.733 & 4.400 \\ 0.000 & 0.000 & 4.069 & -1.525 & -0.506 \\ 0.000 & 0.000 & 0.000 & 3.124 & 9.371 \\ 0.000 & 0.000 & 0.000 & 0.000 & 0.000 \end{bmatrix}_{5 \times 5} = [u_1 \ u_2 \ u_3 \ u_4 \ u_5]_{5 \times 5}, \quad (7.36)$$

where we have labeled its columns as $u_1 \dots u_5$. Working from left to right with the columns of U , we have that because $u_1 \neq 0_{5 \times 1}$, the set $\{u_1\}$ is linearly independent. We next check the set $\{u_1, u_2\}$. To emphasize the beauty of the triangular structure in U , we check if there exist non-trivial solutions to

$$\alpha_1 u_1 + \alpha_2 u_2 = \begin{bmatrix} 6.680 & -1.090 \\ 0.000 & 7.282 \\ 0.000 & 0.000 \\ 0.000 & 0.000 \\ 0.000 & 0.000 \end{bmatrix} \begin{bmatrix} \alpha_1 \\ \alpha_2 \end{bmatrix} = \begin{bmatrix} 0.0 \\ 0.0 \\ 0.0 \\ 0.0 \\ 0.0 \end{bmatrix}.$$

There answer is clearly no⁴, the unique solution is $\alpha_1 = 0$ and $\alpha_2 = 0$, and thus $\{u_1, u_2\}$ is linearly independent. The same reasoning works for $\{u_1, u_2, u_3\}$ and $\{u_1, u_2, u_3, u_4\}$. Indeed, we could have jumped straight to the set of vectors $\{u_1, u_2, u_3, u_4\}$ because

⁴There are three rows of zeros. Hence, starting from the second row, back substitution provides that the unique answer is zero

checking its linear independence comes down to looking for solutions to

$$\begin{bmatrix} \boxed{6.680} & -1.090 & 1.320 & 0.900 \\ 0.000 & \boxed{7.282} & -1.965 & -2.733 \\ 0.000 & 0.000 & \boxed{4.069} & -1.525 \\ 0.000 & 0.000 & 0.000 & \boxed{3.124} \\ 0.000 & 0.000 & 0.000 & 0.000 \end{bmatrix} \begin{bmatrix} \alpha_1 \\ \alpha_2 \\ \alpha_3 \\ \alpha_4 \end{bmatrix} = \begin{bmatrix} 0.0 \\ 0.0 \\ 0.0 \\ 0.0 \\ 0.0 \end{bmatrix}.$$

Ignoring the final row of zeros, we really have a square triangular system with a non-zero diagonal, hence the trivial solution $\alpha_1 = \alpha_2 = \alpha_3 = \alpha_4 = 0$ is, in fact, the unique solution, proving linear independence.

What about $\{u_1, u_2, u_3, u_4, u_5\}$? The answer is no, and to see this we note that

$$\begin{aligned} \alpha_1 u_1 + \alpha_2 u_2 + \alpha_3 u_3 + \alpha_4 u_4 + \alpha_5 u_5 &= 0_{5 \times 1} \\ &\Downarrow \\ \alpha_1 u_1 + \alpha_2 u_2 + \alpha_3 u_3 + \alpha_4 u_4 &= -\alpha_5 u_5 \end{aligned}$$

Writing down this equation yields

$$\begin{bmatrix} \boxed{6.680} & -1.090 & 1.320 & 0.900 \\ 0.000 & \boxed{7.282} & -1.965 & -2.733 \\ 0.000 & 0.000 & \boxed{4.069} & -1.525 \\ 0.000 & 0.000 & 0.000 & \boxed{3.124} \\ 0.000 & 0.000 & 0.000 & 0.000 \end{bmatrix} \begin{bmatrix} \alpha_1 \\ \alpha_2 \\ \alpha_3 \\ \alpha_4 \end{bmatrix} = - \begin{bmatrix} -4.840 \\ 4.400 \\ -0.506 \\ 9.371 \\ 0.000 \end{bmatrix} \alpha_5.$$

If we set $\alpha_5 = -1$, for example, and once again ignore the bottom row of zeros because they do not affect the solution of the equations, we can solve the resulting triangular system of equations for α_1 through α_4 , giving us a non-trivial solution to $\alpha_1 u_1 + \dots + \alpha_5 u_5 = 0$. Hence, $\{u_1, u_2, u_3, u_4, u_5\}$ is linearly dependent. ■

Yes, the triangular structure of U is very helpful, but it still requires a lot of work to check for solutions. Is there anything like our **Pro-Tip for linear independence that we can apply for counting the maximum number of linearly independent vectors?**

Uber Pro-Tip: Number of Linearly Independent Vectors via an Enhanced LU Factorization

Assume that a set of vectors in (7.33) have been stacked to form the columns of an $n \times m$ matrix A as in (7.34), or that the matrix A has been given to us directly. **Fact:** The matrix $A^\top \cdot A$ always has an **LDLT Factorization**

$$P \cdot A^\top \cdot A \cdot P^\top = L \cdot D \cdot L^\top, \quad (7.37)$$

where

- P is a (row) permutation matrix;
- P^\top , the transpose of P , permutes the columns of A ;
- L is uni-lower triangular and L^\top , the transpose of L , is therefore uni-upper triangular; and
- D is diagonal and has non-negative entries.

Moreover,

- **the number of linearly independent columns of A is equal to the number of non-zero entries on the diagonal of D ;** and, if we denote this number by k ,
- **then for the version of the LDLT given below, the first k -columns of $A \cdot P^\top$ are linearly independent, and the remaining $(m - k)$ -columns (if any) are linearly dependent on the first k columns.**
- Because the columns of $A \cdot P^\top$ are simply the columns of A permuted by P^\top (that is, re-ordered by the permutation matrix), **the first k -columns of $A \cdot P^\top$ provide a selection of linearly independent columns of A .**

The algorithm for LDLT Factorization is derived in Chap. 7.11, in case you are curious. You are not responsible for its derivation.

Remarks: (Optional Read) The LDLT factorization may look intimidating, but once you realize that $U := D \cdot L^\top$ is upper triangular, this is really a refined LU Factorization that is possible for matrices of the form $A^\top \cdot A$. The name **LDLT Factorization** comes from $L \cdot D \cdot L^\top$, where the last **T** stands for transpose. It is also called a **Cholesky Factorization**. Whatever you call it, it is simply our well known LU Factorization with U in the special form $U := D \cdot L^\top$, which is possible for special kinds of matrices of the form $A^\top \cdot A$.

Recalling that the matrix transpose of a product is the product **in reverse order** of the matrix transposes, we check that

$$\begin{aligned} (P \cdot A^\top \cdot A \cdot P^\top)^\top &= (P^\top)^\top \cdot (A)^\top \cdot (A^\top)^\top \cdot (P)^\top \\ &= P \cdot A^\top \cdot A \cdot P^\top, \end{aligned}$$

once one notes that $(A^\top)^\top = A$ and $(P^\top)^\top = P$. Hence $P \cdot A^\top \cdot A \cdot P^\top$ is symmetric. Without the P^\top on the right, the term $P \cdot A^\top \cdot A$ alone would not be symmetric in general. A similar computation shows that $L \cdot D \cdot L^\top$ is also symmetric. ■

Here is the Julia code for computing the LDLT Factorization:

```

1 # the LDLT factorization is a special LU Factorization for M=A'A
2 # P A' A P' = L D L', L unitriangular, D diagonal with non-negative entries, P
   permutation matrix
3 # the P' on the right does column permutations to maintain the symmetry of A'*A
4 #
5 function ldl(A::Array{<:Number, 2})
6     epsilon=1e-12
7     M=A'*A
8     n,m= size(A)
9     A_reduced = deepcopy(M)
10    L = Array{Float64, 2}(undef, m, 0)
11    Id=zeros(m, m) + I
12    P=deepcopy(Id)
13    # could make D a vector for efficiency
14    D=zeros(m, m)
15    for i = 1:m
16        # move the biggest entry to the pivot position
17        ii=argmax( diag(A_reduced[i:m, i:m]) );
18        mrow=ii[1]+(i-1)
19        if ~ (i==mrow)
20            # row permutation
21            P[[i, mrow], :] = P[[mrow, i], :];
22            # row and column permutation
23            A_reduced[[i, mrow], :] = A_reduced[[mrow, i], :];
24            A_reduced[:, [i, mrow]] = A_reduced[:, [mrow, i]];
25        end
26        if i>1
27            L[[i, mrow], :] = L[[mrow, i], :];
28        end
29        pivot=A_reduced[i, i]
30        if ~isapprox(pivot, 0, atol=epsilon)
31            D[i, i]=pivot
32            C=A_reduced[:, i]/pivot #normalize all entries by C[i]
33            L=[L C]
34            A_reduced=A_reduced-C*pivot*C'
35        else
36            # Remainder of factorization is trivial
37            L=[L Id[:, i:m]]
38            break

```

```

39         end
40     end
41     diagD=diag (D)
42 return L, P, D, diagD
43 end

```

Example 7.12 We revisit Example 7.11: how many columns of the matrix

$$A = \begin{bmatrix} -0.2 & -0.2 & -0.4 & 0.3 & 0.3 \\ 0.3 & 1.0 & -0.1 & -1.1 & -1.7 \\ 0.7 & -1.9 & 1.5 & -0.0 & -3.0 \\ 0.9 & -1.0 & -0.7 & 0.6 & -1.8 \\ -0.5 & 0.8 & -1.1 & -0.5 & -0.5 \\ -2.0 & -0.9 & -0.5 & 0.2 & 0.3 \\ -1.0 & 0.6 & 0.7 & -0.9 & 0.2 \end{bmatrix}_{7 \times 5} \quad (7.38)$$

are linearly independent?

Solution: We turn to Julia and perform the LDLT Factorization: $\mathbf{P} \cdot \mathbf{A}^\top \cdot \mathbf{A} \cdot \mathbf{P}^\top = \mathbf{L} \cdot \mathbf{D} \cdot \mathbf{L}^\top$, for which we report the diagonal of D

$$\text{diag}(D) = [15.6 \quad 5.2 \quad 4.4 \quad 2.3 \quad 0.000]_{1 \times 5} \quad (7.39)$$

Because there are four non-zero entries on the diagonal of D , we conclude that A has four linearly independent columns. ■

Example 7.13 How many columns of the matrix

$$A = \begin{bmatrix} -0.2 & -0.2 & -0.4 & 0.3 & 0.3 & -0.5 \\ 0.3 & 1.0 & -0.1 & -1.1 & -1.7 & 0.1 \\ 0.7 & -1.9 & 1.5 & -0.0 & -3.0 & 0.3 \\ 0.9 & -1.0 & -0.7 & 0.6 & -1.8 & -0.2 \\ -0.5 & 0.8 & -1.1 & -0.5 & -0.5 & -1.3 \\ -2.0 & -0.9 & -0.5 & 0.2 & 0.3 & -3.2 \\ -1.0 & 0.6 & 0.7 & -0.9 & 0.2 & -0.6 \end{bmatrix}_{7 \times 6} \quad (7.40)$$

are linearly independent? Which ones are they?

Solution: We turn to Julia and perform the LDLT Factorization: $\mathbf{P} \cdot \mathbf{A}^\top \cdot \mathbf{A} \cdot \mathbf{P}^\top = \mathbf{L} \cdot \mathbf{D} \cdot \mathbf{L}^\top$, for which we report the diagonal of D

$$\text{diag}(D) = [15.6 \quad 12.6 \quad 3.6 \quad 2.6 \quad 0.0 \quad 0.0]_{1 \times 6} \quad (7.41)$$

which has four non-zero elements. We conclude that A has four linearly independent columns.

We next compute

$$A \cdot P^\top = \begin{bmatrix} 0.3 & -0.5 & -0.4 & 0.3 & -0.2 & -0.2 \\ -1.7 & 0.1 & -0.1 & -1.1 & 0.3 & 1.0 \\ -3.0 & 0.3 & 1.5 & 0.0 & 0.7 & -1.9 \\ -1.8 & -0.2 & -0.7 & 0.6 & 0.9 & -1.0 \\ -0.5 & -1.3 & -1.1 & -0.5 & -0.5 & 0.8 \\ 0.3 & -3.2 & -0.5 & 0.2 & -2.0 & -0.9 \\ 0.2 & -0.6 & 0.7 & -0.9 & -1.0 & 0.6 \end{bmatrix} \quad (7.42)$$

and conclude that the first four columns of $A \cdot P^\top$ are linearly independent and the last two columns are linearly dependent on the first four.

Here, we can see that the four linearly independent columns correspond to columns 3, 4, 5, and 6 of A , while columns 1 and 2 of A can be written as a linear combination of columns 3, 4, 5, and 6. Indeed, one can show

$$\begin{bmatrix} 0.3 & -0.5 & -0.4 & 0.3 \\ -1.7 & 0.1 & -0.1 & -1.1 \\ -3.0 & 0.3 & 1.5 & 0.0 \\ -1.8 & -0.2 & -0.7 & 0.6 \\ -0.5 & -1.3 & -1.1 & -0.5 \\ 0.3 & -3.2 & -0.5 & 0.2 \\ 0.2 & -0.6 & 0.7 & -0.9 \end{bmatrix} \begin{bmatrix} -0.33 & 0.33 \\ 0.67 & 0.33 \\ -0.33 & -0.67 \\ 0.33 & -1.33 \end{bmatrix} = \begin{bmatrix} -0.2 & -0.2 \\ 0.3 & 1.0 \\ 0.7 & -1.9 \\ 0.9 & -1.0 \\ -0.5 & 0.8 \\ -2.0 & -0.9 \\ -1.0 & 0.6 \end{bmatrix}. \quad (7.43)$$

■

Identifying the Linearly Independent Vectors

Is there a simple way to look at the data from the LDLT Factorization of $A^T A$ and find a specific choice of columns of A to form a linearly independent set? **Yes!** Here is the permutation matrix

$$P = \begin{bmatrix} 0.0 & 0.0 & 0.0 & 0.0 & 1.0 & 0.0 \\ 0.0 & 0.0 & 0.0 & 0.0 & 0.0 & 1.0 \\ 0.0 & 0.0 & 1.0 & 0.0 & 0.0 & 0.0 \\ 0.0 & 0.0 & 0.0 & 1.0 & 0.0 & 0.0 \\ 1.0 & 0.0 & 0.0 & 0.0 & 0.0 & 0.0 \\ 0.0 & 1.0 & 0.0 & 0.0 & 0.0 & 0.0 \end{bmatrix} \quad (7.44)$$

for Example 7.13. We see that, starting with the first row of P , the ones are in columns $\{5, 6, 3, 4, 1, 2\}$, that is,

$$P_{15} = 1.0, P_{26} = 1.0, P_{33} = 1.0, P_{44} = 1.0, P_{51} = 1.0 \text{ and } P_{62} = 1.0.$$

From this information, we can conclude that if we denote the columns of A as $A = [A_1 \ A_2 \ A_3 \ A_4 \ A_5 \ A_6]$, then

$$A \cdot P^T = [A_5 \ A_6 \ A_3 \ A_4 \ A_1 \ A_2].$$

Since we know that $A \cdot P^T$ has four linearly independent columns, we conclude that columns

$$\{ A_5 \ A_6 \ A_3 \ A_4 \}$$

of A are linearly independent.

You have to admit, it's pretty cool to do the analysis in this manner.

Looking Ahead: We will later have a very nice name for the number of linearly independent vectors in a set of vectors: the **dimension** of $\text{span}\{v_1, v_2, \dots, v_m\}$. For now, we're just counting the number of linearly independent vectors.

7.7 Attractive Test for Linear Combinations

Pro Tip! Linear Combination or Not?

Fact: A vector $v_0 \in \mathbb{R}^n$ can be written as a linear combination of $\{v_1, \dots, v_m\} \subset \mathbb{R}^n$ if, and only if, the set $\{v_0, v_1, \dots, v_m\}$ has the same number of linearly independent vectors as $\{v_1, \dots, v_m\}$.

Applying this to determining if a linear system of equations $Ax = b$ has a solution, we first define $A_e := [A \ b]$ by appending b to the columns of A . Then we do the corresponding LDLT Factorizations

- $P \cdot (A^\top \cdot A) \cdot P^\top = L \cdot D \cdot L^\top$
- $P_e \cdot (A_e^\top \cdot A_e) \cdot P_e^\top = L_e \cdot D_e \cdot L_e^\top$.

Fact: $Ax = b$ has a solution if, and only if, D and D_e have the same number of non-zero entries on their diagonals.

Why? We know that $Ax = b$ has a solution if, and only if, b is a linear combination of the columns of A . That is, A and $A_e := [A \ b]$ must have the same number of linearly independent columns.

(Optional Remark): You get the same result by defining $A_e := [b \ A]$; in fact, you can insert b anywhere you want, even in the middle of the columns of A .

Example 7.14 We consider a rectangular system that is on the edge of what we would like to analyze by hand,

$$\underbrace{\begin{bmatrix} 3.5 & 1.0 & 5.0 \\ 5.0 & 2.0 & 6.0 \\ 6.5 & 3.0 & 7.0 \\ 8.0 & 4.0 & 8.0 \end{bmatrix}}_A \underbrace{\begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}}_x = \underbrace{\begin{bmatrix} 4.0 \\ 4.0 \\ 4.0 \\ 4.0 \end{bmatrix}}_b. \quad (7.45)$$

Does it have a solution?

Solution: We seek to determine if (7.45) will have a solution. We do the indicated LDLT Factorizations

$$P \cdot (A^\top A) \cdot P^\top = L \cdot D \cdot L^\top \quad \text{and} \quad P_e \cdot ([A \ b]^\top [A \ b]) \cdot P_e^\top = L_e \cdot D_e \cdot L_e^\top$$

and report the diagonals of D and D_e as row vectors

$$\text{diag}(D) = [174.0 \quad 1.8 \quad 0.0]_{1 \times 3} \quad (7.46)$$

$$\text{diag}(D_e) = [174.0 \quad 1.8 \quad 0.0 \quad 0.0]_{1 \times 4}. \quad (7.47)$$

Based on the above, we see that A and $[A \ b]$ have the same number of linearly independent columns. Hence, b is a linear combination of the columns of A and therefore (7.45) has a solution. In fact, one can compute a solution is

$$x = \begin{bmatrix} 0.0 \\ -1.0 \\ 1.0 \end{bmatrix}. \quad (7.48)$$

We now change the vector b to

$$\begin{bmatrix} 20.0 \\ 11.0 \\ 12.0 \\ 14.0 \end{bmatrix},$$

and repeat the analysis, giving

$$\text{diag}(D) = [174.0 \quad 1.8 \quad 0.0]_{1 \times 3} \quad (7.49)$$

$$\text{diag}(D_e) = [861.0 \quad 28.3 \quad 0.5 \quad 0.0]_{1 \times 4}. \quad (7.50)$$

This time, the analysis shows that $[A \ b]$ has one more linearly independent column than A , and hence b is linearly independent of the columns of A . We conclude, therefore, that the new system of linear equations does not have a solution. ■

7.8 Existence and Uniqueness of Solutions to $Ax=b$

In Chapter 7.4, we showed that $Ax = b$ has a solution if, and only if, b is a linear combination of the columns of A . Here, we will show that uniqueness of a solution follows from the columns of A being linearly independent. **When both of these conditions hold, we have existence and uniqueness of solutions.**

Let A be an $n \times m$ matrix, not necessarily square, and consider

$$\underbrace{\begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1m} \\ a_{21} & a_{22} & \cdots & a_{2m} \\ \vdots & \vdots & \ddots & \vdots \\ a_{n1} & a_{n2} & \cdots & a_{nm} \end{bmatrix}}_A \underbrace{\begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_m \end{bmatrix}}_x = \underbrace{\begin{bmatrix} b_1 \\ b_2 \\ \vdots \\ b_n \end{bmatrix}}_b. \quad (7.51)$$

Suppose that \bar{x} satisfies $A\bar{x} = b$, which is the same thing as saying \bar{x} is a solution of $Ax = b$. Suppose also that $\bar{\bar{x}}$ satisfies $A\bar{\bar{x}} = b$. We give conditions that guarantee $\bar{\bar{x}} = \bar{x}$.

If both \bar{x} and $\bar{\bar{x}}$ satisfy $Ax = b$, we have that

$$A\bar{\bar{x}} - A\bar{x} = A(\bar{\bar{x}} - \bar{x}) = b - b = 0.$$

We define $\alpha := \bar{\bar{x}} - \bar{x}$ and write out its components as

$$\alpha := \begin{bmatrix} \alpha_1 \\ \alpha_2 \\ \vdots \\ \alpha_m \end{bmatrix}$$

Expanding the expression $A\alpha = 0$ in terms of the columns of A and the components of the vector α gives

$$\alpha_1 \underbrace{\begin{bmatrix} a_{11} \\ a_{21} \\ \vdots \\ a_{n1} \end{bmatrix}}_{a_1^{\text{col}}} + \alpha_2 \underbrace{\begin{bmatrix} a_{12} \\ a_{22} \\ \vdots \\ a_{n2} \end{bmatrix}}_{a_2^{\text{col}}} + \cdots + \alpha_m \underbrace{\begin{bmatrix} a_{1m} \\ a_{2m} \\ \vdots \\ a_{nm} \end{bmatrix}}_{a_m^{\text{col}}} = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}. \quad (7.52)$$

We know that $\alpha_1 = 0, \alpha_2 = 0, \dots, \alpha_m = 0$ is the unique solution to

$$\alpha_1 a_1^{\text{col}} + \alpha_2 a_2^{\text{col}} + \cdots + \alpha_m a_m^{\text{col}} = 0,$$

if and only if, the vectors $\{a_1^{\text{col}}, a_2^{\text{col}}, \dots, a_m^{\text{col}}\}$ are linearly independent. Hence, if $Ax = b$ has a solution, it will be unique if, and only if, the columns of A are linearly independent.

Existence and Uniqueness of Solutions to $Ax = b$

The following two statements are equivalent

- (a) The system of linear equations $Ax = b$ has a solution, and, it is unique.
- (b) b is a linear combination of the columns of A , and, the columns of A are linearly independent.

Remark: If b is not a linear combination of the columns of A , then $Ax = b$ does not have a solution. If the columns of A are not linearly independent, then if $Ax = b$ has one solution, it also has an infinite number of solutions.

Remark: Please return to Chapter 1.2 and re-work the examples using this knowledge. Our new results do not even require the system of equations to be square, and moreover, thanks to our Pro and Uber-Pro Tips, our results are not limited to a small number of equations that are amenable to hand calculations!

Example 7.15 Let's analyze an example that most of us could not get right if done by hand! We consider $Ax = b$, where

$$A = \begin{bmatrix} -0.2 & -0.2 & -0.4 & 0.3 & 0.3 \\ 0.3 & 1.0 & -0.1 & -1.1 & -1.7 \\ 0.7 & -1.9 & 1.5 & -0.0 & -3.0 \\ 0.9 & -1.0 & -0.7 & 0.6 & -1.8 \\ -0.5 & 0.8 & -1.1 & -0.5 & -0.5 \\ -2.0 & -0.9 & -0.5 & 0.2 & 0.3 \\ -1.0 & 0.6 & 0.7 & -0.9 & 0.2 \end{bmatrix}_{7 \times 5} \quad \text{and} \quad b = \begin{bmatrix} -0.5 \\ 0.1 \\ 0.3 \\ -0.2 \\ -1.3 \\ -3.2 \\ -0.6 \end{bmatrix}_{7 \times 1} \quad (7.53)$$

Does it have a solution? If it does, is it unique?

Solution: We form $A^T A$ and compute in Julia its LDLT Factorization and report the diagonal of D as a row vector

$$\text{diag}(D) = [15.6 \quad 5.2 \quad 4.4 \quad 2.3 \quad 0.0]_{1 \times 5} \quad (7.54)$$

There is a single zero on the diagonal and four non-zero elements. Thus we know that exactly four of the five columns of A are linearly independent. Hence, if there does exist a solution to $Ax = b$, it will not be unique.

Next, we form $A_e := [A \ b]$ and compute in Julia the LDLT Factorization of $A_e^T A_e$. We report the diagonal of D_e written as a row vector

$$\text{diag}(D_e) = [15.6 \quad 12.6 \quad 3.6 \quad 2.6 \quad 0.0 \quad 0.0]_{1 \times 6} \quad (7.55)$$

and note that it also has four non-zero entries. We deduce that b is a linear combination of the columns of A . Hence, the system of linear equations (7.53) has a solution and it is not unique. Therefore, it has an infinite number of solutions!

We select a different right hand side for (7.53) and report the key result of its LDLT Factorization,

$$\tilde{b} = \begin{bmatrix} 0.6 \\ 0.4 \\ 0.9 \\ 1.0 \\ 0.8 \\ 0.8 \\ 0.9 \end{bmatrix} \implies \text{diag}(\tilde{D}_e) = [15.6 \quad 5.2 \quad 1.2 \quad 2.3 \quad 0.9 \quad 0.0]_{1 \times 6}. \quad (7.56)$$

This time $\text{diag}(\tilde{D}_e)$ has five non-zero entries, whereas $\text{diag}(D)$ had four non-zero entries. Hence, \tilde{b} is not a linear combination of the columns of A , and the system of equations, with this new right hand side, does not have a solution. ■

7.9 When are the Solutions to $Ax = b$ the same as the Solutions to $M \cdot Ax = M \cdot b$?

So far, we have seen two special cases of multiplying (both sides of) a set of equations on the LEFT by a matrix M of compatible size.

a) In Chapter 6.3, we noted that if A is invertible, then

$$Ax = b \iff \underbrace{A^{-1} \cdot Ax}_I = A^{-1}b \iff x = A^{-1}b.$$

b) In Chapter 5.9, for P a permutation matrix, we noted that

$$Ax = b \iff P \cdot Ax = Pb,$$

where we needed this result when applying the LU Factorization $P \cdot A = L \cdot U$ to solve $Ax = b$.

Here, we will generalize these two special cases to a more general situation of multiplying on the left by a size-compatible matrix M that can be **rectangular**.

We assume that A is $n \times m$ and b is $m \times 1$. Then for any $p \times n$ matrix M , the matrix multiplications $M \cdot A$ and Mb both make sense. We now ask the question, for what matrices M is it true that $Ax = b$ and $M \cdot Ax = Mb$ have the same solutions? It turns out that **linear independence** of the columns of M is sufficient for this to hold.

To see this, we define $y = Ax - b$ and note that $x \in \mathbb{R}^m$ is a solution to $Ax = b$ if, and only if, $y = 0_{n \times 1}$; we have simply said that $Ax = b \iff Ax - b = 0_{n \times 1}$. Similarly, we note that $x \in \mathbb{R}^m$ is a solution to $M \cdot Ax = M \cdot b$ if, and only if, $My = 0_{p \times 1}$, which is because $M \cdot Ax = M \cdot b \iff M \cdot (Ax - b) = 0_{p \times 1}$. So far, so good.

Surprisingly, the proof is done! Why? From our proof of the Pro-Tip in Chapter 7.5.5, we have that if the columns of M are linearly independent, then

$$y = 0_{n \times 1} \iff My = 0_{p \times 1}.$$

In other words, if the columns of M are linearly independent, then x is a solution of $Ax = b$ if, and only if, x is a solution of $M \cdot Ax = M \cdot b$.

What about the converse? (that is, the other direction?) If $Ax = b$ and $M \cdot Ax = M \cdot b$ have the same solutions, can we conclude that the columns of M are linearly independent? In general, that statement is false. Consider,

$$\underbrace{\begin{bmatrix} 1 & 0 \\ 1 & 1 \\ 1 & 0 \end{bmatrix}}_A \underbrace{\begin{bmatrix} x_1 \\ x_2 \end{bmatrix}}_b = \underbrace{\begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix}}_b \text{ and define } M := \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix} \text{ so that } \underbrace{\begin{bmatrix} 1 & 0 \\ 1 & 1 \\ 0 & 0 \end{bmatrix}}_{M \cdot A} \underbrace{\begin{bmatrix} x_1 \\ x_2 \end{bmatrix}}_{M \cdot b} = \underbrace{\begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}}_{M \cdot b}.$$

Then $x_1 = 1, x_2 = -1$ is the unique solution to both systems of linear equations, but the columns of M are not linearly independent because of the column of zeros. Finding a counterexample is how we go about showing something is false.

Remark 1 In Chapter 8, we will *deliberately* multiply both sides of an equation that has **no solutions** by a matrix with columns that are linearly **dependent** so as to **create a system of equations that has a unique solution**. It's kind of mind-blowing, we know! That's part of why Linear Algebra is so exciting! Like a good mystery novel, its plot has unexpected twists and turns.

7.10 (Optional Read): Why LDLT and not Simply LU?

This is intended for instructors! Why can't one simply use the LU Factorization to determine the number of linearly independent columns? Why must one use the LDLT Factorization? For those with a driving sense of curiosity or simply a high tolerance for pain, here is an example of a matrix A where the number of non-zero elements on the diagonal of U with a standard LU Factorization does not correspond to the number of linearly independent columns of A , where

$$P \cdot A^T \cdot A = L \cdot U.$$

Consider

$$A = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}.$$

A is symmetric, meaning that $A^\top = A$. Moreover, you can check that $A^\top \cdot A = A$. The following is a valid LU Factorization with row permutations

$$\underbrace{\begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}}_P \cdot \underbrace{\begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}}_{A^\top \cdot A} = \underbrace{\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}}_L \cdot \underbrace{\begin{bmatrix} \boxed{0} & 1 & 0 \\ 0 & \boxed{0} & 0 \\ 0 & 0 & \boxed{0} \end{bmatrix}}_U.$$

We see that U has one linearly independent column (the same as A), while $\text{diag}(U)$ has three zeros. Hence, the number of non-zero elements on the diagonal of U does not always correspond to the number of linearly independent columns of A and U . For the magicians, we note that the eigenvalue zero of U has algebraic multiplicity three and geometric multiplicity two. When we include the necessary column permutations to maintain the symmetry of $P \cdot A^\top \cdot A \cdot P^\top$ at each step of the factorization, we avoid this problem because eigenvalues of symmetric matrices always have equal geometric and algebraic multiplicities. This fact is way beyond the scope of ROB 101. ■

7.11 (Optional Read): LU Factorization for Symmetric Matrices

Recall that a matrix M is symmetric if $M^\top = M$. For such matrices, we seek to refine our way of computing the LU Factorization so that at each stage of the reduction, $\text{Temp}_{k+1} = \text{Temp}_k - C_k \cdot R_k$, we preserve the symmetry of the matrix. We will discover a condition that will allow us to write

$$P \cdot M \cdot P^\top = L \cdot D \cdot L^\top, \quad (7.57)$$

where P is a permutation matrix, L is uni-lower triangular, and D is a diagonal matrix. When we can accomplish such a factorization, it will follow that the number of linearly independent columns of M is precisely equal to the number of non-zero elements on the diagonal of D . A useful application of this result would then be counting the number of linearly independent columns in a rectangular matrix A by counting the number of linearly independent columns of $A^\top A$, which is always symmetric.

To build up our intuition, consider a 3×3 symmetric matrix

$$\text{Temp}_1 := M = \begin{bmatrix} a & b & c \\ b & d & e \\ c & e & f \end{bmatrix}, \quad (7.58)$$

which has six parameters. If $a \neq 0$, our first step in constructing the LU Factorization is to define

$$\tilde{C}_1 = \begin{bmatrix} a \\ b \\ c \end{bmatrix} \frac{1}{a} \text{ and } \tilde{R}_1 = [a \quad b \quad c], \quad (7.59)$$

where we are using a bar over the extracted columns and rows to distinguish them from a slightly different definition we will give shortly. The symmetry in \tilde{C}_1 and \tilde{R}_1 is clearly visible and it comes from the symmetry in M .

We can take better advantage of the symmetry in M if we define instead

$$C_1 = \begin{bmatrix} a \\ b \\ c \end{bmatrix} \frac{1}{a}, \quad D_{11} = a \text{ and } R_1 = C_1^\top. \quad (7.60)$$

It follows that

$$C_1 \cdot D_{11} \cdot C_1^\top = C_1 \cdot D_{11} \cdot R_1 = \tilde{C}_1 \cdot \tilde{R}_1. \quad (7.61)$$

A nice property of $C_1 \cdot D_{11} \cdot C_1^\top$ is that it is clearly symmetric. Also, if we subtract one symmetric matrix from another, the result is another symmetric matrix, so the symmetry property propagates. In our case, we'd have something like

$$\text{Temp}_2 := M - C_1 \cdot D_{11} \cdot C_1^\top = \begin{bmatrix} a & b & c \\ b & d & e \\ c & e & f \end{bmatrix} - \begin{bmatrix} a & b & c \\ b & \frac{b^2}{a} & \frac{bc}{a} \\ c & \frac{bc}{a} & \frac{c^2}{a} \end{bmatrix} = \begin{bmatrix} 0 & 0 & 0 \\ 0 & \alpha & \beta \\ 0 & \beta & \gamma \end{bmatrix}. \quad (7.62)$$

We let the reader fill in the values of α , β , and γ because all we care is that Temp_2 is symmetric and we can repeat the process.

When doing the LU Factorization, there were two other cases to consider. One is when there is an entire column of zeros, as in

$$\text{Temp}_1 := M = \begin{bmatrix} 0 & 0 & 0 \\ 0 & d & e \\ 0 & e & f \end{bmatrix}, \quad (7.63)$$

so that $a = b = c = 0$. Because of the symmetry in Temp_1 , if we have an entire column of zeros, we also have an entire row of zeros! When we had entire column of zeros, we defined

$$\tilde{C}_1 = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} \text{ and } \tilde{R}_1 = [a \ b \ c] = [0 \ 0 \ 0], \quad (7.64)$$

We can take better advantage of the symmetry if we instead define

$$C_1 = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \quad D_{11} = 0 \text{ and } R_1 = C_1^\top. \quad (7.65)$$

It still follows that

$$C_1 \cdot D_{11} \cdot C_1^\top = C_1 \cdot D_{11} \cdot R_1 = \tilde{C}_1 \cdot \tilde{R}_1, \quad (7.66)$$

and $\text{Temp}_2 := M - C_1 \cdot D_{11} \cdot C_1^\top$ is once again symmetric, with the first column and first row identically zero, so the process continues. So far, no permutation matrices have shown up!

A permutation of rows is required when our matrix has a zero in the entry that is supposed to define the pivot, such as here,

$$\text{Temp}_1 := M = \begin{bmatrix} 0 & b & c \\ b & d & e \\ c & e & f \end{bmatrix}, \quad (7.67)$$

where $a = 0$. If $b \neq 0$, we would swap the first and second rows to arrive at

$$P \cdot \text{Temp}_1 := \begin{bmatrix} b & d & e \\ 0 & b & c \\ c & e & f \end{bmatrix}, \quad (7.68)$$

where to swap the first two rows, we have

$$P = \begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}. \quad (7.69)$$

But, in (7.68), we have destroyed the symmetry of our matrix!

Destroyed Symmetry!

A key observation is that we can **RESTORE SYMMETRY** if we also swap the first and second columns, like this,

$$\begin{bmatrix} d & b & e \\ b & 0 & c \\ e & c & f \end{bmatrix}, \quad (7.70)$$

but oops, we no longer have $b \neq 0$ at the pivot position. We've moved d , an element of the diagonal, to the pivot position. If we had instead swapped rows one and three, followed by a swap of columns one and three, we'd end up with f , another element of the diagonal, at the pivot position! Hence, if preserving symmetry is our goal, we need to look at the diagonal of Temp_k for a non-zero element, and then to a double swap, two rows and two columns, to move it to the pivot position and continue the algorithm. The swapping of the columns is achieved by multiplying Temp_k on the right (instead of the left, which we do for row swapping), and the required permutation matrix can be shown to be the transpose of the matrix used to do the row swapping! Go to Julia and play round with it. It'll be fun!

For a general symmetric matrix, if we arrive at a step where $\text{diag}(\text{Temp}_k)$ is all zeros, then the algorithm fails and we cannot factor M as in (7.57). A matrix that will lead to immediate failure is

$$M = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}. \quad (7.71)$$

However, if $M = A^\top \cdot A$ for some matrix A , then M is a *positive semi-definite matrix*. For such matrices,

$$\text{diag}(\text{Temp}_k) = 0 \iff \text{Temp}_k = 0,$$

in which case the algorithm can be completed as we did for the LDLT Factorization presented earlier.

We give the resulting algorithm and then revisit a few of our previous examples. So that a novice user can play around and check whether or not $Q = P^\top$ and $U = D \cdot L^\top$, we compute the additional matrices. You are free to remove them.

```
1 # LU Factorization for Symmetric Matrices
2 # P M P' = L D L', L unitriangular, D diagonal, P permutation matrix
3 #
4 function luJWGSymD(M::Array{<:Number, 2})
5     epsilon=1e-12; K=100
6     if isapprox(norm(M-M'), 0, atol=epsilon)
7         # M is indeed symmetric
8         n, m = size(M)
9         A=deepcopy(M)
10        L = Array{Float64, 2}(undef, n, 0)
11        # Could remove U for efficiency
12        U = Array{Float64, 2}(undef, 0, n)
13        P=zeros(n, n) + I
14        # Could remove Q for efficiency
15        Q=zeros(n, n) + I
16        # could make D a vector for efficiency
17        D=zeros(n, n)
18        for i = 1:n
19            C = A[:, i] # i-th column
20            R = C' # i-th row
21            if maximum(abs.(C)) <= K*epsilon #column of zeros
22                C=0.0*C; C[i]=1.0; R = 0.0*R
23                U=[U; R];
24                L=[L C];
```

```

25     D[i,i]=0.0
26     Areduced=Areduced-C*R;
27 else # move the biggest entry to the pivot position
28     ii=argmax( abs.(diag(Areduced)) );
29     nrow=ii[1]
30     P[[i,nrow],:] = P[[nrow,i],:];
31     Q[:,[i,nrow]] = Q[:,[nrow,i]];
32     Areduced[[i,nrow],:] = Areduced[[nrow,i],:];
33     Areduced[:,[i,nrow]] = Areduced[:,[nrow,i]];
34     if i>1
35         L[[i,nrow],:] = L[[nrow,i],:];
36         U[:,[i,nrow]] = U[:,[nrow,i]];
37         d1=D[i,i]; d2=D[nrow,nrow]
38         D[i,i]=d2; D[nrow,nrow]=d1
39     end
40     C = Areduced[:,i] # i-th column
41     R = C' # i-th row
42     pivot = C[i];
43     if isapprox(pivot,0, atol=epsilon)
44         if isapprox(norm(Areduced),0, atol=K*epsilon)
45             # Remainder of factorization is trivial
46             L=[L Id[:,i:m]]
47             # U=D*L'; Q=P'; # All matrices included for pedagogical reasons
48             return L, U, P, Q, D
49             break
50         else
51             println("Algorithm failed at step $i")
52             println("A symmetric factorization PMP' = L D L' is not possible.")
53             return -1.0
54             break
55         end
56     else
57         D[i,i]=pivot
58         C=C/pivot #normalize all entires by C[i]
59         U=[U;R]; # could remove U for efficiency
60         L=[L C];
61         Areduced=Areduced-C*R;
62     end
63     Areduced[:,i]=0*Areduced[:,i]; Areduced[i,:]=0*Areduced[i,:];
64 end
65 end
66 # U=D*L'; Q=P'; # All matrices included for pedagogical reasons
67 return L, U, P, Q, D
68 else
69     println("Matrix is not symmetric")
70     return 0.0
71 end
72 end

```

Why does the LDLT Factorization work for matrices of the form $A^\top \cdot A$? The proof needs ideas from positive semi-definite matrices as presented in Appendix A.3. Here we'll sketch the proof assuming the material in the Appendix is known, which may help instructors using the book. We suppose that $M := \text{Temp}_k$ is positive semi-definite and we arrive at a step of LDLT where the diagonal of M consists only of zeros. We write the matrix as

$$M = \begin{bmatrix} 0 & M_{12} \\ M_{12}^\top & M_{22} \end{bmatrix},$$

and note that M_{22} must be positive semi-definite by the assumption on M . We claim that necessarily, $M_{12} = 0$, the zero vector. To

show this, we form the product

$$x^\top Mx = \begin{bmatrix} x_1 & x_2^\top \end{bmatrix} \begin{bmatrix} 0 & M_{12} \\ M_{12}^\top & M_{22} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = 2x_1 M_{12}x_2 + x_2^\top M_{22}x_2.$$

If $M_{12} \neq 0$, then there exists x_2 such that $M_{12}x_2 < 0$. It follows that by taking $x_1 > 0$ sufficiently large, that

$$2x_1 M_{12}x_2 + x_2^\top M_{22}x_2 < 0,$$

which contradicts M being positive semi-definite. This argument can be repeated. Applying it by induction to M_{22} , we arrive at M is the zero matrix, which can be easily factored. Otherwise, at each step, we have non-zero elements on the diagonal of $M := \text{Temp}_k$ and the LDLT Factorization can be continued.

7.12 Looking Ahead

In the next Chapter, we show how to measure the “length” of a vector and use that to formulate the problem of finding approximate solutions to linear systems of equations that do not have exact solutions. At first blush, this seems like a strange thing to do. In real engineering, however, the data we collect is never exact. The data is perturbed by various sources of errors, from imprecision in our instruments, to the fact that experiments are hard to repeat! To get around this, we take many more measurements than we need, which gives us sets of what we call “overdetermined equations,” which means more equations than unknown variables. We seek to “smooth” or “average” out the measurement errors by finding approximate solutions so that $e := Ax - b$, the error in the “equation,” is small. We measure the “length of the error vector” by a function called the “Euclidean norm.”

Chapter 8

Euclidean Norm, Least Squared Error Solutions to Linear Equations, and Linear Regression

Learning Objectives

- Learn one way to assign a notion of length to a vector
- The concept of finding approximate solutions to $Ax = b$ when an exact solution does not exist and why this is extremely useful in engineering.

Outcomes

- Euclidean norm and its properties
- If $Ax = b$ does not have a solution, then for any $x \in \mathbb{R}^n$, the vector $Ax - b$ is never zero. We will call $e := Ax - b$ the error vector and search for the value of x that minimizes the norm of the error vector.
- An application of this idea is Linear Regression, one of the “superpowers” of Linear Algebra: fitting functions to data.

8.1 Norm or “Length” of a Vector

Definition: The norm of a vector in \mathbb{R}^n

Let $v = \begin{bmatrix} v_1 \\ v_2 \\ \vdots \\ v_n \end{bmatrix}$ be a vector in \mathbb{R}^n . The **Euclidean norm** of v , denoted $\|v\|$, is defined as

$$\|v\| := \sqrt{(v_1)^2 + (v_2)^2 + \cdots + (v_n)^2} \quad (8.1)$$

Example: The length of the vector $v = \begin{bmatrix} \sqrt{2} \\ -1 \\ 5 \end{bmatrix}$ is

$$\|v\| := \sqrt{(\sqrt{2})^2 + (-1)^2 + (5)^2} = \sqrt{2 + 1 + 25} = \sqrt{28} = \sqrt{4 \cdot 7} = 2\sqrt{7} \approx 5.29.$$

Properties of the Norm of a vector

Let's get used to term **norm of a vector**, which is the correct mathematical term for the length of a vector. There are actually many different ways of defining a notion of “length” to a vector. The particular norm defined in Definition 4.14 is the “Euclidean norm.” All norms satisfy the following properties

- For all vectors $v \in \mathbb{R}^n$, $\|v\| \geq 0$ and moreover, $\|v\| = 0 \iff v = 0$.
- For any real number $\alpha \in \mathbb{R}$ and vector $v \in \mathbb{R}^n$,

$$\|\alpha v\| = |\alpha| \cdot \|v\|.$$

- For any pair of vectors v and w in \mathbb{R}^n ,

$$\|v + w\| \leq \|v\| + \|w\|.$$

The first property says that v has norm zero if, and only if, v is the zero vector. It seems pretty obvious for our notion of norm and it is!

For the second property, we note that we have to take the absolute value of the constant when we “factor it out” of the norm. This is because $\sqrt{a^2} = |a|$ and NOT a when $a < 0$. Of course, when $a \geq 0$, $\sqrt{a^2} = a$.

The third property is often called the **triangle inequality**. It says that the norm of a sum of vectors is upper bounded by the sum of the norms of the individual vectors. Another way to say this is, “the length of $v + w$ can never be strictly larger than the length of v plus the length of w .” What happens if you have $\|v - w\|$? Well,

$$\begin{aligned} \|v - w\| &= \|v + (-w)\| \\ &\leq \|v\| + \|-w\| \\ &= \|v\| + |-1| \cdot \|w\| \\ &= \|v\| + \|w\|. \end{aligned} \quad (8.2)$$

Hence, a minus sign changes nothing!

Remark: Equation 8.2 is correct with the “equals sign” in the last two equations because $\|v\| + \|-w\| = \|v\| + |-1| \cdot \|w\|$ and $\|v\| + |-1| \cdot \|w\| = \|v\| + \|w\|$. Some authors would carry the “less than or equal to symbol” all the way through. With our

notation, you know where the upper bounding took place!

Example We'll take three vectors in \mathbb{R}^4 and check the "triangle inequality."

$$u = \begin{bmatrix} 2 \\ -1 \\ 5 \\ 2 \end{bmatrix}, v = \begin{bmatrix} 1 \\ 2 \\ 0 \\ 3 \end{bmatrix}, w = \begin{bmatrix} 0 \\ 1 \\ 0 \\ 4 \end{bmatrix}$$

We first compute the norms of the three vectors

$$\|u\| = \sqrt{34} \approx 5.83, \|v\| = \sqrt{14} \approx 3.74, \|w\| = \sqrt{17} \approx 4.12$$

and we then form a few sums against which to test the triangle inequality

$$u + v = \begin{bmatrix} 3 \\ 1 \\ 5 \\ 5 \end{bmatrix}, u + v + w = \begin{bmatrix} 3 \\ 2 \\ 5 \\ 9 \end{bmatrix}, v + w = \begin{bmatrix} 1 \\ 3 \\ 0 \\ 7 \end{bmatrix}.$$

Then we can check that

$$\begin{aligned} \|u + v\| &= \sqrt{60} \approx 7.75 \leq 5.8 + 3.7 \leq \|u\| + \|v\| \\ \|u + v + w\| &= \sqrt{119} \approx 10.91 \leq 7.7 + 4.1 \leq \|u + v\| + \|w\| \\ \|u + v + w\| &= \sqrt{119} \approx 10.91 \leq 5.8 + 3.7 + 4.1 \leq \|u\| + \|v\| + \|w\| \end{aligned}$$

8.2 Least Squared Error Solutions to Linear Equations

In this section, we tackle the problem of providing a notion of "best approximate solution" to a system of linear equations that does not have an exact solution. If the system of equations does have an exact solution, our approximate solution will be identical to it. Hence, our development can be viewed as an alternative means of defining solutions to linear equations.

Consider a system of linear equations $Ax = b$ and define the vector

$$e(x) := Ax - b \tag{8.3}$$

as the **error in the solution** for a given value of x . Normally, we'll simply write $e := Ax - b$, but in (8.3), we are emphasizing that the error really is a function of x . Hence, as we vary x , the "length" of $e(x)$ can become bigger or smaller.

If the system of equations $Ax = b$ has a solution, then it is possible to make the error zero. In the previous section, we introduced the Euclidean norm as a means of measuring the "length" of a vector. It is traditional when posing the best approximation problem to use the square of the norm instead of the norm itself, which means we are simply removing the square root operation in the formula for a norm.

Least Squared Error Solution: Consider a linear system of equations expressed in matrix form $Ax = b$, where A is $n \times m$, x is $m \times 1$ and b is $n \times 1$. For a given value of $x \in \mathbb{R}^m$, define the error as in (8.3), an $n \times 1$ vector. The norm squared of the error vector $e(x)$ is then

$$\|e(x)\|^2 := \sum_{i=1}^n (e_i(x))^2 = e(x)^\top e(x) = (Ax - b)^\top (Ax - b) = \|Ax - b\|^2. \quad (8.4)$$

We note that $\|e(x)\|^2 \geq 0$ for any $x \in \mathbb{R}^m$ and hence zero is a lower bound on the norm squared of the error vector. A value $x^* \in \mathbb{R}^m$ is a **Least Squared Error Solution** to $Ax = b$ if it satisfies

$$\|e(x^*)\|^2 = \min_{x \in \mathbb{R}^m} \|Ax - b\|^2 \quad (8.5)$$

If such an $x^* \in \mathbb{R}^m$ exists and is unique, we will write it as

$$x^* := \arg \min_{x \in \mathbb{R}^m} \|Ax - b\|^2. \quad (8.6)$$

With this notation, the value of x that minimizes the error in the solution is what is returned by the function $\arg \min$, while the minimum value of the error is what is returned by the function \min ,

- $x^* = \arg \min_{x \in \mathbb{R}^m} \|Ax - b\|^2$ is the value of x that achieves the minimum value of the squared norm of the error, $\|Ax - b\|^2$, while
- $\|e(x^*)\|^2 = \|Ax^* - b\|^2 = \min_{x \in \mathbb{R}^m} \|Ax - b\|^2$ is the minimum value of the “squared approximation error”.

If $Ax = b$ has a solution, then

$$\min_{x \in \mathbb{R}^m} \|Ax - b\|^2 = 0,$$

because the norm squared error cannot be negative; indeed, zero is the smallest possible value. Hence, while we are most interested in cases where $Ax = b$ does not have a solution, if it does have a unique solution \bar{x} such that $A\bar{x} = b$, then

$$\bar{x} = x^* := \arg \min_{x \in \mathbb{R}^m} \|Ax - b\|^2.$$

We recall that uniqueness of solutions to $Ax = b$ is tied to the columns of A being linearly independent. Let’s also observe that if $Ax = b$, then

- multiplying both sides of $Ax = b$ by A^\top gives

$$A^\top \cdot Ax = A^\top \cdot b.$$

- The indicated matrix multiplications are well defined because A^\top is $m \times n$, A is $n \times m$, and b is $n \times 1$.
- $A^\top \cdot A$ is therefore square and $m \times m$.
- $A^\top \cdot A$ is symmetric because

$$(A^\top \cdot A)^\top = A^\top \cdot (A^\top)^\top = A^\top \cdot A$$

because $(A^\top)^\top = A$.

Least Squared Error Solutions to Linear Equations

Here are the main results on solutions to $Ax = b$ that minimize the squared error $\|Ax - b\|^2$.

(a) $A^T A$ is invertible if, and only if, the columns of A are linearly independent.

(b) If $A^T A$ is invertible, then there is a unique vector $x^* \in \mathbb{R}^m$ achieving $\min_{x \in \mathbb{R}^m} \|Ax - b\|^2$ and it satisfies the equation

$$(A^T A) x^* = A^T b. \tag{8.7}$$

(c) Therefore, if $A^T A$ is invertible,

$$x^* = (A^T A)^{-1} A^T b \iff x^* = \arg \min_{x \in \mathbb{R}^m} \|Ax - b\|^2 \iff (A^T A) x^* = A^T b. \tag{8.8}$$

As you might guess by now, your instructors prefer that for large systems of equations, you solve (8.7) to obtain the least squares solution and avoid doing the inverse. For small systems, we'll cut you some slack.

Useful Remark: Suppose that A is a “tall matrix” (more rows than columns) and suppose that $Ax = b$ has a solution (hence, b is a linear combination of the columns of A). Then, if the columns of A are linearly independent, you can compute the solution using (8.7) and the squared error will be zero, meaning x^* really is a solution to the equation because

$$\text{the squared error is zero} \iff \|Ax^* - b\|^2 = 0 \iff \|Ax^* - b\| = 0 \iff Ax^* - b = 0 \iff Ax^* = b.$$

Try it on your own in Julia!

Example 8.1 Let's consider a system of linear equations, with more equations than unknowns. The extra equations provide more conditions that a solution must satisfy, making non-existence of a solution a common occurrence! We take

$$\underbrace{\begin{bmatrix} 1.0 & 1.0 \\ 2.0 & 1.0 \\ 4.0 & 1.0 \\ 5.0 & 1.0 \\ 7.0 & 1.0 \end{bmatrix}}_A \underbrace{\begin{bmatrix} x_1 \\ x_2 \end{bmatrix}}_x = \underbrace{\begin{bmatrix} 4 \\ 8 \\ 10 \\ 12 \\ 18 \end{bmatrix}}_b, \tag{8.9}$$

and note that the columns of A are linearly independent. If a regular solution exists, find it. If not, then a least squared solution will be fine.

Solution: We'll compute a least squared error solution to the equations, and then we'll evaluate the error; if the error is zero, we'll also have an exact solution to (8.9). We compute

$$A^T \cdot A = \begin{bmatrix} 95.0 & 19.0 \\ 19.0 & 5.0 \end{bmatrix}, \quad A^T \cdot b = \begin{bmatrix} 246.0 \\ 52.0 \end{bmatrix} \implies x^* = \begin{bmatrix} 2.12 \\ 2.33 \end{bmatrix}$$

For the record, $\det(A^T \cdot A) = 114$. We also compute

$$e^* := Ax^* - b = \begin{bmatrix} 0.456 \\ -1.421 \\ 0.825 \\ 0.947 \\ -0.807 \end{bmatrix}, \quad \|e\| = 2.111, \quad \text{and } \|e\|^2 = 4.456$$

Because the “smallest” error vector $e^* = Ax^* - b$ is non-zero, we now know that (8.9) does not have an exact solution. Next, we take

a new vector $b = \begin{bmatrix} 0.30 \\ 1.00 \\ 2.40 \\ 3.10 \\ 4.50 \end{bmatrix}$ and check if there is an exact solution or not. We proceed in the same manner and solve $A^T \cdot Ax^* = A^T b$

and compute

$$x^* = \begin{bmatrix} 0.7000 \\ -0.4000 \end{bmatrix} \text{ and } e := Ax^* - b = \begin{bmatrix} 0.0000 \\ 0.0000 \\ 0.0000 \\ 0.0000 \\ 0.0000 \end{bmatrix},$$

and hence for this new vector b , the rectangular system $Ax = b$ has an exact solution. Moreover, we found it by applying our theory for least squared error solutions to overdetermined equations, which is pretty nifty! ■

The set of linear equations (8.9) and its solution look rather ho hum. They are actually anything but boring. Figure 8.1 shows that the equations and their solution correspond to fitting a line through data, while minimizing the “fitting error”! In the next section, we will develop this idea thoroughly and give you a hint of some of the things you can do with least squared error solutions to linear equations.

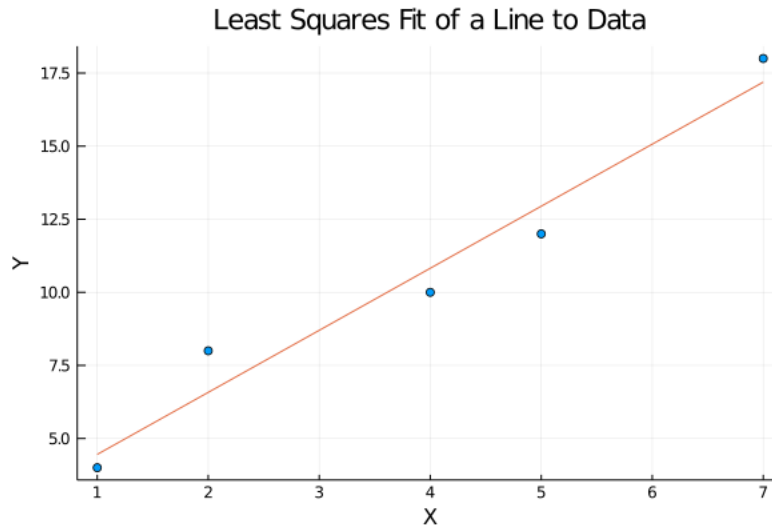


Figure 8.1: The physical basis for the numbers in (8.9), which is really about fitting a line to data with minimum squared error! The dots are the (x, y) values of the raw (measured) data, while the line is a model that summarizes the data in the form $y = mx + b$. One important aspect of the model is that you can now predict values of y for values of x that you did not directly measure. This is called interpolation when x is within the limits of the measured data (here, $1 \leq x \leq 7$) and extrapolation otherwise.

8.3 Linear Regression or Fitting Functions to Data

The goal of this section is explain how to fit functions to data. The following is a prototypical problem: given the data shown in Table 8.1, which is also plotted in Fig. 8.2, find a function that explains the data.

It is clear that the data do NOT lie exactly on a straight line. How can we **approximately** fit a straight line to the data? In particular, how can we find a function that minimizes a meaningful sense of fitting error to the data?

Table 8.1: Data for our first fitting problem.

i	x_i	y_i
1	1	4
2	2	8
3	4	10
4	5	12
5	7	18




Figure 8.2: Plot of data in Table 8.1

Let's suppose that we wish to fit a linear model $\hat{y} = mx + b$ to the data, which has been plotted in Fig. 8.2. We set up the linear equations

$$y_i = mx_i + b = [x_i \quad 1] \begin{bmatrix} m \\ b \end{bmatrix}, \quad 1 \leq i \leq N,$$

where N is the number of data points (five in the case of Table 8.1), and write it out in matrix form

$$\underbrace{\begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_N \end{bmatrix}}_Y = \underbrace{\begin{bmatrix} x_1 & 1 \\ x_2 & 1 \\ \vdots & \vdots \\ x_N & 1 \end{bmatrix}}_\Phi \cdot \underbrace{\begin{bmatrix} m \\ b \end{bmatrix}}_\alpha, \quad (8.10)$$

where Y is the vector of y -data, Φ is called the **regressor matrix** and α is the vector of **unknown coefficients** that parameterize the model.

From the data in Table 8.1, the matrices are

$$Y = \begin{bmatrix} 4 \\ 8 \\ 10 \\ 12 \\ 18 \end{bmatrix}, \quad \Phi = \begin{bmatrix} 1.0 & 1.0 \\ 2.0 & 1.0 \\ 4.0 & 1.0 \\ 5.0 & 1.0 \\ 7.0 & 1.0 \end{bmatrix}, \quad \text{and } \alpha = \begin{bmatrix} m \\ b \end{bmatrix}.$$

Y is the vector of “measured” y -values. α is the vector of unknown coefficients that we seek to estimate. Φ , the regressor matrix, is defined so that the i -th row of $Y = \Phi\alpha$ corresponds to $y_i = mx_i + b$.

The fitting error will be $e_i = y_i - (mx_i + b)$, which when written as a vector gives

$$\underbrace{\begin{bmatrix} e_1 \\ e_2 \\ e_3 \\ e_4 \\ e_5 \end{bmatrix}}_e = \underbrace{\begin{bmatrix} 4 \\ 8 \\ 10 \\ 12 \\ 18 \end{bmatrix}}_Y - \underbrace{\begin{bmatrix} 1.0 & 1.0 \\ 2.0 & 1.0 \\ 4.0 & 1.0 \\ 5.0 & 1.0 \\ 7.0 & 1.0 \end{bmatrix}}_\Phi \cdot \underbrace{\begin{bmatrix} m \\ b \end{bmatrix}}_\alpha,$$

that is, $e := Y - \Phi\alpha$. We propose to choose the coefficients in α so as to minimize the total squared error

$$E_{tot} = \sum_{i=1}^5 (e_i)^2 = e^\top e = \|e\|^2 = \|Y - \Phi\alpha\|^2.$$

Least Squares Fit to Data also called Linear Regression

From Chapter 8.2, if the columns of Φ are linearly independent, or equivalently, $\Phi^\top \Phi$ is invertible, then the following are equivalent

$$\alpha^* = (\Phi^\top \Phi)^{-1} \Phi^\top Y \iff \alpha^* = \arg \min_{\alpha} \|Y - \Phi \alpha\|^2 \iff (\Phi^\top \Phi) \alpha^* = \Phi^\top Y. \quad (8.11)$$

The associated equations are formulated and solved in the Julia code given below. The plot of the fit is given in Fig. 8.1. You can generate your own plots too!

```
1 using Plots
2 gr()
3
4 # Given data
5 X=[1 2 4 5 7]'
6 Y=[4 8 10 12 18]'
7
8 # Scatter plot
9 scatter(X, Y)
10 plot!(xlabel="X", ylabel="Y", title="dataSet1", leg=false)
11 plot!(fmt = :png)
12
13 # Build the regressor matrix
14 Phi=[X ones(5,1)]
15 @show Phi
16
17 Phi = [1.0 1.0; 2.0 1.0; 4.0 1.0; 5.0 1.0; 7.0 1.0]
18 5x2 Array{Float64, 2}:
19  1.0  1.0
20  2.0  1.0
21  4.0  1.0
22  5.0  1.0
23  7.0  1.0
24
25 # Take a shortcut to finding alpha
26 # because the problem is small
27 alphaStar=inv(Phi' * Phi) * Phi' * Y
28 2x1 Array{Float64, 2}:
29  2.1228070175438596
30  2.3333333333333375
31
32 # Extract the physically meaningful parameters
33 m=alphaStar[1]
34 b=alphaStar[2]
35
36 # Build the line for plotting
37 XX=[1, 7]
38 YY=m*XX.+b
39 scatter(X, Y)
40 plot!(xlabel="X", ylabel="Y", title="Least Squares Fit of a Line to Data", leg=false)
41 plot!(fmt = :png)
42
43 # Plot the line over the scatter plot of the data
44 plot!(XX, YY)
```

The above shows you how to formulate a least squares fit to data by working through THE CLASSIC EXAMPLE, fitting a line to data by minimizing the total squared error (square of the norm of the error vector). We'll do another example so that you understand that you are not limited to fitting "linear functions" to data.

Table 8.2: Data for our second fitting problem.

i	x_i	y_i
1	0	1.0
2	0.25	1.0
3	0.5	1.5
4	0.75	2.0
5	1.0	3.0
6	1.25	4.25
7	1.5	5.5
8	1.75	7.0
9	2.0	10.0




Figure 8.3: Scatter plot of the data in Table 8.2. The curve looks nonlinear!

Example 8.2 *Our method of fitting a line to data is more general than it might seem. Consider the data in Table 8.2, which has been plotted in Fig. 8.3. It sure doesn't seem like we should fit a line to it. How about a quadratic?*

Solution: Let's choose a model of the form

$$y = c_0 + c_1x + c_2x^2 = \begin{bmatrix} 1 & x & x^2 \end{bmatrix} \begin{bmatrix} c_0 \\ c_1 \\ c_2 \end{bmatrix}.$$

Note that even though the model is nonlinear in x , it is linear in the unknown coefficients c_0 , c_1 , c_2 . This is what is important!!! Just as before, define $\hat{y}_i = c_0 + c_1x_i + c_2x_i^2$, the i -th term of the error vector is then

$$e_i := y_i - \hat{y}_i = y_i - (c_0 + c_1x_i + c_2x_i^2)$$

and the total squared error is

$$E_{tot} = \sum_{i=1}^N e_i^2.$$

Writing out the equation $y_i = c_0 + c_1x_i + c_2x_i^2, i = 1, \dots, N$ in matrix form yields

$$\underbrace{\begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_N \end{bmatrix}}_Y = \underbrace{\begin{bmatrix} 1 & x_1 & (x_1)^2 \\ 1 & x_2 & (x_2)^2 \\ \vdots & \vdots & \vdots \\ 1 & x_N & (x_N)^2 \end{bmatrix}}_\Phi \underbrace{\begin{bmatrix} c_0 \\ c_1 \\ c_2 \end{bmatrix}}_\alpha,$$

which gives us the equation $Y = \Phi\alpha$. We plug in our numbers and check that $\det(\Phi^\top \cdot \Phi) = 40.6 \neq 0$. The resulting fit is given in Fig. 8.4. The Julia code is also given.

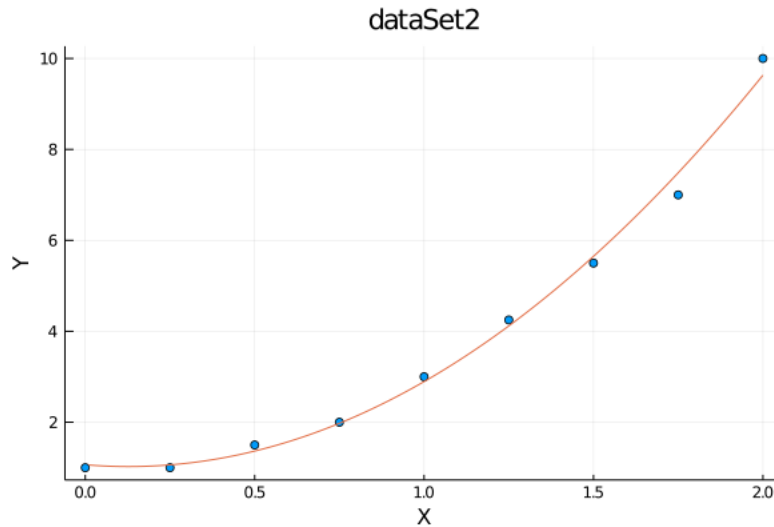


Figure 8.4: A least squares fit of a quadratic curve, $\hat{y} = c_0 + c_1x + c_2x^2$, to data.

```

1 # Data set
2 dataSet2=[
3 1 0.0 1.0
4 2 0.25 1.0
5 3 0.5 1.5
6 4 0.75 2.0
7 5 1.0 3.0
8 6 1.25 4.25
9 7 1.5 5.5
10 8 1.75 7.0
11 9 2.0 10.0]
12
13 # Extract relevant data
14 X=dataSet2[:,2]
15 Y=dataSet2[:,3]
16
17 # Look at data to see what kind of curve it may be
18 using Plots
19 gr()
20 scatter(X, Y)
21 plot!(xlabel="X", ylabel="Y", title="dataSet2", leg=false)
22 plot!(fmt = :png)
23
24 # Build the regressor matrix
25 Phi=[ones(9,1) X X.^2]
26 9 x 3 Array{Float64, 2}:
```

```

27 1.0 0.0 0.0
28 1.0 0.25 0.0625
29 1.0 0.5 0.25
30 1.0 0.75 0.5625
31 1.0 1.0 1.0
32 1.0 1.25 1.5625
33 1.0 1.5 2.25
34 1.0 1.75 3.0625
35 1.0 2.0 4.0
36
37 # Solve Regression Problem
38 alphaStar=inv (Phi' *Phi) *Phi' *Y
39
40 # Plot the curve: first way to do it
41 Yhat=Phi*alphaStar
42 plot! (X, Yhat)
43
44 # Plot the curve with more points in x so that it is smoother
45 temp=1:200
46 XX=vec (temp/100.0) ;
47 N=length (XX)
48 PHI=[ones (N, 1) XX XX.^2]
49 #YY=c0.+ c1.*XX + c2.*(XX).^2;
50 YY=PHI*alphaStar
51
52 # Plot used in the notes
53 scatter (X, Y)
54 plot! (xlabel="X", ylabel="Y", title="dataSet2", leg=false)
55 plot! (fmt = :png)
56 plot! (XX, YY)

```

In HW and in Project 2, we'll explore more refined aspects of regressing functions and surfaces to data. We'll actually divide the data set into two pieces: one to be used for fitting the data and a second portion of the data reserved for checking the quality of the fit. We will be looking to see that the fit to data that was not used in the regression problem is comparable to the fit produced by the regression algorithm. The idea is that the second piece of data better represents how your regressed function (or surface) will work in the real world. **If you have heard of Machine Learning, these ideas are very close to current practice when “training” Supervised Machine Learning Algorithms.**

■

Large Scale Least Squares via the LU Factorization

From Chapter 8.2, if the columns of Φ are linearly independent, or equivalently, $\Phi^T \Phi$ is invertible, we know the following are equivalent

$$\alpha^* = (\Phi^T \Phi)^{-1} \Phi^T Y \iff \alpha^* = \arg \min_{\alpha} \|Y - \Phi \alpha\|^2 \iff (\Phi^T \Phi) \alpha^* = \Phi^T Y. \quad (8.12)$$

The suggested “pipeline” for computing a least squared error solution to $(\Phi^T \Phi) \alpha^* = \Phi^T Y$ is

- factor $P \cdot (\Phi^T \Phi) =: L \cdot U$, that is, do the LU Factorization of $\Phi^T \cdot \Phi$,
- compute $\bar{b} := P \cdot \Phi^T Y$, and then
- solve $Ly = \bar{b}$ via forward substitution, and
- solve $U\alpha^* = y$ via backward substitution.

Reminder on Calling LU Correctly in Julia

```
1 # Find alphaStar by solving Phi' * Y= Phi' * Phi*alphaStar (Ax = b) using LU
  decomposition
2 using LinearAlgebra
3 F = lu(Phi' * Phi)
4 L = F.L
5 U = F.U
6 P = F.P
7 #
8 # Phi' * Y = Phi' * Phi * alphaStar
9 # after LU Factorization of Phi'*Phi, we have
10 # P * Phi' * Y = L * U * alphaStar
11 #
12 y_alpha = forwardsub(L, P*Phi'*Y)
13 alphaStar = backwardsub(U, y_alpha)
```

As a reminder, the following is erroneous:

```
1 # Find alphaStar by solving Phi' * Y= Phi' * Phi*alphaStar (Ax = b) using LU
  decomposition
2 using LinearAlgebra
3 L, U, P = lu(Phi' * Phi)
4 #
5 y_alpha = forwardsub(L, P*Phi'*Y)
6 alphaStar = backwardsub(U, y_alpha)
```

Why? Because P will NOT contain the permutation matrix. It will contain the list of permutation indices. Give a try!

8.4 (Optional Read): How to Derive the Main Regression Formula

The main method is “completing the square”. You probably learned that in High School when you studied the quadratic formula. We recall first the derivation of the quadratic formula via “completing the square” and then give the main steps for a “vector-matrix version of completing the square.”

8.4.1 Completing the Square for the Quadratic Formula:

We assume that $a \neq 0$ and that all coefficients are real numbers. The main “trick” is to recall that $(x + d)^2 = x^2 + 2dx + d^2$. Hence, if you see something like $x^2 + 2dx$, you can complete the square by adding and subtracting d^2 to obtain $x^2 + 2dx = (x + d)^2 - d^2$. Below, we do this for $d = \frac{b}{2a}$.

$$\begin{aligned} ax^2 + bx + c &= 0 \\ \Downarrow \\ x^2 + \frac{b}{a}x + \frac{c}{a} &= 0 \\ \Downarrow \\ x^2 + 2\frac{b}{2a}x + \frac{c}{a} &= 0 \\ \Downarrow \\ \left(x^2 + 2\frac{b}{2a}x + \left(\frac{b}{2a}\right)^2\right) - \left(\frac{b}{2a}\right)^2 + \frac{c}{a} &= 0 \text{ (the square was completed here!)} \end{aligned}$$

Really? Yes, because $x^2 + 2\left(\frac{b}{2a}\right)x + \left(\frac{b}{2a}\right)^2 = \left(x + \frac{b}{2a}\right)^2$. Once we have completed the square, the rest is basic manipulation of terms,

$$\begin{aligned}
 \left(x + \frac{b}{2a}\right)^2 &= \left(\frac{b}{2a}\right)^2 - \frac{c}{a} \\
 &\Downarrow \\
 \left(x + \frac{b}{2a}\right)^2 &= \frac{b^2 - 4ac}{4a^2} \text{ (a perfect square)} \\
 &\Downarrow \\
 \left(x + \frac{b}{2a}\right) &= \pm\sqrt{\frac{b^2 - 4ac}{4a^2}} \text{ (note the plus-minus sign)} \\
 &\Downarrow \\
 \left(x + \frac{b}{2a}\right) &= \pm\frac{\sqrt{b^2 - 4ac}}{2a} \text{ (the rest is "algebra")} \\
 &\Downarrow \\
 x &= -\frac{b}{2a} \pm \frac{\sqrt{b^2 - 4ac}}{2a} \\
 &\Downarrow \\
 x &= \frac{-b \pm \sqrt{b^2 - 4ac}}{2a}.
 \end{aligned}$$

That's a lot of steps. You probably forgot how it went, right? We had to refresh our own thoughts on completing the square.

8.4.2 Completing the Square for Least Squares Solutions to Systems of Linear Equations:

Consider the set of equations $Ax = b$ and suppose that the columns of A are linearly independent, or equivalently, that $A^T A$ is invertible (i.e., $\det(A^T A) \neq 0$). Then the value of x that minimizes the error satisfies

$$A^T A x^* = A^T b.$$

Sketch:

$$\begin{aligned}
 \|Ax - b\|^2 &= (Ax - b)^T (Ax - b) \\
 &= x^T A^T A x - x^T A^T b - b^T A x + b^T b
 \end{aligned} \tag{8.13}$$

This is where completing the square comes in. It is much easier to do if you already know that answer from other techniques! In Linear Algebra, an expression of the form

$$\boxed{(A^T A x - A^T b)^T (A^T A)^{-1} (A^T A x - A^T b)}$$

is the equivalent of a perfect square. We expand all the terms

$$\begin{aligned}
 (A^T A x - A^T b)^T (A^T A)^{-1} (A^T A x - A^T b) &= (x^T A^T A - b^T A) (A^T A)^{-1} (A^T A x - A^T b) \\
 &= \left(x^T - b^T A (A^T A)^{-1}\right) (A^T A x - A^T b) \\
 &= x^T A^T A x - x^T A^T b - b^T A x + b^T A (A^T A)^{-1} A^T b
 \end{aligned}$$

and then relate them to the equation we are trying to minimize.

Substituting this result into (8.13) gives

$$\begin{aligned}
 \|Ax - b\|^2 &= (Ax - b)^T (Ax - b) \\
 &= x^T A^T A x - x^T A^T b - b^T A x + b^T b \\
 &= x^T A^T A x - x^T A^T b - b^T A x + b^T b + b^T A (A^T A)^{-1} A^T b - b^T A (A^T A)^{-1} A^T b \\
 &= (A^T A x - A^T b)^T (A^T A)^{-1} (A^T A x - A^T b) + b^T b - b^T A (A^T A)^{-1} A^T b
 \end{aligned}$$

Here is the *coup de grâce*: We note that $b^\top b - b^\top A (A^\top A)^{-1} A^\top b$ does not depend on x . Hence, the x^* that minimizes (8.13) is that same as the x^* that minimizes

$$(A^\top Ax - A^\top b)^\top (A^\top A)^{-1} (A^\top Ax - A^\top b).$$

Therefore, the solution is

$$(A^\top Ax^* - A^\top b) = 0 \iff A^\top Ax^* = A^\top b.$$

Remark: Uniqueness follows because $A^\top A$ is invertible. A more complete derivation would use properties of positive definite matrices that are covered in Appendix A.3.

8.5 Looking Ahead

The next two Chapters will conclude our introduction to Computational Linear Algebra. We want to leave space for covering some nonlinear topics in Chapters 11 and 12, topics that are very computational in nature and super useful in engineering practice.

Our next chapter will build on our notions of linear combinations and linear independence to introduce a tool, called a dot product, that allows us to study the notion of “orthogonal vectors” (also called perpendicular vectors) in \mathbb{R}^n for any $n \geq 2$. You probably recall “right angles” from planar geometry? Don’t worry, we will not be doing “similar triangles” or the “side-angle-side theorem”. What we will be doing is more on the level of difficulty of the Pythagorean Theorem, $a^2 + b^2 = c^2$. We’ll introduce an algorithm that you will want to program up in Julia for constructing orthogonal vectors from a set of linearly independent vectors. And from such vectors, we build matrices that have the amazing property that their inverse is equal to their transpose! We know, it seems too good to be true. And we’ll learn how to write any matrix A as the product of one of these magical matrices and an upper triangular matrix.

Chapter 9

The Vector Space \mathbb{R}^n : Part 2

Learning Objectives

- A second encounter with some of the essential concepts in Linear Algebra.
- A more abstract view of \mathbb{R}^n as a vector space.

Outcomes

- What is a vector space, a subspace, and the span of a set of vectors.
- Range, column span, and null space of a matrix.
- The dot product and orthogonality.
- Gram Schmidt process for generating a basis consisting of orthogonal vectors
- Orthogonal matrices: they have the magical property that their inverse is the matrix transpose.
- The QR Factorization is the most numerically robust method for solving systems of linear equations.
- Our second recommended “pipeline” (CS-speak for method) for solving $Ax = b$.

Reminder: Additional Study Time May Be Required

For many students, the concepts in Chapters 7, 9 and 10 are significantly more challenging than the material in any of the other chapters of the book. Why? Well, the reasons vary from person to person, but the biggest reason seems to be the level of abstraction. All of our previous work has pretty much dealt with solving equations and that is something most students of ROB 101 can wrap their heads around.

9.1 Motivation

Up to this point, everything we have studied has been built around the notion of solving systems of linear equations. We found that it is very powerful to express linear equations in terms of matrices and vectors, namely in the form $Ax = b$. While x is really a vector of “unknowns”, we’ve become so comfortable with the vector x being an object in and of itself that we’re perfectly content to call x “the unknown”. Likewise, we’re content to think of the matrix A as an object, and even though it may be 50×50 and hence have 2,500 individual entries, we’ve learned that moving our attention away from the individual entries to the bigger object, namely the matrix formed by the entries, is a very useful and powerful idea for solving equations.

Being able to think in terms of vectors and matrices was a big step in the level of abstraction from how we initially expressed equations in Chapter 1. For those of you who had seen matrices in High School, you took it in stride and thought for sure we’d be doing eigenvalues and eigenvectors pretty soon, though you would not have been able to say why studying them would be important for engineers, and for sure, you did not see triangular matrices, LU Factorization, and least-squared-error solutions coming nor how these tools would allow us to solve (exactly or approximately) equations with hundreds, or even thousands of variables. For those of you who had not seen matrices and vectors in High School, it was initially very hard to grasp the idea of vectors and matrices, but once you had accepted them as valid mathematical objects, you got to immediately manipulate them in Julia instead of only doing hand calculations. By working with vectors and matrices in a programming language, you could associate them with doing “math at the scale of life” instead of the tedium, not to mention the high chance of error, of doing hand calculations with vectors and matrices. And when we said “NOPE” to formulas for computing determinants or matrix inverses for general square matrices, but instead took a structured approach through triangular matrices and factorizations, you probably thought, yeah, this is how it is supposed to be done, and you were right!

In this Chapter, we take another big step in the level of abstraction. In fact, it’s a really big step. Instead of looking at a given set of vectors or a particular linear combination formed from a set of vectors, we will look at the “set of all possible linear combinations” that one can generate from the set of vectors! Or, instead of looking at one solution to a set of linear equations, we will look at the “set of all possible solutions” to the set of equations! Why in the world would we ever want to do that? You’re probably saying to yourself, “the payoffs better be worth it, because, just thinking about the term *all possible* is hurting my brain!”

For sure. Now, back in Chapter 2, if we had tried to foreshadow everything we would be able to do with vector and matrices, you would have been overwhelmed. In fact, some of you may have sought a different course altogether. We’ll take a chance that you are now more mathematically sophisticated and can handle the truth. If you are not so sure, you can stop here and begin Chapter 9.2.

- (a) When we form all possible linear combinations from a set of vectors that live in \mathbb{R}^n , we’ll generate a subset of \mathbb{R}^n called a *subspace*. We’ll come to realize that we can sort of go back and forth between this larger set, namely the subspace, and the original set of vectors. Now, once we have accepted the existence of this larger set called a subspace, we will ask the questions, “Can a different set of vectors generate the same subspace that we generated from our original vectors? And if so, could these vectors be simpler to work with, or “better” in some sense?” The answers will be “yes” and “yes”! This will lead us to the *QR Factorization* of a matrix, one of the most recommended¹ methods for computing solutions to large systems of linear equations, the other being the *LU Factorization*.
- (b) So far in ROB 101, we have not come to grips with how to treat systems of linear equations $Ax = b$ that have an infinite number of solutions! We’ll look at the set of all possible solutions and realize that it can be formed by knowing any particular solution to the equation, say $A\bar{x} = b$, and a subspace called the *null space* of A . We’ll then search within the set of all possible solutions for a solution having minimum norm. In Project 3, you’ll use this idea to balance a Segway! Here, we’ll not do anything so cool.

9.2 \mathbb{R}^n as a Vector Space

Let’s recall that an n -tuple is a fancy name for an ordered list of n numbers, (x_1, x_2, \dots, x_n) and that we typically identify them with column vectors, as in

$$(x_1, x_2, \dots, x_n) \longleftrightarrow \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix}.$$

¹The world’s leading experts argue back and forth over which is better: QR or LU? The QR Factorization is often better in the face of the numerical errors that arise from representing numbers in digital computers via a finite set of zeros and ones, while in many situations, the LU Factorization can be faster to compute, and hence allow one to solve larger problems.

Moreover, we identify \mathbb{R}^n with the set of all n -column vectors with real entries

$$\{(x_1, x_2, \dots, x_n) \mid x_i \in \mathbb{R}, 1 \leq i \leq n\} \iff \mathbb{R}^n \iff \left\{ \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix} \mid x_i \in \mathbb{R}, 1 \leq i \leq n \right\}$$

Finally, the choice of identifying n -tuples of numbers with column vectors instead of row vectors is completely arbitrary, and yet, we have to choose one or the other when doing computations, and the most common choice is to use column vectors.

Two absolutely key properties of vectors in \mathbb{R}^n is that we know how to add them and obtain another vector in \mathbb{R}^n , namely

$$\begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix} + \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_n \end{bmatrix} := \begin{bmatrix} x_1 + y_1 \\ x_2 + y_2 \\ \vdots \\ x_n + y_n \end{bmatrix}, \quad (9.1)$$

and we know how to multiply a scalar times a vector and obtain another vector in \mathbb{R}^n , namely

$$\alpha \cdot \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix} := \begin{bmatrix} \alpha x_1 \\ \alpha x_2 \\ \vdots \\ \alpha x_n \end{bmatrix}. \quad (9.2)$$

Equation (9.1) says that the sum of two vectors is DEFINED by the sum of their respective components using the definition of addition of real numbers. Equation (9.2) says that the product of a scalar and a vector is DEFINED by the multiplication of the real numbers constituting the components of the vector by the scalar, which is another real number. It is emphasized that the vector operations are defined in terms the elementary operations of addition and multiplication of real numbers.

Now, (9.1) and (9.2) are special cases of linear combinations. In fact, the following statements are equivalent (means, one holds if, and only if, the other holds)

(a) For all real numbers α and β , and all vectors x and y in \mathbb{R}^n

$$\alpha \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix} + \beta \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_n \end{bmatrix} = \begin{bmatrix} \alpha x_1 + \beta y_1 \\ \alpha x_2 + \beta y_2 \\ \vdots \\ \alpha x_n + \beta y_n \end{bmatrix} \quad (9.3)$$

(b) Both (9.1) and (9.2) hold individually.

Remark: To see the equivalency, note that if in (9.3), you take $\alpha = \beta = 1$, you obtain (9.1), while if you take $\beta = 0$, you obtain (9.2). To go the other way around, we observe that, by (9.2),

$$\alpha \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix} + \beta \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_n \end{bmatrix} = \begin{bmatrix} \alpha x_1 \\ \alpha x_2 \\ \vdots \\ \alpha x_n \end{bmatrix} + \begin{bmatrix} \beta y_1 \\ \beta y_2 \\ \vdots \\ \beta y_n \end{bmatrix},$$

and that by (9.1),

$$\begin{bmatrix} \alpha x_1 \\ \alpha x_2 \\ \vdots \\ \alpha x_n \end{bmatrix} + \begin{bmatrix} \beta y_1 \\ \beta y_2 \\ \vdots \\ \beta y_n \end{bmatrix} = \begin{bmatrix} \alpha x_1 + \beta y_1 \\ \alpha x_2 + \beta y_2 \\ \vdots \\ \alpha x_n + \beta y_n \end{bmatrix}.$$

Therefore, (9.1) and (9.2) together imply (9.3).

