

NICHOLAS HOELL

LINEAR ALGEBRA

MAT223 COURSE NOTES

UNIVERSITY OF TORONTO

Copyright © 2017 Nicholas Hoell

PUBLISHED BY UNIVERSITY OF TORONTO

WWW.MATH.TORONTO.EDU/NHOELL/MAT223

Licensed under the Apache License, Version 2.0 (the “License”); you may not use this file except in compliance with the License. You may obtain a copy of the License at <http://www.apache.org/licenses/LICENSE-2.0>. Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an “AS IS” BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

First printing, September 2017

Contents

<i>Introduction</i>	5
<i>What is this course about?</i>	6
<i>Vocabulary</i>	9
<i>Proofs</i>	10
<i>Sets</i>	12
<i>Expectations</i>	14
<i>Tips on Preparing</i>	15
<i>Crowdsourced Tips for Success</i>	15
<i>Exercises</i>	17
<i>Systems of Linear Equations</i>	19
<i>Linear Equations</i>	20
<i>Systems of Linear Equations</i>	23
<i>The Reduction Algorithm</i>	28
<i>Exercises</i>	36
<i>Geometry of Vectors</i>	39
<i>The Dot Product in \mathbb{R}^2</i>	40
<i>The Dot Product in \mathbb{R}^n</i>	42
<i>Projections and Expansions</i>	47
<i>Exercises</i>	49

<i>The Rank Theorems</i>	53
<i>Subspaces of \mathbb{R}^n</i>	54
<i>Bases</i>	55
<i>Expansions and Orthogonalization</i>	57
<i>Rank Unification</i>	60
<i>Maximal Rank</i>	66
<i>Exercises</i>	68
<i>The Fundamental Theorem of Linear Algebra</i>	71
<i>Prelude: Orthogonal Complements</i>	72
<i>The Fundamental Theorem of Linear Algebra</i>	73
<i>The Diagrams</i>	75
<i>Exercises</i>	79

Introduction

Contents

<i>What is this course about?</i>	6
<i>Where is this material used?</i>	7
<i>Examples</i>	7
<i>Vocabulary</i>	9
<i>Proofs</i>	10
<i>Sets</i>	12
<i>Expectations</i>	14
<i>Tips on Preparing</i>	15
<i>Crowdsourced Tips for Success</i>	15
<i>Exercises</i>	17

THIS chapter provides an overview of what this course is about and the applications in which you'll make use of techniques encountered in this course: *solving systems of linear equations*. As well, there's a primer on mathematical topics and terms we **expect you to know** or to familiarize yourself with as soon as possible. There are some tips from successful students at the end which you may find helpful as well as advice I can offer from from years of experience in observing what works for students.

What is this course about?

This course is about *one, and only one thing*: solving systems of linear algebraic equations. Of course, there's a lot to unpack here since we'll need to formally define the words *system* and *linear* but equally important, we'll need to think clearly about the word *solving*. By that I mean *it's better to deeply understand what's going on* in trying to "solve" systems of linear equations than blindly applying algorithms, because the deep understanding allows a greater facility with modelling and, ironically, with applying known algorithms. Because of this fact, this course is going to involve trying to understand the general character of systems of linear equations through abstracting out the more fundamental features and proving general statements about them rather than working through specific cases. This class is going to force you to think differently than you may be used to thinking about mathematical objects since the emphasis is heavily tilted towards *abstraction* for the purpose of *clarity*.

I should say, the *goal* of the course is for you to understand the following picture which incorporates all of the mathematical structures we'll encounter over the course of the semester.

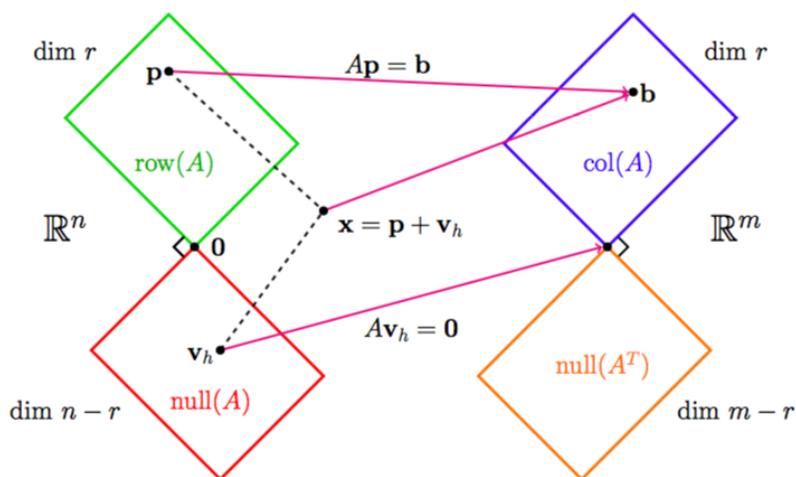


Figure 1: Understanding the objects in this picture and how they relate to one another is the ultimate aim for us this semester.

The picture above has a lot of moving parts and it will take quite awhile to reach the point where we can even precisely state what's happening there. But in case you're curious, the figure is a visual representation of a theorem, the *fundamental theorem of linear algebra* which helps us *quantitatively* understand how linear systems, and their possible solutions, allow for a very

particular geometrical way of dividing up the world where “vectors” live. Of course, if none of this makes sense right now - good! It’s meant only as flag for us to point out where we’re headed as we get into some of the details in solving linear systems of algebraic equations.

Where is this material used?

Linear systems and the *ideas of linearity* are ubiquitous in science. As you’ll discover in your futures in the your respective disciplines, linear algebra is the de facto *lingua franca* of the sciences. Why is that? I can offer two obvious reasons.

1. **Cases where Nature appears to actually *be* linear.** I’m not sure how common this is but it certainly appears more as an exception than as the rule. The most obvious (to me) example of this is in quantum mechanics, the branch of physics devoted to the very smallest bits of Nature. Big parts of the well-known strangeness of the picture of reality painted by quantum mechanics is entirely due to the fact that Nature appears to be obeying linear rules.¹
2. **Cases where Nature appears to not be linear.** I think this is the standard situation. When things deviate too far from linearity people tend to call these cases “non-linear” which is a technical term meaning, basically, “hard”.² It may be surprising to you, but these are precisely the cases in which the tools of linear algebra are often the most useful. This is because often a general technique of **linearization** is used wherein a hard, non-linear problem which cannot be solved directly gets approximated by an easier, linear problem which can be solved. Provided the approximations are done cleverly and carefully enough, the approximations can often give enough of an answer as to satisfy our reasons for asking. A special case of this which should be familiar to you from high school calculus is the local approximation of a function by it’s tangent line: tangent lines are “linear approximations” to messy, complicated non-linear functions. Far more complicated examples abound in applications.

¹ The weirdness has more to do with the interpretations offered for the equations, the equations themselves being perfectly normal, linear differential equations. The other big source of weirdness is the coupling of the linear equations with the so-called “Born rule” for those who are interested.

² Strictly speaking, it often means “so hard it’s not really solvable”.

Examples

I want to give a few examples of places where linear algebra and the tools from this course make an appearance in real-life. The list isn’t exhaustive, it’s just places where I’ve used linear algebra, or I’ve seen linear algebra used, or I know that linear algebra happens to be used. In these examples, linear algebra plays a *big* role.

1. **Physics.** I already mentioned quantum mechanics above. Beyond that, much of classical physics is described (or approximated) by linear laws, at various stages. Special relativity, for example, is basically a clever application of tools from this class. Particle physics, the type of physics

making the news for the discovery of the Higgs particle for example, rests on equations that can't be understood without mastering the material in this course.

2. **Mathematics.** Surprisingly, linear algebra is has applications *within mathematics itself*. In fact, one enormous branch of mathematics "representation theory", is based on massively clever uses of linear algebra. The basic idea is that while the objects in linear algebra are abstract, they have the benefit of being very well-understood. So if one encounters a mathematical object which is *really abstract*³ then we can study these really abstract things by somehow moving them into the world of linear algebra (at the cost of losing certain bits of informations along the way). By studying how abstract objects present themselves in linear form, you can actually learn something about the original abstract objects. This idea is used in many foundational areas of modern mathematics.
3. **Computer Science.** How fast can the fastest possible algorithm reliably multiply two arrays of numbers of increasingly large size? That question, as of this writing, remains an unsolved question in *complexity theory*, the discipline of computer science dealing with optimal, idealized ways of solving abstract problems. How could such an easy operation, multiplying two matrices, be so difficult to understand? Because every time people think they've found the quickest way of doing it, some brilliant computer scientist notices some very clever method for shaving a little bit of work off the total cost. Very small speedups in performance on matrix arithmetic have serious advantages since many algorithms require doing these things on large sets of numbers repeatedly. Small gains in performance time matter, actually. A lot.
4. **Google.** Search engines are a type of "black box linear algebra device". The way they work rests on doing very fast matrix manipulations, some of which are based on things we'll do in this course.
5. **Video Game Design.** The representation of images and objects in video games is often array-based and many of the tools we learn in this course have applications in the design of video games. People in this industry are masters of this course.
6. **Image Recognition.** There is now software that can identify people (or cars, handwriting, whatever) in new photos based on prior images. This turns out to be a linear algebra problem actually.
7. **Artificial Intelligence and Machine Learning.** Much of what's done now in modern applications of artificial intelligence (self-driving cars, self-directed vacuums) and automated learning (there are computers that have "learned" how to play perfect games of Breakout for instance) rests on linear algebra. Things common to both fields, like neural networks, are based on material we'll do in this class.

³ And you'd be shocked at quite how abstract this can be!

8. **Forecasting.** I give you a list of closing prices for a tradable asset for the past 2 weeks and I want you to estimate tomorrow's. How do you do this? There are many ways⁴ to attempt this. All the ones I know of involve linear algebra.
9. **Data Science.** Any manipulations done on large data sets must meet good performance requirements. Keeping things linear is a good way to proceed. As well, big data is usually kept in formats where linear algebra is the obvious weapon of choice.
10. **Statistics.** The last four examples are, in some sense, applications of statistics. Then again, the entire scientific enterprise is kind of "just" statistics in the sense that statistics is the careful and quantitative analysis and inference of data, and science is simply the collection, organization and interpretation of experimental data. You can not get anywhere in statistics without a mastery of linear algebra. Period. It is a discipline drenched in the language of linear algebra and probably the biggest masters of matrices are found in Departments of Statistics.⁵

⁴ And, I hope this is obvious, no reliably perfect ways.

⁵ But don't tell my colleagues in the mathematics department I said so

Vocabulary

An **enormous** amount of difficulty students in this course run up against is the *correct and grammatical* usage of precise terminology. The words have very rigid usages in this course, unlike in natural languages like English. This simple, almost naive, observation will become fundamental as the concepts encountered become more abstract. Moreover, *misuse* of very precisely defined objects indicates a flaw in proper understanding of the underlying concepts. For this reason it is **absolutely paramount** that you are able to convey, in written form, a clear, concise, rigorous, and correct argument using mathematical definitions.

First off, there's a bit of vocabulary invented by mathematicians to help them deal with parsing aspects of the mathematical theories they develop. If you like music, there are various phrases (coda, cadenza, transition, resolution, etc) which help the musicians/composers demarcate the control flow among passages inside a single coherent composition. Here are some of the ones used analogously in mathematics.

1. **Theorem.** This is used to indicate the big result, the ultimate goal of intense mathematical labour. All of the deepest results in mathematics are given this honorific.⁶ The general language in theorems is a statement of *assumptions* (or "hypotheses") e.g. "Let A be an $m \times n$ matrix..." or "Suppose that \mathbf{x} and \mathbf{y} are vectors in \mathbb{R}^n and..." followed by *conclusions* which are guaranteed true provided the hypotheses hold. The language used in statements of theorems is famously *precise* and *technical*.
2. **Lemma.** A lemma is like a micro-theorem. It's used to title results that are somewhat interesting in and of themselves, but whose primary purpose

⁶ Gauss even has one which now bears the name "Theorema Egregium" which means, roughly, "Remarkable Theorem", or "Really big theorem" depending on your tastes.

is to assist in proving theorems.

3. **Proposition.** This is something closer to theorem than to lemma but, well, we can't all be theorems now can we?
4. **Corollary.** This is an important result which follows, using not too much work, from a Theorem or a Proposition.
5. **Axiom.** These are things we just have to take in without being able to prove. We like to keep the list of these as minimal as possible (both in number and in cognitive complexity). Basically these are Propositions we cannot prove but simply assume.⁷
6. **Proof.** This is something which follows a lemma, proposition, theorem, or corollary. It's a *formal argument* designed to be incontrovertible evidence against further scrutiny. If all that were known to someone were existing axioms, and the lemmas, propositions, theorems and corollaries already established using these axioms, then that person would be able to verify, based solely on your proof, that a given statement was true. The famed physicist Richard Feynman once said (about science, not mathematics, but it holds here as well)

"The first principle is that you must not fool yourself – and you are the easiest person to fool."

I encourage you to take that advice seriously. Proofs are arguments made to safeguard against the kind of deception warned against in the quote.

Proofs

Each proof is unique since you're proving a different statement, but there are some common strategies you'll encounter. For general guidelines, here are a few thoughtlets.

- (a) If the statement you're asked to prove is something like "Prove such and such exists." it suffices to simply exhibit an object meeting the requirements described in the statement. In other words, *providing an example* constitutes a proof. Conversely, a *counterexample* is often used to disprove an erroneous claim.
- (b) In some cases (though not in this course) the above cannot be done, and existence is *non-constructive*, namely existence is established without being able to produce a single example of the object proven to exist. Don't worry about this case in this class since we won't see things like this, just be aware that this stuff can be subtle.
- (c) Sometimes we *argue by contradiction*, which is to say, we assume that the result we want to show *isn't true* and use logic to arrive at something we know to be false. Suppose *A* and *B* are statements (called *propositions* but not to be confused with the word Proposition used

⁷ An example of this would be the axiom that for any two sets the collection of all things in either of the two sets is itself a new set we can play with. Another, well-known from Greek geometry, would be that all right angles are equal. People spend entire careers trying to see which axioms imply others in any given system of axioms, in order to possibly reduce the "assumptive burden" of the system. Proving which things can or cannot be proven in any given list of axioms is an entire subfield of mathematics.

before as a kind of theorem!) and we are hoping to show that statement A implies statement B (written $A \implies B$ symbolically). Well, since

$$(\text{not } B) \implies (\text{not } A) \quad \text{if and only if} \quad A \implies B$$

then if we can prove $(\text{not } B) \implies (\text{not } A)$ then we can conclude the claim we wanted to establish must be true. Make sure you understand why $A \implies B$ is equivalent to $(\text{not } B) \implies (\text{not } A)$. It may help to think of examples "If I win the lottery, then I'll be rich" must be the same as "If I'm not rich, I did not win the lottery". But also notice that these are distinct from $B \implies A$. After all, not every rich person won the lottery.⁸

- (d) Sometimes we may *argue inductively*. That is to say, if we want to prove a statement $P(n)$ is true for all natural numbers n we need to show

- $P(0)$ is true. This establishes a "base case", i.e. the case for $n = 0$ (or, often $n = 1$ or whatever).
- If $P(k)$ is true for ($k >$ base case) then $P(k + 1)$ is true.

We won't use this much but it's there if we want it. The following is an example of an inductive argument.

EXAMPLE 1

Prove that $1^2 + 2^2 + \dots + n^2 = \frac{n(n+1)(2n+1)}{6}$ for $n = 1, 2, 3, \dots$

PROOF

Here $P(n)$ is the claim that $1^2 + 2^2 + \dots + n^2 = \frac{n(n+1)(2n+1)}{6}$ holds for positive natural numbers. Suppose $n = 1$. Then $1^2 = \frac{1(1+1)(2 \cdot 1 + 1)}{6}$. This gives the base case. Next, if it were true that for $k \geq 1$ we had that $1^2 + 2^2 + 3^2 + \dots + k^2 = \frac{k(k+1)(2k+1)}{6}$ then we would have

$$\begin{aligned} 1^2 + 2^2 + \dots + k^2 + (k+1)^2 &= \underbrace{\frac{k(k+1)(2k+1)}{6}}_{\text{by inductive hypothesis}} + (k+1)^2 \\ &= \frac{k(k+1)(2k+1)}{6} + \frac{6(k+1)^2}{6} \\ &= \frac{k(k+1)(2k+1) + 6(k+1)^2}{6} \\ &= \frac{(k+1)[k(2k+1) + 6(k+1)]}{6} \\ &= \frac{(k+1)[2k^2 + 7k + 6]}{6} \\ &= \frac{(k+1)(k+2)(2k+3)}{6} \end{aligned}$$

⁸ A little terminology here. If $A \implies B$ we say that A is "sufficient" for B and that B is "necessary" for A . If $A \implies B$ but $B \not\implies A$ then we say " B is necessary but not sufficient for A ".

which equals $\frac{(k+1)((k+1)+1)(2(k+1)+1)}{6}$, the result for $k + 1$. Since $P(k) \implies P(k + 1)$ we necessarily then have that $P(n)$ holds for all n .

CAUTION!

A surprising number of students make serious errors when working through proofs because of mixing up the hypotheses and the conclusion by mistakenly *assuming what is to be proven!* This can often happen in sometimes subtle ways so let's review.

Let C be a claim we wish to prove. For instance the claim might be something like "there are infinitely many prime numbers". We could restate this claim as "For every given prime number, there exists a larger prime number". Stated this way, it's more obvious what the hypothesis and conclusion are, namely the hypothesis here is "If you give me a prime number" and the conclusion is "There will always be a larger prime number". If I call the hypothesis P and the conclusion Q , we want to prove that $P \implies Q$ and we would be wrong to assert Q without having begun at P .

Sets

A lot of the definitions and proofs in the course are phrased in the language of sets. Sets are the primary foundational objects in mathematics. A set S is simply a collection of *elements*. These elements are often indicated as an unordered list $S = \{s_1, s_2, \dots, s_m\}$ say for the case of a finite set (a set with a finite number of things in it). The real numbers \mathbb{R} are an example of a familiar (hopefully) set with an infinite number of elements⁹. So is the set $\mathbb{N} = \{0, 1, \dots\}$ of natural numbers.

We use the notation $\#S$ to denote the number of elements¹⁰ in the set S . For example $\#\{2, 3, 5, 7, 11\} = 5$. By fiat we say that $\#S = \infty$ for sets S with an infinite number of elements.

The primary relationship used to describe sets is *membership*, denoted by the symbol \in or its negation \notin read as "in" or "not in", respectively. This symbol is used to indicate that an element is in a given set: namely, $s \in S$ means the element s is a member of the set S . These symbols can be oriented in a reversed way as \ni and \ni so that

$$s \in S \iff S \ni s$$

In other words, the above are equivalent statements.

Often, the sets we encounter in this course are described by listing condi-

⁹ The real numbers are the continuum of numbers on the number line used for graphing functions in high school. This is the set which contains all integers, fractions, and numbers like π , e and other numbers with non-repeating decimal expansion.

¹⁰ Also called the *cardinality* of the set S .

tions inside the defining brackets of the set. For example

$$\mathbb{Q} = \left\{ \frac{p}{q} \mid p, q \in \mathbb{N}, q \neq 0 \right\}$$

describes the set of *rational numbers*. Notice, in the above, that the vertical line $|$ (some people use colons instead of vertical lines) separates the left hand side, which gives a description of the elements, from the right hand side, which gives a *restriction* on the things appearing on the left hand side. Notice, in the above, p, q *don't really exist* since they represent placeholders for numbers which meet some restrictive criteria: if I were to replace them in all occurrences, I would still have the same set description, namely

$$\left\{ \frac{p}{q} \mid p, q \in \mathbb{N}, q \neq 0 \right\} = \left\{ \frac{s}{r} \mid s, r \in \mathbb{N}, r \neq 0 \right\}$$

are both perfectly valid descriptions of the set \mathbb{Q} . This property of the variables appearing in the sets above, that the variables can be replaced by any other unused variable provided the replacement is done in all occurrences of the variable, is termed **binding** and the variables are said to be **bound variables**. A misunderstanding of the differences between free and bound variables is a source of constant trouble for many MAT223 students. If you are confused - *don't wait to get unconfused* because this issue is crucial in understanding the concepts throughout the semester.

As well, there are a few natural binary operations¹¹ The first one is the *union* operation, denoted \cup . Unioning two sets S_1, S_2 creates a new set $S_1 \cup S_2$ containing all elements which appear in *either* set S_1 or S_2 *or both*. The way we write this fact is

$$S_1 \cup S_2 = \{s \mid s \in S_1 \text{ or } s \in S_2\}$$

Another natural binary operation is the intersection operation \cap which takes two sets S_1, S_2 and creates a new set $S_1 \cap S_2$ containing elements which appear in *both* S_1 and S_2 .

In addition to the above there are binary *inclusion* relationships. For instance $S_1 \subset S_2$ means that S_1 is a *subset* of set S_2 . What that means is that everything in S_1 must also be in S_2 . Namely

$$s \in S_1 \implies s \in S_2, \quad \text{for all } s \in S_1$$

Establishing that implication above, for an *arbitrary* element $s \in S_1$, is all that's required to show that $S_1 \subset S_2$. Here I want to make an important point about **notation**: different people prefer different conventions, and some authors prefer the notation $S_1 \subset S_2$ to generally indicate that S_1 is a *proper* subset of S_2 , namely that S_1 is *not actually equal* to S_2 . Those authors may use notation like $S_1 \subseteq S_2$ to indicate the neutral position, not making assumptions about whether S_1 is a proper subset or not. But many authors (the majority) use the notation \subset and \subseteq interchangeably and will use the

¹¹ Operations which takes two *operands*, in this case two sets.

notation $S_1 \subseteq S_2$ to denote that S_1 is a subset contained in, but not equal to, set S_2 . All of these notations have their reversed counterparts $\supseteq, \supsetneq, \supset$ which simply reverse the inclusion relationship to be read from right to left.

When are two sets equal? Clearly, when they have the same elements. In other words $S_1 = S_2$ must mean that everything in set S_1 must appear in set S_2 , i.e. $S_1 \subseteq S_2$ and everything in set S_2 must appear in set S_1 i.e. $S_2 \subseteq S_1$. In other words

$$S_1 = S_2 \iff S_1 \subseteq S_2 \quad \text{and} \quad S_2 \subseteq S_1$$

This is a standard approach to showing two sets are equal which we'll use through the semester so make sure you understand what it says and why it's valid.

Two final comments about sets. First, there's a (somewhat pathological) set that appears as a subset of every set - namely, the *empty set*. The empty set, denoted \emptyset (or $\{\}$) is what its name implies, it's a set with no elements. Its utility is there for reasons of mathematical consistency well beyond the scope of what we will encounter in this course. Note well that \emptyset is the only set satisfying $\#\emptyset = 0$. Secondly, there's an operation called complementing, which given a set S produces a set S^c called the *complement* of S . S^c is a set whose elements are all elements not appearing¹² in S . For instance, what is \mathbb{Q}^c ? Well, it should be all numbers not expressible as ratios of integers. In other words, it's the set of *irrational* numbers. So we have $\pi \in \mathbb{Q}^c$, $\sqrt{2} \in \mathbb{Q}^c$, etc.

