

LINEAR ALGEBRA

An Introduction to Data Science

Paul A. Jensen
University of Illinois at Urbana-Champaign

Contents

	<i>Introduction</i>	5
	<i>0.1 Notation</i>	5
<i>1</i>	<i>Fields and Vectors</i>	9
	<i>1.1 Algebra</i>	9
	<i>1.2 The Field Axioms</i>	9
	<i>1.2.1 Common Fields in Mathematics</i>	11
	<i>1.3 Vector Addition</i>	12
	<i>1.4 Vector Multiplication is not Elementwise</i>	13
	<i>1.4.1 Do We Need Multiplication?</i>	13
	<i>1.5 Linear Systems</i>	13
	<i>1.6 Vector Norms</i>	14
	<i>1.6.1 Normalized (Unit) Vectors</i>	15
	<i>1.7 Scalar Vector Multiplication</i>	16
	<i>1.8 Inner (Dot) Products</i>	16
	<i>1.8.1 Computing the Dot Product</i>	17
	<i>1.8.2 Dot Product Summary</i>	18
<i>2</i>	<i>Matrices</i>	19
	<i>2.1 Matrix Multiplication</i>	19
	<i>2.1.1 Generalized Multiplication</i>	20

2.2	<i>Identity Matrix</i>	21
2.3	<i>Matrix Transpose</i>	21
2.4	<i>Solving Linear Systems</i>	22
2.5	<i>Gaussian Elimination</i>	25
2.6	<i>Computational Complexity of Gaussian Elimination</i>	26
2.7	<i>Solving Linear Systems in MATLAB</i>	27
3	<i>The Finite Difference Method</i>	29
3.1	<i>Finite Differences</i>	29
3.2	<i>Linear Differential Equations</i>	30
3.3	<i>Discretizing a Linear Differential Equation</i>	30
3.4	<i>Boundary Conditions</i>	32
4	<i>Inverses, Solvability, and Rank</i>	35
4.1	<i>Matrix Inverses</i>	35
4.2	<i>Elementary Matrices</i>	36
4.3	<i>Proof of Existence for the Matrix Inverse</i>	36
4.4	<i>Computing the Matrix Inverse</i>	38
4.5	<i>Numerical Issues</i>	39
4.6	<i>Inverses of Elementary Matrices</i>	41
4.7	<i>Rank</i>	41
4.8	<i>Rank and Matrix Inverses</i>	46
4.9	<i>Summary</i>	47

Introduction

This class has three parts. In Part I, we analyze linear systems that transform an n -dimensional vector (\mathbf{x}) into another n -dimensional vector (\mathbf{y}). This transformation is often expressed as a linear system via matrix multiplication: $\mathbf{y} = \mathbf{A}\mathbf{x}$. In Part II, we expand the type of systems we can solve, including systems that transform vectors from n dimensions to m dimensions. We will also consider solution strategies that use alternative objectives when the system contains either too much or not enough information. Finally, Part III dispenses with linear systems altogether, focusing purely on observations of sets of n -dimensional vectors (matrices). We will learn how to analyze and extract information from matrices without a clear input/output relationship.

0.1 Notation

We will distinguish scalars, vectors, and matrices with the following typographic conventions:

Object	Font and Symbol	Examples
Scalars	italicized, lowercase letters	x, α, y
Vectors	bold lowercase letters	$\mathbf{x}, \mathbf{y}, \mathbf{n}, \mathbf{w}$
Matrices	bold uppercase	$\mathbf{A}, \mathbf{A}^{-1}, \mathbf{B}, \mathbf{\Gamma}$

There are many ways to represent rows or columns of a matrix \mathbf{A} . Since this course uses MATLAB, we think it is convenient to use a matrix addressing scheme that reflects Matlab's syntax. So, the i th row of matrix \mathbf{A} will be $\mathbf{A}(i, :)$, and the j th column will be $\mathbf{A}(:, j)$. Rows or columns of a matrix are themselves vectors, so we choose to keep the boldface font for the matrix \mathbf{A} even when it is subscripted. We could also use Matlab syntax for vectors ($\mathbf{x}(i)$, for example). However, the form x_i is standard across many fields of mathematics and engineering, so we retain the common notation. The lack of boldface font reminds us that elements of vectors are scalars in the field.

One goal of this class is to increase the precision of your mathematical writing. We will regularly use the following symbols to describe mathematical concepts and relations.

Symbol	Read As	Description
\Rightarrow	implies	$p \Rightarrow q$ means that whenever p is true, q must also be true.
\Leftrightarrow	if and only if	A symmetric, stronger version of \Rightarrow . The expression $p \Leftrightarrow q$ means $p \Rightarrow q$ and $q \Rightarrow p$.
\forall	for all	Remember this symbol as an upside down “A”, as in “for All ”.
\exists	there exists	Remember this symbol as a backwards “E”, as in “there E xists”. To say something does not exist, use \nexists .
\in (\notin)	is (not) a member of	Used to state that a single element is a member of a set, i.e. $1 \in \mathbb{Z}$. To say that a set is a subset of another set, use \subset .
s.t.	such that	Other texts use the symbol $ $ (a vertical pipe) instead of “s.t.”. Note that “s.t.” is set in normal, not italicized font.
\mathbb{R}	the real numbers	The numbers along a line. The reals include both rational and irrational numbers.
\mathbb{R}^n	the set of n -dimensional vectors of real numbers	Each value of n is a different set. If $r \in \mathbb{R}^2$ then $r \notin \mathbb{R}^3$. Also, \mathbb{R}^2 is not a subset of \mathbb{R}^3 , etc.
\mathbb{Z}	the integers	The integers contain the “natural numbers” (1, 2, 3, ...), their negatives (-1, -2, -3, ...), and the number zero (0). The symbol comes from the German word for “number” (Zahlen). The word “integer” (Latin for “whole”) is used since integers have no fractional part.
\mathbb{Q}	the rationals	The rational numbers are all numbers that are the quotient of two integers. The symbol derives from the word “quotient”.
\mapsto	maps to	Describes the inputs and outputs of an operation. An operation that maps a vector of reals to a real number is $\mathbb{R}^n \mapsto \mathbb{R}$. An operation that maps two integers to a rational is $\mathbb{Z} \times \mathbb{Z} \mapsto \mathbb{Q}$.
\equiv	is defined as	Two expressions are equivalent because we have defined them as such, not because the relationship can be shown logically. For example, $a/b \equiv a \times b^{-1}$ defines the division operator using the multiplicative inverse.

These symbols can be used to succinctly write mathematical statements. For example, we can formally define the set of rational numbers as

A number is rational if and only if it can be expressed as the quotient of two integers.

with the statement

$$r \in \mathbb{Q} \Leftrightarrow \exists p, q \in \mathbb{Z} \text{ s.t. } r = p/q$$

While the latter statement is shorter, it is more difficult to understand. So whenever possible we recommend writing statements with as few symbols as necessary. Rely on mathematical symbols only when a textual definition would be unwieldy or imprecise, or when brevity is important (like when writing on a chalkboard).

1

Fields and Vectors

1.1 *Algebra*

Algebra is a branch of mathematics that contains symbols and a set of rules to manipulate them.

You are probably familiar with the idea of symbols. We call them variables, and in previous algebra courses you used them to represent unknown (real) numbers. In this course, we will use variables to represent vectors. Vectors are collections of elements (e.g. real numbers). The number of elements in a vector is its dimension. We can write an n -dimensional vector as

$$\mathbf{x} = \begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{pmatrix}$$

While some vectors have elements that are real numbers, vectors themselves are not numbers. An n -dimensional vector does not belong to the set \mathbb{R} of real numbers; it belongs to a special set \mathbb{R}^n of all other vectors of dimension n with real elements.

We use the rules of algebra to manipulate elements. However, only certain sets of elements are amenable to the rules of algebra. These algebra-compatible sets are called *fields*. A set of conditions, or *axioms*, must be true about a set before we can consider it a field. These field axioms define the rules of algebra.

After spending years studying algebra, you might think that there are many byzantine rules that govern fields. In fact, there are only five. The five axioms describe only two operations (addition and multiplication) and define two special elements that must be in every field (0 and 1).

1.2 *The Field Axioms*

Given elements a , b , and c in a field:

When we say vector, we assume a *column vector* – a vertical array of elements. A *row vector* is a horizontal array of elements. We will see that column vectors are more convenient.

We can surround the elements of a vector with either parentheses () or square brackets []. Straight lines || or curly braces { } are not allowed, as these have special meanings.

1. **Associativity.**

$$a + b + c = (a + b) + c = a + (b + c)$$

$$abc = (ab)c = a(bc)$$

2. **Commutativity.**

$$a + b = b + a$$

$$ab = ba$$

3. **Distribution** of multiplication over addition.

$$a(b + c) = ab + ac$$

4. **Identity.** There exist elements 0 and 1, both in the field, such that

$$a + 0 = a$$

$$1 \times a = a$$

5. **Inverses.**

- For all a , there exists an element $(-a)$ in the field such that $a + (-a) = 0$.
- For all $a \neq 0$, there exists an element (a^{-1}) in the field such that $a \times a^{-1} = 1$.

It might surprise you that only five axioms are sufficient to recreate everything you know about algebra. For example, nowhere do we state the special property of zero that $a \times 0 = 0$ for any number a . We don't need to state this property, as it follows from the field axioms:

Theorem. $a \times 0 = 0$

Proof.

$$\begin{aligned} a \times 0 &= a \times (1 - 1) \\ &= a \times 1 + a \times (-1) \\ &= a - a \\ &= 0 \end{aligned}$$

□

Similarly, we can prove corollaries from the field axioms.

Corollary. *If $ab = 0$, then either $a = 0$ or $b = 0$ (or both).*

Proof. Suppose $a \neq 0$. Then there exists a^{-1} such that

$$\begin{aligned} a^{-1}ab &= a^{-1} \times 0 \\ 1 \times b &= 0 \\ b &= 0 \end{aligned}$$

A similar argument follows when $b \neq 0$. □

The fundamental theorem of algebra relies on the above corollary when solving polynomials. If we factor a polynomial into the form $(x - r_1)(x - r_2) \cdots (x - r_k) = 0$, then we know the polynomial has roots r_1, r_2, \dots, r_k . This is only true because the left hand side of the factored expression only reaches zero when one of the factors is zero, i.e. when $x = r_i$.