9.3 Subspaces

Recall that a set V is a **subset** of some other set, say W , if $x \in V \implies x \in W$. One writes $V \subset W$ to denote V is a subset of W . We say that $V = W$ if $V \subset W$ and $W \subset V$ **both hold**.

9.3.1 Subspaces of \mathbb{R}^n

Many students find this material challenging. Yet, it is too important to leave out of any linear algebra course. Chapter 9.3.2 takes a more abstract view of the topic. Surprisingly, for many of us, stepping back for a moment and looking at a topic in a totally different (and perhaps surprising) way can help. For those who wish this experience, please check out Chapter 9.3.2.

Subspace of \mathbb{R}^n

Suppose that $V \subset \mathbb{R}^n$ is nonempty, that is, V is a subset of \mathbb{R}^n and it contains at least one element.

Def. V is a **subspace** of \mathbb{R}^n if any linear combination constructed from elements of V and scalars in \mathbb{R} is once again an element of V . One says that V is **closed under linear combinations**. In symbols, $V \subset \mathbb{R}^n$ is a subspace of \mathbb{R}^n if for all real numbers α and β , and all vectors v_1 and v_2 in V

$$\alpha v_1 + \beta v_2 \in V. \quad (9.4)$$

From the equivalence of (9.3) with the separate conditions given in (9.1) and (9.2), one is free to check that a subset is a subspace by checking individually that it is **closed under vector addition** and **closed under scalar times vector multiplication**.

Being “closed under something” simply means that if you perform the operation “something” on an element of a set, you get a new element that is once again an element of the set. For us, “something” is the operation of “forming linear combinations”, “doing vector addition”, or “doing scalar times vector multiplication”. If you do one of these operations and you end up with something new that is NOT in the subset V , then V is NOT a subspace.

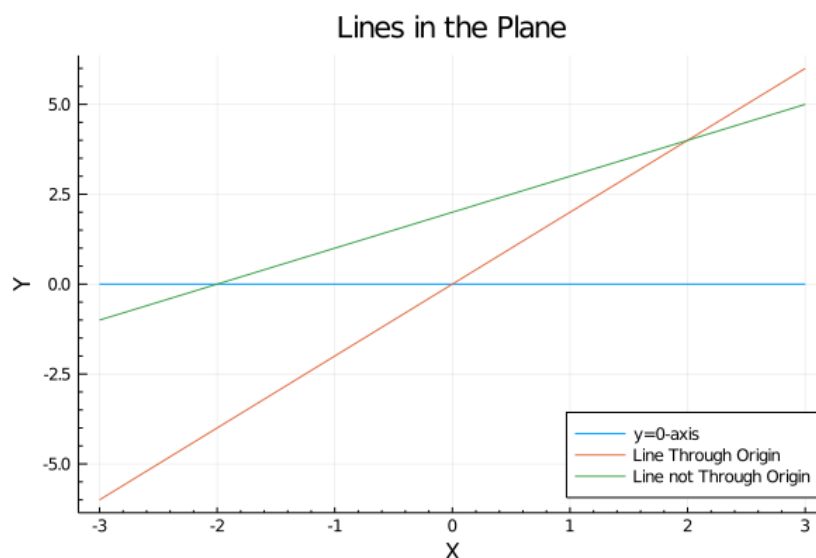


Figure 9.1: If a line does not pass through the origin, then it does not contain the origin, and hence it cannot be a subspace.

Easy First Test for Subspaces:

Every subspace must contain the zero vector. Why? Suppose that $V \subset \mathbb{R}^n$ is a subspace and that $v \in V$. Then $0 \cdot v \in V$ because V is closed under scalar times vector multiplication. But $0 \cdot v = 0$, the zero vector. Figure 9.1 drives home the point that, even in \mathbb{R}^2 , not all lines are subspaces. In fact, almost no lines are subspaces! It's a very special case when the y -intercept equals zero.

Example 9.1 Let $V \subset \mathbb{R}^2$ be the set of all points that lie on a line $y = mx + b$, that is

$$V := \left\{ \begin{bmatrix} x \\ mx + b \end{bmatrix} \mid x \in \mathbb{R} \right\}.$$

Then V is a subspace of \mathbb{R}^2 if, and only if, $b = 0$, that is, the line must pass through the origin.

Solution: V contains the zero vector if, and only, if the y -intercept is zero. But this means that $b = 0$. Now, $0 \in V$ is a *necessary condition*, but not a *sufficient condition* for V to be a subspace. So, let's check if V , with $b = 0$ is closed under vector addition and scalar times vector multiplication. V is then

$$V := \left\{ \begin{bmatrix} x \\ mx \end{bmatrix} \mid x \in \mathbb{R} \right\}.$$

We take

$$v_1 = \begin{bmatrix} x_1 \\ mx_1 \end{bmatrix} \text{ and } v_2 = \begin{bmatrix} x_2 \\ mx_2 \end{bmatrix}$$

for x_1 and x_2 arbitrary real numbers. Then

$$v_1 + v_2 = \begin{bmatrix} x_1 + x_2 \\ mx_1 + mx_2 \end{bmatrix} = \begin{bmatrix} x_1 + x_2 \\ m(x_1 + x_2) \end{bmatrix} \in V,$$

and hence V is closed under vector addition. To be extra clear, we note that

$$v_1 + v_2 = \begin{bmatrix} x \\ mx \end{bmatrix}, \text{ for } x = x_1 + x_2,$$

and that is why $v_1 + v_2 \in V$.

We now let $\alpha \in \mathbb{R}$ be arbitrary and check scalar times vector multiplication.

$$\alpha v_1 = \alpha \begin{bmatrix} x_1 \\ mx_1 \end{bmatrix} = \begin{bmatrix} \alpha x_1 \\ \alpha mx_1 \end{bmatrix} = \begin{bmatrix} \alpha x_1 \\ m(\alpha x_1) \end{bmatrix} \in V,$$

and hence V is closed under scalar times vector multiplication. To be extra clear, we note that

$$\alpha v_1 = \begin{bmatrix} x \\ mx \end{bmatrix}, \text{ for } x = \alpha x_1,$$

and that is why $\alpha v_1 \in V$.

Suppose that $b \neq 0$. Can we show that V is not a subspace without taking the shortcut of first checking that V contains the zero vector? The answer is yes. Let's do vector addition with $b \neq 0$. Then

$$v_1 = \begin{bmatrix} x_1 \\ mx_1 + b \end{bmatrix} \text{ and } v_2 = \begin{bmatrix} x_2 \\ mx_2 + b \end{bmatrix}$$

and

$$v_1 + v_2 = \begin{bmatrix} x_1 + x_2 \\ mx_1 + mx_2 + 2b \end{bmatrix} = \begin{bmatrix} x_1 + x_2 \\ (m(x_1 + x_2) + b) + \boxed{b} \end{bmatrix} \notin V.$$

We see that we cannot write

$$v_1 + v_2 = \begin{bmatrix} x \\ mx + b \end{bmatrix}$$

for some x in \mathbb{R} , when $b \neq 0$, and that is why V is not closed under vector addition. You can also check that it is not closed under scalar times vector multiplication, but for the purpose of showing a subset of \mathbb{R}^n is not a subspace, you only need to violate ONE of the conditions. ■

Example 9.2 $V := \mathbb{R}^n$ is a subspace of \mathbb{R}^n and $W := \{0_{n \times 1}\} \subset \mathbb{R}^n$, the zero vector, is a subspace of \mathbb{R}^n

Solution: We'll let you apply the definition to V and check that it is a subspace. V is the largest subspace of \mathbb{R}^n , because it is all of \mathbb{R}^n , while W is the smallest subspace of \mathbb{R}^n , since it consists only of the zero vector. Is W really a subspace? Well, if we add any two vectors in W , we are adding $n \times 1$ zero vectors, and we'll get the $n \times 1$ zero vector. If we take any real number $\alpha \in \mathbb{R}$ and vector $w \in W$, then we are multiplying the $n \times 1$ zero vector by α , and we get once again the $n \times 1$ zero vector. Hence, W is closed under the operations of vector addition and scalar times vector multiplication, showing that it is a subspace. ■

Example 9.3 $V := \left\{ \begin{bmatrix} x_1 \\ 0 \\ x_1 + 4x_2 \end{bmatrix} \mid x_1, x_2 \in \mathbb{R} \right\}$ is a subspace of \mathbb{R}^3 , while $W := \left\{ \begin{bmatrix} x_1 \\ 5 \\ x_1 + 4x_2 \end{bmatrix} \mid x_1, x_2 \in \mathbb{R} \right\}$ is a NOT subspace of \mathbb{R}^3 .

Solution: W does not contain the zero vector and hence it cannot be a subspace. What about V ? We write

$$v_1 = \begin{bmatrix} \bar{x}_1 \\ 0 \\ \bar{x}_1 + 4\bar{x}_2 \end{bmatrix} \text{ and } v_2 = \begin{bmatrix} \bar{x}_1 \\ 0 \\ \bar{x}_1 + 4\bar{x}_2 \end{bmatrix}.$$

Then,

$$v_1 + v_2 = \begin{bmatrix} \bar{x}_1 + \bar{x}_1 \\ 0 \\ \bar{x}_1 + 4\bar{x}_2 + \bar{x}_1 + 4\bar{x}_2 \end{bmatrix} = \begin{bmatrix} (\bar{x}_1 + \bar{x}_1) \\ 0 \\ (\bar{x}_1 + \bar{x}_1) + 4(\bar{x}_2 + \bar{x}_2) \end{bmatrix} = \begin{bmatrix} x_1 \\ 0 \\ x_1 + 4x_2 \end{bmatrix}$$

for $x_1 = (\bar{x}_1 + \bar{x}_1)$ and $x_2 = (\bar{x}_2 + \bar{x}_2)$. Hence, $v_1 + v_2 \in V$.

Working out $\alpha v_1 \in V$ for all $\alpha \in \mathbb{R}$ follows the same pattern. We conclude that V is a subspace. ■

9.3.2 (Optional Read:) A Broader View of Subspaces

For some readers, this section may be very helpful and for others, it may be a total distraction. You have to decide for yourself!

The key property of vectors is that we know how to form linear combinations of them. Do we know other sets of “things” where we can add them up and get another element of the set and also multiply an element of the set by an arbitrary real number and obtain another element of the set? We submit that you know how to add a finite number of functions to produce another function, and you know how to multiply a function by a real number to define a new function! Hence, the set of all functions from the reals to the reals,

$$X := \{f : \mathbb{R} \rightarrow \mathbb{R}\},$$

is closed under linear combinations and is a vector space. The set includes complicated functions such as $f(x) = x^3 \sin(e^x) - x^5$, radial basis functions $r(x) = e^{\frac{(x-x_c)^2}{s^2}}$, and simpler functions such as $g(x) = 1 + x$.

Example 9.4 What is the zero vector in the vector space X ?

Solution: It would be the zero function, $\text{zero} : \mathbb{R} \rightarrow \mathbb{R}$ by $\text{zero}(x) = 0$ for all $x \in \mathbb{R}$. It is the only function satisfying $\text{zero}(x) + f(x) = f(x)$ for all vectors (functions) in X . ■

Can we define any interesting subsets of this vector space? And would they be subspaces of X ? How about we consider the set of all polynomials with real coefficients? That is,

$$P := \{p(x) \in X \mid p(x) \text{ is a polynomial}\}.$$

P is clearly a subset of X because polynomials are special types of functions. Is it a subspace? Well, the sum of any two polynomials is a polynomial, and when we multiply a polynomial by a real number (that is, we multiply each of its coefficients by the same real number), we obtain another polynomial. Hence, P is a subspace of X .

We could also define

$$P_n := \{p(x) \in P \mid \text{the degree of } p(x) \text{ is less than or equal to } n\}.$$

Do we know specific vectors in P_n ? Sure! How about

$$\{1(x), x, x^2, \dots, x^n\},$$

where $1(x)$ here is denoting the constant function $1(x) := 1 = x^0$ for all $x \in \mathbb{R}$. Wait, those are vectors? They are vectors in the vector space X , the set of all functions from the reals to the reals: $v_k(x) := x^k$, $0 \leq k \leq n$ are perfectly good functions, called the monomials. Moreover, every polynomial consists of linear combinations of monomials!

Using notation that we introduce in the next section,

$$P_n = \text{span}\{x^0, x, x^2, \dots, x^n\},$$

which means that every polynomial of degree less than or equal to n can be written as a linear combination of the monomials from x^0 to x^n . It's also clear that if we remove any elements from this set, then there would be polynomials of degree less than or equal to n that we could not generate. Hence, this set is the smallest one that can generate all polynomials of degree n or less.

Using ideas from Chapter 13.5, we could find a polynomial of a fixed degree that is the best approximation² of a general function in X . That is a useful idea and we have explored a special case of it in homework and in Project 2.

Example 9.5 Is the set $V := \{p(x) \in P \mid \text{the degree of } p(x) \text{ is exactly equal to } n\}$ a subspace, for $n \geq 1$?

Solution: No, because the zero vector does not have degree n and hence is not an element of the set, V . Another way to see it is that if we define $p_1(x) = 2 + 3x^n$ and $p_2(x) = x^n$, then $p_1(x), p_2(x) \in V$, but their linear combination $p_1(x) - 3p_2(x) \notin V$. ■

9.4 Three Sources of Subspaces in \mathbb{R}^n : Matrix Null Space, Span of a set of Vectors, and Column Span of a Matrix

Null Space of a Matrix

For an $n \times m$ matrix A , its **null space** is

$$\text{null}(A) := \{x \in \mathbb{R}^m \mid Ax = 0_{n \times 1}\},$$

the set of all solutions (that is, vectors) that result in Ax being the zero vector or the “null vector”.

Remark: $Ax = 0$ has a unique solution if, and only if, $\text{null}(A) = \{0_{m \times 1}\}$, the zero vector in \mathbb{R}^m . If $Ax = b$ has a solution and $\text{null}(A) \neq \{0_{m \times 1}\}$, then the equation has an infinite number of solutions. We will later find the solution with minimum norm.

Example 9.6 Compute the null space of $A = \begin{bmatrix} 1 & 3 & 0 \\ 0 & 4 & 1 \end{bmatrix}$.

Solution: We note that A is a 2×3 matrix. Hence, its null space is

$$\text{null}(A) := \{x \in \mathbb{R}^3 \mid Ax = 0_{2 \times 1}\} \subset \mathbb{R}^3.$$

Moreover,

$$Ax = 0_{2 \times 1} \iff \begin{bmatrix} 1 & 3 & 0 \\ 0 & 4 & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = 0_{2 \times 1} \iff \begin{cases} x_1 + 3x_2 + 0x_3 = 0 \\ 0x_1 + 4x_2 + 1x_3 = 0 \end{cases} \iff \begin{cases} x_1 = -3x_2 \\ x_3 = -4x_2. \end{cases}$$

²There are some technicalities that we are skipping. The function and its square would have to be Riemann integrable (for us engineers, Lebesgue integrable for mathematicians), and we would have to do the best fit over a bounded interval $[a, b] \subset \mathbb{R}$, $a < b$.

Hence,

$$x \in \text{null}(A) \iff x = \begin{bmatrix} -3x_2 \\ x_2 \\ -4x_2 \end{bmatrix} = \left\{ x_2 \begin{bmatrix} -3 \\ 1 \\ -4 \end{bmatrix} \mid x_2 \in \mathbb{R} \right\} = \left\{ \alpha \begin{bmatrix} -3 \\ 1 \\ -4 \end{bmatrix} \mid \alpha \in \mathbb{R} \right\},$$

where we replaced x_2 with α for purely cosmetic reasons, meaning that we did it to emphasize that x_2 was really just an arbitrary scalar. ■

Remark: Anticipating the next definition, we can express the null space of A as

$$\text{null}(A) = \text{span}\left\{ \begin{bmatrix} -3 \\ 1 \\ -4 \end{bmatrix} \right\} := \left\{ \alpha \begin{bmatrix} -3 \\ 1 \\ -4 \end{bmatrix} \mid \alpha \in \mathbb{R} \right\}.$$

Remark: Additional worked examples are in Chapter 10, Example 10.11.

Example 9.7 For an $n \times m$ matrix A , show that $\text{null}(A)$ is a subspace of \mathbb{R}^m .

Solution: Suppose $v_1, v_2 \in \text{null}(A)$, so that $Av_1 = Av_2 = 0_{n \times 1}$. We need to show that for all real numbers α and β ,

$$\alpha v_1 + \beta v_2 \in \text{null}(A) \quad (\text{closed under linear combinations}). \quad (9.5)$$

But, $A(\alpha v_1 + \beta v_2) = \alpha Av_1 + \beta Av_2 = \alpha 0_{n \times 1} + \beta 0_{n \times 1} = 0_{n \times 1}$, and thus $\alpha v_1 + \beta v_2 \in \text{null}(A)$. ■

Span of a Set of Vectors

Suppose that $S \subset \mathbb{R}^n$, then S is a set of vectors. The set of all possible linear combinations of elements of S is called the span of S ,

$$\text{span}\{S\} := \{\text{all possible linear combinations of elements of } S\}.$$

It follows that $\text{span}\{S\}$ is a subspace of \mathbb{R}^n because, by definition, it is closed under linear combinations. This is true for any subset $S \subset \mathbb{R}^n$.

What is the span good for? The span operation is how one takes an arbitrary set of vectors and generates a subspace from it. If S is a set, $\text{span}\{S\}$ is the smallest subspace that contains all of the elements of the set S . Shortly, we'll find a different set, \tilde{S} , such that $\text{span}\{S\} = \text{span}\{\tilde{S}\}$ AND \tilde{S} is a "nicer" set of vectors.

If a set S is already known to be a subspace of \mathbb{R}^n , then taking its span does not add any new vectors because a subspace is closed under linear combinations. Hence, $S \subset \mathbb{R}^n$ and S a subspace $\implies \text{span}\{S\} = S$. When S is not a subspace, then there is at least one linear combination of elements of S that is not in S itself. In this case, the span is a bigger set, meaning that $S \subset \text{span}\{S\}$ and $S \neq \text{span}\{S\}$.

Example 9.8 Consider the vector space \mathbb{R}^3 and two vectors e_1 and e_2 , where for $k = 1, 2$, e_k has a one in the k -th entry and zeros elsewhere. In other words, e_1 and e_2 are the first and second columns of I_3 , the 3×3 identity matrix. Compute $\text{span}\{e_1 + e_3, e_2 + e_3\}$.

Solution: We have $S = \{e_1 + e_3, e_2 + e_3\}$. The set $\{e_1 + e_3, e_2 + e_3\}$ is not a subspace because the vector $e_1 - e_2 = (e_1 + e_3) -$

$$(e_1 + e_3) \notin S.$$

$$\begin{aligned} \text{span}\{S\} &:= \{\text{all possible linear combinations of elements of } S\} \\ &= \left\{ \alpha_1(e_1 + e_3) + \alpha_2(e_2 + e_3) \mid \alpha_1 \in \mathbb{R}, \alpha_2 \in \mathbb{R} \right\} \\ &= \left\{ \alpha_1 \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix} + \alpha_2 \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix} \mid \alpha_1 \in \mathbb{R}, \alpha_2 \in \mathbb{R} \right\} \\ &= \left\{ \begin{bmatrix} \alpha_1 \\ \alpha_2 \\ \alpha_1 + \alpha_2 \end{bmatrix} \mid \alpha_1 \in \mathbb{R}, \alpha_2 \in \mathbb{R} \right\} \\ &= \left\{ \begin{bmatrix} x \\ y \\ x + y \end{bmatrix} \mid x \in \mathbb{R}, y \in \mathbb{R} \right\} \\ &= \left\{ \begin{bmatrix} x \\ y \\ z \end{bmatrix} \mid z = x + y, x \in \mathbb{R}, y \in \mathbb{R} \right\}. \end{aligned}$$

Hence, $\text{span}\{e_1 + e_3, e_2 + e_3\}$ is the plane given by $z = x + y$ in \mathbb{R}^3 . ■

For this very simple and very special case, we could “recognize” what the span of the set of vectors turned out to be and name it. More generally, we cannot “recognize” what the span looks like. We simply use it as convenient notation when we want to work with the linear combinations associated with a set of vectors and not just the vectors themselves.




Figure 9.2: Where do null spaces come from? (a) A chamber to cool molecules to near absolute zero, courtesy of Imperial College of London. The laser beams cool the center object by pushing on it to reduce its “vibrational energy”. You can think of it as “clamping the molecule in place”. (b) A toy example we use to illustrate null spaces. Arrows 2, 3, 5, and 6 are at 45° .

Example 9.9 (Combining null space, span, and physical intuition) For the image in Fig. 9.2-(b), we imagine the laser beams as applying forces to the center molecule. Labeling the forces as F_1 through F_6 and balancing their x and y components to achieve an equilibrium results in

$$\underbrace{\begin{bmatrix} 1 & \frac{\sqrt{2}}{2} & -\frac{\sqrt{2}}{2} & -1 & -\frac{\sqrt{2}}{2} & \frac{\sqrt{2}}{2} \\ 0 & -\frac{\sqrt{2}}{2} & -\frac{\sqrt{2}}{2} & 0 & \frac{\sqrt{2}}{2} & \frac{\sqrt{2}}{2} \end{bmatrix}}_A \underbrace{\begin{bmatrix} F_1 \\ F_2 \\ F_3 \\ F_4 \\ F_5 \\ F_6 \end{bmatrix}}_x = \begin{bmatrix} 0 \\ 0 \end{bmatrix}. \quad (9.6)$$

Compute the null space of A and interpret it!

Solution: Newton's Laws tell us that for a particle to be in equilibrium, the forces acting on it must "balance out", which is "physics-speak" for saying that when you treat the forces as vectors in \mathbb{R}^2 , their vector sum must add up to zero. This can also be interpreted as the sum of the forces in the x -direction must be zero and the same for the sum in the y -direction. This vector sum has been done for you in Equation (9.6). The purpose of the problem is to understand that equilibrium problems like the one posed in the example correspond to null spaces of a matrix.

Pure mathematical computation of the null space: The matrix A has two rows and six columns. A non-zero vector in the null space gives a non-trivial linear combination of the columns of A that add up to the zero vector. We claim that there are four linearly independent vectors in the null space. Let's compute them!

The first two columns of A are linearly independent; indeed the matrix

$$\underbrace{\begin{bmatrix} 1 & \frac{\sqrt{2}}{2} \\ 0 & -\frac{\sqrt{2}}{2} \end{bmatrix}}_A$$

is triangular and has non-zero entries on its diagonal. We can use the linear independence of these two columns to express each of the remaining four columns as linear combinations of the first two.

We write

$$[b_1 \ b_2 \ b_3 \ b_4] := \begin{bmatrix} -\frac{\sqrt{2}}{2} & -1 & -\frac{\sqrt{2}}{2} & \frac{\sqrt{2}}{2} \\ -\frac{\sqrt{2}}{2} & 0 & \frac{\sqrt{2}}{2} & \frac{\sqrt{2}}{2} \end{bmatrix},$$

so that the vectors b_i , $1 \leq i \leq 4$ are the remaining columns of A . We set up the equations

$$\underbrace{\begin{bmatrix} 1 & \frac{\sqrt{2}}{2} \\ 0 & -\frac{\sqrt{2}}{2} \end{bmatrix}}_A \begin{bmatrix} \alpha_{i,1} \\ \alpha_{i,2} \end{bmatrix} = b_i,$$

and note that because \bar{A} is invertible, each of the equations has a unique solution. After slightly rearranging the above we have

$$\alpha_{i,1} \begin{bmatrix} 1 \\ 0 \end{bmatrix} + \alpha_{i,2} \begin{bmatrix} \frac{\sqrt{2}}{2} \\ -\frac{\sqrt{2}}{2} \end{bmatrix} - b_i = 0_{2 \times 1}.$$

It follows that the four vectors

$$\left\{ v_1 = \begin{bmatrix} \alpha_{1,1} \\ \alpha_{1,2} \\ -1 \\ 0 \\ 0 \\ 0 \end{bmatrix}, v_2 = \begin{bmatrix} \alpha_{2,1} \\ \alpha_{2,2} \\ 0 \\ -1 \\ 0 \\ 0 \end{bmatrix}, v_3 = \begin{bmatrix} \alpha_{3,1} \\ \alpha_{3,2} \\ 0 \\ 0 \\ -1 \\ 0 \end{bmatrix}, v_4 = \begin{bmatrix} \alpha_{4,1} \\ \alpha_{4,2} \\ 0 \\ 0 \\ 0 \\ -1 \end{bmatrix} \right\}$$

are in the null space of A because

$$\begin{aligned} Av_i &= \begin{bmatrix} 1 & \frac{\sqrt{2}}{2} & -\frac{\sqrt{2}}{2} & -1 & -\frac{\sqrt{2}}{2} & \frac{\sqrt{2}}{2} \\ 0 & -\frac{\sqrt{2}}{2} & -\frac{\sqrt{2}}{2} & 0 & \frac{\sqrt{2}}{2} & \frac{\sqrt{2}}{2} \end{bmatrix} v_i \\ &= \alpha_{i,1} \begin{bmatrix} 1 \\ 0 \end{bmatrix} + \alpha_{i,2} \begin{bmatrix} \frac{\sqrt{2}}{2} \\ -\frac{\sqrt{2}}{2} \end{bmatrix} - b_i \\ &= 0_{2 \times 1}. \end{aligned}$$

We'll leave it to the reader to check that the four vectors $\{v_1, v_2, v_3, v_4\}$ are linearly independent. Once you note that the last four entries of $\{v_1, v_2, v_3, v_4\}$ consist of zeros and negative ones, showing independence is very quick!

Could there be more than four linearly independent vectors in the null space of A ? We'll prove in Chapter 10.5 that four is the magic number and that this follows from the much easier observation that A has two linearly independent rows and six columns! Hence,

$$\text{null}(A) = \text{span}\{v_1, v_2, v_3, v_4\}.$$

We next give a completely different method and claim that we can “see” the null space from the diagram. If we increase F_1 and F_4 by the same amount, then they cancel in both the x and y directions. If we increase F_2 and F_5 by the same amount, then they cancel. If we increase F_3 and F_6 by the same amount, then they cancel. And finally, if we increase F_2 and F_6 by the same amount, their y -components cancel but, not their x -components, so we need to offset that by F_4 . Ta da! Pretty cool, right? This gives

$$\text{null}(A) = \text{span} \left\{ \begin{bmatrix} 1 \\ 0 \\ 0 \\ 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 1 \\ 0 \\ 0 \\ 1 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 0 \\ \sqrt{2} \\ 0 \\ 1 \end{bmatrix} \right\}, \quad (9.7)$$

where

$$\begin{array}{l} \begin{bmatrix} 1 \\ 0 \\ 0 \\ 1 \\ 0 \\ 0 \end{bmatrix} \leftrightarrow F_1 \text{ and } F_4 \text{ balancing one another,} \\ \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \\ 1 \\ 0 \end{bmatrix} \leftrightarrow F_2 \text{ and } F_5 \text{ balancing one another} \\ \begin{bmatrix} 0 \\ 0 \\ 1 \\ 0 \\ 0 \\ 1 \end{bmatrix} \leftrightarrow F_3 \text{ and } F_6 \text{ balancing one another,} \\ \begin{bmatrix} 0 \\ 1 \\ 0 \\ \sqrt{2} \\ 0 \\ 1 \end{bmatrix} \leftrightarrow F_2, \sqrt{2}F_4 \text{ and } F_6 \text{ balancing.} \end{array}$$

The null space is then all linear combinations of forces that result in the particle being in equilibrium. This same principle was used to analyze the truss-bridge example in Fig. 5.2.

Statics Problems and Null Spaces

When you take a Statics & Dynamics course, all of the statics problems involve computing forces that lie in null spaces of matrices, where the matrices represent the (vector) sum of the forces (and possibly torques) being zero.

Question: Are we missing the linear combination of F_3 , F_5 and F_1 given by

$$\begin{bmatrix} \sqrt{2} \\ 0 \\ 1 \\ 0 \\ 1 \\ 0 \end{bmatrix} ?$$

Analogously to how F_2 , F_6 , and F_4 work together, this vector will also result in a zero net change to the force equilibrium, won't it? The answer is a resounding yes, but it is already included in the above span, because

$$\begin{bmatrix} \sqrt{2} \\ 0 \\ 1 \\ 0 \\ 1 \\ 0 \end{bmatrix} = \sqrt{2} \begin{bmatrix} 1 \\ 0 \\ 0 \\ 1 \\ 0 \\ 0 \end{bmatrix} + \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \\ 1 \\ 0 \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ 1 \\ 0 \\ 0 \\ 1 \end{bmatrix} - \begin{bmatrix} 0 \\ 1 \\ 0 \\ \sqrt{2} \\ 0 \\ 1 \end{bmatrix}.$$

Question: Is it a problem that the four vectors we found in our mathematical calculation are different than the four vectors we found using physical reasoning based on force balancing? The answer is no. We invite you to go to Julia and check that the two spans are the same, which you can do by writing the vectors in one span as linear combinations of the vectors in the other span. In Chapter 10.2, we will introduce the notion of a *basis for a subspace*. Our two methods of computing the null space of A have produced two different sets of basis vectors. Is one better than the other? No. Do you like one more than the other? Probably! Some of you will like the physical reasoning and some will be hard-core math types. The world needs all types!

Hopefully, this example helps to clarify linear combinations, span, and null space. ■

Example 9.10 Let S be the unit circle in \mathbb{R}^2 . Compute $\text{span}\{S\}$.

Solution: We have that $S := \{(x_1, x_2) \mid (x_1)^2 + (x_2)^2 = 1\} \subset \mathbb{R}^2$, which has an infinite number of elements. Hence, computing its span must be super tricky, right? Not really. We note that e_1 and e_2 are both elements of S because e_1 corresponds to $x_1 = 1$ and $x_2 = 0$, while e_2 corresponds to $x_1 = 0$ and $x_2 = 1$. Hence,

$$\text{span}\{e_1, e_2\} \subset \text{span}\{S\}.$$

But we know that $\text{span}\{e_1, e_2\} = \mathbb{R}^2$ and $\text{span}\{S\} \subset \mathbb{R}^2$, where the last statement is true because every element of S is in \mathbb{R}^2 and a linear combination of vectors in \mathbb{R}^2 gives another vector in \mathbb{R}^2 . Therefore,

$$\mathbb{R}^2 = \text{span}\{e_1, e_2\} \subset \text{span}\{S\} \subset \mathbb{R}^2,$$

which shows that $\text{span}\{S\} = \mathbb{R}^2$. ■

Column Span of a Matrix

Let A be an $n \times m$ matrix. Then its columns are vectors in \mathbb{R}^n . Their span is called the **column span of A** .

$$\text{col span}\{A\} := \text{span}\{a_1^{\text{col}}, \dots, a_m^{\text{col}}\}.$$

We saw in Chapter 7.4 that $Ax = b$ has a solution if, and only if, b is a linear combination of the columns of A . A more elegant way to write this is

$$\mathbf{Ax} = \mathbf{b} \text{ has a solution if, and only if, } \mathbf{b} \in \text{col span}\{\mathbf{A}\}$$

Example 9.11 Suppose $A = \begin{bmatrix} 3 & 2 \\ 1 & -2 \\ -1 & 1 \end{bmatrix}$ and $b = \begin{bmatrix} 0 \\ -8 \\ 5 \end{bmatrix}$. Does $Ax = b$ have a solution?

Solution: This is identical to the problems we solved in Chapter 7.7. We check that

$$b = -2a_1^{\text{col}} + 3a_2^{\text{col}} \in \text{span}\{a_1^{\text{col}}, a_2^{\text{col}}\},$$

and hence b is in the column span of A , and the system of linear equations has a solution, namely,

$$x = \begin{bmatrix} -2 \\ 3 \end{bmatrix}$$

is a solution! ■

Here is an intriguing example.

Example 9.12 Consider the matrices

$$A := \begin{bmatrix} 0.8399 & -1.1136 & -0.7449 \\ -0.8898 & 0.9555 & -0.1578 \\ 0.0069 & 0.8307 & -0.8218 \\ -1.1286 & 2.0774 & -0.2728 \\ -0.0115 & -0.2561 & 0.1850 \end{bmatrix}_{5 \times 3} \quad \text{and } Q := \begin{bmatrix} 0.5046 & 0.1296 & -0.7328 \\ -0.5345 & -0.3516 & -0.6560 \\ 0.0041 & 0.8038 & -0.1712 \\ -0.6780 & 0.3812 & -0.0287 \\ -0.0069 & -0.2612 & -0.0492 \end{bmatrix}_{5 \times 3}. \quad (9.8)$$

Is the vector

$$b := \begin{bmatrix} -0.4524 \\ -0.9162 \\ 0.8603 \\ 0.4963 \\ -0.3618 \end{bmatrix} \quad (9.9)$$

a linear combination of the columns of A ? In other words, is $b \in \text{col span}\{A\}$?

Solution: The matrix Q has been constructed using methods covered in Chapter 9.5 to satisfy $Q^\top Q = I_3$ and

$$\text{col span}\{A\} = \text{col span}\{Q\};$$

hence,

$$b \in \text{col span}\{A\} \iff b \in \text{col span}\{Q\}.$$

For matrices that satisfy $Q^\top \cdot Q = I_n$ for some n , it must be easy to check $b \in \text{col span}\{Q\}$, right? Indeed, we'll later show that

$$b \in \text{col span}\{Q\} \iff b = Q\alpha, \text{ for } \alpha := Q^\top b.$$

In Julia, we verify these conditions, and we're done. We have checked that $Ax = b$ has a solution. ■

(Optional Read) Remark: Is the above method (which we will study in more detail later), better than our previous method with LDLT (recall Chapter 7.7)? It can be faster in Julia because it involves fewer computations. In Example 9.12, once we had Q , we needed to perform two matrix times vector multiplications, namely

$$\alpha := Q^\top b = \begin{bmatrix} 0.5046 & -0.5345 & 0.0041 & -0.6780 & -0.0069 \\ 0.1296 & -0.3516 & 0.8038 & 0.3812 & -0.2612 \\ -0.7328 & -0.6560 & -0.1712 & -0.0287 & -0.0492 \end{bmatrix}_{3 \times 5} \begin{bmatrix} -0.4524 \\ -0.9162 \\ 0.8603 \\ 0.4963 \\ -0.3618 \end{bmatrix} = \begin{bmatrix} -0.0690 \\ 1.2386 \\ 0.7889 \end{bmatrix}$$

and

$$Q\alpha = \begin{bmatrix} 0.5046 & 0.1296 & -0.7328 \\ -0.5345 & -0.3516 & -0.6560 \\ 0.0041 & 0.8038 & -0.1712 \\ -0.6780 & 0.3812 & -0.0287 \\ -0.0069 & -0.2612 & -0.0492 \end{bmatrix}_{5 \times 3} \begin{bmatrix} -0.0690 \\ 1.2386 \\ 0.7889 \end{bmatrix} = \begin{bmatrix} -0.4524 \\ -0.9162 \\ 0.8603 \\ 0.4963 \\ -0.3618 \end{bmatrix},$$

and then verify $Q\alpha - b = 0_{5 \times 1}$. We will learn shortly that Q comes from an algorithm that is similar to a matrix factorization. With LDLT, we would have to perform two matrix products, namely $A^\top \cdot A$ and $A_e^\top \cdot A_e$ for $A_e := [A \ b]$, and two matrix factorizations. A product of a matrix and a vector is generally faster to compute than a matrix times another matrix, and all things being equal, computing one "factorization" is faster than computing two of them. While this is not a rigorous analysis, it gives you a sense that it may be worthwhile to learn about matrices that satisfy $Q^\top \cdot Q = I_n$!

Example 9.13 Check if the following subsets of \mathbb{R}^3 are also subspaces of \mathbb{R}^3 .

- (a) $S_1 = \{x \in \mathbb{R}^3 \mid [\ 2.0 \ -1.0 \ 11.0] x = 1.0\}$
- (b) $S_2 = \{x \in \mathbb{R}^3 \mid [\ 2.0 \ -1.0 \ 11.0] x = 0.0\}$
- (c) $S_3 = \text{span}\{S_1\}$

Solution:

- (a) The zero vector is not contained in S_1 . According to the fact that every subspace must contain the zero vector, S_1 is not a subspace of \mathbb{R}^3 .
- (b) S_2 is the null space of the 1×3 matrix $A := [\ 2.0 \ -1.0 \ 11.0]$, and hence is a subspace.
- (c) The span of any subset of \mathbb{R}^n is always a subspace. This is a sufficient answer.



Example 9.14 For each example, find linearly independent vectors that span the given subspace:

(a) V_1 is defined to be the span of the columns of B_1 , where

$$B_1 = \begin{bmatrix} 1 & 2 \\ 1 & 3 \\ 1 & 4 \end{bmatrix}.$$

(b) V_2 is defined to be the span of the columns of B_2 , where

$$B_2 = \begin{bmatrix} 1 & 2 \\ 1 & 2 \\ 1 & 2 \end{bmatrix}.$$

(c) $V_3 := \text{span}\{S_3\}$, where $S_3 = \{x \in \mathbb{R}^2 \mid (x_1)^2 + (x_2)^2 = 1\}$ (do not miss that $V_3 \subset \mathbb{R}^2$).

(d) $V_4 := \text{span}\{S_4\}$, where $S_4 = \{x \in \mathbb{R}^4 \mid x_1 + 3x_2 + x_3 + 3x_4 = 3\}$ (do not miss that $V_4 \subset \mathbb{R}^4$).

Solutions: The first two problems are easy because you are directly given a finite set of vectors that span the subspace in question. The only issue is whether they are linearly independent or not. The last two problems are more challenging! Let's get to work.

(a) $V_1 := \text{span}\left\{ \underbrace{\begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}}_{v_1}, \underbrace{\begin{bmatrix} 2 \\ 3 \\ 4 \end{bmatrix}}_{v_2} \right\}$. Hence, we already know vectors that span V_1 . We need to check if the given vectors are linearly independent or not. If not, we'll throw away dependent vectors.

We quickly check that $\alpha_1 v_1 + \alpha_2 v_2 = 0_{3 \times 1}$ if, and only if, $\alpha_1 = \alpha_2 = 0$. Thus the columns of B_1 are linearly independent and they span the subspace V_1 .

(b) This time, $V_2 := \text{span}\left\{ \underbrace{\begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}}_{v_1}, \underbrace{\begin{bmatrix} 2 \\ 2 \\ 2 \end{bmatrix}}_{v_2} \right\}$. We observe that $2v_1 - v_2 = 0_{3 \times 1}$ and thus they are linearly dependent. Because each vector is non-zero, either one alone is linearly independent and hence we can choose either one as the linearly independent vector. For example,

- $V_2 = \text{span}\left\{ \underbrace{\begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}}_{v_1} \right\}$, or

- $V_2 = \text{span}\left\{ \underbrace{\begin{bmatrix} 2 \\ 2 \\ 2 \end{bmatrix}}_{v_2} \right\}$.

- We could also write $V_2 = \text{span}\left\{ \begin{bmatrix} 11 \\ 11 \\ 11 \end{bmatrix} \right\}$ because any non-zero multiple of v_1 is also a valid solution.

(c) This problem is quite different from the previous two because we are not given an explicit finite set of vectors that span the subspace. Instead, we are given an implicit description of a set and then we are asked to compute the span of that set of vectors. This is more common in applications.

Before attacking the problem, it's good to establish:

Lemma. If $A \subset B \subset \mathbb{R}^n$, then $\text{span}\{A\} \subset \text{span}\{B\} \subset \text{span}\{\mathbb{R}^n\} = \mathbb{R}^n$.

Proof: If $v \in \text{span}\{A\}$, then $v = \sum_{i=1}^K c_i \cdot v_i$ for some constants $c_i \in \mathbb{R}$ and vectors $v_i \in A$. But because $A \subset B$, we have that $v_i \in B$ for all $1 \leq i \leq K$ and hence $\sum_{i=1}^K c_i \cdot v_i \in \text{span}\{B\}$. Therefore, $v \in \text{span}\{A\} \implies v \in \text{span}\{B\}$, which means that $\text{span}\{A\} \subset \text{span}\{B\}$. The same reasoning can be applied to $B \subset \mathbb{R}^n$. Finally, why is $\text{span}\{\mathbb{R}^n\} = \mathbb{R}^n$? One way to look at it: \mathbb{R}^n is (already) a vector space and is thus closed under linear combinations. Hence, the span operation does not create any new vectors. ■

Let's now find some vectors in S_3 . We note that $v_1 := \begin{bmatrix} 1 \\ 0 \end{bmatrix} \in S_3$ and $v_2 := \begin{bmatrix} 0 \\ 1 \end{bmatrix} \in S_3$. Hence, applying the Lemma, we have that

$$\{v_1, v_2\} \subset S_3 \subset \mathbb{R}^2 \implies \text{span}\{v_1, v_2\} \subset \text{span}\{S_3\} \subset \text{span}\{\mathbb{R}^2\} = \mathbb{R}^2.$$

Next, we note that

$$\text{span}\{v_1, v_2\} = \text{span}\left\{\begin{bmatrix} 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \end{bmatrix}\right\} = \mathbb{R}^2$$

and, therefore,

$$\mathbb{R}^2 = \text{span}\{v_1, v_2\} \subset \text{span}\{S_3\} \subset \mathbb{R}^2.$$

The only way this can hold is if $\text{span}\{S_3\} = \mathbb{R}^2$. Hence, an answer to the problem is

$$\text{span}\{S_3\} = \text{span}\left\{\begin{bmatrix} 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \end{bmatrix}\right\}$$

and we have expressed the subspace as a span of linearly independent vectors.

Are other solutions possible? Yes! $\bar{v}_1 := \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ 1 \end{bmatrix} \in S_3$ and $\bar{v}_2 := \frac{1}{\sqrt{2}} \begin{bmatrix} -1 \\ 1 \end{bmatrix} \in S_3$. Moreover, these vectors are linearly independent. Hence, we can write $\text{span}\{S_3\} = \text{span}\{\bar{v}_1, \bar{v}_2\}$. In fact, there are an infinite number of choices for linearly independent vectors that span V_3 .

- (d) This problem is nearly identical to the previous one. We will enumerate a few vectors in the set and use them to compute the span.

We note that $v_1 := \begin{bmatrix} 3 \\ 0 \\ 0 \\ 0 \end{bmatrix} \in S_4$, $v_2 := \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \end{bmatrix} \in S_4$, $v_3 := \begin{bmatrix} 0 \\ 0 \\ 3 \\ 0 \end{bmatrix} \in S_4$, and $v_4 := \begin{bmatrix} 0 \\ 0 \\ 0 \\ 1 \end{bmatrix} \in S_4$. These four vectors are linearly independent and their span is all of \mathbb{R}^4 . Hence, following the same reasoning as in part (c), we have that $V_4 = \text{span}\{v_1, v_2, v_3, v_4\} = \mathbb{R}^4$. ■

9.5 Dot Product and Orthonormal Vectors

Dot Product or Inner Product

Definition: We let $u \in \mathbb{R}^n$ and $v \in \mathbb{R}^n$ be **column vectors**,

$$u := \begin{bmatrix} u_1 \\ u_2 \\ \vdots \\ u_n \end{bmatrix}, \quad v := \begin{bmatrix} v_1 \\ v_2 \\ \vdots \\ v_n \end{bmatrix}.$$

The **dot product** of u and v is defined as

$$u \bullet v := \sum_{k=1}^n u_k v_k.$$

We note that

$$u^\top \cdot v = [u_1 \quad u_2 \quad \cdots \quad u_n] \cdot \begin{bmatrix} v_1 \\ v_2 \\ \vdots \\ v_n \end{bmatrix} = \sum_{k=1}^n u_k v_k =: u \bullet v$$

For many people, this is how they remember the **dot product**: as $u^\top v$. In fact, you are welcome to use this as the definition of the dot product.

The dot product is also called the **inner product**. The terminology of inner product is very common. In ROB 101, you can choose your preferred terminology in this regard. Your instructors prefer *inner product*.

Example 9.15 Compute the dot product for

$$u = \begin{bmatrix} 1 \\ 0 \\ 3 \end{bmatrix}, \quad v = \begin{bmatrix} 2 \\ 4 \\ 1 \end{bmatrix}.$$

Solution:

$$u \bullet v = (1)(2) + (0)(4) + (3)(1) = 5$$

$$u^\top v = (1)(2) + (0)(4) + (3)(1) = 5.$$

You can use either notation. ■

Example 9.16 Compute the inner product for

$$u = \begin{bmatrix} 1 \\ 0 \\ -1 \\ 0 \end{bmatrix}, \quad v = \begin{bmatrix} 0 \\ 1 \\ 0 \\ 1 \end{bmatrix}.$$

Solution:

$$u \bullet v = (1)(0) + (0)(1) + (-1)(0) + (0)(1) = 0$$

$$u^\top v = (1)(0) + (0)(1) + (-1)(0) + (0)(1) = 0.$$


You can use either notation. ■

Key Use of the Dot (aka Inner) Product of Two Vectors

The inner product will provide us a generalization of a right angle (90 deg angle) between two vectors in \mathbb{R}^n .

$$w_1 \perp w_2 \iff w_1 \bullet w_2 = 0 \iff w_1^\top w_2 = 0$$

(Read it as: w_1 is **orthogonal** to w_2 if, and only if, their inner product is zero. Orthogonal means “at right angle”)



Source: <https://study.com/academy/lesson/the-gram-schmidt-process-for-orthonormalizing-vectors.html>

Reading the values from the graph, we have

$$w_1 = \begin{bmatrix} 3 \\ 4 \end{bmatrix}, w_2 = \begin{bmatrix} -\frac{7}{3} \\ \frac{7}{4} \end{bmatrix} \implies w_1 \bullet w_2 = w_1^\top w_2 = -3\frac{7}{3} + 4\frac{7}{4} = 0$$

It's very useful and amazing that this works for all vector spaces \mathbb{R}^n , as long as $n \geq 2$. Can you picture a right angle in \mathbb{R}^{27} ? Neither can your instructors, but later we'll see how useful the idea can be!

Pythagorean Theorem in \mathbb{R}^n , $n \geq 2$.

Suppose that $w_1 \perp w_2$. Then,

$$\|w_1 + w_2\|^2 = \|w_1\|^2 + \|w_2\|^2.$$

Remark In the above plot, draw the line from w_1 to w_2 and call its length c ; in addition, label the length of w_1 as a and that of w_2 as b . Then yes, we have the classic relationship: $c^2 = a^2 + b^2$.

Example 9.17 *Why is this true?*

Solution: [**Trigger Warning: This is a proof. You may want to skip it.**] Because $w_1 \perp w_2$, we know that $w_1 \bullet w_2 = 0$, which means that $w_1^\top \cdot w_2 = w_2^\top \cdot w_1 = 0$. Finally, we recall that the norm-squared of a vector v is $\|v\|^2 = v^\top \cdot v$. Using these facts we grind out the computation and see

$$\begin{aligned} \|w_1 + w_2\|^2 &:= (w_1 + w_2)^\top \cdot (w_1 + w_2) \\ &= w_1^\top \cdot (w_1 + w_2) + w_2^\top \cdot (w_1 + w_2) \\ &= w_1^\top \cdot w_1 + w_1^\top \cdot w_2 + w_2^\top \cdot w_1 + w_2^\top \cdot w_2 \\ &= \underbrace{w_1^\top \cdot w_1}_{\|w_1\|^2} + \underbrace{w_1^\top \cdot w_2}_0 + \underbrace{w_2^\top \cdot w_1}_0 + \underbrace{w_2^\top \cdot w_2}_{\|w_2\|^2} \\ &= \|w_1\|^2 + \|w_2\|^2. \end{aligned}$$

Example 9.18 Determine which pairs of vectors, if any, are orthogonal

$$u = \begin{bmatrix} 2 \\ 1 \\ -1 \end{bmatrix}, v = \begin{bmatrix} 1 \\ 3 \\ 5 \end{bmatrix}, w = \begin{bmatrix} -5 \\ 0 \\ 1 \end{bmatrix}.$$

Solution

$$\begin{aligned} u \bullet v &= u^\top \cdot v = [2 \quad 1 \quad -1] \cdot \begin{bmatrix} 1 \\ 3 \\ 5 \end{bmatrix} = (2)(1) + (1)(3) + (-1)(5) = 0 \\ u \bullet w &= u^\top \cdot w = [2 \quad 1 \quad -1] \cdot \begin{bmatrix} -5 \\ 0 \\ 1 \end{bmatrix} = (2)(-5) + (1)(0) + (-1)(1) = -11 \\ v \bullet w &= v^\top \cdot w = [1 \quad 3 \quad 5] \cdot \begin{bmatrix} -5 \\ 0 \\ 1 \end{bmatrix} = (1)(-5) + (3)(0) + (5)(1) = 0, \end{aligned}$$

and hence, $u \perp v$, $u \not\perp w$, and $v \perp w$. In words, u is orthogonal to v , u is not orthogonal to w , and v is orthogonal to w . ■

Orthogonal and Orthonormal Vectors

A set of vectors $\{v_1, v_2, \dots, v_n\}$ is **orthogonal** if, for all $1 \leq i, j \leq n$, and $i \neq j$

$$v_i \bullet v_j = 0. \tag{9.10}$$

We can also write this as $v_i^\top v_j = 0$ or $v_i \perp v_j$.

A set of vectors $\{v_1, v_2, \dots, v_n\}$ is **orthonormal** if,

- they are orthogonal, and
- for all $1 \leq i \leq n$, $\|v_i\| = 1$.

Example 9.19 Scale the vector w so that its norm becomes one,

$$w = \begin{bmatrix} -5 \\ 0 \\ 1 \end{bmatrix}.$$

Solution: In general, if $\alpha := \|w\| \neq 0$ and we define $\tilde{w} := \frac{1}{\alpha}w$, then

$$\|\tilde{w}\| = 1.$$

This is true because

$$\begin{aligned} \|\tilde{w}\| &= \left\| \frac{1}{\alpha} \cdot w \right\| \\ &= \left| \frac{1}{\alpha} \right| \cdot \|w\| \quad (\text{property of norms}) \\ &= \frac{1}{\alpha} \cdot \|w\| \quad \left(\frac{1}{\alpha} \text{ is positive} \right) \\ &= \frac{1}{\alpha} \cdot \alpha \quad (\text{definition of } \alpha) \\ &= \frac{\alpha}{\alpha} = 1 \quad (\text{how scalar multiplication and division work}) \end{aligned}$$

Hence, we need to form $\frac{w}{\|w\|}$, which gives

$$\tilde{w} := \frac{1}{\|w\|} \cdot w = \frac{1}{\sqrt{26}} \begin{bmatrix} -5 \\ 0 \\ 1 \end{bmatrix}.$$

■

Example 9.20 From Example 9.18, we already know that the set $\{u, v\}$ is orthogonal, that is, $u \perp v$. Make it an orthonormal set.

$$u = \begin{bmatrix} 2 \\ 1 \\ -1 \end{bmatrix}, v = \begin{bmatrix} 1 \\ 3 \\ 5 \end{bmatrix}.$$

Solution: We need to normalize their lengths to one. We compute

$$\|u\| = \sqrt{(2)^2 + (1)^2 + (-1)^2} = \sqrt{6}$$

$$\|v\| = \sqrt{(1)^2 + (3)^2 + (5)^2} = \sqrt{35}$$

and thus

$$\left\{ \tilde{u} := \frac{1}{\sqrt{6}} \begin{bmatrix} 2 \\ 1 \\ -1 \end{bmatrix}, \tilde{v} := \frac{1}{\sqrt{35}} \begin{bmatrix} 1 \\ 3 \\ 5 \end{bmatrix} \right\}$$

is an orthonormal set of vectors.

■

Orthonormal Vectors are Linearly Independent

For a set of vectors $\{v_1, v_2, \dots, v_k\}$ in \mathbb{R}^n , the following statements are true:

- (a) $\{v_1, v_2, \dots, v_k\}$ orthonormal implies it is linearly independent
- (b) $\{v_1, v_2, \dots, v_k\}$ orthogonal and for all i , $v_i \neq 0_{n \times 1}$, together imply that the set is linearly independent

Remark: In (b), the zero vector would be orthogonal to everything, but we needed to exclude it for us to have a linearly independent set. In (a), the lengths of the vectors being 1.0 excludes the zero vector.

What about the other direction? That is, given linearly independent vectors, can we construct a set of orthonormal vectors that span them? And if we can, would it even be useful? In Chapter 9.7.2, we address these questions.

9.6 Orthogonal Matrices or Why Orthonormal Vectors are Super Useful

This section is super cool! You will first learn about a kind of matrix whose inverse is equal to its transpose!! This is only the second inverse formula that your instructors want you to know and actually use!!!

The hardest aspect of the this section is the vocabulary BECAUSE **a square matrix is orthogonal** if its columns are **orthonormal vectors**. I know, why not call them orthonormal matrices? Because that would be too easy? Because, by making it confusing, we can check who really knows what they are doing and who does not? No, it's another one of those historical accidents that we live with. Fortunately, the terminology *orthonormal matrices* is used for a rectangular version of orthogonal matrices, and in applications, those are very important too.

Let's recall something about the sizes of matrices. If a matrix Q is $n \times m$, then its transpose is $m \times n$. Therefore, we can form $Q^T \cdot Q$

and have an $m \times m$ square matrix and we can form $Q \cdot Q^\top$ and have an $n \times n$ square matrix.

Orthonormal and Orthogonal Matrices

An $n \times m$ rectangular matrix Q is **orthonormal**:

- if $n > m$ (tall matrix), its columns are orthonormal vectors, which is equivalent to $Q^\top \cdot Q = I_m$; and
- if $n < m$ (wide matrix), its rows are orthonormal vectors, which is equivalent to $Q \cdot Q^\top = I_n$.

A square $n \times n$ matrix is **orthogonal** if $Q^\top \cdot Q = I_n$ and $Q \cdot Q^\top = I_n$, and hence, $Q^{-1} = Q^\top$.

Remarks:

- For a **square matrix**, $n = m$, $(Q^\top \cdot Q = I_n) \iff (Q \cdot Q^\top = I_n) \iff (Q^{-1} = Q^\top)$.
- For a tall matrix, $n > m$, $(Q^\top \cdot Q = I_m) \not\iff (Q \cdot Q^\top = I_n)$.
- For a wide matrix, $m > n$, $(Q \cdot Q^\top = I_n) \not\iff (Q^\top \cdot Q = I_m)$.

Determinants of Orthogonal Matrices

Suppose that Q is $n \times n$ and orthogonal so that $Q^\top Q = I_n$. Then $[\det(Q)]^2 = 1$, and hence $\det(Q) = \pm 1$. This follows from the determinant rule for a product of matrices, once you know for a square matrix A that $\det(A^\top) = \det(A)$.

$$Q \text{ orthogonal} \implies \det(Q) = \pm 1.$$

9.7 Constructing Orthonormal Vectors: the Gram-Schmidt Process

We already know how to normalize a set of orthogonal vectors to create a set of orthonormal vectors. Hence, we next look at how we might go about building a set of orthogonal vectors from vectors that are not orthogonal. We'll then learn the Gram-Schmidt Process, which is simply an algorithmic form of our process for building orthogonal vectors from non-orthogonal vectors.

9.7.1 Building Orthogonal Vectors

Example 9.21 Suppose we have two vectors in \mathbb{R}^3 ,

$$u_1 = \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} \text{ and } u_2 = \begin{bmatrix} 1 \\ -1 \\ 2 \end{bmatrix}.$$

It is easy to compute $u_1^\top \cdot u_2 = 2 \neq 0$, and thus the two vectors are not orthogonal. Find, if possible, two vectors v_1 and v_2 such that

- $v_1 \perp v_2$,
- $\text{span}\{v_1\} = \text{span}\{u_1\}$, and
- $\text{span}\{v_1, v_2\} = \text{span}\{u_1, u_2\}$.

In other words, v_1 and v_2 are orthogonal and yet they “generate” the same subspace as u_1 and u_2 , which are not orthogonal.

Solution: If we set $v_1 = u_1$, then we trivially satisfy $\text{span}\{v_1\} = \text{span}\{u_1\}$. Let's see if we can write v_2 as a linear combination of u_1 and u_2 in such a way that $v_2 \bullet v_1 = 0$ and $\text{span}\{v_1, v_2\} = \text{span}\{u_1, u_2\}$, where we have used the fact that

$$v_1 \perp v_2 \iff v_1 \bullet v_2 = 0 \iff v_2 \bullet v_1 = 0.$$

Step 1: $v_1 := u_1 = \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}$.

Step 2: $v_2 := u_2 - \alpha v_1$, where we seek to choose α such that $v_2 \bullet v_1 = 0$. We compute

$$\begin{aligned} v_2 \bullet v_1 &= (u_2 - \alpha v_1) \bullet v_1 \\ &= u_2 \bullet v_1 - \alpha v_1 \bullet v_1. \end{aligned}$$

If $v_1 \bullet v_1 \neq 0$, then we can set $u_2 \bullet v_1 - \alpha v_1 \bullet v_1 = 0$ and solve for α , namely

$$\alpha = \frac{u_2 \bullet v_1}{v_1 \bullet v_1}.$$

Important Formula to Build $v_1 \perp v_2$ from u_1 and u_2 while Preserving Spans

$$v_1 = u_1$$

$$v_2 = u_2 - \left(\frac{u_2 \bullet v_1}{v_1 \bullet v_1} \right) v_1$$

(9.11)

$$\begin{aligned} \text{span}\{v_1\} &= \text{span}\{u_1\} \\ \text{span}\{v_1, v_2\} &= \text{span}\{u_1, u_2\} \end{aligned}$$

In our case,

$$\begin{aligned} u_2 \bullet v_1 &= u_2^\top \cdot v_1 = \begin{bmatrix} 1 & -1 & 2 \end{bmatrix} \cdot \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} = 2 \\ v_1 \bullet v_1 &= u_1 \bullet u_1 = u_1^\top \cdot u_1 = \begin{bmatrix} 1 & 1 & 1 \end{bmatrix} \cdot \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} = 3 \neq 0, \end{aligned}$$

and thus

$$\alpha = \frac{u_2 \bullet v_1}{v_1 \bullet v_1} = \frac{2}{3}.$$

and hence

$$v_2 = u_2 - \alpha v_1 = \begin{bmatrix} 1 \\ -1 \\ 2 \end{bmatrix} - \frac{2}{3} \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} = \begin{bmatrix} \frac{1}{3} \\ -\frac{5}{3} \\ \frac{4}{3} \end{bmatrix}$$

Did it work? Let's check if the vectors are orthogonal, that is, $v_1 \perp v_2$,

$$v_1 \bullet v_2 = v_1^\top \cdot v_2 = \begin{bmatrix} 1 & 1 & 1 \end{bmatrix} \cdot \begin{bmatrix} \frac{1}{3} \\ -\frac{5}{3} \\ \frac{4}{3} \end{bmatrix} = 0.$$

What about the span property? Well, as we noted when we started, $\text{span}\{v_1\} = \text{span}\{u_1\}$ because $v_1 = u_1$.

What about $\text{span}\{v_1, v_2\} = \text{span}\{u_1, u_2\}$? This part becomes a bit technical, so feel free to stop reading here and skip to the next subsection. We first ask, what does it even mean that $\text{span}\{v_1, v_2\} = \text{span}\{u_1, u_2\}$? Well, it means that if we take all linear combinations of the vectors $\{v_1, v_2\}$, we obtain the same "set of vectors" as taking all linear combinations of the vectors $\{u_1, u_2\}$.

We note that $v_1 = u_1$ and hence,

$$v_2 = u_2 - \alpha v_1 \tag{9.12}$$

$$\Downarrow$$

$$v_2 = u_2 - \alpha u_1 \tag{9.13}$$

$$\Downarrow$$

$$u_2 = v_2 + \alpha v_1 \tag{9.14}$$

From (9.13), we have that

$$\text{span}\{v_1, v_2\} = \text{span}\{u_1, u_2 - \alpha u_1\} \subset \text{span}\{u_1, u_2\},$$

while from (9.14),

$$\text{span}\{u_1, u_2\} = \text{span}\{v_1, v_2 + \alpha v_1\} \subset \text{span}\{v_1, v_2\}.$$

But for arbitrary subsets S_1 and S_2 ,

$$S_1 = S_2 \iff (S_1 \subset S_2 \text{ and } S_2 \subset S_1),$$

and thus we have shown that $\text{span}\{v_1, v_2\} = \text{span}\{u_1, u_2\}$.

To be clear, this kind of proof was maybe a bit over the top for ROB 101! ■

9.7.2 Gram-Schmidt Process (aka the Gram-Schmidt Algorithm) for Building Orthonormal Vectors

Let's do one more step of producing orthogonal vectors. Let's assume that we have three linearly independent vectors $\{u_1, u_2, u_3\}$ and that we have already computed $v_1 := u_1$ and $v_2 := u_2 - \left(\frac{u_2 \bullet v_1}{v_1 \bullet v_1}\right) v_1$ so that $v_1 \bullet v_2 = 0$. Let's now work on a third orthogonal vector. We define

$$v_3 := u_3 - \alpha_1 v_2 - \alpha_2 v_1,$$

and try to find values for α_1 and α_2 so that $v_1 \bullet v_3 = 0$ and $v_2 \bullet v_3 = 0$.

We compute

$$\begin{aligned} v_1 \bullet v_3 &= v_1 \bullet (u_3 - \alpha_1 v_1 - \alpha_2 v_2) \\ &= v_1 \bullet u_3 - \alpha_1 v_1 \bullet v_1 - \alpha_2 v_1 \bullet v_2 \\ &= v_1 \bullet u_3 - \alpha_1 v_1 \bullet v_1 \quad (\text{because } v_1 \bullet v_2 = 0) \end{aligned}$$

and

$$\begin{aligned} v_2 \bullet v_3 &= v_2 \bullet (u_3 - \alpha_1 v_1 - \alpha_2 v_2) \\ &= v_2 \bullet u_3 - \alpha_1 v_2 \bullet v_1 - \alpha_2 v_2 \bullet v_2 \\ &= v_2 \bullet u_3 - \alpha_2 v_2 \bullet v_2 \quad (\text{because } v_2 \bullet v_1 = 0). \end{aligned}$$

Solving for α_1 and α_2 and substituting yields

$$v_3 := u_3 - \frac{u_3 \bullet v_1}{v_1 \bullet v_1} v_1 - \frac{u_3 \bullet v_2}{v_2 \bullet v_2} v_2. \tag{9.15}$$




Figure 9.3: Illustration of the Gram Schmidt Process. In the above, $y^i \leftrightarrow u_i$ and $v^i \leftrightarrow v_i$. Image courtesy of Abhishek Venkataraman and Bruce Huang.

Gram-Schmidt Process

Suppose that the set of vectors $\{u_1, u_2, \dots, u_m\}$ is linearly independent and you generate a new set of vectors by

$$\begin{aligned}
 v_1 &= u_1 \\
 v_2 &= u_2 - \left(\frac{u_2 \bullet v_1}{v_1 \bullet v_1} \right) v_1 \\
 v_3 &= u_3 - \left(\frac{u_3 \bullet v_1}{v_1 \bullet v_1} \right) v_1 - \left(\frac{u_3 \bullet v_2}{v_2 \bullet v_2} \right) v_2 \\
 &\vdots \\
 v_k &= u_k - \sum_{i=1}^{k-1} \left(\frac{u_k \bullet v_i}{v_i \bullet v_i} \right) v_i \quad (\text{General Step})
 \end{aligned} \tag{9.16}$$

Then the set of vectors $\{v_1, v_2, \dots, v_m\}$ is

- orthogonal, meaning, $i \neq j \implies v_i \bullet v_j = 0$
- span preserving, meaning that, for all $1 \leq k \leq m$,

$$\text{span}\{v_1, v_2, \dots, v_k\} = \text{span}\{u_1, u_2, \dots, u_k\}, \tag{9.17}$$

and

- linearly independent.

Suggestion: If you do not see the pattern in the steps of the Gram-Schmidt Algorithm, please compare (9.16) to (9.11).

Preview of Basis Vectors

A set of vectors that is linearly independent and spans a given subspace is called a **basis for the subspace**. We are delaying the study of basis vectors until Chapter 10. Equation (9.17) means that if you apply Gram-Schmidt to a set of **basis vectors**, then you will end with a **basis made of orthogonal vectors**. If you then normalize the **orthogonal basis**, you will produce an **orthonormal basis**!

Example 9.22 You are given that the set below is a basis for \mathbb{R}^3 . Produce from it an orthonormal basis.

$$\{u_1, u_2, u_3\} = \left\{ \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix} \right\}$$

Solution:

Step 1 is to apply Gram-Schmidt to produce an orthogonal basis.

$$\begin{aligned}
 v_1 &= u_1 = \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix} \\
 v_1 \bullet v_1 &= (v_1)^\top v_1 = 2; \\
 v_2 &= u_2 - \frac{u_2 \bullet v_1}{v_1 \bullet v_1} v_1 \\
 &= \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix} - \underbrace{[1 \ 1 \ 0]}_3 \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix} \frac{1}{2} \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix} = \begin{bmatrix} -\frac{1}{2} \\ \frac{1}{2} \\ 3 \end{bmatrix} \\
 v_2 \bullet v_2 &= v_2^\top v_2 = \frac{19}{2} \\
 v_3 &= u_3 - \frac{u_3 \bullet v_1}{v_1 \bullet v_1} v_1 - \frac{u_3 \bullet v_2}{v_2 \bullet v_2} v_2 \\
 &= \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix} - \underbrace{[1 \ 1 \ 0]}_1 \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix} \frac{1}{2} \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix} - \underbrace{[-\frac{1}{2} \ \frac{1}{2} \ 3]}_{3\frac{1}{2}} \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix} \frac{1}{\frac{19}{2}} \begin{bmatrix} -\frac{1}{2} \\ \frac{1}{2} \\ 3 \end{bmatrix} \\
 &= \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix} - \begin{bmatrix} \frac{1}{2} \\ \frac{1}{2} \\ 0 \end{bmatrix} - \begin{bmatrix} -\frac{7}{38} \\ \frac{7}{38} \\ \frac{21}{19} \end{bmatrix} = \begin{bmatrix} -\frac{6}{19} \\ \frac{6}{19} \\ -\frac{2}{19} \end{bmatrix}.
 \end{aligned}$$

Collecting the answers, we have

$$\{v_1, v_2, v_3\} = \left\{ \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} -\frac{1}{2} \\ \frac{1}{2} \\ 3 \end{bmatrix}, \begin{bmatrix} -\frac{6}{19} \\ \frac{6}{19} \\ -\frac{2}{19} \end{bmatrix} \right\}$$

Step 2: Normalize to obtain an orthonormal basis (often useful to do this, but not always required).

$$\begin{aligned}
 \tilde{v}_1 &= \frac{v_1}{\|v_1\|} = \frac{\sqrt{2}}{2} \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix} \\
 \tilde{v}_2 &= \frac{v_2}{\|v_2\|} = \frac{\sqrt{38}}{38} \begin{bmatrix} -1 \\ 1 \\ 6 \end{bmatrix} \\
 \tilde{v}_3 &= \frac{v_3}{\|v_3\|} = \frac{\sqrt{19}}{19} \begin{bmatrix} -3 \\ 3 \\ -1 \end{bmatrix}
 \end{aligned}$$

All of this is quite tedious by hand, while being super fast and fun in Julia!



When we program the Gram-Schmidt process in Julia, we typically do the normalization as we go. When we normalize by $v_i \leftarrow v_i / \|v_i\|$, it follows that $v_i \bullet v_i = 1$. Hence, the algorithm simplifies to the following.

Gram-Schmidt Process with Normalization

Suppose that the set of vectors $\{u_1, u_2, \dots, u_m\}$ is linearly independent and you generate a new set of vectors by applying Gram-Schmidt with normalization as we go, via

$$\begin{aligned}v_1 &= u_1 \\v_1 &\leftarrow v_1 / \|v_1\| \\v_2 &= u_2 - (u_2 \bullet v_1) v_1 \\v_2 &\leftarrow v_2 / \|v_2\| \\v_3 &= u_3 - (u_3 \bullet v_1) v_1 - (u_3 \bullet v_2) v_2 \\v_3 &\leftarrow v_3 / \|v_3\| \\&\vdots \\v_k &= u_k - \sum_{i=1}^{k-1} (u_k \bullet v_i) v_i \quad (\text{General Step}) \\v_k &\leftarrow v_k / \|v_k\|,\end{aligned} \tag{9.18}$$

where the symbol \leftarrow means that we reassign the vector v_k to its normalized value.

Then the set of vectors $\{v_1, v_2, \dots, v_m\}$ is

- **orthonormal**, meaning, $i \neq j \implies v_i \bullet v_j = 0$ **and** for all i , $\|v_i\| = 1$,
- span preserving, meaning that, for all $1 \leq k \leq m$,

$$\text{span}\{v_1, v_2, \dots, v_k\} = \text{span}\{u_1, u_2, \dots, u_k\},$$

and

- linearly independent.

9.8 QR Factorization and Solutions of Linear Equations

In Chapter 9.3 we introduced the column span of an $n \times m$ matrix as the subspace of \mathbb{R}^n generated by taking all of the linear combinations of the columns of the matrix. Applying Gram-Schmidt to the columns of a matrix yields the QR Factorization, which is one of the most advanced numerical methods for solving systems of linear equations.

QR Factorization

Suppose that A is an $n \times m$ matrix with linearly independent columns. Then there exists an $n \times m$ matrix Q with orthonormal columns and an upper triangular, $m \times m$, invertible matrix R such that $A = Q \cdot R$. Moreover, Q and R are constructed as follows:

- Let $\{u_1, \dots, u_m\}$ be the columns of A with their order preserved so that

$$A = [u_1 \quad u_2 \quad \cdots \quad u_m]$$

- Q is constructed by applying the Gram-Schmidt Process to the columns of A and normalizing their lengths to one,

$$\{u_1, u_2, \dots, u_m\} \xrightarrow[\text{Process}]{\text{Gram-Schmidt}} \{v_1, v_2, \dots, v_m\}$$

$$Q := \left[\begin{array}{ccc} \frac{v_1}{\|v_1\|} & \frac{v_2}{\|v_2\|} & \cdots & \frac{v_m}{\|v_m\|} \end{array} \right]$$

- Because $Q^T Q = I_m$, it follows that $A = Q \cdot R \iff R := Q^T \cdot A$.
- Recalling that the columns of A are linearly independent, if, and only if $x = 0$ is the unique solution to $Ax = 0$, we have that

$$x = 0 \iff Ax = 0 \iff Q \cdot Rx = 0 \iff Q^T \cdot Q \cdot Rx = Q^T \cdot 0 \iff Rx = 0 \iff \det(R) \neq 0,$$

where the last step follows because R is square.

Remark: Because R is upper triangular, everything below its diagonal is zero. Hence, the calculation of R can be sped up by extracting its coefficients from the Gram-Schmidt Process instead of doing the indicated matrix multiplication. We explore this in HW.

Example 9.23 Compute the QR Factorization of $A = \begin{bmatrix} 1 & 1 & 0 \\ 1 & 2 & 1 \\ 0 & 3 & 1 \end{bmatrix}$.

Solution: We extract the columns of A and obtain

$$\{u_1, u_2, u_3\} = \left\{ \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix} \right\}$$

From Example 9.22, we have that

$$\left\{ \tilde{v}_1 = \frac{v_1}{\|v_1\|} = \frac{\sqrt{2}}{2} \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}, \tilde{v}_2 = \frac{v_2}{\|v_2\|} = \frac{\sqrt{38}}{38} \begin{bmatrix} -1 \\ 1 \\ 6 \end{bmatrix}, \tilde{v}_3 = \frac{v_3}{\|v_3\|} = \frac{\sqrt{19}}{19} \begin{bmatrix} -3 \\ 3 \\ -1 \end{bmatrix} \right\}$$

and therefore,

$$Q \approx \begin{bmatrix} 0.707107 & -0.162221 & -0.688247 \\ 0.707107 & 0.162221 & 0.688247 \\ 0.000000 & 0.973329 & -0.229416 \end{bmatrix}$$

and

$$R = Q^T \cdot A \approx \begin{bmatrix} 1.41421 & 2.12132 & 0.707107 \\ 0.00000 & 3.08221 & 1.13555 \\ 0.00000 & 0.00000 & 0.458831 \end{bmatrix}.$$

As a numerical check, we also compute how close Q^T is to being a matrix inverse of Q ,

$$Q^T \cdot Q - I = \begin{bmatrix} -2.22045e-16 & 9.71445e-17 & 1.11022e-16 \\ 9.71445e-17 & 0.00000e-17 & -2.49800e-16 \\ 1.11022e-16 & -2.49800e-16 & 0.00000e-17 \end{bmatrix}.$$

In addition, we check that $\det(Q) = -1.0$. ■

Suggested Pipeline: Solutions of Linear Equations via the QR Factorization

Suppose that A is $n \times n$ and its columns are linearly independent. Let $A = Q \cdot R$ be its QR Factorization. Then

$$(Ax = b) \iff (Q \cdot Rx = b) \iff (Rx = Q^\top b). \quad (9.19)$$

Hence, whenever $\det(A) \neq 0$, the suggested “pipeline” for solving $Ax = b$ is

- factor $A =: Q \cdot R$,
- compute $\bar{b} := Q^\top b$, and then
- solve $Rx = \bar{b}$ via back substitution.

Example 9.24 Use the suggested pipeline to solve the system of linear equations

$$\underbrace{\begin{bmatrix} 1 & 1 & 0 \\ 1 & 2 & 1 \\ 0 & 3 & 1 \end{bmatrix}}_A x = \underbrace{\begin{bmatrix} 1 \\ 4 \\ 7 \end{bmatrix}}_b.$$

Solution: From Example 9.23,

$$\underbrace{\begin{bmatrix} 1 & 1 & 0 \\ 1 & 2 & 1 \\ 0 & 3 & 1 \end{bmatrix}}_A = \underbrace{\begin{bmatrix} 0.707107 & -0.162221 & -0.688247 \\ 0.707107 & 0.162221 & 0.688247 \\ 0.000000 & 0.973329 & -0.229416 \end{bmatrix}}_Q \cdot \underbrace{\begin{bmatrix} 1.41421 & 2.12132 & 0.707107 \\ 0.00000 & 3.08221 & 1.13555 \\ 0.00000 & 0.00000 & 0.458831 \end{bmatrix}}_R.$$

We form

$$\bar{b} := Q^\top b = \begin{bmatrix} 3.53553 \\ 7.29996 \\ 0.45883 \end{bmatrix}$$

and then use back substitution to solve

$$\underbrace{\begin{bmatrix} 1.41421 & 2.12132 & 0.707107 \\ 0.00000 & 3.08221 & 1.13555 \\ 0.00000 & 0.00000 & 0.458831 \end{bmatrix}}_R x = \underbrace{\begin{bmatrix} 3.535534 \\ 7.299964 \\ 0.458831 \end{bmatrix}}_{\bar{b}},$$

which yields

$$x = \begin{bmatrix} -1 \\ 2 \\ 1 \end{bmatrix}.$$

Least Squares via the QR Factorization

Suppose that A is $n \times m$ and its columns are linearly independent (tall matrix). Let $A = Q \cdot R$ be its QR Factorization. Then $A^T A = R^T \cdot Q^T \cdot Q \cdot R = R^T \cdot R$ and thus

$$(A^T \cdot Ax = A^T b) \iff (R^T \cdot Rx = R^T \cdot Q^T b) \iff (Rx = Q^T b), \quad (9.20)$$

where we have used the fact that R is invertible. Hence, whenever the columns of A are linearly independent, the suggested “pipeline” for computing a least squared error solution to $Ax = b$ is

- factor $A =: Q \cdot R$,
- compute $\bar{b} := Q^T b$, and then
- solve $Rx = \bar{b}$ via back substitution.

Yes! The two pipelines are identical!! Your surprise will be tempered when you go back to our discussion of least squared error solutions of linear equations where we noted that if A is square and its determinant is non-zero, then

$$(Ax = b) \iff (A^T \cdot Ax = A^T b).$$

The only difference in the two cases is that when A is square, Q is an orthogonal matrix whereas, when A is a tall matrix, then Q is an orthonormal matrix. Yes, it’s a subtle difference! Is it an important difference? Not really, as long as you only form $Q^T \cdot Q$ and you avoid $Q \cdot Q^T$.

Example 9.25 We rework Example 8.1 from Chapter 8.2, where we seek a least squared error solution to the system of linear equations

$$\underbrace{\begin{bmatrix} 1.0 & 1.0 \\ 2.0 & 1.0 \\ 4.0 & 1.0 \\ 5.0 & 1.0 \\ 7.0 & 1.0 \end{bmatrix}}_A \underbrace{\begin{bmatrix} x_1 \\ x_2 \end{bmatrix}}_x = \underbrace{\begin{bmatrix} 4 \\ 8 \\ 10 \\ 12 \\ 18 \end{bmatrix}}_b. \quad (9.21)$$

Solution: Since the columns of A are linearly independent, we compute the QR factorization of A and obtain

$$Q = \begin{bmatrix} 0.102598 & 0.730297 \\ 0.205196 & 0.547723 \\ 0.410391 & 0.182574 \\ 0.512989 & 0.000000 \\ 0.718185 & -0.365148 \end{bmatrix}, R = \begin{bmatrix} 9.74679 & 1.94936 \\ 0.0 & 1.09545 \end{bmatrix}, \text{ and } Q^T \cdot b = \begin{bmatrix} 25.23906 \\ 2.556038 \end{bmatrix}.$$

We then use back substitution to solve

$$\begin{bmatrix} 9.74679 & 1.94936 \\ 0.0 & 1.09545 \end{bmatrix} x = \begin{bmatrix} 25.23906 \\ 2.556038 \end{bmatrix}$$

which yields

$$x^* = \begin{bmatrix} 2.1228 \\ 2.3333 \end{bmatrix}.$$

For small made-up problems like this one, there is no real numerical advantage to using a sophisticated solution like our “suggested pipeline”. In real engineering, it makes a huge difference. Prof. Grizzle’s and Ghaffari’s students use the suggested pipeline, either as given with QR or its equivalent via LU, when working with Cassie Blue. ■

LU vs QR: Which is Better?

The answer seems to be highly situational. One of the giants of Linear Algebra education, Prof. Gilbert Strang, writes “The LU vs QR choice comes up for example with least squares equations $A^T \cdot Ax = A^T b$. If we actually form $A^T \cdot A$ and solve by LU, it is a bit faster than using $A = QR$, but less [numerically] stable : We could do $R^T \cdot Q^T \cdot Q \cdot Rx = R^T \cdot Q^T b$, which is just $Rx = Q^T b$ and [is] more stable. Cleve Moler [a co-founder of the Mathworks] and I are testing various methods for underdetermined systems with $n \gg m$ (because deep learning often has more parameters than equations to determine them, and it still works).”

9.9 Underdetermined Equations or What to do When $Ax=b$ has an Infinite Number of Solutions?

We consider $Ax = b$ and recall what we know about its solutions:

- $b \in \text{col span}\{A\} \iff$ a solution exists;
- the solution is unique if, and only if, the columns of A are linearly independent; and thus
- if there exists one solution and the columns of A are linearly dependent, then there exist an infinite number of solutions.

Underdetermined Equations

The columns of A will be linearly dependent when $Ax = b$ has fewer equations than unknowns. In other words, A is $n \times m$ and $m > n$; we’ve been calling these wide matrices: more columns than rows. When dealing with an equation $Ax = b$ with fewer equations than unknowns, one says that it is **underdetermined**. Why? Because, to determine x uniquely, at a minimum, we need as many equations as unknowns.

Is there a difference between being underdetermined and having an infinite number of solutions? Yes. It’s possible to be underdetermined and have no solution at all when $b \notin \text{col span}\{A\}$. If the rows of A are linearly independent, then

$$Ax = b \text{ is underdetermined} \iff Ax = b \text{ has an infinite number of solutions.}$$

The rows of A being linearly independent is equivalent to the columns of A^T being linearly independent.

When $Ax = b$ has an infinite number of solutions, is there a way that we can make one of them appear to be more interesting, more special, or just flat out “better” than all the other solutions? Is there a property that we could associate with each solution and optimize our choice of solution with respect to that property? The most common approach is to choose the solution with minimum norm!

Minimum Norm Solution to Underdetermined Equations

Consider an underdetermined system of linear equations $Ax = b$. If the rows of A are linearly independent (equivalently, the columns of A^\top are linearly independent), then

$$x^* = \arg \min_{Ax=b} \|x\| \iff x^* = A^\top \cdot (A \cdot A^\top)^{-1} b \iff x^* = A^\top \alpha \text{ and } A \cdot A^\top \alpha = b. \quad (9.22)$$

We recommend that the minimum norm solution x^* be computed with the right-hand side of (9.22) so that the matrix inverse is avoided, but for small problems, the middle answer is fine.

Suppose we do the QR Factorization of A^\top instead of A itself, so that

$$A^\top = Q \cdot R.$$

Because the columns of A^\top are linearly independent, R is square and invertible. It follows that $A = R^\top \cdot Q^\top$ and $A \cdot A^\top = R^\top \cdot R$ because $Q^\top \cdot Q = I$. Using these facts, (9.22) can be rewritten as

$$x^* = \arg \min_{Ax=b} \|x\| \iff x^* = Q \cdot (R^\top)^{-1} b \iff x^* = Q\beta \text{ and } R^\top \beta = b. \quad (9.23)$$

We note that R^\top is lower triangular, and thus $R^\top \beta = b$ can be solved via forward substitution. Hence, our suggested “pipeline” for underdetermined problems $Ax = b$ is

- Check that the columns of A^\top are linearly independent and compute $A^\top = Q \cdot R$.
- Solve $R^\top \beta = b$ by forward substitution.
- $x^* = Q\beta$.

Remark: In case you are curious, β in (9.23) is related to α in (9.22) by $\beta = R\alpha$. The two x^* are the same!

Example 9.26 Use the suggested pipeline to determine a minimum norm solution to the system of underdetermined equations

$$\underbrace{\begin{bmatrix} 1 & 1 & 0 \\ 1 & 2 & 1 \end{bmatrix}}_A x = \underbrace{\begin{bmatrix} 1 \\ 4 \end{bmatrix}}_b.$$

Solution: The columns of A^\top are linearly independent and we compute the QR Factorization to be

$$A^\top = \begin{bmatrix} 1 & 1 \\ 1 & 2 \\ 0 & 1 \end{bmatrix} = \underbrace{\begin{bmatrix} 0.707107 & -0.408248 \\ 0.707107 & 0.408248 \\ 0.000000 & 0.816497 \end{bmatrix}}_Q \cdot \underbrace{\begin{bmatrix} 1.41421 & 2.12132 \\ 0.00000 & 1.22474 \end{bmatrix}}_R.$$

We solve $R^\top \beta = b$ and obtain

$$\beta = \begin{bmatrix} 0.70711 \\ 2.04124 \end{bmatrix},$$

and then $x^* = Q\beta$ to arrive at the final answer

$$x^* = \begin{bmatrix} -0.33333 \\ 1.33333 \\ 1.66667 \end{bmatrix}.$$

To verify that x^* is indeed a solution, we substitute it into $Ax - b$ and obtain

$$Ax^* - b = \begin{bmatrix} 0.000000 \\ -4.441e-16 \end{bmatrix},$$

which looks like a pretty good solution! ■

A Quick Check

Is x^* in Example 9.26 the solution of smallest norm? Of course! Don't you trust us? To which you respond, "of course not!"