¹² A **subtle point**: the idea of taking complements raises the question of *to which ambient set* are we referring? For instance, is $\mathbb{Q}^c = \mathbb{N}$ or \mathbb{R} or something else altogether? For the most part, the context is clear and so this isn't an issue. When complete clarity is needed, the notation $S_2 \setminus S_1$ will be used. $S_2 \setminus S_1$ means all elements in S_2 which are not in S_1 . This indicates that S_2 is the ambient set for which to take complements.

Expectations

In general we expect complete facility with the logic of things like " $A \implies B$ is equivalent to $(\text{not}B) \implies (\text{not}A)$ " on tests and quizzes. In other words, basic logic and solid reasoning is what we expect of you.

That said, of course we don't expect you to be able to prove super hard things on quizzes or tests so you shouldn't stress **too** much. Most of what we will ask for in a conceptual or proof-style question on a test can be done in only a few lines, just a short, rigorous explanation which correctly applies definitions or known theorems. Often, amazingly, performing a rigorous argument isn't a whole lot harder than being able to write down a definition and think about what it actually means, and think through the consequences of the assumptions you've been given.

As well, you may be asked to give definitions of things we've gone over in lecture or in homework. These should be easy, free points. All we ask is to restate a definition. It's our way of checking that you are paying attention and really internalizing the concepts. But, of course, the definition has to be precise and correct!

Tips on Preparing

I haven't discovered some new, super-fancy and stress-free way of mastering mathematics. All I can offer are a few, very general suggestions. They do work, but you *actually have to do them* in order for them to work. Most people end up trying to "yeah, yeah" and cut a few corners. What *you* do with my well-meant suggestions is up to you.

1. **Working with friends can help a lot.** It can help because it gives you the opportunity to explain your reasoning, out loud, to another person. You'd be surprised how helpful that can be.
2. **Working all the suggested problems.** Mathematics isn't a spectator sport. The only thing that makes you better is practice, practice, practice.
3. **Not getting behind.** The semester moves fast. It's very easy to procrastinate and fall behind on homework. It's a recipe for disaster in a course like this because cramming won't help much. Staying on schedule and being diligent with the homework requires discipline but it's worth it.
4. **Don't miss tutorials.** Exiting course grades are **strongly correlated** with attendance in tutorials and performance on quizzes. Not attending tutorials is a very reliable predictor of poor final grades.
5. **Use the class Piazza.** It's there for you to ask questions and learn from each other. I wish it had been around when I was a student.
6. **Try to make time each day to work on MAT223 related things.** This could be as easy as reading the book for 20 minutes every other day and doing problems in the days between. If you can find a routine where you reliably spend a portion of each day working problems, thinking through concepts, reading material, etc. you'll find your ability to retain the information greatly enhanced.
7. **Get enough sleep.** You'd be amazed at the benefits of sleep. I tried very hard as an undergrad to sleep a lot, especially before any tests. In general, the mental sharpness a good night's sleep gave far outweighed any little cramming I'd get by staying up too late.
8. **Try reading ahead.** If you read a few pages of the next lecture's topic, it often improves your ability to follow the lecture.

Crowdsourced Tips for Success

An excellent question was asked on Piazza in Fall, 2016. I'm sharing it and a few of the responses it received since I consider the advice therein to be very helpful. Here was the original posting.¹³

¹³ Forgive the small font, it happens to be the only way to have a reasonable looking image but it's admittedly hard to read.

study tips?

Those who receive good grades, how did you study for this course?

exam

~ An instructor (Nicholas Hoell) thinks this is a good question ~

And here's the first answer posted, which indicates the diligence typical of students who perform well on hard tests.

This is how I did it. 12 days before the test, I sat down and planned out my study schedule for this.

Since Professor Hoell is kind enough to give us the topics that are going to be covered, I used it to my advantage.

On the topics sheet, there are 9 topics that are listed (Inverses, Inverse Theorem, ..., Linear Independence)

Each day, I would review and master one topic at a time; for a total of 9 days of review. This way, I avoid cramming and in 1-2 hours a day, I can reach my review goal.

On the 10th day, I attempted the practice test EXACTLY as Professor Hoell recommended. This means that you simulate exam conditions as best as possible (no phone, no computer, no calculator, USE A PEN, give yourself 50 mins and no cheating).

After finishing the practice test, I went over it and identified what my strengths and weaknesses were. For example, I did not know how to express transformations as matrices, which ended up being one of the hardest questions on this test LOL.

For the next 2 days, I focused on my weaknesses and then I knew I was ready.

ESSENTIALLY, it comes down to TWO things: time management (DO NOT GRAM) and effective studying (there's no point of reviewing what you already understand, time is a limited resource, use it wisely). Effective studying also includes UNDERSTANDING how to do the problems vs MEMORIZING how to do them.

If I had just memorized how to convert linear transformations into matrices, I wouldn't have got the projdu question on the test. But since I know I understood that I had to break it down into the standard basis vectors, the rest was easy.

I also had 2 other tests on the same day to study for, so breaking the studying for this down into small manageable chunks helped me.

Good luck on finals!

Then I chimed in with my two cents.

i the instructors' answer, where instructors collectively construct a single answer

I can add one thing to the excellent ideas mentioned by other commenters. When I was a student, I used to read each section that I knew would be covered that week **before** coming to class. As I would read the section I used to read the statements of the theorems and **then see if I could prove them myself** before checking how the book did it. I know it sounds like it's probably a ridiculously hard thing to do, but I never went to a single undergrad course (all 4 years) without first having read the section and attempted the theorems on my own. Sometimes I'd be able to prove it myself, often, I'd fail and not get anywhere and end up seeing how the text proved the result. In this way, when I did read the proof, the cleverness of the ideas wasn't lost on me. Sometimes, when you see something proven, without having tried it yourself, it can appear deceptively easy. Part of the art of proving new things is the way the cleverness comes off as "natural". Beyond that, I also did precisely what the other commenters did. Also, keep in mind, math really is hard for everyone. At some point I realized that even the people for whom it seemed to come naturally, were putting in a lot of time behind the scenes thinking and studying.

edit good answer 3 Updated 10 days ago by Nicholas Hoell

And there was another commenter with additional ideas.

For my own input, I'll tell you exactly what I did (this midterm mark was 83% so whatever I did works for me):

- I went through every single chapter and did one per day starting a week before, and made sure I understood every single proof and could prove it back to myself.
- I did every single assigned question at the back, and I re-did each chapter at least 2x (did 5.1 and 5.2 at least three times)
- Did all the harder questions posted by our prof
- The amount of time spent on chapter 5 was probably double or triple the other chapters since it was weird to me at first
- Spent many hours on Khan Academy/PatrickJMT to understand what was going on
- Did the practice exam, freaked out when question 3 kicked my rear end and studied harder
- When there was a question I didn't understand at any part of the course, I kept working on it until I understood it
- MAT235 augments MAT223 so that helped with chapter 4
- Had a horrible fear of failing because I was ill and missed the first midterm, which was a security blanket almost everyone else had that I didn't, this drove me pretty hard

I did not go to any math aid center or any professor office hours. It was purely self studying. I also missed pretty much every single lecture from 2.4 -> 5.2 because I got the raw deal with health issues this semester, but those are more or less out of the way now fortunately. I do not recommend missing lecture of course (mine was way beyond my control), cause I'm pretty sure it just added to the stress of the midterm. What I'm trying to say is if you fall behind, there are resources out there for anyone to pick up this stuff and succeed. If I can do it with all the crap I faced, anyone else should be able to.

And, lastly, a voice from one of the 3 students who earned a perfect score on midterm 2. As in my own thoughtlets before, there isn't really a "royal road" here. There's just hard work and the ability to stay calm in a test environment. Being prepared is one of the best ways to remain calm.

Here are the things I did to study for this course (got perfect marks on quizzes and tests so far so this definitely is working for me) :

- Review the newly learned concepts weekly, and do the suggested chapter questions and the 'harder questions'. I always review BEFORE diving into the questions, so that I would know what I didn't really understand as I go through the questions. DON'T leave the questions until one week before the test
- Mark the questions that I did wrong and was not because of calculation error (though this indicates that I was being careless). Redo these a week after, and make sure that I actually understand the reasoning behind each step.
- Write down what concepts I had problems with when doing the questions and why my initial reasoning was wrong
- The way I prepared for the test was very similar to the first answer
- Get enough sleep the day before test (VERY important lol)
- During the test, read the questions CAREFULLY, underline important details when reading
When coming across a problem that I am not sure how to solve (after thinking for 1- 2 minutes) , I would tell my self NOT TO PANIC and skip this question for now; after answering most of the questions, I would go back to the unsolved ones
- I also summarized the types of questions (and the general solution) that might be tested (e.g. show that the following two subspaces are equal; determine the shortest distance between two lines, a plane and a line etc.) based on lecture notes and chapter questions

Hope this would help :)

I encourage you to study the success stories carefully and glean what you can from them. The big takeaway is that doing well requires a lot of work, and a strong desire to succeed. Good luck in your studies!

Exercises

Do the following exercises offline.

1. Describe the sets $\{x \in \mathbb{R} \mid \frac{x}{2} = k, \quad k \in \mathbb{N}\}$ and $\{x, y \in \mathbb{R} \mid \frac{x}{y} = 2, y \neq 0\}$
2. Describe the set $\{x \in \mathbb{N} \mid x \neq kl, k, l \in \mathbb{N}, k < x, l < x, k \neq 1, l \neq 1\}$
3. Write down a definition of $S_1 \cap S_2$ using set notation.
4. Show that $(S_1 \cap S_2)^c = S_1^c \cup S_2^c$ holds for all sets S_1 and S_2 .
5. Show that $(S_1 \cup S_2)^c = S_1^c \cap S_2^c$ holds for all sets S_1 and S_2 .
6. Prove that $(S_1 \cap S_2) \subset (S_1 \cup S_2)$ holds for all sets S_1 and S_2 .
7. Are the notations $\{s\} \subset S$ and $s \in S$ interchangeable? Why or why not?
8. Prove that for any subset T of set S , we have $S = T \cup T^c$
9. Here's a claim: $S = T \iff \#S = \#T$. If true prove it. If false, give a counterexample.
10. Use induction to prove that $1 + 2 + \dots + n = \frac{n(n+1)}{2}$ for all positive integers n .
11. Use induction to prove that $2 + 2^2 + \dots + 2^n = 2^{n+1} - 2$ holds for all $n > 0$.
12. Use induction to prove that $7n < 8^{n+1}$ holds for $n \geq 0, n \in \mathbb{N}$
13. Use induction to prove the following claim: Every nonempty subset of \mathbb{N} has a smallest element.

14. Consider a set S . We define a new set, $\mathcal{P}(S) = \{\text{all subsets of } S\}$.
- Write down $\mathcal{P}(\{0, 1\})$.
 - Write down $\mathcal{P}(\{0\})$ and $\mathcal{P}(\mathcal{P}(\{0\}))$
 - Use induction to prove that $\#\mathcal{P}(S) = 2^{\#S}$ holds for all sets with $\#S < \infty$
15. (Harder). Consider the set $S = \{s \mid s \notin s\}$. Is this set well-defined? Why or why not?
16. (Harder). Use induction to prove that $8^n - 3^n$ is divisible by 5 for $n > 0$
17. (Harder). Consider a well-known *incorrect* use of induction used to “prove” the false claim that cars are all the same colour. The “proof” goes like this:

The claim is that cars are all the same colour. This is equivalent to the claim that any set of cars must contain cars of the same colour. The claim is trivially true for the base case of $n = 1$ cars since, after all, a car has only one colour.¹⁴ Next, we’ll assume that the claim is true for all sets of n cars are we’ll now consider a set of $n + 1$ cars, {car 1, car 2, ..., car n , car $n + 1$ }. We’ll consider two subsets of this set $C_1 = \{\text{car 1, car 2, ..., car } n\}$ and $C_2 = \{\text{car 2, ..., car } n, \text{car } n + 1\}$, each of which is a set of n cars. Therefore, by our induction hypothesis, C_1 and C_2 only contain cars of a single colour. But then again, since there’s overlap in the entries of C_1 and C_2 the colour of each must be the same. Therefore {car 1, car 2, ..., car n , car $n + 1$ } must have only cars of a single colour. Thus, cars are all of the same colour.

¹⁴ Though this isn’t true for real cars, we’ll actually pretend this is true for cars here. In other words this line *is not* where the mistake happens.

What went wrong with the above “proof”?

18. (Harder). In a previous question you were asked to use induction to show that every nonempty subset of \mathbb{N} has a smallest element. Now show the converse. Namely show that *if* the claim that every nonempty subset of \mathbb{N} has a smallest element is a true fact about set the \mathbb{N} , *then* induction is a valid method of proof.

Systems of Linear Equations

Contents

<i>Linear Equations</i>	20
<i>Parameterized Solutions</i>	21
<i>General Linear Equations</i>	23
<i>Systems of Linear Equations</i>	23
<i>Equivalence</i>	25
<i>The Reduction Algorithm</i>	28
<i>The Reduction Algorithm</i>	32
<i>Exercises</i>	36

THIS chapter introduces you to the main subject of study in this course: systems of linear equations. We derive some elementary facts about such systems and develop a rigorous procedure for solving the solvable ones.

Linear Equations

This course is about solving systems of linear algebraic equations. We'll be dealing with methods of abstracting the fundamental concepts which are crucial in understanding such systems. In order for this to be comprehensible, we'll need to have a completely thorough understanding of the concrete specifics on which the abstraction is based.

With that in mind, let's first recall that a single linear algebraic equation, in variables x and y is an equation like

$$ax + by = c$$

In the above, x, y are *variables* while the numbers a, b, c are *coefficients*. As the coefficient c doesn't appear as being multiplied by a variable, we give it a special name and refer to it as the **constant term**. Often we place constant terms on the right hand side of the equals sign. If the constant term $c = 0$ we call the equation **homogeneous**.

In general, coefficients are **givens**, in the sense that we cannot change their explicit values whereas the variables are **unknowns** because they are placeholders for numbers we'd like to determine. Therefore, we often refer to the above as one equation in two unknowns. For example,

$$x - 2y = 5$$

has $a = 1, b = -2$ and $c = 5$. All this is just to remind you that the letters we'll use to represent coefficients are **fixed** values whereas generally the letters we use to represent variables are meant to be *unknown but not necessarily fixed* values.

By a **solution** to a linear equation like the above we mean an assignment of **values** to the variables resulting in a valid equality. For instance, the assignment of $x = 3$ and $y = -1$ in the above yields $3 - 2(-1) = 5$. In this example the assignment $x = 3, y = -1$ works as well. In fact, this example actually has *infinitely many* solutions as we can see by rewriting it as

$$x = 5 + 2y$$

In this form notice that each value we assign to the variable y , *by construction* must result in a valid assignment of a value for x . In other words, you can now immediately see that, for instance, if $y = 17$ then $x = 5 + 2(17) = 39$ so $x = 39, y = 17$ must be a valid solution to the original equation.¹⁵

By a **solution set** to a linear equation we mean the collection of all possible solutions. For instance, the solution set to the equation $x - 2y = 5$ is

$$\{(x, y) \mid x = 5 + 2y, y \in \mathbb{R}\}$$

which is the set of all pairs of numbers satisfying the constraint that the first number in the pair is always 5 plus twice the second number in the pair.

¹⁵ Notice that writing it this way is equivalent to writing $x = f(y)$, viewing x as a function of y .

Phrasing things this (overly verbose) way should highlight that we may also write the solution set above as

$$\{(x, y) \mid x = 5 + 2t, y = t, t \in \mathbb{R}\} \quad (0.0.1)$$

a point we focus on in the next section.

A final remark. In high school, you'll recall that the solution set for linear equations like $-mx + y = b$ is visually represented as the *graph* of the function $y = f(x) = mx + b$. While this visual reminder is helpful early on when we increase the number of variables it can become much harder to visualize solution sets.

Parameterized Solutions

We can express the relationship between y and x implicitly expressed in the equation $x = 5 + 2y$ by writing the solution in so-called **parameterized form** as

$$\begin{cases} x = 5 + 2t \\ y = t, \quad t \in \mathbb{R} \end{cases}$$

Where, in the above, we've introduced a *parameter* $t \in \mathbb{R}$, which implicitly varies over an infinite range of values.¹⁶ Namely, $t = 1, t = -3, t = \pi$ give three different values of pairs (x, y) each of which solves the original linear equation. We refer to the parameterized solution above as the **general solution** to the linear equation as every **particular solution** like $x = 3, y = -1$ can easily be obtained from the general solution by using a particular choice of parameter.

¹⁶ Take a moment to compare the above with (0.0.1) to verify that they are equivalent.

DEFINITION 1

A **parameterized solution** to a linear equation is a solution written in terms of one or more **parameters**. Parameters are variables presumed to vary over a given range of values, each of which gives a valid solution to the original linear equation.

As the definition indicates and the next example clarifies, parameterized solutions may well consist of multiple parameters.

EXAMPLE 2

Find the general solution, in parametric form, to the linear equation in three unknowns $2x - 3y + 6z = 19$.

Solution: We pick a variable (in this case x) to stay on the left hand side of the equals sign and move all others to the right giving

$$2x = 19 + 3y - 6z$$

Dividing by the coefficient on x gives us $x = \frac{19}{2} + \frac{3}{2}y - 3z$. Finally, we can introduce two parameters t, s to write the parameterized solution

as

$$\begin{cases} x = \frac{19}{2} + \frac{3}{2}t - 3s \\ y = t \\ z = s \end{cases}$$

for $s, t \in \mathbb{R}$

The utility of the above is that we can set t and s to be any numbers we wish and plugging them into the formula will immediately yield a valid assignment of values to variables x, y, z . For instance, $t = 2, s = 1$ gives, $x = 8, y = 2, z = 1$ as one particular solution to the linear equation.

In the previous example there was nothing canonical about the choice of x as staying on the left hand side of the equals sign. We could just as well have written the parameterized solution as, say,

$$\begin{cases} x = t \\ y = -\frac{19}{3} + \frac{2}{3}t + 2s \\ z = s \end{cases}$$

for $t, s \in \mathbb{R}$. Another important remark is that the parameters are *logically bound variables*¹⁷ which means that they may be replaced by different variables, provided the replacement is made in all occurrences, if you wish. For example, in the above example we'd obtained

$$\begin{cases} x = \frac{19}{2} + \frac{3}{2}t - 3s \\ y = t, \quad t \in \mathbb{R} \\ z = s, \quad s \in \mathbb{R} \end{cases}$$

But we can use a new parameter $\tilde{t} = \frac{3}{2}t$ and $\tilde{s} = -3s$ to give an equivalent general solution to the same problem

$$\begin{cases} x = \frac{19}{2} + \tilde{t} + \tilde{s} \\ y = \frac{2}{3}\tilde{t}, \quad \tilde{t} \in \mathbb{R} \\ z = -\frac{1}{3}\tilde{s}, \quad \tilde{s} \in \mathbb{R} \end{cases}$$

This is a subtle and important point which will recur throughout the course - if t, s are arbitrary numbers, then so must be $\tilde{t} = c_1t$ and $\tilde{s} = c_2s$ for any nonzero c_1, c_2 so there cannot be a *unique* parametric form for the general solution to a linear equation.

In fact, having done such a change of parameters doesn't warrant giving the parameters new symbols! We often prefer to simply write

$$\begin{cases} x = \frac{19}{2} + \frac{3}{2}t - 3s \\ y = t, \quad t \in \mathbb{R} \\ z = s, \quad s \in \mathbb{R} \end{cases} \text{ describes the same solution set as } \begin{cases} x = \frac{19}{2} + t + s \\ y = \frac{2}{3}t, \quad t \in \mathbb{R} \\ z = -\frac{1}{3}s, \quad s \in \mathbb{R} \end{cases}$$

¹⁷ Make sure you understand the discussion about free versus bound variables in the Introduction.

Notice that we use the same symbols t, s in both parameterized forms of the general solution without any problems since, owing to the fact that the parameters are *bound variables*, they are simply placeholders for particular numbers and therefore don't have intrinsic values associated with them. Again, this is often a point of confusion for many students so please review this example carefully.

General Linear Equations

We can generalize the simple examples described in the previous section by increasing the number of variables and coefficients in a linear equations. To do this we often write things like

$$a_1x_1 + a_2x_2 + \cdots + a_nx_n = b$$

for a linear equation in n variables x_1, \dots, x_n . We also call this *one equation in n unknowns*. As we've done before we have the constant term b on the right hand side of the equals sign. If $b = 0$ we call this equation **homogeneous**. We use subscripts on coefficients and variables since otherwise for equations with lots of variables we'd quickly run through the alphabet and not have enough symbols. We say that the equation is **consistent** if it has at least one solution and inconsistent otherwise. For example, the linear equation $x - 2y = 5$ is consistent since, as we already saw, it has infinitely many solutions (which we can describe using a parameterized form for the general solution). You might be wondering how a linear equation could be inconsistent. To see how this could happen consider a simple equation of the form $ax = b$ with $a = 0$ and $b = 1$. Clearly no value of x could solve such an equation so therefore the equation must be inconsistent.

Systems of Linear Equations

A *system* of linear equations is precisely what its name implies - namely, a collection of linear equations *in the same unknown variables*. For example

$$\begin{aligned} x - 2y &= 5 \\ 7x + 15y &= 6 \end{aligned} \tag{0.0.2}$$

is a linear system of two equations in the two variables x, y .

Generalizing this, we have a system of m linear equations in n unknowns, generally takes the form

$$\begin{aligned} a_{11}x_1 + a_{12}x_2 + \cdots + a_{1n}x_n &= b_1 \\ a_{21}x_1 + a_{22}x_2 + \cdots + a_{2n}x_n &= b_2 \\ \vdots & \\ a_{m1}x_1 + a_{m2}x_2 + \cdots + a_{mn}x_n &= b_m \end{aligned} \tag{0.0.3}$$

In the above, we're using *double subscripts* like a_{ij} to denote the coefficient of variable x_j appearing in the i th equation. In general of course, the order

matters - i.e. $a_{31} \neq a_{13}$ etc. We can unambiguously reference the coefficients in such a system by the notation a_{ij} , $1 \leq i \leq m$, $1 \leq j \leq n$. We may refer to the coefficient a_{ij} as the (i, j) 'th coefficient of the system. As well, the coefficients appearing on the right, b_1, \dots, b_m are the **constant terms**. If the constant terms are all zero, then the linear system is **homogeneous**.