1.2.1 Common Fields in Mathematics

The advantage of fields is that once a set is proven to obey the five field axioms, we can operate on elements in the field just like we would operate on real numbers. Besides the real numbers (which the concept of fields was designed to emulate), what are some other fields?

The rational numbers are a field. The numbers 0 and 1 are rational, so they are in the field. Since we add and multiply rational numbers just as we do real numbers, these operations commute, associate, and distribute. All that remains is to show that the rationals have additive and multiplicative inverses in the field. Let us consider a rational number p/q , where p and q are integers.

- We know that $-p/q$ is also rational, since $-p$ is still an integer. The additive inverse of a rational number is in the field of rational numbers.
- The additive inverse of p/q is q/p , which is also rational. The multiplicative inverse of a rational is also in the field.

So the rational numbers are a field. What does this mean? If we are given an algebraic expression, we can solve it by performing any algebraic manipulation and still be assured that the answer will be another rational number.

The integers, by contrast, are not a field. Every integer has a reciprocal ($2 \rightarrow 1/2$, $-100 \rightarrow -1/100$, etc.). However, the reciprocals are themselves not integers, so they are not in the same field. The field axioms require that the inverses for every element are members of the field. When constructing a field, every part of every axiom must be satisfied.

Let's see an example of this. Imagine the simple equation $y = ax + b$, which we solve for x to yield

$$x = \frac{y - b}{a}$$

If we wanted to solve this equation using only rational numbers, we would not need to change anything. So long as the values we input for the variables a , b , and y were rational, the value of x would also be rational. We solved the equation using field algebra, and the rationals constitute a field. Everything works out.

Now imagine you wanted only integer solutions. Even if the values for a , b , and y were integers, there is no guarantee that x would be an integer. ($a = 2$, $b = 4$, and $y = 3$ yields $x = -1/2$, for example). Because the integers are not a field, algebra does not work on them. In particular, the integers do not have integer multiplicative inverses (except for 1 and -1). When we divide by a , we assumed that the value $1/a$ exists in the field, which it does not.

Interestingly, the integers always have integer additive inverses, so the solution to the equation $y = x - b$ is always an integer (for integer y and b) since we could solve the equation with only additive inverses.

1.3 Vector Addition

Addition of two vectors is defined *elementwise*, or element-by-element.

$$\begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{pmatrix} + \begin{pmatrix} y_1 \\ y_2 \\ \vdots \\ y_n \end{pmatrix} = \begin{pmatrix} x_1 + y_1 \\ x_2 + y_2 \\ \vdots \\ x_n + y_n \end{pmatrix}$$

Since this is a direct extension of scalar addition, it is clear that vector addition commutes [$\mathbf{x} + \mathbf{y} = \mathbf{y} + \mathbf{x}$] and is associative [$\mathbf{x} + \mathbf{y} + \mathbf{z} = (\mathbf{x} + \mathbf{y}) + \mathbf{z} = \mathbf{x} + (\mathbf{y} + \mathbf{z})$].

The additive inverse of a vector \mathbf{x} (written as $-\mathbf{x}$) is also constructed elementwise:

$$-\mathbf{x} = \begin{pmatrix} -x_1 \\ -x_2 \\ \vdots \\ -x_n \end{pmatrix}$$

From our elementwise definition of vector addition, we can construct the zero element for the vector space. We know from the field axioms that $\mathbf{x} + \mathbf{0} = \mathbf{x}$, so the zero element must be a vector of the same dimension with all zero entries.

$$\mathbf{x} + \mathbf{0} = \begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{pmatrix} + \begin{pmatrix} 0 \\ 0 \\ \vdots \\ 0 \end{pmatrix} = \begin{pmatrix} x_1 + 0 \\ x_2 + 0 \\ \vdots \\ x_n + 0 \end{pmatrix} = \mathbf{x}$$

Notice that each set of n -dimensional vectors has its own zero element. In \mathbb{R}^2 , $\mathbf{0} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}$. In \mathbb{R}^3 , $\mathbf{0} = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}$.

1.4 Vector Multiplication is not Elementwise

What happens when we try to define multiplication as an elementwise operation? For example

$$\begin{pmatrix} -1 \\ 0 \\ 4 \end{pmatrix} \times \begin{pmatrix} 0 \\ 2 \\ 0 \end{pmatrix} = \begin{pmatrix} -1 \times 0 \\ 0 \times 2 \\ 4 \times 0 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix} = \mathbf{0}$$

This is bad. Very bad. Here we have an example where $\mathbf{xy} = \mathbf{0}$, but neither \mathbf{x} nor \mathbf{y} is the zero element $\mathbf{0}$. This is a direct violation of a corollary of the field axioms, so **elementwise vector multiplication is not a valid algebraic operation**.

Sadly, vectors are not a field. There is no way to define multiplication using only vectors that satisfies the field axioms. Nor is there anything close to a complete set of multiplicative inverses, or even the element $\mathbf{1}$. Instead, we will settle for a weaker result – showing that vectors live in a *normed inner product space*. The concepts of a vector norm and inner product will let us create most of the operations and elements that vectors need to be a field.

On the bright side, if vectors were a field this class would be far too short.

1.4.1 Do We Need Multiplication?

When you were first taught to multiply, it was probably introduced as a “faster” method of addition, i.e. $4 \times 3 = 3 + 3 + 3 + 3$. If so, why do we need multiplication as a separate requirement for fields? Couldn't we simply require the addition operator and construct multiplication from it? The answer is no, for two reasons. First, the idea of multiplication as a shortcut for addition only makes sense when discussing the non-negative integers. What, for example, does it mean to have -2.86×3 ? What does -2.86 groups look like in terms of addition?

Second, we must realize that multiplication is a much stronger relationship between numbers. To understand why, we should start talking about the “linear” part of linear algebra.

Also, the integers are not a field!

1.5 Linear Systems

Linear systems have two special properties.

1. **Proportionality.** If the input to a linear system is multiplied by a scalar, the output is multiplied by the same scalar: $f(kx) = kf(x)$.
2. **Additivity.** If two inputs are added, the result is the sum of the original outputs: $f(x_1 + x_2) = f(x_1) + f(x_2)$.

We can combine both of these properties into a single condition for linearity.

Definition. A system f is linear if and only if

$$f(k_1x_1 + k_2x_2) = k_1f(x_1) + k_2f(x_2)$$

for all inputs x_1 and x_2 and scalars k_1 and k_2 .

Consider a very simple function, $f(x) = x + 3$. Is this function linear? First we calculate the lefthand side of the definition of linearity.

$$f(k_1x_1 + k_2x_2) = k_1x_1 + k_2x_2 + 3$$

We compare this to the righthand side.

$$\begin{aligned} k_1f(x_1) + k_2f(x_2) &= k_1(x_1 + 3) + k_2(x_2 + 3) \\ &= k_1x_1 + k_2x_2 + 3(k_1 + k_2) \\ &\neq f(k_1x_1 + k_2x_2) \end{aligned}$$

This does not follow the definition of linearity. The function $f(x) = x + 3$ is not linear. Now let's look at a simple function involving multiplication: $f(x) = 3x$. Is this function linear?

$$\begin{aligned} f(k_1x_1 + k_2x_2) &= 3(k_1x_1 + k_2x_2) \\ &= k_1(3x_1) + k_2(3x_2) \\ &= k_1f(x_1) + k_2f(x_2) \end{aligned}$$

The function involving multiplication is linear.

These results might not be what you expected, at least concerning the nonlinearity of functions of the form $f(x) = x + b$. This is probably because in earlier math courses you referred to equations of straight lines ($y = mx + b$) as linear equations. In fact, any equation of this form (with $b \neq 0$) is called *affine*, not linear.

Truly linear functions have the property that $f(0) = 0$. Addition is, in a way, not “strong” enough to drive a function to zero. The expression $x + y$ is zero only when both x and y are zero. By contrast, the product xy is zero when either x or y is zero.

This follows from proportionality. If $f(k0) = kf(0)$ for all k , then $f(0)$ must equal zero.

1.6 Vector Norms

One of the nice properties of the real numbers is that they are well ordered. Being well ordered means that for any two real numbers, we can determine which number is larger (or if the two numbers are equal). Well orderedness allows us to make all sorts of comparisons between the real numbers.

Vectors are not well ordered. Consider the vectors $(3, 4)$ and $(5, 2)$. Which one is larger? Each vector has one element that is larger than the other (4 in the first, 5 in the second). There is no unambiguous way to place all vectors in order.

This doesn't stop us from making comparisons between vector quantities. Consider velocity, which, contrary to how many people use the term, is a vector. Since vectors are not ordered, we should not be able to compare velocities. Instead, we often compare speeds, which are the magnitude of the velocity vectors. Traveling 30 mph due north and 30 mph due east are technically two different velocities. However, they have the same magnitude (30 mph), so most people consider them equivalent.

Vector magnitudes are calculated by taking a *norm* of the vector. There are many different kinds of norms, but the most commonly used norm is the 2-norm (or Cartesian or Pythagorean norm). We will refer to the 2-norm as simply “the norm” unless we state otherwise. If we treat a vector as a point in n -dimensional space, the norm is the length of the arrow drawn from the origin to that point.

In 2D, the norm corresponds to the hypotenuse of the right triangle with sides equivalent to the two elements in the vector – hence the name “Pythagorean norm” (since the norm can be calculated by the Pythagorean theorem). In higher dimensions, we simply generalize the Pythagorean definition of the norm to

$$\|\mathbf{x}\| = \sqrt{x_1^2 + x_2^2 + \cdots + x_n^2}$$

In one dimension, taking the 2-norm yields the same result as taking the absolute value:

$$\|-3\| = \sqrt{(-3)^2} = 3 = |-3|$$

There are two useful properties of norms that derive directly from its definition. These properties must be true of all norms, not just the 2-norm.