Let's see about that. What are other solutions? We claim they all have the form $x^* + \bar{x}$, where $A\bar{x} = 0$, because then

$$A(x^* + \bar{x}) = Ax^* + A\bar{x} = b + 0 = b.$$

Indeed, we compute that all solutions to $Ax = 0$ have the form

$$\bar{x} = \gamma \begin{bmatrix} 1 \\ -1 \\ 1 \end{bmatrix},$$

and therefore, all solutions to $Ax = b$ have the form

$$x_{\text{sol}} = x^* + \bar{x}.$$

The next thing we can check is that for all $\gamma \in \mathbb{R}$, $x^* \perp \bar{x}$, and hence, by the Pythagorean Theorem, we have that

$$\|x_{\text{sol}}\|^2 = \|x^* + \bar{x}\|^2 = \|x^*\|^2 + \|\bar{x}\|^2 = \|x^*\|^2 + 3\gamma^2.$$

It follows that

$$\min_{\gamma} \|x_{\text{sol}}\|^2 = \min_{\gamma} (\|x^*\|^2 + 3\gamma^2) = \|x^*\|^2 + \min_{\gamma} (3\gamma^2)$$

and thus the minimum occurs for $\gamma = 0$. Hence, x^* is indeed the minimum norm solution!

Optional Read: The Pythagorean Theorem is a powerful ally when one seeks to establish minimum norm properties. We use it to show that (9.22) has the claimed minimum norm property among all solutions of $Ax = b$. All of the ideas are actually present in our analysis of Example 9.26. Here we sketch the general case.

The proposed minimum norm solution to $Ax = b$ has the form $x^* = A^T \alpha$, which is a linear combination of the columns of A^T . Indeed, for A an $n \times m$ matrix we write

$$A = \begin{bmatrix} a_1^{\text{row}} \\ \vdots \\ a_n^{\text{row}} \end{bmatrix} \quad \text{and} \quad \alpha = \begin{bmatrix} \alpha_1 \\ \vdots \\ \alpha_n \end{bmatrix} \quad \text{so that} \quad A^T \alpha = \begin{bmatrix} (a_1^{\text{row}})^T & \cdots & (a_n^{\text{row}})^T \end{bmatrix} \begin{bmatrix} \alpha_1 \\ \vdots \\ \alpha_n \end{bmatrix} = \alpha_1 (a_1^{\text{row}})^T + \cdots + \alpha_n (a_n^{\text{row}})^T.$$

We next note that $A\bar{x} = 0$ if, and only if, for all $1 \leq i \leq n$, $a_i^{\text{row}} \bar{x} = 0$, which is equivalent to $(a_i^{\text{row}})^T \perp \bar{x}$. Hence, we have that

$$x^* \perp \bar{x}$$

A general solution to $Ax = b$ can be written as $x_{\text{sol}} = x^* + \bar{x}$, where \bar{x} is any solution to $Ax = 0$. Applying the Pythagorean Theorem we have

$$\|x_{\text{sol}}\|^2 = \|x^* + \bar{x}\|^2 = \|x^*\|^2 + \|\bar{x}\|^2.$$

Because $\|x^*\|^2 + \|\bar{x}\|^2$ is smallest for $\bar{x} = 0$, it follows that

$$\min_{\bar{x}} \|x_{\text{sol}}\|^2 = \min_{\bar{x}} (\|x^*\|^2 + \|\bar{x}\|^2) = \|x^*\|^2,$$

which shows that (9.22) really is the minimum norm solution.

9.10 Steering a Mobile Robot as a Practical Example of an Underdetermined System of Linear Equations

We first develop a model of a mobile robot. We assume that the robot moves in \mathbb{R}^2 with its x -position denoted p^x and y -position denoted p^y . We gather these two coordinates together and write them as a vector

$$p := \begin{bmatrix} p^x \\ p^y \end{bmatrix}, \quad (9.24)$$

which is typically called the **state of the robot**. Because we are modeling a mobile robot, its position/state changes with time. For reasons explained in Appendix B of our textbook, we discretize time into uniform samples. So, we let $\delta t > 0$ be some base unit or duration of time, typically small, and define $t_0, t_1, t_2, \dots, t_k, \dots$, where $t_k := k\delta t$. We denote the robot's position at time t_k by

$$p_k := \begin{bmatrix} p_k^x \\ p_k^y \end{bmatrix}; \quad (9.25)$$

in other words, the subscript k keeps track of time. With this notation, p_0 will be the initial position of the robot in the plane (i.e., in \mathbb{R}^2) at time t_0 .



Figure 9.4: Two mobile robots. (a) Michigan M-bot as used in ROB 103, ROB 320, ROB 330, and ROB 550. (b) iRobot's Roomba vacuum cleaning robot.

The next thing we posit is that our mobile robot has “dynamics”, meaning that its state at time t_{k+1} can be expressed as a function of its state at time t_k and any motor commands that we provide. Because we are studying linear algebra, we assume that

$$p_{k+1} = Ap_k + Bu_k, \quad (9.26)$$

where $u_k \in \mathbb{R}^2$ is a pair of motor commands, the 2×2 matrix B distributes the motor commands to effect motion of the robot (that is, change its next position), and A is a 2×2 matrix that in a realistic model would capture the mass of the robot, the inertia of rotating parts, and other effects from physics.

Even though the following values are not so realistic, we'll go ahead and assume that

$$\begin{aligned} \delta t &= 0.1 \\ A &= I_{2 \times 2} + \delta t \begin{bmatrix} 0.0 & -0.5 \\ 0.5 & 0.0 \end{bmatrix} \\ B &= \delta t I_{2 \times 2}, \end{aligned} \quad (9.27)$$

so that we can focus on the process of steering the robot and not the complexity of the robot's model.




Figure 9.5: Controlling the motion of a mobile robot. (a) shows the evolution of the robot if we apply motor commands that are identically zero. With our chosen model, the robot wanders around like a Roomba. (b) shows us deliberately steering the robot to the origin using methods from Linear Algebra. The problem turns out to be one of an underdetermined system of linear equations.

Why do we want a model? So that we can predict the future behavior of the robot. Not only do we have a way to compute the state at the next time instant based on the current state and current input, we can also iterate forward and **predict** the state at some future time, say N , as a function of a hypothesized input sequence, $\{u_0, u_1, \dots, u_{N-1}\}$ as follows

$$\begin{aligned}
 p_0 &= \text{given initial position} \\
 p_1 &= Ap_0 + Bu_0 \\
 p_2 &= Ap_1 + Bu_1 = A(Ap_0 + Bu_0) + Bu_1 = A^2p_0 + ABu_0 + Bu_1 \\
 p_3 &= Ap_2 + Bu_2 = A(A^2p_0 + ABu_0 + Bu_1) + Bu_2 = A^3p_0 + A^2Bu_0 + ABu_1 + Bu_2 \\
 &\vdots \\
 p_N &= A^N p_0 + A^{N-1}Bu_0 + A^{N-2}Bu_1 + \dots + ABu_{N-2} + Bu_{N-1}
 \end{aligned}$$

That's a lot of symbols, but what do they tell us? Suppose we start the robot at

$$p_0 := \begin{bmatrix} p_0^x \\ p_0^y \end{bmatrix} = \begin{bmatrix} 1.0 \\ 1.0 \end{bmatrix}$$

and we set the motor commands to be identically zero. Then we can predict how the robot will move "on its own". Figure 9.5-(a) shows the evolution of our robot for $0 \leq t_k \leq 20$ seconds, that is, $0 \leq k \leq 200$ (because $\delta t = 0.1$). When we are not actively modifying its trajectory, our robot is a bit like a wandering Roomba, spiraling outward trying to get its bearings!

Let's suppose we want to actively steer the robot to the origin in 2 seconds, that is,

$$p_N := \begin{bmatrix} p_N^x \\ p_N^y \end{bmatrix} = \begin{bmatrix} 0.0 \\ 0.0 \end{bmatrix},$$

where $N = 20$. Then we are seeking a solution to the equation

$$\begin{aligned}
 p_N &= A^N p_0 + A^{N-1}Bu_0 + A^{N-2}Bu_1 + \dots + ABu_{N-2} + Bu_{N-1} \\
 &\Downarrow \\
 p_N - A^N p_0 &= A^{N-1}Bu_0 + A^{N-2}Bu_1 + \dots + ABu_{N-2} + Bu_{N-1} \\
 &\Downarrow \\
 p_N - A^N p_0 &= \begin{bmatrix} A^{N-1}B & A^{N-2}B & \dots & AB & B \end{bmatrix} \begin{bmatrix} u_0 \\ u_1 \\ \vdots \\ u_{N-2} \\ u_{N-1} \end{bmatrix}, \tag{9.28}
 \end{aligned}$$

where the unknowns are the control decisions $u_{\text{seq}} := (u_0, u_1, \dots, u_{N-2}, u_{N-1})$. For $N = 20$, we have 40 control values to compute, because each $u_k \in \mathbb{R}^2$,

$$u_{\text{seq}} := \begin{bmatrix} u_0 \\ u_1 \\ \vdots \\ u_{N-2} \\ u_{N-1} \end{bmatrix} \in \mathbb{R}^{2N}. \quad (9.29)$$

The problem is clearly **underdetermined** because we have two equations and $2N = 40$ unknowns. If we view the Euclidean norm of the control sequence as a measure of **control effort** (perhaps, energy drawn from a battery to operate motors on the robot), then it makes sense to solve for

$$u_{\text{seq}}^* := \arg \min_{Mu_{\text{seq}} = (p_N - Sp_0)} \|u_{\text{seq}}\|^2, \quad (9.30)$$

where $M := [A^{N-1}B \quad A^{N-2}B \quad \dots \quad AB \quad B]$ and $S := A^N$.

Minimum Norm Solution to Steering a Robot

The solution to the minimum norm-squared problem in (9.30) was given by (9.22) in the previous big green box! Some of you may see that immediately. For those who don't, we'll relate the "generic" notation used in finding the minimum norm solution of $Ax = b$ to the problem at hand.

Correspondence between the variables:

$$\begin{aligned} x^* = \arg \min_{Ax=b} \|x\|^2 &\longleftrightarrow u_{\text{seq}}^* := \arg \min_{Mu_{\text{seq}} = p_N - Sp_0} \|u_{\text{seq}}\|^2 \\ A &\longleftrightarrow M \\ b &\longleftrightarrow p_N - Sp_0 \\ x, x^* &\longleftrightarrow u_{\text{seq}}, u_{\text{seq}}^* \end{aligned} \quad (9.31)$$

Correspondence between the solutions:

$$\begin{aligned} x^* = A^\top \cdot (A \cdot A^\top)^{-1} b &\longleftrightarrow u_{\text{seq}}^* = M^\top \cdot (M \cdot M^\top)^{-1} (p_N - Sp_0) \text{ for small problems, and} \\ x^* = A^\top \alpha \text{ and } A \cdot A^\top \alpha = b &\longleftrightarrow u_{\text{seq}}^* = M^\top \alpha \text{ and } M \cdot M^\top \alpha = (p_N - Sp_0) \text{ for larger problems.} \end{aligned} \quad (9.32)$$

Solution via the QR Factorization pipeline:

- Check that the columns of M^\top are linearly independent and compute $M^\top = Q \cdot R$.
- Solve $R^\top \beta = (p_N - Sp_0)$ by forward substitution.
- $u_{\text{seq}}^* = Q\beta$.

Just for fun, we'll print out on the next page the resulting optimal control sequence for $N = 20$, namely

$$u_{\text{seq}}^* = \begin{bmatrix} -0.4864 \\ -0.5375 \\ -0.4583 \\ -0.5605 \\ -0.4292 \\ -0.5819 \\ -0.3991 \\ -0.6019 \\ -0.3681 \\ -0.6203 \\ -0.3363 \\ -0.6371 \\ -0.3037 \\ -0.6523 \\ -0.2704 \\ -0.6658 \\ -0.2365 \\ -0.6776 \\ -0.2021 \\ -0.6877 \\ -0.1673 \\ -0.6961 \\ -0.1322 \\ -0.7027 \\ -0.0968 \\ -0.7075 \\ -0.0612 \\ -0.7106 \\ -0.0257 \\ -0.7119 \\ 0.0099 \\ -0.7114 \\ 0.0454 \\ -0.7091 \\ 0.0806 \\ -0.7051 \\ 0.1156 \\ -0.6993 \\ 0.1502 \\ -0.6918 \end{bmatrix}. \quad (9.33)$$

Figure 9.5-(b) shows the evolution of the robot's trajectory in \mathbb{R}^2 as we steer it efficiently to the origin in 2 seconds. Do you think you could do that by hand? Not a chance!

It is possible to steer the robot to the origin in 0.1 seconds. The control sequence is

$$u_{\text{short seq}}^* = \begin{bmatrix} -9.5 \\ -10.5 \end{bmatrix}.$$

Its norm squared is 200.5, while the norm squared of the longer control sequence in (9.33) is 10.26, twenty times smaller. Just for the fun of it, we let the controller have 20 seconds to reach the origin, and then the norm squared of the control sequence drops to 1.27, a further factor of eight smaller. At 2,000 seconds, the norm squared of $u_{\text{very long seq}}^*$ plateaus at 0.5. As an engineer, we would need to trade off speed of response (how quickly we reach a goal state) versus how much it costs us to reach the goal in a given interval of time. Traveling $\sqrt{2}$ meters in 2 seconds is a pretty good pace without being ridiculously expensive.

9.11 (Optional Read): In the QR Factorization, Why R is Upper Triangular and How to Efficiently Obtain its Coefficients from the Gram-Schmidt Process

Going back to Example 9.22, we were given that the set $\{u_1, u_2, u_3\}$ obtained from the columns of A was linearly independent. Applying Gram-Schmidt and normalizing gave the columns of Q , $\{\frac{v_1}{\|v_1\|}, \frac{v_2}{\|v_2\|}, \frac{v_3}{\|v_3\|}\}$, which form an orthonormal basis for $\text{span}\{u_1, u_2, u_3\}$. Moreover, Gram-Schmidt naturally gives us a triangular relationship among the two sets of linearly independent vectors

$$\begin{aligned}\text{span}\{u_1\} &= \text{span}\left\{\frac{v_1}{\|v_1\|}\right\} \\ \text{span}\{u_1, u_2\} &= \text{span}\left\{\frac{v_1}{\|v_1\|}, \frac{v_2}{\|v_2\|}\right\} \\ \text{span}\{u_1, u_2, u_3\} &= \text{span}\left\{\frac{v_1}{\|v_1\|}, \frac{v_2}{\|v_2\|}, \frac{v_3}{\|v_3\|}\right\}.\end{aligned}$$

The triangular structure of R is a reflection of this triangular relationship between the columns of A and the columns of Q . In particular, we can write u_1 as a linear combination of $\frac{v_1}{\|v_1\|}$, u_2 as a linear combination of $\frac{v_1}{\|v_1\|}$ and $\frac{v_2}{\|v_2\|}$, and finally, u_3 as a linear combination of $\frac{v_1}{\|v_1\|}$, $\frac{v_2}{\|v_2\|}$, and $\frac{v_3}{\|v_3\|}$. If we use r_{ij} to denote the coefficients in the linear combinations, we end up with

$$\begin{aligned}u_1 &= r_{11} \frac{v_1}{\|v_1\|} \\ u_2 &= r_{12} \frac{v_1}{\|v_1\|} + r_{22} \frac{v_2}{\|v_2\|} \\ u_3 &= r_{13} \frac{v_1}{\|v_1\|} + r_{23} \frac{v_2}{\|v_2\|} + r_{33} \frac{v_3}{\|v_3\|}.\end{aligned}$$

Writing this out in matrix form then gives $A = Q \cdot R$,

$$\begin{aligned}\underbrace{\begin{bmatrix} u_1 & u_2 & u_3 \end{bmatrix}}_A &= \underbrace{\begin{bmatrix} r_{11} \frac{v_1}{\|v_1\|} & r_{12} \frac{v_1}{\|v_1\|} + r_{22} \frac{v_2}{\|v_2\|} & r_{13} \frac{v_1}{\|v_1\|} + r_{23} \frac{v_2}{\|v_2\|} + r_{33} \frac{v_3}{\|v_3\|} \end{bmatrix}}_{Q \cdot R} \\ &= \underbrace{\begin{bmatrix} \frac{v_1}{\|v_1\|} & \frac{v_2}{\|v_2\|} & \frac{v_3}{\|v_3\|} \end{bmatrix}}_Q \cdot \underbrace{\begin{bmatrix} r_{11} & r_{12} & r_{13} \\ 0 & r_{22} & r_{23} \\ 0 & 0 & r_{33} \end{bmatrix}}_R\end{aligned}$$

In case that last step was too much, too fast, we break it down into Q multiplying the various columns of R ,

$$\begin{aligned}\underbrace{\begin{bmatrix} \frac{v_1}{\|v_1\|} & \frac{v_2}{\|v_2\|} & \frac{v_3}{\|v_3\|} \end{bmatrix}}_Q \cdot \begin{bmatrix} r_{11} \\ 0 \\ 0 \end{bmatrix} &= r_{11} \frac{v_1}{\|v_1\|} \\ \underbrace{\begin{bmatrix} \frac{v_1}{\|v_1\|} & \frac{v_2}{\|v_2\|} & \frac{v_3}{\|v_3\|} \end{bmatrix}}_Q \cdot \begin{bmatrix} r_{12} \\ r_{22} \\ 0 \end{bmatrix} &= r_{12} \frac{v_1}{\|v_1\|} + r_{22} \frac{v_2}{\|v_2\|} \\ \underbrace{\begin{bmatrix} \frac{v_1}{\|v_1\|} & \frac{v_2}{\|v_2\|} & \frac{v_3}{\|v_3\|} \end{bmatrix}}_Q \cdot \begin{bmatrix} r_{13} \\ r_{23} \\ r_{33} \end{bmatrix} &= r_{13} \frac{v_1}{\|v_1\|} + r_{23} \frac{v_2}{\|v_2\|} + r_{33} \frac{v_3}{\|v_3\|}.\end{aligned}$$

More Efficient QR Factorization by Reading R Directly from Gram-Schmidt

Suppose that the columns of $A = [u_1 \ u_2 \ \cdots \ u_m]$ are linearly independent. We can then re-arrange the steps of the Gram-Schmidt Process (9.16) and introduce normalization to obtain

$$\begin{aligned}
 u_1 &= \|v_1\| \frac{v_1}{\|v_1\|} \\
 u_2 &= \left(\frac{u_2 \bullet v_1}{v_1 \bullet v_1} \right) \|v_1\| \frac{v_1}{\|v_1\|} + \|v_2\| \frac{v_2}{\|v_2\|} \\
 u_3 &= \left(\frac{u_3 \bullet v_1}{v_1 \bullet v_1} \right) \|v_1\| \frac{v_1}{\|v_1\|} + \left(\frac{u_3 \bullet v_2}{v_2 \bullet v_2} \right) \|v_2\| \frac{v_2}{\|v_2\|} + \|v_3\| \frac{v_3}{\|v_3\|} \\
 &\vdots \\
 u_k &= \sum_{i=1}^{k-1} \left(\frac{u_k \bullet v_i}{v_i \bullet v_i} \right) \|v_i\| \frac{v_i}{\|v_i\|} + \|v_k\| \frac{v_k}{\|v_k\|}, \quad 3 \leq k \leq m.
 \end{aligned} \tag{9.34}$$

Recognizing that $Q := \left[\begin{array}{ccc} \frac{v_1}{\|v_1\|} & \frac{v_2}{\|v_2\|} & \cdots & \frac{v_m}{\|v_m\|} \end{array} \right]$, we identify that

$$\begin{aligned}
 u_1 &= \underbrace{\|v_1\|}_{r_{11}} \frac{v_1}{\|v_1\|} \\
 u_2 &= \underbrace{\left(\frac{u_2 \bullet v_1}{v_1 \bullet v_1} \right) \|v_1\|}_{r_{12}} \frac{v_1}{\|v_1\|} + \underbrace{\|v_2\|}_{r_{22}} \frac{v_2}{\|v_2\|} \\
 u_3 &= \underbrace{\left(\frac{u_3 \bullet v_1}{v_1 \bullet v_1} \right) \|v_1\|}_{r_{13}} \frac{v_1}{\|v_1\|} + \underbrace{\left(\frac{u_3 \bullet v_2}{v_2 \bullet v_2} \right) \|v_2\|}_{r_{23}} \frac{v_2}{\|v_2\|} + \underbrace{\|v_3\|}_{r_{33}} \frac{v_3}{\|v_3\|} \\
 &\vdots \\
 u_k &= \sum_{i=1}^{k-1} \underbrace{\left(\frac{u_k \bullet v_i}{v_i \bullet v_i} \right) \|v_i\|}_{r_{ik}} \frac{v_i}{\|v_i\|} + \underbrace{\|v_k\|}_{r_{kk}} \frac{v_k}{\|v_k\|}, \quad 3 \leq k \leq m.
 \end{aligned} \tag{9.35}$$

Hence, for $1 \leq i, j \leq m$

$$r_{ij} = \begin{cases} 0 & i > j \\ \|v_i\| & i = j \\ \left(\frac{u_j \bullet v_i}{v_i \bullet v_i} \right) \|v_i\| & i < j \end{cases}$$

9.12 (Optional Read): Modified Gram-Schmidt Algorithm

The classical Gram-Schmidt Process is straightforward to understand, which is why it is taught in courses. Unfortunately, it behaves poorly under the round-off error that occurs in digital computations! Here is a standard example:

$$u_1 = \begin{bmatrix} 1 \\ \varepsilon \\ 0 \\ 0 \end{bmatrix}, \quad u_2 = \begin{bmatrix} 1 \\ 0 \\ \varepsilon \\ 0 \end{bmatrix}, \quad u_3 = \begin{bmatrix} 1 \\ 0 \\ 0 \\ \varepsilon \end{bmatrix}, \quad \varepsilon > 0$$

Let $\{e_1, e_2, e_3, e_4\}$ be the standard basis vectors corresponding to the columns of the 4×4 identity matrix. We note that

$$\begin{aligned}
 u_2 &= u_1 + \varepsilon(e_3 - e_2) \\
 u_3 &= u_2 + \varepsilon(e_4 - e_3)
 \end{aligned}$$

and thus, for $\epsilon \neq 0$,

$$\begin{aligned}\text{span}\{u_1, u_2\} &= \text{span}\{u_1, (e_3 - e_2)\} \\ \text{span}\{u_1, u_2, u_3\} &= \text{span}\{u_1, (e_3 - e_2), (e_4 - e_3)\}\end{aligned}$$

Hence, Gram-Schmidt applied to $\{u_1, u_2, u_3\}$ and $\{u_1, (e_3 - e_2), (e_4 - e_3)\}$ should “theoretically” produce the same orthonormal vectors. To check this, we go to Julia, and for $\epsilon = 0.1$, we do indeed get the same results. You can verify this yourself. However, with $\epsilon = 10^{-8}$,

$$Q_1 = \begin{bmatrix} 1.0000 & 0.0000 & 0.0000 \\ 0.0000 & -0.7071 & -0.7071 \\ 0.0000 & 0.7071 & 0.0000 \\ 0.0000 & 0.0000 & 0.7071 \end{bmatrix}$$

$$Q_2 = \begin{bmatrix} 1.0000 & 0.0000 & 0.0000 \\ 0.0000 & -0.7071 & -0.4082 \\ 0.0000 & 0.7071 & -0.4082 \\ 0.0000 & 0.0000 & 0.8165 \end{bmatrix}$$

where

$$Q_1 = \begin{bmatrix} \frac{v_1}{\|v_1\|} & \frac{v_2}{\|v_2\|} & \frac{v_3}{\|v_3\|} \end{bmatrix}$$

has been computed with Classical-Gram-Schmidt for $\{u_1, u_2, u_3\}$ while

$$Q_2 = \begin{bmatrix} \frac{v_1}{\|v_1\|} & \frac{v_2}{\|v_2\|} & \frac{v_3}{\|v_3\|} \end{bmatrix}$$

has been computed with Classical-Gram-Schmidt for $\{u_1, (e_3 - e_2), (e_4 - e_3)\}$. Hence we do NOT obtain the same result!

Modified Gram-Schmidt has better Numerical Performance

for $k = 1 : n$

$v_k = u_k$ #copy over the vectors

end

for $k = 1 : n$

$v_k = \frac{v_k}{\|v_k\|}$

for $j = (k + 1) : n$

$v_j = v_j - (v_j \bullet v_k)v_k$ #Makes v_j orthogonal to v_k

end

end

At **Step 1**, v_1 is normalized to length one, and then v_2, \dots, v_n are redefined to be orthogonal to v_1 . At **Step 2**: v_2 is normalized to length one, and then v_3, \dots, v_n are redefined to be orthogonal to v_2 . We note that they were already orthogonal to v_1 . At **Step k** : v_k is normalized to length one, and then v_{k+1}, \dots, v_n are redefined to be orthogonal to v_k . We note that they were already orthogonal to v_1, \dots, v_{k-1} .

Hence, if Modified Gram-Schmidt is so great, when applied to $\{u_1, u_2, u_3\}$ and $\{u_1, (e_3 - e_2), (e_4 - e_3)\}$, it should produce the same orthonormal vectors and it does! To check this, we go to Julia for $\varepsilon = 10^{-8}$ and obtain

$$Q_1 = \begin{bmatrix} 1.0000 & 0.0000 & 0.0000 \\ 0.0000 & -0.7071 & -0.7071 \\ 0.0000 & 0.7071 & 0.0000 \\ 0.0000 & 0.0000 & 0.7071 \end{bmatrix}$$

$$Q_2 = \begin{bmatrix} 1.0000 & 0.0000 & 0.0000 \\ 0.0000 & -0.7071 & -0.7071 \\ 0.0000 & 0.7071 & 0.0000 \\ 0.0000 & 0.0000 & 0.7071 \end{bmatrix}$$

where Q_1 and Q_2 are defined above. **When one is equipped with the right Algorithm, the world is truly a marvelous place.**

9.13 (Optional Read) Source of the Definition of Orthogonal Vectors




Figure 9.6: How do we go from right triangles to orthogonal vectors satisfying $v_1 \bullet v_2 = 0$?

Earlier in the Chapter, we defined two vectors to be orthogonal if their dot product was zero. That's fine, we can make any definition we want, but does this really correspond to our notion of perpendicular vectors? It does, and we can prove that here if you accept that the triangle in Fig. 9.6-(a) is a right triangle if, and only if, the Pythagorean Theorem holds, that is, $a^2 + b^2 = c^2 \iff$ the triple (a, b, c) forms a right triangle. In the following we take this as a given. You may need to go back to a High School Geometry book to find this fact.

Fig. 9.6-(b) interprets the three sides of the triangle in terms of vectors v_1, v_2, v_3 and their norms. We seek to understand the relation that must hold between these vectors for the Pythagorean Theorem to hold. We note that

$$\begin{aligned} \|v_1 + v_2\|^2 &:= (v_1 + v_2)^\top (v_1 + v_2) \\ &= v_1^\top v_1 + v_1^\top v_2 + v_2^\top v_1 + v_2^\top v_2 \\ &\quad v_1^\top v_1 + v_2^\top v_2 + 2v_1^\top v_2 \\ &=: \|v_1\|^2 + \|v_2\|^2 + 2v_1 \bullet v_2. \end{aligned} \tag{9.36}$$

Hence,

$$\text{Pythagorean Theorem holds} \iff (\|v_1 + v_2\|^2 = \|v_1\|^2 + \|v_2\|^2) \iff v_1 \bullet v_2 = 0 \iff v_1 \perp v_2,$$

which is what we wanted to show!

How general do you want to go?

The notions of inner products and orthogonality can be greatly extended. We have merely scratched the surface here. ROB 501 explores these topics in great detail. You may also enjoy this YouTube video by Michael Penn: https://youtu.be/Dz_tsaocWek. He has a massive channel full of Math videos: <https://www.youtube.com/@MichaelPennMath>.

9.14 Looking Ahead

We will complete our introduction to Linear Algebra by introducing you to eigenvalues and eigenvectors. We'll also apply to matrices the concepts of subspace and dimension. This will give us a more complete understanding of solutions to systems of linear equations.

Chapter 10

The Vector Space \mathbb{R}^n : Part 3

Learning Objectives

- Learn how to define coordinates in a subspace of \mathbb{R}^n and understand how many coordinates you need.
- An introduction to eigenvalues and eigenvectors of square matrices.
- Applying to matrices some of the essential concepts in Linear Algebra.

Outcomes

- Basis vectors, dimension, and coordinates
- Eigenvalues, eigenvectors, and understanding when when eigenvectors provide a basis of \mathbb{R}^n .
- Range of a matrix and its relation to column span and null space.
- Handy matrix properties dealing with rank and nullity.

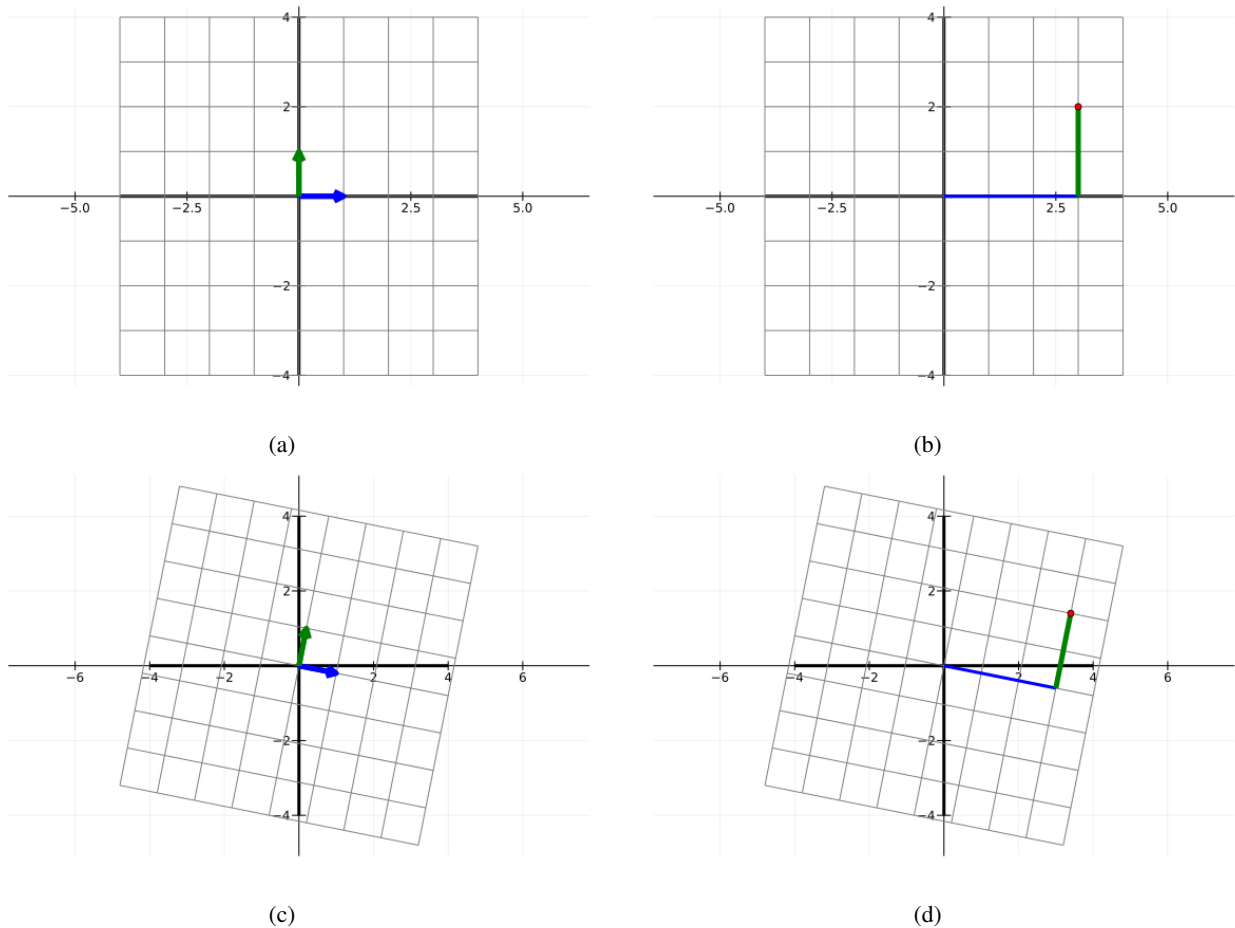


Figure 10.1: Basis vectors $\{u, v\}$ (blue and green) for \mathbb{R}^2 with the corresponding coordinates. The graphs (a) and (b) show coordinates on \mathbb{R}^2 with the natural basis vectors, $\{u = e_1, v = e_2\}$. The red dot in (b) is the point $3u + 2v$, that is, the point with coordinates $(3, 2)$ in the basis $\{u, v\}$. The graphs (c) and (d) show coordinates on \mathbb{R}^2 with basis vectors, $\{u = [1.0, -0.2]^\top, v = [0.2, 1.0]^\top\}$. The red dot in (b) is the point $3u + 2v$, that is, the point with coordinates $(3, 2)$ in the basis $\{u, v\}$. You go along the u -axis for three units and then follow the v -axis for two units. That is what the point $(3, 2)$ means in a basis $\{u, v\}$. The vectors in the second basis are still orthogonal, because $u \bullet v = 0$. They are rotated a few degrees clockwise with respect to the natural basis, however, and their lengths are not equal to one. When we study eigenvectors, you'll see a clear motivation for using bases on \mathbb{R}^n that are distinct from the natural basis. For now, we are just saying that we can use different basis vectors, but not why we might want to do that.

10.1 Motivation

Our first main topic, the notion of coordinates, links the notions of basis vectors, subspace, and dimension in a very tangible manner. Eigenvalues and eigenvectors are required in EECS 442 (Computer Vision), EECS 445 (Machine learning), and ME 561 (Digital Control). The material on rank and nullity collects in one place useful facts that your author had to learn over his *first five or six years* of using Linear Algebra. Having them all in one place like this is almost too nice of a gift!

10.2 Basis Vectors, Coordinates, and Dimension

We consider \mathbb{R}^n again, and define some special vectors. Let I_n be the $n \times n$ identity matrix. Then $e_i := i$ -th column of I_n . For example, when $n = 4$,

$$e_1 = \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \end{bmatrix}, e_2 = \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \end{bmatrix}, e_3 = \begin{bmatrix} 0 \\ 0 \\ 1 \\ 0 \end{bmatrix}, e_4 = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 1 \end{bmatrix}.$$

Looking at \mathbb{R}^2 just to make things definite, we recall that $\{e_1, e_2\}$ is a linearly independent set, because

$$\left(\alpha_1 e_1 + \alpha_2 e_2 = \begin{bmatrix} \alpha_1 \\ \alpha_2 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix} \right) \iff (\alpha_1 = 0, \alpha_2 = 0). \quad (10.1)$$

An important property of the set $\{e_1, e_2\} \subset \mathbb{R}^2$ is that any vector $x \in \mathbb{R}^2$ can be written as a linear combination of e_1, e_2 . Indeed,

$$x =: \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = x_1 \begin{bmatrix} 1 \\ 0 \end{bmatrix} + x_2 \begin{bmatrix} 0 \\ 1 \end{bmatrix} = x_1 e_1 + x_2 e_2. \quad (10.2)$$

Moreover, there is **only one linear combination** of $\{e_1, e_2\}$ that yields the point $x = [x_1, x_2]^\top \in \mathbb{R}^2$.

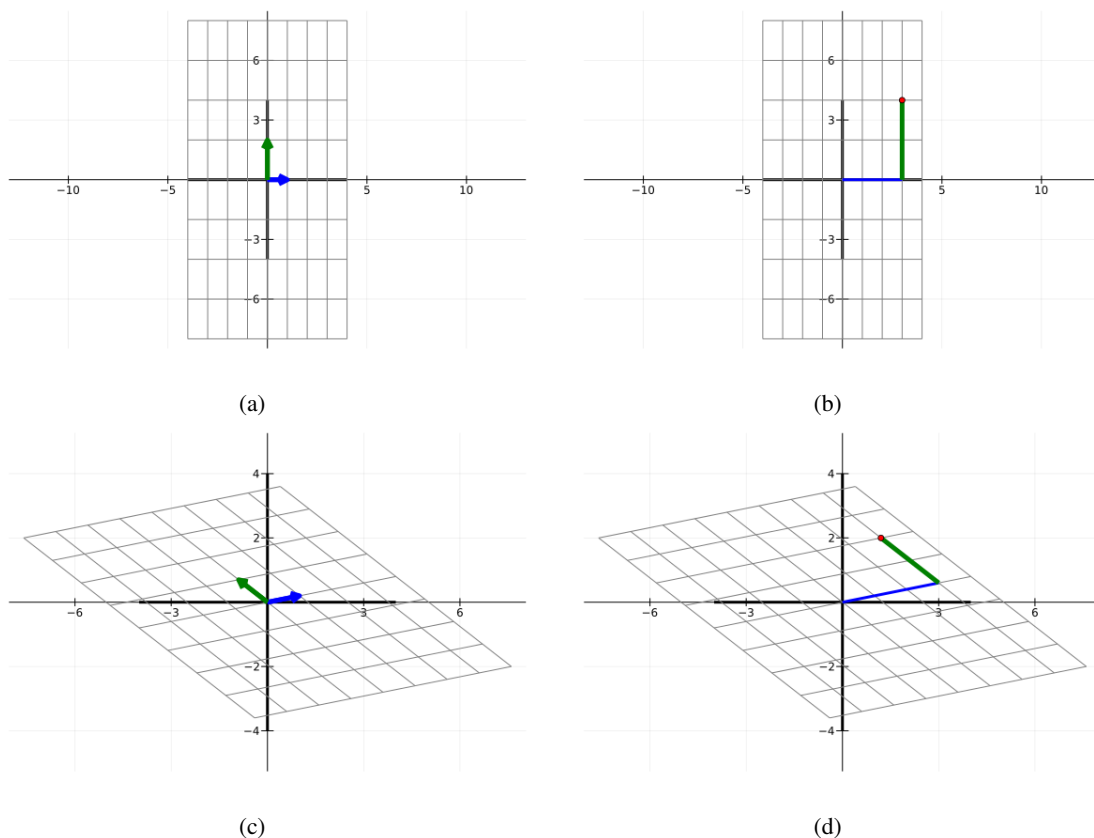


Figure 10.2: Basis vectors $\{u, v\}$ (blue and green) for \mathbb{R}^2 with the corresponding coordinates. The graphs (a) and (b) show coordinates on \mathbb{R}^2 with the “almost natural” basis vectors, $\{u = e_1, v = 2e_2\}$. Note that we now have a rectangular grid instead of a square grid because the lengths of $u = e_1$ and $v = 2e_2$ are not equal. The red dot in (b) is the point $3u + 2v$, that is, the point $(3, 2)$ in the basis $\{u, v\}$. The graphs (c) and (d) this time show coordinates on \mathbb{R}^2 with basis vectors, $\{u = [1.0, 0.2]^\top, v = [-0.9, .7]^\top\}$, which are not orthogonal. **Some of you may see the grid rotated out of the plane of the page...is so, this is an optical illusion. Everything is plotted in the same plane.** The red dot in (d) is still the point $3u + 2v$, that is, the point $(3, 2)$ in the basis $\{u, v\}$. You go along the u -axis for three units and then follow the v -axis for two units.




Figure 10.3: \mathbb{R}^3 with a two-dimensional subspace $z = -(x + y)/2$ (the colored planar surface) that we'll call V in shown in (a), while part (b) shows the natural basis (which gives the (x, y, z) coordinates in \mathbb{R}^3) do not lie in V and hence do not form natural coordinates for the surface. **Hence, expressing locations in V in terms of the “natural coordinates” (x, y, z) from \mathbb{R}^3 is not very natural at all!** It is much simpler, and more natural, to express points in V in terms of basis vectors that lie in the plane, such as the vectors $\{u, v\}$ shown in (c). Here, the vectors $\{u, v\}$ were NOT selected to be *orthogonal*, but we could have applied G-S and produced an orthonormal basis for V . Any pair of linearly independent vectors in V will work, as illustrated in Figures 10.1 and 10.2.

Basis Vectors and Dimension

Suppose that V is a subspace of \mathbb{R}^n . Then $\{v_1, v_2, \dots, v_k\}$ is a **basis for V** if

1. the set $\{v_1, v_2, \dots, v_k\}$ is linearly independent, and
2. $\text{span}\{v_1, v_2, \dots, v_k\} = V$.

The **dimension of V** is k , the number of basis vectors^a.

We note that the above definition applies to \mathbb{R}^n **because** \mathbb{R}^n is a subset of itself and it is closed under linear combinations. In particular, \mathbb{R}^n has dimension n , or we say that \mathbb{R}^n is an n -dimensional vector space.

^aA more correct definition is the maximum number of vectors in any linearly independent set contained in V . For ROB 101, the definition we gave is good enough

Basis Intuition

The essence of a **basis**: a set of vectors that is (a) “small enough” to be linearly independent and yet (b) “big enough” to generate all vectors in a vector space or a subspace by forming linear combinations.

Yes, it's kind of a **“Goldilocks” notion**: a set of vectors that is **not too big** (linearly independent) and **not too small** (spans the subspace). **Just the right size!**

If we add one more vector to a basis for V , then either the new set will become linearly dependent or it will span a set that is larger than V . If we take away even one vector from a basis, then it will no longer span the original set. So yes, a basis really is a “Goldilocks” notion.

Basis vectors are important because they provide a simple means to generate all vectors in a vector space or a subspace by forming linear combinations from a finite list of vectors. The basis and the subspace can be essentially treated as one and the same object when it comes to computations: we can manipulate a subspace in a computer by computing with its basis vectors! **An under appreciated aspect of a basis for a subspace V is that it defines a set of coordinates for V** as illustrated in Figures 10.1, 10.2, and 10.4.

A Basis for a Subspace Provides a Coordinate Grid Adapted to the Subspace




Figure 10.4: This figure is an alternative representation of what we showed in Fig. 10.3, namely \mathbb{R}^3 with a two-dimensional subspace (the gridded planar surface) V , this time with an orthogonal basis $\{u, v\}$. The natural basis $\{e_1, e_2, e_3\}$, which gives the (x, y, z) coordinates in \mathbb{R}^3 , do not lie in V . **Hence, expressing locations in V in terms of the “natural coordinates” from \mathbb{R}^3 is not very natural at all!** Because $V = \text{span}\{u, v\}$, it is much simpler, and more natural, to express points in V in terms of the basis vectors u and v . Here, u and v were selected to be *orthogonal*, but that is not a requirement. Any pair of linearly independent vectors in V will work, as illustrated in Figures 10.1 and 10.2.

Vector Space Coordinates and Vector Representation

Suppose that V is a k -dimensional subspace of \mathbb{R}^n with basis $\{v_1, v_2, \dots, v_k\}$ or all of \mathbb{R}^n itself (in which case, $k = n$). Then each $x \in V$ can be expressed (uniquely) as a linear combination of basis vectors

$$x = \alpha_1 v_1 + \alpha_2 v_2 + \dots + \alpha_k v_k. \quad (10.3)$$

Stacking the coefficients $\alpha_1, \alpha_2, \dots, \alpha_k$ into a column vector yields

$$[x]_{\{v_1, \dots, v_k\}} := \begin{bmatrix} \alpha_1 \\ \alpha_2 \\ \vdots \\ \alpha_k \end{bmatrix}, \quad (10.4)$$

which is called the **representation of x** in the basis $\{v_1, v_2, \dots, v_k\}$. The k -tuple

$$\alpha := (\alpha_1, \alpha_2, \dots, \alpha_k) \quad (10.5)$$

forms the **coordinates of x** associated to the basis $\{v_1, v_2, \dots, v_k\}$.

Remark: We could just as easily written $x = z_1 v_1 + z_2 v_2 + \dots + z_k v_k$, and then denoted our coordinates on V as $z := (z_1, z_2, \dots, z_k)$. When you think of the coefficients in the linear combination as being constants, then denoting them as α_k, c_k or a_k is rather common. If you are thinking of them as being variables, then denoting them as x_k, y_k , or z_k is common. There is not fixed convention.

Example 10.1 We consider the subspace of \mathbb{R}^3 defined by

$$V := \left\{ x = \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} \in \mathbb{R}^3 \mid x_1 + x_2 + x_3 = 0 \right\}.$$

Show that

$$\left\{ v_1 = \begin{bmatrix} 1 \\ 0 \\ -1 \end{bmatrix}, v_2 = \begin{bmatrix} 0 \\ 1 \\ -1 \end{bmatrix} \right\}$$

is a basis for V and hence V is a two dimensional subspace of \mathbb{R}^3 . In addition, show that

$$v := \begin{bmatrix} 3 \\ -4 \\ 1 \end{bmatrix} \in V$$

and find its coordinates on V .

Solution: To show that $\{v_1, v_2\}$ is a basis for V , we need that to check that

- $\{v_1, v_2\} \subset V$,
- the set $\{v_1, v_2\}$ is linearly independent, and
- $\text{span}\{v_1, v_2\} = V$.

We leave the reader to show the first two properties: that v_1 and v_2 are in V and they are linearly independent. The hard part is showing the span property, namely, that all vectors in V can be written as a linear combination of v_1 and v_2 . To do this, we note that

$$x = \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} \in V \iff x_1 + x_2 + x_3 = 0 \iff x_3 = -(x_1 + x_2) \iff x = \begin{bmatrix} x_1 \\ x_2 \\ -(x_1 + x_2) \end{bmatrix}.$$

Taking $x_1 = 1$ and $x_2 = 0$ gives v_1 , while taking $x_1 = 0$ and $x_2 = 1$ gives v_2 .

We have

$$x = \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} \in V \iff x = \begin{bmatrix} x_1 \\ x_2 \\ -(x_1 + x_2) \end{bmatrix} \iff x = x_1 v_1 + x_2 v_2 \iff x \in \text{span}\{v_1, v_2\}.$$

The dimension follows from the number of elements in the basis.

Now, we could just as easily have written

$$x = \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} \in V \iff x_1 + x_2 + x_3 = 0 \iff x_2 = -(x_1 + x_3) \iff x = \begin{bmatrix} x_1 \\ -(x_1 + x_3) \\ x_3 \end{bmatrix}.$$

Then, taking $x_1 = 1$ and $x_3 = -1$ gives v_1 , while taking $x_1 = 0$ and $x_3 = -1$ gives v_2 .

To complete the problem, we first verify that $v^\top = [3 \ -4 \ 1]^\top$ is in V because the sum of its components equals zero. Next, we check that

$$v := \begin{bmatrix} 3 \\ -4 \\ 1 \end{bmatrix} = 3v_1 - 4v_2$$

and hence its coordinates are $(3, -4)$ in the basis $\{v_1, v_2\}$. ■

The point is that V is now rather simple to understand and manipulate as the set of linear combinations of v_1 and v_2 . Moreover, we can navigate within V by using the natural coordinates induced by our choice of basis vectors.

The same idea applies to \mathbb{R}^n itself. We are used to thinking of coordinates (x_1, x_2, \dots, x_n) corresponding to the vector (or point)

$$(x_1, x_2, \dots, x_n) \iff x_1 e_1 + x_2 e_2 + \dots + x_n e_n = \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix},$$

just as we did in (10.2). However, we are not obliged to use the natural basis vectors. We can in fact use any basis $\{v_1, v_2, \dots, v_n\}$ for \mathbb{R}^n and express a vector x in the new basis

$$(x_1, x_2, \dots, x_n) \longleftrightarrow x_1 e_1 + x_2 e_2 + \dots + x_n e_n = z_1 v_1 + z_2 v_2 + \dots + z_n v_n \longleftrightarrow (z_1, z_2, \dots, z_n).$$

Canonical or Natural Basis Vectors

Let $n \geq 1$ and, as before, define $e_i := i$ -th column of the $n \times n$ identity matrix, I_n . Then

$$\{e_1, e_2, \dots, e_n\}$$

is a basis for the vector space \mathbb{R}^n . Its elements e_i are called both **natural basis vectors** and **canonical basis vectors**. The frequency of usage of one name vs the other is about fifty-fifty!

Remark: Showing linear independence is identical to (10.1) and showing that the span is all of \mathbb{R}^n is the same as in (10.2).

Columns of Matrices and Bases of \mathbb{R}^n

We let A be an $n \times n$ matrix. The following statements are equivalent

- (a) $\det(A) \neq 0$.
- (b) The columns of A are linearly independent.
- (c) The columns of A form a basis for \mathbb{R}^n .

Remark: As a special case, we can take $A = I_n$, the columns of which give the canonical basis vectors.

Example 10.2 Determine if the vectors $\{v_1, \dots, v_5\}$ form a basis for \mathbb{R}^5 .

$$v_1 = \begin{bmatrix} 1.0 \\ 2.0 \\ 0.0 \\ 0.0 \\ 2.0 \end{bmatrix}, v_2 = \begin{bmatrix} -1.0 \\ 0.0 \\ 0.0 \\ 1.0 \\ 1.0 \end{bmatrix}, v_3 = \begin{bmatrix} 1.0 \\ 2.0 \\ -2.0 \\ 0.0 \\ 2.0 \end{bmatrix}, v_4 = \begin{bmatrix} 0.0 \\ 2.0 \\ -2.0 \\ 0.0 \\ 0.0 \end{bmatrix}, v_5 = \begin{bmatrix} -1.0 \\ 2.0 \\ 0.0 \\ 0.0 \\ 2.0 \end{bmatrix}. \quad (10.6)$$

Solution: We define

$$A = \begin{bmatrix} 1.0 & -1.0 & 1.0 & 0.0 & -1.0 \\ 2.0 & 0.0 & 2.0 & 2.0 & 2.0 \\ 0.0 & 0.0 & -2.0 & -2.0 & 0.0 \\ 0.0 & 1.0 & 0.0 & 0.0 & 0.0 \\ 2.0 & 1.0 & 2.0 & 0.0 & 2.0 \end{bmatrix}_{5 \times 5}. \quad (10.7)$$

In Julia, we compute $\det(A) = 16.0$ and hence the set of vectors $\{v_1, \dots, v_5\}$ does form a basis for \mathbb{R}^5 . ■

Example 10.3 Compute a QR Factorization of A in (10.7) and relate the vectors in (10.6), that is, the columns of A , to the matrices Q and R .

Solution: We apply the Gram-Schmidt Process with Normalization to $\{v_1, \dots, v_5\}$, the columns of A , and obtain

$$Q = \begin{bmatrix} 0.3333 & -0.6537 & 0.0000 & -0.4193 & -0.5345 \\ 0.6667 & -0.1307 & 0.0000 & 0.7338 & 0.0000 \\ 0.0000 & 0.0000 & -1.0000 & 0.0000 & 0.0000 \\ 0.0000 & 0.5883 & 0.0000 & 0.1048 & -0.8018 \\ 0.6667 & 0.4576 & 0.0000 & -0.5241 & 0.2673 \end{bmatrix}, R = \begin{bmatrix} 3.0000 & 0.3333 & 3.0000 & 1.3333 & 2.3333 \\ 0.0000 & 1.6997 & 0.0000 & -0.2615 & 1.3074 \\ 0.0000 & 0.0000 & 2.0000 & 2.0000 & 0.0000 \\ 0.0000 & 0.0000 & 0.0000 & 1.4676 & 0.8386 \\ 0.0000 & -0.0000 & 0.0000 & 0.0000 & 1.0690 \end{bmatrix}.$$

By construction, the columns of Q form an orthonormal basis for $\text{span}\{v_1, \dots, v_5\} =: \text{col span}\{A\}$. Indeed, we let $\{q_1, \dots, q_5\}$ denote the columns of Q and then verify that

$v_1 = 3q_1, v_2 = 0.33q_1 + 1.7q_2, v_3 = 3q_1 + 3q_3, v_4 = 1.33q_1 - 0.26q_2 + 2.0q_3 + 1.47q_4, v_5 = 2.22q_1 + 1.31q_2 + 0.84q_4 + 1.07q_5$, confirming what we know from Gram-Schmidt, namely that

$$\text{span}\{v_1, \dots, v_5\} = \text{span}\{q_1, \dots, q_5\}.$$

■

(Optional Read): Proof the Facts (a) \iff (b) \iff (c) for Columns of Matrices and Bases of \mathbb{R}^n .

(a) \iff (b). We denote the columns of A by $\{v_1 = a_1^{\text{col}}, v_2 = a_2^{\text{col}}, \dots, v_n = a_n^{\text{col}}\}$. The columns of A are linearly independent if, and only if, the unique solution to

$$\alpha_1 a_1^{\text{col}} + \alpha_2 a_2^{\text{col}} + \dots + \alpha_n a_n^{\text{col}} = 0_{n \times 1} \quad (10.8)$$

is the trivial solution, $\alpha_1 = 0, \alpha_2 = 0, \dots, \alpha_n = 0$. But this is equivalent to

$$\underbrace{\begin{bmatrix} a_1^{\text{col}} & a_2^{\text{col}} & \dots & a_n^{\text{col}} \end{bmatrix}}_A \underbrace{\begin{bmatrix} \alpha_1 \\ \alpha_2 \\ \vdots \\ \alpha_n \end{bmatrix}}_\alpha = 0_{n \times 1}$$

$$\Downarrow$$

$$A\alpha = 0_{n \times 1}$$

When A is square, we know that $\det(A) \neq 0$ if, and only if, the unique solution to $A\alpha = 0$ is the trivial solution $\alpha = 0_{n \times 1}$

(b) \iff (c). The direction (c) \implies (b) is trivial, hence we only need to show that (b) \implies (c). To do so, we assume that A is $n \times n$ and its columns are linearly independent, and must show that they span \mathbb{R}^n , that is, we must show that

$$\text{span}\{v_1, v_2, \dots, v_n\} = \mathbb{R}^n.$$

Well, spans are simply linear combinations, so the question becomes, can every vector in $b \in \mathbb{R}^n$ be written as a linear combination of the columns of A ? Because the dimension of \mathbb{R}^n equals n , we know that the set

$$\{b, v_1, v_2, \dots, v_n\}$$

is linearly dependent. Hence, there there exist coefficients $\alpha_0, \dots, \alpha_n$ not all zero such that

$$\alpha_0 b + \alpha_1 v_1 + \alpha_2 v_2 + \dots + \alpha_n v_n = 0.$$

We observe that if $\alpha_0 \neq 0$, because if it were zero, then

$$\alpha_1 v_1 + \alpha_2 v_2 + \dots + \alpha_n v_n = 0,$$

which is not possible because $\{v_1, v_2, \dots, v_n\}$ is linearly independent. Hence, we have that

$$b = -\frac{\alpha_1}{\alpha_0} v_1 - \frac{\alpha_2}{\alpha_0} v_2 - \dots - \frac{\alpha_n}{\alpha_0} v_n,$$

proving that $b \in \text{span}\{v_1, v_2, \dots, v_n\}$. ■

10.3 Eigenvalues and Eigenvectors

For a first introduction to eigenvalues and eigenvectors, this video by 3Blue1Brown is quite good¹: <https://www.youtube.com/watch?v=PFDu9oVAE-g>. Here, we provide a cursory introduction to the topic, hitting only some of the highlights. A more

¹In the video, when they talk about \hat{i} , they mean our natural basis vector e_1 , and when they say \hat{j} , they mean our natural basis vector e_2

thorough treatment is given in Appendices A.1.5 and A.2.

What does “eigen” even mean? From Quora, “eigen” is a German word that in English means “own”, “unique to”, “peculiar to”, or “characteristic of” the originating matrix. Your author likens it to the English word *self*. A non-zero vector $v \in \mathbb{R}^n$ such that when you multiply it by an $n \times n$ matrix you basically get the vector *itself* back, is called an *eigenvector*, or a self-vector. Exactly the same vector back? No, but almost! An *eigenvector* v satisfies $Av = \lambda v$, with λ being a scalar called the *eigenvalue*, or self-value. The exactly true “itself” part is that whenever $\lambda \neq 0$

$$\text{span}\{Av\} = \text{span}\{v\}.$$

It’s kind of amazing to contemplate: an $n \times n$ matrix has n^2 entries so how is it even possible that there are non-zero $n \times 1$ -vectors such that $Av = \lambda v$, without A being something trivial like a constant times an identity matrix?

In fact, when we multiply a matrix times a vector, we expect the terms to get all jumbled up. For example, we know that

$$\begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} a_{11}x_1 + a_{12}x_2 + a_{13}x_3 \\ a_{21}x_1 + a_{22}x_2 + a_{23}x_3 \\ a_{31}x_1 + a_{32}x_2 + a_{33}x_3 \end{bmatrix},$$

and thus it is very hard to imagine that we could have

$$\begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \lambda \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}$$

for some “magical values” of x_1, x_2 and x_3 .

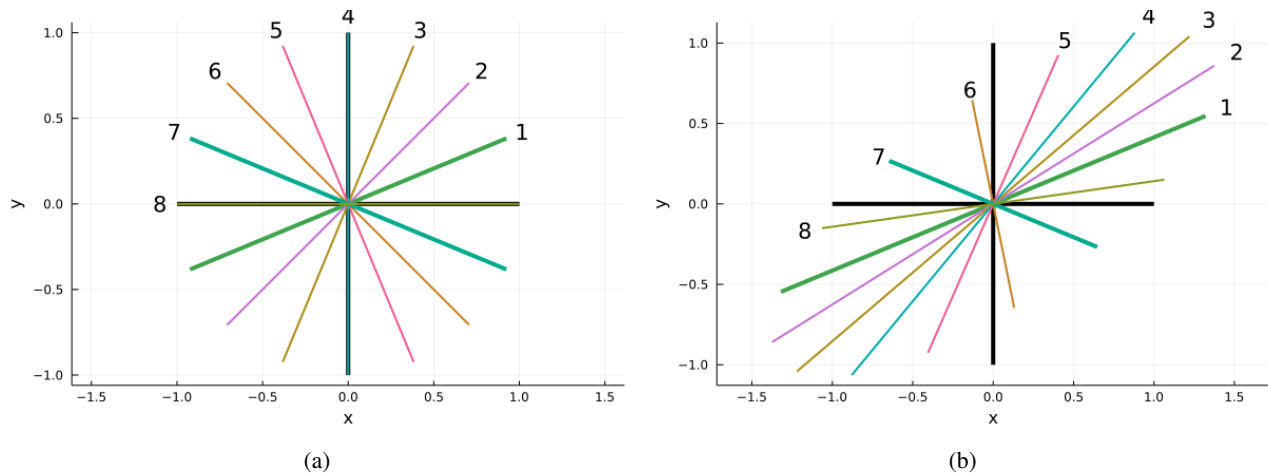


Figure 10.5: Matrices tend to both rotate and scale vectors, and at first blush, the amount of rotation and scaling is seemingly hard to predict. (a) Shows a set of eight vectors $v_i \in \mathbb{R}^2$, with each vector having length one. We’ve also plotted the negative of each vector so that you can more easily visualize their span as one-dimensional subspaces in \mathbb{R}^2 . (b) Shows the vectors Av_i for a 2×2 matrix A given in Example 10.4. If you look carefully, you will see that vectors $\{v_1, v_7\}$ are not rotated by A ; they are only scaled. Indeed $Av_1 = 1.4v_1$ and $Av_7 = 0.7v_7$. The remaining vectors $\{v_2, v_3, v_4, v_5, v_6, v_8\}$ are both rotated and scaled. Non-zero vectors that satisfy $Av = \lambda v$ are called *eigenvectors* of A and λ is called an *eigenvalue*. Comparing Figures (a) and (b), we observe that $\text{span}\{Av_1\} = \text{span}\{v_1\}$ and $\text{span}\{Av_7\} = \text{span}\{v_7\}$.

Example 10.4 For the matrix $A := \begin{bmatrix} 1.0643 & 0.8795 \\ 0.1509 & 1.0643 \end{bmatrix}$ and the vectors $\{v_1, v_2, \dots, v_8\}$ plotted in Fig. 10.5, compute Av_i and check if you can find a real constant λ_i such that $Av_i = \lambda_i v_i$.

Solution: We first give the vectors and immediately below them, their multiplication by A .

$$\underbrace{\begin{bmatrix} 0.9239 \\ 0.3827 \end{bmatrix}}_{v_1}, \underbrace{\begin{bmatrix} 0.7071 \\ 0.7071 \end{bmatrix}}_{v_2}, \underbrace{\begin{bmatrix} 0.3827 \\ 0.9239 \end{bmatrix}}_{v_3}, \underbrace{\begin{bmatrix} 0.0000 \\ 1.0000 \end{bmatrix}}_{v_4}, \underbrace{\begin{bmatrix} -0.3827 \\ 0.9239 \end{bmatrix}}_{v_5}, \underbrace{\begin{bmatrix} -0.7071 \\ 0.7071 \end{bmatrix}}_{v_6}, \underbrace{\begin{bmatrix} -0.9239 \\ 0.3827 \end{bmatrix}}_{v_7}, \underbrace{\begin{bmatrix} -1.0000 \\ 0.0000 \end{bmatrix}}_{v_8} \quad (10.9)$$

$$\underbrace{\begin{bmatrix} 1.3198 \\ 0.5467 \end{bmatrix}}_{Av_1}, \underbrace{\begin{bmatrix} 1.3744 \\ 0.8593 \end{bmatrix}}_{Av_2}, \underbrace{\begin{bmatrix} 1.2198 \\ 1.0410 \end{bmatrix}}_{Av_3}, \underbrace{\begin{bmatrix} 0.8795 \\ 1.0643 \end{bmatrix}}_{Av_4}, \underbrace{\begin{bmatrix} 0.4052 \\ 0.9255 \end{bmatrix}}_{Av_5}, \underbrace{\begin{bmatrix} -0.1307 \\ 0.6459 \end{bmatrix}}_{Av_6}, \underbrace{\begin{bmatrix} -0.6467 \\ 0.2679 \end{bmatrix}}_{Av_7}, \underbrace{\begin{bmatrix} -1.0643 \\ -0.1509 \end{bmatrix}}_{Av_8} \quad (10.10)$$

For there to exist λ_1 such that $Av_1 = \lambda_1 v_1$, we know from the first rows of v_1 and Av_1 that their ratio is $1.3198/0.9239 \approx 1.4$; we further check that the same ratio holds for the second components, and hence $Av_1 = 1.4v_1$. We'll leave it to the reader to apply the same method on the remaining vectors and verify that $Av_7 = 0.7v_7$, while none of the other vectors satisfies $Av_i = \lambda_i v_i$ for some scalar λ_i . ■

Example 10.5 Multiply the matrix

$$A := \begin{bmatrix} -8.0 & 10.0 & 10.0 \\ -2.0 & 5.0 & 2.0 \\ -10.0 & 9.0 & 12.0 \end{bmatrix} \quad (10.11)$$

times each of the vectors $\{v_1, v_2, v_3\}$, where

$$v_1 = \begin{bmatrix} 1.0 \\ 0.0 \\ 1.0 \end{bmatrix}, v_2 = \begin{bmatrix} 0.0 \\ -1.0 \\ 1.0 \end{bmatrix}, \text{ and } v_3 = \begin{bmatrix} 5.0 \\ 2.0 \\ 4.0 \end{bmatrix}. \quad (10.12)$$

Solution: It's more impressive if you do the required multiplications by hand, but turning to Julia we obtain

$$Av_1 = \begin{bmatrix} 2.0 \\ 0.0 \\ 2.0 \end{bmatrix} = 2v_1, Av_2 = \begin{bmatrix} 0.0 \\ -3.0 \\ 3.0 \end{bmatrix} = 3v_2, \text{ and } Av_3 = \begin{bmatrix} 20.0 \\ 8.0 \\ 16.0 \end{bmatrix} = 4v_3. \quad (10.13)$$

Hence, when A acts on this set of vectors, all it does is scale the vector by a factor of 2, 3 or 4, respectively. There is no “rotation” of the vector. That seems kind of magical. ■

Eigen Stuff: Temporary Definitions

Let A be an $n \times n$ matrix with real coefficients. A scalar $\lambda \in \mathbb{R}$ is an **eigenvalue** of A , if there exists a non-zero vector $v \in \mathbb{R}^n$ such that $Av = \lambda v$. Any such vector v is called an **eigenvector** associated with λ . We note that if v is an eigenvector, then so is αv for any $\alpha \neq 0$, and therefore, eigenvectors are not unique. The true definition is given in Appendix A.2.

To find eigenvalues, we need to have conditions under which there exists $v \in \mathbb{R}^n$, $v \neq 0$, such that $Av = \lambda v$. We first note that

$$Av = \lambda v \iff \lambda v - Av = 0_{n \times 1} \iff \lambda I v - Av = 0_{n \times 1} \iff (\lambda I - A)v = 0_{n \times 1}.$$

We then note that there exists $v \neq 0_{n \times 1}$ such that $(\lambda I - A)v = 0_{n \times 1}$ if, and only if

$$\det(\lambda I - A) = 0.$$

Example 10.6 Let A be the 2×2 real matrix $A = \begin{bmatrix} 1 & 2 \\ 3 & 2 \end{bmatrix}$. Determine, if any, its eigenvalues and eigenvectors.

Solution: To find eigenvalues, we need to solve

$$\det(\lambda I - A) = \begin{vmatrix} \lambda - 1 & -2 \\ -3 & \lambda - 2 \end{vmatrix} = (\lambda - 1)(\lambda - 2) - 6 = \lambda^2 - 3\lambda - 4 = 0.$$

We compute the discriminant of this quadratic equation and we find

$$b^2 - 4ac = 9 + 16 = 25 > 0,$$

and therefore there are two real solutions. We compute them with the quadratic formula to be $\lambda_1 = -1$ and $\lambda_2 = 4$.

To determine an eigenvector associated with $\lambda_1 = -1$, we need to find $v_1 \in \mathbb{R}^2$ such that

$$\begin{aligned} (A - \lambda_1 I_2)v_1 &= 0_{2 \times 1} \\ \Downarrow \\ \left(\begin{bmatrix} 1 & 2 \\ 3 & 2 \end{bmatrix} - (-1) \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \right) \begin{bmatrix} v_{1a} \\ v_{1b} \end{bmatrix} &= \begin{bmatrix} 0 \\ 0 \end{bmatrix} \\ \Downarrow \\ \begin{bmatrix} 2 & 2 \\ 3 & 3 \end{bmatrix} \begin{bmatrix} v_{1a} \\ v_{1b} \end{bmatrix} &= \begin{bmatrix} 0 \\ 0 \end{bmatrix} \\ \Downarrow \\ \begin{bmatrix} v_{1a} \\ v_{1b} \end{bmatrix} &= \alpha_1 \begin{bmatrix} 1 \\ -1 \end{bmatrix}, \alpha_1 \neq 0. \end{aligned}$$

Similarly, to determine an eigenvector associated with $\lambda_2 = 4$, we need to find $v_2 \in \mathbb{R}^2$ such that

$$\begin{aligned} (A - \lambda_2 I_2)v_2 &= 0_{2 \times 1} \\ \Downarrow \\ \left(\begin{bmatrix} 1 & 2 \\ 3 & 2 \end{bmatrix} - (4) \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \right) \begin{bmatrix} v_{2a} \\ v_{2b} \end{bmatrix} &= \begin{bmatrix} 0 \\ 0 \end{bmatrix} \\ \Downarrow \\ \begin{bmatrix} -3 & 2 \\ 3 & -2 \end{bmatrix} \begin{bmatrix} v_{2a} \\ v_{2b} \end{bmatrix} &= \begin{bmatrix} 0 \\ 0 \end{bmatrix} \\ \Downarrow \\ \begin{bmatrix} v_{2a} \\ v_{2b} \end{bmatrix} &= \alpha_2 \begin{bmatrix} 2 \\ 3 \end{bmatrix}, \alpha_2 \neq 0. \end{aligned}$$

■

Finding Eigenvectors and Eigenvalues with Julia

Beyond 2×2 matrices, we do not compute eigenvalues or eigenvectors by hand! In ROB 101, we use Julia!

```

1 Random.seed!(876543212345678);
2 B=randn(4,4)
3 A=B'*B # symmetric matrices have real eigenvalues
4
5 E=eigen(A)
6 @show E.values
7 E.vectors

```

Output

E.values = [0.06287200462929299, 0.6813033999332612, 2.9738855645273268, 4.4839915456638]

```

4 x 4 Matrix{Float64}:
 0.339074  0.0385456 -0.71411  -0.61122
 0.551488  0.528452  -0.226828  0.604275
-0.191731  0.824533   0.318536  -0.426521
 0.737651 -0.19849   0.58063  -0.281676

```

Example 10.7 Let A be the 2×2 real matrix $A = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}$. Using our temporary definition, determine, if any, its eigenvalues and eigenvectors.

Solution: To find eigenvalues, we need to solve

$$\det(\lambda I - A) = \begin{vmatrix} \lambda & -1 \\ 1 & \lambda \end{vmatrix} = \lambda^2 + 1 = 0.$$

We compute the discriminant of this quadratic equation and we find

$$b^2 - 4ac = -4 < 0,$$

and therefore there are no real solutions. Hence, by our *temporary definition*, this 2×2 real matrix does not have any eigenvalues, and hence, neither does it have any eigenvectors.

■

Full Story on Eigenstuff

The correct definition of eigenvalues and eigenvectors requires complex numbers. Example 10.7 shows that if we allow eigenvalues to be complex numbers, then we'll have two eigenvalues corresponding to the two complex solutions of the quadratic equation $\lambda^2 + 1 = 0$, namely, $\lambda_1 = \mathbf{i}$ and $\lambda_2 = -\mathbf{i}$. As illustrated in Example 10.8, when seeking solutions to $(A - \lambda_i)v_i = 0$, you'll find that you need to allow the eigenvectors to have complex entries as well.

The full and correct treatment of eigenstuff is given in Appendices A.1.5 and A.2. Here, we are giving you a simplified treatment.

Example 10.8 (Optional Read:) Let A be the 2×2 real matrix that we treated in Example 10.7, namely, $A = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}$. Determine its eigenvalues and eigenvectors in the sense of Appendix A.2.

Solution: As in Example 10.7, to find eigenvalues, we solve

$$\det(\lambda I - A) = \begin{vmatrix} \lambda & -1 \\ 1 & \lambda \end{vmatrix} = \lambda^2 + 1 = 0.$$

We apply the quadratic equation and determine $\lambda_1 = \mathbf{i}$ and $\lambda_2 = -\mathbf{i}$. To find the eigenvectors, we solve

$$(A - \lambda_i I)v_i = 0.$$

The eigenvectors are

$$v_1 = \begin{bmatrix} 1 \\ \mathbf{i} \end{bmatrix}, v_2 = \begin{bmatrix} 1 \\ -\mathbf{i} \end{bmatrix}.$$

Note that the eigenvalues and eigenvectors each form complex conjugate pairs. Indeed,

$$\lambda_2 = \lambda_1^* \text{ and } v_2 = v_1^*.$$

■

When the Eigenvalues are Real and Distinct, the Eigenvectors form a Basis of \mathbb{R}^n

Let A be an $n \times n$ matrix with real coefficients. If the eigenvalues $\{\lambda_1, \dots, \lambda_n\}$ are real and **distinct**, that is, $\lambda_i \neq \lambda_j$ for all $1 \leq i \neq j \leq n$, then the eigenvectors $\{v_1, \dots, v_n\}$ are real and provide a basis of \mathbb{R}^n .

Once again, the full story is given in Appendices A.1.5 and A.2. An interesting tidbit is that symmetric matrices always have real eigenvalues. Moreover, their eigenvectors can always be selected to form an orthogonal matrix; see Appendix A.2.2.

Example 10.9 Using Julia, find the eigenvalues and eigenvectors of the 3×3 (symmetric) matrix below. Furthermore, determine if the eigenvectors form a basis of \mathbb{R}^3 .

$$A := \begin{bmatrix} 2.2216 & 1.6798 & -0.2670 \\ 1.6798 & 0.8457 & -0.1651 \\ -0.2670 & -0.1651 & 0.6391 \end{bmatrix}. \quad (10.14)$$

Solution:

```
1 E=eigen(A)
2 E.values
3 E.vectors
4 det(E.vectors)
```

Using Julia, we compute that

$$\begin{bmatrix} \lambda_1 & \lambda_2 & \lambda_3 \end{bmatrix} = \begin{bmatrix} -0.2817 & 0.6034 & 3.3848 \end{bmatrix} \text{ and } \begin{bmatrix} v_1 & v_2 & v_3 \end{bmatrix} = \begin{bmatrix} 0.5581 & 0.0870 & -0.8252 \\ -0.8297 & 0.0741 & -0.5533 \\ 0.0131 & 0.9934 & 0.1135 \end{bmatrix}. \quad (10.15)$$

Because the eigenvalues are distinct, we know that set $\{v_1, v_2, v_3\}$ forms a basis of \mathbb{R}^3 . To double check this, we determine that $\det(E.vectors) = 1 \neq 0$ so that indeed, the eigenvectors are linearly independent and hence form a basis. ■

Utility of Eigenvalues and Eigenvectors: They Explain how a Square Matrix acts on a Vector

Let A be an $n \times n$ real matrix with real eigenvalues $\{\lambda_1, \dots, \lambda_n\}$ that are **distinct**, that is, $\lambda_i \neq \lambda_j$ for all $1 \leq i \neq j \leq n$. It then follows that the eigenvectors $\{v_1, \dots, v_n\}$ provide a basis for \mathbb{R}^n . Let $x \in \mathbb{R}^n$ be arbitrary and write it as a linear combination of the basis of eigenvectors

$$x = \alpha_1 v_1 + \alpha_2 v_2 + \dots + \alpha_n v_n. \quad (10.16)$$

Then because $Av_i = \lambda_i v_i$,

$$Ax = \alpha_1 \lambda_1 v_1 + \alpha_2 \lambda_2 v_2 + \dots + \alpha_n \lambda_n v_n. \quad (10.17)$$

If we apply A to both sides of (10.17), we obtain

$$\begin{aligned} A^2 x &= \alpha_1 \lambda_1 A v_1 + \alpha_2 \lambda_2 A v_2 + \dots + \alpha_n \lambda_n A v_n \\ &= \alpha_1 (\lambda_1)^2 v_1 + \alpha_2 (\lambda_2)^2 v_2 + \dots + \alpha_n (\lambda_n)^2 v_n, \end{aligned} \quad (10.18)$$

where $A^2 := A \cdot A$ and we have used again, $Av_i = \lambda_i v_i$. Moreover, using this fact iteratively yields that, for all $k \geq 2$,

$$A^k x = \alpha_1 (\lambda_1)^k v_1 + \alpha_2 (\lambda_2)^k v_2 + \dots + \alpha_n (\lambda_n)^k v_n. \quad (10.19)$$

Remarks:

- Equation (10.19) for $k = 1$ says that if we write a vector x in the coordinates provided by the “eigen-basis” of a matrix, then how the matrix acts on the vector is very transparent: it simply scales each component by the corresponding eigenvalue. When a matrix has real eigenvalues, what we perceive as the matrix “rotating a vector” in the natural basis vectors $\{e_1, e_2, \dots, e_n\}$ is an illusion; what is really happening is that the matrix is expanding, contracting, or leaving the same length individual components of the vector along various directions determined by the matrix’s eigenvectors.
- The above property is exploited heavily in the design of feedback control systems.
- Equation (10.19) explains the phenomenon in Fig. 10.6, because

$$\begin{cases} |\lambda_i| < 1 & \iff \lim_{k \rightarrow \infty} \|(\lambda_i)^k v_i\| = \lim_{k \rightarrow \infty} |\lambda_i|^k \|v_i\| = 0 \\ |\lambda_i| > 1 & \iff \lim_{k \rightarrow \infty} \|(\lambda_i)^k v_i\| = \lim_{k \rightarrow \infty} |\lambda_i|^k \|v_i\| = \infty \\ |\lambda_i| = 1 & \iff \|(\lambda_i)^k v_i\| = |\lambda_i|^k \|v_i\| = \|v_i\|, k \geq 0 \end{cases}$$




Figure 10.6: The data for this figure come from Example 10.4. The matrix A is 2×2 and has real eigenvalues and eigenvectors that satisfy $Av_1 = 1.4v_1$ and $Av_2 = 0.7v_2$. A is “expanding” in the direction v_1 and “contracting” in the direction v_2 . (a) Shows a uniform distribution of points, with the two eigenvectors highlighted in red. (b) For each point x in the grid of (a), (b) shows its image Ax . The circle is being squished into an ellipse under the action of A . This phenomenon is accentuated in (c), which shows A^3x , and even more so in (d), which shows A^5x . If we looked at $A^n x$ as $n \rightarrow \infty$, all of the points would lie on the “expanding” eigenvector, v_1 .

- For the matrix in Example 10.4, we see that the component of a vector x along the eigenvector v_2 is “squished” (contracted) by the matrix because $\lambda_2 = 0.7$, while its component along the eigenvector v_1 gets “spread out” (expanded) because $\lambda_1 = 1.4$.
- If the sign of an eigenvalue were negative, then the matrix would also “flip the direction” of a vector’s component along a corresponding eigenvector, in addition to possibly expanding or contracting it.

Example 10.10 Michael Penn poses this problem on YouTube <https://youtu.be/EBdEMIJK6aY> with the title “Just an average recursion...OR IS IT?” Consider a sequence of real numbers defined by

$$a_{n+2} = \frac{a_{n+1} + a_n}{2} \quad (10.20)$$

with $a_0 = \alpha$ and $a_1 = \beta$. What is the limit of the sequence as n tends to infinity? That is, what is

$$L := \lim_{n \rightarrow \infty} a_n? \quad (10.21)$$

Your gut reaction is probably $L = \frac{\alpha + \beta}{2}$, because each term in the sequence is taking the average or mean of the preceding two terms. Is that really the answer?

Solution: It turns out this problem can be analyzed very simply with eigenvalues and eigenvectors! What, you didn't see that coming? We'll set it up as a vector recursion problem. To do that, we define

$$x_n := \begin{bmatrix} a_{n-1} \\ a_n \end{bmatrix} \quad (10.22)$$

and we note that

$$x_{n+1} = \begin{bmatrix} a_n \\ a_{n+1} \end{bmatrix} = \begin{bmatrix} a_n \\ \frac{a_n + a_{n-1}}{2} \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ \frac{1}{2} & \frac{1}{2} \end{bmatrix} \begin{bmatrix} a_{n-1} \\ a_n \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ \frac{1}{2} & \frac{1}{2} \end{bmatrix} x_n.$$

In other words,

$$x_{n+1} = \underbrace{\begin{bmatrix} 0 & 1 \\ \frac{1}{2} & \frac{1}{2} \end{bmatrix}}_A x_n. \quad (10.23)$$

Either by hand or using Julia, we compute that the eigenvalues of A are $\lambda_1 = 1$ and $\lambda_2 = -\frac{1}{2}$ with eigenvectors

$$v_1 = \begin{bmatrix} 1 \\ 1 \end{bmatrix} \text{ and } v_2 = \begin{bmatrix} 2 \\ -1 \end{bmatrix}.$$

We express our initial condition for (10.23) as a linear combination of the eigenvectors v_1 and v_2 ,

$$x_1 = \begin{bmatrix} a_0 \\ a_1 \end{bmatrix} = \begin{bmatrix} \alpha \\ \beta \end{bmatrix} = \frac{\alpha + 2\beta}{3} v_1 + \frac{\alpha - \beta}{3} v_2 \quad (10.24)$$

so that we can apply (10.19). We obtain that

$$\begin{aligned} x_{n+1} &= A^n x_1 \\ &= \frac{\alpha + 2\beta}{3} (\lambda_1)^n v_1 + \frac{\alpha - \beta}{3} (\lambda_2)^n v_2 \\ &= \frac{\alpha + 2\beta}{3} \underbrace{(1)^n}_{(\lambda_1)^n} \underbrace{\begin{bmatrix} 1 \\ 1 \end{bmatrix}}_{v_1} + \frac{\alpha - \beta}{3} \underbrace{\frac{(-1)^n}{2^n}}_{(\lambda_2)^n} \underbrace{\begin{bmatrix} 2 \\ -1 \end{bmatrix}}_{v_2}. \end{aligned} \quad (10.25)$$

Because $(1)^n = 1$ and $\lim_{n \rightarrow \infty} \frac{(-1)^n}{2^n} = 0$, we have that

$$\lim_{n \rightarrow \infty} x_n = \frac{\alpha + 2\beta}{3} \begin{bmatrix} 1 \\ 1 \end{bmatrix}$$

and hence

$$\boxed{\lim_{n \rightarrow \infty} a_n = \frac{\alpha + 2\beta}{3}},$$

which is not equal to $\frac{\alpha + \beta}{2}$! We now understand Michael Penn's title was a pun, "Just an **average recursion** or IS IT" NOT! ■

10.4 Range, Column Span, and Null Space

A Function View of a Matrix Defines two Subspaces: its Null Space and its Range

Let A be an $n \times m$ matrix. We can then define a function $f : \mathbb{R}^m \rightarrow \mathbb{R}^n$ by, for each $x \in \mathbb{R}^m$

$$f(x) := Ax \in \mathbb{R}^n. \quad (10.26)$$

The following **subsets** are naturally motivated by the function view of a matrix.

Definition:

(a) $\text{null}(A) := \{x \in \mathbb{R}^m \mid Ax = 0_{n \times 1}\}$ is the **null space** of A .

(b) $\text{range}(A) := \{y \in \mathbb{R}^n \mid y = Ax \text{ for some } x \in \mathbb{R}^m\}$ is the **range** of A .

In Example 10.15, we show that the null space and range of a matrix are in fact **subspaces**.