The important thing is that the variables x_1, x_2, \dots, x_n are the same in all equations. Each equation in (0.0.3) can be seen as imposing an additional *constraint* on the allowed values for the variables x_1, \dots, x_n , since in general adding new equations can never serve to increase the possible numbers of solutions to the system. As with linear equations, by a **solution** to a linear system of equations we simply mean a particular assignment of values to the variables such that all equations in the system are valid. Of course, the collection of all possible solutions to a given system of equations is called the **solution set** of the system. Notice that this definition allows for the solution set to be empty. If there is one, and only one, solution in the solution set for a given linear system, then we say the solution is **unique**. These are summarized below.

DEFINITION 2

A linear system is said to be

- **Consistent** if there is at least one solution in the solution set
- **Inconsistent** if there are no solutions

Moreover, we say that the solution is **unique** if it is consistent and there is one and only one solution in the solution set. If the system is consistent but there is more than one solution, we say that a given solution is **not unique**.

The following is a fundamental result characterizing the nature of consistency of solutions to systems of linear equations.

THEOREM 1

A linear system of equations

$$\begin{array}{ccccccc} a_{11}x_1 + a_{12}x_2 + & \cdots & + a_{1n}x_n = & b_1 \\ a_{21}x_1 + a_{22}x_2 + & \cdots & + a_{2n}x_n = & b_2 \\ \vdots & & & \vdots \\ a_{m1}x_1 + a_{m2}x_2 + & \cdots & + a_{mn}x_n = & b_m \end{array}$$

can only have 0, 1 or infinitely many solutions.

PROOF

If the system of equations is not solvable then it definitely has 0 solutions and there's nothing left to prove. So suppose that the system of equations is consistent and that the solution set is S . If $\#S = 1$ then, yet again, there's nothing left to prove. So let's consider the case of $\#S > 1$. In other words, suppose that there are two *distinct* elements $s_1, s_2 \in S$, where $s_1 \neq s_2$. Each solution, s_1 say, is just a set of values for the variables in the linear system, i.e. $s_1 = \{x_1 = c_1, x_2 = c_2, \dots, x_n = c_n\}$ and $s_2 = \{x_1 = d_1, x_2 = d_2, \dots, x_n = d_n\}$ where $c_i \neq d_i$ for at least one i . But then now we can define

$$s_3 = \{x_1 = k_1c_1 + k_2d_1, x_2 = k_1c_2 + k_2d_2, \dots, x_n = k_1c_n + k_2d_n\}$$

for any choice of nonzero numbers k_1, k_2 such that $k_1 + k_2 = 1$. You can easily verify that $s_3 \in S$ and $s_3 \neq s_1, s_3 \neq s_2$. Notice we can select infinitely many choices of nonzero k_1, k_2 such that $k_1 + k_2 = 1$ so therefore there are infinitely many more ways to select elements like s_3 . Thus, $\#S = \infty$ and the result is proved. \square

Equivalence

Our strategy in solving systems of linear equations is to perform various manipulations on them in order to transform them into equivalent, but simpler to solve, systems. The following definition gives us the right manipulations to attempt.

DEFINITION 2: EQUIVALENT SYSTEMS

Two systems of linear equations are said to be **equivalent** if, and only if, they have the same solution set.

Making this concrete, let's consider (0.0.2). We already saw that the first equation in (0.0.2), $x - 2y = 5$, has infinitely many solutions, which we were able to produce in a simple parameterized form. However, not all of those solutions will be solutions to the full system (0.0.2). For instance $x = 1, y = -2$ solves the first but not the second of the equations in the system. To try to solve a system like (0.0.2) we recall a few observations about manipulations of simultaneous algebraic equations from high school.

DEFINITION 3: ELEMENTARY MANIPULATIONS OF SYSTEMS

For a system of linear algebraic equations

$$\begin{array}{r} a_{11}x_1 + a_{12}x_2 + \cdots + a_{1n}x_n = b_1 \\ a_{21}x_1 + a_{22}x_2 + \cdots + a_{2n}x_n = b_2 \\ \vdots \qquad \qquad \qquad \vdots \qquad \qquad \qquad \vdots \\ a_{m1}x_1 + a_{m2}x_2 + \cdots + a_{mn}x_n = b_m \end{array}$$

we can always

1. Swap the order of any two equations (called **interchange**)
2. Multiply both sides of any equation by a nonzero constant (called **scaling**)
3. Replace an equation with the result of adding a multiple of it and another equation (called **replacement**)

The set of solutions to the system of equations which results from any of the above elementary manipulations is the same as the set of solutions to the original system (the systems are **equivalent**).

The way that we systematically solve systems of linear equations is by careful and successive application of the allowed equation manipulations listed above so that the system appears in a nicer form than it began. The next example illustrates this in detail.

EXAMPLE 3

Solve system (0.0.2)

$$\begin{array}{l} x - 2y = 5 \\ 7x + 15y = 6 \end{array}$$

by use of elementary equation operations.

Solution Let's begin by replacing the current second equation with

$$\text{new second equation} = -7 \times (\text{first equation}) + \text{second equation}$$

This results in the system

$$\begin{array}{l} x - 2y = 5 \\ 0x + 29y = -29 \end{array}$$

Next we apply the scaling operation

$$\text{new second equation} = \frac{1}{29}(\text{second equation})$$

to get the system

$$\begin{array}{l} x - 2y = 5 \\ y = -1 \end{array}$$

We'll perform one last elementary manipulation of the rows by replacing the first row with the result of adding the first equation and twice the second equation as in

$$\text{new first equation} = \text{first equation} + 2(\text{second equation})$$

The result of which yields

$$x = 3$$

$$y = -1$$

which, as you can check, solves the *original* system of equations.

Make sure you heed the following caution about correct manipulations of linear systems.

CAUTION!

You'll notice that nowhere in the list of elementary equation manipulations is there a rule saying something like "put one variable in terms of another and substitute that variable into another equation and..." Often, students learn to solve a simple system by means of such *substitution* techniques. This is **NOT** how we solve systems of linear equations in general.

For instance, it would be **wrong** (and result in no points on an exam) to try to solve

$$\begin{aligned}x - 2y &= 5 \\ 7x + 15y &= 6\end{aligned}$$

by saying

$$\begin{aligned}x - 2y = 5 & \quad = \quad x = 2y + 5 \\ 7x + 15y = 6 & \quad = \quad 7x + 15y = 6 \\ & \implies 7(2y + 5) + 15y = 6 \\ & \implies (14 + 15)y = 6 - 35 \\ & \implies y = -\frac{29}{29} = -1 \\ & \implies x = 2(-1) + 5 = 3\end{aligned}$$

The above "method" for solving the linear system is **not allowed** in the course. The reasons for why it's not a reliable method will, hopefully, become obvious as the course progresses. The *only* methods we may use to solve systems of linear equations right now are those outlined in the *Elementary Manipulations of Systems*.

Of course, just like with single equations we may encounter parameterized solutions when solving systems of equations as the next example illustrates.

EXAMPLE 4

Solve
$$\begin{cases} x + y + 7z = 0 \\ 5x + y - z = 8 \end{cases}$$

We begin with a replacement operation, subtracting 5 times the first row from the second row resulting in
$$\begin{cases} x + y + 7z = 0 \\ -4y - 36z = 8 \end{cases}$$
. We can then perform a scaling operation on the second equation by dividing it by -4 which gives
$$\begin{cases} x + y + 7z = 0 \\ y + 9z = -2 \end{cases}$$
. Then we use a replacement operation to subtract equation 2 from equation 1 yielding

$$\begin{cases} x - 2z = 0 \\ y + 9z = -2 \end{cases}$$

At this point we can see that both x and y variables can be solved for in terms of the variable z so we can write the solution set in the following parameterized form

$$\begin{cases} x = 2t \\ y = -2 - 9t \\ z = t, \quad t \in \mathbb{R} \end{cases}$$

Since the system is solvable we say the system is consistent. Because the solution set is infinite, the solution is not unique.

The Reduction Algorithm

We're going to systematize the method we introduced in the last section for attempting to solve systems of linear equations. Let's begin by noting that when given a system like

$$\begin{cases} x - 2y - 4z = 5 \\ 7x + 15y + 2z = 6 \\ x + 3y - 12z = 0 \end{cases}$$

we can concentrate on the coefficients by considering the array of numbers making up the system, namely,

$$\left[\begin{array}{ccc|c} 1 & -2 & -4 & 5 \\ 7 & 15 & 2 & 6 \\ 1 & 3 & -12 & 0 \end{array} \right]$$

referred to as the **augmented matrix** of the linear system. Each row of the above matrix¹⁸ corresponds to an equation in the original system. The first 3 columns of the above matrix correspond to variables in the original system, whereas the 4th column corresponds to the constant terms. For the purposes

¹⁸ Matrix in this context simply means a two-dimensional array of numbers.

of trying to solve a linear system it's enough to work with the associated augmented matrix. The array of numbers

$$\begin{bmatrix} 1 & -2 & -4 \\ 7 & 15 & 2 \\ 1 & 3 & -12 \end{bmatrix}$$

is the **coefficient matrix** for the system and $\begin{bmatrix} 5 \\ 6 \\ 0 \end{bmatrix}$ is the **constant matrix** for the system. Thus, the augmented matrix is the coefficient matrix and the constant matrix for a system separated by a vertical line.¹⁹

It's easier to work with numbers rather than with equations. Since each row of the augmented matrix associated to a linear system is an equation of the system, our elementary manipulations of systems from last section can now be restated in terms of operations on rows of matrices.²⁰

DEFINITION 4: ELEMENTARY ROW OPERATIONS

The following are called **elementary row operations** on a matrix.

1. Swap the order of any two rows (called **interchange**)
2. Multiply a row by a nonzero constant (called **scaling**)
3. Add any multiple of one row to a different row (called **replacement**)

Performing any of the above elementary row operations to an augmented matrix for a linear system produces an augmented matrix for an equivalent system.

Any two augmented matrices whose linear systems have the same solution set are said to be **equivalent** (or row equivalent). Clearly elementary row operations produce equivalent matrices. We indicate that two matrices, A and B are equivalent by the notation $A \sim B$.

Let's go back and revisit an earlier example from the point of view of augmented matrices.

EXAMPLE 5

Solve

$$x - 2y = 5$$

$$7x + 15y = 6$$

by use of elementary row operations on the associated augmented matrix.

The corresponding augmented matrix is

$$\left[\begin{array}{cc|c} 1 & -2 & 5 \\ 7 & 15 & 6 \end{array} \right]$$

¹⁹ Many authors choose not to include the vertical line since the constant matrix always appears as the rightmost column of the augmented matrix.

²⁰ The plural form of matrix is *matrices*.

We can apply a replacement operation by adding -7 times row 1 to row 2, which will give us the opportunity to use a nice notation as well.

$$\left[\begin{array}{cc|c} 1 & -2 & 5 \\ 7 & 15 & 6 \end{array} \right] \xrightarrow{R_2 - 7R_1} \left[\begin{array}{cc|c} 1 & -2 & 5 \\ 0 & -29 & -29 \end{array} \right]$$

Then we perform a scaling operation, followed by another replacement operation.

$$\left[\begin{array}{cc|c} 1 & -2 & 5 \\ 0 & 29 & -29 \end{array} \right] \xrightarrow{\frac{1}{29}R_2} \left[\begin{array}{cc|c} 1 & -2 & 5 \\ 0 & 1 & -1 \end{array} \right] \xrightarrow{R_1 + 2R_2} \left[\begin{array}{cc|c} 1 & 0 & 3 \\ 0 & 1 & -1 \end{array} \right]$$

The system above corresponds to $\begin{matrix} x = 3 \\ y = 1 \end{matrix}$ so we can simply read off the solution.

In the last example we used a helpful bookkeeping system to indicate which elementary row operation we performed. Namely we used arrows with row combinations above them. For instance $\xrightarrow{R_4 - 5R_2}$ would indicate replacing row 4 by row 4 minus five times row 2. Swapping two rows could be indicated by, say, $\xrightarrow{R_2 \leftrightarrow R_5}$ which would mean interchange rows two and five.

CAUTION!

Be careful when using elementary row operations to avoid *combining* multiple operations into a single step. Namely, when performing a sequence of elementary row operations, get in the habit of doing them one at a time. Namely, instead of

$$\left[\begin{array}{cc} 1 & 0 \\ 0 & 1 \end{array} \right] \rightarrow \left[\begin{array}{cc} 0 & 5 \\ 7 & 0 \end{array} \right]$$

The above should be indicated as being the following sequence

$$\left[\begin{array}{cc} 1 & 0 \\ 0 & 1 \end{array} \right] \xrightarrow{R_1 \leftrightarrow R_2} \left[\begin{array}{cc} 0 & 1 \\ 1 & 0 \end{array} \right] \xrightarrow{5R_1} \left[\begin{array}{cc} 0 & 5 \\ 1 & 0 \end{array} \right] \xrightarrow{7R_2} \left[\begin{array}{cc} 0 & 5 \\ 7 & 0 \end{array} \right]$$

Whether or not a matrix is considered an *augmented* matrix, we're going to look at an algorithm for solving a system by applying a sequence of row operations to a corresponding matrix. For this, we begin by characterizing a "nice" form for matrices. Namely, in the preceding example, we ended up with the "nice" matrix $\left[\begin{array}{cc|c} 1 & 0 & 3 \\ 0 & 1 & -1 \end{array} \right]$ which has a form allowing us to simply read off the values of the associated variables. Our goal is to develop a systematic procedure for applying elementary row operations to matrices to get them into this kind of "nice" form so we can immediately read off solutions, if they are solvable, or determine if the solution is inconsistent

otherwise.

We say that a row is a zero row if every number in that row is zero. As well, we call the first nonzero entry in a nonzero row the **leading entry** for that row.

DEFINITION 5: ROW ECHELON AND REDUCED ROW ECHELON FORM

A matrix is said to be in **row echelon form** or **REF** if it satisfies

1. Any zero rows appear at the bottom of the matrix
2. The leading entry in a given nonzero row appears in a column to the right of the leading entry in the row above it (if any).
3. All entries (if any) in a column below a leading entry are zero.

Furthermore, a matrix is said to be in **reduced row echelon form** or **RREF** if it's in row echelon form and additionally satisfies

1. All (if any) leading entries are 1.
2. Each leading 1 (if any) is the only nonzero entry in its column.

The above is a lot to absorb so give it some thought. While being cumbersome to state it has the power of being *very precise*. So precise, in fact, that a computer can readily verify whether a given matrix is in RREF.

EXAMPLE 6

We use the convention that ■ is an arbitrary nonzero number (i.e. leading entry) and * shall denote any number. Then, as you can verify,

$$\begin{bmatrix} \blacksquare & * & * & * \\ 0 & 0 & \blacksquare & * \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

and

$$\begin{bmatrix} 0 & \blacksquare & * \\ 0 & 0 & 0 \end{bmatrix}$$

are both in REF.

The matrices

$$\begin{bmatrix} 0 & 1 & * & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \begin{bmatrix} 1 & * & 0 & * \\ 0 & 0 & 1 & * \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

are in RREF. These example illustrate the more general fact that REF and RREF matrices have a "staircase" like pattern to them.

It turns out (though we will not prove it) that every matrix is equivalent to a matrix in reduced echelon form. In other words, given *any* matrix, there

exists a sequence of elementary row operations you can apply to it which will, in the end, result in a matrix in RREF. Notice that this fact implies that to each matrix, there are special positions, called **pivot positions** which are **locations of leading entries in the RREF** of a given matrix. For example, the matrix

$$\begin{bmatrix} 0 & 0 & 4 \\ 1 & 0 & 2 \\ 0 & 0 & 0 \end{bmatrix}$$

has pivot positions, indicated by ■, in the locations $\begin{bmatrix} \blacksquare & 0 & 0 \\ 0 & 0 & \blacksquare \\ 0 & 0 & 0 \end{bmatrix}$, namely in

the (1, 1) and (2, 3) spots. Similarly, **pivot columns** of a matrix are columns containing pivot positions. So, in the preceding example, columns 1 and 3 are pivot columns. Notice that, in this example, 0 can occupy a pivot position since pivot position is defined by the equivalent *reduced form* of a given matrix, not necessarily by the *current* form of the matrix.

The Reduction Algorithm

We can now describe an algorithm which will be used to determine if a given linear system is consistent and, if so, what its solution set is. Namely, the algorithm will find a matrix's unique RREF form by using elementary row operations.

DEFINITION 6: REDUCTION ALGORITHM

Apply the following to A in order to produce the unique RREF matrix equivalent to A .

1. Select the leftmost nonzero column, if one exists. This is a pivot column. If none exists, the matrix is already in RREF, so stop.
2. Select (possibly using row interchanges) a nonzero entry from this pivot column to place into the pivot position. This entry is called a **pivot**.
3. Use replacement operations to produce zeros in the pivot column below the pivot.
4. Apply the above steps 1–3 on the submatrix of rows to which the above has not already been applied.

After applying the above, the matrix will be in REF. Since, at each pass through steps 1–3 a new leftmost column is selected, this is called the **forward phase** of the algorithm. Now to get it into RREF we must apply the **backwards phase** to the output.

1. Select the rightmost pivot position, if one exists. If needed, use

a scaling operation to ensure that the pivot occupying the pivot position is 1. If none exists, stop.

2. Use replacement operations with the pivot you've selected to create zeros above the pivot.
3. Apply the above two steps on the submatrix of rows for which the above has not yet been applied.

The following demonstrates the algorithm.

EXAMPLE 7

The algorithm can be applied to any matrix but for this example, we'll

apply it to the augmented matrix $\left[\begin{array}{ccc|c} 1 & -2 & 0 & 5 \\ 7 & -8 & 1 & 5 \\ 2 & 8 & 5 & -53 \end{array} \right]$. To begin, the

first column happens to be the leftmost nonzero column. Its pivot position already contains 1 as a pivot. We apply replacement operations, using that pivot to create zeros in the pivot column.

$$\left[\begin{array}{ccc|c} 1 & -2 & 0 & 5 \\ 7 & -8 & 1 & 5 \\ 2 & 8 & 5 & -53 \end{array} \right] \xrightarrow{R_2-7R_1} \left[\begin{array}{ccc|c} 1 & -2 & 0 & 5 \\ 0 & 6 & 1 & -30 \\ 2 & 8 & 5 & -53 \end{array} \right] \xrightarrow{R_3-2R_1} \left[\begin{array}{ccc|c} 1 & -2 & 0 & 5 \\ 0 & 6 & 1 & -30 \\ 0 & 12 & 5 & -63 \end{array} \right]$$

Next, we apply the algorithm to the submatrix $\left[\begin{array}{cc|c} 6 & 1 & -30 \\ 12 & 5 & -63 \end{array} \right]$, namely we have

$$\left[\begin{array}{ccc|c} 1 & -2 & 0 & 5 \\ 0 & 6 & 1 & -30 \\ 0 & 12 & 5 & -63 \end{array} \right] \xrightarrow{R_3-2R_2} \left[\begin{array}{ccc|c} 1 & -2 & 0 & 5 \\ 0 & 6 & 1 & -30 \\ 0 & 0 & 3 & -3 \end{array} \right]$$

We next apply to algorithm to the submatrix $[3 \mid -3]$ for which there's nothing to apply. So we have now completed the forward phase.

Notice that our resulting matrix, $\left[\begin{array}{ccc|c} 1 & -2 & 0 & 5 \\ 0 & 6 & 1 & -30 \\ 0 & 0 & 3 & -3 \end{array} \right]$ is in row echelon

form. We proceed to the backwards phase. The rightmost pivot is the element 3 in the third column.

$$\left[\begin{array}{ccc|c} 1 & -2 & 0 & 5 \\ 0 & 6 & 1 & -30 \\ 0 & 0 & 3 & -3 \end{array} \right] \xrightarrow{\frac{1}{3}R_3} \left[\begin{array}{ccc|c} 1 & -2 & 0 & 5 \\ 0 & 6 & 1 & -30 \\ 0 & 0 & 1 & -1 \end{array} \right] \xrightarrow{R_2-R_1} \left[\begin{array}{ccc|c} 1 & -2 & 0 & 5 \\ 0 & 6 & 0 & -29 \\ 0 & 0 & 1 & -1 \end{array} \right]$$

Now we go through the steps again with the next most rightmost

pivot, which happens to be the 6. This gives

$$\left[\begin{array}{ccc|c} 1 & -2 & 0 & 5 \\ 0 & 6 & 0 & -29 \\ 0 & 0 & 1 & -1 \end{array} \right] \xrightarrow{\frac{1}{6}R_2} \left[\begin{array}{ccc|c} 1 & -2 & 0 & 5 \\ 0 & 1 & 0 & -\frac{29}{6} \\ 0 & 0 & 1 & -1 \end{array} \right] \xrightarrow{R_1+2R_2} \left[\begin{array}{ccc|c} 1 & 0 & 0 & 5 - \frac{29}{3} \\ 0 & 1 & 0 & -\frac{29}{6} \\ 0 & 0 & 1 & -1 \end{array} \right]$$

which equals $\left[\begin{array}{ccc|c} 1 & 0 & 0 & -\frac{14}{3} \\ 0 & 1 & 0 & -\frac{29}{6} \\ 0 & 0 & 1 & -1 \end{array} \right]$. Applying the algorithm one last time

to the last remaining pivot we see that it's already a 1 and we can stop.

Notice that the result $\left[\begin{array}{ccc|c} 1 & 0 & 0 & -\frac{14}{3} \\ 0 & 1 & 0 & -\frac{29}{6} \\ 0 & 0 & 1 & -1 \end{array} \right]$ is in RREF.