1. **Nonnegativity.** $\|\mathbf{x}\| \geq 0$
2. **Zero Identity.** $\|\mathbf{x}\| = 0 \Leftrightarrow \mathbf{x} = \mathbf{0}$

1.6.1 Normalized (Unit) Vectors

A vector contains information about both its magnitude and its orientation. We've seen how to extract the magnitude as a scalar from the vector by taking the norm. Is it possible to similarly separate out the vector's orientation? Yes, by *normalizing* the vector. A normalized vector (or unit vector) is any vector with magnitude equal to one. We can convert a vector to a normalized vector by dividing each element by the vector's magnitude. For example

$$\mathbf{x} = \begin{pmatrix} 3 \\ -4 \end{pmatrix} \Rightarrow \|\mathbf{x}\| = \sqrt{3^2 + (-4)^2} = 5$$

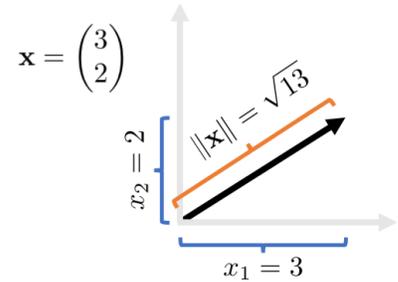


Figure 1.1: The vector norm.

We use a pair of two vertical bars ($\|\cdot\|$) to represent the norm. This differentiates the norm from the absolute value (which is, in fact, the 1-norm). Some texts use a subscript to identify which norm we are taking, i.e. $\|\mathbf{x}\|_2$ is the 2-norm of \mathbf{x} .

Be careful not to confuse the 2-norm and the absolute value, as they are not the same thing. Their equivalence in one dimension is a coincidence. However, the absolute value is a norm – it returns the magnitude of a number, but strips away the direction (negative or positive).

Intuitively, the idea of a normalized unit vector as a direction makes sense. If a vector is a product of both a magnitude and a direction, then the vector divided by the magnitude (the norm) should equal a direction (a unit vector).

The normalized unit vector ($\hat{\mathbf{x}}$) is

$$\hat{\mathbf{x}} = \begin{pmatrix} 3/\|\mathbf{x}\| \\ -4/\|\mathbf{x}\| \end{pmatrix} = \begin{pmatrix} 3/5 \\ -4/5 \end{pmatrix}$$

1.7 Scalar Vector Multiplication

We saw earlier that elementwise multiplication was a terrible idea. In fact, defining multiplication this way violates a corollary of the field axioms ($\mathbf{xy} = \mathbf{0}$ implies that $\mathbf{x} = \mathbf{0}$ or $\mathbf{y} = \mathbf{0}$). However, elementwise multiplication does work in one case – *scalar multiplication*, or the product between a scalar (real number) and a vector:

$$k\mathbf{x} = k \begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{pmatrix} = \begin{pmatrix} kx_1 \\ kx_2 \\ \vdots \\ kx_n \end{pmatrix}$$

where k is a scalar real number. Notice that scalar multiplication does not suffer from the same problem as elementwise vector multiplication. If $k\mathbf{x} = \mathbf{0}$, then either the scalar k equals zero or the vector \mathbf{x} must be the zero vector.

What happens when you multiply a vector by a scalar? For one, the norm changes:

$$\begin{aligned} \|k\mathbf{x}\| &= \sqrt{(kx_1)^2 + (kx_2)^2 + \cdots + (kx_n)^2} \\ &= \sqrt{k^2(x_1^2 + x_2^2 + \cdots + x_n^2)} \\ &= |k| \|\mathbf{x}\| \end{aligned}$$

Scalar multiplication scales the length of a vector by the scalar. If the scalar is negative, the direction of the vector “reverses”.

1.8 Inner (Dot) Products

One way to think of the product of two vectors is to consider the product of their norms (magnitudes). Such operations are common in mechanics. Work, for example, is the product of force and displacement. However, simply multiplying the magnitude of the force vector and the magnitude of the displacement vector disregards the orientation of the vectors. We know from physics that only the component of the force aligned with the displacement should count.

In general, we want an operation that multiplies the magnitude of one vector with the *projection* of a second vector onto the first. We call this operation the *inner product* or the *dot product*. Geometrically,

We use the hat symbol ($\hat{}$) over a unit vector to remind us that it has been normalized.

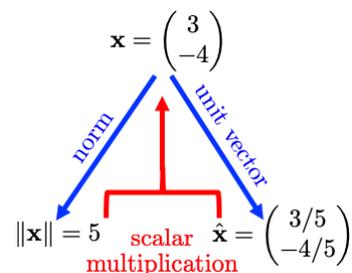


Figure 1.2: Vectors separate into a length (norm) and direction (unit vector). The length and direction can be combined by scalar multiplication

Remember that $\sqrt{k^2} = |k|$, not k itself. We consider the square root to be the positive root.

the dot product is a measure of both the product of the vectors' magnitudes and how well they are aligned. For vectors \mathbf{x} and \mathbf{y} the dot product is defined

$$\mathbf{x} \cdot \mathbf{y} = \|\mathbf{x}\| \|\mathbf{y}\| \cos \theta$$

where θ is the angle between the vectors.

If two vectors are perfectly aligned, $\theta = 0^\circ$ and the dot product is simply the product of the magnitudes. If the two vectors point in exactly opposite directions, $\theta = 180^\circ$ and the dot product is -1 times the product of the magnitudes. If the vectors are *orthogonal*, the angle between them is 90° , so $\cos \theta = 0$ and the dot product is zero. Thus, **the dot product of two vectors is zero if and only if the vectors are orthogonal.**

1.8.1 Computing the Dot Product

We know how to calculate norms, but how do we calculate the angle between two n -dimensional vectors? The answer is that we don't need to. There is an easier way to calculate $\mathbf{x} \cdot \mathbf{y}$ than the formula $\|\mathbf{x}\| \|\mathbf{y}\| \cos \theta$.

First, we need to define a special set of vectors – the unit vectors $\hat{\mathbf{e}}_i$. These are vectors that have only a single nonzero entry, a 1 at element i . For example,

$$\hat{\mathbf{e}}_1 = \begin{pmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{pmatrix}, \hat{\mathbf{e}}_2 = \begin{pmatrix} 0 \\ 1 \\ \vdots \\ 0 \end{pmatrix}, \hat{\mathbf{e}}_n = \begin{pmatrix} 0 \\ 0 \\ \vdots \\ 1 \end{pmatrix}$$

Every vector can be written as a sum of scalar product with unit vectors. For example,

$$\begin{aligned} \begin{pmatrix} -3 \\ 6 \\ 2 \end{pmatrix} &= -3 \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} + 6 \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix} + 2 \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} \\ &= -3\hat{\mathbf{e}}_1 + 6\hat{\mathbf{e}}_2 + 2\hat{\mathbf{e}}_3 \end{aligned}$$

In general

$$\begin{aligned} \mathbf{x} &= x_1 \begin{pmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{pmatrix} + x_2 \begin{pmatrix} 0 \\ 1 \\ \vdots \\ 0 \end{pmatrix} + \cdots + x_n \begin{pmatrix} 0 \\ 0 \\ \vdots \\ 1 \end{pmatrix} \\ &= \sum_{i=1}^n x_i \hat{\mathbf{e}}_i \end{aligned}$$

Now we see why we use the symbol \times for multiplication; the dot (\cdot) is reserved for the dot product.

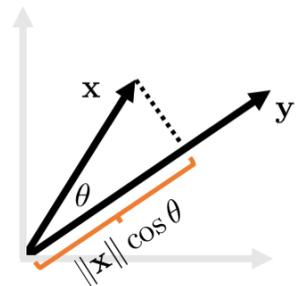


Figure 1.3: The projection of \mathbf{x} onto \mathbf{y} is a scalar equal to $\|\mathbf{x}\| \cos \theta$.

Now let's compute the dot product using the unit vector expansion for \mathbf{x} and \mathbf{y} .

$$\begin{aligned}\mathbf{x} \cdot \mathbf{y} &= (x_1 \hat{\mathbf{e}}_1 + x_2 \hat{\mathbf{e}}_2 + \cdots + x_n \hat{\mathbf{e}}_n) \cdot (y_1 \hat{\mathbf{e}}_1 + y_2 \hat{\mathbf{e}}_2 + \cdots + y_n \hat{\mathbf{e}}_n) \\ &= x_1 \hat{\mathbf{e}}_1 \cdot (y_1 \hat{\mathbf{e}}_1 + y_2 \hat{\mathbf{e}}_2 + \cdots + y_n \hat{\mathbf{e}}_n) \\ &\quad + x_2 \hat{\mathbf{e}}_2 \cdot (y_1 \hat{\mathbf{e}}_1 + y_2 \hat{\mathbf{e}}_2 + \cdots + y_n \hat{\mathbf{e}}_n) \\ &\quad + \cdots \\ &\quad + x_n \hat{\mathbf{e}}_n \cdot (y_1 \hat{\mathbf{e}}_1 + y_2 \hat{\mathbf{e}}_2 + \cdots + y_n \hat{\mathbf{e}}_n)\end{aligned}$$

Consider each of the terms $x_i \hat{\mathbf{e}}_i \cdot (y_1 \hat{\mathbf{e}}_1 + y_2 \hat{\mathbf{e}}_2 + \cdots + y_n \hat{\mathbf{e}}_n)$. By distribution, this is equivalent to

$$x_i y_1 \hat{\mathbf{e}}_i \cdot \hat{\mathbf{e}}_1 + \cdots + x_i y_j \hat{\mathbf{e}}_i \cdot \hat{\mathbf{e}}_j + \cdots + x_i y_n \hat{\mathbf{e}}_i \cdot \hat{\mathbf{e}}_n$$

The only nonzero term in this entire summation is $x_i y_i \hat{\mathbf{e}}_i \cdot \hat{\mathbf{e}}_i$, which equals $x_i y_i$. The dot product reduces to

$$\begin{aligned}\mathbf{x} \cdot \mathbf{y} &= x_1 y_1 + x_2 y_2 + \cdots + x_n y_n \\ &= \sum_{i=1}^n x_i y_i\end{aligned}$$

Although the above formula is convenient for computing dot products, it lacks the intuition of our previous method ($\mathbf{x} \cdot \mathbf{y} = \|\mathbf{x}\| \|\mathbf{y}\| \cos \theta$). Whenever you use the latter method, be sure to remember that you're really calculating the product magnitudes after one vector is projected onto the other.