Example 10.11 Find the null spaces of

$$A_1 = \begin{bmatrix} 1 & 0 & 0 \\ 0 & -1 & 1 \end{bmatrix} \text{ and } A_2 = \begin{bmatrix} 1 & 2 & 3 \\ 0 & 1 & 2 \\ 0 & 0 & 1 \end{bmatrix}.$$

Solution:

$$A_1 x = 0 \iff \begin{bmatrix} 1 & 0 & 0 \\ 0 & -1 & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = 0 \iff \begin{bmatrix} x_1 \\ -x_2 + x_3 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix} \iff x = \begin{bmatrix} 0 \\ \alpha \\ \alpha \end{bmatrix}, \text{ for } \alpha \in \mathbb{R}.$$

Hence,

$$\text{null}(A_1) = \left\{ \alpha \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix} \mid \alpha \in \mathbb{R} \right\}.$$

For the second matrix,

$$A_2 x = 0 \iff \begin{bmatrix} 1 & 2 & 3 \\ 0 & 1 & 2 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix} \iff \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}.$$

Hence,

$$\text{null}(A_2) = \left\{ \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix} \right\}.$$

In passing, we note that

$$\text{null}(A_1) = \text{span} \left\{ \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix} \right\},$$

and hence is a one-dimensional subspace, and that $\text{null}(A_2) = \{0_{3 \times 1}\}$, which is a zero-dimensional subspace. ■

Example 10.12 Find the ranges of

$$A_3 = \begin{bmatrix} 1 & 0 & 0 \\ 0 & -1 & 1 \end{bmatrix} \text{ and } A_4 = \begin{bmatrix} 1 & 2 & 3 \\ 0 & 0 & 2 \\ 0 & 0 & 1 \end{bmatrix}.$$

Solution: We note that A_3 is 2×3 and A_4 is 3×3 . Hence,

$$\begin{aligned} \text{range}(A_3) &= \{A_3 x \mid x \in \mathbb{R}^3\} \\ &= \left\{ \begin{bmatrix} 1 & 0 & 0 \\ 0 & -1 & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} \mid \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} \in \mathbb{R}^3 \right\} \\ &= \left\{ \begin{bmatrix} 1 & 0 & 0 \\ 0 & -1 & 1 \end{bmatrix} \begin{bmatrix} \alpha_1 \\ \alpha_2 \\ \alpha_3 \end{bmatrix} \mid \alpha_1, \alpha_2, \alpha_3 \in \mathbb{R} \right\} \\ &= \left\{ \alpha_1 \begin{bmatrix} 1 \\ 0 \end{bmatrix} + \alpha_2 \begin{bmatrix} 0 \\ -1 \end{bmatrix} + \alpha_3 \begin{bmatrix} 0 \\ 1 \end{bmatrix} \mid \alpha_1, \alpha_2, \alpha_3 \in \mathbb{R} \right\} \\ &= \left\{ \alpha_1 \begin{bmatrix} 1 \\ 0 \end{bmatrix} + \alpha_2 \begin{bmatrix} 0 \\ -1 \end{bmatrix} \mid \alpha_1, \alpha_2 \in \mathbb{R} \right\}, \end{aligned}$$

where the third column of A_3 was eliminated because it is linearly dependent on the first two columns; in fact, it is the negative of the second column.

$$\begin{aligned} \text{range}(A_4) &= \{A_4x \mid x \in \mathbb{R}^3\} \\ &= \left\{ \begin{bmatrix} 1 & 2 & 3 \\ 0 & 0 & 2 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} \mid \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} \in \mathbb{R}^3 \right\} \\ &= \left\{ \begin{bmatrix} 1 & 2 & 3 \\ 0 & 0 & 2 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \alpha_1 \\ \alpha_2 \\ \alpha_3 \end{bmatrix} \mid \alpha_1, \alpha_2, \alpha_3 \in \mathbb{R} \right\} \\ &= \left\{ \alpha_1 \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} + \alpha_2 \begin{bmatrix} 2 \\ 0 \\ 0 \end{bmatrix} + \alpha_3 \begin{bmatrix} 3 \\ 2 \\ 1 \end{bmatrix} \mid \alpha_1, \alpha_2, \alpha_3 \in \mathbb{R} \right\} \\ &= \left\{ \alpha_1 \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} + \alpha_3 \begin{bmatrix} 3 \\ 2 \\ 1 \end{bmatrix} \mid \alpha_1, \alpha_3 \in \mathbb{R} \right\}, \end{aligned}$$

where the second column was eliminated because it is dependent on the first column (in fact, it is twice the first column).

We note that

$$\text{range}(A_3) = \text{span}\left\{ \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ -1 \end{bmatrix} \right\} \text{ and } \text{range}(A_4) = \text{span}\left\{ \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 3 \\ 2 \\ 1 \end{bmatrix} \right\},$$

which are both two-dimensional subspaces. ■

Null Space of A Consists of Vectors Orthogonal to the Rows of A

Let A be an $n \times m$ matrix so that its rows are m -vectors.

$$A = \begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1m} \\ a_{21} & a_{22} & \cdots & a_{2m} \\ \vdots & \vdots & \ddots & \vdots \\ a_{n1} & a_{n2} & \cdots & a_{nm} \end{bmatrix} =: \begin{bmatrix} a_1^{\text{row}} \\ a_2^{\text{row}} \\ \vdots \\ a_n^{\text{row}} \end{bmatrix},$$

with $(a_i^{\text{row}})^\top \in \mathbb{R}^m$ for $1 \leq i \leq n$. Then

$$x \in \text{null}(A) \iff A \cdot x = 0 \iff x \perp (a_i^{\text{row}})^\top, 1 \leq i \leq n \iff x \bullet (a_i^{\text{row}})^\top = 0, 1 \leq i \leq n.$$

Remark: The above fact means we can use Gram-Schmidt to compute the null space of a matrix!

Remark: Note that $A \cdot x = \begin{bmatrix} a_1^{\text{row}} \\ a_2^{\text{row}} \\ \vdots \\ a_n^{\text{row}} \end{bmatrix} \cdot x =$

For an $n \times m$ matrix A , we seek a basis for \mathbb{R}^m where the last part of the basis is orthogonal to $\{(a_1^{\text{row}})^\top, (a_2^{\text{row}})^\top, \dots, (a_n^{\text{row}})^\top\}$. To do this, we apply Gram-Schmidt to

$$\{(a_1^{\text{row}})^\top, (a_2^{\text{row}})^\top, \dots, (a_n^{\text{row}})^\top, e_1, \dots, e_m\}, \quad (10.27)$$

where $\{e_1, \dots, e_m\}$ are the canonical basis vectors for \mathbb{R}^m . **Why does this work?** The set of vectors in (10.27) span all of \mathbb{R}^m , because

$$\text{span}\{(a_1^{\text{row}})^\top, (a_2^{\text{row}})^\top, \dots, (a_n^{\text{row}})^\top, e_1, \dots, e_m\} \supset \text{span}\{e_1, \dots, e_m\} = \mathbb{R}^m.$$

When we apply Gram-Schmidt working from left to right, we'll build an orthogonal (or orthonormal, it's our choice) basis for $\text{span}\{(a_1^{\text{row}})^\top, (a_2^{\text{row}})^\top, \dots, (a_n^{\text{row}})^\top\}$, and then complete it with a set of vectors that are orthogonal to these vectors. These last vectors will be an orthogonal (or orthonormal, it's our choice) basis for the null space of A . Here is the idea implemented in code.

```

1 function NullSpace (A)
2     n, m=size (A)
3     myI=zeros (m, m) +I
4     M=[copy (A')  myI]
5     V=Array{Float64, 2} (undef, m, 0)
6     epsilon=1e-8
7     i=0
8     # Build an orthonormal basis for the column span of transpose(A)
9     # It is not assumed that the columns are linearly independent in R^m
10    for k = 1:n
11        vi=M[:, k]
12        for j=1:i
13            vi= vi- (vi' *V[:, j]) *V[:, j]
14        end
15        norm_vi=sqrt (vi' *vi)
16        if norm_vi > epsilon
17            V=[V  vi/norm_vi]
18            i=i+1
19        end
20    end
21    dimColSpan=i
22    # Now, we complete the above basis for the column span to a basis for all of R^m.
23    # Gram-Schmidt will make sure that these extra vectors are orthonormal to
24    # the column span of transpose(A), and hence they form a basis for null space of A
25    for k = n+1:n+m
26        vi=M[:, k]
27        for j=1:i
28            vi= vi- (vi' *V[:, j]) *V[:, j]
29        end
30        norm_vi=sqrt (vi' *vi)
31        if norm_vi > epsilon
32            V=[V  vi/norm_vi]
33            i=i+1
34        end
35    end
36    # We could easily have combined the two for loops, but we separated them so we could
37    # explain what is being done at each part of the computations
38    dimNullSpace=m-dimColSpan
39    if dimNullSpace > 0
40        nullSpace=V[:, (dimColSpan+1):i]
41    else
42        nullSpace=0.0*myI[:, 1]
43    end
44    return nullSpace, dimColSpan, dimNullSpace, V
45 end
46

```

Example 10.13 Use Gram-Schmidt to compute the null spaces of the matrices in Example 10.11.

Solution: The algorithm returns

$$\text{null}(A_1) = \text{span}\left\{\begin{bmatrix} 0.0000 \\ 0.7071 \\ 0.7071 \end{bmatrix}\right\} \text{ and } \text{null}(A_2) = \text{span}\left\{\begin{bmatrix} 0.0000 \\ 0.0000 \\ 0.0000 \end{bmatrix}\right\},$$

which we can compare to the solutions in Example 10.11. ■

Example 10.14 Find the null space of the 6×5 matrix

$$A := \begin{bmatrix} 1.0857 & -0.8998 & -0.0514 & -2.3708 & 0.6346 \\ -1.2072 & -0.5681 & 0.7294 & 0.5572 & 1.0732 \\ 0.2926 & 0.0825 & 0.8810 & -0.1519 & 1.1934 \\ 0.6916 & -0.1154 & 0.3295 & -0.8731 & 0.6363 \\ -0.8576 & 1.2444 & 0.4510 & 2.6148 & -0.2469 \\ -0.3632 & -0.1200 & -1.0146 & 0.1691 & -1.3654 \end{bmatrix} \quad (10.28)$$

as well as an orthonormal basis for \mathbb{R}^5 .

Solution: We apply our Julia function `NullSpace(A)` and compute that an orthonormal basis for \mathbb{R}^5 is given by the columns of the matrix

$$V = \begin{bmatrix} 0.3835 & -0.5382 & 0.5258 & 0.5356 & 0.0000 \\ -0.3178 & -0.3922 & 0.5123 & -0.6693 & 0.1867 \\ -0.0182 & 0.3847 & 0.4783 & -0.0700 & -0.7861 \\ -0.8375 & 0.0643 & 0.1579 & 0.5092 & 0.1016 \\ 0.2242 & 0.6360 & 0.4555 & 0.0314 & 0.5804 \end{bmatrix}.$$

Our function conveniently returns the dimension of $(A) = 2$, from which we know that the last two columns of V are a basis for the null space of A . Just to drive home the point, we compute

$$A \cdot V = \begin{bmatrix} 2.8310 & 0.0000 & 0.0000 & 0.0000 & 0.0000 \\ -0.5217 & 1.8714 & -0.0000 & -0.0000 & -0.0000 \\ 0.4647 & 0.8983 & 1.1371 & 0.0000 & 0.0000 \\ 1.1698 & 0.1484 & 0.6141 & 0.0000 & 0.0000 \\ -2.9777 & 0.1581 & 0.7027 & -0.0000 & -0.0000 \\ -0.5305 & -1.0053 & -1.3329 & -0.0000 & -0.0000 \end{bmatrix},$$

from which we confirm that the last two columns of V form a basis for the null space of A ,

$$\text{null}(A) = \text{col span}\left\{\begin{bmatrix} 0.5356 & 0.0000 \\ -0.6693 & 0.1867 \\ -0.0700 & -0.7861 \\ 0.5092 & 0.1016 \\ 0.0314 & 0.5804 \end{bmatrix}\right\} =: \text{span}\{v_1, v_2\}.$$

Chapter 10.5 provides more information on the relation of $\dim \text{null}(A)$ and the number of columns of A . In particular, it defines the terms `rank` and `nullity`, which allow one to deduce that the first three columns of V form an orthonormal basis for the column span of A^\top , while the last two provide an orthonormal basis for the null space of A . ■

Remark: We note that

$$Av_1 = 0.535 \begin{bmatrix} 1.0857 \\ -1.2072 \\ 0.2926 \\ 0.6916 \\ -0.8576 \\ -0.3632 \end{bmatrix} - 0.669 \begin{bmatrix} -0.8998 \\ -0.5681 \\ 0.0825 \\ -0.1154 \\ 1.2444 \\ -0.1200 \end{bmatrix} - 0.070 \begin{bmatrix} -0.0514 \\ 0.7294 \\ 0.8810 \\ 0.3295 \\ 0.4510 \\ -1.0146 \end{bmatrix} + 0.509 \begin{bmatrix} -2.3708 \\ 0.5572 \\ -0.1519 \\ -0.8731 \\ 2.6148 \\ 0.1691 \end{bmatrix} + 0.031 \begin{bmatrix} 0.6346 \\ 1.0732 \\ 1.1934 \\ 0.6363 \\ -0.2469 \\ -1.3654 \end{bmatrix} = \begin{bmatrix} 0.0 \\ 0.0 \\ 0.0 \\ 0.0 \\ 0.0 \\ 0.0 \end{bmatrix},$$

showing that vectors in the null space provide linear combinations of the columns that add up to the zero vector.

Range of A Equals Column Span of A

Let A be an $n \times m$ matrix so that its columns are vectors in \mathbb{R}^n ,

$$A = \begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1m} \\ a_{21} & a_{22} & \cdots & a_{2m} \\ \vdots & \vdots & \ddots & \vdots \\ a_{n1} & a_{n2} & \cdots & a_{nm} \end{bmatrix} =: [a_1^{\text{col}} \quad a_2^{\text{col}} \quad \cdots \quad a_m^{\text{col}}]$$

Then

$$\text{range}(A) := \{Ax \mid x \in \mathbb{R}^m\} = \text{span}\{a_1^{\text{col}}, a_2^{\text{col}}, \dots, a_m^{\text{col}}\} =: \text{col span}\{A\}.$$

Remark: $\{Ax \mid x \in \mathbb{R}^m\} = \{x_1 a_1^{\text{col}} + x_2 a_2^{\text{col}} + \cdots + x_m a_m^{\text{col}} \mid (x_1, x_2, \dots, x_m) \in \mathbb{R}^m\} =: \text{col span}\{A\}$

Example 10.15 Show that both the null space and range of an $n \times m$ matrix A are subspaces.

Solution: (a) We suppose that v_1 and v_2 are in $\text{null}(A)$. Hence, $Av_1 = 0$ and $Av_2 = 0$. We form a linear combination $\alpha_1 v_1 + \alpha_2 v_2 \in \mathbb{R}^n$ and check if it is also in $\text{null}(A)$. For the linear combination to be in $\text{null}(A)$, we must have that A multiplying $\alpha_1 v_1 + \alpha_2 v_2$ yields zero. So we check

$$A(\alpha_1 v_1 + \alpha_2 v_2) = \alpha_1 Av_1 + \alpha_2 Av_2 = 0 + 0 = 0.$$

Hence, $\text{null}(A)$ is closed under linear combinations and it is therefore a subspace.

(b) We suppose that v_1 and v_2 are in $\text{range}(A)$. Hence, there exists u_1 and u_2 such that $Au_1 = v_1$ and $Au_2 = v_2$. We form a linear combination $\alpha_1 v_1 + \alpha_2 v_2 \in \mathbb{R}^n$ and check if it is also in $\text{range}(A)$. For the linear combination to be in $\text{range}(A)$, we must produce a $u \in \mathbb{R}^m$ such that $Au = \alpha_1 v_1 + \alpha_2 v_2$. We propose $u = \alpha_1 u_1 + \alpha_2 u_2$ and check that

$$Au = A(\alpha_1 u_1 + \alpha_2 u_2) = \alpha_1 Au_1 + \alpha_2 Au_2 = \alpha_1 v_1 + \alpha_2 v_2,$$

and hence $\alpha_1 v_1 + \alpha_2 v_2 \in \text{range}(A)$. Because it is closed under linear combinations, $\text{range}(A)$ is therefore a subspace. ■

Relation of Null Space and Range to Solutions of Linear Equations

Suppose that A is $n \times m$. Here are the key relations between solutions of $Ax = b$ and the null space and range of A .

- (a) $Ax = b$ has a solution if, and only if, $b \in \text{range}(A)$.
- (b) If $Ax = b$ has a solution, then it is unique if, and only if, $\text{null}(A) = \{0_{m \times 1}\}$.
- (c) Suppose that \bar{x} is a solution of $Ax = b$, so that $A\bar{x} = b$. Then the set of all solutions is

$$\{x \in \mathbb{R}^m \mid Ax = b\} = \bar{x} + \text{null}(A) := \{\bar{x} + \eta \mid \eta \in \text{null}(A)\}.$$

- (d) $Ax = b$ has a unique solution if, and only if $b \in \text{range}(A)$ and $\text{null}(A) = \{0_{m \times 1}\}$.
- (e) When $b = 0_{n \times 1}$, then it is always true that $b \in \text{range}(A)$. Hence we deduce that $Ax = 0_{n \times 1}$ has a unique solution if, and only if, $\text{null}(A) = \{0_{m \times 1}\}$.

In the following, we sketch the steps that prove why the above statements are true.

- Recall that $\text{range}(A) := \{y \in \mathbb{R}^n \mid y = Ax \text{ for some } x \in \mathbb{R}^m\}$. If we simply rename $y \in \mathbb{R}^n$ by $b \in \mathbb{R}^n$, we have that

$$\begin{aligned} \text{range}(A) &:= \{y \in \mathbb{R}^n \mid y = Ax \text{ for some } x \in \mathbb{R}^m\} \\ &= \{b \in \mathbb{R}^n \mid b = Ax \text{ for some } x \in \mathbb{R}^m\} \\ &= \{b \in \mathbb{R}^n \mid Ax = b \text{ for some } x \in \mathbb{R}^m\} \\ &= \{b \in \mathbb{R}^n \mid Ax = b \text{ has a solution}\}. \end{aligned}$$

Hence, $Ax = b$ has a solution if, and only if, $b \in \text{range}(A)$.

- Suppose that \bar{x} is a solution of $Ax = b$, that is, $A\bar{x} = b$, and let $\eta \in \text{null}(A)$. Is $\bar{\bar{x}} = \bar{x} + \eta$ also a solution?

$$A\bar{\bar{x}} = A(\bar{x} + \eta) = \underbrace{A\bar{x}}_b + \underbrace{A\eta}_{0_{n \times 1}} = b.$$

- Does this tell us how to generate all solutions to $Ax = b$? In fact, yes! If \bar{x} and $\bar{\bar{x}}$ are any two solutions of the system of equations, then $A(\bar{x} - \bar{\bar{x}}) = A\bar{x} - A\bar{\bar{x}} = b - b = 0_{n \times 1}$, and thus $\bar{x} - \bar{\bar{x}} \in \text{null}(A)$. Hence, once we know any one solution to $Ax = b$ and we know $\text{null}(A)$, we can generate all solutions of $Ax = b$.
- Moreover, if $Ax = b$ has a solution, then it is unique if, and only if, $\text{null}(A) = \{0_{m \times 1}\}$, the zero vector in \mathbb{R}^m .

10.5 Rank and Nullity

When we think about the columns of a matrix, it seems pretty clear that the sum of the number of linearly independent vectors and the number of linearly dependent vectors has to equal the number of columns in the matrix. When you first learn about null space and range (or column span) of a matrix, it may not be clear at all that their dimensions are related. We will now show that for an arbitrary $n \times m$ matrix A , the number of linearly **dependent** columns of A is equal to $\dim \text{null}(A)$. On the one hand, this is kind of remarkable because the columns of A are vectors in \mathbb{R}^n , while vectors in the null space of A are in \mathbb{R}^m , which are different vector spaces. On the other hand, if the columns of A are (all) linearly independent, then the unique solution to $Ax = 0$ is the zero vectors, which means that $\text{null}(A) = \{0_{m \times 1}\}$, so yeah, there is some connection here!

Definition of Rank and Nullity

For an $n \times m$ matrix A ,

Def. $\text{rank}(A) := \dim \text{range}(A)$.

Def. $\text{nullity}(A) := \dim \text{null}(A)$.

Remark: If a system of equations $Ax = b$ has a solution, \bar{x} , then $A(\bar{x} + \eta) = b$ for all $\eta \in \text{null}A$. Hence, $\text{nullity}(A)$ is measuring the “dimension” of the set of solutions. Because $\text{range}(A) \subset \mathbb{R}^n$, we see that $\text{rank}(A) \leq n$.

Useful Properties of Rank and Nullity

For an $n \times m$ matrix A , the following all hold:

Fact $\text{rank}(A) + \text{nullity}(A) = m$, the number of columns in A .

Fact $\text{rank}(A^\top \cdot A) = \text{rank}(A)$.

Fact $\text{rank}(A^\top) = \text{rank}(A)$.

Fact $\text{nullity}(A^\top \cdot A) = \text{nullity}(A)$.

Fact $\text{nullity}(A^\top) + m = \text{nullity}(A) + n$.

Fact For any $m \times k$ matrix B , $\text{rank}(A \cdot B) \leq \text{rank}(A)$.

Fact For any $p \times n$ matrix C , $\text{rank}(C \cdot A) \leq \text{rank}(A)$.

Rank-Nullity Theorem

For an $n \times m$ matrix A , the property

$$\text{rank}(A) + \text{nullity}(A) = m \text{ number of columns of } A,$$

is so important that it has a name of its own: the **Rank-Nullity Theorem**. Since $\text{rank}(A)$ is equal to the number of linearly independent columns of A , it follows that $\text{nullity}(A)$ is counting the number of linearly dependent columns of A . If all of the columns of A are linearly independent, then none are dependent, and hence $\text{null}(A) = \{0_{m \times 1}\}$.

Example 10.16 Verify the Rank-Nullity Theorem for the matrices of Example 10.12, namely

$$A_3 = \begin{bmatrix} 1 & 0 & 0 \\ 0 & -1 & 1 \end{bmatrix}_{2 \times 3} \quad \text{and} \quad A_4 = \begin{bmatrix} 1 & 2 & 3 \\ 0 & 0 & 2 \\ 0 & 0 & 1 \end{bmatrix}_{3 \times 3}.$$

Solution: From Example 10.12, we have that $\text{rank}(A_3) = 2$ and $\text{rank}(A_4) = 2$. From Example 10.11, $\text{nullity}(A_3) = 1$ and thus $2 + 1 = 3$, the number of columns of A_3 . A quick calculation gives that

$$\text{null}(A_4) = \text{span}\left\{ \begin{bmatrix} 2 \\ -1 \\ 0 \end{bmatrix} \right\}.$$

Hence, $\text{nullity}(A_4) = 1$ and $2 + 1 = 3$, the number of columns of A_4 . ■

10.6 Finding a Basis for the Null Space without Using Orthogonality

In Chapter 10.4, we characterized the null space of a matrix through the lens of the Gram-Schmidt Algorithm. Specifically, a vector $x \in \text{null}(A) \iff x \perp \text{col span}\{A^\top\}$. The link with the Gram-Schmidt Algorithm is that it allows us to compute vectors that are orthogonal to the columns of A^\top . This is developed extensively in Chapter 8 of the Lab Manual.

Here we'll take a different approach. Our hope is that seeing things from multiple perspectives will advance your understanding. You may find it helpful to look back at Example 9.9, where we carried out the process given below on a concrete example.

Suppose that A is an $n \times m$ matrix and we partition its columns as

$$A = [a_1^{\text{col}} \quad \cdots \quad a_r^{\text{col}} \quad a_{r+1}^{\text{col}} \quad \cdots \quad a_m^{\text{col}}], \quad (10.29)$$

where the first r columns are linearly independent and the last $m - r$ columns are linearly dependent on the first r columns. In the language of this Chapter, we have

- $\{a_1^{\text{col}}, \dots, a_r^{\text{col}}\}$ is a basis for $\text{col span}\{A\} := \text{span}\{a_1^{\text{col}}, \dots, a_r^{\text{col}}, a_{r+1}^{\text{col}}, \dots, a_m^{\text{col}}\}$ and we recall that $\text{range}(A) = \text{col span}\{A\}$.
- $r = \text{rank}(A) := \dim \text{range}(A) = \dim \text{col span}\{A\}$
- by the Rank-Nullity Theorem, $m - r = \text{nullity}(A) := \dim \text{null}(A)$.

Hence, to build a basis for the null space, we need to find $m - r$ linearly independent vectors $\{u_1, u_2, \dots, u_{m-r}\}$ such that $Au_i = 0_{n \times 1}$. Toward this goal, we define the following matrix and set of vectors,

- $A_1 := [a_1^{\text{col}} \quad \cdots \quad a_r^{\text{col}}]$
- $b_i := a_{r+i}^{\text{col}}, 1 \leq i \leq (m - r)$
- $A_1 v_i = b_i, 1 \leq i \leq (m - r)$, where we note that the vectors $\{v_1, v_2, \dots, v_{m-r}\} \subset \mathbb{R}^r$ are defined implicitly through the solution of the given equations (the length of the vectors v_i must equal the number of columns of A_1).

By construction, the columns of A_1 are linearly independent and each b_i is a linear combination of the columns of A_1 . These two statements are true because the columns of A_1 are a basis for the column span of A and each vector b_i is a column of A . Hence, there exists a (unique) solution to the equations $A_1 v_i = b_i$, $1 \leq i \leq (m-r)$. We claim that we can use the solutions v_i to build a basis for the null space of A .

Indeed, we define

$$u_1 := \begin{bmatrix} v_1 \\ -1 \\ 0 \\ \vdots \\ 0 \\ 0 \end{bmatrix}, \quad u_2 := \begin{bmatrix} v_2 \\ 0 \\ -1 \\ \vdots \\ 0 \\ 0 \end{bmatrix}, \quad \dots, \quad u_{m-r} := \begin{bmatrix} v_{m-r} \\ 0 \\ 0 \\ \vdots \\ 0 \\ -1 \end{bmatrix}, \quad (10.30)$$

or equivalently,

$$u_1 := \begin{bmatrix} v_1 \\ -e_1 \end{bmatrix}, \quad u_2 := \begin{bmatrix} v_2 \\ -e_2 \end{bmatrix}, \quad \dots, \quad u_{m-r} := \begin{bmatrix} v_{m-r} \\ -e_{m-r} \end{bmatrix}, \quad (10.31)$$

where $\{e_1, e_2, \dots, e_{m-r}\}$ are the natural basis vectors for \mathbb{R}^{m-r} . The linear independence of the vectors in (10.30) follows from

$$\alpha_1 u_1 + \dots + \alpha_{m-r} u_{m-r} = \begin{bmatrix} \alpha_1 v_1 + \dots + \alpha_{m-r} v_{m-r} \\ -\alpha_1 \\ -\alpha_2 \\ \vdots \\ -\alpha_{m-r-1} \\ -\alpha_{m-r} \end{bmatrix} = 0_{m \times 1},$$

if, and only if, $\alpha_1 = \dots = \alpha_{m-r} = 0$. By design, the vectors $\{u_1, u_2, \dots, u_{m-r}\}$ satisfy

$$A u_i = \begin{bmatrix} A_1 & b_1 & b_2 & \dots & b_{m-r} \end{bmatrix} \begin{bmatrix} v_i \\ e_i \end{bmatrix} = A_1 v_i - b_i = 0_{n \times 1},$$

and hence $v_i \in \text{null}(A)$.

Remark 2 Recall that because the columns of A_1 are linearly independent and b_i is a linear combination of the columns of A_1 , we can use least squares to compute the solution to $A v_i = b_i$, even when A_1 is rectangular. Namely,

$$A_1 v_i = b_i \iff A_1^\top \cdot A_1 v_i = A_1^\top b_i \iff v_i = (A_1^\top \cdot A_1)^{-1} \cdot A_1^\top b_i.$$

Please see Chapter 8.2 and review the big green box before Example 8.1. Also, when you look at Example 8.1, ask yourself how the conclusion would change if $\|Ax^* - b\|$ equaled zero.

10.7 (Optional Read): Proofs of the Rank and Nullity Relations

This example is useful in the proof.

Example 10.17 (Basis and Null Space Example Combined): Suppose that $n = n_1 + n_2$, where $n_1 \geq 1$ and $n_2 \geq 1$, and let M be an $n_1 \times n_2$ matrix. Let's note that we can write any vector $x \in \mathbb{R}^n$ by stacking two vectors $x_1 \in \mathbb{R}^{n_1}$ and $x_2 \in \mathbb{R}^{n_2}$ as in

$$x = \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}.$$

Define a subset of \mathbb{R}^n by

$$V := \left\{ x = \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \in \mathbb{R}^n \mid x_1 + M x_2 = 0 \right\}.$$

Show that V is a subspace of \mathbb{R}^n and that

$$\left\{ v_1 = \begin{bmatrix} -M e_1 \\ e_1 \end{bmatrix}, v_2 = \begin{bmatrix} -M e_2 \\ e_2 \end{bmatrix}, \dots, v_{n_2} = \begin{bmatrix} -M e_{n_2} \\ e_{n_2} \end{bmatrix} \right\} \quad (10.32)$$

is a basis for V , where the e_i are the canonical basis vectors for \mathbb{R}^{n_2} . This will show that V is an n_2 -dimensional subspace of \mathbb{R}^n .

Solution: (Optional Read:) V is the null space of the matrix $\begin{bmatrix} I_{n_1 \times n_1} & M \end{bmatrix}$ because

$$\begin{bmatrix} I_{n_1 \times n_1} & M \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = 0 \iff x_1 + Mx_2 = 0.$$

Hence, V is a subspace. Moreover, we see that $x_1 + Mx_2 = 0 \iff x_1 = -Mx_2$ and thus

$$x = \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \in V \iff x = \begin{bmatrix} -Mx_2 \\ x_2 \end{bmatrix}, x_2 \in \mathbb{R}^{n_2}. \quad (10.33)$$

Let $\{e_1, e_2, \dots, e_{n_2}\} \subset \mathbb{R}^{n_2}$ be the canonical basis vectors for \mathbb{R}^{n_2} and define

$$\{v_1, v_2, \dots, v_{n_2}\} \subset \mathbb{R}^n$$

by

$$v_i := \begin{bmatrix} -Me_i \\ e_i \end{bmatrix}. \quad (10.34)$$

By (10.33), v_i is indeed an element of V . To show the vectors in (10.34) form a basis, we need to investigate:

- Is the set $\{v_1, v_2, \dots, v_{n_2}\}$ linearly independent?
- Does the set $\{v_1, v_2, \dots, v_{n_2}\}$ span V ?

We look at these one at a time. Let $\alpha_1, \alpha_2, \dots, \alpha_{n_2}$ be real numbers and consider the linear combination $\alpha_1 v_1 + \alpha_2 v_2 + \dots + \alpha_{n_2} v_{n_2}$. Then,

$$\begin{aligned} 0 &= \alpha_1 v_1 + \alpha_2 v_2 + \dots + \alpha_{n_2} v_{n_2} \\ &\iff \\ \begin{bmatrix} 0 \\ 0 \end{bmatrix} &= \alpha_1 \begin{bmatrix} -Me_1 \\ e_1 \end{bmatrix} + \alpha_2 \begin{bmatrix} -Me_2 \\ e_2 \end{bmatrix} + \dots + \alpha_{n_2} \begin{bmatrix} -Me_{n_2} \\ e_{n_2} \end{bmatrix} \\ &\iff \\ \begin{bmatrix} 0 \\ 0 \end{bmatrix} &= \begin{bmatrix} -\alpha_1 Me_1 \\ \alpha_1 e_1 \end{bmatrix} + \begin{bmatrix} -\alpha_2 Me_2 \\ \alpha_2 e_2 \end{bmatrix} + \dots + \begin{bmatrix} -\alpha_{n_2} Me_{n_2} \\ \alpha_{n_2} e_{n_2} \end{bmatrix} \\ &\iff \text{(one way implication)} \\ 0 &= \alpha_1 e_1 + \alpha_2 e_2 + \dots + \alpha_{n_2} e_{n_2} \\ &\iff \\ \alpha_1 &= 0, \alpha_2 = 0, \dots, \alpha_{n_2} = 0, \end{aligned} \quad (10.35)$$

where the last line follows from the linear independence of the canonical basis vectors. Hence, $\{v_1, v_2, \dots, v_{n_2}\}$ is linearly independent.

Next, we note that any vector $x_2 \in \mathbb{R}^{n_2}$ can be written as a linear combination

$$x_2 = \alpha_1 e_1 + \alpha_2 e_2 + \dots + \alpha_{n_2} e_{n_2},$$

because the canonical basis vectors, as their name suggests, are a basis. Hence, from (10.33), we have that $x \in V$ can be written as

$$\begin{aligned} x &= \begin{bmatrix} -Mx_2 \\ x_2 \end{bmatrix} \\ &= \begin{bmatrix} -M(\alpha_1 e_1 + \alpha_2 e_2 + \dots + \alpha_{n_2} e_{n_2}) \\ \alpha_1 e_1 + \alpha_2 e_2 + \dots + \alpha_{n_2} e_{n_2} \end{bmatrix} \\ &= \alpha_1 \begin{bmatrix} -Me_1 \\ e_1 \end{bmatrix} + \alpha_2 \begin{bmatrix} -Me_2 \\ e_2 \end{bmatrix} + \dots + \alpha_{n_2} \begin{bmatrix} -Me_{n_2} \\ e_{n_2} \end{bmatrix} \\ &= \alpha_1 v_1 + \alpha_2 v_2 + \dots + \alpha_{n_2} v_{n_2}, \end{aligned}$$

and therefore, $x \in V \iff x \in \text{span}\{v_1, v_2, \dots, v_{n_2}\}$. We conclude that (10.32) is a basis for V and hence V has dimension n_2 . ■

Useful Properties of Rank and Nullity

For an $n \times m$ matrix A , the following are true:

- (a) $\text{rank}(A) + \text{nullity}(A) = m$, the number of columns in A .
- (b) $\text{nullity}(A^\top \cdot A) = \text{nullity}(A)$.
- (c) $\text{rank}(A^\top \cdot A) = \text{rank}(A)$.
- (d) $\text{rank}(A \cdot A^\top) = \text{rank}(A^\top)$.
- (e) For any $m \times k$ matrix B , $\text{rank}(A \cdot B) \leq \text{rank}(A)$.
- (f) $\text{rank}(A^\top) = \text{rank}(A)$.
- (g) $\text{rank}(A^\top \cdot A) = \text{rank}(A \cdot A^\top)$.
- (h) $\text{nullity}(A^\top) + m = \text{nullity}(A) + n$.

- (a) If $\text{rank}(A) = m$, then the columns of A are linearly independent, which implies that $x = 0$ is the unique solution of $Ax = 0$. Hence, $\text{null}(A) = \{0_{m \times 1}\}$. It follows that the $\text{nullity}(A) = 0$ and therefore (a) holds. If $\text{rank}(A) = 0$, then A must be the zero matrix, and hence $\text{null}(A) = \mathbb{R}^m$, and we once again verify that (a) holds.

Hence, we define $\rho := \text{rank}(A)$ and suppose that $0 < \rho < m$. Our goal is to determine the dimension of the null space of A .

From the equivalence of the range and column span of a matrix, we know that A has ρ linearly independent columns and $m - \rho$ columns that are dependent on them. Because permuting the order of the columns of A does not change its rank, we assume without loss of generality that the first ρ columns of A are linearly independent and for $\rho + 1 \leq j \leq m$, we have

$$a_j^{\text{col}} \in \{a_1^{\text{col}}, a_2^{\text{col}}, \dots, a_\rho^{\text{col}}\},$$

which means there exist coefficients $\beta_{ij} \in \mathbb{R}$ such that

$$a_j^{\text{col}} = \beta_{1j} a_1^{\text{col}} + \beta_{2j} a_2^{\text{col}} + \dots + \beta_{\rho j} a_\rho^{\text{col}}. \quad (10.36)$$

Based on the above, we partition the columns of A as

$$A = [A_1 \ A_2]$$

where A_1 is given by the first ρ independent columns of A and A_2 consists of the remaining dependent columns. From (10.36), it follows that $A_2 = A_1 B$, where

$$B = \begin{bmatrix} \beta_{1(\rho+1)} & \cdots & \beta_{1m} \\ \vdots & \vdots & \vdots \\ \beta_{\rho(\rho+1)} & \cdots & \beta_{\rho m} \end{bmatrix}.$$

Therefore, we have

$$A = [A_1 \ A_1 B] = A_1 [I_\rho \ B], \quad (10.37)$$

where the columns of A_1 are linearly independent. From the same reasoning in Chapter 7.5.5 that we employed in our Pro Tip, we have

$$Ax = 0 \iff A_1 [I_\rho \ B] x = 0 \iff A_1^\top A_1 [I_\rho \ B] x = 0 \iff [I_\rho \ B] x = 0,$$

where the last equality is because $A_1^\top A_1$ is invertible. Based on the above, we partition x as

$$x = \begin{bmatrix} x_1 \\ x_2' \end{bmatrix},$$

where $x_1 \in \mathbb{R}^\rho$ and $x_2 \in \mathbb{R}^{m-\rho}$, and obtain

$$x = 0 \iff x_1 + Bx_2 = 0.$$

From here, we can apply the result in Example 10.17 and deduce that the dimension of the null space of A is $m - \rho$, which completes the proof.

(b) Our Pro Tip in Chapter 7.5.5 showed that $Ax = 0 \iff A^\top \cdot Ax = 0$. Hence,

$$\text{null}(A^\top \cdot A) = \text{null}(A),$$

and therefore their nullities agree.

(c) Combining (a) and (b) proves (c).

(d) True by starting with A^\top in place of A and then recognizing that $(A^\top)^\top = A$.

(e) From the sum of columns times row form of matrix multiplication, we have that the columns of $A \cdot B$ are a linear combination of the columns of A . Hence, $\text{col span}\{A \cdot B\} \subset \text{col span}\{A\}$, which implies that $\text{rank}(A \cdot B) \leq \text{rank}(A)$.

(f) Combining (e) with (c) we have

$$\text{rank}(A^\top) \leq \text{rank}(A^\top \cdot A) = \text{rank}(A)$$

and then combining (e) with (d) we have

$$\text{rank}(A) \leq \text{rank}(A \cdot A^\top) = \text{rank}(A^\top).$$

Hence, $\text{rank}(A^\top) \leq \text{rank}(A) \leq \text{rank}(A^\top)$, and therefore $\text{rank}(A^\top) = \text{rank}(A)$.

(g) Combining (f) with (c) and (d) proves (g).

(h) Combining (a) with (f) implies (h). Indeed, from (a), the Rank-Nullity Theorem,

$$m + \text{nullity}(A^\top) + \text{rank}(A^\top) = m + n \text{ and } n + \text{nullity}(A) + \text{rank}(A) = n + m;$$

hence

$$m + \text{nullity}(A^\top) + \text{rank}(A^\top) = n + \text{nullity}(A) + \text{rank}(A).$$

Using (f), we can cancel $\text{rank}(A^\top)$ from the left-hand side and $\text{rank}(A)$ from the right-hand side, which yields (h). ■

10.8 Looking Ahead

So far in ROB 101, we've only looked at linear problems. However, it turns out that techniques from Linear Algebra, namely vectors and matrices, can be very useful for some problems involving nonlinear functions. Hence, we will disengage from pure Linear Algebra and explore two very interesting problems:

1. **Root finding:** this is the problem of finding $x \in \mathbb{R}^n$ such that $f(x) = 0$. A special case would be $f(x) = Ax - b$, in which case, "root finding" is the same as solving $Ax = b$, a problem we know a lot about! What we will do is assume a result from Calculus² which says that near a root of $f(x) = 0$, we can approximate the nonlinear function $f(x)$ by an affine function, $Ax - b$. Solving $Ax - b = 0$ will give us an approximation of a solution to $f(x) = 0$. We can then re-approximate the function $f(x)$ near our current estimate of the root and attempt to improve the quality of the solution. Putting this in a for-loop gives us an algorithm.
2. **Minimizing a real-valued function of x :** this is the problem of finding x^* such that

$$x^* = \arg \min_{x \in \mathbb{R}^n} c(x),$$

where $c : \mathbb{R}^n \rightarrow [0, \infty)$. A special case is

$$c(x) := \|Ax - b\|^2,$$

our least-squared error solution to $Ax = b$. Modern engineering is broadly based on "optimization", the process of maximizing the efficiency of some process or minimizing the energy consumed in making a product. In Robotics, we formulate "perception" problems as one of minimizing the error in the estimated position of objects that we "perceive" with a camera or LiDAR, for example.

You will find both of these tools broadly applicable throughout your engineering career, and of course, in your engineering subjects at UofM. In Project 3, we will see how to use optimization to do "balance control" of a Segway! Your project will be a simplified version of an algorithm that could be used on Cassie Blue, the amazing bipedal robot at Michigan.

²We will teach you the result without proof. For theory, you can see Calculus I.

Chapter 11

Changing Gears: Solutions of Nonlinear Equations

Learning Objectives

- Extend our horizons from linear equations to nonlinear equations.
- Appreciate the power of using algorithms to iteratively construct approximate solutions to a problem.
- Accomplish all of this without assuming a background in Calculus.

Outcomes

- Learn that a root is a solution of an equation of the form $f(x) = 0$.
- Learn two methods for finding roots of real-valued functions of a real variable, that is for $f : \mathbb{R} \rightarrow \mathbb{R}$, namely the Bisection Method and Newton's Method
- Become comfortable with the notion of a “local slope” of a function at a point and how to compute it numerically.
- Linear approximations of nonlinear functions.
- Extensions of these ideas to vector-valued functions of several variables, that is $f : \mathbb{R}^m \rightarrow \mathbb{R}^n$, with key notions being the gradient and Jacobian of a function and their use in the Newton-Raphson Algorithm.

11.1 Motivation and Simple Ideas

The focus of ROB 101 has been systems on linear equations. Long before you came to ROB 101, however, you had studied some Algebra and solved a few nonlinear equations. Our goal here is to develop numerical methods for finding solutions to **systems of nonlinear equations**.

We will limit our notion of a solution to the set of real numbers or real vectors. Limiting our search for solutions to the real numbers has consequences. We already know that

$$x^2 + 1 = 0,$$

for example, has no real solutions because its discriminant is $\Delta = b^2 - 4ac = -4 < 0$. Nevertheless, many interesting problems in Engineering and Science can be formulated and solved in terms of “real solutions” to systems of equations¹.

Root of an Equation

Let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be a function. Then $f(x) = 0$ defines an equation. A solution to the equation is also called a **root**^a; that is $x^* \in \mathbb{R}^n$ is a root of $f(x) = 0$ if

$$f(x^*) = 0. \quad (11.1)$$

Just as with quadratic equations, it is possible that (11.1) has multiple real solutions or no real solutions.

You may wonder if we could seek solutions to $f(x) = \pi$, for example, and if we were to do that, would we still call them roots? Technically, the answer is no. The term root is reserved for solutions to $f(x) = 0$. However, if we define a new function, $\bar{f}(x) := f(x) - \pi$, then

$$\bar{f}(x^*) = 0 \iff f(x^*) - \pi = 0 \iff f(x^*) = \pi,$$

and x^* is a root of our new function $\bar{f}(x)$. If this seems like we are splitting hairs, yeah, it’s hard to disagree with that sentiment!

^aIn the 9th century, Arab writers usually called one of the solutions to a polynomial equation **jadhr** (جذر), and their medieval European translators used the Latin word **radix**; see <https://www.britannica.com/science/root-mathematics>.




Figure 11.1: Examples of a continuous function, a discontinuous function, and a graph that is not a function. Yes, in (c), the point $x = 0$ is mapped to the interval $[-1, 2]$. To be a function, each point in the domain can only map to a single point in the range.

We will say very informally that a function $f : \mathbb{R} \rightarrow \mathbb{R}$ is **continuous** if you can draw the graph of $y = f(x)$ on a sheet of paper without lifting your pencil (from the paper)! Figure 11.1-(a) clearly passes this test while Fig. 11.1-(b) does not. Figure 11.1-(c) “seems” to pass the “without lifting your pencil test,” but the graph does not represent a function! Recall that a function is a rule that associates to each element of the domain, a single value in the range, meaning, that for a given $x \in \mathbb{R}$, there can be only one value of $y \in \mathbb{R}$ such that $y = f(x)$. In Fig. 11.1-(c), for $x = 0$, we have $f(x) = y$ for all $y \in [-1, 2]$, which makes it not a function. What about the part of the graph where the “non-function” is constant, does that also make it not a function? No, it is fine for the single value of $y = -1.0$ to be associated with many values of x ; it’s the other way around that is a problem. Functions map points to points

¹In addition, we should not overlook the fact that as a first introduction to the subject of numerical methods for solving systems of nonlinear equations, working with real numbers and vectors is a great place to start!

and not points to non-trivial sets, such as the interval $[-1, 2]$.

In Calculus, you will encounter a formal definition, which goes something like this: $f : \mathbb{R} \rightarrow \mathbb{R}$ is **continuous at a point** $x_0 \in \mathbb{R}$ if for every $\epsilon > 0$, there exists a $\delta > 0$ such that,

$$|x - x_0| < \delta \implies |f(x) - f(x_0)| < \epsilon.$$

And then one says that f is **continuous** if it is continuous at x_0 for all x_0 in its domain of definition! For our purposes, the pencil test is good enough.

11.2 Bisection

We begin with the most straightforward and intuitive method for finding roots of scalar equations

$$f(x) = 0,$$

where $f : \mathbb{R} \rightarrow \mathbb{R}$, that is, f maps real numbers to real numbers. The method is based on the following fact, which, once again, is something you will encounter in Calculus. In ROB 101, we are giving you a reason to pay attention to the result when it is presented in Calculus!

Intermediate Value Theorem

Assume that f is a continuous real valued function and you know two real numbers $a < b$ such that $f(a) \cdot f(b) < 0$. Then there exists a real number c such that

- $a < c < b$ (c is between a and b), and
- $f(c) = 0$ (c is a root).

The values a and b are said to **bracket the root**, c .

Remarkably, this leads to a “method” for approximating roots with arbitrary accuracy! Here is the basic idea.

Bisection Algorithm Pseudo Code

- **Initialize:** define $a < b$ such that $f(a) \cdot f(b) < 0$
- **Start:** compute $c := \frac{a+b}{2}$, the midpoint of the interval $[a, b]$.
- Two things are possible:
 - $f(c) = 0$, in which case, we are done. $x^* = c$.
 - $f(c) \neq 0$, in which case, **either** $f(c) \cdot f(a) < 0$ **or** $f(c) \cdot f(b) < 0$. (We’ll leave to you the task of “proving” that at least one of these statements must be true and they cannot both be true.)
- **If** $f(c) \cdot f(a) < 0$
 - $b = c$; # b updates while a stays the same
- **Else**
 - $a = c$; # a updates while b stays the same
- **End If**
- **Loop Back to Start.** (*Wash, rinse, and repeat!*)

Now, as written, the above is not an effective algorithm because it may never terminate, meaning it could loop for ever and ever. For example, suppose you wanted to solve $x^2 - 2 = 0$. You know that answer is $x^* = \sqrt{2}$, an irrational number. You might think to start with the initial guesses being $a = 0$ and $b = 2$, because then $f(0) \cdot f(2) = (-2) \cdot (2) = -4 < 0$. However, $c = \frac{a+b}{2} = 1$, a rational

number, and because $f(c) \cdot f(b) < 0$, your next step is $a = 1$ and $b = 2$. In fact, you can check that $[a \ c \ b]$ evolves like this

a	c	b
0.0	1.0	2.0
1.0	1.5	2.0
1.0	1.25	1.5
1.25	1.375	1.5
1.375	1.4375	1.5
1.375	1.40625	1.4375
1.40625	1.421875	1.4375
1.40625	1.4140625	1.421875
1.4140625	1.41796875	1.421875
1.4140625	1.416015625	1.41796875
1.4140625	1.4150390625	1.416015625

the point being that c will always be a rational number and hence it will never be true that $f(c) = 0$. Of course, we can get very close to zero and we need to define what does close enough mean!




Figure 11.2: Plot of $y = 0.2x^5 + x^3 + 3x + 1$. There does not exist any formula that provides the roots of general quintic polynomials (no quintic equation)! If we want to find a root, we are forced to use numerical methods.

Example 11.1 Figure 11.2 presents a graph of the function $f(x) = 0.2x^5 + x^3 + 3x + 1$. Find a root of the function, that is, find a solution of

$$0.2x^5 + x^3 + 3x + 1 = 0.$$

Because formulas for the roots of quintic polynomials do not exist², you must use a numerical method.

Solution: We apply the bisection method. Based on Fig. 11.2, we'll bracket the root with $a = -2$ and $b = 1$. We run the algorithm and obtain the following data

a	$c = \frac{a+b}{2}$	b	$\text{sign}(f(a) \cdot f(c))$	$f(c)$
-2.0	-0.5	1.0	+1.0	-0.6312500000000001
-0.5	0.25	1.0	-1.0	1.7658203125
-0.5	-0.125	0.25	-1.0	0.623040771484375
-0.5	-0.3125	-0.125	-1.0	0.031386375427246094
-0.5	-0.40625	-0.3125	+1.0	-0.28801019787788396
-0.40625	-0.359375	-0.3125	+1.0	-0.12573728393763295
-0.359375	-0.3359375	-0.3125	+1.0	-0.046580093697411895
-0.3359375	-0.32421875	-0.3125	+1.0	-0.007453918341161714

²In fact, N. Abel proved in 1826 that formulas for roots do not exist for families of polynomials of degree higher than four. In 1835, while still in his teens, E. Galois was able to determine a necessary and sufficient condition for any given polynomial to be solvable by "roots", thereby resolving a problem that had been open for 350 years. The story goes that he wrote down his solution days before being killed in a duel.

Figure 11.3 shows the evolution of the bracketing points a and b as well as the midpoint c for the first four steps of the algorithm, while Fig. 11.4 zooms in to show more detail. From (11.2), the logic of the algorithm can be pinned down.

Logic of the Algorithm in Detail

- In Step 1, we compute $c = \frac{a+b}{2}$ and $f(a) \cdot f(c) > 0$. Recall that at each step, the Intermediate Value Theorem says we need $f(a) \cdot f(b) < 0$ to ensure that there exists a $c \in (a, b)$ such that $f(c) = 0$. Because $f(a) \cdot f(c) > 0$, we know without checking that $f(b) \cdot f(c) < 0$, and therefore, in the next step, we update $a = c$ and leave b unchanged so that $f(a) \cdot f(b) < 0$. Similar logic applies in the following steps.
- As noted, in Step 2, we have $a_{\text{new}} = c = -0.5$, while $b = 1.0$ is unchanged. This gives $c = \frac{a+b}{2} = 0.25$ and $f(a) \cdot f(c) < 0$.
- Hence, in Step 3, we have $b_{\text{new}} = c = -0.5$, while $a = 0.25$ is unchanged. This gives $c = \frac{a+b}{2} = -0.125$ and $f(a) \cdot f(c) < 0$.
- Hence, in Step 4, we have $b_{\text{new}} = c = -0.125$, while $a = 0.25$ is once again unchanged. This gives $c = \frac{a+b}{2} = -0.3125$ and $f(a) \cdot f(c) < 0$.

From the zooms in Fig. 11.4, we observe that the more we zoom into our function at a point, the more it looks like a straight line! In fact, already at Step 4, we see that if we had the formula for the line that approximates the function, we'd use it to approximate the root instead of doing more iterations with the Bisection Algorithm.

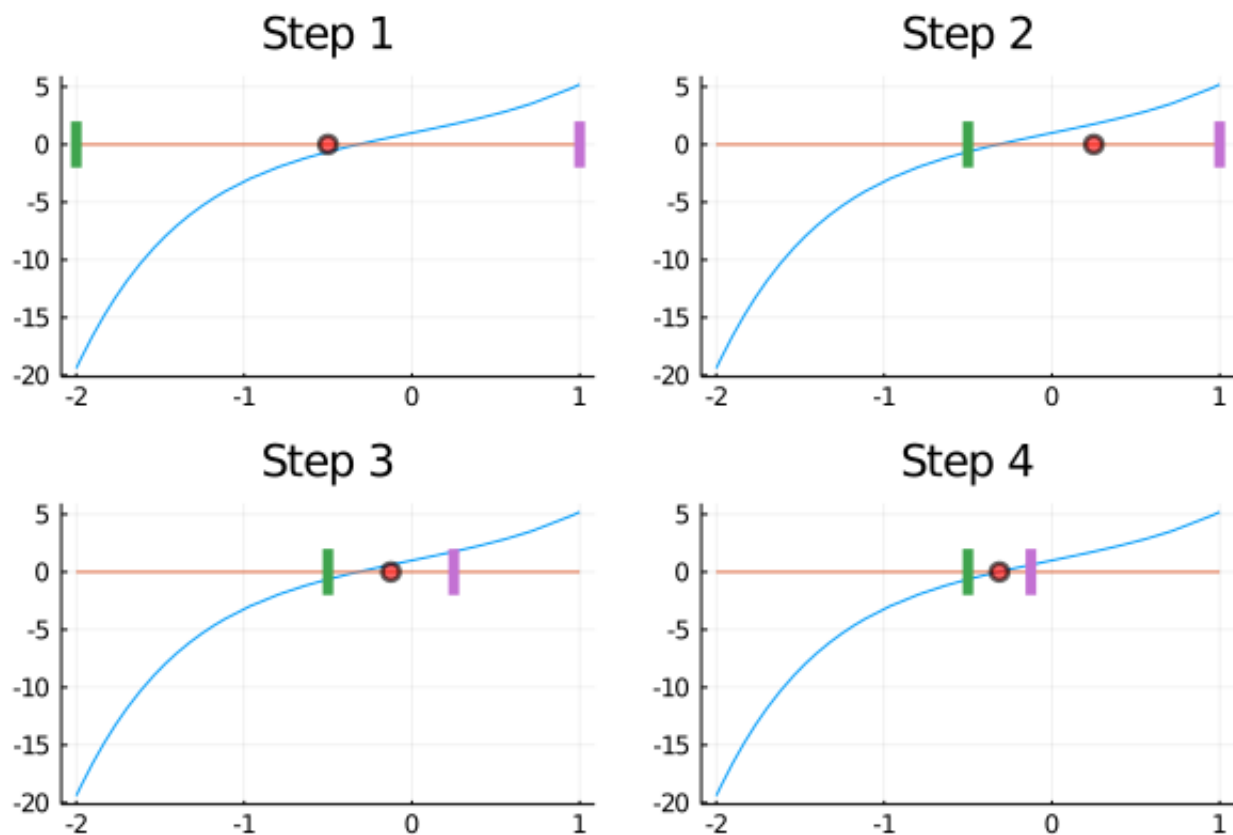


Figure 11.3: Evolution of the bracketing points a and b as well as the midpoint c in the first four steps of the Bisection Algorithm for finding a root of $0.2x^5 + x^3 + 3x + 1 = 0$. It is very clear that the algorithm hones in on a root!

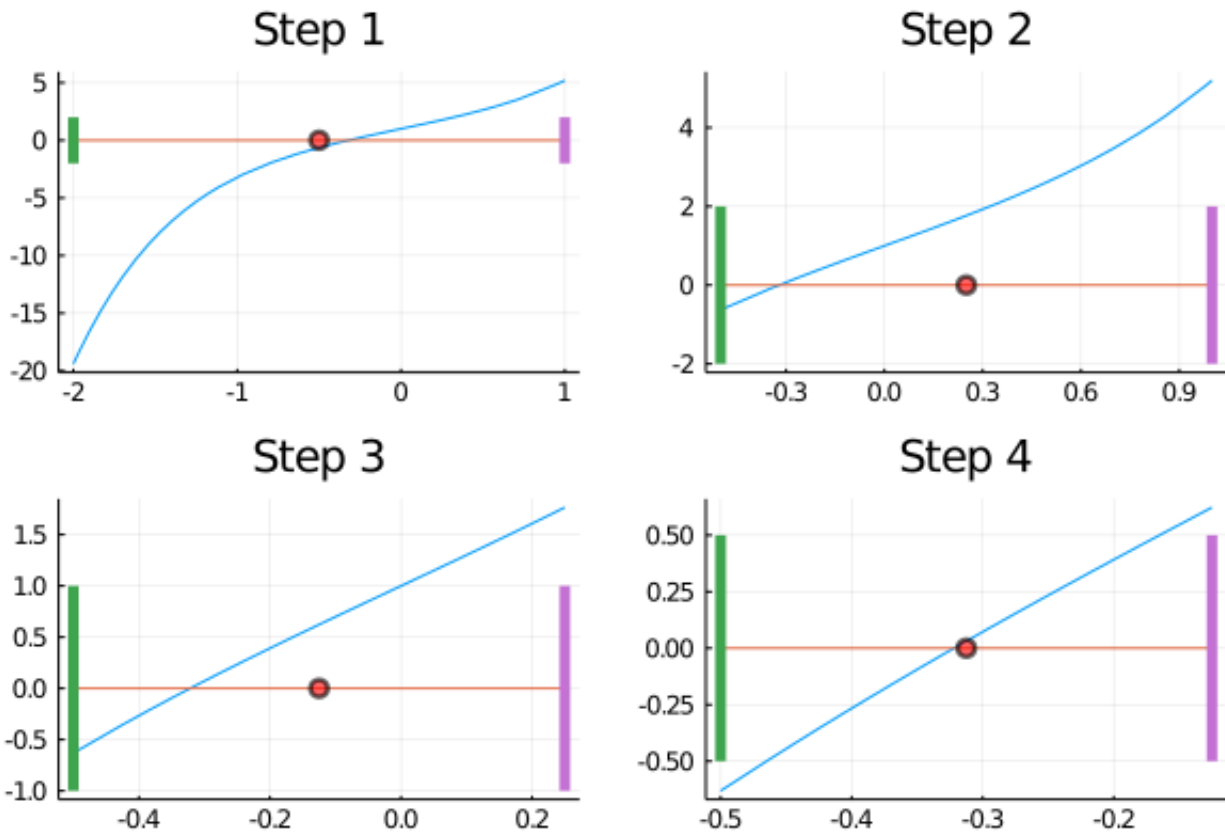


Figure 11.4: Zooms of the first four steps of the Bisection Algorithm for finding a root of $0.2x^5 + x^3 + 3x + 1 = 0$ that lies between -1 and 2 . Observe that as we zoom into the function at a point, it looks more and more like a straight line!

Linear Approximations of Functions can be Very Useful

Let's write the "line" in Step 4 of Fig. 11.4 in the form

$$y = y_c + m(x - c),$$

where y_c is the value of the line at $x = c$ and m is the slope of the line. Using the data in (11.2), and the traditional notion of "rise over run" to define the slope, we obtain

$$\begin{aligned} c &= -0.3125 \\ y_c &= f(c) \approx 0.0313864 \\ m &= \frac{f(b) - f(a)}{b - a} = \frac{0.623041 - (-0.63125)}{-0.125 - (-0.5)} \approx 3.34478 \end{aligned}$$

In Fig. 11.5, we visually illustrate how good of a fit the "line approximation"

$$y = 0.0313864 + 3.34478(x + 0.3125) = 3.34478x + 1.07663$$

provides to the function. To show its utility, we set $y = 0$ and solve for x . Doing so, we obtain

$$x^* = -0.321884 \implies f(x^*) = 0.00030674,$$

an estimate of the root that it is much better than the value given by the Bisection Algorithm at Step 4! In fact, the Bisection Algorithm has to muddle along until its 12-th iteration to better this approximation of the root. **Issac Newton made this same observation back in 1669 and turned it into an algorithm for finding roots of equations.**




Figure 11.5: Linear Approximation (black) of $f(x) = 0.2x^5 + x^3 + 3x + 1$ compared to the function itself (red). The linear approximation is very good in a sufficiently small region.

Bisection Algorithm with Sanity Checks and Tolerance Included

The algorithm takes as input a generic function $f(x)$, bracketing points a and b , and a tolerance value, tol , for terminating the algorithm, where convergence is declared when $|f(c)| \leq \text{tol}$. The algorithm also terminates after 10^4 iterations. The function returns the final values for c and prints out k , the number of iterations it took to meet the convergence criteria.

```

1 function Bisection (f, a, b, tol)
2     # First check the input data makes sense
3     if !(a < b)
4         println("a is not strictly less than b")
5         return NaN
6     end
7     if !( f(a)*f(b) < 0)
8         println("a and b fail the test provided by the Intermediate Value Theorem")
9         return NaN
10    end
11    if tol < 1e-15
12        println("tolerance is too tight")
13        return NaN
14    end
15    c = (a+b) / 2.0
16    fc = f(c)
17    k = 0
18    #
19    # Ready to run the bisection algorithm
20    #
21    while (abs(fc) > tol) & (k < 1e4)
22        if fc*f(a) < 0
23            b = copy(c)
24        else
25            a = copy(c)
26        end
27        c = (a+b) / 2
28        fc = f(c)
29        k = k + 1
30    end
31    println("Root is $c found at iteration $k")
32    return c
33 end

```

```

1 f(x)=0.2*x^5 + x^3 + 3*x + 1
2 a0=-2; b0=1
3 c=Bisection(f, a, b, 1e-10)
4 @show f(c);

```

Output

```

Root is -0.3219763464294374 found at iteration 21
f(c) = -3.854006003223276e-11

```

11.3 The Concept of a Derivative and Its Numerical Approximation




Figure 11.6: (a) The line segments represent the local slope (“rise over run”) of $f(x) = \sin(x)$ at the points $[-\pi, -\frac{3\pi}{4}, \dots, \frac{3\pi}{4}, \pi]$. Notice that each line segment is also a local linear approximation of the function. In a practical sense, what this means is that in a small region about a given point, we can replace the function with a local linear equivalent and then use linear techniques to analyze the function! In Calculus, the “local slope of a function” is called the derivative of the function. (b) The derivative of $f(x)$ is another function, denoted $\frac{df(x)}{dx}$. In Calculus, you will learn that $\frac{d}{dx} \sin(x) = \cos(x)$. Here, we are NOT using Calculus. We have computed the derivative numerically and plotted it! The maximum error in our numerical estimation of the derivative is less than 6.58×10^{-6} .

Another concept that you will learn in Calculus is the **derivative of a function**. Geometrically, it is the slope of the function at a given point, say $x_0 \in \mathbb{R}$. Note that if $x_1 \leq x_0 < x_2$, then the “rise” of the function over the interval (x_1, x_2) would be $df(x_0) := f(x_2) - f(x_1)$, while the “run” would be $dx = x_2 - x_1$, and hence the “slope” would be

$$\text{slope} := \frac{\text{rise}}{\text{run}} = \frac{df(x_0)}{dx} = \frac{f(x_2) - f(x_1)}{x_2 - x_1}.$$

In Fig. 11.6, we have attached short line segments with slopes corresponding to the derivative of the function $\sin(x)$ computed at a number of points. The hope is that this helps you grasp the geometric meaning of a derivative of a function at point as the “local slope” of the function at that point. We see that the “local slope of the function” varies with x . To tie the idea of “slope equals rise over run” to a given point, say x_0 , we let $h \neq 0$ be a small number and then we define x_1 and x_2 in terms of x_0 and h , by $x_1 = x_0$ and $x_2 = x_0 + h$. This leads to

$$\frac{df(x_0)}{dx} = \frac{f(x_2) - f(x_1)}{x_2 - x_1} = \frac{f(x_0 + h) - f(x_0)}{(x_0 + h) - x_0} = \frac{f(x_0 + h) - f(x_0)}{h}. \quad (11.3)$$

In Calculus, one analyzes what happens in the limit when h becomes very small, and in particular, one works hard to understand when the ratio in (11.3) approaches a well defined value as h becomes smaller and smaller. While we’ll explore this a bit in HW, it

is beyond the scope of our effort here.

Numerical Approximations of a Derivative

We will adopt the traditional notation from Calculus for the limiting value in (11.3), namely

$$\frac{df(x_0)}{dx} := \lim_{h \rightarrow 0} \frac{f(x_0 + h) - f(x_0)}{h}. \quad (11.4)$$

In practice, we will use “small values” for h and never compute the exact limit. Hence, we have an **approximation for the derivative** at a point, namely

$$\frac{df(x_0)}{dx} \approx \frac{f(x_0 + h) - f(x_0)}{h}, \quad (11.5)$$

which is called a **forward difference approximation to the derivative**. Note that we have replaced the informal term “slope” with the symbol for the derivative at a point, namely $\frac{df(x_0)}{dx}$.

You can also do a **backward difference approximation to the derivative**,

$$\frac{df(x_0)}{dx} \approx \frac{f(x_0) - f(x_0 - h)}{h}, \quad (11.6)$$

and a **symmetric difference approximation**, where you go both forward and backward from the point x_0 ,

$$\frac{df(x_0)}{dx} \approx \frac{f(x_0 + h) - f(x_0 - h)}{2h}. \quad (11.7)$$

The forward and backward difference approximations to the derivative are in fact exact for linear functions, while the symmetric difference approximation is *exact* for quadratic polynomials. The symmetric difference is also sometimes called a **central difference**.

If the derivative of $f(x)$ at a point x_0 exists, then for h sufficiently small, the forward difference, backward difference, and symmetric difference approximations to the derivative will always agree. If they provide different answers, then the limit in (11.4) does not exist and the function is said to be not differentiable.

```

1 f(x) = sin.(x)
2 #
3 #
4 function SymmetricDifference(f, a, b)
5     # f = generic function
6     # does 100 equally spaced points from a to b
7     # returns x and df/dx using Symmetric Differences
8     if !(b > a)
9         println("You need b > a")
10    end
11    N=100
12    h=(b-a)/(10*N)
13    x=LinRange(a, b, N)
14    x=collect(x)
15    dfdx=0*x;
16    for k=1:N
17        dfdx[k] = (f(x[k]+h) - f(x[k]-h)) / (2*h)
18    end
19    return dfdx, x
20 end
21 #
22 (dfdx, x) = SymmetricDifference(f, pi, -pi)
23 #

```

```

24 p1=plot(x, dfdx, legend=false, linewidth=3, color=:black)
25 plot!(yzero, -3.5, 3.5)
26 x0=[-pi -3*pi/4 -pi/2 -pi/4 0 pi/4 pi/2 3*pi/4 pi]'
27 df(x)=cos.(x) #Known from Calculus
28 #included to make the plot look pretty??
29 y0=df(x0)
30 scatter!(x0, y0, color=:red)
31 plot(p1)
32 plot!(fmt = :png)

```

Linear Approximation at a Point

The importance of being able to approximate a function in a region about a point by a linear function cannot be overstated. When studying the Bisection Method for finding roots, we noted that as we zoomed in on the function near the root, it looked more and more like a straight line. This property holds for all points x_0 at which a function is differentiable, that is, all points at which we can compute a derivative.

The linear function $y(x)$ that passes through the point (x_0, y_0) with slope m can be written as

$$y(x) = y_0 + m(x - x_0).$$

We use this to define the **linear approximation of a function at a point** x_0 by taking $y_0 := f(x_0)$ and $m := \frac{df(x_0)}{dx}$. This gives us

$$\mathbf{f(x) \approx f(x_0) + \frac{df(x_0)}{dx} (x - x_0).} \quad (11.8)$$

Figure 11.7 shows the linear approximation of a cubic about a point. For the record, in Calculus, this is called a First-order Taylor Expansion. You do not need to recall this terminology in ROB 101, but when you see it again in Calculus, you can say, yeah, I know why that is important!

Are all functions differentiable? No. A minimum requirement for a function to be differentiable at a point is that the function be continuous at that point. Are there functions that are continuous and not differentiable? Yes, the classic example is $f(x) = |x|$, which is plotted in Fig. 11.8 and discussed in Example 11.2.

Example 11.2 Explain why the function $f(x) = |x|$ is not differentiable at $x_0 = 0$.

Solution: We compute the forward difference, backward difference, and symmetric difference approximations to the derivative at the point $x_0 = 0$ and see if we obtain similar answers or not. For this, we let $h > 0$ be arbitrarily small. We note that then

$$|h| = h, \text{ and } |-h| = -(-h) = h.$$

Proceeding, we compute

$$\text{forward difference } \frac{df(0)}{dx} \approx \frac{f(0+h) - f(0)}{h} = \frac{|h| - 0}{h} = \frac{h}{h} = \boxed{+1}$$

$$\text{backward difference } \frac{df(0)}{dx} \approx \frac{f(0) - f(0-h)}{h} = \frac{0 - |-h|}{h} = \frac{-h}{h} = \boxed{-1}$$

$$\text{symmetric difference } \frac{df(0)}{dx} \approx \frac{f(0+h) - f(0-h)}{2h} = \frac{|h| - |-h|}{2h} = \frac{h-h}{2h} = \boxed{0}$$

These three methods giving very different approximations to the “slope” at the origin is a strong hint that the function is not differentiable at the origin. What they are telling us is that by following different paths as we approach x_0 , approaching x_0 from the left versus the right for example, gives different answers for the “slope” of the function at x_0 . In Calculus, you’ll learn that this means the function is not differentiable at x_0 . ■




Figure 11.7: The function $f(x) = x^3$ is plotted in cyan. The value of the function at the point $x_0 = 1.5$ is indicated in red. The line in black passing through $f(x_0)$ with slope $m = \frac{df(x_0)}{dx}$ satisfies $y(x) := f(x_0) + \frac{df(x_0)}{dx}(x - x_0)$. The line is called the linear approximation of $f(x)$ at x_0 . The linear approximation represents the function well in a sufficiently small region about x_0 . This approximation can be done for any point x_0 at which the derivative exists.




Figure 11.8: The function $f(x) = |x|$ is not differentiable at the origin ($x = 0$). The slope of the function just to the left of the origin is -1 , the slope just to the right of the origin is $+1$, and the slope at the origin is undefined. Everywhere else, the function is differentiable.

11.4 Newton's Method for Scalar Problems

We consider again the problem of finding roots of scalar equations $f(x) = 0$, where $f : \mathbb{R} \rightarrow \mathbb{R}$. In the Bisection Algorithm, we only required that the function be continuous. The method we develop now uses the “local slope information” of a function, and hence

requires that the function be differentiable, that is, that we can define $\frac{df(x)}{dx}$.

Let x_k be our current approximation of a root of the function f . We write the linear approximation of f about the point x_k as

$$f(x) \approx f(x_k) + \frac{df(x_k)}{dx} \cdot (x - x_k). \quad (11.9)$$

We want to choose x_{k+1} so that $f(x_{k+1}) = 0$. Based on our linear approximation in (11.9), we have that

$$f(x_{k+1}) \approx 0 \iff 0 = f(x_k) + \frac{df(x_k)}{dx} \cdot (x_{k+1} - x_k).$$

If $\frac{df(x_k)}{dx} \neq 0$, we can solve for x_{k+1} , giving us

$$\begin{aligned} \frac{df(x_k)}{dx} x_{k+1} &= \frac{df(x_k)}{dx} x_k - f(x_k) \\ \Downarrow \\ x_{k+1} &= x_k - f(x_k) / \frac{df(x_k)}{dx} \end{aligned}$$

For reasons that will become clear when we attempt a vector version of Newton's Algorithm, let's rewrite the division operation in the above formula as

$$x_{k+1} = x_k - \left(\frac{df(x_k)}{dx} \right)^{-1} f(x_k).$$

The above equation is screaming for us to put it in a loop! One step of Newton's Method is shown in Fig. 11.9.

Newton's Method

The iterative process

$$x_{k+1} = x_k - \left(\frac{df(x_k)}{dx} \right)^{-1} f(x_k) \quad (11.10)$$

for finding a root of a nonlinear equation is called **Newton's Method** or **Newton's Algorithm**. Given the current approximation x_k to a root of $f(x)$, Newton's Method corrects the approximation by the term

$$- \left(\frac{df(x_k)}{dx} \right)^{-1} f(x_k) = -f(x_k) / \frac{df(x_k)}{dx}.$$

The validity of the next approximation x_{k+1} rests upon:

- the function f being differentiable;
- the derivative $\frac{df(x_k)}{dx}$ not vanishing at points generated by the algorithm in (11.10); and
- the linear equation (11.9) is a good approximation to the function.

We boxed this last item because it is easy to overlook and is often a source of failure for the algorithm. Because (11.10) has "total faith" in (11.9) being a good approximation, it sometimes takes very big "steps" (meaning $x_{k+1} - x_k$ is large) when generating x_{k+1} to zero the linear approximation in (11.9). A safer update is to go only "part way" to the linear solution. This leads to the so-called **damped** or **modified Newton Method**,

$$x_{k+1} = x_k - \epsilon \left(\frac{df(x_k)}{dx} \right)^{-1} f(x_k), \quad (11.11)$$

where $0 < \epsilon < 1$. A typical value may be $\epsilon = 0.1$.

The standard way to "ensure" that Newton's Method generates points x_k such that the linear equation (11.9) is a good approximation to $f(x_k)$ is to start the algorithm "near" a root. As you can imagine, this is easier said than done! Hence, the damped version of Newton's Algorithm is very useful in practice.




Figure 11.9: This figure demonstrates one step of Newton's Algorithm. At a point x_k , one uses the derivative to compute a linear approximation to the function. Solving for where the linear approximation (red line) crosses the x -axis gives the next value, x_{k+1} .

Visual Representation of Newton's Algorithm

Some of you are undoubtedly more visually wired than algebraically wired (did you know that was a thing?). Here are some potential visual sources:

- Wikipedia https://upload.wikimedia.org/wikipedia/commons/e/e0/NewtonIteration_Ani.gif (The words in the legend are **function** and **tangent**.) You'll find additional textual information here as well https://en.wikipedia.org/wiki/Newton%27s_method
- Kahn Academy <https://www.youtube.com/watch?v=WuaI5G04Rcw>
- Christine Breiner, MIT Calculus I, https://www.youtube.com/watch?v=ER5B_YBFMJo

Example 11.3 For the same function as treated in Example 11.1, namely, $f(x) = 0.2x^5 + x^3 + 3x + 1$, find a root using Newton's Algorithm.

Solution: We apply the basic Newton's Method in (11.9) (that is, no damping), using each of the derivative methods given in (11.5), (11.6), and (11.7). We take $x_0 = 2$ and $h = 0.01$ for the approximate derivatives. We iterate until $|f(x_k)| < 10^{-4}$ or the algorithm fails by $\left| \frac{df(x_k)}{dx} \right| < 10^{-4}$.

Using the **Symmetric Difference Approximation** for the derivative, Newton's Method converges in five steps

x_k	$f(x_k)$	$\frac{df(x_k)}{dx}$	k
2.0000	21.4000	31.2209	0.0000
1.3146	8.0005	11.1709	1.0000
0.5984	3.0247	4.2025	2.0000
-0.1214	0.6341	3.0445	3.0000
-0.3296	-0.0255	3.3379	4.0000
-0.3220	-0.0001	3.3219	5.0000

(11.12)

Using the **Forward Difference Approximation** for the derivative, Newton's Method **fails** after 45 steps due to the estimated deriva-

tive vanishing

\mathbf{x}_k	$\mathbf{f}(\mathbf{x}_k)$	$\frac{df(\mathbf{x}_k)}{dx}$	\mathbf{k}	
2.0000	21.4000	31.2209	0.0000	
1.3146	8.0005	-1328.6981	1.0000	
1.3206	8.0680	18.1162	2.0000	
0.8752	4.3989	-360.9904	3.0000	
\vdots	\vdots	\vdots	\vdots	
$5.4e + 01$	$8.9e + 07$	$7.6e + 03$	41.0000	
$-1.2e + 04$	$-4.5e + 19$	$-4.5e + 21$	42.0000	
$-1.2e + 04$	$-4.5e + 19$	$-8.1e + 10$	43.0000	
$-5.5e + 08$	$-1.0e + 43$	$-1.0e + 45$	44.0000	
$-5.5e + 08$	$-1.0e + 43$	$0.0e + 00$	45.0000	(11.13)

Using the **Backward Difference Approximation** for the derivative, Newton's Method converges after 23 steps

\mathbf{x}_k	$\mathbf{f}(\mathbf{x}_k)$	$\frac{df(\mathbf{x}_k)}{dx}$	\mathbf{k}	
2.0000	21.4000	31.2209	0.0000	
1.3146	8.0005	1351.0399	1.0000	
1.3086	7.9346	17.5719	2.0000	
0.8571	4.2934	369.8273	3.0000	
0.8455	4.2272	12.2347	4.0000	
0.5000	2.6311	163.4044	5.0000	
0.4839	2.5702	9.8345	6.0000	
0.2225	1.6787	92.2938	7.0000	
0.2043	1.6216	8.8299	8.0000	
0.0207	1.0621	58.9551	9.0000	
0.0027	1.0080	8.4054	10.0000	
-0.1173	0.6466	39.1841	11.0000	
-0.1338	0.5963	8.0873	12.0000	
-0.2075	0.3685	25.9194	13.0000	
-0.2217	0.3239	7.6217	14.0000	
-0.2642	0.1887	16.7398	15.0000	
-0.2755	0.1524	6.8761	16.0000	
-0.2976	0.0803	10.4913	17.0000	
-0.3053	0.0552	5.8084	18.0000	
-0.3148	0.0238	6.4494	19.0000	
-0.3185	0.0116	4.5491	20.0000	
-0.3210	0.0031	4.1765	21.0000	
-0.3218	0.0006	3.5817	22.0000	
-0.3220	0.0000	3.3919	23.0000	(11.14)



Symmetric Difference Makes a Difference

In general, the symmetric difference is a better approximation to the true analytical derivative than are the forward and backward difference approximations. When used in Newton's Method, the big difference in performance of the three approximate derivative methods surprised us as well!

Why do people not use the symmetric difference all the time? Depending on the situation, you may have the value of $f(x_k)$ already at hand, in which case, to determine a forward or backward difference, you only need one additional function evaluation, namely, either $f(x_k + h)$ or $f(x_k - h)$, whereas with the symmetric difference, you must do both additional function evaluations. If f is complicated to evaluate, that may bias you toward the computationally "lighter" methods. On the other hand, as we saw in our example with Newton's Algorithm, if you converge faster, you may still come out way ahead!

The fact that the decision of which numerical differentiation method to use is not obvious and depends on the problem being solved is actually A GREAT THING: it keeps Engineers and Applied Mathematicians employed!

Remark 3 We redo the above example using the forward difference approximation of the derivative with $\epsilon = 0.9$. The results are that Newton's Method with damping converges, though very slowly.

x_k	$f(x_k)$	$\frac{df(x_k)}{dx}$	k
2.0000	21.4000	31.2209	0.0000
1.3831	8.8075	-1246.7601	1.0000
1.3895	8.8867	20.5359	2.0000
1.0000	5.2000	-361.6178	3.0000
1.0129	5.2914	16.3261	4.0000
0.7212	3.5779	-166.4880	5.0000
0.7406	3.6725	14.4315	6.0000
0.5116	2.6755	-95.8236	7.0000
0.5367	2.7735	13.7664	8.0000
0.3554	2.1121	-62.7388	9.0000
0.3857	2.2160	13.8764	10.0000
\vdots	\vdots	\vdots	\vdots
-0.3387	-0.0558	6.8063	368.0000
-0.3313	-0.0311	5.8011	369.0000
-0.3265	-0.0150	4.9292	370.0000
-0.3237	-0.0059	4.2261	371.0000
-0.3225	-0.0017	3.7289	372.0000
-0.3221	-0.0003	3.4495	373.0000
-0.3220	-0.0000	3.3411	374.0000

11.5 Vector Valued Functions: Linear Approximations, Partial Derivatives, Jacobians, and the Gradient

When developing Newton's Method of root finding for functions $f : \mathbb{R} \rightarrow \mathbb{R}$, we started with the notion of a derivative being the local slope of a function at a point, and from there, we were led to the idea of locally approximating a nonlinear function by a line! Once we had the idea of a linear approximation of the function about a given point, Newton's Algorithm basically fell into our lap by solving for a root of the linear approximation.

For the vector case of functions $f : \mathbb{R}^m \rightarrow \mathbb{R}^n$, we'll turn things around a bit and start with the idea of a linear approximation of the function about a point and see how that leads us to the notion of a **partial derivative**. Once we have that concept down, the rest is book keeping, in other words, the rest is developing a nice matrix-vector formulation of a derivative of a function. It sounds harder than it is. Let's do it!

Remark: A more traditional approach that starts by introducing the notion of a partial derivative and, from there, builds the gradient, the Jacobian, and only then, introduces the idea of a linear approximation, maybe better for some readers. That path is followed in Chap. 11.7.

11.5.1 Linear Approximation about a Point: Take 1

Our goal is to generalize the idea of a linear approximation of a (nonlinear) function $f : \mathbb{R}^m \rightarrow \mathbb{R}^n$ at a point x_0 . What we'll do is posit that the appropriate generalization should be

$$f(x) \approx f(x_0) + A(x - x_0), \quad (11.15)$$

where A is an $n \times m$ matrix. We note that the dimensions make sense because $f(x_0)$ is $n \times 1$, $(x - x_0)$ is $m \times 1$, and therefore, $A(x - x_0)$ is $n \times 1$. So far, so good.

Let's now figure out what the columns of A need to be for (11.15) to hold of x "near" x_0 . We write $A = [a_1^{\text{col}} \ a_2^{\text{col}} \ \dots \ a_m^{\text{col}}]$, where a_j^{col} is the j -th column of A . Further, let $\{e_1, e_2, \dots, e_m\}$ the canonical basis vectors for \mathbb{R}^m (which we recall are the columns of the $m \times m$ identity matrix). We next recall that our "sum over columns times rows" method of matrix multiplication gives us that

$$Ae_j = a_j^{\text{col}},$$

which is true because, using "Julia notation",

$$(e_j)[i] = \begin{cases} 1 & i = j \\ 0 & \text{otherwise} \end{cases}$$

implies that

$$Ae_j = \sum_{i=1}^m a_i^{\text{col}}(e_j)[i] = a_j^{\text{col}}.$$

We let $x = x_0 + he_j$ be a small perturbation about the nominal value x_0 . We note that $x = x_0 + he_j$ **holds all components of x constant and equal to x_0 , except for the j -th component, which is perturbed by an amount h** . When $h > 0$ is sufficiently small, (11.15) gives us

$$\begin{aligned} f(x_0 + he_j) &= f(x_0) + A(x_0 + he_j - x_0) \\ &\Downarrow \\ f(x_0 + he_j) &= f(x_0) + hAe_j \\ &\Downarrow \\ f(x_0 + he_j) &= f(x_0) + ha_j^{\text{col}} \\ &\Downarrow \\ f(x_0 + he_j) - f(x_0) &= ha_j^{\text{col}} \\ &\Downarrow \\ \frac{f(x_0 + he_j) - f(x_0)}{h} &= a_j^{\text{col}}. \end{aligned} \quad (11.16)$$

In other words, the j -th column of the matrix A in (11.15) is given by

$$\boxed{a_j^{\text{col}} = \frac{f(x_0 + he_j) - f(x_0)}{h}}, \quad (11.17)$$

which looks suspiciously like the forward difference approximation of a derivative. In fact, it looks like here we are ignoring all variables except the j -th one and computing a derivative of f with respect to x_j . And indeed, that is exactly what we are doing! Calculus has a term for it, the **partial derivative of $f(x)$** with respect to x_j , and it uses a cool symbol,

$$\boxed{\frac{\partial f(x_0)}{\partial x_j} = \lim_{h \rightarrow 0} \frac{f(x_0 + he_j) - f(x_0)}{h}}. \quad (11.18)$$

The symbol ∂ is pronounced "partial". We'd better dig into this!

11.5.2 Partial Derivatives

Partial Derivatives as Motivated by a Linear Approximation to a Function about a Point

If we let $\{e_1, e_2, \dots, e_m\}$ be the natural basis vectors for \mathbb{R}^m , then we have three ways to **numerically approximate a partial derivative**, just as we did with a “scalar” derivative

$$\frac{\partial f(x_0)}{\partial x_j} = \begin{cases} \frac{f(x_0 + he_j) - f(x_0)}{h} & \text{forward difference approximation} \\ \frac{f(x_0) - f(x_0 - he_j)}{h} & \text{backward difference approximation} \\ \frac{f(x_0 + he_j) - f(x_0 - he_j)}{2h} & \text{symmetric difference approximation.} \end{cases} \quad (11.19)$$

Example 11.4 For the function

$$f(x_1, x_2, x_3) := \begin{bmatrix} x_1 x_2 x_3 \\ \log(2 + \cos(x_1)) + x_2^{x_1} \\ \frac{x_1 x_3}{1 + x_2^2} \end{bmatrix}, \quad (11.20)$$

compute the partial derivatives $\frac{\partial f(x_0)}{\partial x_1}$, $\frac{\partial f(x_0)}{\partial x_2}$, and $\frac{\partial f(x_0)}{\partial x_3}$ at the point

$$x_0 = \begin{bmatrix} \pi \\ 1.0 \\ 2.0 \end{bmatrix}.$$

In the next example, we’ll interpret the computed partial derivatives in terms of derivatives of scalar valued functions, which we intuitively understood as slopes of a function at a point.

Solution A We’ll compute the partial derivatives in Julia, two different ways. We only need one of them to click for you.