The prior example illustrates the utility of the reduction algorithm in solving linear systems. Namely, to solve the linear system

$$\begin{aligned} x - 2y &= 5 \\ 7x - 8y + z &= 5 \\ 2x + 8y + 5z &= -53 \end{aligned}$$

we simply apply the reduction algorithm to the associated augmented matrix,

$$\left[\begin{array}{ccc|c} 1 & -2 & 0 & 5 \\ 7 & -8 & 1 & 5 \\ 2 & 8 & 5 & -53 \end{array} \right] \sim \left[\begin{array}{ccc|c} 1 & 0 & 0 & -\frac{14}{3} \\ 0 & 1 & 0 & -\frac{29}{6} \\ 0 & 0 & 1 & -1 \end{array} \right] \text{ which in this case corresponds to}$$

$$x = -\frac{14}{3}, y = -\frac{29}{6}, z = -1.$$

Suppose that we have applied the reduction algorithm to an augmented matrix for a linear system and the result is, say,

$$\left[\begin{array}{ccc|c} 1 & 0 & 2 & 6 \\ 0 & 1 & 5 & 7 \\ 0 & 0 & 0 & 0 \end{array} \right]$$

Which corresponds to the system of equations

$$\begin{aligned} x + 2z &= 6 \\ y + 5z &= 7 \end{aligned}$$

The variables, x and z which appear in the pivot columns are referred to as **basic variables** (sometimes also called a *leading variable*) while the variable in the non-pivot column, y is a **free variable**. We've seen how to deal with cases like this before, namely we expressed the solutions in terms of a parameter by **expressing the basic variables in terms of the free variables**. Namely, the solution set is given by

$$\begin{cases} x = 6 - 2t \\ y = 7 - 5t \\ z = t, \quad t \in \mathbb{R} \end{cases}$$

DEFINITION 7: SOLVING SYSTEMS BY REDUCTION ALGORITHM

Given a linear system

$$\begin{array}{cccc} a_{11}x_1 + a_{12}x_2 + \cdots + a_{1n}x_n = b_1 \\ a_{21}x_1 + a_{22}x_2 + \cdots + a_{2n}x_n = b_2 \\ \vdots \qquad \qquad \qquad \vdots \qquad \qquad \qquad \vdots \\ a_{m1}x_1 + a_{m2}x_2 + \cdots + a_{mn}x_n = b_m \end{array}$$

construct the corresponding augmented matrix $[A \mid \mathbf{b}]$, where A is the coefficient matrix of the system and \mathbf{b} is the constant matrix of the same system. To solve the system do the following:

1. Apply the reduction algorithm to the augmented matrix $[A \mid \mathbf{b}]$ to produce an RREF matrix R
2. If the RREF matrix R has a row of the form $[0 \ 0 \ \cdots \ 0 \mid 1]$ the original system is inconsistent
3. Otherwise express all basic variables in terms of any free variables which may appear. If there are no free variables the solution is unique otherwise express your answer in a parameterized form (with free variables acting as parameters).

We remark that in the above we've observed that if the RREF of the augmented matrix corresponding to a linear system contains a row like $[0 \ 0 \ \cdots \ 0 \mid 1]$ then the original system must be inconsistent. After all, such a row would correspond to an equation of the form $0x_1 + 0x_2 + \cdots + 0x_n = 1$, an impossibility. We can make this a little more concrete in the following theorem.

THEOREM 8

A linear system is consistent if and only if the last column of the associated augmented matrix is not a pivot column

PROOF

First let's assume that the given linear system is consistent and prove that the last column of the associated augmented matrix is not a pivot column. This is logically equivalent to saying that we want to show that if the last column of the associated augmented matrix *is* a pivot column then the system must be inconsistent. Well, if the last column is a pivot column then the RREF of the augmented matrix has a row that looks like $[0 \ 0 \ \cdots \ 0 \mid 1]$, and we've seen above that this corresponds to a nonsolvable equation, so the system must be inconsistent.

Next, if the last column of the associated augmented matrix is not a pivot column there can never be a row like $[0 \ 0 \ \cdots \ 0 \ | \ 1]$ in the reduced form of the augmented matrix and we can proceed to step 3 of the outline above for solving systems. \square

The reduced form of an augmented matrix allows us to determine some answers to questions of the size of the solution set for a linear system. The next theorem makes this precise.

THEOREM 9

Suppose that a linear system of m equations in n variables has an associated augmented matrix which has p pivots positions. Then

1. If $n = p$ the system has a unique solution
2. If $p < n$ the system has infinitely many solutions

PROOF

If $p < n$ then there are $0 < n - p$ free variables, each of which must appear as a parameter in the solution set. If there are any parameters in the solution set there must be infinitely many solutions. If $p = n$ then every variable is a basic variable and there can be no parameters, so the solution is unique. \square

Phrasing the above slightly differently, the solution is unique when every variable is a basic variable.

Exercises

The following are exercises for you to practice offline for more practice. Not working these homework questions is a recipe for disaster.

1. Give three examples of linear equations which have zero, one, and infinitely many solutions respectively.
2. Solve $x + 2y = 3$, leaving your answer in parameterized form.
3. Solve $x + 2y + 3z = 4$ leaving your answer in parameterized form.
4. Prove that the Elementary Manipulations of Systems result in equivalent systems.
5. Use Elementary Manipulations of Systems to solve the following, leaving your answer in parameterized form where appropriate.

(a)
$$\begin{aligned} 4x + y &= 0 \\ 16x + 4y &= 0 \end{aligned}$$

(b) $x + 2y = 3$
 $2x + 5y = 6$

(c) $4x + y + 6z = 0$
 $5x + y = 1$

(d) $4x + y + 6z = 0$
 $2x + 2y + 2z = 9$
 $5x + y = 1$

6. Suppose I have a system of three equations in three unknowns. Can such a system have exactly 3 distinct solutions? Can it have exactly 2 distinct solutions? Explain why or why not.
7. Show that the elementary row operations each have an *inverse* row operation of the same type. In other words, show that any of the elementary row operations can be undone by an elementary row operation of the same type.

Geometry of Vectors

Contents

<i>The Dot Product in \mathbb{R}^2</i>	40
<i>The Dot Product in \mathbb{R}^n</i>	42
<i>Projections and Expansions</i>	47
<i>Exercises</i>	49

IN THIS CHAPTER, we see how to “geometrize” the world in which vectors live, even when we cannot directly visualize them. As the material becomes more abstract, we are going to need to rely more and more on the intuitions we develop in this simple setting which is a generalization of the geometry we’ve done in high school.

The Dot Product in \mathbb{R}^2

To begin with, we recall that $\mathbb{R}^2 = \{(a_1, a_2) \mid a_1, a_2 \in \mathbb{R}\}$, namely *ordered pairs* of real numbers. We often *identify* these ordered pairs of numbers with column vectors, i.e. we can represent \mathbb{R}^2 as $\left\{ \begin{bmatrix} a_1 \\ a_2 \end{bmatrix} \mid a_1, a_2 \in \mathbb{R} \right\}$. In class and

in the text we've defined the **dot product** between two vectors $\mathbf{u} = \begin{bmatrix} u_1 \\ u_2 \end{bmatrix}$,

$\mathbf{v} = \begin{bmatrix} v_1 \\ v_2 \end{bmatrix}$ in \mathbb{R}^2 to be the number given by

$$\mathbf{u} \cdot \mathbf{v} = u_1 v_1 + u_2 v_2$$

I'm going to show you that this dot product has a *geometric interpretation*. Let's first recall the *law of cosines* from high school trigonometry, which is a generalization of the law of Pythagoras²¹

²¹ Which says that for right triangles with sides a, b and hypotenuse c then $c^2 = a^2 + b^2$.

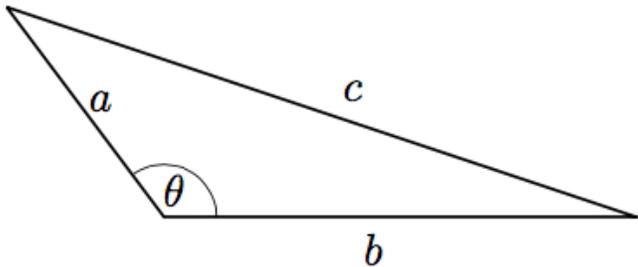
THEOREM 10: LAW OF COSINES

Consider a triangle with side lengths a, b , and c with angle θ opposite to side with length c . Then

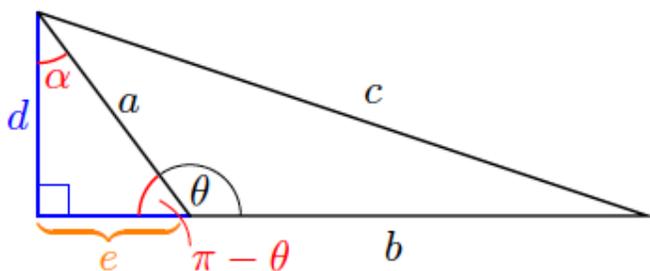
$$c^2 = a^2 + b^2 - 2ab \cos \theta$$

PROOF

To prove the law of cosines, we'll consider only the case of an obtuse angle θ only, the case where θ is acute is proven similarly. Consider the following picture



Then, extend the above triangle to sit inside of a right triangle as in the following image



Notice that $\alpha + (\pi - \theta) = \frac{\pi}{2}$, since the sum of the interior angles must add up to 180 degrees (π radians). This means that $\theta - \frac{\pi}{2} = \alpha$. Then, by definition, we have $\sin \alpha = \frac{e}{a}$ i.e. $e = a \sin \alpha = a \sin(\theta - \frac{\pi}{2}) = -a \cos \theta$. As well, $\cos \alpha = \frac{d}{a}$ gives $d = a \cos(\theta - \frac{\pi}{2}) = a \sin \theta$. Then, since the larger triangle pictured above is a right triangle we have that

$$\begin{aligned} c^2 &= d^2 + (e + b)^2 \\ &= d^2 + e^2 + 2eb + b^2 \\ &= (a \sin \theta)^2 + (-a \cos \theta)^2 + 2b(-a \cos \theta) + b^2 \\ &= a^2(\cos^2 \theta + \sin^2 \theta) + b^2 - 2ab \cos \theta \\ &= a^2 + b^2 - 2ab \cos \theta \end{aligned}$$

The last line followed from the fact that $\sin^2 \theta + \cos^2 \theta = 1$. This is the desired result. \square

We also recall the *length* of a vector in \mathbb{R}^2 is defined as $\|\mathbf{v}\| = \sqrt{\mathbf{v} \cdot \mathbf{v}}$.

CAUTION!

Don't forget that when calculating $\|\mathbf{v}\|$ you must take a square root.

I.e. $\left\| \begin{bmatrix} 1 \\ 2 \end{bmatrix} \right\| = \sqrt{1^2 + 2^2} = \sqrt{5}$. Many students forget to take the square root when doing this calculation.

We can then obtain the following result helping to give the dot product a geometric interpretation.

THEOREM 11

If \mathbf{v} and \mathbf{w} are two vectors in \mathbb{R}^2 then

$$\mathbf{v} \cdot \mathbf{w} = \|\mathbf{v}\| \|\mathbf{w}\| \cos \theta$$

where θ is the angle between the vectors \mathbf{v} and \mathbf{w} .

PROOF

The law of cosines shows that $\|\mathbf{v} - \mathbf{w}\|^2 = \|\mathbf{v}\|^2 + \|\mathbf{w}\|^2 - 2\|\mathbf{v}\|\|\mathbf{w}\|\cos\theta$. But direct calculation reveals

$$\begin{aligned}\|\mathbf{v} - \mathbf{w}\|^2 &= (\mathbf{v} - \mathbf{w}) \cdot (\mathbf{v} - \mathbf{w}) \\ &= \mathbf{v} \cdot \mathbf{v} + \mathbf{w} \cdot \mathbf{w} - \mathbf{v} \cdot \mathbf{w} - \mathbf{w} \cdot \mathbf{v} \\ &= \|\mathbf{v}\|^2 + \|\mathbf{w}\|^2 - 2\mathbf{v} \cdot \mathbf{w}\end{aligned}$$

Comparing these two formulas gives the claimed identity. \square

The above is a nice formula relating the dot product to *geometrical* information about the vectors in question. Namely the dot product encodes information about the lengths of the vectors being “dotted” along with information about the angle between them.

The above is a nice formula relating the dot product to *geometrical* information about the vectors in question. Namely the dot product encodes the lengths of the vectors being “dotted” along with information about the angle between them. Provided that neither $\mathbf{v} \neq \mathbf{0}$, $\mathbf{w} \neq \mathbf{0}$ we have that the angle between \mathbf{v} and \mathbf{w} will satisfy

$$\cos\theta = \frac{\mathbf{v} \cdot \mathbf{w}}{\|\mathbf{v}\|\|\mathbf{w}\|}$$

Therefore, if $\theta \in (0, \frac{\pi}{2}) \cup (\frac{3\pi}{2}, 2\pi)$ we'll have $\mathbf{v} \cdot \mathbf{w} > 0$ and if $\theta \in (\frac{\pi}{2}, \frac{3\pi}{2})$ we'll have $\mathbf{v} \cdot \mathbf{w} < 0$. But also we have the fact that $\mathbf{v} \cdot \mathbf{w}$ will be zero if and only if the angle between \mathbf{v} and \mathbf{w} is right. When \mathbf{v} and \mathbf{w} have a right angle between them we say that they are **orthogonal** to each other. In other words we have proven that

THEOREM 12

Two vectors $\mathbf{v}, \mathbf{w} \in \mathbb{R}^2$ are orthogonal if and only if $\mathbf{v} \cdot \mathbf{w} = 0$.

Notice that this definition shows the zero vector $\mathbf{0}$ is trivially orthogonal to every vector.

The Dot Product in \mathbb{R}^n

In this section we'll derive a few fundamental results for vectors living in the larger geometrical space, \mathbb{R}^n . To begin with, the space \mathbb{R}^n is the collection of ordered lists of n real numbers, $\mathbb{R}^n = \{(a_1, a_2, \dots, a_n) \mid a_1, a_2, \dots, a_n \in \mathbb{R}\}$ which we can represent as $n \times 1$ column vectors, i.e. we think of \mathbb{R}^n as being

the set $\left\{ \begin{bmatrix} a_1 \\ a_2 \\ \vdots \\ a_n \end{bmatrix} \mid a_1, a_2, \dots, a_n \in \mathbb{R} \right\}$. We extend the definition of dot product onto

the space \mathbb{R}^n as follows: if $\mathbf{v} = \begin{bmatrix} v_1 \\ \vdots \\ v_n \end{bmatrix}$, $\mathbf{w} = \begin{bmatrix} w_1 \\ \vdots \\ w_n \end{bmatrix} \in \mathbb{R}^n$ then

$$\mathbf{v} \cdot \mathbf{w} = v_1 w_1 + \cdots + v_n w_n$$

An **important point** for you to keep in mind, which can help in various proofs and/or calculations is that

$$\mathbf{x}^T \mathbf{y} = \mathbf{y}^T \mathbf{x} = [\mathbf{x} \cdot \mathbf{y}]$$

On the LHS above we have matrix multiplication of $1 \times n$ and $n \times 1$ matrices resulting in a 1×1 matrix on the right, whose sole entry is the dot product of the two $n \times 1$ vectors. For this reason, *par abus de langage*, we often simply write

$$\mathbf{x}^T \mathbf{y} = \mathbf{x} \cdot \mathbf{y}$$

CAUTION!

We should **keep in mind** that, strictly speaking, the above cannot be correct since the LHS is a 1×1 matrix whereas the RHS is a number. We will often use this (admittedly somewhat confusing) notation.

And, just as before we have that for all $\mathbf{v}, \mathbf{w} \in \mathbb{R}^n$, $\mathbf{v} \cdot \mathbf{w} = \mathbf{w} \cdot \mathbf{v}$ holds. For all scalars $c \in \mathbb{R}$ we have $(c\mathbf{v}) \cdot \mathbf{w} = \mathbf{v} \cdot (c\mathbf{w}) = c(\mathbf{v} \cdot \mathbf{w})$. And $\mathbf{v} \cdot \mathbf{v} \geq 0$ always holds with $\mathbf{v} \cdot \mathbf{v} = 0$ if and only if $\mathbf{v} = \mathbf{0}$. The last fact allows us to define the **length** (or “norm”) of a vector $\mathbf{v} \in \mathbb{R}^n$ as

$$\|\mathbf{v}\| = \sqrt{\mathbf{v} \cdot \mathbf{v}}$$

You should convince yourself of the fact that $\|c\mathbf{v}\| = |c|\|\mathbf{v}\|$ holds for all $c \in \mathbb{R}$ and all $\mathbf{v} \in \mathbb{R}^n$ as a check on your understanding of the notations. In addition, we define the **unit vector parallel to $\mathbf{v} \neq \mathbf{0}$** as $\hat{\mathbf{v}} = \frac{\mathbf{v}}{\|\mathbf{v}\|}$. You should verify that, in fact, $\hat{\mathbf{v}}$ has length 1. The procedure used to create $\hat{\mathbf{v}}$ from \mathbf{v} , namely taking a nonzero vector and dividing by its length to produce a new nonzero vector parallel to the original one, is called *unitizing* or *normalizing* the vector and $\hat{\mathbf{v}}$ is called the **unit vector** in direction \mathbf{v} .

EXAMPLE 8

We calculate $\widehat{\begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix}}$. Notice that $\left\| \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix} \right\| = \sqrt{1^2 + 2^2 + 3^2} = \sqrt{14}$. Thus

$$\widehat{\begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix}} = \begin{bmatrix} \frac{\sqrt{14}}{14} \\ \frac{2\sqrt{14}}{14} \\ \frac{3\sqrt{14}}{14} \end{bmatrix}$$

Our first result is an extension of the law of cosines to \mathbb{R}^n and is considered one of the most important inequalities in mathematics.²²

²² The Cauchy-Schwarz inequality is foundational in mathematics as it allows one to derive many other important inequalities used in a host of areas.

THEOREM 13: CAUCHY-SCHWARZ INEQUALITY

For all vectors $\mathbf{v}, \mathbf{w} \in \mathbb{R}^n$ we have the following inequality holds

$$|\mathbf{v} \cdot \mathbf{w}| \leq \|\mathbf{v}\| \|\mathbf{w}\| \quad (0.0.4)$$

Equality holds when and only when \mathbf{v} and \mathbf{w} are scalar multiples of each other.

PROOF

First of all, if either (or both) of \mathbf{v}, \mathbf{w} are equal to the zero vector then the result is immediate. Therefore, we can assume that both \mathbf{v} and \mathbf{w} are not equal to $\mathbf{0}$. Then, since length of vectors is always non-negative, we have

$$\begin{aligned} 0 &\leq \|(\mathbf{w}\|\mathbf{v}\| \pm \mathbf{v}\|\mathbf{w}\|)\|^2 \\ &= (\mathbf{w}\|\mathbf{v}\| \pm \mathbf{v}\|\mathbf{w}\|) \cdot (\mathbf{w}\|\mathbf{v}\| \pm \mathbf{v}\|\mathbf{w}\|) \\ &= (\mathbf{w}\|\mathbf{v}\|) \cdot (\mathbf{w}\|\mathbf{v}\|) + (\mathbf{w}\|\mathbf{v}\|) \cdot (\pm \mathbf{v}\|\mathbf{w}\|) \\ &\quad + (\pm \mathbf{v}\|\mathbf{w}\|) \cdot (\mathbf{w}\|\mathbf{v}\|) + (\mathbf{v}\|\mathbf{w}\|) \cdot (\mathbf{v}\|\mathbf{w}\|) \\ &= \|\mathbf{w}\|^2 \|\mathbf{v}\|^2 \pm 2\mathbf{w} \cdot \mathbf{v} \|\mathbf{v}\| \|\mathbf{w}\| \pm 2\mathbf{v} \cdot \mathbf{w} \|\mathbf{w}\| \|\mathbf{v}\| + \|\mathbf{v}\|^2 \|\mathbf{w}\|^2 \\ &= 2\|\mathbf{v}\|^2 \|\mathbf{w}\|^2 \pm 2(\mathbf{v} \cdot \mathbf{w}) \|\mathbf{v}\| \|\mathbf{w}\| \end{aligned}$$

Which gives the following two inequalities,

$$(\mathbf{v} \cdot \mathbf{w}) \|\mathbf{v}\| \|\mathbf{w}\| \leq \|\mathbf{v}\|^2 \|\mathbf{w}\|^2, \quad -(\mathbf{v} \cdot \mathbf{w}) \|\mathbf{v}\| \|\mathbf{w}\| \leq \|\mathbf{v}\|^2 \|\mathbf{w}\|^2$$

Since we've assumed that neither \mathbf{v} nor \mathbf{w} is zero we can divide by the terms $\|\mathbf{v}\| \|\mathbf{w}\|$ to get

$$\mathbf{v} \cdot \mathbf{w} \leq \|\mathbf{v}\| \|\mathbf{w}\|, \quad -\mathbf{v} \cdot \mathbf{w} \leq \|\mathbf{v}\| \|\mathbf{w}\|$$

But that's the same as saying $-\|\mathbf{v}\| \|\mathbf{w}\| \leq \mathbf{v} \cdot \mathbf{w} \leq \|\mathbf{v}\| \|\mathbf{w}\|$, i.e. $|\mathbf{v} \cdot \mathbf{w}| \leq \|\mathbf{v}\| \|\mathbf{w}\|$, as claimed.

Next, the claim is the inequality is actually an equality when the vectors are scalar multiples of one another. In other words, if we first assume that $\mathbf{v} = c\mathbf{w}$ for some choice of $c \in \mathbb{R}$ then

$$\begin{aligned} |\mathbf{v} \cdot \mathbf{w}| &= |(c\mathbf{w}) \cdot \mathbf{w}| \\ &= |c(\mathbf{w} \cdot \mathbf{w})| \\ &= |c\|\mathbf{w}\|^2| \\ &= |c|\|\mathbf{w}\|^2 \\ &= |c|\|\mathbf{w}\| \|\mathbf{w}\| \\ &= \|c\mathbf{w}\| \|\mathbf{w}\| \\ &= \|\mathbf{v}\| \|\mathbf{w}\| \end{aligned}$$

which is the claimed result.

If, on the other hand, if we knew that $\mathbf{v} \neq c\mathbf{w}$ for all choices of $c \in \mathbb{R}$, then in particular we would have that $\|\mathbf{v} - c\mathbf{w}\|^2 \neq 0$. Writing this out means that there is no c which solves the quadratic equation

$$\|\mathbf{v}\|^2 - c(2\mathbf{v} \cdot \mathbf{w}) + c^2\|\mathbf{w}\|^2 = 0$$

The quadratic formula tells us the only way for this to happen is for the discriminant of this quadratic polynomial $4(\mathbf{v} \cdot \mathbf{w})^2 - 4\|\mathbf{v}\|^2 \|\mathbf{w}\|^2$ to be negative. Thus, $(\mathbf{v} \cdot \mathbf{w})^2 < \|\mathbf{v}\|^2 \|\mathbf{w}\|^2$ and taking square roots gives $|\mathbf{v} \cdot \mathbf{w}| < \|\mathbf{v}\| \|\mathbf{w}\|$ so the inequality is strict. \square

By analogy to \mathbb{R}^2 , we now define two vectors \mathbf{v} and \mathbf{w} in \mathbb{R}^n to be **orthogonal** if $\mathbf{v} \cdot \mathbf{w} = 0$. It turns out (but we will not prove this) that this definition is well-founded in the sense that vectors in \mathbb{R}^n are at right angles to each other if and only if $\mathbf{v} \cdot \mathbf{w} = 0$. We only proved this for $n = 2$ but it turns out to still be true in \mathbb{R}^n generally.

We define the obvious extension of orthogonality and orthonormality to *sets* of vectors.

DEFINITION 14: ORTHOGONAL AND ORTHONORMAL SETS

$S = \{\mathbf{v}_1, \dots, \mathbf{v}_k\} \subset \mathbb{R}^n$ is said to be an *orthogonal set* of vectors in \mathbb{R}^n if $\mathbf{v}_i \cdot \mathbf{v}_j = 0$ whenever $i \neq j$. The set S is said to be an *orthonormal set* of vectors if it's orthogonal and all vectors \mathbf{v}_i have length 1, i.e. $\|\mathbf{v}_i\| = 1$.