1.8.2 Dot Product Summary

- Dot products are defined between two vectors with the same dimension.
- Dot products return a scalar from two vectors. This is the projected product of the two magnitudes.
- $\mathbf{x} \cdot \mathbf{y} = \|\mathbf{x}\| \|\mathbf{y}\| \cos \theta = x_1 y_1 + x_2 y_2 + \cdots + x_n y_n$
- $\mathbf{x} \cdot \mathbf{y} = 0 \Leftrightarrow \mathbf{x}$ and \mathbf{y} are orthogonal.
- Dot products are associative [$\mathbf{x} \cdot \mathbf{y} \cdot \mathbf{z} = (\mathbf{x} \cdot \mathbf{y}) \cdot \mathbf{z} = \mathbf{x} \cdot (\mathbf{y} \cdot \mathbf{z})$], commutative [$\mathbf{x} \cdot \mathbf{y} = \mathbf{y} \cdot \mathbf{x}$], and distributive over addition [$\mathbf{x} \cdot (\mathbf{y} + \mathbf{z}) = \mathbf{x} \cdot \mathbf{y} + \mathbf{x} \cdot \mathbf{z}$].

Think about the dot product $\hat{\mathbf{e}}_i \cdot \hat{\mathbf{e}}_j$. If $i = j$, this product is 1 since $\|\hat{\mathbf{e}}_i\| = \|\hat{\mathbf{e}}_j\| = 1$ and $\theta = 0^\circ$. However, if $i \neq j$, the vectors are always orthogonal and the dot product is 0.

2

Matrices

2.1 Matrix Multiplication

Let's take stock of the operations we've defined so far.

- The **norm** (magnitude) maps a vector to a scalar. ($\mathbb{R}^n \mapsto \mathbb{R}$)
- The **scalar product** maps a scalar and a vector to a new vector ($\mathbb{R} \times \mathbb{R}^n \mapsto \mathbb{R}^n$), but can only scale the magnitude of the vector (or flip it if the scalar is negative).
- The **dot product** maps two vectors to a scalar ($\mathbb{R}^n \times \mathbb{R}^n \mapsto \mathbb{R}$) by projecting one onto the other and multiplying the resulting magnitudes.

All of these operations appeared consistent with the field axioms. Unfortunately, we still do not have a true multiplication operation – one that can transform any vector into any other vector. Can we construct such an operation using only the above methods?

Let's construct a new vector \mathbf{y} from vector \mathbf{x} . To be as general as possible, we should let each element in \mathbf{y} be an arbitrary linear combination of the elements in \mathbf{x} . This implies that

$$\begin{aligned}y_1 &= a_{11}x_1 + a_{12}x_2 + \cdots + a_{1n}x_n \\y_2 &= a_{21}x_1 + a_{22}x_2 + \cdots + a_{2n}x_n \\&\vdots \\y_n &= a_{n1}x_1 + a_{n2}x_2 + \cdots + a_{nn}x_n\end{aligned}$$

where the scalars a_{ij} determine the relative weight of x_j when constructing y_i . There are n^2 scalars required to unambiguously map \mathbf{x} to \mathbf{y} . For convenience, we collect the set of weights into an n by n numeric grid called a *matrix*.

If \mathbf{A} is a real-valued matrix with dimensions $m \times n$, we say $\mathbf{A} \in \mathbb{R}^{m \times n}$ and $\dim(\mathbf{A}) = m \times n$.

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

What we have been calling “vectors” all along are really just matrices with only one column. Thinking of vectors as matrices lets us write a simple, yet powerful, definition of multiplication.

Definition. The product of matrices \mathbf{AB} is a matrix \mathbf{C} where each element c_{ij} in \mathbf{C} is the dot product between the i th row in \mathbf{A} and the j th column in \mathbf{B} :

$$c_{ij} = \mathbf{A}(i, :) \cdot \mathbf{B}(:, j)$$

Using this definition of matrix multiplication, the previous system of n equations becomes the matrix equation

$$\begin{pmatrix} y_1 \\ \vdots \\ y_n \end{pmatrix} = \begin{pmatrix} a_{11} & \cdots & a_{1n} \\ \vdots & \ddots & \vdots \\ a_{n1} & \cdots & a_{nn} \end{pmatrix} \begin{pmatrix} x_1 \\ \vdots \\ x_n \end{pmatrix}$$

or, more succinctly

$$\mathbf{y} = \mathbf{Ax}$$

2.1.1 Generalized Multiplication

In the previous example, both \mathbf{x} and \mathbf{y} were n -dimensional. This does not need to be the case. In general, the vector \mathbf{y} could have $m \neq n$ dimensions. The matrix \mathbf{A} would have m rows, each used to construct an element y_i in \mathbf{y} . However, the matrix \mathbf{A} would still need n columns to match the n rows in \mathbf{x} . (Each row in \mathbf{A} is “dotted” with the n -dimensional vector \mathbf{x} , and dot products require the two vectors have the same dimension.)

Any matrices \mathbf{A} and \mathbf{B} are *conformable* for multiplication if the number of columns in \mathbf{A} matches the number of rows in \mathbf{B} . If the dimensions of \mathbf{A} are $m \times n$ and the dimensions of \mathbf{B} are $n \times p$, then the product will be a matrix of dimensions $m \times p$.

Matrix multiplication is associative [$\mathbf{ABC} = (\mathbf{AB})\mathbf{C} = \mathbf{A}(\mathbf{BC})$] and distributive over addition [$\mathbf{A}(\mathbf{B} + \mathbf{C}) = \mathbf{AB} + \mathbf{AC}$], provided \mathbf{A} , \mathbf{B} , and \mathbf{C} are all conformable. However, it is **not** commutative. To see why, consider $\mathbf{A} \in \mathbb{R}^{m \times n}$ and $\mathbf{B} \in \mathbb{R}^{n \times p}$. The product \mathbf{AB} is an $m \times p$ matrix, but the product \mathbf{BA} is not conformable since $p \neq m$. Even if \mathbf{BA} were conformable, it is not the same as the product \mathbf{AB} .

$$\mathbf{A} = \begin{pmatrix} 1 & 2 \\ 3 & 4 \end{pmatrix}, \quad \mathbf{B} = \begin{pmatrix} 0 & 1 \\ -1 & 2 \end{pmatrix}$$

MATLAB returns an error that “matrix dimensions must agree” when multiplying non-conformable objects.

For the system $\mathbf{y} = \mathbf{Ax}$, if $\dim(\mathbf{A}) = m \times n$ and $\dim(\mathbf{x}) = n \times 1$, $\dim(\mathbf{y}) = m \times 1$, i.e. \mathbf{y} is a column vector in \mathbb{R}^m .

$$\mathbf{AB} = \begin{pmatrix} 1 \times 0 + 2 \times (-1) & 1 \times 1 + 2 \times 2 \\ 3 \times 0 + 4 \times (-1) & 3 \times 1 + 4 \times 2 \end{pmatrix} = \begin{pmatrix} -2 & 5 \\ -4 & 11 \end{pmatrix}$$

$$\mathbf{BA} = \begin{pmatrix} 0 \times 1 + 1 \times 3 & 0 \times 2 + 1 \times 4 \\ -1 \times 1 + 2 \times 3 & -1 \times 2 + 2 \times 4 \end{pmatrix} = \begin{pmatrix} 3 & 4 \\ 5 & 6 \end{pmatrix}$$

2.2 Identity Matrix

We need to find an element that serves as $\mathbf{1}$ for vectors. The field axioms define this element by the property that $1 \times x = x$ for all x in the field. For vectors, we defined multiplication to involve matrices, so the element $\mathbf{1}$ will be a matrix, which we will call the *identity matrix* \mathbf{I} . We require that

$$\mathbf{Ix} = \mathbf{xI} = \mathbf{x}$$

for all \mathbf{x} . Assuming that \mathbf{x} is n -dimensional, \mathbf{I} must have n columns to be conformable. Also, the output of \mathbf{Ix} has n elements, so \mathbf{I} must have n rows. Therefore, we know that \mathbf{I} is a square $n \times n$ matrix whenever \mathbf{x} has dimension n .

Consider the first row of \mathbf{I} , i.e. $\mathbf{I}(1, :)$. We know from the definition of \mathbf{I} that $\mathbf{I}(1, :) \cdot \mathbf{x} = x_1$, so $\mathbf{I}(1, :) = (1\ 0\ 0 \cdots 0)$. For the second row, $\mathbf{I}(2, :) \cdot \mathbf{x} = x_2$, so $\mathbf{I}(2, :) = (0\ 1\ 0 \cdots 0)$. In general, the i th row of \mathbf{I} has a 1 at position i and zeros everywhere else

$$\mathbf{I} = \begin{pmatrix} 1 & 0 & 0 & \cdots & 0 \\ 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \cdots & 0 \\ \vdots & & & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & 1 \end{pmatrix}$$

The identity matrix \mathbf{I} for a vector in \mathbb{R}^n is an $n \times n$ matrix with ones along the diagonal and zeroes everywhere else.

Our definition of the identity matrix also works for matrix multiplication. For any square matrix \mathbf{A}

$$\mathbf{IA} = \mathbf{AI} = \mathbf{A}$$

The identity matrix also works for non-square matrices; however, the dimensions of the identity matrix change if the multiplication is on the left or right side. If \mathbf{A} is an $m \times n$ matrix, then $\mathbf{IA} = \mathbf{A}$ if \mathbf{I} is an $m \times m$ identity matrix, and $\mathbf{AI} = \mathbf{A}$ if \mathbf{I} is an $n \times n$ identity matrix.

2.3 Matrix Transpose

The transpose operator flips the rows and columns of a matrix. The element a_{ij} in the original matrix becomes element a_{ji} in the transposed matrix. The transpose operator is a superscript “T”, as in \mathbf{A}^T .

In \mathbb{R}^1 , $\mathbf{I} = (1)$, which behaves like the real number 1.

Other notations for the matrix transpose include \mathbf{A}^t and \mathbf{A}' . The latter is used in MATLAB.