In the first solution, we write the function given in (11.20) as $f(x)$, where $x = [x_1; x_2; x_3]$. We can then apply the numerical approximations in (11.19) directly. Because $f(x) \in \mathbb{R}^3$, the partial derivatives will also be vectors in \mathbb{R}^3 . This follows from (11.19), where each numerator is a vector in \mathbb{R}^3 , while the denominators are scalars.

```

1 x0=[pi; 1.0; 2.0]
2 # function defined in terms of x as a vector with components [x1; x2; x3].
3 function f(x)
4     x1=x[1]
5     x2=x[2]
6     x3=x[3]
7     f=[x1*x2*x3; log(2 + cos(x1)) + x2^x1; (x1*x3)/(1+x2^2)]
8     return f
9 end
10 h=0.001
11 Id=zeros(3,3)+I
12 e1=Id[:,1]; e2=Id[:,2]; e3=Id[:,3]
13 # Partial derivatives via symmetric differences
14 dfdx1=( f(x0+h*e1) - f(x0-h*e1) ) / (2*h)
15 dfdx2=( f(x0+h*e2) - f(x0-h*e2) ) / (2*h)
16 dfdx3=( f(x0+h*e3) - f(x0-h*e3) ) / (2*h)
17
18

```

Using the above code, we determine

$$\frac{\partial f(x_0)}{\partial x_1} = \begin{bmatrix} 2.0 \\ 0.0 \\ 1.0 \end{bmatrix}, \quad \frac{\partial f(x_0)}{\partial x_2} = \begin{bmatrix} 6.2832 \\ 3.1416 \\ -3.1416 \end{bmatrix}, \quad \frac{\partial f(x_0)}{\partial x_3} = \begin{bmatrix} 3.1416 \\ 0.0000 \\ 1.5708 \end{bmatrix}. \quad (11.21)$$

Solution B In the second method, we express the function exactly as it is written in (11.20). We then have to recognize that

$$x_0 + he_1 = (x_{01} + h, x_{02}, x_{03}), \quad x_0 + he_2 = (x_{01}, x_{02} + h, x_{03}), \quad \text{and} \quad x_0 + he_3 = (x_{01}, x_{02}, x_{03} + h).$$

The point is, in Mathematics, we write a function that depends on several variables like this

$$f(x) = f(x_1, x_2, x_3),$$

and never like this

$$f(x) = f\left(\begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}\right). \quad (11.22)$$

However, when we program, it is often easier to work with a function as if it were written as in (11.22), with x a column vector; as an example,

$$f(x + he_2) = f\left(\begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} + \begin{bmatrix} 0 \\ h \\ 0 \end{bmatrix}\right) = f(x_1, x_2 + h, x_3).$$

You will learn quickly enough that it is easier to “vectorize” (that is, put operations in a loop) expressions such as $f(x + he_2) = f(x + hId[:, 2])$ than it is expressions such as $f(x_1, x_2 + h, x_3)$, but we digress.

```

1 x0=[pi; 1.0; 2.0]
2 f(x1, x2, x3)=[x1*x2*x3; log(2 + cos(x1) + x2^x1); (x1*x3)/(1+x2^2)]
3 h=0.001
4
5 dfdx1=( f(pi+h, 1.0, 2.0) - f(pi-h, 1, 2) ) / (2*h)
6 dfdx2=( f(pi, 1.0+h, 2.0) - f(pi, 1-h, 2) ) / (2*h)
7 dfdx3=( f(pi, 1.0, 2.0+h) - f(pi, 1, 2-h) ) / (2*h)
8

```

The results match those in (11.21). The code for the second solution looks simpler, doesn't it? But imagine writing that out if you have 25 variables! On the other hand, the code segment

```

1 # As a loop
2 n=3
3 dfdx=Array{Float64, 2}(undef, n, 0)
4 for k = 1 : n
5     dfdxk=( f(x0+h*Id[:, k]) - f(x0-h*Id[:, k]) ) / (2*h)
6     dfdx=[dfdx dfdxk]
7 end
8 dfdx
9
10 3×3 Array{Float64, 2}:
11  2.0   6.28319  3.14159
12  0.0   3.14159  0.0
13  1.0  -3.14159  1.5708

```

is very easy to scale up! **Vectors and matrices are really about careful bookkeeping. It's kind of sad to say it that way, but it's also kind of true.**

Example 11.5 For the function in Example 11.4, interpret the components of its partial derivatives in terms of “ordinary derivatives”.

Solution Let's quite arbitrarily focus on x_2 . We define a function $g : \mathbb{R} \rightarrow \mathbb{R}^3$ by

$$g(x_2) := f(\pi, x_2, 2) = \begin{bmatrix} \pi x_2^2 \\ \log(2 + \frac{\cos(\pi)}{1+(x_2)^2}) + (x_2)^\pi \\ \frac{2\pi}{1+(x_2)^2} \end{bmatrix} = \begin{bmatrix} 2\pi x_2 \\ (x_2)^\pi \\ \frac{2\pi}{1+(x_2)^2} \end{bmatrix},$$

where we have set $x_1 = \pi$ and $x_3 = 2$. Because the components of g only depend on the single variable x_2 , we can compute their ordinary derivatives about the point $x_{02} = 1$ using symmetric differences. We do so and determine that

$$\frac{dg(x_{02})}{dx_2} \approx \frac{g(1+h) - g(1-h)}{2h} = \begin{bmatrix} 6.2832 \\ 3.1416 \\ -3.1416 \end{bmatrix}.$$

We observe that $\frac{dg(x_{02})}{dx_2} = \frac{\partial f(x_0)}{\partial x_2}$. We also note that $g(x_2)$ being a column vector with three components does not really change anything: we are simply computing the slope of each component of g . ■

Partial Wisdom

Partial derivatives are simply ordinary derivatives that we perform one variable at a time, while holding all other variables constant.

$$\begin{aligned} \frac{\partial f(x_0)}{\partial x_j} &\approx \frac{f(x_{01}, \dots, \mathbf{x}_{0j} + \mathbf{h}, \dots, x_{0m}) - f(x_{01}, \dots, \mathbf{x}_{0j}, \dots, x_{0m})}{\mathbf{h}} \\ &= \frac{f(x_0 + h e_j) - f(x_0)}{\mathbf{h}} \\ &= \frac{df(x_0 + x_j e_j)}{dx_j}, \end{aligned} \tag{11.23}$$

where the last expression is kind of awesome: it underlines that we have really fixed all of the components of x EXCEPT for x_j when we compute the partial derivative with respect to the j -th component of x ; therefore, $\bar{f}(x_j) := f(x_0 + x_j e_j)$ is now a function of the scalar variable x_j to which we can apply the ordinary derivative. The fact that $\bar{f}(x_j) \in \mathbb{R}^n$ means it has n -components instead of one has not really changed anything. For us, working with vectors or scalars, it's all the same.

11.5.3 The Jacobian and Linear Approximation of a Function about a Point

We now turn to functions $f : \mathbb{R}^m \rightarrow \mathbb{R}^n$. Based on the compact notation introduced in (11.19), we define the **Jacobian of a function** as

$$\frac{\partial f(x)}{\partial x} := \begin{bmatrix} \frac{\partial f(x)}{\partial x_1} & \frac{\partial f(x)}{\partial x_2} & \dots & \frac{\partial f(x)}{\partial x_m} \end{bmatrix} \tag{11.24}$$

by packaging the column vectors $\frac{\partial f(x)}{\partial x_j}$ into a matrix. There is no mystery here as we discovered partial derivatives as the columns of a matrix back in (11.17). We need to keep in mind that, for each value of $x \in \mathbb{R}^m$, the compact and innocent looking object

$$\frac{\partial f(x)}{\partial x}$$

is really an $n \times m$ matrix: once again, there are m columns of the form $\frac{\partial f(x)}{\partial x_j}$, and each column is an n -vector. When computing the Jacobian numerically, we typically build it up one column at a time as we did in Solution A to Example 11.4.

Just for the record, we will write out $\frac{\partial f(x)}{\partial x}$ as an $n \times m$ matrix. We write

$$f(x) = \begin{bmatrix} f_1(x) \\ f_2(x) \\ \vdots \\ f_n(x) \end{bmatrix} = \begin{bmatrix} f_1(x_1, x_2, \dots, x_m) \\ f_2(x_1, x_2, \dots, x_m) \\ \vdots \\ f_n(x_1, x_2, \dots, x_m) \end{bmatrix}.$$

Writing out all of the entries in the $n \times m$ Jacobian matrix gives

$$\frac{\partial f(x)}{\partial x} = \begin{bmatrix} \frac{\partial f_1(x)}{\partial x_1} & \frac{\partial f_1(x)}{\partial x_2} & \dots & \frac{\partial f_1(x)}{\partial x_m} \\ \frac{\partial f_2(x)}{\partial x_1} & \frac{\partial f_2(x)}{\partial x_2} & \dots & \frac{\partial f_2(x)}{\partial x_m} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial f_n(x)}{\partial x_1} & \frac{\partial f_n(x)}{\partial x_2} & \dots & \frac{\partial f_n(x)}{\partial x_m} \end{bmatrix}. \tag{11.25}$$

In other symbols, the ij component of $\frac{\partial f(x)}{\partial x}$ is

$$\left[\frac{\partial f(x)}{\partial x} \right]_{ij} = \frac{\partial f_i(x)}{\partial x_j},$$

which is much more intimidating than (11.24).

Perhaps it is better not to read any further? If you want to toss out the window all of the benefits of vector-matrix notation, you can compute each entry of the Jacobian matrix one by one, as in

$$\frac{\partial f_i(x)}{\partial x_j} \approx \frac{f_i(x_1, \dots, x_j + h, \dots, x_m) - f_i(x_1, \dots, x_j - h, \dots, x_m)}{2h}. \quad (11.26)$$

In case you are wondering, your instructors almost never do this when doing real robotics! We use the “vector version” of the Jacobian where we compute each column of the matrix in a loop. However, for simple examples, such as $n = m = 2$, the above very explicit scalar (means non-vector) calculations can be informative!

Linear Approximation at a Point for Functions of Vectors

The linear approximation of a (nonlinear) function $f : \mathbb{R}^m \rightarrow \mathbb{R}^n$ at a point x_0 is defined to be

$$f(x) \approx f(x_0) + A(x - x_0) = f(x_0) + \frac{\partial f(x_0)}{\partial x}(x - x_0), \quad (11.27)$$

where the $n \times m$ matrix A is the Jacobian of f at the point x_0 . As we did previously, we note that the dimensions make sense because $f(x_0)$ is $n \times 1$, $(x - x_0)$ is $m \times 1$, and therefore, $\frac{\partial f(x_0)}{\partial x}(x - x_0)$ is $n \times 1$.

Example 11.6 For the function

$$f(x_1, x_2, x_3) := \begin{bmatrix} x_1 x_2 x_3 \\ \log(2 + \cos(x_1)) + x_2^{x_1} \\ \frac{x_1 x_3}{1 + x_2^2} \end{bmatrix}, \quad (11.28)$$

compute its Jacobian at the point

$$x_0 = \begin{bmatrix} \pi \\ 1.0 \\ 2.0 \end{bmatrix}$$

and evaluate the “accuracy” of its linear approximation.

Solution From (11.21) in Example 11.4, we have that

$$\frac{\partial f(x_0)}{\partial x_1} = \begin{bmatrix} 2.0 \\ 0.0 \\ 1.0 \end{bmatrix}, \quad \frac{\partial f(x_0)}{\partial x_2} = \begin{bmatrix} 6.2832 \\ 3.1416 \\ -3.1416 \end{bmatrix}, \quad \frac{\partial f(x_0)}{\partial x_3} = \begin{bmatrix} 3.1416 \\ 0.0000 \\ 1.5708 \end{bmatrix}.$$

Hence, packaging up the columns correctly gives the Jacobian at x_0 ,

$$A := \frac{\partial f(x_0)}{\partial x} = \begin{bmatrix} 2.0000 & 6.2832 & 3.1416 \\ 0.0000 & 3.1416 & 0.0000 \\ 1.0000 & -3.1416 & 1.5708 \end{bmatrix},$$

and the linear approximation is

$$f(x) \approx f(x_0) + A(x - x_0) = \begin{bmatrix} 6.2832 \\ 1.0000 \\ 3.1416 \end{bmatrix} + \begin{bmatrix} 2.0000 & 6.2832 & 3.1416 \\ 0.0000 & 3.1416 & 0.0000 \\ 1.0000 & -3.1416 & 1.5708 \end{bmatrix} \begin{bmatrix} x_1 - \pi \\ x_2 - 1.0 \\ x_3 - 2.0 \end{bmatrix}.$$

1 **Jacf** = [dfdx1 dfdx2 dfdx3]

One way to think about the question of assessing the quality of the linear approximation is to measure the error defined as

$$e(x) := \|f(x) - f_{\text{lin}}(x)\|,$$

where $f_{\text{lin}}(x) := f(x_0) + \frac{\partial f(x_0)}{\partial x}(x - x_0)$. We will seek to estimate the maximum value of $e(x)$ over a region containing the point x_0 . Define

$$S(x_0) := \{x \in \mathbb{R}^3 \mid |x_i - x_{0i}| \leq d, i = 1, 2, 3\}$$

and

$$\text{Max Error} := \max_{x \in S(x_0)} e(x) = \max_{x \in S(x_0)} \|f(x) - f_{\text{lin}}(x)\|. \quad (11.29)$$

For $d = 0.25$, we used a “random search” routine and estimated that

$$\text{Max Error} = 0.12.$$

To put this into context,

$$\max_{x \in S(x_0)} \|f(x)\| = 8.47,$$

and thus the relative error is about 1.5%. ■

11.5.4 The Gradient and Linear Approximation of a Function about a Point

The **gradient** is simply the special name given to the Jacobian of a function $f : \mathbb{R}^m \rightarrow \mathbb{R}$, that is, for each $x \in \mathbb{R}^m$, $f(x) \in \mathbb{R}$ is a scalar. Along with its special name, it comes with a special symbol!

The Gradient and Linear Approximations

The **gradient** of $f : \mathbb{R}^m \rightarrow \mathbb{R}$ is simply the partial derivatives of f with respect to x_i arranged to form a **row vector**,

$$\nabla f(x_0) := \left[\frac{\partial f(x_0)}{\partial x_1} \quad \frac{\partial f(x_0)}{\partial x_2} \quad \dots \quad \frac{\partial f(x_0)}{\partial x_m} \right], \quad (11.30)$$

which we can also call a $1 \times m$ matrix. The cool symbol ∇ is usually pronounced as “**grad**” and one says “**grad f**” when speaking of ∇f .

An important use of the gradient of a function of several variables is to form a **linear approximation of the function about a point**

$$f(x) \approx f(x_0) + \nabla f(x_0)(x - x_0). \quad (11.31)$$

Comparing this to (11.15), we see that $A = \nabla f(x_0)$, a $1 \times m$ matrix. Expanding (11.31) into its components gives

$$\begin{aligned} f(x) &\approx f(x_0) + \nabla f(x_0)(x - x_0) \\ &= f(x_0) + \underbrace{\left[\frac{\partial f(x_0)}{\partial x_1} \quad \frac{\partial f(x_0)}{\partial x_2} \quad \dots \quad \frac{\partial f(x_0)}{\partial x_m} \right]}_A \cdot \underbrace{\begin{bmatrix} x_1 - x_{01} \\ x_2 - x_{02} \\ \vdots \\ x_m - x_{0m} \end{bmatrix}}_{(x-x_0)} \\ &= f(x_0) + \sum_{i=1}^m \frac{\partial f(x_0)}{\partial x_i} (x_i - x_{0i}). \end{aligned} \quad (11.32)$$

The linear approximation in (11.32) looks just like our linear approximation for functions of a single variable x , namely $f(x) \approx f(x_0) + \frac{df(x_0)}{dx}(x - x_0)$, where $a = \frac{df(x_0)}{dx}$ is 1×1 .

Example 11.7 Compute (approximately) the gradient of $f : \mathbb{R}^2 \rightarrow \mathbb{R}$, for $f(x_1, x_2) = x_1 \cos(x_2)$ and $x_0 = [2 \ \pi/4]^\top$.




Figure 11.10: (a) Close up view. (b) A more global view. The function $f(x_1, x_2) = x_1 \cos(x_2)$ is plotted in blue, while in red is shown its linear approximation about the point $x_0 = [2 \ \pi/4]^\top$, that is, $f_{\text{lin}}(x) := f(x_0) + \nabla f(x_0)(x - x_0)$. This approximation can be done for any point x_0 at which the partial derivatives exists. In Calculus, the red plane is also called the **tangent plane at x_0** . These linear approximations accurately represent the nonlinear function in a small enough region about a given point, allowing us to use our Linear Algebra skills.

Solution: Let's get the formulas for a general $h > 0$ and then we'll build a Table comparing the results for several values of h . For the partial derivative with respect to x_1 , we perturb x_1 about 2 while holding x_2 constant and equal to $\pi/4$. This gives

$$\begin{aligned} \frac{\partial f(x_0)}{\partial x_1} &\approx \frac{f(2+h, \pi/4) - f(2-h, \pi/4)}{2h} \\ &= \frac{(2+h) \cos(\pi/4) - (2-h) \cos(\pi/4)}{2h} \\ &= \frac{2h \cos(\pi/4)}{2h} \\ &= \cos(\pi/4) = \frac{\sqrt{2}}{2}, \end{aligned}$$

which is independent of h and hence there is nothing more to compute. For the next partial derivative, we perturb x_2 about $\pi/4$ while holding x_1 constant and equal to 2. This gives

$$\begin{aligned} \frac{\partial f(x_0)}{\partial x_2} &\approx \frac{f(2, \pi/4+h) - f(2, \pi/4-h)}{2h} \\ &= \frac{2 \cos(\pi/4+h) - 2 \cos(\pi/4-h)}{2h}, \end{aligned} \tag{11.33}$$

which, though we could simply it with some trigonometric identities, we'll stop here and turn to Julia. Doing so, leads to the following results

$\frac{f(2, \pi/4+h) - f(2, \pi/4-h)}{2h}$	h	
-1.41421356001592	0.0001	
-1.414213326670799	0.001	
-1.4141899922649026	0.01	
-1.4118577179998826	0.1	
-1.4048043101898116	0.2	
-1.3768017528243548	0.4	

(11.34)

The true answer is

$$\frac{\partial f(2, \pi/4)}{\partial x_2} = -\sqrt{2} \approx -1.4142135623730951$$

Example 11.8 Compute a linear approximation of $f(x_1, x_2) = x_1 \cos(x_2)$ at the point $x_0 = [2 \ \pi/4]^\top$.

Solution: We know both of the partial derivatives of f at the point $x_0 = [2 \ \pi/4]^\top$. Hence, we have

$$\begin{aligned} f(x) &\approx f(x_0) + \nabla f(x_0)(x - x_0) \\ &= f(x_0) + \begin{bmatrix} \frac{\partial f(x_0)}{\partial x_1} & \frac{\partial f(x_0)}{\partial x_2} \end{bmatrix} \begin{bmatrix} x_1 - x_{01} \\ x_2 - x_{02} \end{bmatrix} \\ &= 2 \cos(\pi/4) + \begin{bmatrix} \frac{\sqrt{2}}{2} & -\sqrt{2} \end{bmatrix} \begin{bmatrix} x_1 - 2 \\ x_2 - \pi/4 \end{bmatrix} \end{aligned} \quad (11.35)$$

Figure 11.10 compares the linear approximation to the nonlinear function in a region about x_0 . ■

Knowledge is Power

For root finding, an accuracy of a few percent is probably good enough. That said, **learning how to compute the partial derivatives analytically will make your code faster and will eliminate the question of how to choose the small perturbation parameter $h > 0$.**

11.5.5 Summary on Partial Derivatives

From slopes of lines \rightarrow slopes of functions at points $\rightarrow \frac{df(x_0)}{dx} \rightarrow \nabla f(x_0) \rightarrow \frac{\partial f(x_0)}{\partial x}$

Derivatives, gradients, and Jacobians are all generalizations of the notion of the slope of a line being its “rise over run”. The derivative of a function $f : \mathbb{R} \rightarrow \mathbb{R}$ at a point x_0 is the “local” slope of the function at that point. We compute it by “rise over run”, where, for example

$$\text{slope} = \frac{f(x_0 + h) - f(x_0)}{h} \xrightarrow{h > 0 \text{ small}} \frac{df(x_0)}{dx}, \quad (11.36)$$

and to make it local, we take $h > 0$ small. The gradient recognizes that a function $f : \mathbb{R}^m \rightarrow \mathbb{R}$ has a local slope in the x_1 -direction, the x_2 -direction, all the way up to the x_m -direction. If we let $\{e_1, e_2, \dots, e_m\}$ be the natural basis vectors for \mathbb{R}^m , then we compute each component of the gradient by

$$\text{slope}_j = \frac{f(x_0 + he_j) - f(x_0)}{h} \xrightarrow{h > 0 \text{ small}} \frac{\partial f(x_0)}{\partial x_j}, \quad (11.37)$$

and we assemble the gradient by

$$\nabla f(x_0) := \begin{bmatrix} \frac{\partial f(x_0)}{\partial x_1} & \frac{\partial f(x_0)}{\partial x_2} & \dots & \frac{\partial f(x_0)}{\partial x_m} \end{bmatrix}. \quad (11.38)$$

Finally, for $f : \mathbb{R}^m \rightarrow \mathbb{R}^n$, each of the n components of $f(x)$ has a local slope with respect to each component of x . The bookkeeping is easiest if we leave $f(x)$ as a vector and write the Jacobian so that it **looks just like** the gradient,

$$\frac{\partial f(x_0)}{\partial x} := \begin{bmatrix} \frac{\partial f(x_0)}{\partial x_1} & \frac{\partial f(x_0)}{\partial x_2} & \dots & \frac{\partial f(x_0)}{\partial x_m} \end{bmatrix}, \quad (11.39)$$

but now, because $f(x)$ is an n -vector, we have

$$\begin{bmatrix} \text{slope}_{1j} \\ \vdots \\ \text{slope}_{ij} \\ \vdots \\ \text{slope}_{nj} \end{bmatrix} = \frac{f(x_0 + he_j) - f(x_0)}{h} \xrightarrow{h > 0 \text{ small}} \frac{\partial f(x_0)}{\partial x_j}. \quad (11.40)$$

Compact Way to Numerically Approximate the Jacobian

If we let $\{e_1, e_2, \dots, e_m\}$ be the natural basis vectors for \mathbb{R}^m , then vector notation allows us to compute each column of the Jacobian by

$$\frac{\partial f(x_0)}{\partial x_j} = \frac{f(x_0 + he_j) - f(x_0)}{h}. \quad (11.41)$$

Typically, this is easier to program up than (11.26). But of course, your experience may vary! For a symmetric difference, we'd use

$$\frac{\partial f(x_0)}{\partial x_j} = \frac{f(x_0 + he_j) - f(x_0 - he_j)}{2h}. \quad (11.42)$$

11.6 Newton-Raphson for Vector Functions

We consider functions $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ and seek a root $f(x_0) = 0$. Note that the domain and range are both \mathbb{R}^n and thus this is the nonlinear equivalent of solving a square linear equation $Ax - b = 0$. We recall that $\det(A) \neq 0$ was our magic condition for the existence and uniqueness of solutions to $Ax - b = 0$.

The **Newton-Raphson** Algorithm is precisely a vector version of Newton's Algorithm. Let x_k be our current approximation of a root of the function f . We write the linear approximation of f about the point x_k as

$$f(x) \approx f(x_k) + \frac{\partial f(x_k)}{\partial x} \cdot (x - x_k). \quad (11.43)$$

We want to choose x_{k+1} so that $f(x_{k+1}) = 0$. Based on our linear approximation in (11.43), we have that

$$f(x_{k+1}) \approx 0 \iff 0 \approx f(x_k) + \frac{\partial f(x_k)}{\partial x} \cdot (x_{k+1} - x_k). \quad (11.44)$$

If $\det\left(\frac{\partial f(x_k)}{\partial x}\right) \neq 0$ we could naively solve for x_{k+1} , giving us

$$x_{k+1} = x_k - \left(\frac{\partial f(x_k)}{\partial x}\right)^{-1} f(x_k).$$

By now, you know that we advise against blindly computing inverses of matrices unless you really need that matrix inverse. In our case, what we really want is x_{k+1} , and because we know to avoid computing unnecessary matrix inverses, we'll write the algorithm

down in a different way. As in the scalar case, the equation is screaming for us to put it in a loop!

Newton-Raphson Algorithm

Based on (11.44), we define^a $x_{k+1} := x_k + \Delta x_k$, where Δx_k is our update to x_k . We can then break the algorithm into two steps,

$$\left(\frac{\partial f(x_k)}{\partial x}\right) \Delta x_k = -f(x_k) \quad (\text{solve for } \Delta x_k) \quad (11.45)$$

$$x_{k+1} = x_k + \Delta x_k \quad (\text{use } \Delta x_k \text{ to update our estimate of the root}). \quad (11.46)$$

While for toy problems, we can use the matrix inverse to solve (11.45) for Δx_k , for larger problems, we recommend using an LU Factorization or a QR Factorization. Once (11.45) has been solved, x_{k+1} is updated in (11.46) and the process repeats.

A **damped Newton-Raphson Algorithm** is obtained by replacing (11.46) with

$$x_{k+1} = x_k + \epsilon \Delta x_k, \quad (11.47)$$

for some $\epsilon > 0$. The validity of the Newton-Raphson Algorithm rests upon:

- the function f being differentiable;
- the Jacobian $\frac{\partial f(x_k)}{\partial x}$ having a non-zero determinant at points generated by (11.45) and (11.46); and
- the linear equation $f_{\text{lin}}(x) = f(x_k) + \frac{\partial f(x_k)}{\partial x}(x - x_k)$ being a good approximation to the function.

^aNote that $\Delta x_k = x_{k+1} - x_k$.

Example 11.9 Find a root of $F : \mathbb{R}^4 \rightarrow \mathbb{R}^4$ near $x_0 = [-2.0 \quad 3.0 \quad \pi \quad -1.0]$ for

$$F(x) = \begin{bmatrix} x_1 + 2x_2 - x_1(x_1 + 4x_2) - x_2(4x_1 + 10x_2) + 3 \\ 3x_1 + 4x_2 - x_1(x_1 + 4x_2) - x_2(4x_1 + 10x_2) + 4 \\ 0.5 \cos(x_1) + x_3 - (\sin(x_3))^7 \\ -2(x_2)^2 \sin(x_1) + (x_4)^3 \end{bmatrix}.$$

Solution: We programmed up (11.45) and (11.46) in Julia and used a symmetric difference approximation for the derivatives, with $h = 0.1$. Below are the first five results from the algorithm:

$$x_k = \begin{bmatrix} k=0 & k=1 & k=2 & k=3 & k=4 & k=5 \\ -2.0000 & -3.0435 & -2.4233 & -2.2702 & -2.2596 & -2.2596 \\ 3.0000 & 2.5435 & 1.9233 & 1.7702 & 1.7596 & 1.7596 \\ 3.1416 & 0.6817 & 0.4104 & 0.3251 & 0.3181 & 0.3181 \\ -1.0000 & -1.8580 & -2.0710 & -1.7652 & -1.6884 & -1.6846 \end{bmatrix}$$

and

$$f(x_k) = \begin{bmatrix} k=0 & k=1 & k=2 & k=3 & k=4 & k=5 \\ -39.0000 & -6.9839 & -1.1539 & -0.0703 & -0.0003 & -0.0000 \\ -36.0000 & -6.9839 & -1.1539 & -0.0703 & -0.0003 & -0.0000 \\ 2.9335 & 0.1447 & 0.0323 & 0.0028 & 0.0000 & -0.0000 \\ 15.3674 & -5.1471 & -4.0134 & -0.7044 & -0.0321 & -0.0001 \end{bmatrix}.$$

By iteration five, we have a good approximation of a root because $\|f(x_5)\| \approx 10^{-4}$. We also provide the Jacobians at the initial and final steps,

$$\frac{\partial f(x_0)}{\partial x} = \begin{bmatrix} -19.0000 & -42.0000 & 0.0000 & 0.0000 \\ -17.0000 & -40.0000 & 0.0000 & 0.0000 \\ 0.4539 & 0.0000 & 1.0000 & 0.0000 \\ 7.4782 & 10.9116 & 0.0000 & 3.0100 \end{bmatrix} \quad \text{and} \quad \frac{\partial f(x_5)}{\partial x} = \begin{bmatrix} -8.5577 & -15.1155 & 0.0000 & 0.0000 \\ -6.5577 & -13.1155 & 0.0000 & 0.0000 \\ 0.3854 & 0.0000 & 0.9910 & 0.0000 \\ 3.9296 & 5.4337 & 0.0000 & 8.5616 \end{bmatrix}$$

so that it is clear that as x_k evolves, so does the Jacobian of f at x_k . ■

11.7 (Optional Read): From the Gradient or Jacobian of a Function to its Linear Approximation

This covers the same material as in Chap. 11.5, but in reverse order. Some may find it more digestible.

Consider a function $f : \mathbb{R}^m \rightarrow \mathbb{R}^n$. We seek a means to build a linear approximation of the function near a given point $x_0 \in \mathbb{R}^m$. When $m = n = 1$, we were able to approximate a function by

$$f(x) \approx f(x_0) + \frac{df(x_0)}{dx}(x - x_0).$$

In the above, $\frac{df(x_0)}{dx}$ is a scalar. For reasons that will become clear shortly, let's denote that scalar by $a := \frac{df(x_0)}{dx}$, so that we can rewrite the linear approximation as

$$f(x) \approx f(x_0) + a(x - x_0). \quad (11.48)$$

We do this and note that a can be viewed as a 1×1 matrix, that is, an $n \times m$ matrix for $n = m = 1$. We now ask the question, for n and m not necessarily equal to one, and for $f : \mathbb{R}^m \rightarrow \mathbb{R}^n$, can we find an $n \times m$ matrix A such that

$$f(x) \approx f(x_0) + A(x - x_0). \quad (11.49)$$

The answer is (mostly) yes. To do this, we need to extend the notion of a derivative to the case of vectors. We do this first for $n = 1$ and general $m \geq 1$.

11.7.1 The Gradient

We restrict ourselves to functions $f : \mathbb{R}^m \rightarrow \mathbb{R}$. Hence, for $x \in \mathbb{R}^m$, we have $f(x) \in \mathbb{R}$ and we will make the components of x explicit in the function by writing $f(x) = f(x_1, \dots, x_m)$. One obvious way to extend our notion of a derivative is to perturb the components of x one at a time. In Calculus, these are called **partial derivatives**. We won't try to justify the terminology; it is what it is.

We continue to use $h \in \mathbb{R}$ to denote a small non-zero real number. With this notation, we define the **partial derivative of f with respect to x_i at a point x_0** to be

$$\frac{\partial f(x_0)}{\partial x_i} \approx \frac{f(x_{01}, \dots, x_{0i} + h, \dots, x_{0m}) - f(x_{01}, \dots, x_{0i}, \dots, x_{0m})}{h}, \quad (11.50)$$

where we have highlighted that the increment is applied to x_i and only to x_i . What we are doing is holding constant all coordinates except the i -th one, and then applying the "usual" definition of a derivative of a function that depends on the scalar variable x_i .

Equation (11.50) is a **forward difference approximation** of the partial derivative with respect to x_i about the point x_0 . Just as with our previous treatment of the derivative, we can use a backward approximation or a **symmetric difference approximation**, such as

$$\frac{\partial f(x_0)}{\partial x_i} \approx \frac{f(x_{01}, \dots, x_{0i} + h, \dots, x_{0m}) - f(x_{01}, \dots, x_{0i} - h, \dots, x_{0m})}{2h}. \quad (11.51)$$

Example 11.10 Compute (approximately) the partial derivatives of $f(x_1, x_2) = x_1 \cos(x_2)$ with respect to both x_1 and x_2 about the point $x_0 = [2 \ \pi/4]^T$.

Solution: Let's get the formulas for a general $h > 0$ and then we'll build a Table comparing the results for several values of h . For the partial derivative with respect to x_1 , we perturb x_1 about 2 while holding x_2 constant and equal to $\pi/4$. This gives

$$\begin{aligned} \frac{\partial f(x_0)}{\partial x_1} &\approx \frac{f(2 + h, \pi/4) - f(2 - h, \pi/4)}{2h} \\ &= \frac{(2 + h) \cos(\pi/4) - (2 - h) \cos(\pi/4)}{2h} \\ &= \frac{2h \cos(\pi/4)}{2h} \\ &= \cos(\pi/4) = \frac{\sqrt{2}}{2}, \end{aligned}$$

which is independent of h and hence there is nothing more to compute. For the next partial derivative, we perturb x_2 about $\pi/4$ while holding x_1 constant and equal to 2. This gives

$$\begin{aligned}\frac{\partial f(x_0)}{\partial x_2} &\approx \frac{f(2, \pi/4 + h) - f(2, \pi/4 - h)}{2h} \\ &= \frac{2 \cos(\pi/4 + h) - 2 \cos(\pi/4 - h)}{2h},\end{aligned}\tag{11.52}$$

which, though we could simply it with some trigonometric identities, we'll stop here and turn to Julia. Doing so, leads to the following results

$\frac{f(2, \pi/4+h) - f(2, \pi/4-h)}{2h}$	h	
-1.41421356001592	0.0001	
-1.414213326670799	0.001	
-1.4141899922649026	0.01	
-1.4118577179998826	0.1	
-1.4048043101898116	0.2	
-1.3768017528243548	0.4	(11.53)

The true answer is

$$\frac{\partial f(2, \pi/4)}{\partial x_2} = -\sqrt{2} \approx -1.4142135623730951$$



Knowledge is Power

For root finding, an accuracy of a few percent is probably good enough. That said, **learning how to compute the partial derivatives analytically will make your code faster and will eliminate the question of how to choose the small perturbation parameter $h > 0$.**

The Gradient and Linear Approximations

The **gradient** of $f : \mathbb{R}^m \rightarrow \mathbb{R}$ is simply the partial derivatives of f with respect to x_i arranged to form a **row vector**,

$$\nabla f(x_0) := \left[\frac{\partial f(x_0)}{\partial x_1} \quad \frac{\partial f(x_0)}{\partial x_2} \quad \dots \quad \frac{\partial f(x_0)}{\partial x_m} \right],\tag{11.54}$$

which we can also call a $1 \times m$ matrix. The cool symbol ∇ is usually pronounced as “**grad**” and one says “**grad f**” when speaking of ∇f .

An important use of the gradient of a function of several variables is to form a **linear approximation of the function about a point**

$$f(x) \approx f(x_0) + \nabla f(x_0)(x - x_0).\tag{11.55}$$

Comparing this to (11.49), we see that $A = \nabla f(x_0)$, a $1 \times m$ matrix. Expanding (11.55) into its components gives

$$\begin{aligned}f(x) &\approx f(x_0) + \nabla f(x_0)(x - x_0) \\ &= f(x_0) + \underbrace{\left[\frac{\partial f(x_0)}{\partial x_1} \quad \frac{\partial f(x_0)}{\partial x_2} \quad \dots \quad \frac{\partial f(x_0)}{\partial x_m} \right]}_A \cdot \underbrace{\begin{bmatrix} x_1 - x_{01} \\ x_2 - x_{02} \\ \vdots \\ x_m - x_{0m} \end{bmatrix}}_{(x-x_0)} \\ &= f(x_0) + \sum_{i=1}^m \frac{\partial f(x_0)}{\partial x_i} (x_i - x_{0i}).\end{aligned}\tag{11.56}$$

The linear approximation in (11.56) looks just like our linear approximation for functions of a single variable x , namely $f(x) \approx f(x_0) + \frac{df(x_0)}{dx}(x - x_0)$, where $a = \frac{df(x_0)}{dx}$ is 1×1 .




Figure 11.11: (a) Close up view. (b) A more global view. The function $f(x_1, x_2) = x_1 \cos(x_2)$ is plotted in blue, while in red is shown its linear approximation about the point $x_0 = [2 \ \pi/4]^\top$, that is, $f_{\text{lin}}(x) := f(x_0) + \nabla f(x_0)(x - x_0)$. This approximation can be done for any point x_0 at which the partial derivatives exists. In Calculus, the red plane is also called the **tangent plane at x_0** . These linear approximations accurately represent the nonlinear function in a small enough region about a given point, allowing us to use our Linear Algebra skills.

Example 11.11 Compute a linear approximation of $f(x_1, x_2) = x_1 \cos(x_2)$ at the point $x_0 = [2 \ \pi/4]^\top$.

Solution: From Example 11.10, we know both of the partial derivatives of f at the point $x_0 = [2 \ \pi/4]^\top$. Hence, we have

$$\begin{aligned}
 f(x) &\approx f(x_0) + \nabla f(x_0)(x - x_0) \\
 &= f(x_0) + \begin{bmatrix} \frac{\partial f(x_0)}{\partial x_1} & \frac{\partial f(x_0)}{\partial x_2} \end{bmatrix} \begin{bmatrix} x_1 - x_{01} \\ x_2 - x_{02} \end{bmatrix} \\
 &= 2 \cos(\pi/4) + \begin{bmatrix} \frac{\sqrt{2}}{2} & -\sqrt{2} \end{bmatrix} \begin{bmatrix} x_1 - 2 \\ x_2 - \pi/4 \end{bmatrix}
 \end{aligned} \tag{11.57}$$

Figure 11.11 compares the linear approximation to the nonlinear function in a region about x_0 . ■

11.7.2 Expanding on Vector Notation

We have carefully defined derivatives and partial derivatives of functions at given points. We used the notation x_0 for the given point to make it seem like some concrete value, a fixed scalar in \mathbb{R} or vector in \mathbb{R}^m . However, in practice, such as in Newton's Algorithm, we keep updating the point x_0 at which we are computing derivatives and linear approximations of functions. Hence, because the point is not really fixed, we can just call it x . Then, for $f : \mathbb{R} \rightarrow \mathbb{R}$, we have

$$\frac{df(x)}{dx} \approx \frac{f(x+h) - f(x-h)}{2h}, \tag{11.58}$$

and for $f : \mathbb{R}^m \rightarrow \mathbb{R}$, where $x = (x_1, x_2, \dots, x_m)$, we have

$$\frac{\partial f(x)}{\partial x_i} \approx \frac{f(x_1, \dots, x_i + h, \dots, x_m) - f(x_1, \dots, x_i - h, \dots, x_m)}{2h}. \tag{11.59}$$

With this notation in mind, we now define **derivatives and partial derivatives for vector valued functions**. When $f : \mathbb{R} \rightarrow \mathbb{R}^n$, we have that x is a scalar and $f(x)$ is a vector with n components, as in

$$f(x) = \begin{bmatrix} f_1(x) \\ f_2(x) \\ \vdots \\ f_n(x) \end{bmatrix}. \tag{11.60}$$

We define its derivative with respect to the scalar variable x as

$$\frac{df(x)}{dx} := \begin{bmatrix} \frac{df_1(x)}{dx} \\ \frac{df_2(x)}{dx} \\ \vdots \\ \frac{df_n(x)}{dx} \end{bmatrix}, \quad (11.61)$$

by differentiating each of the components of $f(x)$. It's the obvious thing to do. The key is that you must keep track that when $f(x)$ is vector valued, so is its derivative, $\frac{df(x)}{dx}$.

In terms of our symmetric difference approximation to the derivative, the formula,

$$\frac{df(x)}{dx} \approx \frac{f(x+h) - f(x-h)}{2h}, \quad (11.62)$$

is still valid and is even written in Julia the same way! Yes, you could re-write the above as

$$\frac{df(x)}{dx} \approx \begin{bmatrix} \frac{f_1(x+h) - f_1(x-h)}{2h} \\ \frac{f_2(x+h) - f_2(x-h)}{2h} \\ \vdots \\ \frac{f_n(x+h) - f_n(x-h)}{2h} \end{bmatrix}, \quad (11.63)$$

but that's a lot of extra programming. The expression (11.62) is much more compact and convenient. This is the power of good notation.

Similarly, when $f : \mathbb{R}^m \rightarrow \mathbb{R}^n$, we know that x is a vector with m components and $f(x)$ is a vector with n components, as in

$$f(x) = \begin{bmatrix} f_1(x_1, \dots, x_m) \\ f_2(x_1, \dots, x_m) \\ \vdots \\ f_n(x_1, \dots, x_m) \end{bmatrix}. \quad (11.64)$$

We define its partial derivative with respect to component x_i as

$$\frac{\partial f(x)}{\partial x_i} := \begin{bmatrix} \frac{\partial f_1(x)}{\partial x_i} \\ \frac{\partial f_2(x)}{\partial x_i} \\ \vdots \\ \frac{\partial f_n(x)}{\partial x_i} \end{bmatrix}, \quad (11.65)$$

by doing the partial differentiation of each of the components of $f(x)$. It's the obvious thing to do.

In terms of our symmetric difference approximation to the derivative, the formula,

$$\frac{\partial f(x)}{\partial x_i} \approx \frac{f(x_1, \dots, x_i+h, \dots, x_m) - f(x_1, \dots, x_i-h, \dots, x_m)}{2h}, \quad (11.66)$$

is still valid and is even written in Julia the same way! Once again, you could re-write the above as

$$\frac{\partial f(x)}{\partial x_i} \approx \begin{bmatrix} \frac{f_1(x_1, \dots, x_i+h, \dots, x_m) - f_1(x_1, \dots, x_i-h, \dots, x_m)}{2h} \\ \frac{f_2(x_1, \dots, x_i+h, \dots, x_m) - f_2(x_1, \dots, x_i-h, \dots, x_m)}{2h} \\ \vdots \\ \frac{f_n(x_1, \dots, x_i+h, \dots, x_m) - f_n(x_1, \dots, x_i-h, \dots, x_m)}{2h} \end{bmatrix}, \quad (11.67)$$

but that's a lot of extra programming. The expression (11.66) is much more compact and convenient. **This is again the power of good notation. But to use the notation effectively, we have to do the bookkeeping and remember that the innocent expression**

$$\frac{\partial f(\mathbf{x})}{\partial \mathbf{x}_i}$$

is really a vector with n components.

11.7.3 The Jacobian

We now turn to functions $f : \mathbb{R}^m \rightarrow \mathbb{R}^n$. Based on the compact notation covered in Chapter 11.5.5, we can define the **Jacobian of a function** as

$$\frac{\partial f(x)}{\partial x} := \left[\frac{\partial f(x)}{\partial x_1} \quad \frac{\partial f(x)}{\partial x_2} \quad \dots \quad \frac{\partial f(x)}{\partial x_m} \right]. \quad (11.68)$$

Comparing the above to (11.54), we see that the **Jacobian** of a function $f : \mathbb{R}^m \rightarrow \mathbb{R}$ is the **gradient** of the function. **This is the same as saying an $n \times m$ matrix reduces to a row vector when $n = 1$.**

We need to keep in mind that, for each value of $x \in \mathbb{R}^m$, the compact and innocent looking object

$$\frac{\partial f(x)}{\partial x}$$

is really an $n \times m$ matrix. When computing it numerically, we typically build it up one column at a time as in the following Julia code.

```

1 F(x1, x2, x3)=[x1.*x2.*x3; log.(2+cos.(x1)) .+ x2.^x1; x1.*x3/(1 .+ x2.^2)]
2 h=0.01
3 x0=[pi; 1.0; 2.0]
4 dfdx1 = (F(x0[1]+h, x0[2], x0[3]) - F(x0[1]-h, x0[2], x0[3])) / (2*h)
5 dfdx2 = (F(x0[1], x0[2]+h, x0[3]) - F(x0[1], x0[2]-h, x0[3])) / (2*h)
6 dfdx3 = (F(x0[1], x0[2], x0[3]+h) - F(x0[1], x0[2], x0[3]-h)) / (2*h)
7 dfdx=[dfdx1 dfdx2 dfdx3]
8
9 3x3 Array{Float64, 2}:
10  2.0    6.28319   3.14159
11  0.0    3.14172   0.0
12  1.0   -3.14159   1.5708
13

```

Just for the record, we will write out $\frac{\partial f(x)}{\partial x}$ as an $n \times m$ matrix. In case you are wondering, your instructors never do this when doing real robotics! We use the “vector version” of the Jacobian where we compute each column of the matrix. For $f : \mathbb{R}^m \rightarrow \mathbb{R}^n$,

$$\frac{\partial f(x)}{\partial x} = \begin{bmatrix} \frac{\partial f_1(x)}{\partial x_1} & \frac{\partial f_1(x)}{\partial x_2} & \dots & \frac{\partial f_1(x)}{\partial x_m} \\ \frac{\partial f_2(x)}{\partial x_1} & \frac{\partial f_2(x)}{\partial x_2} & \dots & \frac{\partial f_2(x)}{\partial x_m} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial f_n(x)}{\partial x_1} & \frac{\partial f_n(x)}{\partial x_2} & \dots & \frac{\partial f_n(x)}{\partial x_m} \end{bmatrix}. \quad (11.69)$$

Hence, the ij component of $\frac{\partial f(x)}{\partial x}$ is $\frac{\partial f_i(x)}{\partial x_j}$. If one wishes, the Jacobian can be computed one element at a time via

$$\frac{\partial f_i(x)}{\partial x_j} \approx \frac{f_i(x_1, \dots, x_j + h, \dots, x_m) - f_i(x_1, \dots, x_j - h, \dots, x_m)}{2h}. \quad (11.70)$$

From slopes of lines \rightarrow slopes of functions at points $\rightarrow \frac{df(x_0)}{dx} \rightarrow \nabla f(x_0) \rightarrow \frac{\partial f(x_0)}{\partial x}$

Derivatives, gradients, and Jacobians are all generalizations of the notion of the slope of a line being its “rise over run”. The derivative of a function $f : \mathbb{R} \rightarrow \mathbb{R}$ at a point x_0 is the “local” slope of the function at that point. We compute it by “rise over run”, where, for example

$$\text{slope} = \frac{f(x_0 + h) - f(x_0)}{h} \xrightarrow{h > 0 \text{ small}} \frac{df(x_0)}{dx}, \quad (11.71)$$

and to make it local, we take $h > 0$ small. The gradient recognizes that a function $f : \mathbb{R}^m \rightarrow \mathbb{R}$ has a local slope in the x_1 -direction, the x_2 -direction, all the way up to the x_m -direction. If we let $\{e_1, e_2, \dots, e_m\}$ be the natural basis vectors for \mathbb{R}^m , then we compute each component of the gradient by

$$\text{slope}_j = \frac{f(x_0 + he_j) - f(x_0)}{h} \xrightarrow{h > 0 \text{ small}} \frac{\partial f(x_0)}{\partial x_j}, \quad (11.72)$$

and we assemble the gradient by

$$\nabla f(x_0) := \left[\frac{\partial f(x_0)}{\partial x_1} \quad \frac{\partial f(x_0)}{\partial x_2} \quad \dots \quad \frac{\partial f(x_0)}{\partial x_m} \right]. \quad (11.73)$$

Finally, for $f : \mathbb{R}^m \rightarrow \mathbb{R}^n$, each of the n components of $f(x)$ has a local slope with respect to each component of x . The bookkeeping is easiest if we leave $f(x)$ as a vector and write the Jacobian so that it **looks just like** the gradient,

$$\frac{\partial f(x_0)}{\partial x} := \left[\frac{\partial f(x_0)}{\partial x_1} \quad \frac{\partial f(x_0)}{\partial x_2} \quad \dots \quad \frac{\partial f(x_0)}{\partial x_m} \right], \quad (11.74)$$

but now, because $f(x)$ is an n -vector, we have

$$\begin{bmatrix} \text{slope}_{1j} \\ \vdots \\ \text{slope}_{ij} \\ \vdots \\ \text{slope}_{nj} \end{bmatrix} = \frac{f(x_0 + he_j) - f(x_0)}{h} \xrightarrow{h > 0 \text{ small}} \frac{\partial f(x_0)}{\partial x_j}. \quad (11.75)$$

Compact Way to Numerically Approximate the Jacobian

If we let $\{e_1, e_2, \dots, e_m\}$ be the natural basis vectors for \mathbb{R}^m , then a more compact notation allows us to compute each column of the Jacobian by

$$\frac{\partial f(x_0)}{\partial x_j} = \frac{f(x_0 + he_j) - f(x_0)}{h}. \quad (11.76)$$

Typically, this is easier to program up than (11.70). But of course, your experience may vary! For a symmetric difference, we’d use

$$\frac{\partial f(x_0)}{\partial x_j} = \frac{f(x_0 + he_j) - f(x_0 - he_j)}{2h}. \quad (11.77)$$

11.7.4 Linear Approximation of Vector Valued Functions

Our goal remains to build linear approximations of the form $f(x) \approx f(x_0) + A(x - x_0)$. Just as with our previous investigations, the matrix A is associated with derivatives of the function. In fact, we have

$$f(x) \approx f(x_0) + \underbrace{\frac{\partial f(x_0)}{\partial x}}_A (x - x_0), \quad (11.78)$$

in other words, $A := \left. \frac{\partial f(x)}{\partial x} \right|_{x=x_0}$.

Example 11.12 For the function

$$f(x_1, x_2, x_3) := \begin{bmatrix} x_1 x_2 x_3 \\ \log(2 + \cos(x_1)) + x_2^{x_1} \\ \frac{x_1 x_3}{1+x_2^2} \end{bmatrix}, \quad (11.79)$$

compute its Jacobian at the point

$$x_0 = \begin{bmatrix} \pi \\ 1.0 \\ 2.0 \end{bmatrix}$$

and evaluate the “accuracy” of its linear approximation.

Solution Using Symmetric Differences, the Jacobian at x_0 is

$$A = \frac{\partial f(x_0)}{\partial x} \approx \begin{bmatrix} 2.00 & 6.28 & 3.14 \\ 0.00 & 3.14 & 0.00 \\ 1.00 & -3.14 & 1.57 \end{bmatrix} \quad (11.80)$$

Figure 11.11 is the limit of what we can show in plots. Another way to think about the question of assessing the quality of the linear approximation is to measure the error defined as

$$e(x) := \|f(x) - f_{\text{lin}}(x)\|,$$

where $f_{\text{lin}}(x) := f(x_0) + \frac{\partial f(x_0)}{\partial x}(x - x_0)$. We will seek to estimate the maximum value of $e(x)$ over a region containing the point x_0 . Define

$$S(x_0) := \{x \in \mathbb{R}^3 \mid |x_i - x_{0i}| \leq d, i = 1, 2, 3\}$$

and

$$\text{Max Error} := \max_{x \in S(x_0)} e(x) = \max_{x \in S(x_0)} \|f(x) - f_{\text{lin}}(x)\|. \quad (11.81)$$

The for $d = 0.25$, we used a “random search” routine and estimate that

$$\text{Max Error} = 0.12.$$

To put this into context,

$$\max_{x \in S(x_0)} \|f(x)\| = 8.47,$$

and thus the relative error is about 1.5%. ■

If you made it to here, you should loop back to Chap. 11.6.

11.8 Looking Ahead

We’ve seen that an interesting idea from Calculus, called the derivative, allows nonlinear functions to be approximated by linear functions. The matrices resulting from a derivative, gradient, or Jacobian were instrumental in building algorithms to find roots of nonlinear equations.

In this next Chapter, we’ll use the derivative and the gradient to understand algorithms for finding a minimum or a maximum of a function. We’ll be able to greatly extend the ideas we developed in Chapters 8.2 for least squared error solutions to $Ax = b$ when it had no solutions, Chapter 8.3 for regression of functions to data, and Chapter 9.9, where we were able to identify a unique solution of minimum norm when $Ax = b$ had an infinite number of solutions.

Chapter 12

Changing Gears Again: Basic Ideas of Optimization

Learning Objectives

- Mathematics is used to describe physical phenomena, pose engineering problems, and solve engineering problems. We close our Y1-introduction to Computational Linear Algebra by showing how linear algebra and computation allow you to use a notion of “optimality” as a criterion for selecting among a set of solutions to an engineering problem.

Outcomes

- Arg min should be thought of as another function in your toolbox,

$$x^* = \arg \min_{x \in \mathbb{R}^m} f(x).$$

- Extrema of a function occur at places where the function’s first derivative vanishes.
- The gradient of a function points in the direction of maximum rate of growth.
- We will add to our knowledge of derivatives, specifically, second derivative.
- Second-order optimization methods are based on root finding
- Convex functions have global minima
- Quadratic programs are special kinds of least squares problems
- Optimization packages in Julia provide amazing tools for optimizing functions

12.1 Motivation and Basic Ideas




Figure 12.1: (a) A “simple” cost function with a global minimum at $x^* = 2$ and (b) a “less simple” cost function where there are two local minima, one at $x^* \approx 0.68$ and one at $x^* \approx 3.14$, as well as a local maximum at ≈ 2.1 .

Optimization is the process of finding one or more values $x \in \mathbb{R}^m$ that minimize a function $f : \mathbb{R}^m \rightarrow \mathbb{R}$, where the scalar-valued function f is called the **cost function**. The cost should be minimum at a point x^* that is of interest to you, it should be “small” for values of x that are “near” x^* and it should be “large” for values of x that are “far” from your preferred value, x^* . In Machine Learning, a cost function for minimization is also called a **regret function** in the sense that you regret being far from your preferred value, x^* , and your regret is minimum at x^* . We will abuse notation¹ and write

$$x^* = \arg \min_{x \in \mathbb{R}^m} f(x) \quad (12.1)$$

to denote the value of x achieving the minimum, even when there may be more than one such x .

One can also do the opposite of minimization, which is to **maximize a cost function**. In this case, the Machine Learning community calls the cost function a **reward**, because hey, who does not like to maximize reward and minimize regret! We will stick to minimization throughout the Chapter until the very end, when we’ll discuss briefly how maximization and minimization are very closely related.

When $m = 1$, that is, the cost function depends on a scalar variable, it is easy to graphically understand what minimization is all about. Figure 12.1a shows a “bowl-shaped” function that has a minimum at $x^* = 2$. Moreover, if the function outside of the interval $[-1, 5]$ continues growing as shown, then $x^* = 2$ is a **global minimum**, meaning that for all $x \neq x^*$,

$$f(x^*) < f(x).$$

Even more special, if you imagine setting a marble on the curve at any point other than $x^* = 2$, you can easily imagine the marble rolling and settling at the global minimum (assuming some friction).

On the other hand, Fig. 12.1b presents a more complicated situation, where if the function outside of the interval $[-1, 5]$ continues growing as shown, then there is a global minimum at $x_a^* \approx 0.68$, but if you set a marble near $x = 4$, you can imagine it getting stuck at the local bowl at $x_b^* \approx 3.14$, while if you start the marble at $x = 1.5$, you can imagine it settling in at the local bowl $x_a^* \approx 0.68$. The local bowls in the cost function are called **local minima** in that for all x sufficiently near one of the x_i^* ,

$$f(x_i^*) \leq f(x),$$

for $i \in \{a, b\}$.

¹Abusing notation means that one is being a bit sloppy with one’s use of symbols, which is what notation is! One often abuses notation when doing the right thing is a bit painful and would cause more of a distraction to explain the good notation than to caution about using poor notation!




Figure 12.2: The derivatives of the cost functions have been added at strategic points. (a) A “simple” cost function with a global minimum at $x^* = 2$ and (b) a “less simple” cost function where there are two local minima, one at $x^* \approx 0.68$ and one at $x^* \approx 3.14$, and a local maximum at 2.1.

Derivative of the Cost Tells You How to Move Toward a Local Minimum

In Fig. 12.2, we have plotted the derivatives of the cost functions at several points. Recall that the slope of the line segment at each point is equal to the derivative at that point. The plots tell us some really important things:

- the derivative of the cost is zero at each local minimum;
- to the left of a local minimum, the derivative is negative (slope angles downward), while to the right of a local minimum, the derivative is positive (slope angles upward);
- hence, if you are at a point x_k , defining the next value as

$$x_{k+1} = x_k - s \frac{df(x_k)}{dx}, \quad (12.2)$$

where $s > 0$ is called the **step size**, moves you in the direction of a local minimum (of course, if $s > 0$ is too large, you can overshoot the local minimum);

- because the derivative is zero at a local minimum, we can use the value of the derivative as a stopping criterion for the algorithm in (12.2); and
- the derivative also vanishes at local maxima, and hence if you start the algorithm in (12.2) at a local maxima, you will be stuck there!

Example 12.1 Implement (12.2) in Julia to find local minima of $f : \mathbb{R} \rightarrow \mathbb{R}$ by

$$f(x) = (x - 2)^2 + 5(\sin(x - 2))^2 + 0.03(x + 1)^3 + 4.$$

Solution: We wrote the Julia code given below and selected a set of initial conditions

$$x_0 \in \{ -1.00 \quad 0.00 \quad 0.68 \quad 2.10 \quad 2.15 \quad 3.14 \quad 4.00 \quad 5.00 \} \quad (12.3)$$

for the Algorithm (12.2) such that it would be started to the left and right of local minima and very close to the local maxima. The results are reported in (12.4) for a step size of $h = 0.1$ and in (12.5) for a step size of $h = 0.2$. The Algorithm converged quickly in all cases for a step size of $h = 0.1$ and mostly failed for a step size of $h = 0.2$ (we set a limit of 10^3 iterations before terminating).

You have to chose your step sizes carefully! for $h = 0.1$, when started just to the left of the local maximum, the algorithm converged to the local minimum at 0.68 and when started just to the right of the local maximum, it covered to the local minimum at 3.14.

```

1 #Optimization
2 cost2(x)=(x.-2).^2 .+ 1 .- 5*(sin.(x.-2)).^2 .+ 3 .+ 0.03*(x.+1).^3
3 titre="Less Simple Example"
4 p2=plot(cost2,xmin,xmax,legend=false, title=titre, linewidth=3, color=:red )
5 s=0.2
6 delta=0.01
7 x0=[-1;0;0.68;2.1;2.15;3.14;4;5]
8 y0=cost2(x0)
9 IntermediateValues=Array{Float64}(undef, 0, 6)
10 for k =1:length(x0) #try various initial values of x0
11     xk=x0[k]
12     dcostdxk =( cost2(xk+delta)-cost2(xk-delta) )/(2*delta)
13     fk=cost2(xk)
14     j=0
15     while (abs(dcostdxk)>1e-5) & (j < 1e3)
16         j=j+1
17         xk=xk-s*dcostdxk
18         dcostdxk =( cost2(xk+delta)-cost2(xk-delta) )/(2*delta)
19         fk=cost2(xk)
20     end
21     IntermediateValues=[IntermediateValues; j s x0[k] xk fk dcostdxk]
22 end
23 display(IntermediateValues)
24 # Show how to get latex code for printing matrices
25 using Latexify
26 set_default(fmt = "%.3f", convert_unicode = false)
27 latexify(IntermediateValues) |> print

```

N0. Iterations	s	x_0	x^*	$f(x^*)$	$\frac{df(x^*)}{dx}$
6	0.100	-1.000	0.678	1.193	0.000
6	0.100	0.000	0.678	1.193	0.000
4	0.100	0.680	0.678	1.193	0.000
15	0.100	2.100	0.678	1.193	0.000
12	0.100	2.150	3.137	3.300	0.000
4	0.100	3.140	3.137	3.300	0.000
7	0.100	4.000	3.137	3.300	0.000
8	0.100	5.000	3.137	3.300	0.000

(12.4)

N0. Iterations	s	x_0	x^*	$f(x^*)$	$\frac{df(x^*)}{dx}$
FAILED	0.200	-1.000	XX	XX	XX
FAILED	0.200	0.000	XX	XX	XX
FAILED	0.200	0.680	XX	XX	XX
FAILED	0.200	2.100	XX	XX	XX
65	0.200	2.150	3.137	3.300	0.000
47	0.200	3.140	3.137	3.300	-0.000
FAILED	0.200	4.000	XX	XX1	XX
69	0.200	5.000	3.137	3.300	-0.000

(12.5)

■

We give next a more analytical take on our update law in (12.2). Which derivation is better? The intuition from plots or an analysis via linear approximations? The answer depends on how your brain is wired. I would say the best derivation is the one that gives you

that “ah ha” moment!

Linear Approximation to the Rescue!

Let’s recall our linear approximation to a function $f : \mathbb{R} \rightarrow \mathbb{R}$ near a point x_k by

$$f(x) \approx f(x_k) + \frac{df(x_k)}{dx} (x - x_k). \quad (12.6)$$

We seek to define x_{k+1} so that $f(x_{k+1}) < f(x_k)$, that is, $f(x_{k+1}) - f(x_k) < 0$. We define $\Delta x_k := x_{k+1} - x_k$ and note that (12.6) gives

$$f(x_{k+1}) - f(x_k) \approx \frac{df(x_k)}{dx} \Delta x_k. \quad (12.7)$$

If we believe in the approximation (which is fine as long as x_{k+1} is “near” x_k), then we can replace the approximation sign with an equals sign so that

$$f(x_{k+1}) - f(x_k) < 0 \iff \frac{df(x_k)}{dx} \Delta x_k < 0. \quad (12.8)$$

We see that if $\frac{df(x_k)}{dx} = 0$, there is no choice of Δx_k that yields $\frac{df(x_k)}{dx} \Delta x_k < 0$, which is why the derivative vanishing is our stopping criterion for a local extremum. We assume therefore $\frac{df(x_k)}{dx} \neq 0$, in which case

$$\Delta x_k = -s \frac{df(x_k)}{dx} \implies \frac{df(x_k)}{dx} \Delta x_k = -s \left[\frac{df(x_k)}{dx} \right]^2 < 0 \text{ for all } s > 0,$$

and our update law becomes

$$x_{k+1} = x_k + \Delta x_k = x_k - s \frac{df(x_k)}{dx},$$

which agrees with (12.2).

12.2 Contour Plots and the Gradient of the Cost

We acquired some good intuition by looking at plots of cost functions depending on a scalar variable $x \in \mathbb{R}$. We’ll now augment our intuition by exploring the case $x \in \mathbb{R}^2$. We’ll look at **contour plots of cost functions**. We’ll also **encounter an important property of the gradient of a cost function that will generalize to arbitrary dimensions**.

We introduced the norm of a vector as a means to measure the length of vectors. We used the square of the norm in Chapter 8.2 to find a *least squared error solution* to a system of linear equations². In engineering, the square of the norm is a very common cost function. In Fig. 12.3a we plot

$$f(x) := \|x\|^2 = (x_1)^2 + (x_2)^2$$

on the z -axis versus x_1 and x_2 . It’s hard to imagine anything simpler to optimize! From the plot, we see a “generalized bowl shape”, though it’s hard to clearly identify where the minimum value occurs (yes, we know it occurs at $x = 0$). To the right of the plot, in Fig. 12.3b, we show a **contour plot** of $f(x_1, x_2)$. The circles about the origin show values of $x = (x_1, x_2)$ where the cost function is constant; in the inner circle, $f(x) = \|x\|^2 = 1$, then $f(x) = \|x\|^2 = 4, \dots$, until $f(x) = \|x\|^2 = 25$. In the contour plot, it is easy to see where the minimum value occurs.

Similarly, in Fig. 12.3c, we show a plot of

$$f(x) = \|x - x_0\|^2 = (x_1 - x_{01})^2 + (x_2 - x_{02})^2 \quad (12.9)$$

for $x_0 = [1; 2]$. In the corresponding contour plot, Fig. 12.3d, it is equally easy to see where the minimum occurs. We will use the contour plots to show the correct direction to follow from a given point x_k so that one moves efficiently toward a value of x^* that is a local minimum.

²Recall that used the square of the norm when assessing the error $e := Ax - b$ in equations that do not have exact solutions.

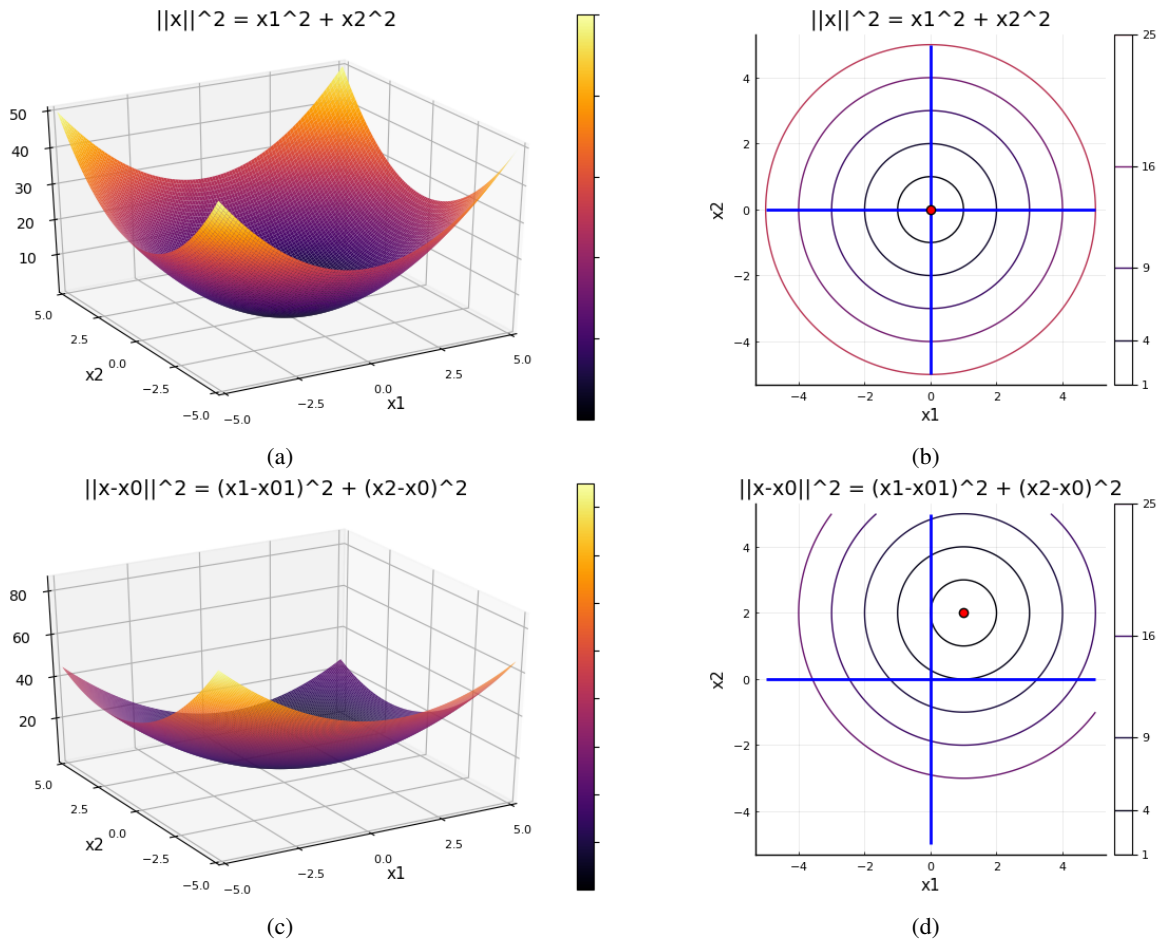


Figure 12.3: The graphs of two functions are shown in (a) and (c), with their corresponding contour plots shown in (b) and (d). Contour plots are lines where a function is constant. If you have ever used a topographical map when hiking, the lines on those maps indicate constant terrain height. In our case, the lines are constant $z = f(x_1, x_2)$ values as (x_1, x_2) vary. Because the cost function we use here is very simple, its lines of constant contour are circles. More “interesting” examples will follow.

Recall that the gradient of (12.9) is a row³ vector of the form $\nabla f(x) = \left[\frac{\partial f(x)}{\partial x_1} \quad \frac{\partial f(x)}{\partial x_2} \right]$. Hence, the transpose of the gradient is a column vector in \mathbb{R}^2 , namely

$$[\nabla f(x)]^\top = \begin{bmatrix} \frac{\partial f(x)}{\partial x_1} \\ \frac{\partial f(x)}{\partial x_2} \end{bmatrix}. \quad (12.10)$$

Figure 12.4a shows two contour plots. In addition, we have plotted in green arrows the gradients of the cost functions at several points and overlaid them on the contour plots. In both cases, the gradients are pointing in the direction of maximum increase of the function. It follows that the negative of the gradient is the direction of maximum decrease.

In Fig. 12.4, we could not plot the gradients at the minima because the length of the gradient vector is zero at a (local) minimum! In case you want to see why this is true, from Calculus, one obtains that

$$\nabla ((x_1 - x_{01})^2 + (x_2 - x_{02})^2) = 2(x_1 - x_{01}) + 2(x_2 - x_{02}) \quad (12.11)$$

$$\nabla ((Ax - b)^\top \cdot (Ax - b)) = 2(Ax - b)^\top \cdot A. \quad (12.12)$$

The minimum of $(x_1 - x_{01})^2 + (x_2 - x_{02})^2$ occurs at $x_1^* = x_{01}$ and $x_2^* = x_{02}$, and indeed, the gradient vanishes at x^* . The minimum of $\|Ax - b\|^2 = (Ax - b)^\top \cdot (Ax - b)$ satisfies $A^\top \cdot Ax^* = A^\top b$, and direct substitution into the corresponding gradient shows that it vanishes as well. While this is not a proof, it gives you a hint that the gradient vanishing at a local minimum may be true.

³From Calculus, in this case, the exact formula is $\nabla f(x) = [2(x_1 - x_{01}) \quad 2(x_2 - x_{02})]$




Figure 12.4: Contour Plots plus Gradients. (a) shows a very simple cost function, $\|x - [1; 2]\|^2$, where the contours of constant cost are circles. (b) shows a more typical situation, where the contours are squeezed in one direction. In both cases, the gradients, plotted in green arrows, are pointing in the direction of maximum increase of the function. Hence, the negative of the gradient is the direction of maximum decrease.

12.3 Gradient Descent

Gradient Descent

Equations (12.11) and (12.12) along with Fig. 12.4 tell us some really important things:

- the gradient vanishes at local minima;
- the gradient points in the direction of fastest growth of the cost function, and the negative of the gradient points in the direction of fastest decrease;
- hence, if you are at a point x_k , defining the next value as

$$x_{k+1} = x_k - s [\nabla f(x_k)]^\top, \quad (12.13)$$

where $s > 0$ is called the **step size**, moves you in the direction of a local minimum (of course, if $s > 0$ is too large, you can overshoot the local minimum);

- because the gradient is zero at a local minimum, we can use the value of the norm of the gradient as a stopping criterion for the algorithm in (12.13); and
- the gradient also vanishes at local maxima, and hence if you start the algorithm in (12.13) at a local maxima, you will be stuck there.

Before doing examples, we'll also show we can derive (12.13) from the linear approximation of $f: \mathbb{R}^m \rightarrow \mathbb{R}$ near a point x_k , namely,

$$f(x) \approx f(x_k) + \nabla f(x_k) (x - x_k), \quad (12.14)$$

where $\nabla f(x_k)$ is the gradient of f at x_k and is a $1 \times m$ row vector. We seek to define x_{k+1} so that $f(x_{k+1}) < f(x_k)$, that is, $f(x_{k+1}) - f(x_k) < 0$. We define $\Delta x_k := x_{k+1} - x_k$ and note that (12.14) gives

$$f(x_{k+1}) - f(x_k) \approx \nabla f(x_k) \Delta x_k. \quad (12.15)$$

If we believe in the approximation (which is fine as long as x_{k+1} is “near” x_k), then we can replace the approximation sign with an equals sign so that

$$f(x_{k+1}) - f(x_k) < 0 \iff \nabla f(x_k) \Delta x_k < 0. \quad (12.16)$$

We see that if $\nabla f(x_k) = 0$, there is no choice of Δx_k that yields $\nabla f(x_k) \Delta x_k < 0$, which is why the gradient vanishing is our stopping criterion for a local extremum. We assume therefore $\nabla f(x_k) \neq 0$, in which case selecting

$$\Delta x_k = -s [\nabla f(x_k)]^\top \implies \nabla f(x_k) \Delta x_k = -s \|\nabla f(x_k)\|^2 < 0 \text{ for all } s > 0.$$

Our update law is then

$$x_{k+1} = x_k + \Delta x_k = x_k - s [\nabla f(x_k)]^\top,$$

which agrees with (12.13).

Remark: There are easy variations of the Gradient Descent Algorithm, [Plestan-2020](#). Let's note that our key condition is

$$\nabla f(x_k) \Delta x_k < 0.$$

If we define the i -th component of Δx_k by

$$(\Delta x_k)_i := \begin{cases} 0 & \frac{\partial f(x_k)}{\partial x_i} = 0 \\ -s \operatorname{sign}\left(\frac{\partial f(x_k)}{\partial x_i}\right) & \text{otherwise} \end{cases},$$

then it is still true that $\nabla f(x_k) \Delta x_k < 0 \iff \nabla f(x_k) \neq 0_{1 \times m}$, and thus we are moving in a descent direction.

Example 12.2 We'll re-visit the least squares problem from Chapter 8.3, and use gradient descent as given in (12.13) to solve

$$x^* = \arg \min_{x \in \mathbb{R}^m} \|Ax - b\|^2$$

and compare it to the closed-form solution $(A^\top A) x^* = A^\top b$, for

$$A = \begin{bmatrix} 1.0000 & 0.0000 & 0.0000 \\ 1.0000 & 0.2500 & 0.0625 \\ 1.0000 & 0.5000 & 0.2500 \\ 1.0000 & 0.7500 & 0.5625 \\ 1.0000 & 1.0000 & 1.0000 \\ 1.0000 & 1.2500 & 1.5625 \\ 1.0000 & 1.5000 & 2.2500 \\ 1.0000 & 1.7500 & 3.0625 \\ 1.0000 & 2.0000 & 4.0000 \end{bmatrix} \quad \text{and} \quad b = \begin{bmatrix} 1.0000 \\ 1.0000 \\ 1.5000 \\ 2.0000 \\ 3.0000 \\ 4.2500 \\ 5.5000 \\ 7.0000 \\ 10.0000 \end{bmatrix}.$$

Solution: We apply gradient descent as given in (12.13) to the optimization problem,

$$x^* = \arg \min_{x \in \mathbb{R}^3} (Ax - b)^\top (Ax - b).$$

We use symmetric differences to compute the gradient of

$$f(x) := (Ax - b)^\top (Ax - b),$$

set the step size to $s = 0.01$, and the tolerance to $\|\nabla f(x_k)\| < 10^{-5}$. After 2,673 iterations, we obtain

$$x^{\text{approx}} = \begin{bmatrix} 1.065144e + 00 \\ -6.257368e - 01 \\ 2.454536e + 00 \end{bmatrix}. \tag{12.17}$$

For comparison purposes, we recall that the true answer is

$$x^* = \begin{bmatrix} 1.065152e + 00 \\ -6.257576e - 01 \\ 2.454545e + 00 \end{bmatrix}. \tag{12.18}$$

The true power of optimization becomes clear when we use it to solve problems that do NOT have closed-form solutions. Check out the next section! ■

12.4 Extrinsic Calibration Using Gradient Descent

Extrinsic calibration is the problem of finding a rotation matrix and a translation vector that allows one to transform the 3D- (x, y, z) points measured by a LiDAR so as to align them with the data captured by a camera. The problem is interesting because it is easy to show visually what one wants to accomplish, but yet, how to formulate the objective mathematically is not clear at first glance, and for sure, there is no obvious “equation” to solve in order to compute a solution.

Here, we’ll go into enough detail that we can formulate an optimization problem for “aligning” the scenes observed by the LiDAR and the camera. We’ll first address the optimization problem using gradient descent and later, we’ll use a more efficient technique that uses “higher-order” information about the cost function. The difference in the rates of convergence is mind blowing.

12.4.1 Introduction

Currently, basic cameras provide many more data points (pixels) for a given surface size at a given distance than even high-end LiDARs. However, cameras rely on the consistency of the ambient lighting and provide less accurate depth estimation. On the other hand, LiDARs give accurate distance measurement, and rapidly changing ambient lighting will not affect measurements of a LiDAR. Due to these different inherent properties of LiDARs and cameras, it is necessary to “fuse” the LiDAR and the camera on Cassie’s torso, as shown in Fig. 12.7. By fusing, we mean overlaying a point cloud from a LiDAR to an image from a camera. The process of fusing two or more different types of data from different sensors is called an “extrinsic calibration”. The problem requires high accuracy and precision; a little error in the process leads to an unusable result, as shown in Fig. 12.5b.




Figure 12.5: a shows good alignment of a LiDAR point cloud projected onto a camera image. b shows that a calibration result is not usable if it has a few degrees of rotation error and a few percent of translation error.




Figure 12.6: a shows a normal image and b illustrates the result of edge detection applied on a.

When we calibrate sensors, we need to find corresponding “features” captured from different sensors. Features are some specific structures in an environment, such as points, edges or corners, etc. Figure 12.6 considers edges as features and extracts them out from an image. Once we find “enough” corresponding features, we want to estimate a “rigid-body transformation” between each sensor that we want to calibrate. A rigid-body transformation H consists of a rotation matrix R and a translation vector t and is defined as

$$H := \begin{bmatrix} R & t \\ 0 & 1 \end{bmatrix}. \quad (12.19)$$

When calibrating two sensors, it is good practice to indicate how the transformation is defined, by specifying that is a transformation from which sensor to which other sensor! We will use the notation $R_{\text{from sensor1}}^{\text{to sensor2}}$ and $t_{\text{from sensor1}}^{\text{to sensor2}}$ to represent the rotation and translation from the sensor 1 to sensor 2.

In the following, we will illustrate how to perform an extrinsic calibration between a LiDAR and a camera. We will walk you through the calibration process and provide some insight on how we try to take advantage of their relative strengths and avoid their weaknesses. All the images and data are collected from the torso of Cassie Blue, as shown in Fig. 12.5a.



Figure 12.7: This figure shows the torso of Cassie Blue. We use it in practice to do our [autonomy adventures!](#)

12.4.2 Problem Setup and Initial Solution

In this problem, we take corners as our features, as shown in Fig. 12.8. In particular, we want to align the LiDAR vertices (green dots) with the camera corners (red dots) in Fig. 12.8b. Additionally, we assume all corresponding corners from the two sensors are already given⁴, and the process of feature extraction is uploaded to our [YouTube Channel!](#) Let X and Y be the features from a LiDAR and a camera, respectively. When overlaying the two sets of features, we want their respective coordinates in the camera frame to be as close as possible. Therefore, the problem can be formulated as

$$\begin{aligned} (R_L^{C*}, t_L^{C*}) &:= \arg \min_{R,t} f(R, t, X, Y) \\ &:= \arg \min_{R,t} \sum_{i=1}^{4n} \|\Pi(X_i; R, t) - {}_C Y_i\|_2^2, \end{aligned} \quad (12.20)$$

where R_L^C and t_L^C are the rotation and translation from the LiDAR frame to the camera frame, f is the cost function, and Π is a projection map, which we will not dive into in here, but is provided in Appendix C.8. For now, you can consider it as a black box that takes points in 3D space and maps them to their corresponding points on the 2D image-plane of the camera. We apply the gradient descent algorithm to solve (12.20) and the update function is:

$$H_{k+1} = H_k - s[\nabla f(H_k, X, Y)]^\top. \quad (12.21)$$

⁴To understand how to obtain corresponding corners, we encourage the readers to read “[Improvements to Target-Based 3D LiDAR to Camera Calibration](#),” by Jiunn-Kai Huang and Jessy W. Grizzle.

After 5745 iterations with a step size of $1e-8$, the cost drops from 49750 to 12.12, and the result is shown in Fig. 12.9. Later, we will introduce a second-order method, the ‘‘Hessian,’’ to solve this problem and you will be amazed when the Hessian-based algorithm converges in 14 iterations! Knowledge is power!


Remark: If you look into the implementation, you will find out that we do not represent the rotation matrix with nine elements even though it is a 3×3 matrix,

$$R = \begin{bmatrix} r_{11} & r_{12} & r_{13} \\ r_{21} & r_{22} & r_{23} \\ r_{31} & r_{32} & r_{33} \end{bmatrix}. \quad (12.22)$$


Instead, we use a fact you would learn in an upper level math course or a mechanics course, which says that any rotation matrix can be represented as the ‘‘matrix exponential’’ of a 3×3 skew symmetric matrix of the form

$$\Omega := \begin{bmatrix} 0 & -\omega_3 & \omega_2 \\ \omega_3 & 0 & -\omega_1 \\ -\omega_2 & \omega_1 & 0 \end{bmatrix},$$

Hence, a rotation matrix depends on only three parameters, $(\omega_1, \omega_2, \omega_3)$. (It’s hard to learn too much Linear Algebra; there is always one more useful fact!). Knowledge is power. Keep digging, and you will keep having fun.



(a)



(b)



(c)

Figure 12.8: a shows the LiDAR returns on the targets in black and estimated vertices in red. b illustrates the initial state of the optimization in (12.20). Our goal is to move the green dots onto the red dots, or you can also imagine moving the yellow box to the magenta box. c demonstrates the result of calibration: the LiDAR vertices (green dots) are successfully projected onto a camera corners (red dots).




Figure 12.9: This figure demonstrates the calibration result of the gradient descent algorithm. It looks pretty good; there is still some misalignment. The main drawback of the algorithm would that a very small step size is often required to make it work.

12.5 Optimization as a Root Finding Problem: the Hessian

Gradient descent is a powerful method to solve optimization problems and we've seen it in two contexts so far: finding a (local) minimum of a cost function that depends on a scalar variable x in (12.2), and finding a (local) minimum of a cost function that depends on a vector variable $x = (x_1, x_2, \dots, x_m)$ in (12.13). Our stopping criterion was $|\frac{df(x_k)}{dx}|$ "sufficiently small" for scalar problems and $\|\nabla f(x_k)\|$ "sufficiently small" for vector problems. In other words, a (locally) optimal solution x^* satisfies $df(x^*)/dx = 0$ for scalar problems and $\nabla f(x^*) = 0$ for vector problems. Said yet another way, our locally minimizing solutions are **roots of the derivative of the cost function**.

Relation to Root Finding

Suppose we seek a point $x^* \in \mathbb{R}$ achieving a local minimum (or maximum) of a function $f : \mathbb{R} \rightarrow \mathbb{R}$. We know that a necessary condition is the derivative of f vanishes at x^* , that is,

$$\left. \frac{df(x)}{dx} \right|_{x=x^*} = 0.$$

We note that the first derivative is just another real-valued function, namely,

$$\frac{df(x)}{dx} : \mathbb{R} \rightarrow \mathbb{R}.$$

Hence, we can say that for x^* to achieve a minimum (or maximum) of the function $f : \mathbb{R} \rightarrow \mathbb{R}$, it must be a **root of** $\frac{df(x)}{dx} : \mathbb{R} \rightarrow \mathbb{R}$. If we apply Newton's Method to the function $g(x) := \frac{df(x)}{dx}$ to find its roots, we'll need to find the first derivative of $g(x)$, which will lead us to the second derivative of the original function $f(x)$. This may sound overwhelming or intimidating, but shortly we will disabuse you of such concerns!

Next, suppose we seek a vector $x^* \in \mathbb{R}^m$ that achieves a local minimum (or maximum) of a function $f : \mathbb{R}^m \rightarrow \mathbb{R}$. We know that a necessary condition is the **gradient** of f vanishes at x^* , that is,

$$\nabla f(x^*) = 0_{1 \times m}.$$

We note that the **transpose of the gradient** is just another vector-valued function, namely,

$$\nabla f(x)^\top : \mathbb{R}^m \rightarrow \mathbb{R}^m;$$

we use the transpose to turn a row vector into a column vector. Because the gradient must vanish at a local min (or max), we can say that for x^* to be a local minimum (or maximum) of the function $f : \mathbb{R}^m \rightarrow \mathbb{R}$, it must be a **root of** $\nabla f(x)^\top : \mathbb{R}^m \rightarrow \mathbb{R}^m$. If we apply the Newton-Raphson Method to the function $g(x) := \nabla f(x)^\top$ to find its roots, we'll need to find the Jacobian of $g(x)$, which will lead us to a vector-version of the second derivative of the original function $f(x)$. This may sound overwhelming or intimidating, but shortly, we will once again disabuse you of such concerns! The Jacobian of the gradient has a cool name, **the Hessian**.