We use the Cauchy-Schwarz inequality to prove the next geometric fact about \mathbb{R}^n which is illustrated in the following figure

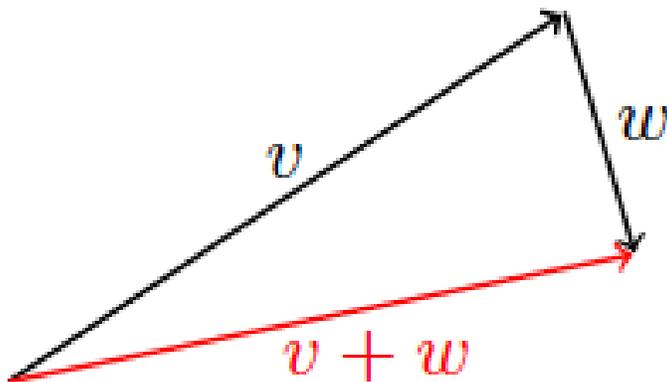


Figure 2: An illustration of one variant of the triangle inequality

THEOREM 15: TRIANGLE INEQUALITY

For all $\mathbf{v}, \mathbf{w} \in \mathbb{R}^n$ we have that $\|\mathbf{v} + \mathbf{w}\| \leq \|\mathbf{v}\| + \|\mathbf{w}\|$

PROOF

Consider

$$\begin{aligned} \|\mathbf{v} + \mathbf{w}\|^2 &= (\mathbf{v} + \mathbf{w}) \cdot (\mathbf{v} + \mathbf{w}) \\ &= \mathbf{v} \cdot \mathbf{v} + \mathbf{v} \cdot \mathbf{w} + \mathbf{w} \cdot \mathbf{v} + \mathbf{w} \cdot \mathbf{w} \\ &= \|\mathbf{v}\|^2 + 2(\mathbf{v} \cdot \mathbf{w}) + \|\mathbf{w}\|^2 \\ &\stackrel{\leq}{\leq} \|\mathbf{v}\|^2 + 2\|\mathbf{v}\| \|\mathbf{w}\| + \|\mathbf{w}\|^2 \\ &\text{by Cauchy Schwarz} \\ &= (\|\mathbf{v}\| + \|\mathbf{w}\|)^2 \end{aligned}$$

In other words, $\|\mathbf{v} + \mathbf{w}\|^2 \leq (\|\mathbf{v}\| + \|\mathbf{w}\|)^2$. The quantities which appear here being squared are *non-negative* which means we can take the positive square root to conclude that

$$\|\mathbf{v} + \mathbf{w}\| \leq \|\mathbf{v}\| + \|\mathbf{w}\|$$

as advertised.

As we've already seen the parallelogram rule for vector addition in lecture, we can give this a geometric interpretation by introducing a **distance** between two vectors $\mathbf{v}, \mathbf{w} \in \mathbb{R}^n$ as

$$d(\mathbf{v}, \mathbf{w}) = \|\mathbf{v} - \mathbf{w}\|$$

You should verify that in \mathbb{R}^2 this does precisely correspond to the distance between the tips of two vectors connecting the geometric points described by \mathbf{v}, \mathbf{w} to the origin. In other words, the definition of distance is very well motivated.

We then can use our work thus far to help conclude the following variation of the triangle inequality above, which is illustrated in the figure above.

THEOREM 16: TRIANGLE INEQUALITY (SECOND VERSION)

For all vectors $\mathbf{u}, \mathbf{v}, \mathbf{w}$ in \mathbb{R}^n we have that

$$d(\mathbf{v}, \mathbf{w}) \leq d(\mathbf{v}, \mathbf{u}) + d(\mathbf{u}, \mathbf{w})$$

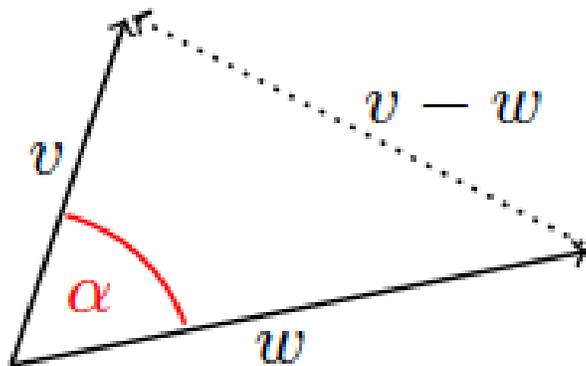


Figure 3: An illustration of another variant of the triangle inequality

PROOF

Consider

$$\begin{aligned}
 d(\mathbf{v}, \mathbf{w}) &= \|\mathbf{v} - \mathbf{w}\| \\
 &= \|(\mathbf{v} - \mathbf{u}) + (\mathbf{u} - \mathbf{w})\| \\
 &\stackrel{\leq}{\leq} \|\mathbf{v} - \mathbf{u}\| + \|\mathbf{u} - \mathbf{w}\| \\
 &\text{By Theorem 2.2} \\
 &= d(\mathbf{v}, \mathbf{u}) + d(\mathbf{u}, \mathbf{w})
 \end{aligned}$$

as advertised. \square

Projections and Expansions

In this section we explore an important use of the dot product which will help us later when we deal with subspaces and bases. We see how the dot product allows us to tell how “similar” two vectors are to one another in a way we’ll make precise.

To begin, consider two vectors \mathbf{v} and $\mathbf{d} \neq \mathbf{0}$ in \mathbb{R}^n . Let’s see if we can *construct* a new vector, $\mathbf{proj}_{\mathbf{d}}\mathbf{v}$, the *projection of vector \mathbf{v} onto vector \mathbf{d}* in such a way that

1. $\mathbf{proj}_{\mathbf{d}}\mathbf{v}$ will be a vector parallel to \mathbf{d}
2. $\mathbf{proj}_{\mathbf{d}}\mathbf{v}$ will have a tip at the *closest* point to \mathbf{v} along the line in direction \mathbf{d} .

Notice, as in Figure 5, we can find where $\mathbf{proj}_{\mathbf{d}}\mathbf{v}$ ends by imagining a line connecting the tip of \mathbf{v} to the line parallel to \mathbf{d} in such a way that these meet at right angles. Rephrasing this symbolically, we can pose this as an

algebraic constraint on the vector $\mathbf{proj}_d \mathbf{v}$ as demanding that

$$\mathbf{d} \cdot (\mathbf{v} - \mathbf{proj}_d \mathbf{v}) = 0$$

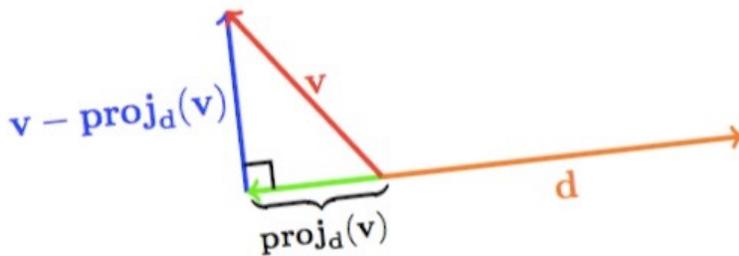
Subject to

$$\mathbf{proj}_d \mathbf{v} = c\mathbf{d}$$

for an, as yet unknown, constant c .²³

²³ After all, this is simply a *precise* restatement, in mathematical form, of the verbal description in the preceding paragraph, or in figure 5.

Figure 4: An illustration of the vector projection of two vectors in \mathbb{R}^2 .



We can substitute the second of the above into the first to get

$$\begin{aligned} 0 &= \mathbf{d} \cdot (\mathbf{v} - \mathbf{proj}_d \mathbf{v}) \\ &= \mathbf{d} \cdot (\mathbf{v} - c\mathbf{d}) \\ &= \mathbf{d} \cdot \mathbf{v} - c\|\mathbf{d}\|^2 \end{aligned}$$

Solving for c then gives that $c = \frac{\mathbf{d} \cdot \mathbf{v}}{\|\mathbf{d}\|^2}$. Since $\mathbf{proj}_d \mathbf{v} = c\mathbf{d}$ this then gives a formula for the projection in the following definition.

DEFINITION 17: VECTOR PROJECTION

Let \mathbf{v} and $\mathbf{d} \neq \mathbf{0} \in \mathbb{R}^n$. The **projection** of \mathbf{v} onto \mathbf{d} is given by

$$\mathbf{proj}_d \mathbf{v} = \frac{\mathbf{d} \cdot \mathbf{v}}{\|\mathbf{d}\|^2} \mathbf{d}$$

and is the unique vector parallel to \mathbf{d} such that \mathbf{d} and $\mathbf{v} - \mathbf{proj}_d \mathbf{v}$ are orthogonal. We call the scalar $c = \frac{\mathbf{d} \cdot \mathbf{v}}{\|\mathbf{d}\|^2}$ the **component of \mathbf{v} along \mathbf{d}** .

The next theorem verifies that our choice in the above actually *does* minimize the distance from the line parallel to \mathbf{d} to \mathbf{v} .

THEOREM 18

Let \mathbf{v} and $\mathbf{d} \neq \mathbf{0} \in \mathbb{R}^n$. The vector projection $\mathbf{proj}_{\mathbf{d}}\mathbf{v}$ minimizes the distance from \mathbf{v} to the line parallel to \mathbf{d} in the sense that

$$\|\mathbf{v} - \mathbf{proj}_{\mathbf{d}}\mathbf{v}\|^2 < \|\mathbf{v} - c\mathbf{d}\|^2$$

holds for all $c \neq \frac{\mathbf{d} \cdot \mathbf{v}}{\|\mathbf{d}\|^2}$.

PROOF

Consider $\|\mathbf{v} - c\mathbf{d}\|^2 = \|\mathbf{v}\|^2 - 2c\mathbf{v} \cdot \mathbf{d} + c^2\|\mathbf{d}\|^2$. This is a polynomial expression in variable c . Since

$$\|\mathbf{d}\|^2 > 0$$

the polynomial is an upwards opening parabola. The roots of the polynomial are

$$\begin{aligned} c_{\pm} &= \frac{2\mathbf{d} \cdot \mathbf{v} \pm \sqrt{4(\mathbf{d} \cdot \mathbf{v})^2 - 4\|\mathbf{d}\|^2\|\mathbf{v}\|^2}}{2\|\mathbf{d}\|^2} \\ &= \frac{\mathbf{d} \cdot \mathbf{v} \pm \sqrt{(\mathbf{d} \cdot \mathbf{v})^2 - \|\mathbf{d}\|^2\|\mathbf{v}\|^2}}{\|\mathbf{d}\|^2} \end{aligned}$$

The polynomial therefore is minimized at the value of c lying in between the polynomial's two roots i.e. the value c which minimizes the function $f(t) = t^2\|\mathbf{d}\|^2 - 2t\mathbf{d} \cdot \mathbf{v} + \|\mathbf{v}\|^2$ is

$$\begin{aligned} c &= \frac{c_+ + c_-}{2} \\ &= \frac{\mathbf{d} \cdot \mathbf{v}}{\|\mathbf{d}\|^2} \end{aligned}$$

as claimed. \square

Exercises

You should do the following problems offline.

1. Verify that $\|\mathbf{v}\| \geq 0$ with $\|\mathbf{v}\| = 0$ if and only if $\mathbf{v} = \mathbf{0}$.
2. Let $\|\mathbf{v}\| = 2$ and $\|\mathbf{w}\| = 4$ and $\mathbf{w} \cdot \mathbf{v} = 2$. Compute $\|2\mathbf{v} + 17\mathbf{w}\|$.
3. Let $\|\mathbf{v}\| = 2$ and $\|\mathbf{w}\| = 3$ and $\mathbf{v} \cdot \mathbf{w} = 5$. Compute $\|12\mathbf{v} - 7\mathbf{w}\|$.
4. In introducing the Cauchy-Schwarz inequality, it was alleged to generalize the dot product formula involving cosine which we found in \mathbb{R}^2 . Explain how.

5. Explain why the triangle inequality (Theorem 2.2) has its name.
6. Prove that $d(c\mathbf{v}, \mathbf{w}) = |c|d(\mathbf{v}, \frac{1}{c}\mathbf{w})$ holds for all $0 \neq c \in \mathbb{R}$ and $\mathbf{v}, \mathbf{w} \in \mathbb{R}^n$
7. Show that $\mathbf{v} \cdot \mathbf{w} = \frac{1}{4}(\|\mathbf{v} + \mathbf{w}\|^2 - \|\mathbf{v} - \mathbf{w}\|^2)$ for all $\mathbf{v}, \mathbf{w} \in \mathbb{R}^n$.
8. Show that $\|\mathbf{v}\|^2 + \|\mathbf{w}\|^2 = \frac{1}{2}(\|\mathbf{v} + \mathbf{w}\|^2 + \|\mathbf{v} - \mathbf{w}\|^2)$ for all $\mathbf{v}, \mathbf{w} \in \mathbb{R}^n$.
9. Pythagorean theorem in \mathbb{R}^n . Namely, if $\mathbf{v}, \mathbf{w} \in \mathbb{R}^n$ then $\|\mathbf{v} + \mathbf{w}\|^2 = \|\mathbf{v}\|^2 + \|\mathbf{w}\|^2$ if and only if \mathbf{v} and \mathbf{w} are orthogonal.
10. True/False: $\mathbf{v}, \mathbf{w} \in \mathbb{R}^n$ are orthogonal if and only if $\|\mathbf{v} - \mathbf{w}\| = \|\mathbf{v} + \mathbf{w}\|$
11. True/False: $\mathbf{v} + \mathbf{w}$ and $\mathbf{v} - \mathbf{w} \in \mathbb{R}^n$ are orthogonal if and only if $\|\mathbf{v}\| = \|\mathbf{w}\|$
12. Suppose $S = \{\mathbf{v}_1, \dots, \mathbf{v}_k\} \subset \mathbb{R}^n$ is a collection of nonzero mutually orthogonal vectors, i.e. $i \neq j \implies \mathbf{v}_i \cdot \mathbf{v}_j = 0$. Prove that the vectors in S are linearly independent.
13. True/False: All linearly independent sets of vectors in \mathbb{R}^n are orthogonal.
14. Let A be an $m \times n$ matrix. Prove that $A^T A = I_n$ if and only if the columns of A are orthogonal.

15. Show that,²⁴ for all non-negative numbers $x, y \in \mathbb{R}$ we have $\sqrt{xy} \leq \frac{1}{2}(x + y)$.

²⁴ The left hand side of the inequality is called the *geometric mean* of the numbers x , and y , whereas the right hand side is the *arithmetic mean* of the same numbers. This result shows that the geometric mean is bounded by the arithmetic mean. Sometimes this fact is called the "AM-GM inequality".

16. If a_1, \dots, a_n are real numbers then

$$\frac{|a_1 + \dots + a_n|}{\sqrt{n}} \leq \sqrt{a_1^2 + \dots + a_n^2}$$

17. Let a_1, \dots, a_n be real and b_1, \dots, b_n be *positive*. Then prove that

$$\frac{a_1^2}{b_1} + \frac{a_2^2}{b_2} + \dots + \frac{a_n^2}{b_n} \geq \frac{(a_1 + \dots + a_n)^2}{b_1 + \dots + b_n}$$

and that equality holds when and only when $\frac{a_1}{b_1} = \frac{a_2}{b_2} = \dots = \frac{a_n}{b_n}$.

18. Let x, y, z be positive numbers. Prove that

$$\sqrt{x(3x + y)} + \sqrt{y(3y + z)} + \sqrt{z(3z + x)} \leq 2(x + y + z)$$

19. Find $\mathbf{proj}_{\begin{bmatrix} 1 & 2 \end{bmatrix}^T} \begin{bmatrix} 3 \\ 4 \end{bmatrix}$

20. Find $\mathbf{proj}_{\begin{bmatrix} 2 & 1 & 2 \end{bmatrix}^T} \begin{bmatrix} 3 \\ 4 \\ 5 \end{bmatrix}$

21. Find $\mathbf{proj}_{\begin{bmatrix} 3 & -1 & 2 \end{bmatrix}^T} \begin{bmatrix} -3 \\ 4 \\ 7 \end{bmatrix}$

22. Let $\mathbf{v}, \mathbf{d} \in \mathbb{R}^n$ and $\mathbf{d} \neq \mathbf{0}$. Show that $\mathbf{proj}_{\mathbf{d}}(\mathbf{proj}_{\mathbf{d}}\mathbf{v}) = \mathbf{proj}_{\mathbf{d}}\mathbf{v}$
23. Let $\mathbf{0} \neq \mathbf{d} \in \mathbb{R}^n$. Prove that $\mathbf{proj}_{\mathbf{d}}\mathbf{v} = \mathbf{0}$ if and only if \mathbf{v} and \mathbf{d} are orthogonal.
24. (Harder). Suppose $S = \{\mathbf{v}_1, \dots, \mathbf{v}_k\} \subset \mathbb{R}^n$ is an orthogonal set. Let $\mathbf{v} \in \mathbb{R}^n$ be arbitrary and let c_i be the component of \mathbf{v} along \mathbf{v}_i i.e. $c_i = \frac{\mathbf{v}_i \cdot \mathbf{v}}{\|\mathbf{v}_i\|^2}$ for $i = 1, \dots, k$. Show that

$$\|\mathbf{v} - (c_1\mathbf{v}_1 + \dots + c_k\mathbf{v}_k)\| \leq \|\mathbf{v} - (a_1\mathbf{v}_1 + \dots + a_k\mathbf{v}_k)\|$$

holds for all $a_1, \dots, a_k \in \mathbb{R}$

25. (Harder). Suppose $S = \{\mathbf{v}_1, \dots, \mathbf{v}_k\} \subset \mathbb{R}^n$ is an orthonormal set. Let $\mathbf{v} \in \mathbb{R}^n$ be arbitrary and let c_i be the component of \mathbf{v} along \mathbf{v}_i i.e. $c_i = \frac{\mathbf{v}_i \cdot \mathbf{v}}{\|\mathbf{v}_i\|^2}$ for $i = 1, \dots, k$. Show that

$$c_1^2\|\mathbf{v}_1\|^2 + \dots + c_k^2\|\mathbf{v}_k\|^2 \leq \|\mathbf{v}\|^2$$

The Rank Theorems

IN THIS CHAPTER we're going to see how a number, the *rank* of a matrix, can help characterize how "reliable" or "defective" a matrix is. To do this though, there's a lot of technical background needed to make precise what we intuitively mean when we refer to "sizes" of spaces in \mathbb{R}^n .

Subspaces of \mathbb{R}^n

A **subspace** of \mathbb{R}^n is a collection of vectors closed under the algebraic vector operations. Namely $S \subseteq \mathbb{R}^n$ is a subspace if and only if $c_1\mathbf{u} + c_2\mathbf{v} \in S$ holds for all numbers c_1, c_2 whenever $\mathbf{u}, \mathbf{v} \in S$. Spans of vectors are natural examples of subspaces. Since \mathbb{R}^n and $\{\mathbf{0}\}$ (called the **zero subspace**) are always subspaces of \mathbb{R}^n they are “improper” subspaces and any other subspace is said to be *proper*. A subset $T \subset S$ of a subset S is a *subspace of subspace* S if it’s a subset of S and satisfies the subspace criteria.

EXAMPLE 9

We prove that $S = \{t\mathbf{v} \mid \mathbf{0} \neq \mathbf{v} \in \mathbb{R}^n\}$ is a subspace. Consider $\mathbf{u}, \mathbf{w} \in S$. Then $\mathbf{u} = a\mathbf{v}$ for some value of $a \in \mathbb{R}$ and $\mathbf{w} = b\mathbf{v}$ for some $b \in \mathbb{R}$. So $c_1\mathbf{u} + c_2\mathbf{w} = (c_1a + c_2b)\mathbf{v} = \tilde{c}\mathbf{v}$ which must be in S . As c_1, c_2 were arbitrary this means that S is a subspace.

The above example is generalizable in the sense that **spans of vectors are always subspaces**. The reader is encouraged to sit down and prove this claim, it’s time well spent.

In addition there are a few canonical subspaces that arise when working with matrices.

EXAMPLE 10

Let A be an $m \times n$ matrix. The following are canonical subspaces.

1. $col(A) = \text{span}\{\text{columns of } A\}$ is a subspace of \mathbb{R}^m called the **column space** of A .
2. $row(A) = col(A^T)$ is a subspace of \mathbb{R}^n called the **row space** of A .
3. $Nul(A) = \{\mathbf{x} \mid A\mathbf{x} = \mathbf{0}\}$ is a subspace of \mathbb{R}^n called the **null space** of A . This is also denoted $null(A)$ or $\ker(A)$.
4. $E_\lambda(A) = \ker(A - \lambda I)$ for *square* matrices A is the **eigenspace** of A for **eigenvalue** λ . The nonzero elements (if there are any) of $E_\lambda(A)$ are called **eigenvectors** of A with eigenvalue λ .

Each of the above is left as an exercise to the reader to verify that they are in fact subspaces.

Notice that, in the above definition of eigenspace $E_\lambda(A)$ that certainly a priori one should expect there to be constraints on numbers λ for which $E_\lambda(A) \neq \{\mathbf{0}\}$.

We begin with an **extremely fundamental result**.²⁵

²⁵ Although this looks innocuous enough, it will serve as the underpinning of all of our future results on subspaces. Notice that, at its heart, the proof only uses facts about homogeneous systems of linear equations.

THEOREM 19: SUBSPACE THEOREM

Let S be a subspace of \mathbb{R}^n . If S is spanned by m vectors and contains k linearly independent vectors, then $k \leq m$.

PROOF

Suppose that $\{\mathbf{v}_1, \dots, \mathbf{v}_m\}$ is a spanning set and $\{\mathbf{w}_1, \dots, \mathbf{w}_k\}$ is a set of linearly independent vectors. Since the first set is a spanning set we have that $\mathbf{w}_j = c_1\mathbf{v}_1 + \dots + c_m\mathbf{v}_m$ holds for $j = 1, \dots, k$ where for each j we have a new set of constants c_1, \dots, c_m . We can write this as

$$\mathbf{w}_j = c_{1j}\mathbf{v}_1 + c_{2j}\mathbf{v}_2 + \dots + c_{mj}\mathbf{v}_m$$

for constants c_{ij} where $1 \leq i \leq m, 1 \leq j \leq k$. Viewing the c_{ij} 's as entries of an $m \times k$ matrix C we have that the above is equivalent to $W = VC$ for $W = [\mathbf{w}_1 \dots \mathbf{w}_k]$ and $V = [\mathbf{v}_1 \dots \mathbf{v}_m]$. If $k > m$ then C has more columns than rows so $C\mathbf{x} = \mathbf{0}$ has a nontrivial solution \mathbf{x} . From this we have that $W\mathbf{x} = \mathbf{0}$ has a nontrivial solution which implies the columns of W are linearly dependent contrary to assumption. Thus $k \leq m$. \square

The above theorem tells us that for subspaces of \mathbb{R}^n we have that

$$\#(\text{independent vectors in } S) \leq \#(\text{spanning vectors in } S)$$

In the above (and below) I'm using the notation that $\#U$ denotes the number of elements in a finite set U (where $\#U = \infty$ if the set has infinitely many elements).