A transposed matrix is reflected about a diagonal drawn from the upper left to the lower right corner.

$$\begin{pmatrix} 1 & 2 \\ 3 & 4 \end{pmatrix}^T = \begin{pmatrix} 1 & 3 \\ 2 & 4 \end{pmatrix}$$

Transposing an $m \times n$ matrix creates an $n \times m$ matrix.

$$\begin{pmatrix} 1 & 2 \\ 3 & 4 \\ 5 & 6 \end{pmatrix}^T = \begin{pmatrix} 1 & 3 & 5 \\ 2 & 4 & 6 \end{pmatrix}$$

Transposing a column vector creates a row vector, and vice versa.

$$\begin{pmatrix} 1 \\ 2 \\ 3 \end{pmatrix}^T = (1 \ 2 \ 3), \quad (1 \ 2 \ 3)^T = \begin{pmatrix} 1 \\ 2 \\ 3 \end{pmatrix}$$

2.4 Solving Linear Systems

Remember back to algebra, when you were asked to solve small systems of equations like

$$\begin{aligned} a_{11}x_1 + a_{12}x_2 &= y_1 \\ a_{21}x_1 + a_{22}x_2 &= y_2 \end{aligned}$$

Your strategy was to manipulate the equations until they reach the form

$$\begin{aligned} x_1 &= y'_1 \\ x_2 &= y'_2 \end{aligned}$$

In matrix form, this process transforms a matrix \mathbf{A} into the identity matrix

$$\begin{pmatrix} a_{11} & a_{21} \\ a_{21} & a_{22} \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \begin{pmatrix} y_1 \\ y_2 \end{pmatrix} \rightarrow \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \begin{pmatrix} y'_1 \\ y'_2 \end{pmatrix}$$

This leads us to our first strategy for solving linear systems of the form $\mathbf{Ax} = \mathbf{y}$. We manipulate both sides of the equation (\mathbf{A} and \mathbf{y}) until \mathbf{A} becomes the identity matrix. The vector \mathbf{x} then equals the transformed vector \mathbf{y}' . Because we will be applying the same transformations to both \mathbf{A} and \mathbf{y} , it is convenient to collect them both into an *augmented matrix* $(\mathbf{A} \ \mathbf{y})$. For 2×2 system above, the augmented matrix is

$$\begin{pmatrix} a_{11} & a_{12} & y_1 \\ a_{21} & a_{22} & y_2 \end{pmatrix}$$

What operations can we use to transform \mathbf{A} into the identity matrix? There are three operations, called the *elementary row operations*, or EROs.

We often use the prime symbol (') to indicate that an unspecified new value is based on an old one. For example, y'_1 is a new value calculated from y_1 . In this case,

$$y'_1 = \frac{y_1 a_{22} - a_{12} y_2}{a_{11} a_{22} - a_{12} a_{21}}$$

1. Exchanging two rows. Since the order of the equations in our system is arbitrary, we can re-order the rows of the augmented matrix at will. By working with the augmented matrix, we ensure that both the left- and right-hand sides move together.

$$\begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{pmatrix} \xrightarrow{R_2 \leftrightarrow R_3} \begin{pmatrix} 1 & 2 & 3 \\ 7 & 8 & 9 \\ 4 & 5 & 6 \end{pmatrix}$$

2. Multiply any row by a scalar. Again, since we are working with the augmented matrix, multiplying a row by a scalar multiplies both the left- and right-hand sides of the equation by the same factor.

$$\begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{pmatrix} \xrightarrow{3R_2} \begin{pmatrix} 1 & 2 & 3 \\ 12 & 15 & 18 \\ 7 & 8 & 9 \end{pmatrix}$$

3. Add a scalar multiple of any row to any other row.

$$\begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{pmatrix} \xrightarrow{R_1 + 3R_2} \begin{pmatrix} 13 & 17 & 21 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{pmatrix}$$

Let's solve the system of equations

$$4x_1 + 8x_2 - 12x_3 = 44$$

$$3x_1 + 6x_2 - 8x_3 = 32$$

$$-2x_1 - x_2 = -7$$

This is a linear system of the form $\mathbf{Ax} = \mathbf{y}$ where the matrix \mathbf{A} and the vector \mathbf{y} are

$$\mathbf{A} = \begin{pmatrix} 4 & 8 & -12 \\ 3 & 6 & -8 \\ -2 & -1 & 0 \end{pmatrix}, \quad \mathbf{y} = \begin{pmatrix} 44 \\ 32 \\ -7 \end{pmatrix}$$

The augmented matrix is therefore

$$\begin{pmatrix} 4 & 8 & -12 & 44 \\ 3 & 6 & -8 & 32 \\ -2 & -1 & 0 & -7 \end{pmatrix}$$

Now we apply the elementary row operations.

$$\begin{aligned} & \xrightarrow{\frac{1}{4}R_1} \begin{pmatrix} 1 & 2 & -3 & 11 \\ 3 & 6 & -8 & 32 \\ -2 & -1 & 0 & -7 \end{pmatrix} \\ & \xrightarrow{R_2-3R_1} \begin{pmatrix} 1 & 2 & -3 & 11 \\ 0 & 0 & 1 & -1 \\ -2 & -1 & 0 & -7 \end{pmatrix} \\ & \xrightarrow{R_3+2R_1} \begin{pmatrix} 1 & 2 & -3 & 11 \\ 0 & 0 & 1 & -1 \\ 0 & 3 & -6 & 15 \end{pmatrix} \end{aligned}$$

Notice that after three steps we have a zero at position (2,2). We need to move this row farther down the matrix to continue; otherwise we can't cancel out the 3 below it. This operation is called a "pivot".

$$\begin{aligned} & \xrightarrow{R_2 \leftrightarrow R_3} \begin{pmatrix} 1 & 2 & -3 & 11 \\ 0 & 3 & -6 & 15 \\ 0 & 0 & 1 & -1 \end{pmatrix} \\ & \xrightarrow{\frac{1}{3}R_2} \begin{pmatrix} 1 & 2 & -3 & 11 \\ 0 & 1 & -2 & 5 \\ 0 & 0 & 1 & -1 \end{pmatrix} \end{aligned}$$

At this point we have a matrix in *row echelon form*. The bottom triangle looks like the identity matrix. We could stop here and solve the system using back substitution:

$$\begin{aligned} x_3 &= -1 \\ x_2 + -2(-1) &= 5 \Rightarrow x_2 = 3 \\ x_1 + 2(3) - 3(-1) &= 11 \Rightarrow x_1 = 2 \end{aligned}$$

Or, we could keep going and place the augmented matrix into *reduced row echelon form*.

$$\begin{aligned} & \xrightarrow{R_1-2R_2} \begin{pmatrix} 1 & 0 & 1 & 1 \\ 0 & 1 & -2 & 5 \\ 0 & 0 & 1 & -1 \end{pmatrix} \\ & \xrightarrow{R_1-R_3} \begin{pmatrix} 1 & 0 & 0 & 2 \\ 0 & 1 & -2 & 5 \\ 0 & 0 & 1 & -1 \end{pmatrix} \\ & \xrightarrow{R_2+2R_3} \begin{pmatrix} 1 & 0 & 0 & 2 \\ 0 & 1 & 0 & 3 \\ 0 & 0 & 1 & -1 \end{pmatrix} \end{aligned}$$

The left three columns are the identity matrix, so the resulting system of equations has been simplified to

$$\begin{aligned}x_1 &= 2 \\x_2 &= 3 \\x_3 &= -1\end{aligned}$$

2.5 Gaussian Elimination

Using EROs to transform the augmented matrix into the identity matrix is called *Gaussian elimination*. Let's develop an algorithm for Gaussian elimination for a general system of equations $\mathbf{Ax} = \mathbf{y}$ when \mathbf{A} is an $n \times n$ matrix. We begin with the augmented matrix

$$\begin{pmatrix} a_{11} & a_{12} & \cdots & a_{1n} & y_1 \\ a_{21} & a_{22} & \cdots & a_{2n} & y_2 \\ \vdots & & \ddots & & \vdots \\ a_{n1} & a_{n2} & \cdots & a_{nn} & y_n \end{pmatrix}$$

We need a 1 in the a_{11} position.

$$\xrightarrow{a_{11}^{-1}R_1} \begin{pmatrix} 1 & a_{11}^{-1}a_{12} & \cdots & a_{11}^{-1}a_{1n} & a_{11}^{-1}y_1 \\ a_{21} & a_{22} & \cdots & a_{2n} & y_2 \\ \vdots & & \ddots & & \vdots \\ a_{n1} & a_{n2} & \cdots & a_{nn} & y_n \end{pmatrix}$$

Now we zero out the a_{21} position using the first row multiplied by $-a_{21}$.

$$\xrightarrow{R_2 - a_{21}R_1} \begin{pmatrix} 1 & a_{11}^{-1}a_{12} & \cdots & a_{11}^{-1}a_{1n} & a_{11}^{-1}y_1 \\ 0 & a_{22} - a_{21}a_{11}^{-1}a_{12} & \cdots & a_{2n} - a_{21}a_{11}^{-1}a_{1n} & y_2 - a_{21}a_{11}^{-1}y_1 \\ \vdots & & \ddots & & \vdots \\ a_{n1} & a_{n2} & \cdots & a_{nn} & y_n \end{pmatrix}$$

We keep zeroing out the entries a_{31} through a_{n1} using the first row.