In Chapter 11, we learned a lot about finding roots of nonlinear equations. We will now apply that knowledge to optimization and thereby arrive at a more powerful optimization algorithm that uses information about the second derivative of the cost function! That sounds pretty wild and scary, but you'll see that it's very do-able.

12.5.1 Second Derivatives

Calculus has developed some good notation over the past 300 plus years and we will use it. The second derivative of a function $f : \mathbb{R} \rightarrow \mathbb{R}$ is the first derivative of the function's first derivative, and is written like this,

$$\frac{d^2 f(x)}{dx^2} := \frac{d}{dx} \frac{df(x)}{dx}. \quad (12.23)$$

Let's write down the derivative of the derivative using the symmetric difference approximation to the derivative. We first recall that

$$\frac{df(x)}{dx} \approx \frac{f(x+h) - f(x-h)}{2h}, \quad (12.24)$$

where $h > 0$ is a small number. Let's use $\delta > 0$ instead of h for a small positive change in x when writing down the second derivative as a symmetric difference

$$\frac{d^2 f(x)}{dx^2} \approx \frac{\frac{df(x+\delta)}{dx} - \frac{df(x-\delta)}{dx}}{2\delta}. \quad (12.25)$$

We now substitute (12.24) into (12.25) to obtain

$$\begin{aligned} \frac{d^2 f(x)}{dx^2} &\approx \frac{\left[\frac{f(x+\delta+h) - f(x+\delta-h)}{2h} \right] - \left[\frac{f(x-\delta+h) - f(x-\delta-h)}{2h} \right]}{2\delta} \\ &= \frac{[f(x+\delta+h) - f(x+\delta-h)] - [f(x-\delta+h) - f(x-\delta-h)]}{4h\delta} \\ &= \frac{f(x+\delta+h) - f(x+\delta-h) - f(x-\delta+h) + f(x-\delta-h)}{4h\delta}. \end{aligned} \quad (12.26)$$

Equation (12.26) is a perfectly fine expression for the second derivative. At least for your author, keeping δ and h separate made it easier to evaluate the individual terms and see how they appeared in the equation. Your experience may vary! It's more traditional to take $\delta = h$, which we now do so as to arrive at our final expression,

$$\boxed{\frac{d^2 f(x)}{dx^2} \approx \frac{f(x+2h) - 2f(x) + f(x-2h)}{4h^2}}. \quad (12.27)$$

In terms of coding, computing a numerical approximation to the second derivative is not much different than approximating the first derivative. Comparing (12.27) to (11.7), we see there is one more term to compute and instead of dividing by $2h$, we divide by $4h^2$.

Note that if you took $h > 0$ to be something relatively "small", such as 10^{-3} , then h^2 is now really small, such as 10^{-6} . Be careful that you do not approach "machine epsilon" when doing your approximate derivatives!

12.5.2 The Hessian is the Jacobian of the Transpose of the Gradient

As the title of this section suggests, we define the **Hessian** of a function $f : \mathbb{R}^m \rightarrow \mathbb{R}$ as the Jacobian of the transpose of the gradient of the function. Using notation that comes to us from Calculus, we have that

$$\nabla^2 f(x) := \frac{\partial}{\partial x} [\nabla f(x)]^\top, \quad (12.28)$$

where $\nabla^2 f(x)$ is the notation for the Hessian of f at the point x . We underline that here, the function f depends on a vector $x \in \mathbb{R}^m$ and that $f(x) \in \mathbb{R}$ is a scalar. If $f(x) \in \mathbb{R}^n$, for $n > 1$, then the above formula makes no sense...it is just a bunch of meaningless symbols.

Equation (12.28) is a lot of notation packed into one tiny formula! Let's unpack it so as to understand it bit by bit. For a function $f : \mathbb{R}^m \rightarrow \mathbb{R}$, the transpose of its gradient is

$$[\nabla f(x)]^\top := \begin{bmatrix} \frac{\partial f(x)}{\partial x_1} \\ \vdots \\ \frac{\partial f(x)}{\partial x_k} \\ \vdots \\ \frac{\partial f(x)}{\partial x_m} \end{bmatrix}. \quad (12.29)$$

Moreover, we can approximate the indicated partial derivatives using symmetric differences,

$$\frac{\partial f(x)}{\partial x_k} = \frac{f(x + he_k) - f(x - he_k)}{2h}, \quad (12.30)$$

where $\{e_1, \dots, e_k, \dots, e_m\}$ is the natural basis vectors for \mathbb{R}^m . For a function $g : \mathbb{R}^m \rightarrow \mathbb{R}^m$, we recall that its Jacobian is

$$\frac{\partial g(x)}{\partial x} := \begin{bmatrix} \frac{\partial g(x)}{\partial x_1} & \dots & \frac{\partial g(x)}{\partial x_j} & \dots & \frac{\partial g(x)}{\partial x_m} \end{bmatrix}. \quad (12.31)$$

Moreover, we know how to numerically approximate the partial derivatives in (12.31) using symmetric differences and selecting $\delta > 0$ by

$$\frac{\partial g(x)}{\partial x_j} = \frac{g(x + \delta e_j) - g(x - \delta e_j)}{2\delta}. \quad (12.32)$$

To put all of this together, we take

$$g(x) := [\nabla f(x)]^\top$$

and apply (12.32) to obtain a first numerical approximation for the Hessian, namely

$$\nabla^2 f(x) \approx \frac{1}{2\delta} \left[(\nabla f(x + \delta e_1) - \nabla f(x - \delta e_1))^\top \quad \cdots \quad (\nabla f(x + \delta e_k) - \nabla f(x - \delta e_k))^\top \quad \cdots \quad (\nabla f(x + \delta e_m) - \nabla f(x - \delta e_m))^\top \right]. \quad (12.33)$$

Going one step further, if we let $[\nabla^2 f(x)]_{ij}$ be the ij -entry of the Hessian matrix (that is, the entry for its i -th row and j -th column), using (12.30), we have that

$$[\nabla^2 f(x)]_{ij} \approx \frac{1}{4h\delta} (f(x + he_i + \delta e_j) - f(x + he_i - \delta e_j) - f(x - he_i + \delta e_j) + f(x - he_i - \delta e_j)). \quad (12.34)$$

In this case, setting $\delta = h$ does not really simplify the expression. In Julia, we suspect that you will find (12.33) the easiest to implement.

Remark: In Calculus, we also use the notation

$$\frac{\partial^2 f(x)}{\partial x_i \partial x_j} := [\nabla^2 f(x)]_{ij}.$$

In (12.34), when you take $h = \delta$, you can see that

$$[\nabla^2 f(x)]_{ji} = [\nabla^2 f(x)]_{ij}$$

and therefore the Hessian is a symmetric matrix.

12.5.3 Use of the Second Derivative and Hessian in Optimization

The most accessible forms of second-order optimization methods are based on applying Newton's method to the first derivative of a cost function that depends on a scalar variable $x \in \mathbb{R}$ or the Newton-Raphson Algorithm to the (transpose of the) gradient of a cost function that depends on a vector variable $x \in \mathbb{R}^m$.

Optimization as a Form of Root Finding: Scalar Decision Variable

For $f : \mathbb{R} \rightarrow \mathbb{R}$, the iterative process

$$x_{k+1} = x_k - \left(\frac{d^2 f(x_k)}{dx^2} \right)^{-1} \frac{df(x_k)}{dx} \quad (12.35)$$

will converge to a local extremum of f if the initial value x_0 is "well chosen". Because the algorithm is looking for **roots** of the first derivative, it is *indifferent* to whether the root is a local minimum or a local maximum. **In Calculus, you will learn that if the second derivative is positive at a point where the first derivative vanishes, then you have found a local minimum.** Similarly, a negative second derivative implies a local maximum.

The damped version of the algorithm

$$x_{k+1} = x_k - s \left(\frac{d^2 f(x_k)}{dx^2} \right)^{-1} \frac{df(x_k)}{dx}, \quad (12.36)$$

where $0 < s < 1$, typically performs better in practice. We note that the second derivative, $\frac{d^2 f(x)}{dx^2}$, can be approximated as in (12.27).

All of the remarks made in Chapter 11 about the validity of Newton's Algorithm apply here as well.

Optimization as a Form of Root Finding: Vector Decision Variables

For $f : \mathbb{R}^m \rightarrow \mathbb{R}$, the iterative process

$$\nabla^2 f(x_k) \Delta x_k = - [\nabla f(x_k)]^\top \quad (\text{solve for } \Delta x_k) \quad (12.37)$$

$$x_{k+1} = x_k + \Delta x_k \quad (\text{use } \Delta x_k \text{ to update our estimate of the optimal value}) \quad (12.38)$$

will converge to a local extremum of f if the initial value x_0 is “well chosen”. Because the algorithm is looking for **roots** of the gradient, it is *indifferent* to whether the root is a local minimum, a local maximum, or what is called a “saddle point”. **In Calculus, you will learn that if the Hessian is positive definite at a point where the gradient vanishes, then you have found a local minimum.** Similarly, a negative definite Hessian implies a local maximum. A discussion of these nice properties is beyond the scope of ROB 101.

While for toy problems, we can use the matrix inverse to solve (12.37) for Δx_k , for larger problems, we recommend using an LU Factorization or a QR Factorization. Once (12.37) has been solved, x_{k+1} is updated in (12.38) and the process repeats. In practice, the **damped** version of the algorithm often works better, where one replaces (12.38) with

$$x_{k+1} = x_k + s\Delta x_k, \quad (12.39)$$

for some $0 < s < 1$. We note that the Hessian, $\nabla^2 f(x)$, can be approximated with either (12.33) or (12.34).

All of the remarks made in Chapter 11 about the validity of the Newton-Raphson Algorithm apply here as well.

Don't be Intimidated by the Notation!

While the equation

$$\nabla^2 f(x_k) \Delta x_k = - [\nabla f(x_k)]^\top$$

may look intimidating, it is just another linear equation $Ax = b$ where $A \leftrightarrow \nabla^2 f(x_k)$, a square matrix, $x \leftrightarrow \Delta x_k$, a column vector of unknowns, and $b \leftrightarrow - [\nabla f(x_k)]^\top$ is a column vector on the right-hand side of the equation. Hence, computing Δx_k is cake for you by now.

We re-do several of the previous examples.

Example 12.3 Based on Example 12.1, but this time, we implement (12.35) and (12.36) in Julia to find local extrema of $f : \mathbb{R} \rightarrow \mathbb{R}$,

$$f(x) = (x - 2)^2 + 5(\sin(x - 2))^2 + 0.03(x + 1)^3 + 4.$$

Solution: We run the algorithm from the same initial conditions as in Example 12.1 and for two values of the step size, $s \in \{0.25, 1.0\}$. The reader should note that we got somewhat lucky, and each time the algorithm converged to a root of the first derivative! We've highlighted in blue where the algorithm actually converged to a local maximum instead of a local minimum. How could we tell? The second derivative being positive implies a local minimum while it being negative implies a local maximum.

We note that the rate of convergence is much faster than for gradient descent (here, faster means fewer iterations of the algorithm). Finally, we note that the point to which the algorithm converges does depend on the initial condition. In fact, for $s = 0.25$ and $s = 1.0$, the algorithm sometimes converges to different roots of the first derivative, even when initialized at the same point.

If you can design a cost function so that it has a single extrema and it is your desired minimum, then you can avoid many of these problems. We talk a little about this in the section on “Local vs Global”.

N0. Iterations	s	x_0	x^*	$f(x^*)$	$\frac{df(x^*)}{dx}$	$\frac{d^2f(x^*)}{dx^2}$
$2.2000e + 01$	$1.0000e + 00$	$-1.0000e + 00$	$3.1369e + 00$	$3.3002e + 00$	$-4.5648e - 09$	$9.2092e + 00$
$4.0000e + 00$	$1.0000e + 00$	$0.0000e + 00$	$6.7838e - 01$	$1.1926e + 00$	$-5.9785e - 07$	$1.1085e + 01$
$2.0000e + 00$	$1.0000e + 00$	$6.8000e - 01$	$6.7838e - 01$	$1.1926e + 00$	$6.9886e - 10$	$1.1085e + 01$
$2.0000e + 00$	$1.0000e + 00$	$2.1000e + 00$	$2.1099e + 00$	$4.8543e + 00$	$-1.6541e - 08$	$-7.1982e + 00$
$2.0000e + 00$	$1.0000e + 00$	$2.1500e + 00$	$2.1099e + 00$	$4.8543e + 00$	$5.2104e - 07$	$-7.1982e + 00$
$2.0000e + 00$	$1.0000e + 00$	$3.1400e + 00$	$3.1369e + 00$	$3.3002e + 00$	$-2.8692e - 09$	$9.2092e + 00$
$5.0000e + 00$	$1.0000e + 00$	$4.0000e + 00$	$3.1369e + 00$	$3.3002e + 00$	$-2.6010e - 09$	$9.2092e + 00$
$2.3000e + 01$	$1.0000e + 00$	$5.0000e + 00$	$-2.4271e + 01$	$3.1198e + 02$	$1.2562e - 09$	$4.3032e + 00$

(12.40)

N0. Iterations	s	x_0	x^*	$f(x^*)$	$\frac{df(x^*)}{dx}$	$\frac{d^2f(x^*)}{dx^2}$
$1.2500e + 02$	$2.5000e - 01$	$-1.0000e + 00$	$-2.4271e + 01$	$3.1198e + 02$	$-9.2124e - 06$	$4.3033e + 00$
$4.7000e + 01$	$2.5000e - 01$	$0.0000e + 00$	$6.7838e - 01$	$1.1926e + 00$	$-9.7541e - 06$	$1.1085e + 01$
$2.7000e + 01$	$2.5000e - 01$	$6.8000e - 01$	$6.7838e - 01$	$1.1926e + 00$	$7.6022e - 06$	$1.1085e + 01$
$3.1000e + 01$	$2.5000e - 01$	$2.1000e + 00$	$2.1099e + 00$	$4.8543e + 00$	$9.5934e - 06$	$-7.1982e + 00$
$3.6000e + 01$	$2.5000e - 01$	$2.1500e + 00$	$2.1099e + 00$	$4.8543e + 00$	$-8.9863e - 06$	$-7.1982e + 00$
$2.8000e + 01$	$2.5000e - 01$	$3.1400e + 00$	$3.1369e + 00$	$3.3002e + 00$	$9.0254e - 06$	$9.2093e + 00$
$4.8000e + 01$	$2.5000e - 01$	$4.0000e + 00$	$3.1369e + 00$	$3.3002e + 00$	$9.9328e - 06$	$9.2093e + 00$
$6.8000e + 01$	$2.5000e - 01$	$5.0000e + 00$	$2.1099e + 00$	$4.8543e + 00$	$9.0945e - 06$	$-7.1982e + 00$

(12.41)




Figure 12.10: A cost function along with its first and second derivatives. We note that the extrema (both minima and maxima) in (a) correspond to the zero crossings (aka roots) of the first derivative in (b). Moreover, the sign of the second derivative at roots of the first derivative (aka, extrema) provides information on whether we have a local min, max, or neither.

An instantiation in Julia is given below.

```

1 #Data for Optimization, with the second derivative
2 #
3 # cost function
4 cost(x) = (x.-2).^2 .+ 1 .- 5*(sin.(x.-2)).^2 .+ 3 .+ 0.03*(x.+1).^3
5 yzero(x) = 0.0*cost(x)
6 # x-perturbation for derivatives
7 delta = 0.01
8 xmin = -1.0; xmax = 5.0
9 # first derivative of cost
10 dcost(x) = (cost(x+delta) - cost(x-delta)) / (2*delta)
11 # second derivative of cost
12 ddcost(x) = (dcost(x+delta) - dcost(x-delta)) / (2*delta)

```

```

13
14 # Plotting Commands for Figures
15 titre="First Derivative Less Simple Example"
16 p2=plot (dcost, xmin, xmax, legend=false, title=titre, linewidth=3, color=:red )
17 plot! (yzero, xmin, xmax, linewidth=2, color=:blue)
18 xlabel! ("x")
19 ylabel! ("dcost (x) /dx")
20 titre="Second Derivative Less Simple Example"
21 p3=plot (ddcost, xmin, xmax, legend=false, title=titre, linewidth=3, color=:red )
22 plot! (yzero, xmin, xmax, linewidth=2, color=:blue)
23 xlabel! ("x")
24 ylabel! ("d^2cost (x) /dx^2")
25
26 display (p2)
27 display (p3)
28
29 png (p2, "DerivativeLessSimpleCost")
30 png (p3, "SecondDerivativeLessSimpleCost")

```

```

1 # Second Order Optimization
2
3 s=0.25 # step size
4 x0=[-1;0;0.68;2.1;2.15;3.14;4;5] # Vector of initial conditions to show that different
5                                     # initial values lead to different local extrema
6 y0=cost (x0)
7 IntermediateValues=Array{Float64} (undef, 0, 7)
8 for j =1 :length (x0)
9     xk=x0 [j]
10    dcostdk = dcost (xk)
11    ddcostdk = ddcost (xk)
12    fk=cost (xk)
13    k=0
14    while (abs (dcostdk) >1e-5) & (k < 1e3)
15        k=k+1
16        xk=xk-s*dcostdk/ddcostdk
17        fk=cost (xk)
18        dcostdk = dcost (xk)
19        ddcostdk = ddcost (xk)
20    end
21    IntermediateValues=[IntermediateValues; k s x0 [j] xk fk dcostdk ddcostdk]
22 end
23 display (IntermediateValues)

```

Example 12.4 We re-visit the least squares problem from Example 12.2, which, in turn, came from Chapter 8.3. This time we use the Hessian and second order methods given in (12.37) through (12.39) to solve

$$x^* = \arg \min_{x \in \mathbb{R}^m} \|Ax - b\|^2,$$

where

$$A = \begin{bmatrix} 1.0000 & 0.0000 & 0.0000 \\ 1.0000 & 0.2500 & 0.0625 \\ 1.0000 & 0.5000 & 0.2500 \\ 1.0000 & 0.7500 & 0.5625 \\ 1.0000 & 1.0000 & 1.0000 \\ 1.0000 & 1.2500 & 1.5625 \\ 1.0000 & 1.5000 & 2.2500 \\ 1.0000 & 1.7500 & 3.0625 \\ 1.0000 & 2.0000 & 4.0000 \end{bmatrix} \quad \text{and} \quad b = \begin{bmatrix} 1.0000 \\ 1.0000 \\ 1.5000 \\ 2.0000 \\ 3.0000 \\ 4.2500 \\ 5.5000 \\ 7.0000 \\ 10.0000 \end{bmatrix}.$$

Solution: We apply Newton-Raphson to the gradient of the cost function

$$f(x) := (Ax - b)^T (Ax - b),$$

set the step size $s \in \{0.25, 1.0\}$, and the tolerance to $\|\nabla f(x_k)\| < 10^{-5}$. All derivatives are computed using symmetric differences with $h = 0.001$. For comparison purposes, we recall that the true answer is

$$x^* = \begin{bmatrix} 1.065152e + 00 \\ -6.257576e - 01 \\ 2.454545e + 00 \end{bmatrix}. \quad (12.42)$$

For a step size of $s = 1.0$, and starting from the randomly generated initial condition

$$x_0 := \begin{bmatrix} 1.923764e + 00 \\ 5.579350e + 00 \\ 3.273492e - 01 \end{bmatrix}, \quad (12.43)$$

the algorithm converges in one step to

$$x^* \approx \begin{bmatrix} 1.065152e + 00 \\ -6.257576e - 01 \\ 2.454545e + 00 \end{bmatrix}. \quad (12.44)$$

Starting from the same initial condition and using a step size of $s = 0.25$, the algorithm converges in 58 iterations to

$$x^* \approx \begin{bmatrix} 1.065152e + 00 \\ -6.25757e - 01 \\ 2.454545e + 00 \end{bmatrix}. \quad (12.45)$$

Julia code associated with the above is given below. ■

```

1 # Hessian on least squares
2 dataSet2=[
3 1 0 1.0
4 2 0.25 1.0
5 3 0.5 1.5
6 4 0.75 2.0
7 5 1.0 3.0
8 6 1.25 4.25
9 7 1.5 5.5
10 8 1.75 7.0
11 9 2.0 10.0]
12 X=dataSet2[:,2]
13 Y=dataSet2[:,3]
14 Phi=[ones(9,1) X X.^2]
15 alphaStar=(Phi'*Phi)\(Phi'*Y)

```

```

16 @show alphaStar # known optimal solution
17 #
18 F(x) = ( (Phi*[x[1];x[2];x[3]]-Y)'*(Phi*[x[1];x[2];x[3]]-Y) )

```

Output

```
alphaStar = [1.065151515151514, -0.6257575757575752, 2.4545454545454546]
```

```
F (generic function with 1 method)
```

```

1 function gradHess(f, x0, h=1e-3, delta=1e-3)
2     n=length(x0)
3     m=length(f(x0))
4     if m>1
5         return 0 # f does not map into R
6     end
7     H=zeros(n, n)
8     myGgrad=zeros(1, n)
9     Id=zeros(n, n)+I
10    for i=1:n
11        ei = Id[:,i]
12        myGgrad[i]=(f(x0+ h*ei) - f(x0 - h*ei)) [1]/(2*h)
13        for j=1:n
14            ej = Id[:,j]
15            H[i, j]=(f(x0+h*ei+delta*ej) -f(x0+h*ei-delta*ej) -f(x0-h*ei+delta*ej) +f(x0-h*ei-delta*ej) )
16            [1]/(4*h*delta)
17        end
18    end
19    return myGgrad, H

```

Output

```
gradHess (generic function with 3 methods)
```

```

1 xk = [1.9237640987871218; 5.579349855694035; 0.32734915035269596]
2 myGrad_xk, Hess_xk = gradHess(F, xk)
3
4 s=.25; #step size for Hessian search
5 k=0
6 while (k<1e3) & (norm(myGrad_xk) >1e-5)
7     Dxx=Hess_xk\(-myGrad_xk')
8     #Dxx=inv(Hess_xk)*(myGrad_xk') # Less desirable alternative
9     xk = xk + s*Dxx
10    myGrad_xk, Hess_xk = gradHess(F, xk)
11    k = k + 1
12 end
13
14 display([k F(xk) norm(myGrad_xk) det(Hess_xk)])
15 xk

```

Output

```
1x4 Matrix{Float64}:
 58.0  0.450379  9.69047e-6  324.844
```

```
3x1 Matrix{Float64}:
```


1.0651515638315818
 -0.6257572239523093
 2.454545333941787

Example 12.5 We revisit the extrinsic calibration problem of Chapter 12.4.1, but this time we use the Hessian. The update equation in (12.21) is replaced with

$$\begin{aligned} [\nabla^2 f(H_k, X, Y)] \Delta H_k &= -\nabla f(H_k, X, Y)^\top \\ H_{k+1} &= H_k + s\Delta H_k. \end{aligned} \quad (12.46)$$

After 14 iterations (compared to 5745 iterations when using the gradient descent method) with a step size of 1.5, the problem converges, and the cost drops to 12.113. The result is shown in Fig. 12.11. Compared to the first-order (gradient descent) method, the second-order method (using the Hessian) is much faster!




Figure 12.11: This figure demonstrates the calibration result of the Hessian method. The results look similar to the gradient descent algorithm in Fig 12.9 but the convergence speed when using the Hessian is 400 times faster than when using gradient descent.

The code to generate the results and figures for this problem is available on [GitHub](#)⁵ in MATLAB; how to do it in Julia is left as a homework problem. Given the hints already provided in this Chapter, we expect you will have no trouble with it.

Example 12.6 Fitting with Radial Basis Functions, $\text{rbf}(x) := ae^{-(x-x_c)^2/(2s^2)}$, where a is a weight, x_c is called a center, and s is the width. We re-visit an example that was included in Project 2. We are given noisy measurements of the function

$$f(x) := \cos(2\pi x)e^{-1}, \quad (12.47)$$

as shown in Fig. 12.12. From Project 2, we have a fit with three radial basis functions,

$$\hat{f}(x) = 0.260 - 0.409e^{\frac{(x-2.39)^2}{2(0.25)^2}} - 0.464e^{\frac{(x-1.59)^2}{2(0.25)^2}} + 0.137e^{\frac{(x-2.06)^2}{2(0.25)^2}},$$

where the width parameter set to $s = 0.25$ for each function and the three centers were fixed as $x_c := [2.39, 1.59, 2.06]$. Here, we'll optimize the width parameters, centers, and weights. The data is as follows:

⁵A platform we use to share our code. Just like you share your photos on Instagram!

$$\begin{aligned}
 [x_{\text{measured}} \quad y_{\text{measured}}] = & \begin{bmatrix} 2.390 & -0.104 \\ 1.870 & 0.112 \\ 2.130 & 0.074 \\ 1.470 & -0.240 \\ 1.000 & 0.381 \\ 2.090 & 0.088 \\ 1.020 & 0.334 \\ 2.160 & 0.069 \\ 2.040 & 0.132 \\ 1.440 & -0.204 \\ 2.030 & 0.128 \\ 2.170 & 0.017 \\ 1.290 & -0.020 \\ 1.720 & -0.032 \\ 1.260 & -0.004 \\ 1.160 & 0.172 \\ 2.400 & -0.085 \\ 1.850 & 0.068 \\ 1.990 & 0.137 \\ 1.340 & -0.162 \\ 1.700 & -0.068 \\ 1.590 & -0.194 \\ 2.240 & -0.006 \\ 2.110 & 0.109 \\ 2.440 & -0.081 \\ 2.270 & 0.004 \\ 1.170 & 0.185 \\ 1.600 & -0.134 \\ 1.760 & -0.024 \\ 1.900 & 0.137 \\ 1.920 & 0.145 \\ 2.000 & 0.142 \\ 2.250 & -0.007 \\ 1.680 & -0.074 \\ 1.660 & -0.147 \\ 2.480 & -0.051 \\ 2.450 & -0.049 \\ 1.250 & 0.022 \\ 2.420 & -0.099 \\ 1.740 & -0.027 \\ 1.620 & -0.128 \\ 1.300 & -0.051 \\ 1.280 & -0.052 \\ 2.060 & 0.115 \end{bmatrix}
 \end{aligned}
 \tag{12.48}$$

Solution: We will fit a function

$$f_{\alpha}(x) := a_0 + a_1 e^{-(x-x_{c,1})^2/(2s_1^2)} + a_2 e^{-(x-x_{c,2})^2/(2s_2^2)} + a_3 e^{-(x-x_{c,3})^2/(2s_3^2)},
 \tag{12.49}$$




Figure 12.12: The orange dots are the measured data. The green line is a fit via regression from Project 2.

where α , the vector of unknowns is

$$\alpha = \begin{bmatrix} s_1 \\ s_2 \\ s_3 \\ x_{c,1} \\ x_{c,2} \\ x_{c,3} \\ a_0 \\ a_1 \\ a_2 \\ a_3 \end{bmatrix}. \quad (12.50)$$

We define $Y = y_{\text{measured}}$ and

$$\hat{Y}_\alpha := \begin{bmatrix} a_0 + a_1 e^{-(x_1 - x_{c,1})^2 / (2s_1^2)} + a_2 e^{-(x_1 - x_{c,2})^2 / (2s_2^2)} + a_3 e^{-(x_1 - x_{c,3})^2 / (2s_3^2)} \\ a_0 + a_1 e^{-(x_2 - x_{c,1})^2 / (2s_1^2)} + a_2 e^{-(x_2 - x_{c,2})^2 / (2s_2^2)} + a_3 e^{-(x_2 - x_{c,3})^2 / (2s_3^2)} \\ \vdots \\ a_0 + a_1 e^{-(x_N - x_{c,1})^2 / (2s_1^2)} + a_2 e^{-(x_N - x_{c,2})^2 / (2s_2^2)} + a_3 e^{-(x_N - x_{c,3})^2 / (2s_3^2)} \end{bmatrix}, \quad (12.51)$$

where x_1, x_2, \dots, x_N are values from x_{measured} . We these definitions, the function to be minimized is

$$g(\alpha) := (Y - \hat{Y}_\alpha)^\top (Y - \hat{Y}_\alpha). \quad (12.52)$$

We apply the Hessian to find a minimizing set of parameters for g , namely

$$\nabla^2 g(\alpha_k) \Delta \alpha_k = -[\nabla g(\alpha_k)]^\top \quad (\text{solve for } \Delta \alpha_k) \quad (12.53)$$

$$\alpha_{k+1} = \alpha_k + \eta \Delta \alpha_k \quad (\text{use } \Delta \alpha_k \text{ to update our estimate of the optimal value}) \quad (12.54)$$

with $\eta = 0.2$. The result is

$$\alpha^* = \begin{bmatrix} 0.330 \\ 0.270 \\ 0.167 \\ 2.421 \\ 1.485 \\ 1.724 \\ 0.522 \\ -0.597 \\ -0.735 \\ 0.000 \end{bmatrix} \quad (12.55)$$

with a fitting error (minimum value of g) equal to 0.0165. Because $a_3 = 0$, the fit only uses two of the RBFs.




Figure 12.13: The orange dots are the measured data. The green line is a fit via regression from Project 2, and the violet line is our fit obtained with nonlinear optimization. All fits are given the freedom to use three RBFs.

Remark 4 In the code below, the function $g : \mathbb{R}^{10} \rightarrow \mathbb{R}$ is called *myFunErrorSquared*.

```

1 ThreeRBFs(x, weights, s_vec, centers) = weights[1] + weights[2]*rbf(x, centers[1], s_vec[1]) +
2   weights[3]*rbf(x, centers[2], s_vec[2]) + weights[4]*rbf(x, centers[3], s_vec[3])
3
4 function myFunErrorSquared(alpha, x=x_measured, y=y_measured)
5     Y=y_measured
6     N = length(Y)
7     Y_fit = zeros(N, 1)
8     #
9     s_vec = alpha[1:3]
10    centers=alpha[4:6]
11    weights=alpha[7:10]
12    #
13    for i = 1:N
14        Y_fit[i] = ThreeRBFs(x[i], weights, s_vec, centers)
15    end
16    #
17    y_error = Y-Y_fit
18    errorSquared = y_error' * y_error
19    return errorSquared[1]
20 end

```

Output

myFunErrorSquared (generic function with 3 methods)

```

1 h = .2
2 aTol = 1e-5
3 alphak=[.25*ones(3,1); centers; a_star]
4 errorSquared = myFunErrorSquared(alphak)

```

```

5 println (" Initial fitting error is $errorSquared")
6 k=0
7 errorSquared = myFunErrorSquared (alphak)
8 myGgrad, H = gradHess (myFunErrorSquared, alphak)
9 while (k<1e3) & (norm (myGgrad) >aTol)
10     myGgrad, H = gradHess (myFunErrorSquared, alphak)
11     Delta_alphak = -H\ (myGgrad')
12     alphak = alphak + s * Delta_alphak
13     k = k + 1
14 end
15 errorSquared = myFunErrorSquared (alphak)
16 println ("Final fitting error is $errorSquared")
17
18 @show k
19 @show norm (myGgrad) ;

```

Output

```

Initial fitting error is 0.1389839670139185
Final fitting error is 0.01651529922971189
k = 50
norm(myGgrad) = 7.899062216538802e-6

```

```

1 s_vec = alphak [1:3]
2 centers=alphak [4:6]
3 weights=alphak [7:10]
4
5 myRBFfit (x) =ThreeRBFs (x, weights, s_vec, centers)
6 xMin = minimum (x_measured)
7 xMax = maximum (x_measured)
8
9 p1=plot! (myRBFfit, xMin, xMax, label="NL Optim via Hessian", ylims=(-.25, 0.45))
10 png (p1, "OptimizationViaHessian")
11 display (p1)

```

Output See Fig. 12.13.

12.6 Local vs Global

Consider a function $f : \mathbb{R} \rightarrow \mathbb{R}$. Let's recall that the **graph** of the function is the collection of points

$$\mathbf{graph\ of\ f} := \{(x, f(x)) \mid x \in \mathbb{R}\}.$$

Alternatively, you may be more comfortable thinking of it as

$$\mathbf{graph\ of\ f} := \{(x, y) \mid x \in \mathbb{R}, y = f(x)\}.$$

We revisit the functions used in Fig. 12.1, where this time, in Fig. 12.14, we have also indicated line segments (in black) that connect various points on the graphs of the two functions (shown in red). The lines are given by

$$y = f(x_1) + \frac{f(x_2) - f(x_1)}{x_2 - x_1}(x - x_1), \quad (12.56)$$

for $x_1 \neq x_2$ in the domain of definition of the function.

If we can choose x_1 and x_2 such that the graph of the function is ever strictly above the corresponding line, as in Fig. 12.14b, then **the function is not convex and you can have local minima**. If, on the other hand, the graph of the function always lies below or just touches the line, for all $x_1 \neq x_2$, as in Fig. 12.14a, then **the function is convex and there are no local minima, only global minima!**




Figure 12.14: (a) This is a graph of our “simple” cost function with a global minimum at $x^* = 2$, while (b) a graph of our “less simple” cost function, where there are two local minima, one at $x^* \approx 0.68$ and one at $x^* \approx 3.14$, as well as a local maximum at ≈ 2.1 . In each case, we have overlaid line segments that are used to check for the mathematical property called **convexity**. A convex function only always has global minima. It does not have any local minima.




Figure 12.15: This is a plot of a convex function where there are many points achieving the global minimum of zero. In fact, the function is identically zero for $x \in [-0.5, 0.5]$.

Now, the value of x achieving the global minimum may not be unique. This happens when the “bottom bowl” of the function is a line with zero slope, as shown in Fig. 12.15.

Convex Functions

A function $f : \mathbb{R}^m \rightarrow \mathbb{R}$ is convex if for all $0 \leq \alpha \leq 1$, and for all $x, y \in \mathbb{R}^m$, the function satisfies

$$f(\alpha x + (1 - \alpha)y) \leq \alpha f(x) + (1 - \alpha)f(y). \quad (12.57)$$

In practice, checking this property is relatively hard at the present time, even for the case that $n = 1$, and is beyond the scope of our introduction. However, as you advance in your mathematical education, if ever a lecture is offered on **convex sets** or **convex functions**, you should pay attention! You might learn some cool and very useful stuff!

Convex optimization is not a typical subject for undergraduates in the present day and age, but maybe you guys will change that! It is becoming super important in engineering practice. Our purpose here is to let you know that such a subject as convex optimization exists so that you can be on the lookout for a course on the topic!

12.7 Maximization as Minimizing the Negative of a Function

The two main facts are summarized in the following equations.

$$\arg \max_{x \in \mathbb{R}^m} f(x) = \arg \min_{x \in \mathbb{R}^m} -f(x) \quad (12.58)$$

$$\max_{x \in \mathbb{R}^m} f(x) = - \min_{x \in \mathbb{R}^m} -f(x) \quad (12.59)$$

You should be able to convince yourself they are true by studying Fig. 12.16.




Figure 12.16: Maximization and minimization are almost the same thing! (a) Shows a function that we wish to maximize while (b) shows the negative of the function. From these plots, you can convince yourself that (12.58) and (12.59) hold!

12.8 (Optional Read): Quadratic Programs: Our first Encounter with Constraints

A **Quadratic Program** is a special kind of optimization problem with **constraints**. The cost to be minimized is supposed to be quadratic, meaning that $f : \mathbb{R}^m \rightarrow \mathbb{R}$ has the form

$$f(x) = \frac{1}{2}x^\top Qx + qx, \quad (12.60)$$

where Q is an $m \times m$ **symmetric matrix**, meaning that $Q^\top = Q$, and where q is a $1 \times m$ row vector. Moreover, instead of optimizing over all of \mathbb{R}^m as in our previous problems, we are allowed to seek solutions that lie in a subset of \mathbb{R}^m defined by **linear inequality** and **linear equality** constraints that are typically written in the form

$$A_{in}x \preceq b_{in} \quad (12.61)$$

$$A_{eq}x = b_{eq}. \quad (12.62)$$

The symbol \preceq is a way to define “less than or equal to” for vectors; it means that each component of the vector on the left hand side is less than or equal to the corresponding component of the vector on the right hand side. As an example

$$\begin{bmatrix} 3 \\ 2 \\ 4 \end{bmatrix} \preceq \begin{bmatrix} 4 \\ 3 \\ 4 \end{bmatrix},$$

though

$$\begin{bmatrix} 3 \\ 2 \\ 4 \end{bmatrix} \not\preceq \begin{bmatrix} 1 \\ 3 \\ 4 \end{bmatrix};$$

and

$$\begin{bmatrix} 3 & 1 \\ 2 & 4 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \preceq \begin{bmatrix} 0 \\ 9 \end{bmatrix},$$

means that x_1 and x_2 must satisfy

$$\begin{aligned} 3x_1 + x_2 &\leq 0 \\ 2x_1 + 4x_2 &\leq 9. \end{aligned}$$

What if you really wanted $2x_1 + 4x_2 \geq 9$? Then you need to remember that when you multiply both sides by a minus sign, the inequality sign flips. Hence,

$$\begin{aligned} 3x_1 + x_2 \leq 0 \\ 2x_1 + 4x_2 \geq 9 \end{aligned} \iff \begin{aligned} 3x_1 + x_2 \leq 0 \\ -2x_1 - 4x_2 \leq -9 \end{aligned} \iff \begin{bmatrix} 3 & 1 \\ -2 & -4 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \preceq \begin{bmatrix} 0 \\ -9 \end{bmatrix}.$$

In addition, most QP solvers allow one to specify lower and upper bounds on x of the form

$$lb \preceq x \preceq ub. \quad (12.63)$$

While such constraints could always be rewritten in the form of (12.61), using (12.63) is more convenient, intuitive, and less error prone. The inclusion of constraints allows for very interesting and practical optimization problems to be posed.

Useful Fact about QPs

We consider the QP

$$\begin{aligned} x^* = \arg \min_{x \in \mathbb{R}^m} \frac{1}{2} x^\top Q x + q x \quad (12.64) \\ A_{in} x \preceq b_{in} \\ A_{eq} x = b_{eq} \\ lb \preceq x \preceq ub \end{aligned}$$

and assume that Q is symmetric ($Q^\top = Q$) and **positive definite**^a ($x \neq 0 \implies x^\top Q x > 0$), and that the subset of \mathbb{R}^m defined by the constraints is non empty, that is

$$C := \{x \in \mathbb{R}^m \mid A_{in} x \preceq b_{in}, A_{eq} x = b_{eq}, lb \preceq x \preceq ub\} \neq \emptyset. \quad (12.65)$$

Then x^* exists and is unique.

^aPositive definite matrices are treated in Chapter A.3.

Example 12.7 *The very first optimization problem we came across in ROB 101 was least-squared-error solutions to systems of linear equations, $Ax = b$, back in Chapter 8.2, namely*

$$x^* = \arg \min_{x \in \mathbb{R}^m} \|Ax - b\|^2. \quad (12.66)$$

We used this formulation to solve regression problems in Chapter 8.3. Show that (12.66) is equivalent to the QP

$$x^* = \arg \min_{x \in \mathbb{R}^m} \frac{1}{2} x^\top Q x + q x, \quad (12.67)$$

where

$$\begin{aligned} Q &:= A^\top \cdot A \\ q &:= -b^\top \cdot Ax. \end{aligned} \quad (12.68)$$

In particular, it is a very simple QP, with no inequality constraints and no equality constraints.

Solution: We first note that $\|Ax - b\|^2 = (Ax - b)^\top \cdot (Ax - b)$, where we have used the fact that the norm squared of a vector $v \in \mathbb{R}^n$ is equal⁶ to $v^\top \cdot v$. Hence, multiplying out the terms, we have that

$$\begin{aligned} \|Ax - b\|^2 &= (Ax - b)^\top \cdot (Ax - b) \\ &= (x^\top A^\top - b^\top) \cdot (Ax - b) \\ &= x^\top A^\top \cdot Ax - b^\top \cdot Ax - x^\top A^\top \cdot b + b^\top \cdot b \\ &= x^\top A^\top \cdot Ax - 2b^\top \cdot Ax + b^\top \cdot b, \end{aligned} \quad (12.69)$$

⁶In our case, $v = (Ax - b)$.

where we have used the fact that

$$x^\top A^\top \cdot b = b^\top \cdot Ax$$

because each side is a scalar and the transpose of a scalar is itself. Next, we note that

$$\begin{aligned} x^* &= \arg \min_{x \in \mathbb{R}^m} \|Ax - b\|^2 \\ &= \arg \min_{x \in \mathbb{R}^m} \frac{1}{2} \|Ax - b\|^2 \\ &= \arg \min_{x \in \mathbb{R}^m} \left(\frac{1}{2} x^\top A^\top \cdot Ax - b^\top \cdot Ax + \frac{1}{2} b^\top \cdot b \right) \\ &= \arg \min_{x \in \mathbb{R}^m} \left(\frac{1}{2} x^\top A^\top \cdot Ax - b^\top \cdot Ax \right) \end{aligned}$$

because

- scaling the function to be minimized by a positive constant does not change **where the minimum occurs**, and hence does not change the value of $\arg \min$,
- and shifting the function to be minimized up or down by a constant does not change **where the minimum occurs**, and hence does not change the value of $\arg \min$!

■

Example 12.8 The second optimization problem we came across in ROB 101 was for underdetermined systems of linear equations, $Ax = b$, back in Chapter 9.9, namely

$$x^* = \arg \min_{Ax=b} \|x\|^2 = \arg \min_{Ax=b} x^\top x. \quad (12.70)$$

This too is a rather simple QP, with $Q = I_m$, the $m \times m$ identity matrix, $q = 0_{1 \times m}$, and no inequality constraints.

■




Figure 12.17: The plot shows in red the rectangle function $\text{rect}(x, a, b)$, for the values $a = 0.5$ and $b = 1.75$. The function takes on the value 1.0 for $a \leq x < b$ and equals 0.0 elsewhere. The dotted blue line is the monomial x^2 , which is being applied over the interval $[-1, 2]$. In black is the monomial x^2 **multiplied** by $\text{rect}(x, a, b)$, for the same values of a and b . The function $x^2 \cdot \text{rect}(x, a, b)$ takes on the value x^2 for all $a \leq x < b$ and 0.0 elsewhere. In other words, the action of the function is now localized to the set $[a, b) \subset \mathbb{R}$. This is the basic notion of a spline, being able to localize the action of a function to a subset of x values instead of having the function extend over the entire set of x values, as with the standard monomial x^2 . Splines typically give you more control in the fitting process than using high-order monomials.

Example 12.9 We'll now put the first two examples together, while introducing you to a new way to choose “basis functions” for regression, called splines; see Fig 12.17.

You are given the noisy data shown in Fig. 12.18. The objective is to fit a function to the data, much as we did in Chapter 8.3. We'll add a new twist by introducing a type of function called a “spline”, where even though we are using our familiar monomials, they will be localized to disjoint regions of the data. The solution will provide the details!




Figure 12.18: The true function (which you should pretend not to know) and a set of imperfect measurements from which you are to estimate the function.

Solution: We introduce a new function, called rectangle,

$$\text{rect}(x, a, b) = \begin{cases} 1 & a \leq x < b \\ 0 & \text{otherwise,} \end{cases}$$

that can be used to limit the action of a function to an interval of the form $[a, b)$. In Fig. 12.17, we plot the rectangle function itself and the monomial x^2 multiplied by the rectangle function.

We will use the rectangle function to divide the set $[-1, 3]$ into three subsets and fit low-degree polynomials on each subset. To do this, we define $x_{\min} := 1.0$, $x_{\max} := 2.0$, and further define

$$\Delta x := \frac{x_{\min} + x_{\max}}{3}$$

$$a_k := x_{\min} + (k - 1)\Delta x, \quad k \in 1, 2, 3, 4$$

so that

$$a = [-1.0, \frac{1}{3}, \frac{5}{3}, 3.0].$$

The components of a are called **knot points**; see [https://en.wikipedia.org/wiki/Spline_\(mathematics\)](https://en.wikipedia.org/wiki/Spline_(mathematics)) for more information. We note that $[-1, \frac{1}{3}) \cup [\frac{1}{3}, \frac{5}{3}) \cup [\frac{5}{3}, 3) = [-1, 3)$, so technically, we have left the point $x = 3$ out of consideration. If you wish to include it, just replace the last knot point with something like 3.001.

Inspired by Fig. 12.17, we define a regressor matrix which uses the standard monomials up to degree three on each set $[a_i, a_{i+1})$,

$$\Phi(x, a) := [1 \ x \cdot \text{rect}(x, a_1, a_2) \ \dots \ x^3 \cdot \text{rect}(x, a_1, a_2) \ \dots \ x \cdot \text{rect}(x, a_3, a_4) \ \dots \ x^3 \cdot \text{rect}(x, a_3, a_4)], \quad (12.71)$$

where, x is a (column) vector of (measurement) points and a is the set of knot points.

Figure 12.19 shows the resulting fit, which is a standard least squares problem

$$\alpha^* = \arg \min_{\alpha \in \mathbb{R}^{10}} \|Y - \Phi\alpha\|^2,$$

where Y is the vector of measured function values. We note right away the “jumps” in the fit at the two interior knot points, $a_2 = 1/3$ and $a_3 = 5/3$. The discontinuities arise from jumps in the rectangle function at the spline boundaries, as was seen in Fig. 12.17. We next show how to impose continuity at the boundaries of the splines.




Figure 12.19: The resulting fit of a polynomial spline of degree three. A discontinuity is clear at the two interior knot points, $a_2 = 1/3$ and $a_3 = 5/3$. We next show how to achieve continuity.

To impose continuity at the interior knot points, we will use a linear constraint on the regression coefficients, α . Let $\epsilon = 10^{-3}$ and define

$$\begin{aligned} A_1 &:= \Phi(a_2 - \epsilon, a) - \Phi(a_2 + \epsilon, a) \\ A_2 &:= \Phi(a_3 - \epsilon, a) - \Phi(a_3 + \epsilon, a) \\ A_{\text{eq}} &:= \begin{bmatrix} A_1 \\ A_2 \end{bmatrix}. \end{aligned} \quad (12.72)$$

It follows that $A_{\text{eq}}\alpha = 0_{2 \times 1}$ forces the splines to match up at the boundary points. Indeed, if we denote $\hat{y}(x) := \Phi(x, a)\alpha$ for an arbitrary $x \in \mathbb{R}$, then

$$A_{\text{eq}}\alpha = 0_{2 \times 1} \iff \hat{y}(a_i - \epsilon) = \hat{y}(a_i + \epsilon), i \in \{2, 3\},$$

which is what continuity is all about.

We then solve the least squares problem with a linear constraint, namely

$$\begin{aligned} \alpha^* &= \arg \min_{\alpha \in \mathbb{R}^{10}} \|Y - \Phi\alpha\|^2. \\ A_{\text{eq}}\alpha &= 0 \end{aligned} \quad (12.73)$$




Figure 12.20: The plot shows a spline fit with continuity imposed at the interior knot points. For comparison purposes, a polynomial with the same number of free parameters is shown.

Equation 12.73 is a quadratic program; indeed, one has $Q := \frac{1}{2}\Phi^T\Phi$, $q := -Y^T\Phi$, $B_{eq} = 0_{2 \times 1}$, and A_{eq} as above. Figure 12.20 shows the resulting fit. We note that indeed, we have removed the discontinuities in the regressed function. ■

Example 12.10 (A Graphical Example) We provide a graphical QP example to understand the relationship between the cost function and the constraints. Consider the cost function $J(x_1, x_2) = (x_1 - 2)^2 + (x_2 - 1)^2$ and the following constraints:

$$\begin{aligned} x_1 + 2x_2 &\leq 12 \\ 3x_1 + 3x_2 &\leq 25 \\ x_1 &\leq 7 \\ x_2 &\leq 5 \\ x_1 &\geq 0 \\ x_2 &\geq 0. \end{aligned}$$

Write the problem in the standard form of (12.64), In addition, provide a contour plot of the cost function with an overlay of the constraints.

Solution: We expand the cost function as

$$J(x_1, x_2) = (x_1 - 2)^2 + (x_2 - 1)^2 = x_1^2 + x_2^2 - 4x_1 - 2x_2 + 5.$$

The constant term, 5, has no effect on the optimal solution of this problem, and therefore, it is common to drop it.

Remark: Some software packages for solving QPs might not include a constant term. Remember to take it into account in the end if the actual value of the cost function at the optimal solution is required. In this particular example, because we can visualize the cost function and the constraints, we will keep it.




Figure 12.21: The contour plot of the cost function and constraints. The feasible region is visible in the left bottom corner of the figure.

We can now rearrange everything in the form of (12.64):

$$\begin{aligned} \min_{x \in \mathbb{R}^m} \quad & \frac{1}{2} \begin{bmatrix} x_1 & x_2 \end{bmatrix} \begin{bmatrix} 2 & 0 \\ 0 & 2 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \\ \text{subject to} \quad & \begin{bmatrix} 1 & 2 \\ 3 & 3 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \preceq \begin{bmatrix} 12 \\ 25 \end{bmatrix} \\ & \begin{bmatrix} 0 \\ 0 \end{bmatrix} \preceq \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \preceq \begin{bmatrix} 7 \\ 5 \end{bmatrix}. \end{aligned}$$

At this point, we are almost done. Find your favorite QP solver (such as the one in Chap. 12.9), enter your problem according to the requested format and press run! Next, enjoy the results!

Because we only have two variables here, we can visualize the cost and constraints. A similar graphical approach would not be viable for large-scale problem where we might have thousands of variables. Figure 12.21 shows a contour plot of the cost function and our constraints. The region where all constraints are satisfied is called the **feasible region**. The optimal value of our problem must lie within the feasible region. Applying the QP solver in Chap. 12.9, we obtain $\begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} 2 \\ 1 \end{bmatrix}$. You can check that it does not violate the constraints. ■

12.9 (Optional Read): QP Solver in Julia

We've had success with the QP solver at <https://osqp.org>, called OSQP. The standard form used by OSQP is a bit different than (12.64), though it is every bit as general,

$$\begin{aligned} x^* = \quad & \arg \min_{x \in \mathbb{R}^m} \quad \frac{1}{2} x^\top Q x + q^\top x, & (12.74) \\ & \text{lb} \preceq A x \preceq \text{ub} \end{aligned}$$

where $x \in \mathbb{R}^m$ is the optimization variable. The objective function is defined by a positive semidefinite $m \times m$ matrix Q and vector $q \in \mathbb{R}^m$. The linear inequality and equality constraints as well as upper and lower bounds are grouped together and defined by a single $n \times m$ matrix A and two $n \times 1$ vectors lb and ub , where each component of $lb_i \in \mathbb{R} \cup \{-\infty\}$ and $ub_i \in \mathbb{R} \cup \{+\infty\}$, $i \in \{1, \dots, n\}$. To impose equality constraints, one sets the corresponding entries of l and u to be equal to one another. We provide a script and an illustration below to transform a problem in the form of (12.64) to that of (12.74).

```

1 using Pkg
2 pkg.add("OSQP")
3 Pkg.add("Compat")
4 using OSQP
5 using SparseArrays
6
7 # Define problem data
8 P = sparse([4. 1.; 1. 2.])
9 q = [1.; 1.]
10 A = sparse([1. 1.; 1. 0.; 0. 1.])
11 l = [1.; 0.; 0.]
12 u = [1.; 0.7; 0.7]
13
14 # Create OSQP object
15 prob = OSQP.Model()
16
17 # Setup workspace and change alpha parameter
18 OSQP.setup!(prob; P=P, q=q, A=A, l=l, u=u)
19
20 # Solve problem
21 results = OSQP.solve!(prob)

```

Here is a function that takes QPs formulated as in (12.64).

```

1 using LinearAlgebra
2 using OSQP
3 using SparseArrays
4
5 function quadProg(Q, q, Ain, bin, Aeq, beq, lb, ub, tol=1e-6)
6     # Begin wrapper to make QP solver in OSQP similar to quadprog in Matlab
7     # Need to ensure that matrices are sparse and any n x 1 "matrices" are turned into
8     # vectors
9     dimX = length(q)
10    myI = sparse(zeros(dimX, dimX) + I)
11    tolEq = tol
12    # Define problem data
13    P = sparse(Q)
14    q = q[:]
15    A = sparse(Ain) #
16    u = bin[:] #how to force objects to be vectors
17    l = bin[:] .-Inf
18    if (length(lb) > 0) || (length(ub) > 0)
19        A = [A; myI]
20    end
21    if (length(ub) > 0) & (length(lb) > 0)
22        u = [u; ub[:]]
23        l = [l; lb[:]]
24    elseif (length(ub) > 0) & (length(lb) == 0)
25        u = [u; ub[:]]
26        l = [l; ub[:] .-Inf]
27    elseif (length(ub) == 0) & (length(lb) > 0)

```

```

27     l=[I; lb[:]]
28     u=[u; lb[:].+Inf]
29 end
30 (nreq, nceq)=size(Aeq)
31 if nreq > 0
32     A=[A; sparse(Aeq)]
33     l=[I; beq[:].-tolEq]
34     u=[u; beq[:].+tolEq]
35 end
36 # End wrapper
37
38 # Create OSQP object
39 prob = OSQP.Model()
40
41 # Setup workspace and change alpha parameter
42 OSQP.setup!(prob; P=P, q=q, A=A, l=l, u=u, verbose=false)
43
44 # Solve problem
45 results = OSQP.solve!(prob)
46 return results.x
47 end

```

```

1
2 # Example problem data (same as above)
3 P = sparse([2. 0.; 0. 2.])
4 q = [-4.; -2.]
5 A = sparse([1. 2.; 3. 3.; 1. 0.; 0. 1.])
6 l = [0.; 0.; 0.; 0.]
7 u = [12.; 25.; 7.; 5.]
8
9 dimX=length(q)
10 Aeq = Array{SparseMatrixCSC, 2}(undef, 0, dimX)
11 Beq = Vector{Float64}(undef, 0)
12 lb = Vector{Float64}(undef, dimX).-Inf
13 ub = Vector{Float64}(undef, dimX).+Inf
14
15 xStar = quadProg(P, q, [A; -A], [u; -l], Aeq, Beq, lb, ub)

```

Output

```

2-element Vector{Float64}:
 1.9999978269253138
 0.9999972171335297

```

12.10 (Optional Read): Optimization Packages: The Sky is the Limit

Once you’ve coded up a few optimization algorithms, it’s time to move over and let the pros handle the programming while you focus on problem formulation. Currently, the best source for optimization packages in Julia is <https://jump.dev/>. From the **JuMP** homepage we learn that “JuMP is a modeling language and supporting packages for mathematical optimization in Julia. JuMP makes it easy to formulate and solve linear programming, semidefinite programming, integer programming, convex optimization, constrained nonlinear optimization, and related classes of optimization problems. You can use it to route school buses, schedule trains, plan power grid expansion, or even optimize milk output.”

What are you waiting for? There is a lot of knowledge out there. Slowly but surely, you can master it!

12.11 Looking Ahead

The next Chapter is designed to prepare you for CS courses at the University of Michigan that require Linear Algebra, specifically, Machine Learning, EECS 445, and Computer Vision, EECS 442.

Chapter 13

Background for Classification and Machine Learning

Learning Objectives

- Introduce material that is assumed in UofM Computer Science courses that have Math 214 as a prerequisite.
- Provide a resource for use after you leave ROB 101.

Outcomes

- Learn how to separate \mathbb{R}^n into two halves via hyperplanes
- What is the “signed” distance from a point to a hyperplane and how to compute it
- An example of a max-margin classifier, a common tool in Machine Learning
- Learn the Orthogonal Projection Theorem, which is “the geometric tool” that underlies most least squares problems

13.1 Separating Hyperplanes

We continue with a geometric development that is a natural accompaniment to Chapter 9: linear structures that can be used to divide \mathbb{R}^n into two pieces. The notes are based on lectures by Prof. Maani Ghaffari. This material is used in EECS 445, the undergraduate Machine Learning course, where one seeks to separate observations of a process into two categories, such as spam versus regular email, images of cats versus dogs, or a smooth walking surface versus one that undulates. The observations are typically given in the form of n -vectors and are called *object features*. Once you can handle the task of separating two categories, you are on the road to handling multiple categories, as in Fig. 13.1.




Figure 13.1: This awesome figure shows multiple (hyper)planes dividing \mathbb{R}^3 into disjoint regions, where each region could contain features describing a different object. In this book, we will content ourselves with a single hyperplane and do it in general for \mathbb{R}^n . Image courtesy of Kilom691 <https://commons.wikimedia.org/w/index.php?curid=37508909>.

We will develop the notion of a “hyperplane” as a linear object that is big enough to divide \mathbb{R}^n into two halves, easy to manipulate, and can take on “any orientation or position.” In \mathbb{R}^2 , any line can divide the space into two half spaces. In \mathbb{R}^3 , a line is not “big enough” to divide the space into two parts, though the the classical notion of a plane does the job perfectly well. In \mathbb{R}^n , the appropriate generalization is called a **hyperplane**!

Before we dig into the details, we firm up concepts in \mathbb{R}^2 . While we skip the case of the real line, \mathbb{R} , it does provide some insight because a single point $x_c \in \mathbb{R}$ can be used to divide the real line into two halves, $H^- := \{x \in \mathbb{R} \mid x < x_c\}$ and $H^+ := \{x \in \mathbb{R} \mid x > x_c\}$. Moreover, the object being used to divide the vector space \mathbb{R} into two halves has dimension zero, which is one less than the dimension of \mathbb{R} ! Hold this thought.

13.1.1 Lines in \mathbb{R}^2 as Separating Hyperplanes

While we are very used to describing lines as things that satisfy formulas of the form $x_2 = mx_1 + b$, let’s note that this description leaves out the x_2 -axis, which is a perfectly fine line in \mathbb{R}^2 . It also leaves out all lines parallel to the x_2 -axis. Why is the x_2 -axis not covered by this description? Because it’s slope would be infinity, which is not allowed! A better way to describe a line is actually as a special kind of subset of \mathbb{R}^2 , such as

$$\text{Line} := \{(x_1, x_2) \in \mathbb{R}^2 \mid a_0 + a_1x_1 + a_2x_2 = 0\},$$

where at least one of a_1 and a_2 is non-zero. Indeed, with this formulation,

$$x_2\text{-axis} = \{(x_1, x_2) \in \mathbb{R}^2 \mid 0 + x_1 + 0x_2 = 0\}.$$

Claim 1: Every line in \mathbb{R}^2 can be written as the zero set of $y(x_1, x_2) = a_0 + a_1x_1 + a_2x_2$, where at least one of a_1 and a_2 is non-zero.




Figure 13.2: We know that subspaces must contain the origin. Lines in \mathbb{R}^2 can be viewed as translations of subspaces (loosely speaking, this means you slide them up, down, or sideways, without rotating them). In both figures, the lines corresponding to subspaces are in red while their translations are in green and blue. The blue and green lines are parallel to the red line, but are offset or translated. In (a), the lines correspond to $0 = a_0 + 1.0x_1 - 2.0x_2$, where $a_0 = 0.0$ (red), $a_0 = 3.0$ (blue), and $a_0 = -2.0$ (green) (b) The lines correspond to $0 = a_0 + 1.0x_1 + 0.0x_2$, where $a_0 = 0.0$ (red), $a_0 = 2.0$ (blue), and $a_0 = -3.0$ (green)

Proof: If the line is given by $x_2 = mx_1 + b$, then it is the zero set of $y(x_1, x_2) = b + mx_1 - x_2$, that is, $a_0 = b$, $a_1 = m$ and $a_2 = -1$. If the line is parallel to the x_2 -axis, as in $\{(x_1, x_2) \mid x_1 = a_0 = \text{a constant}, x_2 \in \mathbb{R}\}$, then we can take $y = a_0 - x_1 + 0x_2$, that is, $a_1 = -1$ and $a_2 = 0$. ■

While it may not be apparent that writing a line as the zero set of a function has any value, we next note that the function

$$y(x_1, x_2) = a_0 + a_1x_1 + a_2x_2$$

can also be used to divide \mathbb{R}^2 into two halves. Indeed, we define

$$H^+ := \{(x_1, x_2) \in \mathbb{R}^2 \mid y(x_1, x_2) = a_0 + a_1x_1 + a_2x_2 > 0\}$$

$$H^- := \{(x_1, x_2) \in \mathbb{R}^2 \mid y(x_1, x_2) = a_0 + a_1x_1 + a_2x_2 < 0\}.$$

This is illustrated in Fig 13.3, where the line in red is the set where $y(x_1, x_2) = a_0 + a_1x_1 + a_2x_2 = 0$, showing the utility of thinking of a line as a zero set of a function.

H^+ and H^- are Called Half Spaces

H^+ and H^- are called Half Spaces because they divide \mathbb{R}^2 into two halves. Is that really possible? In Fig 13.3, the red lines are the sets where $y(x_1, x_2) = a_0 + a_1x_1 + a_2x_2 = 0$. We indicated H^+ and H^- as being on opposite sides of the red lines. Does it have to be this way? Can these sets be mixed up, meaning parts of H^+ and H^- can be on the same side of the red line? The answer is a **DEFINITIVE NO!** One side of the red line will be H^+ and the other will necessarily be H^- . To determine which is which, just sample one point on one of the sides and check if y at that point is positive or negative. It's that simple.

Following this box, we give the optional proof.

(Optional) Proof that Half Spaces Work as we Claim: Here is why the line $y(x_1, x_2) = 0$ separates \mathbb{R}^2 into **two half spaces**, H^+ and H^- .

Suppose that $(x_1^+, x_2^+) \in \mathbb{R}^2$ is such that $y^+ := y(x_1^+, x_2^+) > 0$ and similarly, $(x_1^-, x_2^-) \in \mathbb{R}^2$ is such that $y^- := y(x_1^-, x_2^-) < 0$. Let $\alpha \in \mathbb{R}$ and define a new point $(x_1(\alpha), x_2(\alpha)) \in \mathbb{R}^2$ by

$$\begin{bmatrix} x_1(\alpha) \\ x_2(\alpha) \end{bmatrix} := (1 - \alpha) \begin{bmatrix} x_1^+ \\ x_2^+ \end{bmatrix} + \alpha \begin{bmatrix} x_1^- \\ x_2^- \end{bmatrix}.$$

We note that varying $\alpha \in \mathbb{R}$ traces out a line in \mathbb{R}^2 that passes through (x_1^+, x_2^+) when $\alpha = 0$ and through (x_1^-, x_2^-) when $\alpha = 1$.

We next note that we can write $y(x_1, x_2)$ as

$$y(x_1, x_2) = a_0 + a_1x_1 + a_2x_2 = a_0 + \begin{bmatrix} a_1 & a_2 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$$

and therefore,

$$\begin{aligned} y(x_1(\alpha), x_2(\alpha)) &= a_0 + \begin{bmatrix} a_1 & a_2 \end{bmatrix} \begin{bmatrix} x_1(\alpha) \\ x_2(\alpha) \end{bmatrix} \\ &= (1 - \alpha)a_0 + (1 - \alpha) \begin{bmatrix} a_1 & a_2 \end{bmatrix} \begin{bmatrix} x_1^+ \\ x_2^+ \end{bmatrix} + \alpha a_0 + \alpha \begin{bmatrix} a_1 & a_2 \end{bmatrix} \begin{bmatrix} x_1^- \\ x_2^- \end{bmatrix} \\ &= (1 - \alpha)y(x_1^+, x_2^+) + \alpha y(x_1^-, x_2^-) \\ &= (1 - \alpha)y^+ + \alpha y^-. \end{aligned}$$

Solving for α^* to set $y(x_1(\alpha^*), x_2(\alpha^*)) = 0$ yields

$$\alpha^* = \frac{y^+}{y^+ - y^-} = \frac{y^+}{y^+ + |y^-|},$$

where we have used the fact that $y^- < 0 \implies -y^- = |y^-|$. Because $|y^-| > 0$ it follows that $0 < \alpha^* < 1$. Hence, there is a unique point $(x_1(\alpha), x_2(\alpha))$ **strictly between** (x_1^+, x_2^+) and (x_1^-, x_2^-) where $y(x_1(\alpha), x_2(\alpha)) = 0$. All points where $y(x_1, x_2)$ vanishes lie on a red line. Hence, the point $(x_1(\alpha^*), x_2(\alpha^*))$ lies on a red line. The only way this can happen is if (x_1^+, x_2^+) and (x_1^-, x_2^-) lie on opposite sides of a red line. ■

In the next subsection, we want to extend these ideas to \mathbb{R}^n for $n > 2$. While we could stick with formulas of the form

$$y(x_1, \dots, x_n) = a_0 + a_1x_1 + \dots + a_nx_n, \tag{13.1}$$

a more insightful analysis comes about from working directly with subspaces, which was hinted at in Fig. 13.1 and 13.2.

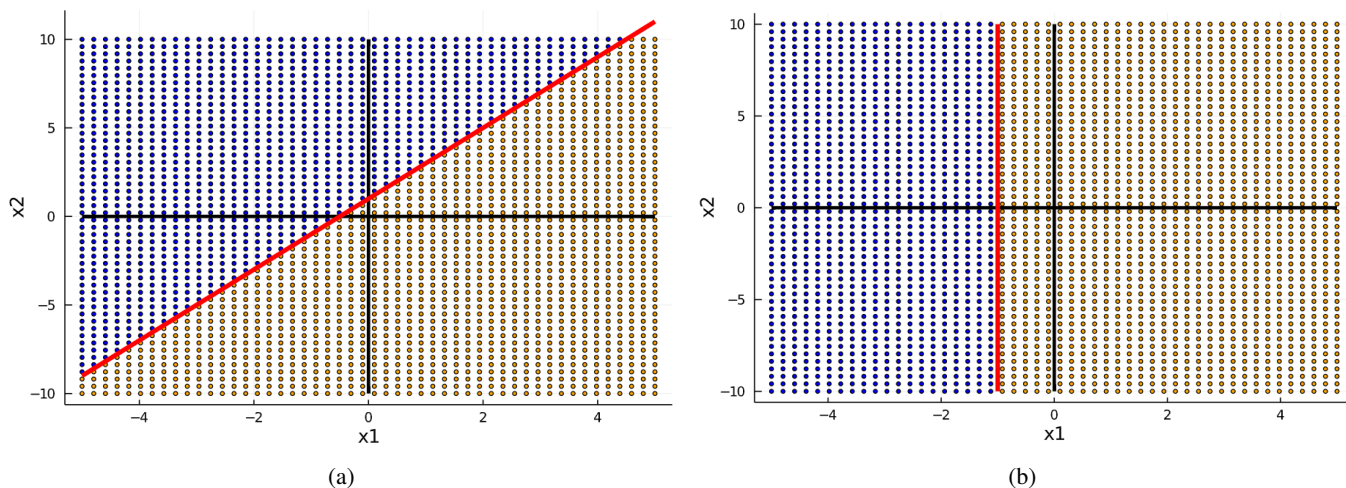


Figure 13.3: Two examples of half spaces corresponding to (a) $y = -1.0 - 2.0x_1 + x_2$ (b) $y = -1.0 - 1.0x_1 + 0.0x_2$. The line $y = a_0 + a_1x_1 + a_2x_2 = 0$ is shown in red, while in blue is plotted the set where $y > 0$ and in brown and the set where $y < 0$. The x_1 -axis and x_2 -axis are in black.

13.1.2 Hyper Subspaces

Consider the vector space \mathbb{R}^n and let A be a $1 \times n$ matrix (you can also call it a row vector). We assume that A is non-zero, meaning that at least one of its components is non-zero. Viewing A as a matrix, we know that its null space

$$N := \text{null}(A) = \{x \in \mathbb{R}^n \mid Ax = 0\}$$

is a subspace of \mathbb{R}^n . By the **Rank-Nullity Theorem**, the dimension of N is equal to $n - 1$, one less than the dimension of \mathbb{R}^n . Why, because $\text{rank}(A) = 1$ due to our assumption that at least one of its columns¹ is nonzero and

$$\dim(N) = \text{nullity}(A) = \dim(\mathbb{R}^n) - \text{rank}(A) = n - 1.$$

A subspace with dimension one less than the ambient space in which it lives, which in our case is \mathbb{R}^n , is called a **co-dimension one** subspace. Though less common, you can also call it a **hyper-subspace**!

We've just seen that the null space of a rank one matrix gives rise to a co-dimension one subspace. Are all co-dimension one subspaces the null space of some matrix? The answer is yes and the easiest way to show it is by using the dot product and the Gram-Schmidt process! What? You did not see that coming?

We write $A =: a^\top$ where

$$a := \begin{bmatrix} a_1 \\ a_2 \\ \vdots \\ a_n \end{bmatrix} \in \mathbb{R}^n.$$

We do this because

$$x \in \text{null}(A) \iff Ax = 0 \iff a^\top x = 0 \iff a \bullet x = 0 \iff x \bullet a = 0 \iff x \perp a.$$

Hence, our question of whether every co-dimension one subspace can be expressed as the null space of a rank one matrix can be rephrased as "is every co-dimension one subspace equal to the set of all vectors that are orthogonal to a non-zero vector $a \in \mathbb{R}^n$?" To answer this question, we can invoke Gram-Schmidt.

We let $N \subset \mathbb{R}^n$ be a co-dimension one subspace, meaning $\dim(N) = n - 1$. We let $\{u_1, \dots, u_{n-1}\}$ be a basis for N . Because N is not all of \mathbb{R}^n , there must exist a non-zero vector $u_n \in \mathbb{R}^n$ such that $u_n \notin N$. We skip the details, but you can then show that

$$\{u_1, \dots, u_{n-1}, u_n\}$$

is a linearly independent set. We apply Gram-Schmidt to produce an orthogonal basis $\{v_1, \dots, v_{n-1}, v_n\}$. By (9.17), we have that

$$N = \text{span}\{u_1, \dots, u_{n-1}\} = \text{span}\{v_1, \dots, v_{n-1}\}.$$

Moreover,

$$x \in N \iff x = \alpha_1 v_1 + \dots + \alpha_{n-1} v_{n-1} \iff x \perp v_n \iff v_n \bullet x = 0$$

Hyper Subspaces and Dot Products

The following are equivalent for a subspace $N \subset \mathbb{R}^n$:

- $\dim(N) = n - 1$, that is, N is a co-dimension one subspace;
- there exists $a \in \mathbb{R}^n$ not equal to zero such that $x \in N \iff x \perp a$; and
- there exists a $1 \times n$ matrix A such that $A \neq 0_{1 \times n}$ and $N = \text{null}(A)$.

We note that the matrix A and the vector a are related by $A = a^\top$.

Example 13.1 Consider a matrix $B = \begin{bmatrix} 1 & -1 \\ -1 & 2 \\ 0 & 1 \end{bmatrix}$ and let $N := \text{col span}\{B\}$. It is clear that $N \subset \mathbb{R}^3$ and $\dim(N) = 2$. Find a

vector $a \in \mathbb{R}^3$ such that

$$N = \{x \in \mathbb{R}^3 \mid a \bullet x = 0\}.$$

¹For a $1 \times n$ matrix, elements and columns are the same thing!

Solution: We define $u_1 = \begin{bmatrix} 1 \\ -1 \\ 0 \end{bmatrix}$, $u_2 = \begin{bmatrix} -1 \\ -2 \\ 1 \end{bmatrix}$, and note that $u_3 = \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}$ is linearly independent of $\{u_1, u_2\}$. Applying Gram Schmidt with normalization to $\{u_1, u_2, u_3\}$ yields

$$[v_1 \ v_2 \ v_3] = \begin{bmatrix} 0.707107 & 0.408248 & 0.57735 \\ -0.707107 & 0.408248 & 0.57735 \\ 0.000000 & 0.816497 & -0.57735 \end{bmatrix}.$$

Hence, we can take $a = \begin{bmatrix} 1 \\ 1 \\ -1 \end{bmatrix}$. It is easily checked that $a \bullet u_1 = 0$ and $a \bullet u_2 = 0$, and thus $N = \{x \in \mathbb{R}^3 \mid a \bullet x = 0\}$. ■




Figure 13.4: Let $S \subset \mathbb{R}^2$ be the star-shaped object in (a) and let x_c be the vector $[2; 3]$. Then $x_c + S$ is the star-shaped object in (b), where each and every point of the object has been shifted by x_c . If you can handle this, then you should have no trouble handling the translation of a line or a plane! Image courtesy of Tribhi Kathuria.

13.1.3 Translations of Sets, Hyper Subspaces, and Hyperplanes

Definition Let $S \subset \mathbb{R}^n$ be a subset and let $x_c \in \mathbb{R}^n$ be a vector. We define the translation of S by x_c as

$$x_c + S := \{x_c + x \mid x \in S\}.$$

Note that, because S consists of vectors in \mathbb{R}^n , the addition in the above formula makes sense. Figure 13.4 provides an illustration.

Claim: Let $N = \{x \in \mathbb{R}^n \mid a \bullet x = 0\} \subset \mathbb{R}^n$ be a hyper subspace (means that $a \neq 0_{n \times 1}$) and let $x_c \in \mathbb{R}^n$ be a vector. Then their sum has a simple description as

$$x_c + N = \{x \in \mathbb{R}^n \mid a \bullet (x - x_c) = 0\} = \{x \in \mathbb{R}^n \mid a \perp (x - x_c)\}. \quad (13.2)$$

The proof is not important, but we give it for those who are interested. N consists of everything in \mathbb{R}^n that is orthogonal to a . Hence,

if $(x - x_c) \perp a$, then $(x - x_c) \in N$. Adding x_c to both sides, we have that $x \in x_c + N$. The other direction is similar. ■

Hyperplanes are Translations of Hyper Subspaces

Let $N = \{x \in \mathbb{R}^n \mid a \bullet x = 0\} \subset \mathbb{R}^n$ be a hyper subspace (means that $a \neq 0_{n \times 1}$) and let $x_c \in \mathbb{R}^n$ be a vector. Then

$$H := x_c + N \quad (13.3)$$

is called a **hyperplane**. Moreover, by (13.2), the real-valued function $y : \mathbb{R}^n \rightarrow \mathbb{R}$ defined by

$$y(x) := a \bullet (x - x_c) \quad (13.4)$$

vanishes on $H = x_c + N$ (because $H = \{x \in \mathbb{R}^n \mid y(x) = a \bullet (x - x_c) = 0\}$). It follows that $y(x)$ can be used to divide \mathbb{R}^n into two halves

$$\begin{aligned} H^+ &:= \{x \in \mathbb{R}^n \mid y(x) > 0\} \\ H^- &:= \{x \in \mathbb{R}^n \mid y(x) < 0\}. \end{aligned} \quad (13.5)$$

A fanciful illustration is given in Fig. 13.5.

We note that H^+ is the set of all vectors $x \in \mathbb{R}^n$ such that the dot product $\langle a, x - x_c \rangle = a \bullet (x - x_c) > 0$, while H^- is the set of all vectors $x \in \mathbb{R}^n$ such that $\langle a, x - x_c \rangle = a \bullet (x - x_c) < 0$.

Remarks:

- Without loss of generality, it is always possible to take $x_c = \alpha a$, for $\alpha \in \mathbb{R}$. Indeed, one can go back and forth between (13.4) and (13.1) by

$$a_0 = -a \bullet x_c \quad \text{and} \quad x_c = -a_0 \frac{a}{\|a\|^2} \quad (13.6)$$

- Vectors such that $\langle a, x - x_c \rangle = a \bullet (x - x_c) > 0$ form an **acute angle (less than 90°)** with respect to the vector a , while vectors such that $\langle a, x - x_c \rangle = a \bullet (x - x_c) < 0$ form an **oblique angle (greater than 90°)** with respect to the vector a . Vectors such that $\langle a, x - x_c \rangle = a \bullet (x - x_c) = 0$ form a **right angle (exactly $\pm 90^\circ$)** with respect to the vector a , and hence are on the hyperplane itself.




Figure 13.5: A separating hyperplane where the features are smiles versus frowns. In Example 13.6, we show how to design the parameters of the hyperplane, that is, $a \in \mathbb{R}^n$ and $x_c \in \mathbb{R}^n$, so as to achieve separation for given data on the features.

13.2 Signed Distance to a Hyperplane

The function $y(x) = a \cdot (x - x_0)$ has another amazing feature: its absolute value is proportional to the distance of a point x from the hyperplane H^0 defined by $H^0 := \{x \in \mathbb{R}^n \mid y(x) = 0\}$. Moreover, $y(x)$ gives rise to the notion of the signed distance of x to H^0 because, depending on which of the half-planes x lies, the sign of $y(x)$ will be either $+1$ or -1 . To make sense of this, we must first define the distance of a point to a subspace, and then we will specialize to the case that the subspace is a hyper-subspace.

From Norms to Distances

Let $V \subset \mathbb{R}^n$ be a subspace and let $x_0 \in \mathbb{R}^n, y_0 \in \mathbb{R}^n, v_c \in \mathbb{R}^n$ be points. Then

- **Definition** $d(x_0, y_0) := \|x_0 - y_0\|$ is called the **distance** from x_0 to y_0 .
- **Definition** $d(x_0, V) := \min_{v \in V} \|x_0 - v\|$ is called the **distance** from x_0 to V .

The translation of a subspace by vector is called a **linear variety**. Let $W := v_c + V$. Then one can also define

- **Definition** $d(x_0, W) := \min_{w \in W} \|x_0 - w\|$ is called the **distance** from x_0 to W .

Fact: If $W = v_c + V$, then $d(x_0, W) = d(x_0 - v_c, V)$.

The somewhat amazing fact is that when V is a hyperplane, the minimization problem defining the distance of a point to the hyperplane has a very simple solution! **This only works for hyperplanes, that is, linear varieties that are translates of hyper subspaces.**

Signed Distance to a Hyperplane

Suppose that $y : \mathbb{R}^n \rightarrow \mathbb{R}$ is given by (13.1) or (13.4) and that $a \neq 0_{n \times 1}$ (recall that one can go back and forth between the two representations by (13.6)). Let $H^0 := \{x \in \mathbb{R}^n \mid y(x) = 0\}$ be the hyperplane defined by y . Then, for all $x_0 \in \mathbb{R}^n$,

$$|y(x_0)| = \|a\| d(x_0, H^0),$$

that is,

$$d(x_0, H^0) = \frac{|y(x_0)|}{\|a\|}$$

For this reason,

$$\frac{y(x)}{\|a\|} \tag{13.7}$$

is called the **signed distance** from x to H^0 .

Example 13.2 Compute the signed distance for a hyperplane defined by

$$P := \{(x_1, x_2) \in \mathbb{R}^2 \mid 1.5x_1 - 2.0x_2 = -4.0\}.$$

Solution We have the hyperplane is defined by $y : \mathbb{R}^2 \rightarrow \mathbb{R}$, where

$$y(x_1, x_2) = 1.5x_1 - 2.0x_2 + 4.0.$$

Hence, $a = [1.5; -2.0]$ and

$$\frac{1}{\|a\|} y(x_1, x_2) = \frac{1.5x_1 - 2.0x_2 + 4.0}{\sqrt{6.25}}$$

is the signed distance from $x = \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$ to the hyperplane P . ■

Example 13.3 Compute the signed distance for a hyperplane defined by

$$P := x_c + \{x \in \mathbb{R}^n \mid x \perp a\}.$$

Solution We know that the hyperplane is defined by $y : \mathbb{R}^n \rightarrow \mathbb{R}$, where

$$y(x) = a \bullet (x - x_c).$$

Hence, $y(x) = \frac{a}{\|a\|} \bullet (x - x_c)$ gives the signed distance. ■

Example 13.4 For a d -dimensional hyperplane that passes through the origin and is defined by the normal vector $[a_1; a_2; \dots, a_d]$, compute the signed distance function.

Solution We know that the hyperplane is defined by $y : \mathbb{R}^d \rightarrow \mathbb{R}$, where

$$y(x) = a \bullet x$$

Hence, $y(x) = \frac{a \bullet x}{\|a\|}$ is the signed distance function. ■

Example 13.5 Prove the signed distance formula, namely,

$$|y(x_0)| = \|a\| d(x_0, H^0) \quad (13.8)$$

Solution The proof uses a few ideas we have not covered in ROB 101, so will only sketch it. So that the hyperplane is well defined, we assume that $a \in \mathbb{R}^n$ is not the zero vector and that

$$H^0 := \{x \in \mathbb{R}^n \mid a \bullet (x - x_c) = 0\} = x_c + N, \text{ where } N := \{x \in \mathbb{R}^n \mid a \bullet x = 0\}. \quad (13.9)$$

Because $a \neq 0_{n \times 1}$ and N is the set of all vectors (points) orthogonal to a , it follows that every vector in \mathbb{R}^n can be written as a multiple of a and a vector in N . In particular, we have that

$$\begin{aligned} x_c &= \alpha a + \bar{x} \text{ for some } \alpha \in \mathbb{R} \text{ and } \bar{x} \in N \\ &\text{and for all } x \in \mathbb{R}^n \\ x &= \beta a + \bar{x} \text{ for some } \beta \in \mathbb{R} \text{ and } \bar{x} \in N. \end{aligned}$$

Using the above relations, we first evaluate

$$\begin{aligned} |y(x)| &= |a \bullet (x - x_c)| \\ &= |a \bullet (\alpha a + \bar{x} - \beta a - \bar{x})| \\ &= |\alpha - \beta| \|a\|^2 \end{aligned}$$

where we used two facts: (i) $a \bullet \bar{x} = 0$ and $a \bullet \bar{x} = 0$ because $\bar{x}, \bar{x} \in N$ and (ii) $a \bullet a = \|a\|^2$ for the Euclidean norm.

Next, we observe that

$$\begin{aligned} d^2(x, H^0) &= \min_{\tilde{x} \in H^0} \|x - \tilde{x}\|^2 \\ &= \min_{\tilde{x} \in N} \|x - (\tilde{x} + x_c)\|^2 \\ &= \min_{\tilde{x} \in N} \|(\beta a + \bar{x}) - (\tilde{x} + \alpha a + \bar{x})\|^2 \\ &= \min_{\tilde{x} \in N} \|(\beta - \alpha)a + (\bar{x} - \tilde{x} - \bar{x})\|^2 \\ &= \min_{\tilde{x} \in N} [\|(\beta - \alpha)a\|^2 + \|\bar{x} - \tilde{x} - \bar{x}\|^2] \\ &= |\beta - \alpha|^2 \|a\|^2 + \min_{\tilde{x} \in N} \|\bar{x} - \tilde{x} - \bar{x}\|^2 \\ &= |\beta - \alpha|^2 \|a\|^2, \end{aligned}$$

where we used the Pythagorean Theorem to arrive at

$$\|(\beta - \alpha)a + (\bar{x} - \tilde{x} - \bar{x})\|^2 = \|(\beta - \alpha)a\|^2 + \|\bar{x} - \tilde{x} - \bar{x}\|^2 =$$

and we noted that

$$\min_{\tilde{x} \in N} \|\bar{x} - \tilde{x} - \bar{x}\|^2 = 0,$$

because $\bar{x} - \bar{x} \in N$. Hence, taking square roots,

$$d(x, H^0) = |\beta - \alpha| \|a\|.$$

Comparing the formulas for $d(x, H^0)$ and $|y(x)|$ we arrive at

$$|y(x)| = \|a\| d(x, H^0).$$

■




Figure 13.6: Raw data for a maximum margin classifier The blue data will be Class 1, labeled with a +1, and the red data will be Class 2, labeled with a -1.

13.3 Max-margin Classifier

This material is from lectures in ROB 101 given by Prof. Maani Ghaffari.