The above theorem is remarkably powerful. Notice that it is a counting result: one number (number of independent vectors) is guaranteed to be bounded by some other number (number of spanning vectors). The above also tells us that if we could find a collection of linearly independent spanning vectors in a subspace that we would "saturate" the above inequality, i.e. the inequality would be an equality. Bases are sets which have this property.

Bases

As mentioned, a **basis** \mathcal{B} for a subspace $S \subseteq \mathbb{R}^n$ is a collection of linearly independent spanning vectors. So $S = \text{span}\{\mathcal{B}\}$ and all elements of \mathcal{B} are linearly independent. Our first result is to convince you that such things *exist*. If you're not used to thinking this way it might be surprising that this needs to be proven at all, but just because we can *define* something doesn't mean that there's anything guaranteed to meet the criteria in our definition.

THEOREM 20

Let $S \neq \{\mathbf{0}\}$ be a subspace of \mathbb{R}^n . Then there's a basis for S .

PROOF

Since $S \neq \{\mathbf{0}\}$ there's a nonzero vector in S and therefore there are sets of linearly independent vectors within S . Suppose that $\{\mathbf{v}_1, \dots, \mathbf{v}_k\}$ is a set of independent vectors in S chosen to be as large as possible. Set $W = \text{span}\{\mathbf{v}_1, \dots, \mathbf{v}_k\}$. By the subspace property of S we have that $W \subseteq S$. There are two options.

Option 1. $W = S$. There's nothing left to prove since then $\{\mathbf{v}_1, \dots, \mathbf{v}_k\}$ is a basis for S .

Option 2. $W \neq S$. In this case there's a vector $\mathbf{w} \in S$ but not in W . Consider the equation $c_1\mathbf{v}_1 + \dots + c_k\mathbf{v}_k + c_{k+1}\mathbf{w} = \mathbf{0}$. Either $c_{k+1} = 0$ or not. If $c_{k+1} = 0$ then $c_1\mathbf{v}_1 + \dots + c_k\mathbf{v}_k = \mathbf{0}$ which then gives $c_1 = \dots = c_k = 0$ since the set $\{\mathbf{v}_1, \dots, \mathbf{v}_k\}$ is linearly independent. But, in that case we see that $\{\mathbf{v}_1, \dots, \mathbf{v}_k, \mathbf{w}\}$ is a larger linearly independent set than $\{\mathbf{v}_1, \dots, \mathbf{v}_k\}$ contrary to our assumptions on how $\{\mathbf{v}_1, \dots, \mathbf{v}_k\}$ was chosen. Therefore $c_{k+1} \neq 0$. But in that case, then $c_1\mathbf{v}_1 + \dots + c_k\mathbf{v}_k + c_{k+1}\mathbf{w} = \mathbf{0}$ gives $\mathbf{w} \in \text{span}\{\mathbf{v}_1, \dots, \mathbf{v}_k\}$, i.e. $\mathbf{w} \in W$, which is contrary to assumption. So this option can never occur. \square

EXAMPLE 11

Consider \mathbf{e}_i the i 'th columns of the $n \times n$ identity matrix I_n . Then,

since any $\mathbf{x} = \begin{bmatrix} x_1 \\ \vdots \\ x_n \end{bmatrix} \in \mathbb{R}^n$ satisfies $\mathbf{x} = I_n\mathbf{x} = x_1\mathbf{e}_1 + \dots + x_n\mathbf{e}_n$ we

see that $\{\mathbf{e}_1, \dots, \mathbf{e}_n\}$ spans \mathbb{R}^n . Moreover, since the \mathbf{e}_j 's are mutually orthogonal, they form a linearly independent set. So $\{\mathbf{e}_1, \dots, \mathbf{e}_n\}$ is a basis for \mathbb{R}^n .

Our next result is that bases (plural of "basis") have a fixed size.

THEOREM 21

A basis for a subspace $S \subseteq \mathbb{R}^n$ can only have one size.

PROOF

Consider two bases for S , \mathcal{A} and \mathcal{B} . Since \mathcal{A} is a basis and hence spanning we have $S = \text{span}\{\mathcal{A}\}$. But since \mathcal{B} is a basis its vectors are linearly independent so we have, by the subspace theorem, that $\#\mathcal{B} \leq \#\mathcal{A}$. Switching the roles in this argument of \mathcal{A} and \mathcal{B} shows that $\#\mathcal{A} \leq \#\mathcal{B}$. Therefore $\#\mathcal{A} = \#\mathcal{B}$ \square

The above result then allows us to unambiguously define the **dimension** of a subspace.

DEFINITION 3

The **dimension** $\dim S$ of a nonzero subspace $S \subseteq \mathbb{R}^n$ is $\#\mathcal{B}$ for any basis \mathcal{B} of S .

Since the zero subspace cannot admit any linearly independent vectors (see exercises), we simply *declare* that

$$\dim\{\mathbf{0}\} = 0$$

Notice that we have freedom in choosing whatever basis we like when calculating the dimension of a given subspace since the amount of vectors in *every* basis must necessarily be the same.

As another example, since, as we've seen, $\{\mathbf{e}_1, \dots, \mathbf{e}_n\}$ is a basis for \mathbb{R}^n we have that

$$\dim(\mathbb{R}^n) = n$$

EXAMPLE 12

Let $S = \left\{ \begin{bmatrix} s \\ t+s \\ t \end{bmatrix} \mid s, t \in \mathbb{R} \right\}$. Prove that S is a subspace and find its dimension.

Solution: Notice that $S = \left\{ \begin{bmatrix} s \\ s \\ 0 \end{bmatrix} + \begin{bmatrix} 0 \\ t \\ t \end{bmatrix} \mid s, t \in \mathbb{R} \right\} = \left\{ s \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix} + t \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix} \mid$

$s, t \in \mathbb{R}$ so $S = \text{span} \left\{ \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix} \right\}$. Since spans of vectors are subspaces,

this means S is a subspace and $\left\{ \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix} \right\}$ is a spanning set. Notice

that you can easily verify $c_1 \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix} + c_2 \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix} = \mathbf{0} \implies c_1 = c_2 = 0$

so $\left\{ \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix} \right\}$ are linearly independent. Thus $\left\{ \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix} \right\}$ is a basis

for the subspace S . Since there are 2 vectors in this basis we have $\dim S = 2$.

Expansions and Orthogonalization

Given a basis $\mathcal{B} = \{\mathbf{b}_1, \dots, \mathbf{b}_k\}$ for a subspace $S \subseteq \mathbb{R}^n$ how do we actually *represent* a given vector $\mathbf{v} \in S$ in terms of the vectors $\mathbf{b}_1, \dots, \mathbf{b}_k$? Well, we can make this problem much more easily tractable if we assume that the

vectors $\mathbf{b}_1, \dots, \mathbf{b}_k$ are mutually orthogonal. In that case, notice that since $\mathbf{v} \in S = \text{span}\{\mathcal{B}\}$ we must have that there are numbers b_1, \dots, b_k such that

$$\mathbf{v} = b_1\mathbf{b}_1 + \cdots + b_k\mathbf{b}_k$$

It remains to determine the numbers b_1, \dots, b_k . But that's simple since we can just use orthogonality of the vectors \mathbf{b}_j as in

$$\begin{aligned} \mathbf{b}_j \cdot \mathbf{v} &= \mathbf{v}_j \cdot (b_1\mathbf{b}_1 + \cdots + b_k\mathbf{b}_k) \\ &= b_1\mathbf{v} \cdot \mathbf{b}_1 + b_2\mathbf{v} \cdot \mathbf{b}_2 + \cdots + b_k\mathbf{v} \cdot \mathbf{b}_k \\ &= 0 + 0 + \cdots + 0 + b_j\mathbf{b}_j \cdot \mathbf{b}_j + 0 + \cdots + 0 \\ &= b_j\|\mathbf{b}_j\|^2 \end{aligned}$$

This means that the coefficient must be given by

$$b_j = \frac{\mathbf{b}_j \cdot \mathbf{v}}{\|\mathbf{b}_j\|^2}$$

Therefore, given the basis vectors $\mathbf{b}_1, \dots, \mathbf{b}_k$ for a subspace S we can uniquely express a vector $\mathbf{v} \in S$ in terms of these basis vectors as

$$\mathbf{v} = \frac{\mathbf{b}_1 \cdot \mathbf{v}}{\|\mathbf{b}_1\|^2}\mathbf{b}_1 + \frac{\mathbf{b}_2 \cdot \mathbf{v}}{\|\mathbf{b}_2\|^2}\mathbf{b}_2 + \cdots + \frac{\mathbf{b}_k \cdot \mathbf{v}}{\|\mathbf{b}_k\|^2}\mathbf{b}_k$$

The above should look familiar. After all, we can rewrite it as

$$\mathbf{v} = \mathbf{proj}_{\mathbf{b}_1}\mathbf{v} + \mathbf{proj}_{\mathbf{b}_2}\mathbf{v} + \cdots + \mathbf{proj}_{\mathbf{b}_k}\mathbf{v}$$

where you'll recall that $\mathbf{proj}_{\mathbf{b}_j}\mathbf{v} = \frac{\mathbf{b}_j \cdot \mathbf{v}}{\|\mathbf{b}_j\|^2}\mathbf{b}_j$ defines the vector projection of \mathbf{v} onto \mathbf{b}_j .

The above representations of \mathbf{v} in terms of basis vectors are called *expansion formulas* since they "expand" the vector \mathbf{v} in terms of the basis vectors \mathbf{b}_j .

All of the above was predicated on the assumption of orthogonality of the basis vectors $\{\mathbf{b}_1, \dots, \mathbf{b}_k\}$. The formula clearly²⁶ won't work if the basis vectors fail to be orthogonal. What should we do if this happens to be the case? The approach we'll take is the following: given any basis \mathcal{B}_1 for a subspace $S \subseteq \mathbb{R}^n$ we can produce, via an algorithm, another basis \mathcal{B}_2 for the same subspace which has the property that its vectors are orthonormal. The algorithm which we use to produce an orthonormal basis from an arbitrary basis is called the *Gram-Schmidt orthonormalization* procedure.

²⁶ Try to verify why

THEOREM 22: GRAM-SCHMIDT ORTHONORMALIZATION PROCEDURE

Let $\{\mathbf{b}_1, \dots, \mathbf{b}_k\} \subseteq \mathbb{R}^n$ be a basis for a subspace S . Define

$$\mathbf{w}_1 = \hat{\mathbf{b}}_1$$

$$\mathbf{w}_2 = \hat{\mathbf{x}}_2, \quad \mathbf{x}_2 \doteq \mathbf{b}_2 - \text{proj}_{\mathbf{w}_1} \mathbf{b}_2$$

$$\mathbf{w}_3 = \hat{\mathbf{x}}_3, \quad \mathbf{x}_3 \doteq \mathbf{b}_3 - \text{proj}_{\mathbf{w}_2} \mathbf{b}_3 - \text{proj}_{\mathbf{w}_1} \mathbf{b}_3$$

$$\vdots$$

$$\mathbf{w}_k = \hat{\mathbf{x}}_k, \quad \mathbf{x}_k \doteq \mathbf{b}_k - \text{proj}_{\mathbf{w}_{k-1}} \mathbf{b}_k - \text{proj}_{\mathbf{w}_{k-2}} \mathbf{b}_k - \dots - \text{proj}_{\mathbf{w}_2} \mathbf{b}_k - \text{proj}_{\mathbf{w}_1} \mathbf{b}_k$$

The vectors $\{\mathbf{w}_1, \dots, \mathbf{w}_k\}$ produced as above are an orthonormal basis for S .

PROOF

To prove the claim requires showing two things: first, that $\text{span}\{\mathbf{w}_1, \dots, \mathbf{w}_k\} = \text{span}\{\mathbf{b}_1, \dots, \mathbf{b}_k\}$ and second, that $\{\mathbf{w}_1, \dots, \mathbf{w}_k\}$ is orthonormal.

1. **Proof that $\text{span}\{\mathbf{w}_1, \dots, \mathbf{w}_k\} = S$.** For this we use induction.

Obviously, $\text{span}\{\mathbf{w}_1\} = \text{span}\{\mathbf{b}_1\}$ since $\mathbf{w}_1 = \hat{\mathbf{b}}_1$. Suppose then that $\text{span}\{\mathbf{w}_1, \dots, \mathbf{w}_l\} = \text{span}\{\mathbf{b}_1, \dots, \mathbf{b}_l\}$ for $l < k$. Then, first of all, notice that $\text{span}\{\mathbf{w}_1, \dots, \mathbf{w}_l, \mathbf{w}_{l+1}\} = \text{span}\{\mathbf{w}_1, \dots, \mathbf{w}_l, \mathbf{b}_{l+1} - (\text{proj}_{\mathbf{w}_1} \mathbf{b}_{l+1} + \dots + \text{proj}_{\mathbf{w}_l} \mathbf{b}_{l+1})\}$. Thus, $\text{span}\{\mathbf{w}_1, \dots, \mathbf{w}_l, \mathbf{w}_{l+1}\} = \text{span}\{\mathbf{w}_1, \dots, \mathbf{w}_l, \mathbf{b}_{l+1} - (\text{proj}_{\mathbf{w}_1} \mathbf{b}_{l+1} + \dots + \text{proj}_{\mathbf{w}_l} \mathbf{b}_{l+1})\} = \{c_1 \mathbf{w}_1 + \dots + c_l \mathbf{w}_l + \dots + c_{l+1} \mathbf{b}_{l+1} - c_{l+1} (\text{proj}_{\mathbf{w}_1} \mathbf{b}_{l+1} + \dots + \text{proj}_{\mathbf{w}_l} \mathbf{b}_{l+1}) \mid c_1, \dots, c_{l+1} \in \mathbb{R}\}$ which is equal to $\{\tilde{c}_1 \mathbf{w}_1 + \dots + \tilde{c}_l \mathbf{w}_l + c_{l+1} \mathbf{b}_{l+1} \mid \tilde{c}_1, \dots, \tilde{c}_l, c_{l+1} \in \mathbb{R}\}$. But since, by the inductive assumption, $\text{span}\{\mathbf{w}_1, \dots, \mathbf{w}_l\} = \text{span}\{\mathbf{b}_1, \dots, \mathbf{b}_l\}$ we have $\{\tilde{c}_1 \mathbf{w}_1 + \dots + \tilde{c}_l \mathbf{w}_l + c_{l+1} \mathbf{b}_{l+1} \mid \tilde{c}_1, \dots, \tilde{c}_l, c_{l+1} \in \mathbb{R}\} = \{c_1 \mathbf{b}_1 + \dots + c_l \mathbf{b}_l + c_{l+1} \mathbf{b}_{l+1} \mid c_1, \dots, c_l, c_{l+1} \in \mathbb{R}\}$. In other words $\text{span}\{\mathbf{w}_1, \dots, \mathbf{w}_{l+1}\} = \text{span}\{\mathbf{b}_1, \dots, \mathbf{b}_{l+1}\}$. Therefore, by induction, we have $\text{span}\{\mathbf{w}_1, \dots, \mathbf{w}_k\} = S$.

2. **Proof of orthonormality.** The vectors \mathbf{w}_i are normalized by construction, so all that needs to be verified is that they are mutually orthogonal. For this, we can again apply induction. Notice that, as a base case, we have $\mathbf{w}_1 \cdot \mathbf{w}_2 = 0$ since $\mathbf{b}_1 \cdot (\mathbf{b}_2 - \text{proj}_{\mathbf{b}_1} \mathbf{b}_2) = 0$. To justify this we simply recall that it's generically true that $\mathbf{d} \cdot (\mathbf{v} - \text{proj}_{\mathbf{d}} \mathbf{v}) = 0$. Assume that $\{\mathbf{w}_1, \dots, \mathbf{w}_l\}$ is an orthonormal set for $l < k$. Then, for $j \leq l$ we have

$$\begin{aligned}
\mathbf{w}_j \cdot \mathbf{w}_{l+1} &= \mathbf{w}_j \cdot (\mathbf{b}_{l+1} - \text{proj}_{\mathbf{w}_1} \mathbf{b}_{l+1} - \dots - \text{proj}_{\mathbf{w}_l} \mathbf{b}_{l+1}) \\
&= \mathbf{w}_j \cdot (\mathbf{b}_{l+1} - \frac{\mathbf{w}_1 \cdot \mathbf{b}_{l+1}}{\|\mathbf{w}_1\|^2} \mathbf{w}_1 - \dots - \frac{\mathbf{w}_l \cdot \mathbf{b}_{l+1}}{\|\mathbf{w}_l\|^2} \mathbf{w}_l) \\
&= \mathbf{w}_j \cdot \mathbf{b}_{l+1} - \frac{\mathbf{w}_1 \cdot \mathbf{b}_{l+1}}{\|\mathbf{w}_1\|^2} \mathbf{w}_1 \cdot \mathbf{w}_j - \frac{\mathbf{w}_2 \cdot \mathbf{b}_{l+1}}{\|\mathbf{w}_2\|^2} \mathbf{w}_2 \cdot \mathbf{w}_j - \dots \\
&\quad - \frac{\mathbf{w}_l \cdot \mathbf{b}_{l+1}}{\|\mathbf{w}_l\|^2} \mathbf{w}_l \cdot \mathbf{w}_j \\
&\stackrel{\text{as } \mathbf{w}_i \cdot \mathbf{w}_j = 0}{=} \underbrace{\mathbf{w}_j \cdot \mathbf{b}_{l+1} - 0 - 0 - \dots - 0}_{\mathbf{w}_j \cdot \mathbf{b}_{l+1}} - \frac{\mathbf{w}_j \cdot \mathbf{b}_{l+1}}{\|\mathbf{w}_j\|^2} \mathbf{w}_j \cdot \mathbf{w}_j - 0 - \dots - 0 \\
&= \mathbf{w}_j \cdot \mathbf{b}_{l+1} - \mathbf{w}_j \cdot \mathbf{b}_{l+1} \\
&= 0
\end{aligned}$$

Where we've used that $\mathbf{w}_i \cdot \mathbf{w}_j = 0$ for $i \neq j, i \leq l$, from the inductive hypothesis. The above calculation shows that \mathbf{w}_{l+1} is orthogonal to $\{\mathbf{w}_1, \dots, \mathbf{w}_l\}$, so then $\{\mathbf{w}_1, \dots, \mathbf{w}_l, \mathbf{w}_{l+1}\}$ is an orthogonal set. Therefore, by induction, $\{\mathbf{w}_1, \dots, \mathbf{w}_k\}$ is an orthonormal set, as advertised.

□

Rank Unification

We next define something which has been used many times throughout the course but, until now, wasn't named.

DEFINITION 4

The **rank** of a matrix A , denoted $\text{rank}(A)$ is the number of pivots in A .

We've seen that, say, for an $m \times n$ matrix if $\text{rank}(A) = m$ then the columns of A span \mathbb{R}^m and if $\text{rank}(A) = n$ then the columns of A are linearly independent. If you don't see why please work it out on your own as an exercise. Rank is the glue that binds the row picture (number of equations) to the column picture (number of variables). It also is the link between the geometrical point of view (spanning, independence, etc) and the algebraic point of view (consistency, unique solvability, etc).

The following theorem tidily expresses the connection between rank and dimension. As well, it provides a description of how to produce bases for certain subspaces.

THEOREM 23: RANK THEOREM

Let A be an $m \times n$ matrix with rank r . Then

$$\dim(\text{col}(A)) = \dim(\text{row}(A)) = r$$

Moreover, if $A \sim R$ where R is in row-echelon form then

1. The r nonzero rows of R form a basis for $\text{row}(A)$
2. The r pivot columns of A form a basis for $\text{col}(A)$

PROOF

First of all if $A \sim R$ then we must have $\text{row}(A) = \text{row}(R)$ since the elementary row operations necessarily don't alter the span of the corresponding rows (see exercises). But the r nonzero rows of R are necessarily linearly independent since each row has a leading entry appearing to the right of the row preceding it (if any). Therefore the r nonzero rows of R are a basis for $\text{row}(R) = \text{row}(A)$. Notice this also shows that $\dim(\text{row}(A)) = r$.

Next, the pivot columns of R are obviously linearly independent (make sure you verify why!) and they span the subspace of \mathbb{R}^m given by

$$S = \left\{ \begin{bmatrix} a_1 \\ a_2 \\ \vdots \\ a_r \\ 0 \\ \vdots \\ 0 \end{bmatrix} \mid a_1, a_2, \dots, a_r \in \mathbb{R} \right\}$$

Notice the above is, essentially, \mathbb{R}^r with $m - r$ extra 0's. Since $A \sim R$ there's an invertible matrix U such that $UA = R$, in other words $R = [U\mathbf{a}_1 \dots U\mathbf{a}_n]$. Denote the pivot columns of A by indices j_1, \dots, j_r . Since pivot columns of R are a basis for $\text{col}(R)$ we have that $\{U\mathbf{a}_{j_k}\}_{k=1}^r$ is a basis for $\text{col}(R)$.

The theorem will be proved if we can show that $\{\mathbf{a}_{j_k}\}_{k=1}^r$ is a basis for $\text{col}(A)$.

1. **Proof that $\{\mathbf{a}_{j_k}\}_{k=1}^r$ is independent.** First of all, if $\mathbf{0} = c_1\mathbf{a}_{j_1} + c_2\mathbf{a}_{j_2} + \dots + c_r\mathbf{a}_{j_r}$ then multiplying by U gives $\mathbf{0} = c_1U\mathbf{a}_{j_1} + \dots + c_rU\mathbf{a}_{j_r}$. But $\{U\mathbf{a}_{j_k}\}_{k=1}^r$ is a basis for $\text{col}(R)$ so in particular is linearly independent and thus $c_1 = c_2 = \dots = c_r = 0$. In other words the set $\{\mathbf{a}_{j_k}\}_{k=1}^r$ is linearly independent.
2. **Proof that $\{\mathbf{a}_{j_k}\}_{k=1}^r$ spans $\text{col}(A)$.** Pick any column of A , say

\mathbf{a}_q . Since $\{U\mathbf{a}_{j_k}\}_{k=1}^r$ is a basis for $\text{col}(R)$ we have that there are constants c_1, \dots, c_r such that $U\mathbf{a}_q = c_1U\mathbf{a}_{j_1} + \dots + c_rU\mathbf{a}_{j_r}$. But U is invertible, so $\mathbf{a}_q = c_1\mathbf{a}_{j_1} + \dots + c_r\mathbf{a}_{j_r}$. In other words $\mathbf{a}_q \in \text{span}\{\mathbf{a}_{j_k}\}_{k=1}^r$, i.e. any column of A may be expressed uniquely as a linear combination of the r pivot columns of A .