We end up with the matrix

$$\xrightarrow{R_n - a_{n1}R_1} \begin{pmatrix} 1 & a_{11}^{-1}a_{12} & \cdots & a_{11}^{-1}a_{1n} & a_{11}^{-1}y_1 \\ 0 & a_{22} - a_{21}a_{11}^{-1}a_{12} & \cdots & a_{2n} - a_{21}a_{11}^{-1}a_{1n} & y_2 - a_{21}a_{11}^{-1}y_1 \\ \vdots & & \ddots & & \vdots \\ 0 & a_{n2} - a_{n1}a_{11}^{-1}a_{12} & \cdots & a_{nn} - a_{n1}a_{11}^{-1}a_{1n} & y_n - a_{n1}a_{11}^{-1}y_1 \end{pmatrix}$$

This is looking a little complicated, so let's rewrite the matrix as

$$\begin{pmatrix} 1 & a'_{12} & \cdots & a'_{1n} & y'_1 \\ 0 & a'_{22} & \cdots & a'_{2n} & y'_2 \\ \vdots & & \ddots & & \vdots \\ 0 & a'_{n2} & \cdots & a'_{nn} & y'_n \end{pmatrix}$$

The first column looks like the identity matrix, which is exactly what we want. Our next goal is to put the lower half of an identity matrix in column 2 by setting $a'_{22} = 1$ and $a_{32}, \dots, a_{n2} = 0$. Notice that this is the same as applying the above procedure to the sub-matrix

$$\begin{pmatrix} a'_{22} & \cdots & a'_{2n} & y'_2 \\ \vdots & \ddots & & \vdots \\ a'_{n2} & \cdots & a'_{nn} & y'_n \end{pmatrix}$$

After that, we can continue recursively until the left part of the augmented matrix is the identity matrix. We can formalize the Gaussian elimination algorithm as follows:

```

function GAUSSIAN ELIMINATION
  for  $j = 1$  to  $n$  do                                ▷ For every column
     $a_{jj}^{-1} R_j$                                        ▷ Set the element on the diagonal to 1
    for  $i = j + 1$  to  $n$  do                             ▷ For every row below the diagonal
       $R_i - a_{ij} R_j$                                    ▷ Zero the below-diagonal element
    end for
  end for
end function

```

We're ignoring pivoting here. In general, we need to check that $a_{jj} \neq 0$; if it is, we swap the row for one below that has a nonzero term in the j th column.

2.6 Computational Complexity of Gaussian Elimination

How many computational operations are needed to perform Gaussian elimination on an $n \times n$ matrix? Let's start by counting operations when reducing the first column. The scaling of the first row ($a_{11}^{-1} R_1$) requires n operations. (There are $n + 1$ entries in the first row of the augmented matrix if you include the value y_1 ; however, we know the result in the first column, $a_{11}^{-1} a_{11}$, will always equal 1, so we don't need to compute it.) Similarly, zeroing out a single row below requires n multiplications and n subtractions. (Again, there are $n + 1$ columns, but we know the result in column 1 will be zero.) In the first column, there are $n - 1$ rows below to zero out, so the total number of operations is

$$\underbrace{n}_{a_{11}^{-1} R_1} + \underbrace{2(n-1)n}_{R_i - a_{i1} R_1} = 2n^2 - n$$

After we zero the bottom of the first row, we repeat the procedure on the $(n - 1) \times (n - 1)$ submatrix, and so on until we reach the 1×1 "submatrix" that includes only a_{nn} . We add up the number of

Remember that

$$\begin{aligned} \sum_{k=1}^n 1 &= n \\ \sum_{k=1}^n k &= \frac{n(n+1)}{2} \\ \sum_{k=1}^n k^2 &= \frac{n(n+1)(2n+1)}{6} \end{aligned}$$

operations for each of these n submatrices

$$\begin{aligned}
 &= \sum_{k=1}^n (2k^2 - k) \\
 &= 2 \sum_{k=1}^n k^2 - \sum_{k=1}^n k \\
 &= \frac{n(n+1)(2n+1)}{3} - \frac{n(n+1)}{2} \\
 &= \mathcal{O}(n^3)
 \end{aligned}$$

Thus the number of operations required to bring an $n \times n$ matrix into row-echelon form is on the order of n^3 . The number of operations needed to perform back substitution and solve the system is $\mathcal{O}(n^2)$. This raises two important points.

1. Gaussian elimination scales cubically. A system with twice as many equations takes eight times longer to solve.
2. The computational bottleneck is generating the row echelon matrix. Back substitution (or creating the reduced row echelon matrix) is significantly faster.

The \mathcal{O} , or “big-O” notation indicates the rate of growth of a function for large values. Any polynomial of degree d is $\mathcal{O}(d)$.

2.7 Solving Linear Systems in MATLAB

MATLAB has multiple functions for solving linear systems. Given variables \mathbf{A} and \mathbf{y} , you can use

- `linsolve(A,y)` to solve using LU decomposition, a variant of Gaussian elimination.
- `(A \ y)` to let MATLAB choose the best algorithm based on the size and structure of \mathbf{A} .
- `rref([A y])` to compute the reduced row echelon form of the augmented matrix.

The Finite Difference Method

When you first learned about derivatives, they were probably introduced as the limit of a finite difference

$$\frac{du}{dx} = \lim_{a \rightarrow b} \frac{u(b) - u(a)}{b - a}$$

As the distance between points a and b shrinks to zero, the finite difference becomes infinitesimal. The resulting differential equations must be solved with integral calculus. This chapter presents an alternative, numerical method for solving differential equations. Rather than shrink the finite differences all the way to zero, we leave a small but finite gap. The resulting algebraic equations approximate the differential equation and can be solved without any tools from calculus.

3.1 Finite Differences

Solving a differential equation analytically yields a solution over the entire domain (the region between the boundary conditions). By contrast, numerical methods find solutions to a differential equation at only a discrete set of points, or *nodes*, in the domain. The derivatives in the differential equation are discretized by converting them into *finite differences*. For example, we can approximate a first derivative as the change between two nodes divided by the distance between the nodes:

$$\frac{du}{dx} = \frac{u^{(k+1)} - u^{(k)}}{\Delta x}$$

where $u^{(k)}$ is the value of the variable at the k th node. To approximate a second derivative, we calculate the finite difference between nodes $u^{(k+1)}$ and $u^{(k)}$; and nodes $u^{(k)}$ and $u^{(k-1)}$. We divide this “difference between differences” by the distance between the centers of the nodes:

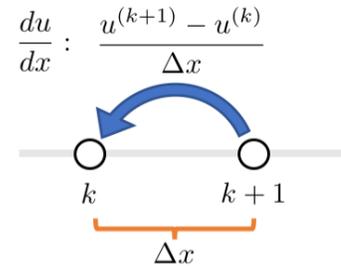


Figure 3.1: First-order finite difference approximation.

We write the value of u at node k as $u^{(k)}$. Using a superscript with parentheses avoids confusion with expressions u^k (the k th power of u) and u_k (the k th element of a vector named \mathbf{u}).

$$\begin{aligned}\frac{d^2u}{dx^2} &\approx \frac{\left(\frac{du}{dx}\right)^{(k+1)} - \left(\frac{du}{dx}\right)^{(k)}}{\Delta x} \\ &= \frac{\frac{u^{(k+1)} - u^{(k)}}{\Delta x} - \frac{u^{(k)} - u^{(k-1)}}{\Delta x}}{\Delta x} \\ &= \frac{u^{(k+1)} - 2u^{(k)} + u^{(k-1)}}{\Delta x^2}\end{aligned}$$

3.2 Linear Differential Equations

In order to generate a set of linear algebraic equations, the starting ODE or PDE must be linear. For differential equations, linearity means that the dependent variable (i.e. u) only appears in linear expressions; there can be nonlinearities involving only the independent variables. Remember that differentiation is a linear operator, so all derivatives of u are linear!

For example, consider the PDE with dependent variable $u(t, x)$:

$$t^2 \frac{\partial u}{\partial t} = c_1 \frac{\partial^2 u}{\partial x^2} + c_2 \sin(x) \frac{\partial u}{\partial x} + c_3 e^{tx}$$

This PDE is linear in u . However, the ODE

$$u \frac{du}{dx} = 0$$

is not linear. In general, for a variable $u(t, x)$, any terms of the form

$$f(t, x) \frac{\partial^n u}{\partial t^n} \quad \text{or} \quad f(t, x) \frac{\partial^n u}{\partial x^n}$$

are linear.

You might be wondering why we only require that a PDE be linear in the dependent variable. Why do nonlinearities in the independent variables not lead to nonlinear algebraic equations? Remember that in the finite difference method we discretize the independent variables across a grid and solve for the dependent variable at each node. Before solving, the value of the dependent variable is unknown. However, the value of all the independent variables are known, i.e. we know the location of every node in space and time. So, we can evaluate all terms involving the independent variables when setting up our equations.

3.3 Discretizing a Linear Differential Equation

Converting a linear differential equation into a set of linear algebraic equations requires three steps.

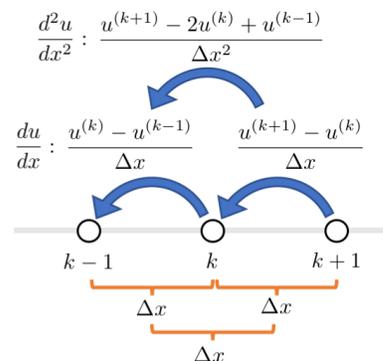


Figure 3.2: Second-order finite difference approximation.

The finite difference method will still work on nonlinear PDEs; however, the resulting set of equations are nonlinear.

As a rule of thumb, you can tell if a PDE is linear in u by ignoring the derivative operators and seeing if the resulting algebraic equation is linear in u .