In the next example, we will formulate a (linear) classifier to separate points that belong to two different categories called **class labels**. For example, you can think of this as a model that predicts whether an email is spam or not. Our model is linear and separates the space into two half-spaces, as described in Chap. 13.1. As shown in Chap. 13.1.1, in the 2D plane, a line divides the space into two half-spaces. In general, we will be designing a **separating hyperplane**, that is, a translation of a co-dimension one subspace.

Definition: Consider a labeled data set $\mathcal{D} = \{(x_i, \ell_i)\}_{i=1}^n$, where $x_i \in \mathbb{R}^m$ and $\ell_i = \pm 1$, and suppose that $H := \{x \in \mathbb{R}^m \mid y(x) = 0\} \subset \mathbb{R}^m$ is a hyperplane. Then H **separates the data**, if for each $1 \leq i \leq n$, (i) $y(x_i) \neq 0$ (no data points lie on the hyperplane) and (ii) $\text{sign } y(x_i) = \ell_i$ (the sign of the label ℓ_i determines on which side of the hyperplane the data point x_i lies). Then the **margin** is defined to be

$$\text{margin} := \min_{x_i, \ell_i = +1} d(x_i, H) + \min_{x_j, \ell_j = -1} d(x_j, H) \quad (13.10)$$

■

Remark 5 Intuitively, the larger the margin, the more robust the separation (aka classification) of the data. We can evaluate the distances used to define the margin by $|y(x_i)|/||a||$. Hence,

$$\text{margin} := \min_{x_i, \ell_i=+1} \frac{|y(x_i)|}{||a||} + \min_{x_j, \ell_j=-1} \frac{|y(x_j)|}{||a||} = \frac{1}{||a||} \left(\min_{x_i, \ell_i=+1} |y(x_i)| + \min_{x_j, \ell_j=-1} |y(x_j)| \right). \quad (13.11)$$

We note that if $|y(x_i)| \geq 1$ for all $1 \leq i \leq n$, then

$$\text{margin} \geq \frac{2}{||a||}.$$

And if there is at least one point in each class such that $|y(x_i)| = 1$, then the margin is exactly equal to $\frac{2}{||a||}$. In that case, maximizing the margin is the same as minimizing $||a||$. Vectors such that $|y(x_i)| = ||a|| d(x_i, H)$ are called **support vectors**.

Example 13.6 (Maximum Margin Classifier) Given a labeled data set $\mathcal{D} = \{(x_i, \ell_i)\}_{i=1}^n$, find, if possible, a separating hyperplane that maximizes the margin. This problem appears in machine learning and is called the max-margin classifier. The goal is to build a linear model $y(x) = a \bullet x + a_0$ that can separate the two classes of data with the maximum margin possible.

Solution:

Figure 13.6 shows a synthetic data set (means we generated it in a computer) with red circles and blue crosses. To generate the data, we defined the line $\mathbf{x}_2 = 1.5\mathbf{x}_1 + 0.4$ as our **ground truth** and randomly generated vectors in \mathbb{R}^2 : if they landed above the line, we labeled them with red circles; and if they fell below the line, we labeled them with blue circles. Synthetic data generated in this manner is how one tests a problem formulation and solution in practice!

Problems such as this one are called toy examples. They are simple enough that we can visualize and track the solution to ensure our software works as expected. The green line in Fig. 13.7 is the hyperline $\mathbf{x}_2 = 1.493\mathbf{x}_1 + 0.398$ we computed to separate the two classes of data. We let you know this ahead of time so that you will read on and see how we did it!

Our dataset consists of 2D vectors (called inputs), $x_i \in \mathbb{R}^2$, and class labels (called target or output), $\ell_i \in \{-1, +1\}$. If we have n data points, we write the dataset as

$$\mathcal{D} = \{(x_i, \ell_i)\}_{i=1}^n.$$

From Chap. 13.1, the line (hyperplane) that separates the data can be written as

$$y(x) = a^\top x + a_0 = 0,$$

for $a_0 \in \mathbb{R}$ and $a \in \mathbb{R}^2$. We can also combine the normal vector, a , and the bias, a_0 , into

$$w := \begin{bmatrix} a \\ a_0 \end{bmatrix}^\top \in \mathbb{R}^3,$$

and append a one to the inputs as $\bar{x}_i := [x_i; 1]$. Then the side of the line (hyperplane in general) on which each data point lies can be written as

$$\begin{aligned} w^\top \bar{x}_i &\geq 1 & \text{if } \ell_i = 1, \\ w^\top \bar{x}_i &\leq -1 & \text{if } \ell_i = -1; \end{aligned} \quad (13.12)$$

moreover, with this assignment, $|y(x_i)| \geq 1$ for all $1 \leq i \leq n$, and therefore,

$$d(H, x_i) \geq \frac{1}{||a||}$$

by (13.7). Hence, the margin is maximized by minimizing $||a||$, or equivalently, minimizing $||a||^2 = a^\top a$.

The inequality constraints (13.12) state that the data points for each class must lie on the correct side of the hyperplane for it to be a separating hyperplane! We can combine the two constraints into

$$\ell_i(w^\top \bar{x}_i) \geq 1, \quad 1 \leq i \leq n. \quad (13.13)$$

Finally, the problem can be formulated as the following QP:

$$\min \frac{1}{2} ||a||^2$$

$$w = \begin{bmatrix} a \\ a_0 \end{bmatrix} \in \mathbb{R}^3$$

$$\text{subject to } \ell_i(w^\top \bar{x}_i) \geq 1, \quad i = 1, \dots, n$$




Figure 13.7: The two classes are now separated by a hyperplane that provides the maximum margin, that is, the gap between the data points and the hyperplane. The code given below provides the optimal parameters $w^* = [-12.88, 8.42, -2.76]$, and thus $a^* = [-12.88, 8.42]$, $a_0^* = -2.76$, and the margin is $2/\|a^*\| = 0.31$. Moreover, any vectors that lie on either of the two dotted lines are the support vectors. There must be at least one vector in each class that lies on the lines, for if not, the line can be moved to increase the margin.

After solving the problem, suppose $w^* = [a^*; a_0^*]$ is the optimal solution. We can predict the class label of a new input (called a query point), x_{data} , by simply checking which side of the hyperplane it lies on

$$\text{class label is } \begin{cases} (+1) \text{ blue circles} & a^* \bullet x_{\text{data}} + a_0^* > 0 \\ (-1) \text{ red circles} & a^* \bullet x_{\text{data}} + a_0^* < 0 \end{cases} .$$

The results are shown in Fig. 13.7.

Highlights of the Max Margin Classifier for Two Classes

- The process starts with a labeled data set $\mathcal{D} = \{(x_i, \ell_i)\}_{i=1}^n$. Here we assume that $\ell_i \in \{-1, +1\}$ and $x_i \in \mathbb{R}^m$.
- The separating hyperplane is parameterized by $y(x) = a^\top x + a_0 = 0$, for $a_0 \in \mathbb{R}$ and $a \in \mathbb{R}^m$. The hyperplane is $H := \{x \in \mathbb{R}^m \mid y(x) = 0\}$
- We write the constraints representing a point x_i belonging to one of the two classes $\{-1, +1\}$ by

$$\begin{aligned} w^\top \bar{x}_i &\geq 1 & \text{if } \ell_i = 1, \\ w^\top \bar{x}_i &\leq -1 & \text{if } \ell_i = -1, \end{aligned} \iff \ell_i w^\top \bar{x}_i \geq 1, \quad 1 \leq i \leq n,$$

where

$$w = \begin{bmatrix} a \\ a_0 \end{bmatrix}.$$

- With the constraints written as above, the margin is greater than or equal to $\frac{2}{\|a\|}$, where $a \in \mathbb{R}^m$ is the normal vector defining the hyperplane. If there is a point in the data set such that $|y(x_i)| = 1$, then $d(H, x_i) = 1/\|a\|$, and thus our estimate for the margin is tight. The bias term $a_0 \in \mathbb{R}$ provides the offset of the hyperplane so that it does not have to pass through the origin. Hence, to maximize the margin, we minimize $\|a\|$, subject to the classification constraints.
- Putting all of this together leads to a Quadratic Program or QP as presented in Chapter 12.8:

$$\begin{aligned} &\min \frac{1}{2} a^\top a \\ &\begin{bmatrix} \ell_1 \bar{x}_1^\top \\ \vdots \\ \ell_n \bar{x}_n^\top \end{bmatrix} \begin{bmatrix} a \\ a_0 \end{bmatrix} \geq \mathbf{1}_{n \times 1} \end{aligned}$$

- Given a new data point $x_{\text{new}} \in \mathbb{R}^m$, how do we determine its class? We evaluate $y(x_{\text{new}}) = a^\top x + a_0$ and check its sign! If $y(x_{\text{new}}) > 0$, it is in Class 1 and if $y(x_{\text{new}}) < 0$, it is in Class 2.

```

1 # # New Packages for solving QPs
2 # using Pkg
3 # Pkg.add("OSQP")
4 # Pkg.add("Compat")
5 using OSQP
6 using SparseArrays
7
8 # Standard Packages for ROB 101
9 using LinearAlgebra
10 using Random
11 Random.seed!(123456);
12
13 # generate a dataset
14 N = 200
15 k1 = 0; # number of 1
16 k2 = 0; # number of -1
17 X = zeros(2*N, 2) # input matrix
18 ell = zeros(2*N, 1) # target values
19 i = 1;
20 marginDes = .3
21 while minimum([k1 k2]) < N
22     x = rand(1, 2) * 5. .+ .05;

```

```

23 # separating line is  $x_2 = 1.5 x_1 + 0.4$ 
24  $y = (x[2] - 1.5 * x[1] - 0.4) / \text{norm}([-1.5 \ 1]);$  # norm of a = 1
25 # generate target values
26 if (y > marginDes/2.) && k1 < N
27     ell[i] = 1;
28     X[i, :] = x;
29     k1 += 1;
30     i += 1;
31 elseif (y < -marginDes/2.) && k2 < N
32     ell[i] = -1;
33     X[i, :] = x;
34     k2 += 1;
35     i += 1;
36 end
37 end
38
39 # Class +1 IDs
40 class1_id = ell.== 1;
41
42 using Plots
43 gr() # Set the backend to GR
44
45 plot(X[class1_id[:,1], X[class1_id[:,2], seriestype = :scatter,
46     aspectratio=:equal, legend=false)
47 plot!(X[.!class1_id[:,1], X[.!class1_id[:,2], seriestype = :scatter, xlims = (0,6), ylims = (0,6))
48 xlabel!("x1")
49 ylabel!("x2")
50 plot!(fmt = :png)

```

Output See Fig. 13.6.

```

1 # Data for Max Margin in R^2
2
3 # Define problem data
4 Q = (zeros(3,3) + I); Q[3,3]=0
5 q = zeros(3,1);
6 Ain = -([ell ell ell] .* [X ones(size(X,1),1)]);
7 bin = -ones(size(X,1),1);
8 dimX=length(q)
9 Aeq = Array{SparseMatrixCSC,2}(undef,0,dimX) # Empty Matrix
10 beq = Vector{Float64}(undef,0) # Empty Matrix
11 lb = Vector{Float64}(undef,dimX).-Inf # - infinity means no hard lower bound
12 ub = Vector{Float64}(undef,dimX).+Inf # + infinity means no hard upper bound
13
14 wStar = quadProg(Q, q, Ain, bin, Aeq, beq, lb, ub)
15 @show wStar
16 @show margin = 2.0/norm(wStar[1:2])
17 aStar = wStar/norm(wStar[1:2])

```

Output

```

wStar = [-5.3085002098319585, 3.5545545440849087, -1.4143558513794734]
margin = 2.0 / norm(wStar[1:2]) = 0.31305447944860626

```

```

1 # Our lines have to be plotted as  $y = mx + b$  and not
2 # as  $a[1]*x + a[2] * y + a[3];$ 
3 # Hence, we must solve for  $x_2$ 

```

```

4
5 x_line = collect(0:0.1:3.1)
6 y_line = -(aStar[1] * x_line .+ aStar[3]) / aStar[2]
7
8
9 pMMC = plot!(x_line, y_line, lw = 3)
10
11
12 y_marginPlus = -(aStar[1] * x_line .+ aStar[3] .+ margin/2) / aStar[2]
13 y_marginMinus = -(aStar[1] * x_line .+ aStar[3] .- margin/2) / aStar[2]
14 plot!(x_line, y_marginPlus, lw=2, ls=:dot, color=:black)
15 plot!(x_line, y_marginMinus, lw=2, ls=:dot, color=:black)
16 plot!(fmt = :png)
17
18 display(pMMC)

```

Output See Fig. 13.7.

Now that you are warmed up ...

There are excellent tutorial videos available for many aspects of Machine Learning (ML). Here are a few related to SVMs:

- <https://youtu.be/-Z4aojJ-pdg> (Under the hood of SVM)
- https://youtu.be/vMmG_7JcfIc (Intro to the Kernel Trick)
- <https://youtu.be/OKFMZQyDROI> (More advanced view of the Kernel Trick)
- https://youtu.be/bM4_AstaBZo (Math behind SVM)
- <https://youtu.be/kb4apnc2imA> (Multi-class SVM)

13.4 Remarks on Soft Margin Classifiers

This material is from lectures in ROB 101 given by Prof. Maani Ghaffari.

In real life, the data are rarely so nicely separated. There is almost always some overlaps, perhaps due to random errors and outliers, or perhaps because some valid emails look a lot like spam!

In such cases, the best one can do is to seek a hyperplane that roughly minimizes the number of data points that are misclassified. This is done with a “soft margin classifier”, where the hard constraint in (13.13) is replaced with

$$\ell_i(w^\top \bar{x}_i) - \xi_i \geq 1, \quad 1 \leq i \leq n, \quad (13.14)$$

where for $\xi_i > 1$, a data point is allowed to be in the wrong class. To make sure this is the exception rather than the rule, we try to make the vector ξ have small norm. This gives rise to the QP

$$\min_{\xi, w} \frac{1}{2} \xi^\top \xi + \frac{\lambda}{2} a^\top a$$

$$\begin{bmatrix} \ell_1 \bar{x}_1^\top \\ \vdots \\ \ell_n \bar{x}_n^\top \end{bmatrix} \begin{bmatrix} a \\ a_0 \end{bmatrix} \geq \begin{bmatrix} 1 - \xi_1 \\ \vdots \\ 1 - \xi_n \end{bmatrix}, \quad (13.15)$$

where $\lambda > 0$ trades off the separation property versus the soft margin.

```

1 # generate a dataset
2 N = 100 # Desired number of data points in each class
3 n = 2*N # Total number of points
4 k1 = 0; # number of 1

```

```

5 k2 = 0; # number of -1
6 X = zeros(n,2); # input matrix
7 ell = zeros(n,1); # target values
8 i = 1;
9 while minimum([k1 k2]) < N
10     x = rand(1,2) * 10.;
11     # separating line is x2 = 1.5 x1 + 0.4
12     y = (x[2] - 1.5 * x[1] - 0.4)
13     # generate target values
14     if (y > -3.5) && k1 < N
15         ell[i] = 1;
16         X[i,:] = x;
17         k1 += 1;
18         i += 1;
19     elseif (y < 3.5) && k2 < N
20         ell[i] = -1;
21         X[i,:] = x;
22         k2 += 1;
23         i += 1;
24     end
25 end
26
27 # Class +1 IDs
28 class1_id = ell.== 1;
29

```




Figure 13.8: Raw data where there is some overlap between the two classes. This is quite typical in practice.


```

30 using Plots
31 gr() # Set the backend to GR
32
33 plot(X[class1_id[:,1], X[class1_id[:,2], seriestype = :scatter, legend = false)
34 plot!(X[.!class1_id[:,1], X[.!class1_id[:,2], seriestype = :scatter)
35 xlabel!("x1")
36 ylabel!("x2")

```

Output See Fig. 13.8.

```

1 # Define problem data using the native Julia form for the QP solver
2 m = 3;
3 lambda = 1; # tunable parameter (called hyperparameter because it's not like w the
   parameter of our model)
4 P = sparse([lambda*(zeros(m,m) + I) zeros(m,n); zeros(n,m) (zeros(n,n) + I)]);
5 q = zeros(n+m, 1);
6 A = [-sparse([ell ell ell] .* [X ones(n,1)]) -(zeros(n,n) + I); zeros(n,m) (zeros(n,n) + I)]
7 l = [zeros(n,1) .- Inf; zeros(n,1)];
8 u = [-ones(n,1); zeros(n,1) .+ Inf];
9
10 # Crate OSQP object
11 prob = OSQP.Model()
12
13 # Setup workspace and change alpha parameter
14 OSQP.setup!(prob; P=P, q=q[:,], A=A, l=l[:,], u=u[:,])

```




Figure 13.9: A “separating” hyperplane that allows a few misclassifications.

```

15
16 # Solve problem
17 results = OSQP.solve!(prob);
18
19 w_line = -results.x ./ results.x[2];
20
21 x_line = collect(0:0.1:10)
22 y_line = w_line[1] * x_line .+ w_line[3]
23 plot!(x_line, y_line, lw=3)

```

Output See Fig. 13.9.

```

1 zeta = results.x[4:end]
2 @show maximum(zeta)
3 @show minimum(zeta)
4 println(" ")
5 println("These are misclassified data points or they are outliers in your data set.")
6 indicesBigZeta=findall(x->x>1,zeta) # These are misclassified

```

Output

```

maximum(zeta) = 3.2141869873099136
minimum(zeta) = 2.3479433444366696e-8

```

These are misclassified data points or they are outliers in your data set.
13-element Vector{Int64}:

```

18
43
56
66
87
94
131
144
150
194
195
196
198

```

We next look at a problem where it seems impossible to separate the data with a hyperplane, such as shown in Fig. 13.10. The trick is to add more features to the data by adding nonlinear terms. For example, instead of working in \mathbb{R}^2 with $[x_1 \ x_2]$ as features, we could work in \mathbb{R}^5 with the feature vector being

$$\begin{bmatrix} x_1 \\ x_2 \\ 1 \\ (x_1)^2 \\ x_1x_2 \\ (x_2)^2 \end{bmatrix}.$$

This gives more possibilities for the data to be separated. The above choice is motivated by the blue data seemingly belonging to a disc. We are not obliged, however, to use monomials. In Project 2, we learned about radial basis functions, or RBFs for short. As shown in Fig. 13.11, this provides a lot of flexibility for separating the data into two classes.

```

1 # generate a dataset
2 N = 100 # Desired number of data points in each class
3 n = 2*N # Total number of points
4 k1 = 0; # number of 1
5 k2 = 0; # number of -1

```

```

6 X = zeros(n,2); # input matrix
7 ell = zeros(n,1); # target values
8 i = 1;
9 while minimum([k1 k2]) < N
10     x = rand(1,2) * 10.;
11     y = (x[1]-5)^2 + (x[2]-5)^2;
12     # generate target values
13     if (y < 5.5) && k1 < N
14         ell[i] = 1;
15         X[i,:] = x;
16         k1 += 1;
17         i += 1;
18     elseif (y > 4) && k2 < N
19         ell[i] = -1;
20         X[i,:] = x;
21         k2 += 1;
22         i += 1;
23     end
24 end
25
26 # Class +1 IDs
27 class1_id = ell .== 1;
28
29 using Plots
30 gr() # Set the backend to GR

```




Figure 13.10: Raw data where it seems impossible to separate the data with a hyperplane.

```

31
32 plot(X[class1_id[:,1], X[class1_id[:,2], seriestype = :scatter, legend=false)
33 plot!(X[.!class1_id[:,1], X[.!class1_id[:,2], seriestype = :scatter)
34 plot!(aspectratio=:equal)

```

Output See Fig. 13.10.

```

1 # Functions from Project 2
2
3 # Radial basis function
4 s = 1;
5 rbf(x, z, s) = exp.(-norm(x-z)^2 / (2*s^2));
6
7 function calc_phi_row(x, z, s)
8     NumBasisElements = size(z,1) + 2
9     # plus two above because we also include a x1 and x2
10    phi_row = zeros(1, NumBasisElements)
11    phi_row[1:2] = [x[1] x[2]]
12    for i in 3:NumBasisElements
13        phi_row[i] = rbf(x, z[i-2, :], s)
14    end
15    return phi_row
16 end

```




Figure 13.11: A Gaussian soft-margin classifier, where radial basis functions have been used to lift the data to \mathbb{R}^{n+2} , where n is the number of data points. When we evaluate the sign of the classifier on our data in \mathbb{R}^2 , we obtain an approximation of a circle. In \mathbb{R}^{202} , there is a separating hyperplane.

```

17
18 function regressor_matrix(X, centers, s)
19     ### BEGIN SOLUTION
20     N = size(X,1)
21     M = size(centers,1)
22     Phi = Array{Float64, 2}(undef, N, M+2)
23     for i = 1:N
24         Phi[i, :] = calc_phi_row(X[i, :], centers, s)
25     end
26     return Phi
27     ### END SOLUTION
28 end

1 # Define problem data
2 m = n+2;
3 lambda = 0.01; # tunable parameter (called hyperparameter because it's not like w the
   parameter of our model)
4 P = sparse([lambda*(zeros(m,m) + I) zeros(m,n); zeros(n,m) (zeros(n,n) + I)]);
5 q = zeros(n+m,1);
6 Phi = regressor_matrix(X, X, s);
7 A = sparse([-\ell.*Phi -(zeros(n,n) + I); zeros(n,m) (zeros(n,n) + I)]);
8 l = [zeros(n,1) .- Inf; zeros(n,1)];
9 u = [-ones(n,1); zeros(n,1) .+ Inf];
10
11 # Create OSQP object
12 prob = OSQP.Model()
13
14 # Setup workspace and change alpha parameter
15 OSQP.setup!(prob; P=P, q=q[:], A=A, l=l[:], u=u[:])
16
17 # Solve problem
18 results = OSQP.solve!(prob);

1 # create test data
2 x1 = 0:0.1:10;
3 x2 = 0:0.1:10;
4 X1 = zeros(length(x2), length(x1));
5 X2 = zeros(length(x2), length(x1));
6 for j=1:length(x1)
7     for i=1:length(x2)
8         X1[i,j] = x1[j]
9         X2[i,j] = x2[i]
10    end
11 end
12 X_test = [X1[:] X2[:]];
13 # get model weights
14 alpha = results.x[1:m, :];
15 # query
16 Phi_test = regressor_matrix(X_test, X, s);
17 Y_test = Phi_test * alpha;
18
19 plot(X[class1_id[:,1]], X[class1_id[:,2]], seriestype = :scatter)
20 plot!(X[.!class1_id[:,1]], X[.!class1_id[:,2]], seriestype = :scatter)
21
22 Z = sign.(reshape(Y_test, (length(x2), length(x1)))); # +1 or -1

```

```

23 # plot the margins
24 contour!(x1, x2, Z, lw=3, color=:red, legend=false) # The contour line separates class -1 from
    class +1
25 plot!(aspectratio=:equal)

```

Output See Fig. 13.11.

Remark 6 What is the classifier for the data? It is

$$\alpha = \text{results.x}[1:m, :]$$

$$\text{Classifier}(x) = (\text{calc_phi_row}(x, X, s) * \alpha)[1]$$

The data from \mathbb{R}^2 have been lifted to $\mathbb{R}^{n+2} = \mathbb{R}^{202}$ via

$$\text{Classifier}(x) = \text{Classifier}(x_1, x_2) = \begin{bmatrix} x_1 & x_2 & e^{-\frac{\|x-z_1\|^2}{2}} & \dots & e^{-\frac{\|x-z_n\|^2}{2}} \end{bmatrix} \begin{bmatrix} \alpha_a^* \\ \alpha_b^* \\ \alpha_1^* \\ \vdots \\ \alpha_n^* \end{bmatrix},$$

where z_1, z_2, \dots, z_n are the two-dimensional data points in Fig. 13.10. The code block below shows that there are no misclassified points!

```

1 alpha = results.x[1:m, :]
2 Classifier(x) = (calc_phi_row(x, X, s) * alpha)[1]
3
4 for i in 1:n
5     test = Classifier(X[i, :])
6     if sign(test) != ell[i]
7         @show [i ell[i] test] # misclassified data points
8     end
9 end

```

Output Nothing! There are no misclassified data points!

13.5 Orthogonal Projection

We extend the importance of the dot product (aka, inner product) by showing its fundamental role in least squares problems.




Figure 13.12: A vector (in red) is orthogonally projected (in black) onto a subspace (in cyan). The error vector (in solid orange) is orthogonal to the plane. This characterizes the orthogonal projection process. The vector in dashed orange is the error vector drawn to highlight that the error forms a right angle with the projection of the vector.

13.5.1 Orthogonal Projection for Subspaces

Review

Consider the vector space \mathbb{R}^n , which we view as the set of all $n \times 1$ column vectors of real numbers. Let $v, w \in \mathbb{R}^n$ and let $V \subset \mathbb{R}^n$ be a subspace.

- $v \bullet w := v^\top w$.
- $w \bullet v = v \bullet w$.
- $v \perp w \iff v \bullet w = 0$.
- $v \perp w \implies \|v + w\|^2 = \|v\|^2 + \|w\|^2$ (Pythagorean Theorem).
- Let $\{u_1, u_2, \dots, u_m\}$ be a basis for V . Then Gram-Schmidt produces an **orthogonal basis** that also satisfies, for all $1 \leq k \leq m$,

$$\text{span}\{u_1, u_2, \dots, u_k\} = \text{span}\{v_1, v_2, \dots, v_k\}.$$

Moreover, by the simple step of adding normalization to Gram-Schmidt, we can assume that $\{v_1, v_2, \dots, v_m\}$ is an **orthonormal basis** for V .

Projection Theorem: The Super Tool that Solves all Least Squares Problems

Let V be a subspace of \mathbb{R}^n and let x_0 be an arbitrary point in \mathbb{R}^n . Then there exists a unique vector $x^* \in V$ such that

$$\|x_0 - x^*\| = \min_{x \in V} \|x_0 - x\|;$$

as before, we denote this vector by $x^* = \arg \min_{x \in V} \|x_0 - x\|$ or by $x^* = \arg \min_{x \in V} \|x_0 - x\|^2$. Moreover, the solution to the least squared error problem is uniquely characterised by

$$x^* = \arg \min_{x \in V} \|x_0 - x\|^2 \iff (x_0 - x^*) \perp V \text{ and } x^* \in V. \quad (13.16)$$

The vector $x_0 - x^*$ is called the **error vector**. The vector x^* is called the **orthogonal projection of x_0 onto V** precisely because the error vector is orthogonal to V . Recalling the Pythagorean Theorem, we have that

$$\|x_0 - x^*\|^2 + \|x^*\|^2 = \|x_0\|^2;$$

once again emphasizing that x^* , $x_0 - x^*$, and x_0 form a “generalized right triangle”.

You already know one way to compute x^* from the Projection Theorem! Really? Yes, Gram Schmidt. If $x_0 \in V$, the solution to the problem is trivial, namely $x^* = x_0$, because then the error is zero, which is as small as it gets. Hence, suppose $x_0 \notin V$ and let $\{u_1, u_2, \dots, u_m\}$ be a basis for V . Then we leave it to you to show that

$$x_0 \notin V \iff x_0 \notin \text{span}\{u_1, u_2, \dots, u_m\} \iff \{u_1, u_2, \dots, u_m, x_0\} \text{ is linearly independent.}$$

We apply Gram-Schmidt² to the set $\{u_1, u_2, \dots, u_m, x_0\}$. The last step gives that

$$v_{m+1} = x_0 - \sum_{k=1}^m \frac{x_0 \bullet v_k}{v_k \bullet v_k} v_k,$$

and moreover, we know that

$$v_{m+1} \perp \text{span}\{v_1, v_2, \dots, v_m\} = V.$$

Hence, by the Projection Theorem,

$$x^* = \sum_{k=1}^m \frac{x_0 \bullet v_k}{v_k \bullet v_k} v_k, \quad (13.17)$$

because $x_0 - x^* = v_{m+1}$ and $v_{m+1} \perp V$.

Remark: Once we know that (13.17) is true, we can simply apply Gram-Schmidt to any basis of V to produce an orthogonal basis and then apply (13.17). If we produce an **orthonormal basis**, then we know that $v_k \bullet v_k = 1$ and the formula simplifies to

$$x^* = \sum_{k=1}^m (x_0 \bullet v_k) v_k = \sum_{k=1}^m \langle x_0, v_k \rangle v_k, \quad (13.18)$$

where we have recalled our alternative notation for an inner product.

A second way to compute the solution follows from (13.16) and leads to the **Normal Equations**. Once again, let $\{u_1, u_2, \dots, u_m\}$ be any basis for V . Because we know that $x^* \in V$, we can pose

$$x^* = \alpha_1 u_1 + \alpha_2 u_2 + \dots + \alpha_m u_m \quad (13.19)$$

as a linear combination of basis vectors for V and seek the conditions on the coefficients $\alpha_1, \alpha_2, \dots, \alpha_m$ so that

$$x_0 - x^* \perp V.$$

²We do not assume normalization, but you can also do that.

You can quickly convince yourself that

$$x_0 - x^* \perp V \iff x_0 - x^* \perp u_k, 1 \leq k \leq m.$$

The above constitutes m -equations, one for each k , and leads to the famous Normal Equations,

$$\underbrace{\begin{bmatrix} u_1 \bullet u_1 & u_1 \bullet u_2 & \cdots & u_1 \bullet u_m \\ u_2 \bullet u_1 & u_2 \bullet u_2 & \cdots & u_2 \bullet u_m \\ \vdots & \vdots & \ddots & \vdots \\ u_m \bullet u_1 & u_m \bullet u_2 & \cdots & u_m \bullet u_m \end{bmatrix}}_G \underbrace{\begin{bmatrix} \alpha_1 \\ \alpha_2 \\ \vdots \\ \alpha_m \end{bmatrix}}_\alpha = \underbrace{\begin{bmatrix} u_1 \bullet x_0 \\ u_2 \bullet x_0 \\ \vdots \\ u_m \bullet x_0 \end{bmatrix}}_\beta. \quad (13.20)$$

The matrix G is called the **Gram matrix** and is invertible if, and only if, the set $\{u_1, u_2, \dots, u_m\}$ is linearly independent. We note the the ij -entry of it is

$$G_{ij} = u_i \bullet u_j = u_i^\top u_j.$$

We'll let you work out that if you take a basis for V that is orthogonal, then G is a diagonal matrix, and if you take an orthonormal basis for V , then G is the identity matrix!

We summarize the various solutions in the following:

Computing the Solution Given by the Projection Theorem

Let V a subspace of \mathbb{R}^n and $x_0 \in \mathbb{R}^n$ be given. Then $x^* = \arg \min_{x \in V} \|x_0 - x\|^2$, the orthogonal projection of x_0 onto V , can be computed by

- $x^* = \sum_{k=1}^m (x_0 \bullet v_k) v_k$ if $\{v_1, \dots, v_m\}$ is an **orthonormal basis** for V ;
- $x^* = \sum_{k=1}^m \frac{x_0 \bullet v_k}{v_k \bullet v_k} v_k$ if $\{v_1, \dots, v_m\}$ is an **orthogonal basis** for V ;
- $x^* = \alpha_1 u_1 + \alpha_2 u_2 + \dots + \alpha_m u_m$, where $G\alpha = \beta$ are the Normal Equations given in (13.20), if $\{u_1, \dots, u_m\}$ is **any basis** for V .

You instructors use all of these forms of the solution at various times when solving problems.

Example 13.7 We'll warm up on a simple example. Consider a subspace given by $V = \text{span}\{u_1, u_2\}$, where

$$u_1 = \begin{bmatrix} 1.0 \\ 1.0 \\ 0.0 \end{bmatrix}, u_2 = \begin{bmatrix} 2.5 \\ 0.0 \\ 1.0 \end{bmatrix}.$$

Compute the orthogonal projection of $x_0 = \begin{bmatrix} 4.0 \\ 4.0 \\ 4.0 \end{bmatrix}$ onto V . Moreover, compute

$$x^* = \arg \min_{x \in V} \|x_0 - x\|^2$$

in at least two different ways.

Solution A: We apply the normal equations

$$G = \begin{bmatrix} u_1^\top u_1 & u_1^\top u_2 \\ u_2^\top u_1 & u_2^\top u_2 \end{bmatrix} = \begin{bmatrix} 2.00 & 2.50 \\ 2.50 & 7.25 \end{bmatrix}$$

$$\beta = \begin{bmatrix} u_1^\top x_0 \\ u_2^\top x_0 \end{bmatrix} = \begin{bmatrix} 8.00 \\ 14.00 \end{bmatrix}$$

$$G\alpha = \beta \implies \alpha = \begin{bmatrix} 2.79 \\ 0.97 \end{bmatrix}$$

$$x^* = \alpha_1 u_1 + \alpha_2 u_2 = \begin{bmatrix} 5.21 \\ 2.79 \\ 0.97 \end{bmatrix}.$$

The results are illustrated in Fig. 13.12.

Solution B: We find an orthonormal basis for V and apply (13.18). We use Gram-Schmidt with normalization and find that $V = \text{span}\{v_1, v_2\}$, for

$$v_1 = \begin{bmatrix} 0.707 \\ 0.707 \\ 0.000 \end{bmatrix} \quad \text{and} \quad v_2 = \begin{bmatrix} 0.615 \\ -0.615 \\ 0.492 \end{bmatrix}.$$

Hence,

$$x^* = (v_1^\top x_0) v_1 + (v_2^\top x_0) v_2 = 5.657v_1 + 1.969v_2 = \begin{bmatrix} 5.212 \\ 2.788 \\ 0.970 \end{bmatrix}.$$

Solution C: Finally, we use an orthogonal basis for V and apply (13.17). For our orthogonal basis, we apply Gram-Schmidt without normalization and obtain $V = \text{span}\{v_1, v_2\}$, for

$$v_1 = \begin{bmatrix} 1.0 \\ 1.0 \\ 0.0 \end{bmatrix} \quad \text{and} \quad v_2 = \begin{bmatrix} 1.25 \\ -1.25 \\ 1.00 \end{bmatrix}.$$

Hence,

$$x^* = \frac{v_1^\top x_0}{v_1^\top v_1} v_1 + \frac{v_2^\top x_0}{v_2^\top v_2} v_2 = 4.0v_1 + 0.970v_2 = \begin{bmatrix} 5.212 \\ 2.788 \\ 0.970 \end{bmatrix}.$$

■

In this next example, we apply the Normal Equations to our very first least squares problem in (8.8)! The example is essentially a proof showing how to derive our original result from the Projection Theorem. **Trigger Warning: This is not for the faint of heart. Your instructors did not learn this as undergraduates, much less, first-year undergraduates!** If you skip the example, we highly recommend the summary that follows it.

Example 13.8 Consider a system of linear equations $Ax = b$, where A is $n \times m$ and its columns are linearly independent. Define

$$V := \text{col span}\{A\}.$$

Relate the following two least squares problems

- $x^* = \arg \min_{x \in \mathbb{R}^m} \|Ax - b\|^2$
- $v^* = \arg \min_{v \in V} \|b - v\|^2,$

where we renamed the solution of the second optimization problem as v^* to avoid confusion later on. (Yikes! It must not be easy.)

Solution The first least squares problem is well known to us from (8.8), which we repeat here for clarity

$$x^* = \arg \min_{x \in \mathbb{R}^m} \|Ax - b\|^2 \iff A^\top Ax^* = A^\top b,$$

which we've always interpreted as the least squared error solution to over determined equations. Moreover, this provided the basis for our work on regression, which was, we recall, pretty awesome.

The second least squares problem is still kind of a mystery to us. If we believe what we were told about its solution, then v^* is the orthogonal projection of b onto the column span of the matrix A . What could that possibly mean? Well, let's find out!

We write A out in terms of its columns, $A = [A_1 \ A_2 \ \dots \ A_m]$, so that

$$\text{col span}\{A\} = \text{span}\{A_1, A_2, \dots, A_m\}.$$

When we compute the Gram matrix, we recall that $G_{ij} = A_i \bullet A_j = A_i^\top A_j$, which is precisely the ij -entry of $A^\top A$, and thus

$$G = A^\top A,$$

an amazing coincidence! We'll let you work out a few examples to see that this is true or we'll let you work out a proof! Moreover, when we compute the i -th entry of β we obtain $\beta_i = A_i \bullet b = A_i^\top b$, so that

$$\beta = A^\top b,$$

another amazing coincidence! (Or, perhaps not!). Finally, we note (aka, "let you work out") that

$$\sum_{k=1}^m \alpha_k A_k = A\alpha,$$

which should either be setting off alarms in your head because this many coincidences should never happen.....except for a reason!

The two least squares problems are really one and the same problem.

To see this, we summarize what we have

- $v^* = \arg \min_{v \in \text{col span}\{A\}} \|b - v\|^2 \iff (A^\top A\alpha = A^\top b \text{ and } v^* = A\alpha).$
- $x^* = \arg \min_{x \in \mathbb{R}^m} \|Ax - b\|^2 \iff A^\top Ax^* = A^\top b.$
- Hence, $\alpha = x^*$, and $v^* = Ax^*$ is the orthogonal projection of b onto the column span of A .
- By projecting b onto the column span of A , we have that

$$v^* \in \text{span}\{A_1, A_2, \dots, A_m\},$$

and hence $Ax = v^*$ has a solution. Kind of clever, isn't it!

- The fact that all of that is being accomplished by the simple equation $A^\top Ax^* = A^\top b$ is one of the **Wonders of Linear Algebra**. It's a really beautiful subject and we hope you will want to learn more about it. There is so much more to it than what we have covered in the main parts of this book.

■

Remark: This is a heavy result, so it will take you some time to wrap your head around it. The main message is that the Projection Theorem and the Normal Equations are your main tools when you approach new least squares problems. They have extensions to settings that you cannot even imagine right now. But they are always there, providing theoretical and computational support for solving least squares problems of so many kinds.

It is clear why we did not broach this result in our main treatment of least squares problems. We would have sent you running

and screaming for any course but ROB 101!

Summary of the Previous Example

Suppose that $\{u_1, u_2, \dots, u_m\}$ is a basis for a subspace $V \subset \mathbb{R}^n$ and $x_0 \in \mathbb{R}^n$. Form a matrix U with the basis vectors as its columns, that is,

$$U = \begin{bmatrix} u_1 & u_2 & \dots & u_m \end{bmatrix}.$$

Then the solution to $x^* = \arg \min_{x \in V} \|x_0 - x\|^2$ is given by

$$U^T U \alpha^* = U^T x_0, \quad x^* = U \alpha^*.$$

13.5.2 Orthogonal Projection onto Linear Varieties (translations of subspaces)

Let $V \subset \mathbb{R}^n$ be a subspace (of any dimension) and $v_c \in \mathbb{R}^n$ be a point. We define the linear variety, $W := v_c + V$, as the translation of the subspace V by the vector v_c . Though it is not very common, one can consider the **orthogonal projection** of a vector $x_0 \in \mathbb{R}^n$ onto the linear variety W . The key idea is to once again pose a best approximation problem and to consider the properties that define its solution, by properly interpreting the Projection Theorem.

Small Extension of the Projection Theorem

Consider a linear variety $W := v_c + V$, where $V \subset \mathbb{R}^n$ is a subspace (of any dimension) and $v_c \in \mathbb{R}^n$ is a point. For $x_0 \in \mathbb{R}^n$ arbitrary, the following are equivalent:

- (a) $w^* = \arg \min_{w \in W} \|x_0 - w\|^2$.
- (b) $w^* = v^* + v_c$, where $v^* = \arg \min_{v \in V} \|(x_0 - v_c) - v\|^2$.
- (c) $w^* = v^* + v_c$, where $v^* \in V$ and $((x_0 - v_c) - v^*) \perp V$.
- (d) $w^* \in W$ and $((x_0 - v_c) - (w^* - v_c)) \perp V$.
- (e) $w^* \in W$ and $(x_0 - w^*) \perp V$.

The last condition, (e), emphasizes that the error term, $x_0 - w^*$, is orthogonal to V . The second condition, (b), shows how to compute w^* : orthogonally project $x_0 - v_c$ onto V , and then add v_c to the answer.

We gave (a) through (e) in the order we would use them in a proof, were we to give it! The order we chose should help you to see how each fact is a small variation of the previous one, while going straight from (a) to (e) would be rather daunting.

Example 13.9 Compute the orthogonal projection of $x_0 = \begin{bmatrix} 4.0 \\ 4.0 \\ 4.0 \end{bmatrix}$ onto $W := v_c + V$, where $V = \text{span}\{u_1, u_2\}$,

$$u_1 = \begin{bmatrix} 1.0 \\ 1.0 \\ 0.0 \end{bmatrix}, \quad u_2 = \begin{bmatrix} 2.5 \\ 0.0 \\ 1.0 \end{bmatrix}$$

and $v_c = \begin{bmatrix} 1.0 \\ 2.0 \\ 3.0 \end{bmatrix}$. Moreover, compute the norm of the error term, $x_0 - w^*$, which we now know is the distance of x_0 from the linear variety W .

Solution: Our strategy is we form $\bar{x}_0 := x_0 - v_c = \begin{bmatrix} 3.0 \\ 2.0 \\ 1.0 \end{bmatrix}$ and compute v^* , its orthogonal projection onto V . We then have $w^* = v^* + v_c$ is the orthogonal projection of x_0 onto $W = V + v_c$.

Using our work from Example 13.7, Solution B, we have that $V = \text{span}\{v_1, v_2\}$, for

$$v_1 = \begin{bmatrix} 0.707 \\ 0.707 \\ 0.000 \end{bmatrix} \quad \text{and} \quad v_2 = \begin{bmatrix} 0.615 \\ -0.615 \\ 0.492 \end{bmatrix}.$$

Hence,

$$v^* = (v_1^\top \bar{x}_0) v_1 + (v_2^\top \bar{x}_0) v_2 = -1.5v_1 - 0.182v_2 = \begin{bmatrix} -1.727 \\ -1.273 \\ -0.182 \end{bmatrix}.$$

Hence,

$$w^* = v^* + v_c = \begin{bmatrix} -0.727 \\ 0.727 \\ 2.818 \end{bmatrix}.$$

We next compute

$$d(x_0, W) := \min_{w \in W} \|x_0 - w\| = \|x_0 - w^*\| = 2.08893$$

■

Example 13.10 As a natural continuation of Example 13.9, we note that $W \subset \mathbb{R}^3$ is a hyperplane. Compute the signed distance of x_0 from W .

Solution: We need to write the hyperplane as the zero set of a function $y : \mathbb{R}^3 \rightarrow \mathbb{R}$, where

$$y(x) = a \bullet (x - v_c),$$

and $a \in \mathbb{R}^3$ has norm one. Once again, appealing to Gram-Schmidt, we have that

$$a = \begin{bmatrix} -0.348 \\ 0.348 \\ 0.870 \end{bmatrix}.$$

Doing the required computation yields that the signed distance is

$$y(x_0) = -2.08893.$$

Comparing to our result in Example 13.9, we see that

$$y(x_0) = -d(x_0, W),$$

in other words, the terminology “signed distance” is justified!

■

Appendix A

To Learn on Your Own (if you want to): Cool and Important Things We Omitted From our Linear Algebra Introduction

Learning Objectives

- Introduce material that is commonly included in a second or third year Linear Algebra Course
- Provide a resource for use after you leave ROB 101.

Outcomes

- Complex numbers obey the same rules of arithmetic as the real numbers, if you really understand the real numbers!
- Eigenvalues and eigenvectors of square matrices
- Symmetric matrices have real eigenvalues and admit orthonormal eigenvectors
- Positive definite matrices allow one to generalize the Euclidean norm
- The Singular Value Decomposition (SVD) allows one to quantify the degree to which vectors are linearly independent. This is super useful in engineering practice.
- Matrices are good for other things than representing systems of equations: they also allow one to transform vectors in interesting ways, giving rise to the concept of *linear transformations*.
- Many more facts about basis vectors.

A.1 Complex Numbers and Complex Vectors

Here are some video resources that you may enjoy consulting:

- <https://youtu.be/T647CGsuOVU>
- <https://youtu.be/2HrSG0fdxLY>
- <https://youtu.be/N9QOLrffcKNc>
- <https://youtu.be/DThAoT3q2V4>
- <https://youtu.be/65wYmy8Pf-Y>

The story of complex numbers begins with the quadratic equation $x^2 + 1 = 0$, which has no real solutions! After much soul searching, the mathematics community finally embraced the notion of an **imaginary quantity** \mathfrak{i} defined by

$$(\mathfrak{i})^2 := -1. \quad (\text{A.1})$$

More commonly, we write this as

$$\mathfrak{i} = \sqrt{-1}. \quad (\text{A.2})$$

The set of **complex numbers** is then defined as

$$\mathbb{C} := \{x + \mathfrak{i}y \mid x \in \mathbb{R}, y \in \mathbb{R}\}. \quad (\text{A.3})$$

If $z = x + \mathfrak{i}y \in \mathbb{C}$, then we define

$$\begin{aligned} x &:= \text{real}(z) \text{ the } \mathbf{real\ part} \text{ of } z \\ y &:= \text{imag}(z) \text{ the } \mathbf{imaginary\ part} \text{ of } z. \end{aligned} \quad (\text{A.4})$$

We note that both x and y are real numbers. Complex numbers of the form $0 + \mathfrak{i}y$ are called **imaginary numbers**. We view a **real number** $x \in \mathbb{R}$ as being a complex number of the form $x + \mathfrak{i}0$. In other words, we view $\mathbb{R} \subset \mathbb{C}$. In addition, we define the **complex conjugate** of $z = x + \mathfrak{i}y$ to be

$$z^* := x - \mathfrak{i}y, \quad (\text{A.5})$$

that is, $\text{imag}(z^*) = -\text{imag}(z)$, while $\text{real}(z^*) = \text{real}(z)$.

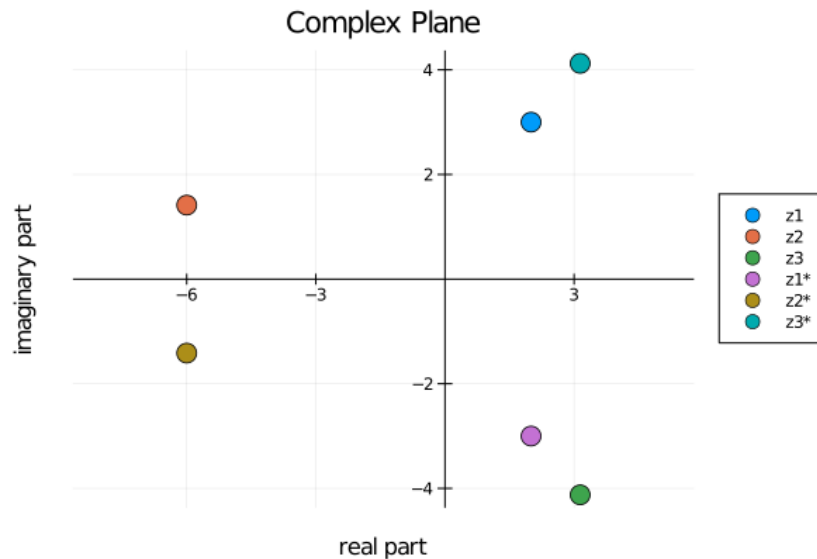


Figure A.1: The complex plane has x -axis given by the real part of a complex number and y -axis given by the imaginary part of a complex number. Here, we plot z_1, z_2, z_3 from Example A.1 and their complex conjugates.

Example A.1 For the following are complex numbers, compute their real and imaginary parts as well as their complex conjugates. Also, plot them in the complex plane.

$$\begin{aligned} z_1 &= 2 + \mathfrak{i} 3 \\ z_2 &= -6 + \mathfrak{i} \sqrt{2} \\ z_3 &= \pi - \mathfrak{i} \sqrt{17}. \end{aligned}$$

Solution

$$\begin{aligned} \text{real}(z_1) &= 2 & \text{imag}(z_1) &= 3 & z_1^* &= 2 - \mathfrak{i} 3 \\ \text{real}(z_2) &= -6 & \text{imag}(z_2) &= \sqrt{2} & z_2^* &= -6 - \mathfrak{i} \sqrt{2} \\ \text{real}(z_3) &= 2\pi & \text{imag}(z_3) &= -\sqrt{17} & z_3^* &= \pi + \mathfrak{i} \sqrt{17}. \end{aligned}$$

All of these values are plotted in Fig. A.1. ■

A.1.1 Arithmetic of Complex Numbers: Enough to Get You By

We'll define all of the major arithmetic operations. Just like operations with vectors and matrices, however, it's much more fun to do the calculations in Julia than by hand!

The **addition of two complex numbers** is defined by adding their respective real and imaginary parts,

$$(x_1 + \mathfrak{i} y_1) + (x_2 + \mathfrak{i} y_2) := (x_1 + x_2) + \mathfrak{i} (y_1 + y_2). \quad (\text{A.6})$$

This is very similar to how we add two vectors

$$\begin{bmatrix} x_1 \\ y_1 \end{bmatrix} + \begin{bmatrix} x_2 \\ y_2 \end{bmatrix} := \begin{bmatrix} x_1 + x_2 \\ y_1 + y_2 \end{bmatrix}$$

by adding their respective components.

The **multiplication** of two complex numbers is defined by

$$(x_1 + \mathfrak{i} y_1) \cdot (x_2 + \mathfrak{i} y_2) := (x_1 x_2 - y_1 y_2) + \mathfrak{i} (x_1 y_2 + y_1 x_2). \quad (\text{A.7})$$

This formula comes from applying basic algebra to the “symbolic expression”

$$(a_1 + b_1)(a_2 + b_2) = a_1 a_2 + b_1 b_2 + a_1 b_2 + b_1 a_2$$

and then substituting in

$$\begin{aligned} a_1 &:= x_1 \\ a_2 &:= x_2 \\ b_1 &:= \mathfrak{i} y_1 \\ b_2 &:= \mathfrak{i} y_2. \end{aligned}$$

The term $b_1 b_2 = (\mathfrak{i} y_1) \cdot (\mathfrak{i} y_2) = (\mathfrak{i})^2 y_1 y_2 = (-1) y_1 y_2$, which explains how the minus sign appears!

Let's note that if we multiply a complex number by its complex conjugate, then we obtain a real number. Indeed,

$$z \cdot z^* = (x + \mathfrak{i} y) \cdot (x - \mathfrak{i} y) = ((x)(x) - (y)(-y)) + \mathfrak{i} ((x)(-y) + (y)(x)) = x^2 + y^2. \quad (\text{A.8})$$

The **magnitude of a complex number** $z = x + \mathfrak{i} y$ is denoted by $|z|$ and is defined by

$$|z| := \sqrt{x^2 + y^2}, \quad (\text{A.9})$$

or equivalently, by

$$|z| := \sqrt{z \cdot z^*}. \quad (\text{A.10})$$

Both definitions are common and we note that the square root makes sense¹ because the magnitude is a non-negative real number.

Using the complex conjugate, the **division of one complex number by another** can be defined, and subsequently, understood. We define

$$\frac{x_1 + \mathfrak{i} y_1}{x_2 + \mathfrak{i} y_2} := \frac{(x_1 x_2 + y_1 y_2) + \mathfrak{i} (y_1 x_2 - x_1 y_2)}{(x_2)^2 + (y_2)^2}, \quad (\text{A.11})$$

and note that the denominator is real, and thus the indicated division can be treated as multiplication by one over the denominator. It follows that when $|z_2| \neq 0$, z_1/z_2 is a well-defined complex number. The formula (A.11) is best understood from an alternative definition of complex division

$$\frac{z_1}{z_2} := \frac{z_1 \cdot z_2^*}{z_2 \cdot z_2^*} = \frac{z_1 \cdot z_2^*}{|z_2|^2}. \quad (\text{A.12})$$

Personally, we try to avoid using either one of these formulas and do the computations in Julia! Multiplication and division of complex numbers by hand is very error prone. For probably a century, engineering faculty have been torturing students by making them do such calculations by hand; at some point, it has to stop!

Example A.2 For the following complex numbers, compute their sum, product, division, and magnitudes,

$$\begin{aligned} z_1 &= 2 + \mathfrak{i} 3 \\ z_2 &= -6 + \mathfrak{i} \sqrt{2} \end{aligned}$$

Solution

$$\begin{aligned} z_1 + z_2 &= -4.0000 + \mathfrak{i} 4.4142 \\ z_1 \cdot z_2 &= -16.2426 - \mathfrak{i} 15.1716 \\ \frac{z_1}{z_2} &= -0.2041 - \mathfrak{i} 0.5481 \\ |z_1| &= 3.6056 \\ |z_2| &= 6.1644 \end{aligned}$$

■

A.1.2 Angles of Complex Numbers and Euler's Formula: More Advanced Aspects




Figure A.2: It is often helpful to understand complex numbers as having a magnitude and an angle. This is similar to using polar coordinates in \mathbb{R}^2 . Here, the angle of z is denoted by $\angle z$ instead of θ . Both conventions are common.

¹Julia will recognize $\text{real}(z)^2 + \text{imag}(z)^2$ as being a real number. It does not recognize $z \cdot z^*$ as a real number. In Julia, the command is `abs(z)`, just as with a real number.

For a pair of real numbers (x, y) , we were taught in High School how to express them in **polar coordinates** (ρ, θ) , where

$$\rho := \sqrt{x^2 + y^2}$$

$$\theta := \begin{cases} \arctan(y/x) & x > 0 \\ \pi - \arctan(y/|x|) & x < 0 \\ \text{sign}(y) \frac{\pi}{2} & x = 0, y \neq 0 \\ \text{undefined} & x = 0, y = 0. \end{cases} \quad (\text{A.13})$$

From polar coordinates (ρ, θ) , we computed the **Cartesian Coordinates** (x, y) as

$$\begin{aligned} x &= \rho \cos(\theta) \\ y &= \rho \sin(\theta). \end{aligned} \quad (\text{A.14})$$

Moving beyond High School, we can also express the above in terms of the canonical basis vectors $\{e_1, e_2\}$ for \mathbb{R}^2 as

$$\begin{bmatrix} x \\ y \end{bmatrix} = \rho \cos(\theta)e_1 + \rho \sin(\theta)e_2. \quad (\text{A.15})$$

In fact, any (non-zero) vector $v \in \mathbb{R}^2$ can be expressed as

$$v = \|v\| \cos(\theta)e_1 + \|v\| \sin(\theta)e_2. \quad (\text{A.16})$$

Equation (A.16) hints at a natural way of expressing complex numbers.

Polar Coordinates Meet Complex Numbers and their Multiplication

For a non-zero complex number $z = x + \mathfrak{i}y$, we define its angle as in (A.13). Doing so allows us to express every (non-zero) $z \in \mathbb{C}$ as

$$z = |z| \cos(\theta) + \mathfrak{i} |z| \sin(\theta). \quad (\text{A.17})$$

The real and imaginary parts of z are playing the role of the basis vectors $\{e_1, e_2\}$. One can also think of $\{1.0, \mathfrak{i}\}$ as being a basis for \mathbb{C} , though this is beyond our scope. Equation (A.17) leads to a very nice way to understand the multiplication of two complex numbers.

Fact: Suppose $z_1 = |z_1| \cos(\theta_1) + \mathfrak{i} |z_1| \sin(\theta_1)$ and $z_2 = |z_2| \cos(\theta_2) + \mathfrak{i} |z_2| \sin(\theta_2)$ are non-zero complex numbers. Then,

$$\boxed{\begin{aligned} z_1 \cdot z_2 &= |z_1||z_2| \cos(\theta_1 + \theta_2) + \mathfrak{i} |z_1||z_2| \sin(\theta_1 + \theta_2) \\ \frac{z_1}{z_2} &= \frac{|z_1|}{|z_2|} \cos(\theta_1 - \theta_2) + \mathfrak{i} \frac{|z_1|}{|z_2|} \sin(\theta_1 - \theta_2). \end{aligned}} \quad (\text{A.18})$$

When expressed in “polar form”, multiplying two complex numbers is equivalent to multiplying their magnitudes and adding their angles (or phases), while dividing two complex numbers is equivalent to dividing their magnitudes and subtracting their angles (or phases). Proving (A.18) involves some trigonometric identities. You may want to give it a go. We’ll provide a simpler way to understand it in a few more lines!

Remark: A positive real number has angle zero, while a negative real number has angle π (or, equivalently, $-\pi$). The angle of \mathfrak{i} is $\pi/2$ and the angle of $-\mathfrak{i}$ is $-\pi/2$ (or, equivalently, $3\pi/2$).

The exponential function of a real number x is defined by the infinite series

$$\begin{aligned} e^x &:= \sum_{n=0}^{\infty} \frac{x^n}{n!} \\ &= 1 + x + \frac{1}{2!}x^2 + \frac{1}{3!}x^3 + \frac{1}{4!}x^4 + \dots \\ &= 1 + x + \frac{1}{2}x^2 + \frac{1}{6}x^3 + \frac{1}{24}x^4 + \dots, \end{aligned} \quad (\text{A.19})$$

where $n! := 1 \cdot 2 \cdot 3 \cdots (n-1) \cdot n$, the product of all integers from 1 to n .

Euler's Formula

A very famous result due to the Swiss Mathematician Leonhard Euler (https://en.wikipedia.org/wiki/Leonhard_Euler), asserts that for a given real number θ ,

$$e^{i\theta} = \cos(\theta) + i \sin(\theta) \quad \text{Euler's Formula.} \quad (\text{A.20})$$

Hence, every complex number can be written as

$$z = |z|e^{i\theta}, \quad (\text{A.21})$$

which leads to

Fact: Suppose $z_1 = |z_1|e^{i\theta_1}$ and $z_2 = |z_2|e^{i\theta_2}$ are non-zero complex numbers. Then,

$$\begin{aligned} z_1 \cdot z_2 &= |z_1||z_2|e^{i(\theta_1+\theta_2)} \\ \frac{z_1}{z_2} &= \frac{|z_1|}{|z_2|}e^{i(\theta_1-\theta_2)}. \end{aligned} \quad (\text{A.22})$$

Deriving this result for multiplication and division is much easier than (A.18), but Euler's Formula (A.20) assures us they are the same.

Remark: Deriving Euler's formula is not that hard. When you substitute $i\theta$ into (A.19), you must first note that $(i)^{2n} = (-1)^n$ and $(i)^{2n+1} = (-1)^n i$. If you then separate the power series into its real and imaginary parts, you will recognize the power series for $\cos(\theta)$ and $\sin(\theta)$.

A.1.3 Iterating with Complex Numbers: Background for Eigenvalues

Consider the equation

$$z_{k+1} = az_k, \quad (\text{A.23})$$

with $a \in \mathbb{C}$ and $z_0 \in \mathbb{C}$. Equation (A.23) is technically called a **scalar linear difference equation**, but for us, it looks not so different than iterating with the bisection method or Newton's Algorithm. We compute a few steps until the general pattern of its solution becomes clear:

$$\begin{aligned} z_1 &= az_0 \\ z_2 &= az_1 = a^2 z_0 \\ z_3 &= az_2 = a^3 z_0 \\ &\vdots \\ z_k &= a^k z_0. \end{aligned} \quad (\text{A.24})$$

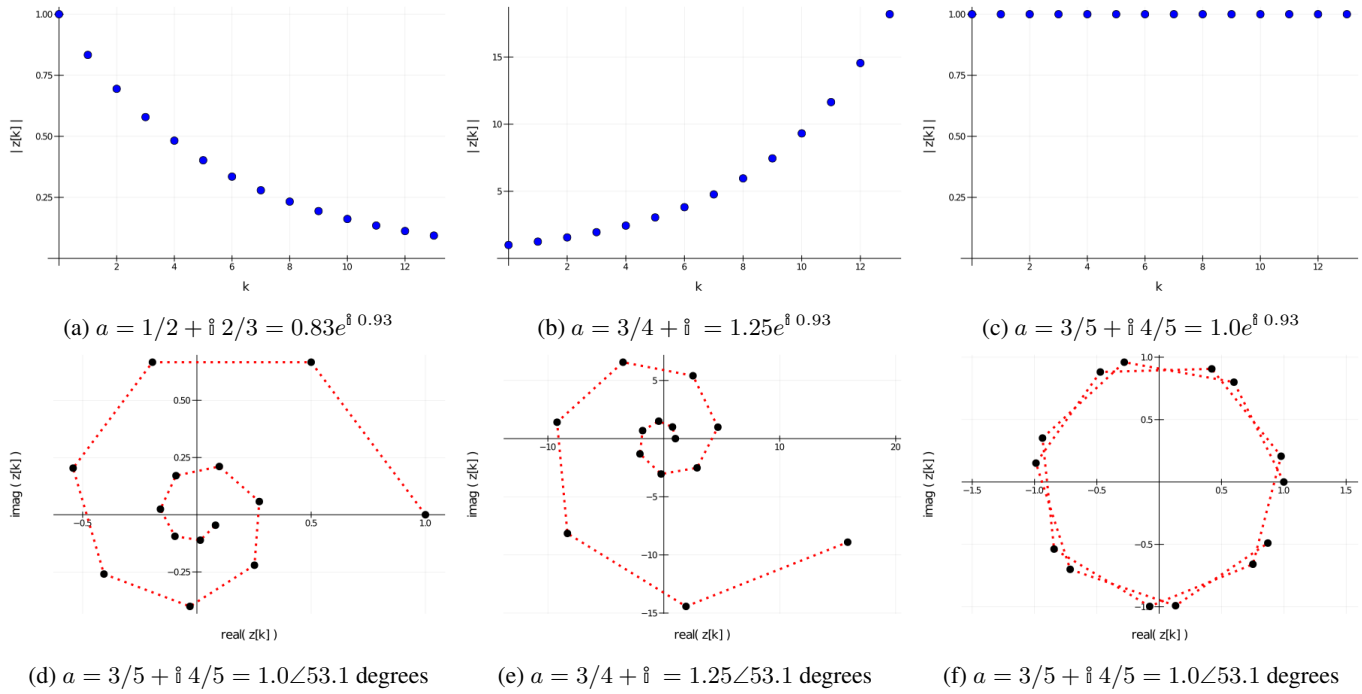


Figure A.3: The solid dots illustrate the evolution of z_k in (A.23) and (A.24) when a has magnitude less than one, greater than one, and equal to one, respectively. In each case, $z_0 = 1.0 + \mathfrak{i} 0.0$ and the angle of a was selected to be 53.1 degrees; therefore the dots rotate counterclockwise. The red dashes are present to guide the eye in connecting the dots. In (f), the points lie on a circle of radius one.

Scalar Linear Difference Equation

The general solution to $z_{k+1} = az_k$, $z_0 \in \mathbb{C}$ is $z_k = a^k z_0$. We write $a = |a|e^{\mathfrak{i}\theta}$, where $\theta = \angle a$, the angle of a as computed in (A.13). Then from (A.22), we conclude that

$$z_k = |a|^k e^{\mathfrak{i} k \angle a} z_0. \quad (\text{A.25})$$

Below, we analyze three cases and show the following for $z_0 \neq 0$

- $|a| < 1 \implies |z_k| \xrightarrow{k \rightarrow \infty} 0$
- $|a| > 1 \implies |z_k| \xrightarrow{k \rightarrow \infty} \infty$
- $|a| = 1 \implies |z_k| = |z_0|, k \geq 0.$

See also Fig. A.3.

The following analysis supports the illustrations in Fig. A.3.

Case 1: $|a| < 1$ We note that $\log(|a|^k) = k \log(|a|)$, and that $|a| < 1 \implies \log(|a|) < 0$. Hence,

$$\lim_{k \rightarrow \infty} |a|^k = \lim_{k \rightarrow \infty} e^{k \log(|a|)} = 0.$$

Case 2: $|a| > 1$ We note that $\log(|a|^k) = k \log(|a|)$, and that $|a| > 1 \implies \log(|a|) > 0$. Hence,

$$\lim_{k \rightarrow \infty} |a|^k = \lim_{k \rightarrow \infty} e^{k \log(|a|)} = \infty.$$

Case 3: $|a|=1$ We note that when $|a| = 1$, then $|a|^k = 1$ for all $k \geq 1$, and thus this case is clear.

A.1.4 \mathbb{C}^n , the Space of Complex Vectors

\mathbb{C}^n sounds harder than it is. It's exactly \mathbb{R}^n where the scalars are complex numbers instead of real numbers. All the definitions of vector addition, linear combinations, spans, and linear independence are the same. We'll cover just a few of the basic ideas so that you get the idea.

Recall that we started by defining \mathbb{R}^n as n -tuples of real numbers and then we identified it with column vectors of length n . We do that same here.

$$\mathbb{C}^n := \{(\alpha_1, \alpha_2, \dots, \alpha_n) \mid \alpha_i \in \mathbb{C}, 1 \leq i \leq n\} \iff \left\{ \begin{bmatrix} \alpha_1 \\ \alpha_2 \\ \vdots \\ \alpha_n \end{bmatrix} \mid \alpha_i \in \mathbb{C}, 1 \leq i \leq n \right\} =: \mathbb{C}^n \quad (\text{A.26})$$

Consider two vectors $v_1 \in \mathbb{C}^n$ and $v_2 \in \mathbb{C}^n$. We define their **vector sum** by

$$v_1 + v_2 = \begin{bmatrix} \alpha_1 \\ \alpha_2 \\ \vdots \\ \alpha_n \end{bmatrix} + \begin{bmatrix} \beta_1 \\ \beta_2 \\ \vdots \\ \beta_n \end{bmatrix} := \begin{bmatrix} \alpha_1 + \beta_1 \\ \alpha_2 + \beta_2 \\ \vdots \\ \alpha_n + \beta_n \end{bmatrix},$$

that is, we sum their respective components or entries. Let γ be a complex number. Then we define

$$\gamma v := \begin{bmatrix} \gamma \alpha_1 \\ \gamma \alpha_2 \\ \vdots \\ \gamma \alpha_n \end{bmatrix},$$

that is, to **multiply a complex vector by a complex number**, we multiply each of the components of the vector by the number, just as we do for real vectors.

Let $\{v_1, v_2, \dots, v_k\}$ be a collection of vectors in \mathbb{C}^n . Then we define their **span** as

$$\text{span}\{v_1, v_2, \dots, v_k\} := \{\alpha_1 v_1 + \alpha_2 v_2 + \dots + \alpha_k v_k \mid \alpha_i \in \mathbb{C}, 1 \leq i \leq k\}.$$

The set of vectors $\{v_1, v_2, \dots, v_k\}$ is **linearly independent** in the vector space \mathbb{C}^n if the only solution to

$$\alpha_1 v_1 + \alpha_2 v_2 + \dots + \alpha_k v_k = 0$$

is $\alpha_1 = 0 + \mathfrak{i}0, \alpha_2 = 0 + \mathfrak{i}0, \dots, \alpha_k = 0 + \mathfrak{i}0$.

The **norm** of a complex vector

$$v = \begin{bmatrix} \alpha_1 \\ \alpha_2 \\ \vdots \\ \alpha_n \end{bmatrix}$$

is

$$\|v\| := \sqrt{\sum_{k=1}^n |\alpha_k|^2},$$

where, to be extra clear, $|\alpha_k|^2 = \alpha_k \cdot \alpha_k^*$. Moreover, if one defines the **complex conjugate of a vector** by taking the complex conjugate of each of its components, then

$$\|v\|^2 = (v^*)^\top \cdot v,$$

and yes, the transpose simply takes the column vector to a row vector.

A.1.5 Iterating with Matrices: The Case for Eigenvalues and Eigenvectors

We now attempt to analyze the matrix versions of (A.23) and (A.24). Recall that you saw matrix difference equations in Project 3. Our real goal is to understand

$$x_{k+1} = Ax_k, \quad (\text{A.27})$$

with A an $n \times n$ real matrix and $x_0 \in \mathbb{R}^n$. But we'll see that allowing the entries of A to be complex and $x_0 \in \mathbb{C}^n$ does not change anything.

With this in mind, we rewrite (A.27) first as

$$x[k+1] = Ax[k],$$

with the time index denoted in square brackets, Julia style! This will allow us to use a subscript for the components of x . Next, we replace $x[k]$ with $z[k]$ to emphasize that we allow $z[k]$ to be a complex vector. We compute a few steps of $z[k+1] = Az[k]$ until the general pattern of its solution becomes clear:

$$\begin{aligned} z[1] &= Az[0] \\ z[2] &= Az[1] = A^2z[0] \\ z[3] &= Az[2] = A^3z[0] \\ &\vdots \\ z[k] &= A^kz[0]. \end{aligned} \quad (\text{A.28})$$

So far so good! Now, our challenges are:

- give conditions on A so that $\|z[k]\|$ contracts, blows up, or stays bounded as k tends to infinity;
- even better, for a given initial condition $z[0]$, describe in detail the evolution of $z[k]$ for $k > 0$.

We'll start with a diagonal $n \times n$ matrix A , and for reasons that will become clear in the next section, we'll denote the entries on the diagonal by λ ,

$$A = \begin{bmatrix} \lambda_1 & 0 & 0 & 0 \\ 0 & \lambda_2 & 0 & 0 \\ 0 & 0 & \ddots & 0 \\ 0 & 0 & 0 & \lambda_n \end{bmatrix}. \quad (\text{A.29})$$

We leave it as an exercise to compute that

$$A^2 = \begin{bmatrix} (\lambda_1)^2 & 0 & 0 & 0 \\ 0 & (\lambda_2)^2 & 0 & 0 \\ 0 & 0 & \ddots & 0 \\ 0 & 0 & 0 & (\lambda_n)^2 \end{bmatrix}, \quad (\text{A.30})$$

and once you have established (A.30), you will have no trouble believing that

$$A^k = \begin{bmatrix} (\lambda_1)^k & 0 & 0 & 0 \\ 0 & (\lambda_2)^k & 0 & 0 \\ 0 & 0 & \ddots & 0 \\ 0 & 0 & 0 & (\lambda_n)^k \end{bmatrix}. \quad (\text{A.31})$$

One thing we can do is observe that

$$\begin{bmatrix} z_1[k] \\ z_2[k] \\ \vdots \\ z_n[k] \end{bmatrix} = \begin{bmatrix} (\lambda_1)^k & 0 & 0 & 0 \\ 0 & (\lambda_2)^k & 0 & 0 \\ 0 & 0 & \ddots & 0 \\ 0 & 0 & 0 & (\lambda_n)^k \end{bmatrix} \begin{bmatrix} z_1[0] \\ z_2[0] \\ \vdots \\ z_n[0] \end{bmatrix} \quad (\text{A.32})$$

results in n scalar equations of the form (A.23), namely,

$$z_j[k] = (\lambda_j)^k z_j[0], \quad 1 \leq j \leq n. \quad (\text{A.33})$$

Linear Difference Equation with a Diagonal Matrix

The general solution to $z[k+1] = Az[k]$, $z[0] \in \mathbb{C}^n$ is $z[k] = A^k z[0]$. When A is diagonal, the solution is given in (A.32) and (A.33). Based on these results and Chapter A.1.3, we analyze three cases for $z_j[0] \neq 0$,

- $|\lambda_j| < 1 \implies |z_j[k]| \xrightarrow[k \rightarrow \infty]{} 0$
- $|\lambda_j| > 1 \implies |z_j[k]| \xrightarrow[k \rightarrow \infty]{} \infty$
- $|\lambda_j| = 1 \implies |z_j[k]| = |z_j[0]|, k \geq 0.$

Being a ROB 101 student, having “real” matrices replaced by diagonal matrices must be a bit disconcerting! You’ll be glad to know that it is really just a step to something kind of magical: most matrices can be factored as $A = M\Lambda M^{-1}$, where $\det(M) \neq 0$ and Λ is diagonal, as in (A.29). But we get ahead of ourselves!

Key features of a Diagonal Matrix: One way that Eigenvalues and Eigenvectors come about

Let $v_j = e_j$, where e_j are the canonical basis vectors for either \mathbb{R}^n or \mathbb{C}^n (aka, columns of the $n \times n$ identity matrix). We have noted before that $Ae_j = a_j^{\text{col}}$. In our case, $a_j^{\text{col}} = \lambda_j e_j$. Hence, substituting in $v_j = e_j$, we arrive at the equation

$$Av_j = \lambda_j v_j, \quad 1 \leq j \leq n. \quad (\text{A.34})$$

Equation (A.34) is the defining relation for eigenvalues (denoted here by λ_j) and eigenvectors (denoted here by v_j). We further note that the set of vectors $\{v_1, v_2, \dots, v_n\}$ is linearly independent and spans both \mathbb{R}^n and \mathbb{C}^n . Having a set of eigenvectors that forms a basis turns out to be a defining characteristic of matrices that are related to a diagonal matrix Λ by a transformation of the form $A = M\Lambda M^{-1}$.

Remark: If $v \in \mathbb{C}^n$ is an eigenvector, meaning $v \neq 0$ and there exists a $\lambda \in \mathbb{C}$ such that (A.34) holds, then we have that

$$\begin{aligned} Av &= \lambda v \\ A^2 v &= A(\lambda v) = \lambda Av = (\lambda)^2 v \\ &\vdots \\ A^k v &= (\lambda)^k v \end{aligned}$$

and hence we can analyze convergence for the difference equation $z[k+1] = Az[k]$, $z[0] = v$, even when A is not diagonal.

Remark: Suppose that A is real and that $\lambda \in \mathbb{C}$ and $v \in \mathbb{C}^n$, satisfy $Av = \lambda v$ and $v \neq 0$. Even though the eigenvalue and eigenvector are complex, their real and imaginary parts are very relevant to computations in \mathbb{R}^n . Decompose λ and v into their real and imaginary parts, viz

$$\begin{aligned} \lambda &=: \lambda_{\text{Re}} + \mathfrak{i} \lambda_{\text{Im}} \\ v &=: v_{\text{Re}} + \mathfrak{i} v_{\text{Im}}. \end{aligned} \quad (\text{A.35})$$

Then

$$\begin{aligned} Av_{\text{Re}} &= \text{real}(Av) = \lambda_{\text{Re}} \cdot v_{\text{Re}} - \lambda_{\text{Im}} \cdot v_{\text{Im}} \\ Av_{\text{Im}} &= \text{imag}(Av) = \lambda_{\text{Im}} \cdot v_{\text{Re}} + \lambda_{\text{Re}} \cdot v_{\text{Im}}. \end{aligned} \quad (\text{A.36})$$

Hence,

$$\begin{bmatrix} Av_{\text{Re}} \\ Av_{\text{Im}} \end{bmatrix} = \begin{bmatrix} \lambda_{\text{Re}} I_n & -\lambda_{\text{Im}} I_n \\ \lambda_{\text{Im}} I_n & \lambda_{\text{Re}} I_n \end{bmatrix} \begin{bmatrix} v_{\text{Re}} \\ v_{\text{Im}} \end{bmatrix}. \quad (\text{A.37})$$

If we write $\lambda = |\lambda|e^{i\theta} = |\lambda| \cos(\theta) + i |\lambda| \sin(\theta)$, then (A.37) can be rewritten as

$$\begin{bmatrix} Av_{\text{Re}} \\ Av_{\text{Im}} \end{bmatrix} = |\lambda| \underbrace{\begin{bmatrix} \cos(\theta)I_n & -\sin(\theta)I_n \\ \sin(\theta)I_n & \cos(\theta)I_n \end{bmatrix}}_{R(\theta)} \begin{bmatrix} v_{\text{Re}} \\ v_{\text{Im}} \end{bmatrix}, \quad (\text{A.38})$$

where $R(\theta)^\top \cdot R(\theta) = R(\theta) \cdot R(\theta)^\top = I_{2n}$, and hence $R(\theta)$ is an orthogonal matrix. This shows how the complex aspect of the eigenvalue and eigenvector manifests itself as a “kind of rotation” of vectors in the two dimensional subspace

$$\text{span}\{v_{\text{Re}}, v_{\text{Im}}\}$$

by $R(\theta)$, in addition to the scaling by $|\lambda|$. A second benefit of the latter expression is that we then have

$$\begin{bmatrix} A^k v_{\text{Re}} \\ A^k v_{\text{Im}} \end{bmatrix} = |\lambda|^k \underbrace{\begin{bmatrix} \cos(k\theta)I_n & -\sin(k\theta)I_n \\ \sin(k\theta)I_n & \cos(k\theta)I_n \end{bmatrix}}_{R(k\theta)} \begin{bmatrix} v_{\text{Re}} \\ v_{\text{Im}} \end{bmatrix}. \quad (\text{A.39})$$

Figure A.4 illustrates a case where $\lambda = 0.9803 \pm i 0.0965 = 0.985 \angle 5.6$ degrees. The rotating and decaying nature of the solution is clearly seen in the figure. The reader should compare Figs. A.3-(a) and -(d) with Fig. A.4.




Figure A.4: The eigenvalues of a real 3×3 matrix are computed to be $0.9803 \pm i 0.0965$ and 1.100 . The initial condition in green was chosen to be a linear combination of the real and imaginary parts of an eigenvector corresponding to the complex pair of eigenvalues. The resulting solution of (A.28) evolves in the plane defined by $\text{span}\{v_{\text{Re}}, v_{\text{Im}}\}$ as indicated by (A.37) through (A.39). This is very analogous to how complex numbers, when iterated, evolve in the Complex Plane.

A.2 Eigenvalues and Eigenvectors

The study of eigenvalues and eigenvectors is very traditional in Linear Algebra courses. We skipped very important aspects of them in the main portion of the book for a few reasons: (1) time is limited; (2) they get complicated really fast; and (3) their most important applications are the evolution of linear difference equations and the Singular Value Decomposition (SVD), neither of which were covered in the main portion of the text. The usual illustrative application of eigenvalues and eigenvectors is to “diagonalize” a matrix,

which we treated indirectly in Chapter 10.17. In the context of Chapter A.1.5 and your Segway Project, it does make sense.

Just in case you are starting here and skipped Appendix A.1 entirely, we start from the beginning.

A.2.1 General Square Matrices

Temporary Def. Let A be an $n \times n$ matrix with real coefficients. A scalar $\lambda \in \mathbb{R}$ is an **eigenvalue** (e-value) of A , if there exists a non-zero vector $v \in \mathbb{R}^n$ such that $A \cdot v = \lambda v$. Any such vector v is called an **eigenvector** (e-vector) associated with λ .

We note that if v is an e-vector, then so is αv for any $\alpha \neq 0$, and therefore, e-vectors are not unique. To find eigenvalues, we need to have conditions under which there exists $v \in \mathbb{R}^n, v \neq 0$, such that $A \cdot v = \lambda v$. Here they are,

$$A \cdot v = \lambda v \iff (\lambda I - A) \cdot v = 0 \xrightarrow{v \neq 0} \det(\lambda I - A) = 0.$$

Example A.3 Let A be the 2×2 real matrix $A = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}$. Determine, if any, its e-values and e-vectors.

Solution: To find e-values, we need to solve

$$\det(\lambda I - A) = \begin{vmatrix} \lambda & -1 \\ 1 & \lambda \end{vmatrix} = \lambda^2 + 1 = 0.$$

We compute the discriminant of this quadratic equation and we find

$$b^2 - 4ac = -4 < 0,$$

and therefore there are no real solutions. Hence, by our *temporary definition*, this 2×2 real matrix does not have any e-values, and hence, neither does it have any e-vectors.

If we were to allow e-values to be complex numbers, then we'd have two e-values corresponding to the two complex solutions of the quadratic equation $\lambda^2 + 1 = 0$, namely, $\lambda_1 = \mathfrak{i}$ and $\lambda_2 = -\mathfrak{i}$.

We'll see shortly that we'll also need to allow the e-vectors to have complex entries. Hence, we need to generalize our temporary definition. ■

Permanent Definition of Eigenvalues and Eigenvectors

Let A be an $n \times n$ matrix with real or complex coefficients. A scalar $\lambda \in \mathbb{C}$ is an **eigenvalue** (e-value) of A , if there exists a non-zero vector $v \in \mathbb{C}^n$ such that $Av = \lambda v$. Any such vector v is called an **eigenvector** (e-vector) associated with λ .

Eigenvectors are not unique.

- To find e-values, we solve $\det(\lambda I - A) = 0$ because

$$A \cdot v = \lambda v \iff (\lambda I - A) \cdot v = 0 \xrightarrow{v \neq 0} \det(\lambda I - A) = 0. \tag{A.40}$$

- To find e-vectors, we find any non-zero $v \in \mathbb{C}^n$ such that

$$(\lambda I - A) \cdot v = 0. \tag{A.41}$$

Of course, if you prefer, you can solve $(A - \lambda I)v = 0$ when seeking e-vectors.

Fundamental Theorem of Algebra (and a bit More)

Let A be an $n \times n$ matrix with real or complex coefficients. Then the following statements are true

- $\det(\lambda I - A) = \lambda^n + \alpha_{n-1}\lambda^{n-1} + \cdots + \alpha_1\lambda + \alpha_0$, and if A is real, so are the coefficients $\alpha_{n-1}, \dots, \alpha_0$.
- The degree n polynomial $\lambda^n + \alpha_{n-1}\lambda^{n-1} + \cdots + \alpha_1\lambda + \alpha_0$ has n roots $\lambda_1, \dots, \lambda_n \in \mathbb{C}$ such that

$$\det(\lambda I - A) = (\lambda - \lambda_1)(\lambda - \lambda_2) \cdots (\lambda - \lambda_n).$$

Each of the roots λ_i , $1 \leq i \leq n$, is an e-value of A .

- The e-values $\{\lambda_1, \dots, \lambda_n\}$ are said to be **distinct** if $\lambda_i \neq \lambda_k$ for all $i \neq k$.
- If $\lambda_i = \lambda_k$ for some $i \neq k$, then λ_i is a **repeated e-value**. The e-values can then be grouped into $1 \leq p \leq n$ sets of distinct roots $\{\lambda_1, \dots, \lambda_p\}$ such that

$$\det(\lambda I - A) = (\lambda - \lambda_1)^{m_1}(\lambda - \lambda_2)^{m_2} \cdots (\lambda - \lambda_p)^{m_p}.$$

The integer m_i is called the **algebraic multiplicity** of λ_i and their sum satisfies $m_1 + m_2 + \cdots + m_p = n$.

- An e-vector associated with λ_i is computed by finding non-zero solutions to (A.41).
- If the matrix A is real, then the e-values occur in **complex conjugate pairs**, that is, if λ_i is an e-value then so is λ_i^* .
- If the matrix A is real and the e-value λ_i is real, then the e-vector v_i can always be chosen to be real, that is, $v_i \in \mathbb{R}^n$ instead of $v_i \in \mathbb{C}^n$.
- There will always be at least one non-zero solution to (A.41), and because any non-zero multiple of a solution is also a solution, there will always be an infinite number of solutions to (A.41).
- If λ_i is a repeated e-value with algebraic multiplicity m_i , then the number of linearly independent e-vectors associated with λ_i is upper bounded by m_i . Another way to say this is, $1 \leq \dim(\text{Null}(A - \lambda_i I)) \leq m_i$.
- **In Julia**, after using `LinearAlgebra`, the commands are `Λ = eigvals(A)` and `V = eigvecs(A)`

Example A.4 Let A be the 2×2 real matrix that we treated in Example A.3, namely, $A = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}$. Determine its e-values and e-vectors in the sense of our “permanent” definition.