This proves the theorem. \square

NB: Note the asymmetry in the two statements above. For a basis of $\text{row}(A)$ we can use the nonzero rows of the corresponding row-equivalent echelon matrix to A whereas for a basis of $\text{col}(A)$ we must use the pivot columns **of A itself!** Study the following example very carefully.

EXAMPLE 13

Consider $A = \begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix}$. Obviously, $A \sim \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$. The theorem above tells us that $\left\{ \begin{bmatrix} 0 \\ 1 \end{bmatrix} \right\}$ must be a basis for $\text{col}(A)$ and $\left\{ \begin{bmatrix} 1 \\ 0 \end{bmatrix} \right\}$ must be a basis for $\text{row}(A)$.

Notice we would be **wrong** to think that $\left\{ \begin{bmatrix} 1 \\ 0 \end{bmatrix} \right\}$ must be a basis for $\text{col}(A)$ and $\left\{ \begin{bmatrix} 0 \\ 0 \end{bmatrix} \right\}$ is a basis for $\text{row}(A)$ which is what we would get if we (incorrectly!) used the pivot column of the reduced matrix as the basis vector of the column space and the pivot row of the original matrix as the basis vector of the row space.

The next example provides an illustration of the utility of the rank theorem.

EXAMPLE 14

Find a basis for (and dimension of) the subspace $S = \text{span}\left\{ \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix}, \begin{bmatrix} 4 \\ 5 \\ 6 \end{bmatrix}, \begin{bmatrix} 7 \\ 8 \\ 9 \end{bmatrix} \right\}$.

Solution: We can view S as the row space of matrix $A = \begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{bmatrix}$.

Applying the reduction algorithm to matrix A yields $A \sim R = \begin{bmatrix} 1 & 0 & -1 \\ 0 & 1 & 2 \\ 0 & 0 & 0 \end{bmatrix}$. The rank theorem says that a basis for $\text{row}(A)$, and therefore of S , will be given by the nonzero rows of R . Therefore, a

basis for S is given by $\left\{ \begin{bmatrix} 1 \\ 0 \\ -1 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 2 \end{bmatrix} \right\}$. We can also conclude from this that $\dim(S) = 2$.

In the preceding example, we could just as well have chosen to view S as the column space of $\begin{bmatrix} 1 & 4 & 7 \\ 2 & 5 & 8 \\ 3 & 6 & 9 \end{bmatrix}$. If we did, we would then find that $A \sim$

$\begin{bmatrix} 1 & 0 & -1 \\ 0 & 1 & 2 \\ 0 & 0 & 0 \end{bmatrix}$ so the first two columns of A are pivot columns and therefore

$\left\{ \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix}, \begin{bmatrix} 4 \\ 5 \\ 6 \end{bmatrix} \right\}$ is a basis for S . Notice that this produces a *different basis* than we obtained previously, and that when using the column space, we must select pivot columns from the original matrix rather than its reduced form.

Rank and Nullity

There are $\text{rank}(A)$ pivot columns in a given $m \times n$ matrix A and the $n - r$ non-pivot columns correspond to free variables. Each such free variable corresponds to a basis element for $\ker(A)$. This follows from the fact that the reduction algorithm will produce linearly independent vectors for the general solution to the homogeneous problem $A\mathbf{x} = \mathbf{0}$ (see exercises). This means we have proven the following theorem.

THEOREM 24: RANK-NULLITY THEOREM

Let A be an $m \times n$ matrix. Then

$$\text{rank}(A) + \dim(\ker(A)) = n$$

The above is called the **rank-nullity theorem** since $\dim(\ker(A))$ is also called the **nullity** of A . If you think of the nullity (**very loosely speaking**) as characterizing how severe the “loss of information” is when multiplying vectors by A then the rank-nullity theorem tells us that $\text{rank}(A)$ and the nullity are tradeoffs where the $\text{rank}(A)$ controls the information-preservation of the linear map induced by A .²⁷

This class is about solving systems of linear equations. The rank-nullity theorem is a way to quantitatively characterize how far a given matrix might be from having $A\mathbf{x} = \mathbf{b}$ be uniquely solvable. It’s one of the most important results in a first course on linear algebra and understanding it is a major task for you this term.

²⁷ Again, these are metaphorical statements because I haven’t actually defined “information” etc. To clarify slightly, the dimension of a subspace, in measuring the number of independent vectors allowed in a space, encodes the possible “information” available to store in a subspace.

EXAMPLE 15

Find a basis for (and dimension of) the null space of $\begin{bmatrix} 1 & 2 & 3 \\ 5 & 1 & 0 \end{bmatrix}$.

Solution: We apply row reduction and find that $\begin{bmatrix} 1 & 2 & 3 \\ 5 & 1 & 0 \end{bmatrix} \sim$

$\begin{bmatrix} 1 & 0 & -\frac{1}{3} \\ 0 & 1 & \frac{5}{3} \end{bmatrix}$. Before going further, we notice that we can immedi-

ately deduce that $\text{rank}\left(\begin{bmatrix} 1 & 2 & 3 \\ 5 & 1 & 0 \end{bmatrix}\right) = 2$ so the rank-nullity theorem

allows us to know that $\dim(\ker\left(\begin{bmatrix} 1 & 2 & 3 \\ 5 & 1 & 0 \end{bmatrix}\right)) = 1$. To find a basis for this space we write all basic variables in terms of free variables, i.e.

$x_1 = \frac{1}{3}x_3, x_2 = -\frac{5}{3}x_3$ so $\ker\left(\begin{bmatrix} 1 & 2 & 3 \\ 5 & 1 & 0 \end{bmatrix}\right) = \text{span}\left\{\begin{bmatrix} 1 \\ -5 \\ 3 \end{bmatrix}\right\}$. Thus $\left\{\begin{bmatrix} 1 \\ -5 \\ 3 \end{bmatrix}\right\}$

is a basis for $\ker\left(\begin{bmatrix} 1 & 2 & 3 \\ 5 & 1 & 0 \end{bmatrix}\right)$.

Rank Inequalities

Throughout this section I also use the fact (which is left to you in the exercises) that if $S \subseteq W$ where S, W are subspaces of \mathbb{R}^n then $\dim S \leq \dim W$ holds, with equality if and only if $W = S$. It's not a simple exercise so you should prove it slowly and carefully to be sure.

To begin with a simple result which we shall get a lot of mileage out of, notice that

$$\begin{aligned} \text{rank}(A) &= \dim(\text{row}(A)) \\ &= \dim(\text{col}(A^T)) \\ &= \text{rank}(A^T) \end{aligned}$$

Therefore

$$\text{rank}(A) = \text{rank}(A^T)$$

A fact which we'll get to use shortly.

Since $\text{col}(A) \subseteq \mathbb{R}^m$ is a subspace we have that $\text{rank}(A) = \dim(\text{col}(A)) \leq \dim(\mathbb{R}^m) = m$. By similar reasoning we have that $\text{rank}(A) = \text{rank}(A^T) \leq n$. So then, we have that

$$\text{rank}(A) \leq \min(m, n)$$

must hold for all matrices. The case of *maximal rank* therefore occurs when equality is met above, namely

DEFINITION 5

An $m \times n$ matrix has **full rank** or **maximal rank** when $\text{rank}(A) = \min(m, n)$

Notice that we know, from the invertible matrix theorem in the textbook, that a **square matrix with full rank must be invertible**. Our first theorem shows us that multiplication by an invertible matrix “preserves information”.

THEOREM 25

Let A be an $m \times n$ matrix and M be an $m \times m$ invertible matrix. Then

$$\text{rank}(A) = \text{rank}(MA)$$

PROOF

If M is invertible then it can be written as a product of elementary matrices $M = E_1 \cdots E_k$. Then $MA = E_1 \cdots E_k A$ represents a sequence of elementary row operations performed on the matrix A , each of which preserves the row space. Therefore

$$\text{row}(A) = \text{row}(E_1 \cdots E_k A) = \text{row}(MA)$$

And if the spaces are the same the dimensions must be equal in which case the result immediately follows. \square

Using the above we can also verify that if P is an invertible $n \times n$ matrix we have

$$\begin{aligned} \text{rank}(AP) &= \text{rank}((AP)^T) \\ &= \text{rank}(P^T A^T) \\ &\underbrace{=} \text{rank}(A^T) \\ &\text{since } P^T \text{ is invertible} \\ &= \text{rank}(A) \end{aligned}$$

In other words we’ve shown that

$$\text{rank}(A) = \text{rank}(AP)$$

for all invertible matrices P . So again, multiplication by invertible matrix leaves rank “invariant”. The above two results are generalized in the following theorem.

THEOREM 26

Let A, B, C be matrices such that the products below are well-defined. Then

1. $\text{col}(AB) \subseteq \text{col}(A)$

$$2. \text{ row}(CA) \subseteq \text{row}(A)$$

In the above \subseteq will simply be $=$ when B or C respectively is invertible.

PROOF

We only prove the first statement since the proof of the other can be shown to follow by use of transposes. Let B have columns $\{\mathbf{b}_j\}$. Then the matrix AB has columns $\{A\mathbf{b}_j\}$, each term $A\mathbf{b}_j$ is a linear combination of the columns of A so therefore we have that the columns of AB are in the column space of A . In other words $\text{col}(AB) \subseteq \text{col}(A)$. To prove the case of equality we can appeal to the prior results we've already seen. Suppose B is invertible. Then, since $\text{col}(AB) \subseteq \text{col}(A)$ and $\text{rank}(AB) = \text{rank}(A)$ if B is invertible we see $\text{col}(AB)$, a subspace of $\text{col}(A)$, has the same dimension as $\text{col}(A)$. Therefore if B is invertible then $\text{col}(AB) = \text{col}(A)$. \square

A sometimes useful result is the following inequality.

THEOREM 27

If A and B are two matrices whose product is defined then

$$\text{rank}(AB) \leq \min(\text{rank}(A), \text{rank}(B))$$

PROOF

Since the prior theorem showed us that $\text{col}(AB) \subseteq \text{col}(A)$ we then have that $\text{rank}(AB) = \dim(\text{col}(AB)) \leq \dim(\text{col}(A)) = \text{rank}(A)$. Similarly we have $\text{row}(AB) \subseteq \text{row}(B)$ gives that $\text{rank}(AB) = \dim(\text{row}(AB)) \leq \dim(\text{row}(B)) = \text{rank}(B)$. The result follows. \square

Maximal Rank

In this section, we're going to look at the cases where the rank of an $m \times n$ matrix A is as large as possible. We've already seen that the case of maximal rank occurs when $\text{rank}(A) = \min(m, n)$. If A were square, then A having full rank ensures that $A\mathbf{x} = \mathbf{b}$ is always *uniquely solvable*. Since the nullity captures the "size" of the subspace controlling the homogeneous solution it controls the "size" of the obstruction to uniqueness if the equation is solvable but not *uniquely solvable*. The following theorems describe both possibilities for maximal rank.

THEOREM 28: RANK = # OF COLUMNS

Let A be $m \times n$. The following are equivalent statements.

1. $\text{rank}(A) = n$
2. rows of A span \mathbb{R}^n
3. columns of A are independent
4. $A^T A$ is invertible
5. There's an $n \times m$ matrix C such that $CA = I_n$. (The matrix C is called a "left inverse" of A)
6. $A\mathbf{x} = \mathbf{0}$ has only the trivial solution

PROOF

I only prove (3) \implies (4) \implies (5) leaving the cases (5) \implies (6) \implies (1) \implies (2) \implies (3) as an exercise. So, suppose that (3) holds and the columns of A are independent. Then that means that $A\mathbf{x} = \mathbf{0} \implies \mathbf{x} = \mathbf{0}$. So consider

$$\begin{aligned} A^T A\mathbf{x} = \mathbf{0} &\implies \mathbf{x}^T A^T A\mathbf{x} = \mathbf{0} \\ &\implies (A\mathbf{x})^T A\mathbf{x} = \mathbf{0} \\ &\implies [\|A\mathbf{x}\|^2] = 0 \\ &\implies \|A\mathbf{x}\|^2 = 0 \\ &\implies A\mathbf{x} = \mathbf{0} \\ &\implies \mathbf{x} = \mathbf{0} \end{aligned}$$

So the square matrix $A^T A$ admits only the trivial solution for its associated homogeneous problem. By the invertible matrix theorem we have that $A^T A$ must be invertible. Thus (3) \implies (4).

Next, if (4) holds then $A^T A$ is invertible. In that case, set $C = (A^T A)^{-1} A^T$. Obviously this choice of C gives the desired conclusion and so (4) \implies (5). □

The other possibility of full rank is dealt with in the next theorem.

THEOREM 29: RANK = # OF ROWS

Let A be $m \times n$. The following are equivalent statements.

1. $\text{rank}(A) = m$
2. columns of A span \mathbb{R}^m
3. rows of A are independent

4. AA^T is invertible
5. There's an $n \times m$ matrix C such that $AC = I_m$
6. $A\mathbf{x} = \mathbf{b}$ is solvable for all $\mathbf{b} \in \mathbb{R}^m$

PROOF

I only prove (3) \implies (4) \implies (5) \implies (6) leaving the cases (6) \implies (1) \implies (2) \implies (3) as an exercise. So, suppose that (3) holds and the rows of A are independent. Then that means that the m rows are a collection of m independent vectors inside of \mathbb{R}^n . But that means that $\dim(\text{row}(A)) = \text{rank}(A) = m$. From this we see that $\text{rank}(A) = m$, but also since $\text{rank}(A^T) = \text{rank}(A)$ we have that $\text{rank}(A^T) = m$ as well. In other words, the matrix A^T , has rank equal to its amount of *columns*. Therefore, our prior theorem, applied to the matrix A^T , says that $(A^T)^T A^T = AA^T$ is invertible. Thus (3) \implies (4). Next, assuming (4) we can set $C = A^T(AA^T)^{-1}$ and easily verify that $AC = I$. So (4) \implies (5).

Finally, assuming (5) means that there's a matrix C whose columns \mathbf{c}_j , $j = 1, \dots, m$ solve $A\mathbf{c}_j = \mathbf{e}_j$. Then, of course, for $\mathbf{b} \in \mathbb{R}^m$, $\mathbf{b} = b_1\mathbf{e}_1 + \dots + b_m\mathbf{e}_m$ we have

$$\begin{aligned} A(b_1\mathbf{c}_1 + \dots + b_m\mathbf{c}_m) &= b_1A\mathbf{c}_1 + \dots + b_mA\mathbf{c}_m \\ &= b_1\mathbf{e}_1 + \dots + b_m\mathbf{e}_m \\ &= \mathbf{b} \end{aligned}$$

which means that $A\mathbf{x} = \mathbf{b}$ is solvable by setting $\mathbf{x} = b_1\mathbf{c}_1 + \dots + b_m\mathbf{c}_m$ so (5) \implies (6). □

Exercises

You should be able to do the following.

1. Prove that spans of vectors are always subspaces.
2. True/False: If S_1 and S_2 are subspaces of \mathbb{R}^n then so is $S_1 \cup S_2$. If true, prove it, otherwise give a counterexample.
3. True/False: If S_1 and S_2 are subspaces of \mathbb{R}^n then so is $S_1 \cap S_2$. If true, prove it, otherwise give a counterexample.
4. Find $E_3\left(\begin{bmatrix} 1 & 1 \\ 4 & 1 \end{bmatrix}\right)$ and $E_{-1}\left(\begin{bmatrix} 1 & 1 \\ 4 & 1 \end{bmatrix}\right)$.
5. Find $E_\lambda(I_n)$ for all values of $\lambda \in \mathbb{R}$.

6. Prove that invertible linear transformations map bases of \mathbb{R}^n onto bases of \mathbb{R}^n . In other words if A is an invertible $n \times n$ matrix, prove that $\{A\mathbf{v}_j\}_{j=1}^n$ is a basis whenever $\{\mathbf{v}_j\}_{j=1}^n$ is.
7. Prove that if $A \sim B$ then $\text{row}(A) = \text{row}(B)$.
8. Prove that the reduction algorithm applied to $A\mathbf{x} = \mathbf{0}$ will produce a basis for $\text{null}(A)$ whenever $\text{null}(A) \neq \{\mathbf{0}\}$. *Hint: think about reduced row-echelon form.*
9. In the proof of Theorem 1.2, how do I know it's possible to choose a largest set of linearly independent vectors $\{\mathbf{v}_1, \dots, \mathbf{v}_k\}$?
10. It was claimed that the zero subspace admits no basis. Explain why.
11. Suppose that AB is invertible. Is A ? If so prove it. If not, give a counterexample.
12. Suppose A, B are $n \times n$ matrices. Use rank inequalities to show that if A is not invertible then AB must be non-invertible as well.
13. Prove that if U, V are subspaces of \mathbb{R}^n then $U \subseteq V$ implies $\dim U \leq \dim V$.
14. Use a theorem on maximal rank to prove that the polynomial $p(x) = 1 + x^2$ has no real roots.
15. Let $A = \begin{bmatrix} 2 & 2 & 3 \end{bmatrix}$. Find a basis for $\ker(A)$ and a basis for $\ker(A^T)$.
16. Suppose that A, B are square matrices of the same size and that $\text{col}(B) \subseteq \ker A$. What can I say about AB ?
17. Let $A = [a_{ij}]$ where $a_{ij} = i + j$ for $1 \leq i, j \leq n$. Determine $\text{rank}(A)$.
18. Suppose that A is $m \times n$. Show that $\ker A = \ker(UA)$ for all invertible $m \times m$ matrices U . Then also show that $\dim(\ker(A)) = \dim(\ker(AB))$ for all invertible $n \times n$ matrices B .

19. Let

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

- (a) Find a matrix B such that $AB = I_2$ or show that such a B doesn't exist.
- (b) Find a matrix C such that $CA = I_3$ or show that such a C doesn't exist.

20. Let

$$A = \begin{bmatrix} a & 1 & a & 0 & 0 & 0 \\ 0 & b & 1 & b & 0 & 0 \\ 0 & 0 & c & 1 & c & 0 \\ 0 & 0 & 0 & d & 1 & d \end{bmatrix}$$

where a, b, c, d are unspecified real numbers.

- (a) Prove that $\text{rank}(A) > 2$.
- (b) Prove that if $a = d = 0$ and $bc = 1$ then $\text{rank}(A) = 3$.
21. Let A be an $n \times n$ matrix.
- (a) Show that $A^2 = 0$ if and only if $\text{col}(A) \subseteq \text{ker}(A)$.
- (b) Show that $A^2 = 0$ implies $\text{rank}(A) \leq \frac{n}{2}$.
- (c) Find a matrix A such that $\text{col}(A) = \text{null}(A)$.
22. Let $\mathbf{c} \neq \mathbf{0}$ be in \mathbb{R}^m and $\mathbf{r} \in \mathbb{R}^n$ and define $A = \mathbf{c}\mathbf{r}^T$.
- (a) Show that $\text{col}(A) = \text{span}\{\mathbf{c}\}$ and $\text{row}(A) = \text{span}\{\mathbf{r}\}$.
- (b) Find $\dim(\text{ker}(A))$.
- (c) Show that $\text{ker } A = \text{null}(\mathbf{r})$.
23. Prove that if $\dim S = m$ then any set of m linearly independent vectors in S must be a basis.
24. Prove that if $\dim S = m$ then any set of m spanning vectors in S must be a basis.
25. Prove that the only *proper* subspaces of \mathbb{R}^3 are lines through the origin and planes through the origin.
26. Prove the remaining equivalent statements in the theorems on full rank (Theorems 9 & 10).
27. Let B be $m \times n$ and AB be $k \times n$. Suppose that $\text{rank}(AB) = \text{rank}(B)$. Show that $\text{ker } B = \text{ker}(AB)$.
28. Show that, for every $n \times m$ matrix A , the matrix $I_m + A^T A$ is invertible.
29. Suppose that $S \subset \mathbb{R}^8$ where $\dim S = 5$. Is there a subspace $V \subseteq \mathbb{R}^8$ such that $\dim V = 2$ and $V \cap S = \{\mathbf{0}\}$. Justify your answer.

The Fundamental Theorem of Linear Algebra

WE ARE GOING to present a striking result, fundamental to understanding how the concepts we've seen all semester relate to one another. But, equally important, we're going to see a remarkable picture which illustrates the theorem.

Contents

<i>Prelude: Orthogonal Complements</i>	72
<i>The Fundamental Theorem of Linear Algebra</i>	73
<i>The Diagrams</i>	75
<i>Decomposition of \mathbb{R}^n</i>	75
<i>Decomposition of \mathbb{R}^m</i>	76
<i>An application</i>	76
<i>Exercises</i>	79

Prelude: Orthogonal Complements

Before we get to the main theorem, a bit on the terminology it uses. We'll begin with a natural definition.

DEFINITION 6

Let $S \subseteq \mathbb{R}^n$ be a subspace. The set $S^\perp \subseteq \mathbb{R}^n$ is defined to be the collection of all vectors $\mathbf{v} \in \mathbb{R}^n$ orthogonal to S and called the **orthogonal complement of S** in \mathbb{R}^n . In other words,

$$S^\perp = \{\mathbf{u} \in \mathbb{R}^n \mid \mathbf{u} \cdot \mathbf{v} = 0, \mathbf{v} \in S\}$$

You should be able to prove that S^\perp is a subspace of \mathbb{R}^n (see exercises)

The purpose of these notes is to establish a version of the **Fundamental Theorem of Linear Algebra**. The result can be thought of as a type of *representation* theorem, namely, it tells us something about *how vectors are* by describing the canonical subspaces of a matrix A in which they live. To understand this we consider the following representation theorem.

THEOREM 30

Let $\mathbf{v} \in \mathbb{R}^n$ and let $S \subseteq \mathbb{R}^n$ be a subspace. Then there are vectors $\mathbf{s} \in S$, $\mathbf{s}_\perp \in S^\perp$ such that

$$\mathbf{v} = \mathbf{s} + \mathbf{s}_\perp$$

In other words, vectors can be expressed in terms of pieces living in orthogonal spaces.

The above theorem is sometimes expressed in the notation $\mathbb{R}^n = S \oplus S^\perp$ where the notation is called the “**direct sum**” of S and S^\perp .

PROOF

The proof needs a small fact we saw early in the section on Gram-Schmidt. Here's the fact: not only does every subspace S of \mathbb{R}^n have a basis, it actually has an *orthonormal* basis, namely a basis \mathcal{B} consisting of vectors \mathbf{s}_j where $\|\mathbf{s}_j\| = 1$ holds for all j and $\mathbf{s}_i \cdot \mathbf{s}_j = 0$ when $i \neq j$. How do we know? We've already proven that subspaces all have bases. From there all that's left is to use the Gram-Schmidt algorithm to turn a given basis into one in which all the vectors are made orthogonal and of unit length.