1. Divide the domain into n equally-sized intervals. Creating n intervals requires $n + 1$ points, or *nodes*, labeled $0, 1, \dots, n$. The spacing between each node is $\Delta x = l/n$ where l is the length of the domain.
2. Starting with the interior nodes $(1, 2, \dots, n - 1)$, we rewrite the differential equation at each node using finite differences.

$$\begin{aligned} u &\rightarrow u^{(k)} \\ \frac{du}{dx} &\rightarrow \frac{u^{(k+1)} - u^{(k)}}{\Delta x} \\ \frac{d^2u}{dx^2} &\rightarrow \frac{u^{(k+1)} - 2u^{(k)} + u^{(k-1)}}{\Delta x^2} \end{aligned}$$

3. Add equations to enforce the boundary conditions at the boundary nodes.

For example, consider the ordinary differential equation

$$\frac{d^2u}{dx^2} + \frac{du}{dx} - 6u = 0, \quad u(0) = 0, \quad u(1) = 3$$

If we divide the domain $[0, 1]$ into four sections, then $n = 4$ and $\Delta x = l/n = 1/4 = 0.25$. We have five nodes $(0, 1, 2, 3, 4)$, three of which are interior nodes $(1, 2, 3)$. We rewrite the ODE using finite differences at each of the interior nodes.

$$\begin{aligned} \frac{u^{(2)} - 2u^{(1)} + u^{(0)}}{(0.25)^2} + \frac{u^{(2)} - u^{(1)}}{0.25} - 6u^{(1)} &= 0 \quad [\text{node 1}] \\ \frac{u^{(3)} - 2u^{(2)} + u^{(1)}}{(0.25)^2} + \frac{u^{(3)} - u^{(2)}}{0.25} - 6u^{(2)} &= 0 \quad [\text{node 2}] \\ \frac{u^{(4)} - 2u^{(3)} + u^{(2)}}{(0.25)^2} + \frac{u^{(4)} - u^{(3)}}{0.25} - 6u^{(3)} &= 0 \quad [\text{node 3}] \end{aligned}$$

which simplifies to the equations

$$\begin{aligned} 16u^{(0)} - 42u^{(1)} + 20u^{(2)} &= 0 \\ 16u^{(1)} - 42u^{(2)} + 20u^{(3)} &= 0 \\ 16u^{(2)} - 42u^{(3)} + 20u^{(4)} &= 0 \end{aligned}$$

Now we can add equations for the boundary nodes. Node 0 corresponds to $x = 0$, so the boundary condition tells us that $u^{(0)} = 0$. Similarly, node 4 corresponds to $x = 1$, so $u^{(4)} = 3$. Combining these two boundary equations with the equations for the interior nodes yields the final linear system

$$\begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ 16 & -42 & 20 & 0 & 0 \\ 0 & 16 & -42 & 20 & 0 \\ 0 & 0 & 16 & -42 & 20 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} u^{(0)} \\ u^{(1)} \\ u^{(2)} \\ u^{(3)} \\ u^{(4)} \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 3 \end{pmatrix}$$

Solving by Gaussian elimination yields

$$\begin{pmatrix} u^{(0)} \\ u^{(1)} \\ u^{(2)} \\ u^{(3)} \\ u^{(4)} \end{pmatrix} = \begin{pmatrix} 0.00 \\ 0.51 \\ 1.07 \\ 1.84 \\ 3.00 \end{pmatrix}$$

We can compare our numerical solution to the exact solution for this expression, which is

$$u(x) = \frac{3}{e^2 - e^{-3}} (e^{2x} - e^{-3x})$$

x	Numerical Solution	Exact Solution	Error
0	0.00	0.00	0.0%
0.25	0.51	0.48	5.7%
0.50	1.07	1.02	4.7%
0.75	1.84	1.79	2.6%
1	3.00	3.00	0.0%

We used only five nodes to solve this ODE, yet our relative error is less than 6%!

3.4 Boundary Conditions

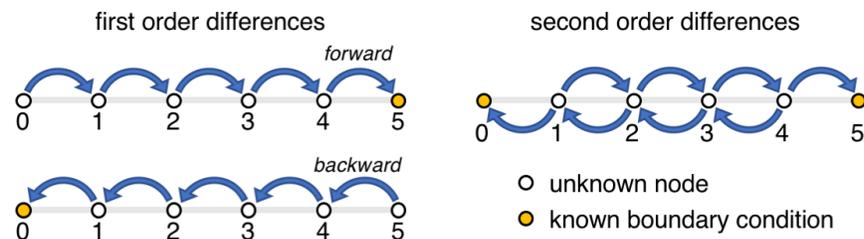
There are two ways to write a finite difference approximation for the first derivative at node k . We can write the *forward difference* using node $k + 1$:

$$\frac{du}{dx} \rightarrow \frac{u^{(k+1)} - u^{(k)}}{\Delta x}$$

or we can use the *backward difference* using the previous node at $k - 1$:

$$\frac{du}{dx} \rightarrow \frac{u^{(k)} - u^{(k-1)}}{\Delta x}$$

The choice of forward or backward differences depends on the boundary conditions. For a first order ODE, we are given a boundary condition at either $x = 0$ or $x = l$. If we are given $u(0)$, we use backward differences. If we are given $u(l)$, we use forward differences. Otherwise, we run out of nodes when writing equations for the finite differences.



You cannot avoid the issue by using the difference $u^{(0)} - u^{(1)}$ for node 0 and $u^{(1)} - u^{(0)}$ for node 1. These equations are *linearly dependent*, which leaves us with too little information when solving the system.

Second order finite differences require nodes on both sides. This is related to the necessity of two boundary conditions, since we cannot write finite difference approximations for nodes at either end. If you ODE comes with two boundary conditions, you can choose either forward or backward differences to approximate any first order derivatives.

4

Inverses, Solvability, and Rank

4.1 *Matrix Inverses*

So far we've demonstrated how Gaussian elimination can solve linear systems of the form $\mathbf{Ax} = \mathbf{y}$. Gaussian elimination involves a series of elementary row operations to transform the coefficient matrix \mathbf{A} into the identity matrix. While Gaussian elimination works well, our initial goal of defining an algebra for vectors requires something stronger – a multiplicative inverse. For vectors, the multiplicative inverse is called a *matrix inverse*. For any square matrix, a matrix inverse (if it exists) is a square matrix \mathbf{A}^{-1} such that

$$\mathbf{A}^{-1}\mathbf{A} = \mathbf{I} = \mathbf{A}\mathbf{A}^{-1}$$

If we could prove the existence of a *matrix inverse* for \mathbf{A} , we could solve a wide variety of linear systems, including $\mathbf{Ax} = \mathbf{y}$.

$$\begin{aligned}\mathbf{Ax} &= \mathbf{y} \\ \mathbf{A}^{-1}\mathbf{Ax} &= \mathbf{A}^{-1}\mathbf{y} \\ \mathbf{Ix} &= \mathbf{A}^{-1}\mathbf{y} \\ \mathbf{x} &= \mathbf{A}^{-1}\mathbf{y}\end{aligned}$$

The existence of the matrix inverse and being amenable to Gaussian elimination are related. While the end result is the same (a transformation of the coefficient matrix into the identity matrix), the processes are different. The Gaussian elimination algorithm applies a series of elementary row operations, including pivoting to avoid numerical issues. For a matrix inverse to exist, it must be able to capture all of these operations in a single matrix multiplication. This condensing of Gaussian elimination is not a trivial task.

Our first goal for this chapter is to prove the existence of the matrix inverse for any coefficient matrix that can be solved by Gaussian elimination. Then we will derive a method to construct the matrix inverse

This definition is analogous to the field axiom that there exists a^{-1} such that $a^{-1}a = 1$ for all nonzero a . Since scalar multiplication always commutes, $a^{-1}a = aa^{-1}$. Since matrix multiplication doesn't commute, we need to state this property separately.

if it exists. Finally, this chapter formalizes the requirements for solvability of a linear system of equations, which is related to the existence of the matrix inverse.

4.2 Elementary Matrices

Before proving the existence of the matrix inverse, we need to add another matrix manipulation tool to our arsenal – the *elementary matrix*. An elementary matrix is constructed by applying any single elementary row operation to the identity matrix. For example, consider swapping the second and third rows of the 3×3 identity matrix:

$$\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \xrightarrow{R_2 \leftrightarrow R_3} \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix} = \mathbf{E}_{R_2 \leftrightarrow R_3}$$

Notice what happens when we left multiply a matrix with an elementary matrix.

$$\mathbf{E}_{R_2 \leftrightarrow R_3} \begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{pmatrix} = \begin{pmatrix} 1 & 2 & 3 \\ 7 & 8 & 9 \\ 4 & 5 & 6 \end{pmatrix}$$

Multiplication by the elementary matrix exchanges the second and third rows – the same operation that created the elementary matrix. The same idea works for other elementary row operations, such as scalar multiplication

$$\mathbf{E}_{3R_2} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 3 & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

$$\mathbf{E}_{3R_2} \begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{pmatrix} = \begin{pmatrix} 1 & 2 & 3 \\ 12 & 15 & 18 \\ 7 & 8 & 9 \end{pmatrix}$$

and addition by a scaled row

$$\mathbf{E}_{R_2+2R_3} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 2 \\ 0 & 0 & 1 \end{pmatrix}$$

$$\mathbf{E}_{R_2+2R_3} \begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{pmatrix} = \begin{pmatrix} 1 & 2 & 3 \\ 18 & 21 & 24 \\ 7 & 8 & 9 \end{pmatrix}$$

4.3 Proof of Existence for the Matrix Inverse

We are now ready to prove the existence of the inverse for any coefficient matrix that is solvable by Gaussian elimination, i.e. any square

We use the notation \mathbf{E}_r to denote an elementary matrix constructed using the elementary row operation r .

Multiplication by an elementary matrix on the right applies the same operation to the columns of the matrix.

matrix that can be transformed into the identity matrix with elementary row operations. We will prove existence in three steps.

1. Construct a matrix \mathbf{P} that looks like a left inverse ($\mathbf{PA} = \mathbf{I}$).
2. Show that this left inverse is also a right inverse ($\mathbf{AP} = \mathbf{I}$).
3. Show that the matrix inverse is unique, implying that \mathbf{P} must be the inverse of \mathbf{A} .