Solution: As in Example A.3, to find e-values, we solve

$$\det(\lambda I - A) = \begin{vmatrix} \lambda & -1 \\ 1 & \lambda \end{vmatrix} = \lambda^2 + 1 = 0.$$

We apply the quadratic equation and determine $\lambda_1 = \mathbf{i}$ and $\lambda_2 = -\mathbf{i}$. To find the eigenvectors, we solve

$$(A - \lambda_i I)v_i = 0.$$

The eigenvectors are

$$v_1 = \begin{bmatrix} 1 \\ \mathbf{i} \end{bmatrix}, v_2 = \begin{bmatrix} 1 \\ -\mathbf{i} \end{bmatrix}.$$

Note that the eigenvalues and eigenvectors each form complex conjugate pairs. Indeed,

$$\lambda_2 = \lambda_1^* \text{ and } v_2 = v_1^*.$$

■

Example A.5 Let A be the $n \times n$ identity matrix. Determine its e-values and e-vectors.

Solution: $\det(\lambda I - I) = \det((\lambda - 1)I) = 0 \iff \lambda = 1$. Alternatively, you can compute that $\det(\lambda I - I) = (\lambda - 1)^n$. Hence, the e-value $\lambda = 1$ is repeated n times, that is, $m_1 = n$. What are the e-vectors? We seek to solve

$$(A - \lambda I) \cdot v = 0 \iff (I - 1 \cdot I) \cdot v = 0 \iff 0_n \cdot v = 0,$$

where 0_n is the $n \times n$ matrix of all zeros! Hence, any non-zero vector $v \in \mathbb{R}^n$ is an e-vector. Moreover, if $\{v_1, \dots, v_n\}$ is a basis for \mathbb{R}^n , then $\{v_1, \dots, v_n\}$ is a set of n linearly independent e-vectors associated with $\lambda_1 = 1$. ■

Example A.6 Let $a \in \mathbb{R}$ be a constant and let A be the 4×4 matrix below. Determine its e-values and e-vectors.

$$A = \begin{bmatrix} a & 1 & 0 & 0 \\ 0 & a & 1 & 0 \\ 0 & 0 & a & 1 \\ 0 & 0 & 0 & a \end{bmatrix}.$$

Solution: To find the e-values, we solve

$$\det(\lambda I - A) = \det \left(\begin{bmatrix} (\lambda - a) & -1 & 0 & 0 \\ 0 & (\lambda - a) & -1 & 0 \\ 0 & 0 & (\lambda - a) & -1 \\ 0 & 0 & 0 & (\lambda - a) \end{bmatrix} \right) = (\lambda - a)^4 = 0,$$

and hence there is one distinct e-value $\lambda_1 = a$. To solve for e-vector(s) we consider

$$0 = (A - aI) \cdot v = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \end{bmatrix} \cdot v$$

and we find that the only solutions are multiples of

$$v = \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \end{bmatrix}.$$

■

We've seen the extremes! A matrix with a single distinct e-value and a complete set of e-vectors (there were n linearly independent e-vectors associated with the e-value), and another matrix with a single distinct e-value, but only one linearly independent e-vector associated with it.

When the e-values are Distinct, the e-vectors form a Basis

Let A be an $n \times n$ matrix with coefficients in \mathbb{R} or \mathbb{C} . If the e-values $\{\lambda_1, \dots, \lambda_n\}$ are distinct, that is, $\lambda_i \neq \lambda_j$ for all $1 \leq i \neq j \leq n$, then the e-vectors $\{v_1, \dots, v_n\}$ are linearly independent in $(\mathbb{C}^n, \mathbb{C})$.

Restatement of the result: If $\{\lambda_1, \dots, \lambda_n\}$ are distinct, then $\{v_1, \dots, v_n\}$ is a basis for $(\mathbb{C}^n, \mathbb{C})$. If A is real and its e-values are real, then the e-vectors can be chosen to be real and they form a basis for \mathbb{R}^n .

A.2.2 Real Symmetric Matrices

We recall that a real $n \times n$ matrix A is **symmetric** if $A^T = A$. E-values and e-vectors of symmetric matrices have nicer properties than those of general matrices.

E-values and E-vectors of Symmetric Matrices

- The e-values of a symmetric matrix are real. Because the e-values are real and the matrix is real, we can always choose the e-vectors to be real. Moreover, we can always normalize the e-vectors to have **norm one**.
- Just as with general matrices, the e-values of a symmetric matrix may be distinct or repeated. However, even when an e-value λ_i is repeated m_i times, there are always m_i linearly independent e-vectors associated with it. By applying Gram-Schmidt, we can always choose these e-vectors to be **orthonormal**.
- E-vectors associated with distinct e-values are automatically orthogonal. To be clear,

$$(A^T = A, Av_i = \lambda_i v_i, Av_k = \lambda_k v_k, \text{ and } \lambda_i \neq \lambda_k) \implies v_i \perp v_k.$$

Since we can assume they have length one, we have that the e-vectors are **orthonormal**.

- **In summary**, when A is symmetric, there is always an orthonormal basis $\{v_1, v_2, \dots, v_n\}$ for \mathbb{R}^n consisting of e-vectors of A . In other words, for all $1 \leq i \leq n$, $Av_i = \lambda_i v_i$, $\|v_i\| = 1$, and for $k \neq i$, $v_k \perp v_i$.

Factoring a Symmetric Matrix

For every real $n \times n$ symmetric matrix A , there exists an $n \times n$ diagonal matrix Λ and an $n \times n$ orthogonal matrix Q such that

$$A = Q \cdot \Lambda \cdot Q^T. \quad (\text{A.42})$$

Moreover,

$$\Lambda = \begin{bmatrix} \lambda_1 & 0 & \cdots & 0 \\ 0 & \lambda_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \lambda_n \end{bmatrix} \quad \text{and} \quad Q = [v_1 \ v_2 \ \cdots \ v_n]$$

are constructed from the e-values of A and a corresponding set of orthonormal e-vectors.

Remark 01: From (A.42), $\det(A) = \det(Q) \cdot \det(\Lambda) \cdot \det(Q^T) = \det(\Lambda) = \lambda_1 \cdot \lambda_2 \cdots \lambda_n$. Hence, a symmetric real matrix A is invertible if, and only if, all of its e-values are non-zero. Moreover, in this case

$$A^{-1} = Q \cdot \Lambda^{-1} \cdot Q^T.$$

While a similar result holds for general square matrices, it requires inverting the matrix formed by stacking the e-vectors as columns, and hence is not numerically attractive. For symmetric matrices, the corresponding inverse is computed via a matrix transpose.

Remark 02: Using the fact that matrix multiplication can be realized by summing over the product of columns times rows, (A.42) can be rewritten as

$$A = \sum_{i=1}^n \lambda_i (v_i \cdot v_i^T). \quad (\text{A.43})$$

Equations (A.42) and (A.43) parallel results we will develop for the Singular Value Decomposition or (SVD). Equation (A.42) factors A into a product of three terms consisting of two orthogonal matrices and a diagonal matrix, while (A.43) is an expansion of A into “rank one” matrices.

A.3 Positive Definite Matrices

Some Definitions and Facts

Def. Let P be an $n \times n$ real matrix and $x \in \mathbb{R}^n$. Then $x^\top Px$ is called a **quadratic form**.

Def. An $n \times n$ matrix S is **skew symmetric** if $S^\top = -S$.

Fact If S is skew symmetric, then $x^\top Sx = 0$ for all $x \in \mathbb{R}^n$.

Fact Let P an $n \times n$ real matrix and write

$$P = \frac{1}{2}(P + P^\top) + \frac{1}{2}(P - P^\top).$$

Then $(P + P^\top)$ is symmetric, $(P - P^\top)$ is skew symmetric, and we see that every (real) square matrix can be written as the sum of a symmetric matrix and a skew symmetric matrix.

Fact Let P an $n \times n$ real matrix. Then, for all $x \in \mathbb{R}^n$

$$x^\top Px = \frac{1}{2}x^\top (P + P^\top)x.$$

Hence, a quadratic form only depends on the symmetric part of a matrix.

Consequence: When working with a quadratic form, $x^\top Px$, one **ALWAYS** assumes that the matrix P is symmetric. Allowing the matrix to be non-symmetric does not increase the generality of the notion of a quadratic form. This is because $x^\top Px = \frac{1}{2}x^\top (P + P^\top)x$ implies that one can always replace P with its symmetric part!

Fact For an $n \times n$ symmetric real matrix P with e-values $\lambda_1, \dots, \lambda_n$, let $\lambda_{\max} := \max_{1 \leq i \leq n} \lambda_i$ and $\lambda_{\min} := \min_{1 \leq i \leq n} \lambda_i$ be the max and min, respectively over the e-values. Then, for all $x \in \mathbb{R}^n$,

$$\lambda_{\min} x^\top x \leq x^\top Px \leq \lambda_{\max} x^\top x. \quad (\text{A.44})$$

Because $x^\top x = \|x\|^2$, the above expression is also commonly written as

$$\lambda_{\min} \|x\|^2 \leq x^\top Px \leq \lambda_{\max} \|x\|^2.$$

Both are useful.

Equation (A.44) is established by choosing an orthonormal set of e-vectors for P , $\{v_1, \dots, v_n\}$, which we know forms a basis for \mathbb{R}^n . Hence, for all $x \in \mathbb{R}^n$, there exist coefficients $\alpha_1, \dots, \alpha_n$ such that $x = \alpha_1 v_1 + \dots + \alpha_n v_n$. Then, using the two facts we have at our disposal, namely (a) $\{v_1, \dots, v_n\}$ is orthonormal and (b), $Av_i = \lambda_i v_i$, we compute

$$\begin{aligned} x^\top x &= \alpha_1^2 + \dots + \alpha_n^2 \\ x^\top Px &= \lambda_1 \alpha_1^2 + \dots + \lambda_n \alpha_n^2. \end{aligned}$$

It follows that

$$\lambda_{\min} x^\top x = \lambda_{\min} \alpha_1^2 + \dots + \lambda_{\min} \alpha_n^2 \leq \lambda_1 \alpha_1^2 + \dots + \lambda_n \alpha_n^2 \leq \lambda_{\max} \alpha_1^2 + \dots + \lambda_{\max} \alpha_n^2 = \lambda_{\max} x^\top x,$$

showing that (A.44) holds.

Positive Definite and Semidefinite Matrices

Def. A real symmetric matrix P is **positive definite**, if for all $x \in \mathbb{R}^n$, $x \neq 0 \implies x^\top P x > 0$. The common notation for such matrices is $P > 0$.

Def. A real symmetric matrix P is positive semidefinite, if for all $x \in \mathbb{R}^n \implies x^\top P x \geq 0$. The common notation for such matrices is $P \geq 0$.

From (A.44), we arrive at the following facts.

Fact A symmetric matrix P is positive definite if, and only if, all of its eigenvalues are greater than 0.

Fact A symmetric matrix P is positive semidefinite if, and only if, all of its eigenvalues are greater than or equal to 0.

Example A.7 Determine, which, if any, of the following matrices are positive definite or positive semidefinite.

$$P_1 = \begin{bmatrix} 2 & -1 \\ -1 & 2 \end{bmatrix}, P_2 = \begin{bmatrix} 1 & 2 \\ 2 & 1 \end{bmatrix}, P_3 = \begin{bmatrix} 4 & 2 \\ 2 & 1 \end{bmatrix}, P_4 = \begin{bmatrix} 4 & 1 \\ 3 & 4 \end{bmatrix}$$

Solution: Because P_4 is not symmetric, it cannot be positive definite or positive semidefinite! Using Julia, we compute the e-values of P_1 , P_2 , and P_3

$$P_1 \implies \lambda_1 = 1, \lambda_2 = 3 \implies P > 0$$

$$P_2 \implies \lambda_1 = -1, \lambda_2 = 3 \implies P \not\geq 0 \text{ (neither positive semidefinite nor positive definite)}$$

$$P_3 \implies \lambda_1 = 0, \lambda_2 = 5 \implies P \geq 0.$$

We note that P being positive definite does NOT mean that all of its entries have to be positive! P can have entries with negative values and still be positive definite. We note that all of the entries of P being positive does NOT imply that P is even positive semidefinite. ■

Computing e-values is a terrible way to determine if a matrix is positive definite or not. The following facts imply that LDLT Factorization can be applied to positive definite and positive semidefinite matrices.

More on Positive Definite and Semidefinite Matrices

Fact A symmetric $n \times n$ matrix P is positive semidefinite if, and only if, there exists a $k \times n$ matrix N such that $N^\top \cdot N = P$.

Fact A symmetric $n \times n$ matrix P is positive definite if, and only if, there exists an $n \times n$ matrix N with linearly independent columns such that $N^\top \cdot N = P$.

Turning these ideas into algorithmic form gives the following:

LDLT or LU Factorization to Check $P > 0$

Here are two better ways to test whether a matrix is positive definite, positive semidefinite, or neither:

- From Chapter 7.6, we can do the LDLT factorization of P , namely

$$Q \cdot P \cdot Q^T = L \cdot D \cdot L^T,$$

where we have used Q to denote the row permutation matrix because P is being used for a symmetric matrix. Then $P > 0 \iff D > 0$, where D is diagonal.

- Straight up LU with no permutations at all: The key fact is that, if an $n \times n$ matrix P is symmetric and invertible, then it can be written as

$$P = L \cdot U;$$

you can do the factorization without permuting any of the rows of P . Moreover, there is always a diagonal matrix D such that

$$U = D \cdot L^T.$$

Determining D from U is trivial: you just normalize by the diagonal of U . Then, $P = L \cdot D \cdot L^T$ and

$$P > 0 \iff L \cdot D \cdot L^T > 0 \iff D > 0, \text{ that is, all of the entries on the diagonal of } D \text{ are positive.}$$

- If the LU Factorization without permutations fails, then P is not positive definite, but could be positive semidefinite. To rule out the latter, you need to run the full LDLT algorithm and check that the diagonal of D has at least one negative entry or not.

Example A.8 Using the LU Factorization without row permutations, determine if the randomly generated symmetric matrix P is positive definite or not.

$$P = \begin{bmatrix} 8.241e-01 & 1.171e+00 & 1.117e+00 & 1.706e+00 & 1.021e+00 \\ 1.171e+00 & 1.574e+00 & 8.547e-01 & 1.102e+00 & 2.871e-01 \\ 1.117e+00 & 8.547e-01 & 1.238e+00 & 7.506e-01 & 1.291e+00 \\ 1.706e+00 & 1.102e+00 & 7.506e-01 & 2.943e-01 & 9.570e-01 \\ 1.021e+00 & 2.871e-01 & 1.291e+00 & 9.570e-01 & 1.448e+00 \end{bmatrix}.$$

Solution: We do the LU Factorization without permutations and obtain

$$L = \begin{bmatrix} 1.000e+00 & 0.000e+00 & 0.000e+00 & 0.000e+00 & 0.000e+00 \\ 1.421e+00 & 1.000e+00 & 0.000e+00 & 0.000e+00 & 0.000e+00 \\ 1.356e+00 & 8.151e+00 & 1.000e+00 & 0.000e+00 & 0.000e+00 \\ 2.070e+00 & 1.469e+01 & 1.616e+00 & 1.000e+00 & 0.000e+00 \\ 1.239e+00 & 1.294e+01 & 1.649e+00 & 5.893e-01 & 1.000e+00 \end{bmatrix}$$

$$U = \begin{bmatrix} 8.241e-01 & 1.171e+00 & 1.117e+00 & 1.706e+00 & 1.021e+00 \\ 0.000e+00 & -8.990e-02 & -7.328e-01 & -1.321e+00 & -1.164e+00 \\ 0.000e+00 & 0.000e+00 & 5.696e+00 & 9.203e+00 & 9.393e+00 \\ 0.000e+00 & 0.000e+00 & 0.000e+00 & 1.296e+00 & 7.636e-01 \\ 0.000e+00 & 0.000e+00 & 0.000e+00 & 0.000e+00 & -6.909e-01 \end{bmatrix}$$

We extract the diagonal of U and we form $D \cdot L^\top$

$$D = \begin{bmatrix} 8.241e-01 & 0.000e+00 & 0.000e+00 & 0.000e+00 & 0.000e+00 \\ 0.000e+00 & -8.990e-02 & 0.000e+00 & 0.000e+00 & 0.000e+00 \\ 0.000e+00 & 0.000e+00 & 5.696e+00 & 0.000e+00 & 0.000e+00 \\ 0.000e+00 & 0.000e+00 & 0.000e+00 & 1.296e+00 & 0.000e+00 \\ 0.000e+00 & 0.000e+00 & 0.000e+00 & 0.000e+00 & -6.909e-01 \end{bmatrix}$$

$$D \cdot L^\top = \begin{bmatrix} 8.241e-01 & 1.171e+00 & 1.117e+00 & 1.706e+00 & 1.021e+00 \\ 0.000e+00 & -8.990e-02 & -7.328e-01 & -1.321e+00 & -1.164e+00 \\ 0.000e+00 & 0.000e+00 & 5.696e+00 & 9.203e+00 & 9.393e+00 \\ 0.000e+00 & 0.000e+00 & 0.000e+00 & 1.296e+00 & 7.636e-01 \\ 0.000e+00 & 0.000e+00 & 0.000e+00 & 0.000e+00 & -6.909e-01 \end{bmatrix}$$

and we recognize, that indeed, $U = D \cdot L^\top$.

Back to the question of determining whether P is positive definite? We see that D has non-positive entries and therefore P is not positive definite. We only formed $D \cdot L^\top$ to illustrate that $P = L \cdot D \cdot L^\top$. ■

The following results are primarily of use for “hand calculations” or proving results about positive definite matrices. We include them for completeness.

Schur Complement Theorem: A way to Decompose the Test for Being Positive Definite

Suppose that A is $n \times n$, symmetric, and invertible, B is $n \times m$, C is $m \times m$, symmetric, and invertible, and

$$M := \begin{bmatrix} A & B \\ B^\top & C \end{bmatrix},$$

which is then $(n + m) \times (n + m)$ and symmetric. Under these conditions, the following three statements are equivalent:

- (a) $M > 0$.
- (b) $A > 0$, and $C - B^\top \cdot A^{-1} \cdot B > 0$.
- (c) $C > 0$, and $A - B \cdot C^{-1} \cdot B^\top > 0$.

Remarks:

- $C - B^\top \cdot A^{-1} \cdot B$ is called the Schur Complement of A in M .
- $A - B \cdot C^{-1} \cdot B^\top$ is called the Schur Complement of C in M .

A.4 Singular Value Decomposition or SVD

The material here is inspired by a handout prepared by Prof. James Freudenberg, EECS, University of Michigan.

A.4.1 Motivation

In abstract linear algebra, a set of vectors is either linearly independent or not. There is nothing in between. For example, the set of vectors

$$\left\{ v_1 = \begin{bmatrix} 1 \\ 1 \end{bmatrix}, v_2 = \begin{bmatrix} 0.999 \\ 1 \end{bmatrix} \right\}$$

is linearly independent. In this case, one looks at the set of vectors and says, yes, BUT, the vectors are “almost” dependent because when one computes the determinant

$$\det \begin{bmatrix} 1 & 0.999 \\ 1 & 1 \end{bmatrix} = 0.001,$$

the result is pretty small, so it should be fine to call them dependent.

Well, what about the set

$$\left\{ v_1 = \begin{bmatrix} 1 \\ 0 \end{bmatrix}, v_2 = \begin{bmatrix} 10^4 \\ 1 \end{bmatrix} \right\}?$$

When you form the matrix and check the determinant, you get

$$\det \begin{bmatrix} 1 & 10^4 \\ 0 & 1 \end{bmatrix} = 1,$$

which seems pretty far from zero. So are these vectors “adequately” linearly independent?

Maybe not! Let’s note that

$$\begin{bmatrix} 1 & 10^4 \\ 0 & 1 \end{bmatrix} + \begin{bmatrix} 0 & 0 \\ 10^{-4} & 0 \end{bmatrix} = \begin{bmatrix} 1 & 10^4 \\ 10^{-4} & 1 \end{bmatrix},$$

and its determinant is zero! Hence, it’s possible to add a very small perturbation to one of the vectors and make the set linearly dependent! This cannot be good.

A.4.2 Definition and Main Theorem

Rectangular Diagonal Matrix

An $n \times m$ matrix Σ is a **Rectangular Diagonal Matrix** if

$$\Sigma_{ij} = 0 \text{ for } i \neq j.$$

Alternative and equivalent way to define Rectangular Diagonal is

(a) (tall matrix) $n > m$ $\Sigma = \begin{bmatrix} \Sigma_d \\ 0 \end{bmatrix}$, where Σ_d is an $m \times m$ diagonal matrix.

(b) (wide matrix) $n < m$ $\Sigma = \begin{bmatrix} \Sigma_d & 0 \end{bmatrix}$, where Σ_d is an $n \times n$ diagonal matrix.

The **diagonal** of Σ is defined to be the diagonal of Σ_d .

Singular Value Decomposition (Main Theorem)

Every $n \times m$ real matrix A can be factored as

$$A = U \cdot \Sigma \cdot V^T,$$

where U is an $n \times n$ orthogonal matrix, V is an $m \times m$ orthogonal matrix, Σ is an $n \times m$ rectangular diagonal matrix, and the diagonal of Σ ,

$$\text{diag}(\Sigma) = [\sigma_1, \sigma_2, \dots, \sigma_p],$$

satisfies $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_p \geq 0$, for $p := \min(n, m)$.

Moreover, the columns of U are eigenvectors of $A \cdot A^T$, the columns of V are eigenvectors of $A^T \cdot A$, and $\{\sigma_1^2, \sigma_2^2, \dots, \sigma_p^2\}$ are eigenvalues of both $A^T \cdot A$ and $A \cdot A^T$.

The **Singular Values of A** are the elements $\{\sigma_1, \dots, \sigma_p\}$ from the diagonal of Σ .

Another way to write the SVD of A is

$$A = \sigma_1 u_1 \cdot v_1^T + \sigma_2 u_2 \cdot v_2^T + \dots + \sigma_p u_p \cdot v_p^T,$$

where u_i and v_i are columns of U and V respectively.

$$U = [u_1 \quad u_2 \quad \dots \quad u_n] \quad \text{and} \quad V = [v_1 \quad v_2 \quad \dots \quad v_m]. \quad (\text{A.45})$$

This formula follows from our matrix multiplication formulation through the sum over columns times rows, where we note that the columns of V are the rows of V^T .

Rank and Nullity of a Matrix

The **rank** of an $n \times m$ matrix A is the dimension of its column span and the **nullity** of A is the dimension of its null space. Let r be the number of non-zero singular values of A . Then

Fact $\text{rank}(A) := \dim \text{col span}\{A\} = r$.

Fact $\text{nullity}(A) := \dim \text{null}(A) = m - r$.

Example A.9 Determine the SVD of A as well as its rank and nullity,

$$A = \begin{bmatrix} 1 & 10^4 \\ 0 & 1 \end{bmatrix}.$$

Solution: Using the LinearAlgebra package in Julia, we find

$$U = \begin{bmatrix} 1.000e+00 & -1.000e-04 \\ 1.000e-04 & 1.000e+00 \end{bmatrix}$$

$$\Sigma = \begin{bmatrix} 1.000e+04 & 0.000e+00 \\ 0.000e+00 & 1.000e-04 \end{bmatrix}$$

$$V = \begin{bmatrix} 1.000e-04 & -1.000e+00 \\ 1.000e+00 & 1.000e-04 \end{bmatrix}$$

There are two non-zero singular values, and thus $r = 2$. It follows that $\text{rank}(A) = 2$ and $\text{nullity}(A) = 0$.

Information about the “near” linear dependence of the columns of A is in the diagonal matrix Σ . There are two singular values, $\sigma_1 = 10^4$ and $\sigma_2 = 10^{-4}$. Their ratio is 10^8 , which is an indicator that these vectors are “nearly linearly dependent”. “Numerically”, one would say that $r = 1$ and hence $\text{rank}(A) = r = 1$ and $\text{nullity}(A) = 2 - r = 1$. ■

A.4.3 Numerical Linear Independence

Illustration: 5×5 matrix. For

$$A = \begin{bmatrix} -32.57514 & -3.89996 & -6.30185 & -5.67305 & -26.21851 \\ -36.21632 & -11.13521 & -38.80726 & -16.86330 & -1.42786 \\ -5.07732 & -21.86599 & -38.27045 & -36.61390 & -33.95078 \\ -36.51955 & -38.28404 & -19.40680 & -31.67486 & -37.34390 \\ -25.28365 & -38.57919 & -31.99765 & -38.36343 & -27.13790 \end{bmatrix},$$

and the Julia commands

```
1 using LinearAlgebra
2
3 A=[-32.57514 -3.89996 -6.30185 -5.67305 -26.21851;
4 -36.21632 -11.13521 -38.80726 -16.86330 -1.42786;
5 -5.07732 -21.86599 -38.27045 -36.61390 -33.95078;
6 -36.51955 -38.28404 -19.40680 -31.67486 -37.34390;
7 -25.28365 -38.57919 -31.99765 -38.36343 -27.13790 ]
8
9 (U ,Sigma, V) = svd(A)
```

one obtains

$$U = \begin{bmatrix} -2.475e-01 & -5.600e-01 & 4.131e-01 & 5.759e-01 & 3.504e-01 \\ -3.542e-01 & -5.207e-01 & -7.577e-01 & -1.106e-02 & -1.707e-01 \\ -4.641e-01 & 6.013e-01 & -1.679e-01 & 6.063e-01 & -1.652e-01 \\ -5.475e-01 & -1.183e-01 & 4.755e-01 & -3.314e-01 & -5.919e-01 \\ -5.460e-01 & 1.992e-01 & -2.983e-02 & -4.369e-01 & 6.859e-01 \end{bmatrix}$$

$$\Sigma = \begin{bmatrix} 1.325e+02 & 0.000e+00 & 0.000e+00 & 0.000e+00 & 0.000e+00 \\ 0.000e+00 & 3.771e+01 & 0.000e+00 & 0.000e+00 & 0.000e+00 \\ 0.000e+00 & 0.000e+00 & 3.342e+01 & 0.000e+00 & 0.000e+00 \\ 0.000e+00 & 0.000e+00 & 0.000e+00 & 1.934e+01 & 0.000e+00 \\ 0.000e+00 & 0.000e+00 & 0.000e+00 & 0.000e+00 & 7.916e-01 \end{bmatrix}$$

$$V = \begin{bmatrix} 4.307e-01 & 8.839e-01 & -5.303e-02 & 8.843e-02 & -1.503e-01 \\ 4.309e-01 & -2.207e-01 & -1.961e-01 & 7.322e-01 & 4.370e-01 \\ 4.617e-01 & -8.902e-02 & 7.467e-01 & -3.098e-01 & 3.539e-01 \\ 4.730e-01 & -3.701e-01 & 7.976e-02 & 1.023e-01 & -7.890e-01 \\ 4.380e-01 & -1.585e-01 & -6.283e-01 & -5.913e-01 & 1.968e-01 \end{bmatrix}$$

Because the **smallest singular value** $\sigma_5 = 0.7916$ is less than 1% of the largest singular value $\sigma_1 = 132.5$, in many cases, one would say that the numerical rank of A was 4 instead of 5.

This notion of numerical rank can be formalized by asking the following question: Suppose $\text{rank}(A) = r$. How far away is A from a matrix of rank strictly less than r ?

The numerical rank of a matrix is based on the expansion in (A.4.2), which is repeated here for convenience,

$$A = U \cdot \Sigma \cdot V^T = \sum_{i=1}^p \sigma_i u_i \cdot v_i^T = \sigma_1 u_1 \cdot v_1^T + \sigma_2 u_2 \cdot v_2^T + \cdots + \sigma_p u_p \cdot v_p^T,$$

where $p = \min\{m, n\}$, and once again, the singular values are ordered such that $\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_p \geq 0$. Each term $u_i \cdot v_i^T$ is a rank-one matrix. The following will help you understand the expansion.

Exercises or Facts:

- $A \cdot A^\top = U \cdot \Sigma \cdot \Sigma^\top \cdot U^\top = \sum_{i=1}^p (\sigma_i)^2 u_i \cdot u_i^\top$
- $A^\top \cdot A = V \cdot \Sigma^\top \cdot \Sigma \cdot V^\top = \sum_{i=1}^p (\sigma_i)^2 v_i \cdot v_i^\top$
- $(u_i \cdot v_i^\top) \cdot v_j = \begin{cases} u_i & j = i \\ 0 & j \neq i \end{cases}$ and hence $\text{rank}(u_i \cdot v_i^\top) = 1$ and $\text{nullity}(u_i \cdot v_i^\top) = m - 1$
- $(u_i \cdot u_i^\top) \cdot u_j = \begin{cases} u_i & j = i \\ 0 & j \neq i \end{cases}$ and hence $\text{rank}(u_i \cdot u_i^\top) = 1$ and $\text{nullity}(u_i \cdot u_i^\top) = n - 1$
- $(v_i \cdot v_i^\top) \cdot v_j = \begin{cases} v_i & j = i \\ 0 & j \neq i \end{cases}$ and hence $\text{rank}(v_i \cdot v_i^\top) = 1$ and $\text{nullity}(v_i \cdot v_i^\top) = m - 1$
- $v_i \cdot v_i^\top$, and $u_i \cdot u_i^\top$ have e-values $\lambda_1 = 1$ distinct and $\lambda_2 = 0$ repeated $m - 1$ and $n - 1$ times, respectively.
- **Hint:** $(u_i \cdot v_i^\top) \cdot v_j = u_i \cdot (v_i^\top \cdot v_j) = \begin{cases} u_i & j = i \\ 0 & j \neq i \end{cases}$ because the $\{v_1, v_2, \dots, v_m\}$ are orthonormal.

So far, we have only defined the norm of a vector. However, it is also useful to measure the “length” of matrices.

Def. (Induced Matrix Norm) Given an $n \times m$ real matrix A , the **matrix norm induced by the Euclidean vector norm** is given by:

$$\|A\| := \max_{x^\top x=1} \|Ax\| = \sqrt{\lambda_{\max}(A^\top A)}$$

where $\lambda_{\max}(A^\top A)$ denotes the largest eigenvalue of the matrix $A^\top A$. (**Recall that the matrices of the form $A^\top A$ are at least positive semidefinite and hence their e-values are real and non-negative.**) Therefore, the square root exists.

Numerical Rank

Facts: Suppose that $\text{rank}(A) = r$, so that σ_r is the smallest non-zero singular value of A .

- If an $n \times m$ matrix E satisfies $\|E\| < \sigma_r$, then $\text{rank}(A + E) \geq r$.
- There exists an $n \times m$ matrix E with $\|E\| = \sigma_r$ and $\text{rank}(A + E) < r$.
- In fact, for $E = -\sigma_r u_r v_r^\top$, $\text{rank}(A + E) = r - 1$.
- Moreover, for $E = -\sigma_r u_r v_r^\top - \sigma_{r-1} u_{r-1} v_{r-1}^\top$, $\text{rank}(A + E) = r - 2$.

Corollary: Suppose A is square and invertible. Then σ_r measures the distance from A to the nearest singular matrix.

Illustration Continued

```

1 u5=U[:,5]; v5=V[:,5]; sig5=Sigma[5]
2 E=-sig5*u5*v5'
3 # Induced Norm
4 M=E'*E
5 SquareRootEigs=(abs.(eigvals(E'*E))).^0.5
6 #
7 (U ,Sigma2, V) = svd(A+E)

```

$$E = \begin{bmatrix} 4.169e-02 & -1.212e-01 & -9.818e-02 & 2.189e-01 & -5.458e-02 \\ -2.031e-02 & 5.906e-02 & 4.784e-02 & -1.066e-01 & 2.659e-02 \\ -1.966e-02 & 5.716e-02 & 4.629e-02 & -1.032e-01 & 2.574e-02 \\ -7.041e-02 & 2.048e-01 & 1.658e-01 & -3.697e-01 & 9.220e-02 \\ 8.160e-02 & -2.373e-01 & -1.922e-01 & 4.284e-01 & -1.068e-01 \end{bmatrix}$$

$$\sqrt{\lambda_i(E^\top \cdot E)} = \begin{bmatrix} 7.376e-09 \\ 2.406e-09 \\ 1.977e-09 \\ 4.163e-09 \\ 7.916e-01 \end{bmatrix}$$

$$\Sigma_2 = \begin{bmatrix} 1.325e+02 & 0.000e+00 & 0.000e+00 & 0.000e+00 & 0.000e+00 \\ 0.000e+00 & 3.771e+01 & 0.000e+00 & 0.000e+00 & 0.000e+00 \\ 0.000e+00 & 0.000e+00 & 3.342e+01 & 0.000e+00 & 0.000e+00 \\ 0.000e+00 & 0.000e+00 & 0.000e+00 & 1.934e+01 & 0.000e+00 \\ 0.000e+00 & 0.000e+00 & 0.000e+00 & 0.000e+00 & 1.775e-15 \end{bmatrix}$$

We added a matrix with norm 0.7916 and made the (exact) rank drop from 4 to 5! How cool is that? This example shows that SVD can exactly measure how close a matrix is to being singular. We also see that $E^\top \cdot E$ has rank one: there is one non-zero e-value and the rest are (essentially) zero as the theory promised.

Other Interesting and Useful Facts

- (a) **Null space:** $\text{null}(A) := \{x \in \mathbb{R}^m \mid Ax = 0\}$
- (b) **Range:** $\text{range}(A) := \{y \in \mathbb{R}^n \mid \text{such that } y = Ax \text{ for some } x \in \mathbb{R}^m\}$
- (c) **Fact:** Suppose $A = U \cdot \Sigma \cdot V^\top$. Then the columns of U corresponding to non-zero singular values are a basis for $\text{range}(A)$ and the columns of V corresponding to zero singular values are a basis for $\text{null}(A)$, viz

$$\begin{aligned} \text{range}(A) &:= \text{span}\{u_1, \dots, u_r\}, \text{ and} \\ \text{null}(A) &:= \text{span}\{v_{r+1}, \dots, v_m\}. \end{aligned}$$

- (d) The SVD can also be used to compute an “effective” range and an “effective” null space of a matrix.
- (e) **Fact:** Suppose that $\sigma_1 \geq \dots \geq \sigma_r > \delta \geq \sigma_{r+1} \geq \dots \geq \sigma_p \geq 0$, so that r is the “effective” or “numerical rank” of A . (Note the δ inserted between σ_r and σ_{r+1} to denote the break point.)
- (f) **Fact:** Let $\text{range}_{\text{eff}}(A)$ and $\text{null}_{\text{eff}}(A)$ denote the effective range and effective null space of A , respectively. Then we can calculate bases for these subspaces by choosing appropriate singular vectors:

$$\begin{aligned} \text{range}_{\text{eff}}(A) &:= \text{span}\{u_1, \dots, u_r\}, \text{ and} \\ \text{null}_{\text{eff}}(A) &:= \text{span}\{v_{r+1}, \dots, v_m\}. \end{aligned}$$

A.5 Linear Transformations and Matrix Representations

Def. A function $L : \mathbb{R}^m \rightarrow \mathbb{R}^n$ is a **linear transformation** if for all $x, z \in \mathbb{R}^m$, $\alpha, \beta \in \mathbb{R}$,

$$L(\alpha x + \beta z) = \alpha L(x) + \beta L(z). \tag{A.46}$$

Just as with checking the subspace property, one can break the condition (A.46) into two separate properties,

$$\begin{aligned} L(x + z) &= L(x) + L(z) \\ L(\alpha x) &= \alpha L(x). \end{aligned}$$

You already know at least one linear transformation from \mathbb{R}^m to \mathbb{R}^n , namely, let A be an $n \times m$ real matrix and for $x \in \mathbb{R}^m$, define $L : \mathbb{R}^m \rightarrow \mathbb{R}^n$ by

$$L(x) = Ax. \quad (\text{A.47})$$

It is straightforward to show that (A.46) holds and thus we leave that to you. If you are trying to think of a clever linear transformation that is not built from a matrix, but you cannot come up with one, well, there is a reason for that: there are none!

Linear Transformations from \mathbb{R}^m to \mathbb{R}^n are Kind of Boring

Fact: Let $L : \mathbb{R}^m \rightarrow \mathbb{R}^n$ be a linear transformation. Then there exists an $n \times m$ real matrix A such that, for every $x \in \mathbb{R}^m$

$$L(x) = Ax.$$

The matrix A has a name: it is called the **matrix representation of L** . It's computation is relatively easy. Let $\{e_1, e_2, \dots, e_m\}$ be the natural basis vectors for \mathbb{R}^m . Define

$$a_j^{\text{col}} := L(e_j) \quad \text{and} \quad A := [a_1^{\text{col}} \quad a_2^{\text{col}} \quad \cdots \quad a_m^{\text{col}}] = [L(e_1) \quad L(e_2) \quad \cdots \quad L(e_m)].$$

Then $a_j^{\text{col}} \in \mathbb{R}^n$ because $L(x) \in \mathbb{R}^n$ for all $x \in \mathbb{R}^m$. Now, because $\{e_1, e_2, \dots, e_m\}$ is a basis, we have

$$x \in \mathbb{R}^m \iff x = x_1e_1 + x_2e_2 + \cdots + x_me_m.$$

Because L is a linear transformation,

$$\begin{aligned} L(x) &= L(x_1e_1 + x_2e_2 + \cdots + x_me_m) \\ &= x_1L(e_1) + x_2L(e_2) + \cdots + x_mL(e_m) \\ &= x_1a_1^{\text{col}} + x_2a_2^{\text{col}} + \cdots + x_ma_m^{\text{col}} \\ &= Ax. \end{aligned}$$

We will give you just a hint that **there are interesting linear transformations**, but to do that, we need to build an interesting vector space, the set of polynomials of degree less than or equal to n , namely

$$P_n(t) := \{a_0 + a_1t + a_2t^2 + \cdots + a_nt^n \mid a_i \in \mathbb{R}, 0 \leq i \leq n\}.$$

We note that everything in $P_n(t)$ appears to be a linear combination of the set of monomials, $\{1, t, t^2, \dots, t^n\}$. Indeed, if you take a bit more Linear Algebra, such as Math 217, you can make sense of the monomials as being vectors, and not just any vectors, but **basis vectors** for $P_n(t)$. Hence,

$$P_n(t) = \text{span}\{1, t, t^2, \dots, t^n\}.$$

It is clear that if you add any two polynomials of degree less than or equal to n , you get another polynomial of degree less than or equal to n . If you multiply a polynomial of degree less than or equal to n by a real number, you get another polynomial of degree less than or equal to n . This tells you that $P_n(t)$ satisfies all the properties of being a **vector space**: it is **closed under linear combinations**! And yes, in abstract Linear Algebra, we call the elements of $P_n(t)$ vectors.

We define $L : P_n(t) \rightarrow P_n(t)$ by $L(p(t)) := \frac{dp(t)}{dt}$, the first derivative of the polynomial $p(t)$, that is, for those of you who have taken Calculus I,

$$L(a_0 + a_1t + a_2t^2 + \cdots + a_nt^n) := a_1 + 2a_2t + 3a_3t^2 + \cdots + na_nt^{n-1}. \quad (\text{A.48})$$

The rules of differentiation imply that L satisfies (A.46), that is, the rules of being a linear transformation. We note that L defined in (A.48) does not look like it comes from a matrix.

Remark: Bummer! Since L does not appear to come from a matrix, does that mean that we cannot apply to it any of the computational tools that we have developed in ROB 101? The emphatic answer is: we absolutely can apply them! How? We'll show below that there is a matrix hiding inside of L ! While this is quite a bit beyond the scope of ROB 101, we feel that it might provide additional motivation for you to take a more advanced Linear Algebra course!

Definition: Let $\{v\} := \{v_1, v_2, \dots, v_k\}$ be a basis for a (real) vector space V . We know that for any vector $x \in V$, there exist² real coefficients $\alpha_1, \alpha_2, \dots, \alpha_k$ such that

$$x = \alpha_1 v_1 + \alpha_2 v_2 + \dots + \alpha_k v_k. \quad (\text{A.49})$$

The column vector built by expressing x as a linear combination of the basis vectors in $\{v\}$,

$$\begin{bmatrix} \alpha_1 \\ \alpha_2 \\ \vdots \\ \alpha_k \end{bmatrix} \in \mathbb{R}^k, \quad (\text{A.50})$$

is called the **representation of x with respect to the basis** $\{v_1, v_2, \dots, v_k\}$. A nice notation for it is

$$[x]_{\{v\}} := \begin{bmatrix} \alpha_1 \\ \alpha_2 \\ \vdots \\ \alpha_k \end{bmatrix} \in \mathbb{R}^k. \quad (\text{A.51})$$

The representation of x in V is a vector in \mathbb{R}^k .

Let's apply this idea to the vector space $V := P_n(t)$ and its very nice basis

$$\{v\} = \{1, t, t^2, \dots, t^n\}.$$

We noted earlier that any vector in " $x \in V$ " (that is, $p(t) \in P_n(t)$) can be written as

$$x = p(t) = a_0 + a_1 t + a_2 t^2 + \dots + a_n t^n,$$

and hence,

$$[x]_{\{v\}} := \begin{bmatrix} a_0 \\ a_1 \\ a_2 \\ \vdots \\ a_n \end{bmatrix} \in \mathbb{R}^{n+1}. \quad (\text{A.52})$$

It's kind of obvious, we add two polynomials of the same degree by adding their coefficients, and (A.52) is simply enticing us to do just that. [In other words, do not overthink it!]

Hence, if polynomials of degree less than or equal to n can be represented by columns of numbers, can differentiation be represented by a matrix? The answer is yes. As in (A.48), we define $L : V \rightarrow V$ by if $x = a_0 + a_1 t + a_2 t^2 + \dots + a_n t^n \in V$,

$$L(x) := a_1 + 2a_2 t + 3a_3 t^2 + \dots + n a_n t^{n-1}. \quad (\text{A.53})$$

We now break this up and apply L to each of the basis elements

$$\{v\} := \{v_1 = 1, v_2 = t, v_3 = t^2, \dots, v_{n+1} = t^n\},$$

yielding

$$L(v_1) = 0, L(v_2) = v_1, L(v_3) = 2v_2, \dots, L(v_{n+1}) = n v_n, \quad (\text{A.54})$$

and note that

$$[L(v_1)]_{\{v\}} = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \\ 0 \\ 0 \end{bmatrix}, [L(v_2)]_{\{v\}} = \begin{bmatrix} 1 \\ 0 \\ \vdots \\ 0 \\ 0 \\ 0 \end{bmatrix}, [L(v_3)]_{\{v\}} = \begin{bmatrix} 0 \\ 2 \\ \vdots \\ 0 \\ 0 \\ 0 \end{bmatrix}, \dots, [L(v_n)]_{\{v\}} = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ n-1 \\ 0 \\ 0 \end{bmatrix}, [L(v_{n+1})]_{\{v\}} = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ n \\ n \\ 0 \end{bmatrix}. \quad (\text{A.55})$$

²The coefficients are actually unique. You should you prove that if you have two such sets of coefficients that they have to be equal. It follows from the definition of linear Independence!

We use the above vectors as the columns of a matrix A , where $a_j^{\text{col}} := [L(v_j)]_{\{v\}}$,

$$A = \begin{bmatrix} 0 & 1 & 0 & \cdots & 0 & 0 \\ 0 & 0 & 2 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & n-1 & 0 \\ 0 & 0 & 0 & \cdots & 0 & n \\ 0 & 0 & 0 & \cdots & 0 & 0 \end{bmatrix}. \quad (\text{A.56})$$

We next note that for $x = a_0 + a_1t + \cdots + a_{n-2}t^{n-2} + a_{n-1}t^{n-1} + a_nt^n \in V$, its representation is given in (A.52) and that

$$\begin{bmatrix} 0 & 1 & 0 & \cdots & 0 & 0 \\ 0 & 0 & 2 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & n-1 & 0 \\ 0 & 0 & 0 & \cdots & 0 & n \\ 0 & 0 & 0 & \cdots & 0 & 0 \end{bmatrix} \begin{bmatrix} a_0 \\ a_1 \\ \vdots \\ a_{n-2} \\ a_{n-1} \\ a_n \end{bmatrix} = \begin{bmatrix} a_1 \\ 2a_2 \\ \vdots \\ (n-1)a_{n-1} \\ na_n \\ 0 \end{bmatrix}; \quad (\text{A.57})$$

moreover, the right hand side of the equation is the representation of

$$\frac{d}{dt} (a_0 + a_1t + \cdots + a_{n-2}t^{n-2} + a_{n-1}t^{n-1} + a_nt^n) = a_1 + 2a_2t + \cdots + (n-1)a_{n-1}t^{n-2} + na_nt^{n-1}$$

with respect to the monomials. In other symbols,

$$A[x]_{\{v\}} = [L(x)]_{\{v\}}. \quad (\text{A.58})$$

A is called the **matrix representation of L with respect to the basis $\{v\}$** . In Robotics, we often need to differentiate signals in real-time on our robots. We do it using versions of (A.57) and (A.58), which transform differentiation into matrix multiplication!

A.6 Affine Transformations

All functions $f : \mathbb{R}^m \rightarrow \mathbb{R}^n$ that fail (A.46) are nonlinear! That does not seem very discerning. There are a few more classes of functions called out, and one in particular is very important in Linear Algebra is **Definition:** A function $f : \mathbb{R}^m \rightarrow \mathbb{R}^n$ is **affine** if there exists a constant vector $b \in \mathbb{R}^n$ and an $n \times m$ real matrix A such that

$$f(x) = Ax + b. \quad (\text{A.59})$$

A bit more abstract definition is $f : \mathbb{R}^m \rightarrow \mathbb{R}^n$ is **affine** if there exists a constant vector $b \in \mathbb{R}^n$ such that the function $L : \mathbb{R}^m \rightarrow \mathbb{R}^n$ by

$$L(x) := f(x) - b$$

is linear. When dealing with \mathbb{R}^n and \mathbb{R}^m , the two definitions are equivalent, while the more abstract definition extends to more general settings.

Appendix B

What is an Ordinary Differential Equation?

Learning Objectives

- Provide motivation for considering equations with derivatives.
- Provide a numerical means to approximate a solution to a differential equation.
- Prepare you a bit for Math 216.

Outcomes

- See $F = ma$ in a new light
- Define an ordinary differential equation, aka an ODE
- Learn how to pass from a differential equation to a difference equation
- Relate the process of iterating a difference equation to the process of solving an ODE

Equations with a derivative in them are called **differential equations**. In Robotics, we use differential equations to understand the motion of robots, such as Cassie Blue, or in the case of Project #3, a Segway. The topic of differential equations is typically delayed until a fourth-semester Calculus course. Moreover, such courses focus almost exclusively on closed-form solutions to differential equations. As you can imagine, we're not a big fan of that. Here, you will see that tools we have developed so far in ROB 101 allow us to develop elementary numerical tools for analyzing the solutions of interesting differential equations.

B.1 Preliminaries: Expanding our Concept of an Equation

Equations in a scalar variable $x \in \mathbb{R}$ probably seem pretty trivial to you by now. Examples we have seen include:

- Linear equation: $ax = b$;
- Quadratic equation: $ax^2 + bx + c = 0$;
- Cubic equation: $ax^3 + bx^2 + cx + d = 0$; or
- Trigonometric equation: $A \cos(x) + B \sin(x) + C = 0$.

One obvious way to increase the generality of our analysis tools for equations is to include several variables, $x \in \mathbb{R}^n$, $n > 1$, and we've done a bunch of that in ROB 101 as well, where we studied how to find solutions to

- System of linear equations: $Ax = b$; and
- System of nonlinear equations: $F(x) = 0$.

In the case of linear systems, we explored map building and linear regression (easy version of Machine Learning) as cool application problems. For nonlinear systems, we looked at the position of a robot gripper in \mathbb{R}^2 and sought values for the angles of its joints so that the gripper could be placed in a desired position.

Another way to increase the generality of our analysis tools is to expand our concept of equations to include time as a variable! This is a big conceptual step: instead of the values in our equation depending on constants or other powers of our variable x , as in the examples we listed above, we will have the value of our variable x at time t depend on x at some other time, say $t - \delta t$, for some $\delta t > 0$.

B.2 Time in a Digital Computer is Discrete

To get our heads around functions that depend on time, we'll start with "time" as it is treated in a digital computer, namely, time is a sequential variable, such as

$$k = 1, 2, 3, \dots \quad (\text{B.1})$$

Digital time is similar to our *human* notion of time in that it is strictly increasing, but whereas our human notion of time is continuous (it can be divided into ever smaller increments as long as we ignore the rules of Quantum Mechanics), *digital* time is fundamentally discrete in that it increases by non-zero increments.

Suppose we denote the value of our variable x at time k by $x[k]$, and we define an *equation for x at the next time, $k + 1$* , by

$$x[k + 1] = \frac{1}{2}x[k] + 1. \quad (\text{B.2})$$

Equation (B.2) says that if we know the value of our variable x at time k , we can find the value of x at time $k + 1$ by multiplying $x[k]$ by one-half and adding one to it. That seems easy enough. OK, then, what is the value of x at time $k = 4$?

To compute $x[4]$, we need to know $x[3]$. To compute $x[3]$, we need to know $x[2]$. To compute $x[2]$, we need to know $x[1]$. If we suppose that time starts at $k = 1$, then we cannot go back any further and we realize that somehow $x[1]$ must be given to us!




Figure B.1: Plot of the function $x : [1, 2, \dots, 20] \rightarrow \mathbb{R}$ at discrete instances of time.




Figure B.2: Plot of the function $x : [1, 2, \dots, 20] \rightarrow \mathbb{R}$, with the “dots” connected, giving us the impression that time is almost continuous.

Equation (B.2) is called a **difference equation** and $x[1]$ is called its **initial value**. If we are given that $x[1] = 4$, we’d then compute that

$$\begin{aligned}
 x[1] &= \frac{1}{2}x[0] + 1 = 3 \\
 x[2] &= \frac{1}{2}x[1] + 1 = \frac{7}{2} \\
 x[3] &= \frac{1}{2}x[2] + 1 = \frac{11}{4} \\
 x[4] &= \frac{1}{2}x[3] + 1 = \frac{19}{8} \\
 &\vdots = \quad \quad \quad \vdots \\
 x[k+1] &= \frac{1}{2}x[k] + 1
 \end{aligned}$$

In Fig. B.1, we plot the function x for $k = 1 : 20$. We observe that it seems to converge to 2.0, which is intriguing. In Fig. B.2, we have “cheated” and made time look quasi-continuous by “connecting the dots”. The Julia code for generating x and the plots is given in Sec. B.6.1.

It would seem that if we could put the “dots closer together”, the effect in Fig. B.2 would be even better. Let’s try!

B.3 Digital Time may be Discrete, but We Can Make the Time Increment δt Quite Small

In our previous model and plot, we implicitly took $\delta t = 1$ “unit” of time, where “unit” could have been seconds, milliseconds, months, or fortnights. We just plotted the index k and had no idea how it was directly related to a physical notion of “time”. This

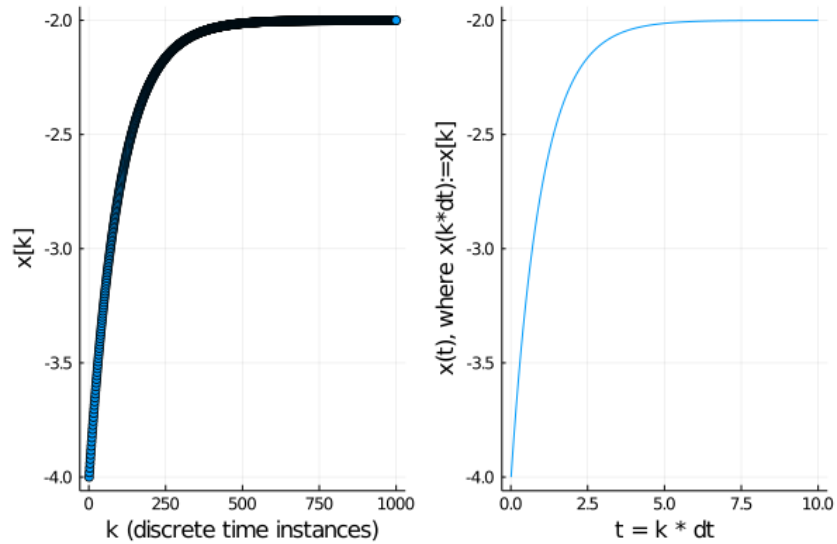


Figure B.3: Plot of the function $x : [0, 10] \rightarrow \mathbb{R}$ that solves (B.3) when its initial condition is $x[1] = 4$. The discrete interval of time $\delta t = 0.01$ is small enough that it looks continuous, and in fact, if one connects the dots, our function versus time looks like a continuous curve.

time we'll define $\delta t = 0.01$ seconds, that is, 10 milliseconds, and we'll think of $x[k]$ as representing $x(k\delta t)$. To be super explicit,

$$\begin{aligned}
 x[1] &:= x(t)|_{t=\delta t} = x(0.01) \\
 x[2] &:= x(t)|_{t=2\delta t} = x(0.02) \\
 x[3] &:= x(t)|_{t=3\delta t} = x(0.03) \\
 &\vdots \\
 x[99] &:= x(t)|_{t=99\delta t} = x(0.99) \\
 &\vdots \\
 x[301] &:= x(t)|_{t=301\delta t} = x(3.01) \\
 &\text{etc.}
 \end{aligned}$$

We introduce δt into the model in a particular manner that will become clear in Sec. B.4,

$$x[k + 1] = x[k] - \delta t \cdot (x[k] + 2). \tag{B.3}$$

Fig B.3 shows two plots side by side. One shows the function x plotted against the index k , while the second plot shows x plotted against time, for $t = k\delta t$. The Julia code for computing the function and making the plots is given in Sec. B.6.2.

B.4 Equations with Derivatives in them are called Differential Equations

All of us hear about “ $F = ma$ ” in High School, one of Newton’s laws for how a body of mass m accelerates under the action of applied forces F . We might also learn that “acceleration”, a , is the rate of change of “velocity”, v , with respect to time, t .

The notation $\frac{d}{dt}$

In Calculus, the rate of change of one quantity, say v , with respect to another quantity, say t , is denoted

$$\frac{dv(t)}{dt}.$$

Hence, in the language of Calculus, acceleration is related to velocity by

$$a(t) := \frac{dv(t)}{dt},$$

which, once again, is just shorthand for “acceleration is the rate of change of velocity with respect to time”. This course does not assume knowledge of Calculus, so please don’t sweat the introduction of this notation. For us, it is just a shorthand way of saying “rate of change with respect to time”.

Using this shorthand notation, we can express Newton’s Law as

$$m \frac{dv(t)}{dt} = F(t). \quad (\text{B.4})$$

If we suppose that the body is falling in air, we might model the total force $F(t)$ acting on the body as being due to gravity and air friction,

$$F(t) = -mg - k_d v(t).$$

Doing so leads us to the equation

$$\boxed{m \frac{dv(t)}{dt} = -k_d v(t) - g}. \quad (\text{B.5})$$

This equation says that the rate of change of the velocity as a function of time is given by a sum of two terms, one of which corresponds to gravity and the other to air resistance. Equation (B.5) is called a **differential equation**, that is, an equation that depends on “derivatives.”

In (B.5), let’s now replace the shorthand Calculus symbol for rate of change, $\frac{dv(t)}{dt}$, with the same kind of numerical approximation we used in our studies of root finding and optimization, namely

$$\frac{dv(t)}{dt} \approx \frac{v(t + \delta t) - v(t)}{\delta t}. \quad (\text{B.6})$$

Substituting (B.6) into (B.5) and assuming the approximation is “good enough” (that we can change \approx into $=$) give

$$\begin{aligned} \frac{dv(t)}{dt} &= \frac{1}{m} (-k_d v(t) - g) \\ \Downarrow \\ \frac{v(t + \delta t) - v(t)}{\delta t} &= \frac{1}{m} (-k_d v(t) - g) \\ \Downarrow \\ v(t + \delta t) - v(t) &= \delta t \frac{1}{m} (-k_d v(t) - g) \\ \Downarrow \\ v(t + \delta t) &= v(t) + \delta t \frac{1}{m} (-k_d v(t) - g). \end{aligned} \quad (\text{B.7})$$

If we then define $t = k \cdot \delta t$ and $v[k] := v(k \cdot \delta t)$, the equation looks just like our *difference equations* in Sec. B.3. Indeed, we have $v(t + \delta t) = v(k \cdot \delta t + \delta t) = v((k + 1) \cdot \delta t) = v[k + 1]$, and (B.7) becomes

$$v[k + 1] = v[k] - \delta t \frac{k_d}{m} v[k] - \delta t \frac{g}{m}, \quad (\text{B.8})$$

which is very much like (B.3) and hence, we now know one way to (approximately) solve the differential equation (B.5): we set an initial condition and iterate based on (B.8).

In fact, if the mass of our falling body is $m = \frac{2}{g}$ and its drag is $k_d = m$, then (B.8) is exactly the same as (B.3). Moreover, we can physically interpret the solution of the differential equation (B.5), which is plotted in Fig. B.3, as a package is tossed out of plane. Our model tracks its evolution from the moment the parachute deploys, greatly increasing its drag, thereby slowing its velocity from an initial value of -4 to -2 units of distance per second (we never specified the units).

B.5 Discretization of ODEs of Higher Dimension

ODE

ODE is short for Ordinary Differential Equation. The word “ordinary” is used because there are other kinds of differential equations, which apparently, are less “ordinary”. Everyone in the know uses the terminology ODE, hence you will too!

ODEs can be vector valued too. A linear ODE may look like this,

$$\frac{dx(t)}{dt} = Ax(t) + b.$$

Just as before, we select $\delta t > 0$ and define $x[k] := x(k \cdot \delta t)$, and re-write the differential equation as a difference equation

$$\begin{aligned} \frac{dx(t)}{dt} &\approx \frac{x(t + \delta t) - x(t)}{\delta t} \\ &\Downarrow \\ \frac{x(t + \delta t) - x(t)}{\delta t} &\approx Ax(t) + b \\ &\Downarrow \\ x(t + \delta t) &= x(t) + \delta t (Ax(t) + b) \\ &\Downarrow \\ x[k + 1] &= x[k] + \delta t Ax[k] + \delta t b \end{aligned}$$

A two-dimensional example is

$$\underbrace{\begin{bmatrix} x_1[k + 1] \\ x_2[k + 1] \end{bmatrix}}_{x[k+1]} = \underbrace{\begin{bmatrix} x_1[k] \\ x_2[k] \end{bmatrix}}_{x[k]} + \delta t \cdot \underbrace{\begin{bmatrix} 0.0 & 1.0 \\ -2.0 & -1.0 \end{bmatrix}}_A \underbrace{\begin{bmatrix} x_1[k] \\ x_2[k] \end{bmatrix}}_{x[k]} + \delta t \cdot \underbrace{\begin{bmatrix} 0.0 \\ 1.0 \end{bmatrix}}_b. \quad (\text{B.9})$$

This time, our $x[1]$ is a 2×1 vector. We’ll arbitrarily specify it as $x[1] = [2 \ 0]^\top$. The plots of x as a function of index k and of time $t = k \cdot \delta t$ are given in Fig. B.5. The associated Julia code is given in Sec. B.6.4.

Finally, not only can the function x of time can be vector valued, it can have nonlinear terms. Here is a common example from physics: a pendulum of mass m and length ℓ swinging from a frictionless pivot satisfies the equation

$$\begin{aligned} \frac{dx_1(t)}{dt} &= x_2(t) \\ \frac{dx_2(t)}{dt} &= -\frac{g}{\ell} \sin(x_1(t)), \end{aligned} \quad (\text{B.10})$$

where x_1 is the angle of the pendulum and x_2 is its angular velocity. We’ll rewrite the model in vector form by defining

$$x := \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}, \quad (\text{B.11})$$




Figure B.4: Plot of the function $x : [0, 10] \rightarrow \mathbb{R}$ that solves (B.3) when its initial condition is $x[1] = 4$. The discrete interval of time $\delta t = 0.01$ is small enough that it looks continuous, and in fact, if one connects the dots, our function versus time looks like a continuous curve.

so that

$$\frac{dx(t)}{dt} = \begin{bmatrix} \frac{dx_1(t)}{dt} \\ \frac{dx_2(t)}{dt} \end{bmatrix} := \underbrace{\begin{bmatrix} x_2(t) \\ -\frac{g}{\ell} \sin(x_1(t)) \end{bmatrix}}_{f(x(t))} =: f(x(t)) \quad (\text{B.12})$$

To find a solution, we discretize the model with $\delta t > 0$. To show the flexibility we have in approximating the derivative, we'll use a symmetric difference approximation this time, namely

$$\begin{aligned} \frac{dx(t)}{dt} &\approx \frac{x(t + \delta t) - x(t - \delta t)}{2\delta t} \quad \text{and} \quad \frac{dx(t)}{dt} = f(x(t)) \\ &\Updownarrow \\ x(t + \delta t) &= x(t - \delta t) + 2\delta t f(x(t)) \quad \text{and} \quad t = k \cdot \delta t \\ &\Updownarrow \\ x[k + 1] &= x[k - 1] + \delta t f(x[k]) \end{aligned} \quad (\text{B.13})$$

We use the last line of (B.13) to iterate and compute a solution to the pendulum ODE (B.10). The solution is plotted in Fig. B.6. The associated Julia code is given in Sec. B.6.5.

Our transformation of ODEs into difference equations are examples of numerical integration methods. When we use the approximation

$$\frac{dx(t)}{dt} \approx \frac{x(t + \delta t) - x(t)}{\delta t}$$

we end up with what is called a first-order forward Euler integration method. In general, it's a pretty bad way to solve ODEs, but for a first introduction, it is great! We can do a lot with it.




Figure B.5: Plot of the vector valued function $x : [0, 10] \rightarrow \mathbb{R}^2$ solving the ODE $\frac{dx(t)}{dt} = Ax + b$, with the solution approximated in (B.9). The discrete interval of time is small enough that it looks continuous!




Figure B.6: Plot of a vector valued function $x : [0, 10] \rightarrow \mathbb{R}^2$ that oscillates back and forth because it corresponds to the solution the pendulum ODE given in (B.10). The blue lines are the angle of the pendulum and the brown(ish) lines are its angular velocity.

B.6 Julia Code for Generating the Figures

We provide code in Julia 1.41 for the figures appearing in this Appendix.

B.6.1 For Figures B.1 and B.2

```
1
2 # Model
3 dt=0.01
4 T=10
5 N=floor(Int, T/dt);
6 K=1:N
7 time = K*dt
8 x=zeros(N, 1)
9 x[1]=-4
10 for k=1:(N-1)
11     x[k+1]=(1-dt)*x[k]-2*dt
12 end
13 # Plotting
14 plot1=scatter(K, x, xlabel="k (discrete time instances)", ylabel="x[k]", leg=false)
15 plot2=plot(time, x, xlabel="t = k * dt", ylabel="x(t), where x(k*dt) :=x[k]", leg=false)
16 plot(plot1, plot2, layout = (1, 2), legend = false)
17 #Turn the plot into an image that one can copy
18 plot!(fmt = :png)
```

B.6.2 Code for Figure B.3

```
1 using Plots
2 gr()
3 # Model
4 N=20
5 time=1:N
6 x=zeros(N, 1)
7 x[1]=4
8 for k=1:(N-1)
9     x[k+1]=0.5*x[k]+1
10 end
11 # Plotting
12 scatter(time, x)
13 #plot!(xlabel="k (discrete time instances)", ylabel="x[k]",
14 #title="Plot of as a function of time: Discrete-time Model", leg=false)
15 plot!(xlabel="k (discrete time instances)", ylabel="x[k]", leg=false)
16 #Turn the plot into an image that one can copy
17 plot!(fmt = :png)
18 #
19 # Second plot
20 #
21 plot(time, x)
22 #plot!(xlabel="k (discrete time instances)", ylabel="x[k]",
23 #title="Plot of as a function of time: Discrete-time Model", leg=false)
24 plot!(xlabel="k (discrete time instances)", ylabel="x[k]", leg=false)
25 scatter!(time, x)
26 #Turn the plot into an image that one can copy
27 plot!(fmt = :png)
```

B.6.3 Code for Figure B.4

```
1 # Model
2 dt=0.01
3 T=10
4 N=floor(Int, T/dt);
5 K=1:N
6 time = K*dt
7 x=zeros(N, 1)
8 x[1]=4
9 for k=1:(N-1)
10     x[k+1]=(1-dt)*x[k]+2*dt
11 end
12 # Plotting
13 plot1=scatter(K, x, xlabel="k (discrete time instances)", ylabel="x[k]", leg=false)
14 plot2=plot(time, x, xlabel="t = k * dt", ylabel="x(t), where x(k*dt):=x[k]", leg=false)
15 plot(plot1, plot2, layout = (1, 2), legend = false)
16 #Turn the plot into an image that one can copy
17 plot!(fmt = :png)
```

B.6.4 Code for Figure B.5

```
1 # Model
2 dt=0.01
3 T=10
4 N=floor(Int, T/dt);
5 K=1:N
6 time = K*dt
7 A=[0 1; -2 -1]
8 b=[0;1]
9 x=zeros(N, 2)
10 x[1, :]=[2 0]
11 for k=1:(N-1)
12     x[k+1, :]=x[k, :] + dt*A*x[k, :] + dt*b
13 end
14 # Plotting
15 plot1=scatter(K, x, xlabel="k (discrete time instances)", ylabel="x[k]", leg=false)
16 plot2=plot(time, x, xlabel="t = k * dt", ylabel="x(t), where x(k*dt):=x[k]", leg=false)
17 plot(plot1, plot2, layout = (1, 2), legend = false)
18 #Turn the plot into an image that one can copy
19 plot!(fmt = :png)
```

B.6.5 Code for Figure B.6

```
1 # Model
2 dt=.01
3 T=50
4 N=floor(Int, T/dt);
5 K=1:N
6 time = K*dt
7 g=9.81 #m/s^2
8 l=4 # m
9 x=zeros(N, 2)
```

```

10 #Pendulum equations
11 f (x1, x2)=[x2; -(g/l) *sin (x1) ]
12 #Initial conditions
13 x[1, :]=[pi/4 0]
14 x[2, :]=x[1, :] + dt*f(x[1, 1], x[1, 2])
15 #Euler integration based on Symmetric Difference for dx/dt
16 for k=2:(N-1)
17 x[k+1, :]=x[k-1, :] +2*dt*f(x[k, 1], x[k, 2])
18 end
19 # Plotting
20 plot1=scatter (K, x, xlabel="k (discrete time instances)", ylabel="x[k]", leg=false)
21 plot2=plot (time, x, xlabel="t = k * dt", ylabel="x(t), where x(k*dt) :=x[k]", leg=false)
22 plot (plot1, plot2, layout = (1, 2), legend = false)
23 #Turn the plot into an image that one can copy
24 plot! (fmt = :png)

```

**From here on, the book is a work in progress.
Please pardon our mess during construction!**

Appendix C

Camera and LiDAR Models for Students of Robotics

Learning Objectives

- [JWG: ???].

Outcomes

- [JWG: ???]

C.1 Pinhole Camera Model

Camera calibration is the process of finding the quantities internal to the camera that affect the imaging process. Today's cheap camera lenses or camera production errors may introduce a lot of distortion to images. Precise camera calibration is important when we need to deal with 3D interpretation of images, reconstruction of world models, and robot interaction with the real world. We will apply what we have learnt so far to discover the charm of camera calibration and its usage in our project.

C.2 Preliminaries: Geometrical Transformations

C.2.1 Homogeneous Coordinates

Homogeneous coordinates are widely used in computer vision and graphics. They are a nice extension of standard 3D vectors and allow us to simplify and compute various vector operations, such as translation, rotation, scaling and perspective projection. The sequence of such operations can be multiplied out into a single matrix with simple and efficient processing. Besides, homogeneous coordinates also allow to represent infinity. The Euclidean coordinate system denotes the location of an object by a triplet of numbers. However, we can't treat infinity as a regular number. Homogeneous coordinates allows by denoting infinity by $[x, y, z, 0] = \frac{[x, y, z]}{0}$ in 3D.

How to transfer between Cartesian Coordinates and Homogeneous Coordinates?

- Cartesian coordinates \rightarrow Homogeneous coordinates : Multiply the coordinates by a non-zero scalar and add an extra coordinate equal to that scalar. For example, $[x, y] \rightarrow [x \cdot z, y \cdot z, z], z \neq 0$. We usually multiply the coordinates by 1.
- Homogeneous coordinates \rightarrow Cartesian coordinates: Divide the Cartesian coordinates by the last coordinate and eliminate it. For example, $[x, y, z], z \neq 0 \rightarrow [x/z, y/z]$.

C.2.2 2D Translation




Figure C.1: Example of 2D Translation

If point $\mathbf{P}'(x', y')$ is obtained by translating \mathbf{P} by $\mathbf{t}(t_x, t_y)$, then the relation between \mathbf{P}' and \mathbf{P} could be written as:

$$\mathbf{P}' = \begin{bmatrix} x + t_x \\ y + t_y \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} + \begin{bmatrix} t_x \\ t_y \end{bmatrix} = \mathbf{I} \cdot \mathbf{P} + \mathbf{t}$$

If we use homogeneous coordinates, $P = (x, y, 1)$, and $t = (t_x, t_y, 1)$. The the relation between P' and P becomes:

$$\mathbf{P}' = \begin{bmatrix} x + t_x \\ y + t_y \\ 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & t_x \\ 0 & 1 & t_y \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = \begin{bmatrix} \mathbf{I} & \mathbf{t} \\ 0 & 1 \end{bmatrix} \cdot \mathbf{P} = \mathbf{T} \cdot \mathbf{P}$$

C.2.3 2D Scaling




Figure C.2: Example of 2D Scaling

If point $\mathbf{P}'(x', y')$ is obtained by scaling \mathbf{P} by $\mathbf{s}(s_x, s_y)$, then the relation between \mathbf{P}' and \mathbf{P} could be written as:

$$\mathbf{P}' = \begin{bmatrix} s_x x \\ s_y y \end{bmatrix} = \begin{bmatrix} s_x & 0 \\ 0 & s_y \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = \mathbf{S}' \cdot \mathbf{P}$$

Similar before, the relation between \mathbf{P}' and \mathbf{P} could be expressed in homogeneous coordinates:

$$\mathbf{P}' = \begin{bmatrix} s_x x \\ s_y y \\ 1 \end{bmatrix} = \begin{bmatrix} s_x & 0 & 0 \\ 0 & s_y & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = \begin{bmatrix} \mathbf{S}' & 0 \\ 0 & 1 \end{bmatrix} \mathbf{P} = \mathbf{S} \cdot \mathbf{P}$$

C.2.4 2D Rotation




Figure C.3: Example of 2D Rotation

If point $\mathbf{P}'(x', y')$ is obtained by rotating \mathbf{P} counterclockwise by θ degree, then the relation between \mathbf{P}' and \mathbf{P} could be written as:

$$\mathbf{P}' = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = \mathbf{R}' \cdot \mathbf{P}$$

The relation between \mathbf{P}' and \mathbf{P} could be expressed in homogeneous coordinates:

$$\mathbf{P}' = \begin{bmatrix} \cos \theta & -\sin \theta & 0 \\ \sin \theta & \cos \theta & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = \begin{bmatrix} \mathbf{R}' & 0 \\ 0 & 1 \end{bmatrix} \mathbf{P} = \mathbf{R} \cdot \mathbf{P}$$

\mathbf{R} is an orthogonal matrix and it satisfies the following properties:

- $\mathbf{R} \cdot \mathbf{R}^\top = \mathbf{R}^\top \cdot \mathbf{R} = \mathbf{I}$
- $\det(\mathbf{R}) = 1$

C.2.5 Other 2D Geometrical Transformations in Homogeneous Coordinates

- Shear in x:
$$\begin{bmatrix} 1 & k_x & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

- Shear in y:
$$\begin{bmatrix} 1 & 0 & 0 \\ k_y & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

- Reflect about y:
$$\begin{bmatrix} 1 & 0 & 0 \\ k_y & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

- General Affine Transformation:
$$\begin{bmatrix} a & b & c \\ d & e & f \\ 0 & 0 & 1 \end{bmatrix}$$

If there are more than one type of geometric transformation, one can multiply the corresponding transformation matrix to get the final transformation matrix.

For example, if we want to find the transformation matrix after scaling P by s firstly, then rotating counterclockwise by θ degree, and finally translate by t. The position of P' could be compute by:

$$\begin{aligned} P' = \mathbf{T} \cdot \mathbf{R} \cdot \mathbf{S} \cdot P &= \begin{bmatrix} 1 & 0 & t_x \\ 0 & 1 & t_y \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \cos\theta & -\sin\theta & 0 \\ \sin\theta & \cos\theta & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} s_x & 0 & 0 \\ 0 & s_y & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} \\ &= \begin{bmatrix} \cos\theta & -\sin\theta & t_x \\ \sin\theta & \cos\theta & t_y \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} s_x & 0 & 0 \\ 0 & s_y & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} \\ &= \begin{bmatrix} R' & t \\ 0 & 1 \end{bmatrix} \begin{bmatrix} S & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} \\ &= \begin{bmatrix} R'S & t \\ 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} \end{aligned}$$

C.3 Pinhole Model

Last section we mainly discussed the affine transformation, which preserves the parallel lines, ratio of areas, ratio of lengths on colinear lines and a few others. However, the affine transformation is not a valid assumption in general for images, nor is it valid around the image edges. The mapping from 3D object to 2D image is a perspective projection as shown in the following image:




Figure C.4: Perspective schematic diagram



Figure C.5: Real-life perspective example: parallel rails

The **pinhole camera model** describes the mathematical relationship between the coordinates of a point in 3D space and its projection onto the image plane of an ideal pinhole camera. The following diagram illustrates the pinhole model and defines some intrinsic parameters of pinhole camera:




Figure C.6: Pinhole camera model

There are some key assumptions of the pinhole model:

- Aperture, i.e. the pinhole, is described as a point (point C in the image)
- No lenses are used to focus light, therefore no geometric distortions (radial distortion)
- Does not take into account that most practical cameras have only discrete image coordinates
- This model can only be used as a first order approximation of the mapping from 3D scene to 2D image. This is a perspective projection
- Ideal model that we use has identity matrix for the perspective portion of the projection matrix (**projection matrix** = **Perspective**|0) = $[\mathbf{I}_{3 \times 3} | 0]_{3 \times 4}$ which assumes no perspective since the surface is close to the camera)

If $\mathbf{P}(X, Y, Z)$ is a 3D scene point, and it is projected into the 2D image by pinhole. $\mathbf{P}'(x, y)$ is the projective point of \mathbf{P} on the 2D virtual image plane. We could find the following equations derived using similar triangles.

Remark: f is the focal length of the camera.




Figure C.7: Pinhole model

$$\frac{f}{Z} = \frac{y}{Y} = \frac{x}{X} \implies \begin{cases} x = \frac{fX}{Z} \\ y = \frac{fY}{Z} \end{cases} \quad (\text{C.1})$$

The image plane is usually read as $n \times m$ pixel matrix in computer. Points in the image is then represented by pixel coordinates.




Figure C.8: Relation of image plane and pixel matrix

The following picture is an example of image plane, which indicates the relation between camera coordinate origin and image coordinate origin




Figure C.9: Pixel coordinates plane

By expanding equation C.1, the relation between positions on the image scene (u, v) and positions in the real world (X, Y, Z) could be written as:

$$u = f s_x x^c + u_0 = \alpha \frac{X}{Z} + u_0$$

$$v = f s_y y^c + v_0 = \beta \frac{Y}{Z} + v_0$$

We assume the world frame and camera frame is aligned. s_x, s_y are the scale in x and y direction to transfer the units from metric to pixel. x^c, y^c are projective points of 3D objects in camera coordinates.

However, there are other considerations:

- Optical axis is not centered on sensor
- Pixel coordinate system origin is in the corner of the sensor
- Pixels aren't square
- Pixel size is unknown

Therefore, we need to find a transformation matrix to consider all the possible situations to transform the camera frame to the image plane. Review 2D geometric transformations in previous section and recall the general affine transformation $\begin{bmatrix} a & b & c \\ d & e & f \\ 0 & 0 & 1 \end{bmatrix}$. We will explore more about it in the next section.

C.4 Geometric and Mathematical Relations

We would like to represent image plane coordinates in terms of real world coordinates. This is exactly the same problem as finding the conversion of 3D world coordinates to image coordinates and then inverting this transformation. In this section, we will examine the three components of this transformation.

C.4.1 Intrinsic Matrix K

Intrinsic matrix is the affine transformation matrix we mentioned before to represent the linear properties of the pinhole camera model (focal length to visual image, scale, shear, translation, and so on). The intrinsic matrix can be used to transform 2D camera coordinates (x^c, y^c) into image plane coordinates (u, v) :

$$\begin{bmatrix} u \\ v \\ 1 \end{bmatrix} = K \begin{bmatrix} x^c \\ y^c \\ 1 \end{bmatrix} = \begin{bmatrix} fs_x & ks_y & u_0 \\ 0 & fs_y & v_0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x^c \\ y^c \\ 1 \end{bmatrix} = \begin{bmatrix} \alpha & s & u_0 \\ 0 & \beta & v_0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x^c \\ y^c \\ 1 \end{bmatrix}$$

Where s_x is the scale in x, s_y is the scale in y, f is the focal length, k is the ration of shear in the y direction to that in x, and (u_0, v_0) is the distance from the image center to the image plane coordinate system origin.

C.4.2 Projection Matrix

This matrix transforms the coordinates from 3D to 2D. The perspective portion of this matrix (left) is equal to the identity because we make the assumption that all captured positions are close.

$$\begin{bmatrix} x^c \\ y^c \\ 1 \end{bmatrix} = \frac{1}{Z^c} [P_{\text{perspective}} \quad 0] \begin{bmatrix} X^c \\ Y^c \\ Z^c \\ 1 \end{bmatrix} = \frac{1}{Z^c} [I_{3 \times 3} \quad 0]_{3 \times 4} \begin{bmatrix} X^c \\ Y^c \\ Z^c \\ 1 \end{bmatrix}$$

The positions are now shown in lower case and lie now on a 2D plane.

C.4.3 Extrinsic Matrix

This matrix belongs in SE(3) and it represents the transformation of position coordinates from the real world frame to the camera coordinate frame.

$$\begin{bmatrix} X^c \\ Y^c \\ Z^c \\ 1 \end{bmatrix} = [R \quad t] \begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix}$$

C.4.4 Full Projection Matrix Workspace to Camera

Grouping everything together we get a final full projection matrix, P, such that

$$\begin{aligned} \begin{bmatrix} u \\ v \\ 1 \end{bmatrix} &= P \begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix} \\ P &= K [I \quad 0] [R \quad t] \\ &= \begin{bmatrix} \alpha & s & u_0 \\ 0 & \beta & v_0 \\ 0 & 0 & 1 \end{bmatrix} [I \quad 0] [R \quad t] \end{aligned}$$

Or with the general affine transformation as the intrinsic matrix

$$P = \begin{bmatrix} a & b & c \\ d & e & f \\ 0 & 0 & 1 \end{bmatrix} [I \quad 0] [R \quad t]$$

C.5 Intrinsic Matrix Calibration

Solve the simpler problem of only finding the intrinsic matrix if the extrinsic is known. If you are already given the 2D camera coordinates (post extrinsic and projective transformation, then you can select at least 3 known points in the real world and image plane (which do not all lie within the same plane). Remember the general form

$$\begin{bmatrix} u \\ v \\ 1 \end{bmatrix} = \begin{bmatrix} a & b & c \\ d & e & f \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x^c \\ y^c \\ 1 \end{bmatrix}$$

For example, let's say that we know two different points then

$$\begin{bmatrix} u_1 \\ v_1 \\ 1 \end{bmatrix} = \begin{bmatrix} a & b & c \\ d & e & f \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_1^c \\ y_1^c \\ 1 \end{bmatrix}, \quad \begin{bmatrix} u_2 \\ v_2 \\ 1 \end{bmatrix} = \begin{bmatrix} a & b & c \\ d & e & f \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_2^c \\ y_2^c \\ 1 \end{bmatrix}$$

The system of equations become

$$\begin{aligned} ax_1 + by_1 + c &= u_1 \\ dx_1 + ey_1 + f &= v_1 \\ ax_2 + by_2 + c &= u_2 \\ dx_2 + ey_2 + f &= v_2 \end{aligned}$$

Rearranging into the form $Ax = b$

$$\begin{bmatrix} x_1 & y_1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & x_1 & y_1 & 1 \\ x_2 & y_2 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & x_2 & y_2 & 1 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \end{bmatrix} \begin{bmatrix} a \\ b \\ c \\ d \\ e \\ f \end{bmatrix} = \begin{bmatrix} u_1 \\ v_1 \\ u_2 \\ v_2 \\ \vdots \end{bmatrix}$$

A quick and less rigorous derivation of the solution: We want to approximate b with the subspace of A . This means through the projection theorem that $b - Ax$ is perpendicular to Ax

$$\begin{aligned} b - Ax &\perp Ax \\ (b - Ax) \cdot (Ax) &= 0 \\ (Ax) \cdot (Ax) &= (Ax) \cdot b \\ x^T A^T Ax &= x^T A^T b \\ x &= (A^T A)^{-1} A^T b \end{aligned}$$

Then plug the corresponding values into the complete intrinsic matrix.

C.6 Nonlinear Phenomena in Calibration

In real life, because of the imperfect lens or inaccurate placement of the lens, the image would be distorted. We can use nonlinear root finding methods like Newton-Raphson to account for nonlinear phenomena. ($x_{distorted} = x_{linear} + g(x)$)

- Radial Distortion

This phenomenon occurs when light rays bend more near the edges of a lens than they do at this optical center. Smaller lens can lead to greater distortion.




Figure C.10: Types of radial distortion

- Tangential Distortion This phenomenon occurs when the lens and the image plane are not parallel




Figure C.11: Tangential distortion

C.7 Projection Map from LiDAR to Camera

C.8 Projection Map

The projection map is a mapping between a 3D and a 2D world, and is used to project a LiDAR point cloud (3D) into an image (2D). Let X be the LiDAR points and Y be the projected points on the image-plane of a camera. The standard relations are ¹

$$\begin{bmatrix} u' \\ v' \\ w' \end{bmatrix} = \underbrace{\begin{bmatrix} f_x & s & c_x \\ 0 & f_y & c_y \\ 0 & 0 & 1 \end{bmatrix}}_{\text{intrinsic parameters}} \begin{bmatrix} \mathbb{1}_{3 \times 3} \\ \mathbf{0}_{1 \times 3} \end{bmatrix}^T \underbrace{\begin{bmatrix} R_L^C & t_L^C \\ \mathbf{0}_{1 \times 3} & 1 \end{bmatrix}}_{\text{extrinsic parameters}} \underbrace{\begin{bmatrix} x_i \\ y_i \\ z_i \\ 1 \end{bmatrix}}_{\text{LiDAR points}} \quad (\text{C.2})$$

$$= K \begin{bmatrix} \mathbb{1}_{3 \times 3} \\ \mathbf{0}_{1 \times 3} \end{bmatrix}^T H_L^C \begin{bmatrix} x_i \\ y_i \\ z_i \\ 1 \end{bmatrix} \quad (\text{C.3})$$

$$Y_i = [u \quad v \quad 1]^T = \begin{bmatrix} u' \\ v' \\ w' \end{bmatrix} \quad (\text{C.4})$$

where (C.2) includes the camera's intrinsic parameters (K) and the extrinsic parameters (R_L^C, T_L^C).

¹More information about the projection map, we refer the readers to "Multiple view geometry in computer vision second edition," by Hartley Richard and Zisserman Andrew, and "Computer vision: a modern approach," by Forsyth David A and Ponce Jean.

For later use, we combine (C.2) and (C.4) to define

$$\Pi(X_i; R, t) := Y_i, \tag{C.5}$$

the projection map from LiDAR coordinates to image coordinates. Note that it is a function of both the extrinsic variables and the LiDAR vertices.