So S has an orthonormal basis as described. Then set

$$\mathbf{s} = (\mathbf{v} \cdot \mathbf{s}_1)\mathbf{s}_1 + \cdots + (\mathbf{v} \cdot \mathbf{s}_{\dim S})\mathbf{s}_{\dim S}$$

which must clearly be an element of S . Then define $\mathbf{s}_\perp = \mathbf{v} - \mathbf{s}$. It's obvious that $\mathbf{v} = \mathbf{s} + \mathbf{s}_\perp$ and you can verify that $\mathbf{s}_\perp \in S^\perp$ giving the result. \square

EXAMPLE 16

Consider the subspace $S = \text{span}\{\mathbf{e}_1\} \subset \mathbb{R}^4$. Then $S^\perp = \{\mathbf{v} \in \mathbb{R}^4 \mid \mathbf{v} \cdot \mathbf{e}_1 = 0\}$. A basis for S^\perp is provided by $\{\mathbf{e}_2, \mathbf{e}_3, \mathbf{e}_4\}$. Clearly any vector $\mathbf{w} \in \mathbb{R}^4$ can be written as

$$\mathbf{w} = c_1\mathbf{e}_1 + c_2\mathbf{e}_2 + c_3\mathbf{e}_3 + c_4\mathbf{e}_4 = \mathbf{s} + \mathbf{s}_\perp$$

as claimed by the theorem.

Where this is going: We'll soon see that for an $m \times n$ matrix A we have $(\ker A)^\perp = \text{col}(A^T)$. If this is true, then that means (using the notation above) that $\mathbb{R}^n = \ker(A) \oplus \text{col}(A^T)$. In this case a vector $\mathbf{x} \in \mathbb{R}^n$ can be written as $\mathbf{x} = \mathbf{p} + \mathbf{v}$ where $A\mathbf{v} = \mathbf{0}$ and $\mathbf{p} \in \text{row}(A)$ since $\text{row}(A) = \text{col}(A^T)$. In other words if $\mathbf{b} \in \text{col}(A)$ we can solve $A\mathbf{x} = \mathbf{b}$ and since $\mathbf{x} \in \mathbb{R}^n = \ker(A) \oplus \text{col}(A^T)$ we can write \mathbf{x} as

$$\mathbf{x} = \mathbf{p} + \mathbf{v}_h$$

In the above $\mathbf{p} \in \text{row}(A)$ since $\text{row}(A) = \text{col}(A^T)$. The space of possible \mathbf{v}_h 's is $\dim(\ker(A))$ -dimensional, so the nullity describes the lack of unique solvability for the linear system $A\mathbf{x} = \mathbf{b}$.

The Fundamental Theorem of Linear Algebra

We can now get on with proving the main theorem in this course, the capstone to our understanding what it means to solve systems of linear equations.

PROPOSITION 1

Suppose A is an $m \times n$ matrix. Then

$$\text{col}(A^T)^\perp = \ker A$$

PROOF

Suppose $\mathbf{v} \in \ker A$, then $A\mathbf{v} = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix} = \begin{bmatrix} (\text{row}_1(A)) \cdot \mathbf{v} \\ (\text{row}_2(A)) \cdot \mathbf{v} \\ \vdots \\ (\text{row}_m(A)) \cdot \mathbf{v} \end{bmatrix}$ so \mathbf{v} is or-

thogonal to $\text{row}_i(A)$ for $i = 1, \dots, m$. Therefore $\mathbf{v} \in (\text{row}(A))^\perp = (\text{col}(A^T))^\perp$, i.e. $\ker A \subseteq (\text{col}(A^T))^\perp$. As well, if $\mathbf{v} \in (\text{col}(A^T))^\perp$ then, in particular, we must have that $\mathbf{v} \cdot (\text{row}_i(A)) = 0$ for $i = 1, \dots, m$. But in this case we'd again have $\mathbf{v} \in \ker A$. Thus $(\text{col}(A^T))^\perp \subseteq \ker A$. Since the sets contain each other they must be equal and we're done.

□

If $\mathbf{v} \in (S^\perp)^\perp$ then clearly $\mathbf{v} \cdot \mathbf{s}_\perp = 0$ for all $\mathbf{s}_\perp \in S^\perp$. If $\{\mathbf{s}_1, \dots, \mathbf{s}_{\dim S}\}$ is an orthonormal basis for S and $\{\mathbf{s}_1^\perp, \dots, \mathbf{s}_{\dim S^\perp}^\perp\}$ is an orthonormal basis for S^\perp we can write out (from the expansion formula we saw earlier)

$$\begin{aligned}\mathbf{v} &= (\mathbf{v} \cdot \mathbf{s}_1)\mathbf{s}_1 + \dots + (\mathbf{v} \cdot \mathbf{s}_{\dim S})\mathbf{s}_{\dim S} + (\mathbf{v} \cdot \mathbf{s}_1^\perp)\mathbf{s}_1^\perp + \dots + (\mathbf{v} \cdot \mathbf{s}_{\dim S^\perp}^\perp)\mathbf{s}_{\dim S^\perp}^\perp \\ &= (\mathbf{v} \cdot \mathbf{s}_1)\mathbf{s}_1 + \dots + (\mathbf{v} \cdot \mathbf{s}_{\dim S})\mathbf{s}_{\dim S} \in S\end{aligned}$$

So $(S^\perp)^\perp \subseteq S$. As well, clearly for $\mathbf{s} \in S$ we have $\mathbf{s} \in (S^\perp)^\perp$ since $\mathbf{s} \cdot \mathbf{s}_\perp = 0$ must be true for all $\mathbf{s}_\perp \in S^\perp$. Thus, $S \subseteq (S^\perp)^\perp$. From these facts it follows that

$$S = (S^\perp)^\perp$$

must hold for all subspaces $S \subseteq \mathbb{R}^n$. (Make sure you can verify this statement). Applying this general fact to the previous theorem we immediately get the following.

THEOREM 31: FUNDAMENTAL THEOREM OF LINEAR ALGEBRA

Suppose A is an $m \times n$ matrix. Then

$$\text{col}(A^T) = (\ker A)^\perp$$

So that

$$\mathbb{R}^n = \ker(A) \oplus \text{col}(A^T)$$

gives an orthogonal decomposition of \mathbb{R}^n into the null space and the row space of matrix A . Therefore, for $\mathbf{b} \in \text{col}(A)$ we have that $A\mathbf{x} = \mathbf{b}$ is solved by

$$\mathbf{x} = \mathbf{p} + \mathbf{v}_h$$

for $\mathbf{p} \in \text{row}(A)$ a particular solution, $A\mathbf{p} = \mathbf{b}$, and $\mathbf{v}_h \in \ker A$ a generic vector in $\ker A$.

EXAMPLE 17

Consider the matrix $A = \begin{bmatrix} 1 & 0 & 7 & 6 \\ 0 & 1 & 5 & 2 \\ 0 & 0 & 0 & 0 \end{bmatrix}$. The fundamental theorem,

applied to A^T , tells us that $\mathbb{R}^m = \ker(A^T) \oplus \text{col}(A)$, which in this case gives $\mathbb{R}^3 = \ker(A^T) \oplus \text{col}(A)$. Let's see if this holds.

Since the matrix is in reduced form, we know that a basis for $\text{col}(A)$

is given by the pivot columns, namely $\left\{ \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix} \right\}$. Next, look at

$\ker(A^T)$. Since $A^T = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 7 & 5 & 0 \\ 6 & 2 & 0 \end{bmatrix}$ we use row reduction to find a basis

for the null space. Namely, $A^T \sim \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}$ and therefore, a basis for

$\ker(A^T)$ is given by $\left\{ \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \right\}$. Notice that, indeed, $\ker(A^T) = (\text{col}(A))^\perp$ and that $\mathbb{R}^3 = \ker(A^T) \oplus \text{col}(A)$ as the theorem dictates.

The Diagrams

The content of this theorem, the fundamental theorem of linear algebra, is encapsulated in the following figure. If a picture is worth a thousand words, these figures are worth at least several hours' thought. The figures should be thought of as a visualization of the effect of the linear transformation T_A , induced by $m \times n$ matrix A with rank r .

Decomposition of \mathbb{R}^n

Here we consider the orthogonal decomposition described by the fundamental theorem of the domain of T_A , namely $\mathbb{R}^n = \ker A \oplus \text{col}(A^T)$.

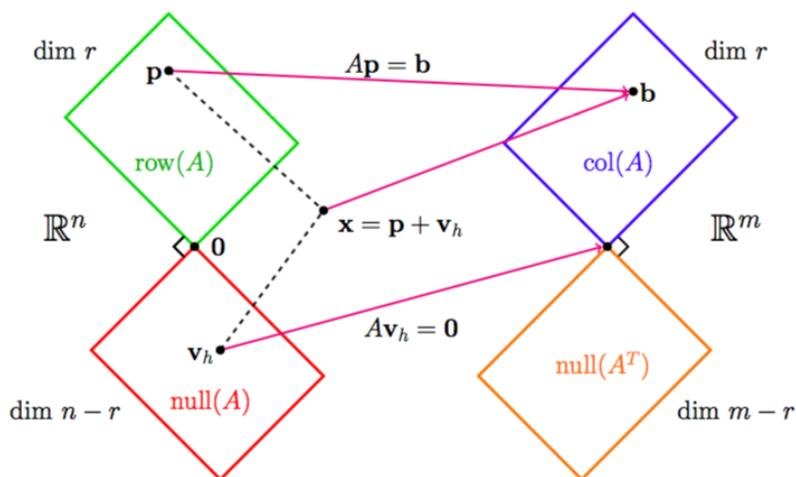


Figure 5: Solving $A\mathbf{x} = \mathbf{b}$ for an $m \times n$ matrix with $\text{rank}(A) = r$. The subspaces on both sides meet at right angles. Notice that the subspaces appearing on the left correspond to their "dual subspaces" for corresponding transpose of the matrix A appearing on the right hand side.

Notice this picture has all the major players we've seen in this course: there are vectors, subspaces, dimensions, and they are all playing a role in the

way we conceptualize solving a linear system of equations. Moreover, note the duality in the picture: the subspaces on the right-hand side, those living in \mathbb{R}^m , are the corresponding subspaces appearing on the left-hand side for the *transpose* of A rather than A (e.g. $\text{col}(A)$ appearing on the upper right is just $\text{row}(A^T)$ whereas we see $\text{row}(A)$ on the upper left-hand side). The subspaces meet at right angles because they're orthogonal complements of each other and the intersection of the spaces is the zero subspace of the respective ambient space (\mathbb{R}^n or \mathbb{R}^m). Try to visualize how the picture changes as we look at limiting values of possible rank for A . Namely, how would this above change in the case where $\text{rank}(A) = n$ say? Or m ? Or if A is square and invertible? If A is symmetric? Etc.

Decomposition of \mathbb{R}^m

Next, we consider a nice application of the orthogonal decomposition of the codomain of the transformation T_A described by the fundamental theorem of linear algebra, namely $\mathbb{R}^m = \ker A^T \oplus \text{col}(A)$.

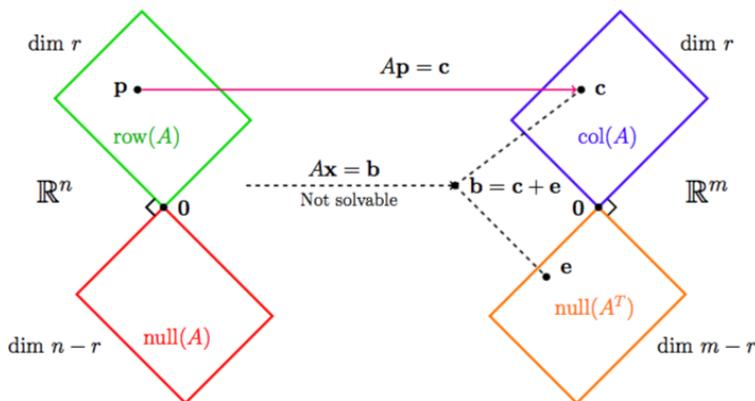


Figure 6: Finding the best \mathbf{x} when $A\mathbf{x} = \mathbf{b}$ isn't solvable since $\mathbf{b} \notin \text{col}(A)$ for an $m \times n$ matrix with $\text{rank}(A) = r$. The subspaces on both sides meet at right angles. Since $\mathbf{b} \notin \text{col}(A)$ the equation $A\mathbf{x} = \mathbf{b}$ isn't solvable. But the orthogonal decomposition of \mathbb{R}^m allows a natural approach to finding the "best" \mathbf{x} .

An application

We're going to expand on the last picture a bit. Consider the following situation. We've collected m output numbers y_1, \dots, y_m from a measuring device at various times labeled t_1, t_2, \dots, t_m . In other words, we have collected data of the form (t_i, y_i) for $i = 1, \dots, m$. The values y_i could be the reported valuation on a stock we are interested in purchasing, say, where the "device" is a computer ticker displaying the values. The values y_i could be the median selling price of a home in a given neighbourhood at a given time t_i , say. The "device" in this instance, might again be an annual financial report or internal

publication. The point is, the situation I've described is quite general. Maybe the data we've collected looks something like the following

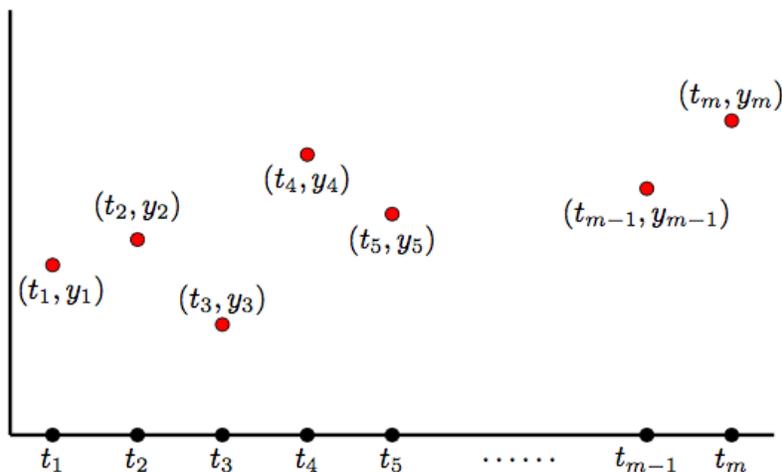


Figure 7: A plot of data points (t_i, y_i) for various t values.

It would be of obvious financial value to be able to use the data we've gathered to be able to *predict* new data at times in the future. This would be doable if, say, we had a *model* for how the data were produced, namely if we had a function $f : \mathbb{R} \rightarrow \mathbb{R}$ for which

$$y_i = f(t_i), \quad i = 1, \dots, m$$

Simplicity being preferable to complexity in modelling, the easiest thing to consider as a first approximation would be a *linear* model. Namely, we assume f takes the form $f(w) = mw + b$ for fixed numbers m, b . In that case we're searching for a model to be able to use for predictive purposes which takes the form

$$y_i = at_i + b \quad i = 1, \dots, m \quad (0.0.5)$$

The problem is that you can see from the figure above that the data points clearly cannot live on a line. On the other hand, they aren't *that* far off so maybe it's not hopeless either. Since, as modellers, we're allowed to *choose* the parameters a, b as we like, we can view them as unknowns in the problem and try to use the data set to determine the best possible choice.

We rewrite equation (0.0.5) as the matrix equation

$$\begin{array}{c} \begin{bmatrix} 1 & t_1 \\ 1 & t_2 \\ 1 & t_3 \\ \vdots & \vdots \\ 1 & t_m \end{bmatrix} \mathbf{x} = \begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ \vdots \\ y_m \end{bmatrix} \\ \text{A} \qquad \qquad \mathbf{b} \end{array} \quad (0.0.6)$$

Since we've already noticed that the points (t_i, y_i) *don't* live on a line we know that, in the above, we have $\mathbf{b} \notin \text{col}(A)$, so there's no chance of actually solving the above matrix equation to actually determine the parameters a, b . But a glance at Figure 2, together with the discussion preceding it, indicates that we shouldn't give up. Namely, instead of solving $A\mathbf{x} = \mathbf{b}$ which we know we cannot, we find the vector $\mathbf{y} \in \text{col}(A)$ which is as close as possible to \mathbf{b} . As before, this vector must satisfy $A^T(\mathbf{y} - \mathbf{b}) = \mathbf{0}$ which means we want to solve

$$A^T A \mathbf{x} = A^T \mathbf{b}$$

Then again, in this particular case, since $t_1 < t_2 < \dots < t_m$ we have that $\text{rank}(A) = 2$. Our theorem on maximal rank guarantees then that $A^T A$ must be invertible. Therefore, we have found an optimal solution

$$\mathbf{x} = (A^T A)^{-1} A^T \mathbf{b}$$

In other words, the parameters m, b which provide us with a "best linear fit" for the data set are

$$\begin{bmatrix} b \\ a \end{bmatrix} = \begin{pmatrix} \begin{bmatrix} 1 & 1 & \dots & 1 \\ t_1 & t_2 & \dots & t_m \end{bmatrix} \\ \begin{bmatrix} 1 & t_1 \\ 1 & t_2 \\ 1 & t_3 \\ \vdots & \vdots \\ 1 & t_m \end{bmatrix} \end{pmatrix}^{-1} \begin{bmatrix} 1 & 1 & \dots & 1 \\ t_1 & t_2 & \dots & t_m \end{bmatrix} \begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ \vdots \\ y_m \end{bmatrix}$$

We can clean this up a little by defining $\mathbb{1} = [1 \ 1 \ \dots \ 1]^T$ and $\mathbf{t} = [t_1 \ t_2 \ \dots \ t_m]^T$. This allows us to rewrite the above as

$$\begin{aligned} \begin{bmatrix} b \\ a \end{bmatrix} &= \begin{pmatrix} \begin{bmatrix} 1 & 1 & \dots & 1 \\ t_1 & t_2 & \dots & t_m \end{bmatrix} \\ \begin{bmatrix} 1 & t_1 \\ 1 & t_2 \\ 1 & t_3 \\ \vdots & \vdots \\ 1 & t_m \end{bmatrix} \end{pmatrix}^{-1} \begin{bmatrix} 1 & 1 & \dots & 1 \\ t_1 & t_2 & \dots & t_m \end{bmatrix} \begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ \vdots \\ y_m \end{bmatrix} \\ &= \begin{bmatrix} m & \mathbb{1} \cdot \mathbf{t} \\ \mathbb{1} \cdot \mathbf{t} & \|\mathbf{t}\|^2 \end{bmatrix}^{-1} \begin{bmatrix} 1 & 1 & \dots & 1 \\ t_1 & t_2 & \dots & t_m \end{bmatrix} \begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ \vdots \\ y_m \end{bmatrix} \\ &= \frac{1}{m\|\mathbf{t}\|^2 - (\mathbb{1} \cdot \mathbf{t})} \begin{bmatrix} \|\mathbf{t}\|^2 & -\mathbb{1} \cdot \mathbf{t} \\ -\mathbb{1} \cdot \mathbf{t} & m \end{bmatrix} \begin{bmatrix} \mathbb{1} \cdot \mathbf{b} \\ \mathbf{t} \cdot \mathbf{b} \end{bmatrix} \end{aligned}$$

So that $b = \left(\frac{\|\mathbf{t}\|^2 \mathbb{1} - (\mathbb{1} \cdot \mathbf{t}) \mathbf{t}}{m\|\mathbf{t}\|^2 - (\mathbb{1} \cdot \mathbf{t})} \right) \cdot \mathbf{b}$ and $a = \left(\frac{m\mathbf{t} - (\mathbb{1} \cdot \mathbf{t}) \mathbb{1}}{m\|\mathbf{t}\|^2 - (\mathbb{1} \cdot \mathbf{t})} \right) \cdot \mathbf{b}$ give the best choice of parameters for fitting the data to a line.

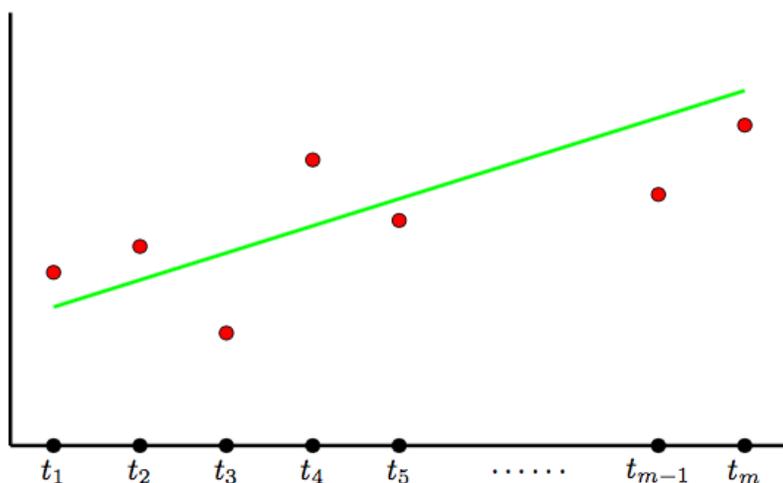


Figure 8: A plot of data points (t_i, y_i) together with the line $y = ax + b$ where a, b are the best possible parameters. Notice that while the points don't fall on the line, they are not far off.

Exercises

You should be able to do the following.

1. Calculate the orthogonal complements of the improper subspaces of \mathbb{R}^n .
2. Let $\mathbf{v} \in \mathbb{R}^3$ be a nonzero vector. Describe $(\text{span}\{\mathbf{v}\})^\perp$.
3. Prove that S^\perp is a subspace of \mathbb{R}^n whenever S is.
4. Prove that for $S \subseteq \mathbb{R}^n$ a subspace we have $\dim S + \dim S^\perp = n$.
5. If S, W are subspaces of \mathbb{R}^n show that $S \subseteq W \implies W^\perp \subseteq S^\perp$
6. Let $S = \{\mathbf{x} \in \mathbb{R}^4 \mid x_1 + x_2 + x_3 + x_4 = 0\}$. Prove that S is a subspace and find a basis for S^\perp .
7. Prove that every $m \times n$ matrix A defines a linear transformation from $\text{row}(A)$ onto $\text{col}(A)$, i.e.

$$T_A : \text{row}(A) \rightarrow \text{col}(A)$$

8. Consider A , an $m \times n$ matrix and $\mathbf{b} \in \mathbb{R}^m$. Consider the following claim²⁸:
One and only one of the following systems is solvable
 - (a) $A\mathbf{x} = \mathbf{b}$
 - (b) $A^T\mathbf{y} = \mathbf{0}$ with $\mathbf{y}^T\mathbf{b} \neq 0$

Prove that the above options *cannot BOTH hold*. In other words, if one of the above holds the other one mustn't.

²⁸ A variation of this strange-sounding claim is very important in numerous applications including, very importantly, in differential equations.