Theorem. *Suppose matrix \mathbf{A} can be reduced to the identity matrix \mathbf{I} by elementary row operations. Then there exists a matrix \mathbf{P} such that $\mathbf{PA} = \mathbf{I}$.*

Proof. We assume that reducing \mathbf{A} to \mathbf{I} requires k elementary row operations. Let $\mathbf{E}_1, \dots, \mathbf{E}_k$ be the associated elementary matrices. Then

$$\begin{aligned}\mathbf{E}_k \mathbf{E}_{k-1} \cdots \mathbf{E}_2 \mathbf{E}_1 \mathbf{A} &= \mathbf{I} \\ (\mathbf{E}_k \mathbf{E}_{k-1} \cdots \mathbf{E}_2 \mathbf{E}_1) \mathbf{A} &= \mathbf{I} \\ \mathbf{PA} &= \mathbf{I}\end{aligned}$$

where $\mathbf{P} = \mathbf{E}_k \mathbf{E}_{k-1} \cdots \mathbf{E}_2 \mathbf{E}_1$. □

Theorem. *If $\mathbf{PA} = \mathbf{I}$ then $\mathbf{AP} = \mathbf{I}$.*

Proof.

$$\begin{aligned}\mathbf{PA} &= \mathbf{I} \\ \mathbf{PAP} &= \mathbf{IP} \\ \mathbf{P}(\mathbf{AP}) &= \mathbf{P}\end{aligned}$$

Since \mathbf{P} multiplied by \mathbf{AP} gives \mathbf{P} back again, \mathbf{AP} must equal the identity matrix ($\mathbf{AP} = \mathbf{I}$). □

Theorem. *The inverse of a matrix is unique.*

Proof. Let \mathbf{A}^{-1} be the inverse of \mathbf{A} , i.e. $\mathbf{A}^{-1}\mathbf{A} = \mathbf{I} = \mathbf{AA}^{-1}$. Suppose there exists a matrix $\mathbf{P} \neq \mathbf{A}^{-1}$ such that $\mathbf{PA} = \mathbf{I} = \mathbf{AP}$.

$$\begin{aligned}\mathbf{PA} &= \mathbf{I} \\ \mathbf{PAA}^{-1} &= \mathbf{IA}^{-1} \\ \mathbf{PI} &= \mathbf{A}^{-1} \\ \mathbf{P} &= \mathbf{A}^{-1}\end{aligned}$$

This contradicts our supposition that $\mathbf{P} \neq \mathbf{A}^{-1}$, so \mathbf{A}^{-1} must be unique. □

4.4 Computing the Matrix Inverse

Our proof of existence of the matrix inverse also provided a method of construction. We performed Gaussian elimination on a matrix, constructing an elementary matrix for each step. These elementary matrices were multiplied together to form the matrix inverse. In practice, this method would be wildly inefficient. Transforming an $n \times n$ matrix to reduced row echelon form requires $\mathcal{O}(n^2)$ elementary row operations, so we would need to construct and multiply $\mathcal{O}(n^2)$ elementary matrices. Since naive matrix multiplication requires $\mathcal{O}(n^3)$ operations per matrix, constructing a matrix inverse with this method requires $\mathcal{O}(n^5)$ operations! Since Gaussian elimination is $\mathcal{O}(n^3)$, we would be far better off avoiding the matrix inverse entirely.

Fortunately, there are better methods for constructing matrix inverses. One of the best is called the *side-by-side method*. To see how the side-by-side method works, consider constructing an inverse for the square matrix \mathbf{A} . We can use Gaussian elimination to transform \mathbf{A} into the identity matrix, which we can represent with a series of k elementary matrices.

$$\underbrace{\mathbf{E}_k \mathbf{E}_{k-1} \cdots \mathbf{E}_2 \mathbf{E}_1}_{\mathbf{A}^{-1}} \mathbf{A} = \mathbf{I}$$

What would happen if we simultaneously apply the same elementary row operations to another identity matrix?

$$\underbrace{\mathbf{E}_k \mathbf{E}_{k-1} \cdots \mathbf{E}_2 \mathbf{E}_1}_{\mathbf{A}^{-1}} \mathbf{I} = \mathbf{A}^{-1} \mathbf{I} = \mathbf{A}^{-1}$$

In the side-by-side method, we start with an augmented matrix containing the $n \times n$ matrix \mathbf{A} and an $n \times n$ identity matrix. Then we apply Gaussian elimination to transform \mathbf{A} into \mathbf{I} . The augmented matrix ensure that the same elementary row operations will be applied to the identity matrix, yielding the inverse of \mathbf{A} :

$$(\mathbf{A} \quad \mathbf{I}) \xrightarrow{\text{EROs}} (\mathbf{I} \quad \mathbf{A}^{-1})$$

Let's solve the following system by constructing the matrix inverse.

$$\begin{aligned} 3x_1 + 2x_2 &= 7 \\ x_1 + x_2 &= 4 \end{aligned}$$

In matrix form,

$$\mathbf{A} = \begin{pmatrix} 3 & 2 \\ 1 & 1 \end{pmatrix}, \quad \mathbf{y} = \begin{pmatrix} 7 \\ 4 \end{pmatrix}$$

Gaussian elimination requires $\mathcal{O}(n^2)$ row operations but $\mathcal{O}(n^3)$ total operations. Each row operation requires $\mathcal{O}(n)$ individual operations – one for each column.

The fastest matrix multiplication algorithm is $\mathcal{O}(n^{2.7373})$, although this is not a big help in practice.

Like the augmented matrix $(\mathbf{A} \mathbf{y})$, there is no direct interpretation of $(\mathbf{A} \mathbf{I})$. It is simply a convenient way to apply the same EROs to both \mathbf{A} and \mathbf{I} .

We start with the augmented matrix for the side-by-side method.

$$\begin{aligned} \begin{pmatrix} 3 & 2 & 1 & 0 \\ 1 & 1 & 0 & 1 \end{pmatrix} &\xrightarrow{R_1 \leftrightarrow R_2} \begin{pmatrix} 1 & 1 & 0 & 1 \\ 3 & 2 & 1 & 0 \end{pmatrix} \\ &\xrightarrow{R_2 - 3R_1} \begin{pmatrix} 1 & 1 & 0 & 1 \\ 0 & -1 & 1 & -3 \end{pmatrix} \\ &\xrightarrow{-R_2} \begin{pmatrix} 1 & 1 & 0 & 1 \\ 0 & 1 & -1 & 3 \end{pmatrix} \\ &\xrightarrow{R_1 - R_2} \begin{pmatrix} 1 & 0 & 1 & -2 \\ 0 & 1 & -1 & 3 \end{pmatrix} \end{aligned}$$

Thus, the matrix inverse is

$$\mathbf{A}^{-1} = \begin{pmatrix} 1 & -2 \\ -1 & 3 \end{pmatrix}$$

and we can compute the solution by matrix multiplication.

$$\mathbf{x} = \mathbf{A}^{-1}\mathbf{y} = \begin{pmatrix} 1 & -2 \\ -1 & 3 \end{pmatrix} \begin{pmatrix} 7 \\ 4 \end{pmatrix} = \begin{pmatrix} -1 \\ 5 \end{pmatrix}$$

We can verify that \mathbf{A}^{-1} is a matrix inverse

$$\begin{aligned} \mathbf{A}^{-1}\mathbf{A} &= \begin{pmatrix} 1 & -2 \\ -1 & 3 \end{pmatrix} \begin{pmatrix} 3 & 2 \\ 1 & 1 \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} = \mathbf{I} \\ \mathbf{A}\mathbf{A}^{-1} &= \begin{pmatrix} 3 & 2 \\ 1 & 1 \end{pmatrix} \begin{pmatrix} 1 & -2 \\ -1 & 3 \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} = \mathbf{I} \end{aligned}$$

The nice thing about having a matrix inverse is that if only the right hand side of a system of equations change, we do not need to retransform the coefficient matrix. For example, to solve

$$\begin{aligned} 3x_1 + 2x_2 &= 1 \\ x_1 + x_2 &= 3 \end{aligned}$$

we can re-use \mathbf{A}^{-1} since \mathbf{A} is unchanged (only \mathbf{y} is different).

$$\mathbf{x} = \mathbf{A}^{-1}\mathbf{y} = \begin{pmatrix} 1 & -2 \\ -1 & 3 \end{pmatrix} \begin{pmatrix} 1 \\ 3 \end{pmatrix} = \begin{pmatrix} -5 \\ 8 \end{pmatrix}$$

4.5 Numerical Issues

Matrix inverses are a powerful method for solving linear systems. However, calculating a matrix inverse should always be your last resort. There are far more efficient and numerically stable methods for solving linear systems. Reasons against using a matrix inverse include:

1. **Computation time.** While the side-by-side method is more efficient for constructing inverses than multiplying elementary matrices, it is still slower than Gaussian elimination for solving linear systems of the form $\mathbf{Ax} = \mathbf{y}$. Both side-by-side and Gaussian elimination reduce an augmented matrix. For an $n \times n$ matrix \mathbf{A} , the augmented matrix for Gaussian elimination ($\mathbf{A} \mathbf{y}$) is $n \times (n+1)$. The augmented matrix for the side-by-side method ($\mathbf{A} \mathbf{I}$) is $n \times 2n$. Solving for the inverse requires nearly twice the computations as solving the linear system directly. Having the inverse allows us to “resolve” the system for a new right hand side (\mathbf{y}) for only the cost of a matrix multiplication. However, there are variants of Gaussian elimination – such as LU decomposition – that allow resolving without repeating the entire reduction of the coefficient matrix \mathbf{A} .

2. **Memory.** Most large matrices in engineering are *sparse*. Sparse matrices contain very few nonzero entries; matrices with less than 0.01% nonzero entries are not uncommon. Examples of sparse matrices include matrices generated from finite difference approximations or matrices showing connections between nodes in large networks. Computers store sparse matrices by only storing the nonzero entries and their locations. However, there is no guarantee that the inverse of a sparse matrix will also be sparse. Consider the arrow matrix, a matrix with ones along the diagonal and last column and row.

$$\mathbf{A} = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \end{pmatrix}$$

An $n \times n$ arrow matrix has n^2 entries but only $3n - 2$ nonzeros. However, the inverse of an arrow matrix always has 100% nonzeros. For example, the inverse of the 8×8 matrix above is

$$\mathbf{A}^{-1} = \frac{1}{6} \begin{pmatrix} 5 & -1 & -1 & -1 & -1 & -1 & -1 & 1 \\ -1 & 5 & -1 & -1 & -1 & -1 & -1 & 1 \\ -1 & -1 & 5 & -1 & -1 & -1 & -1 & 1 \\ -1 & -1 & -1 & 5 & -1 & -1 & -1 & 1 \\ -1 & -1 & -1 & -1 & 5 & -1 & -1 & 1 \\ -1 & -1 & -1 & -1 & -1 & 5 & -1 & 1 \\ -1 & -1 & -1 & -1 & -1 & -1 & 5 & 1 \\ 1 & 1 & 1 & 1 & 1 & 1 & 1 & -1 \end{pmatrix}$$

Imagine the connection matrix for Facebook. It would have hundreds of millions of rows and columns, but each person (row) would only have nonzero entries for a few hundred people (columns) that they knew.

A 1000×1000 arrow matrix has less than 0.3% nonzeros.

Calculating the inverse of a large, sparse matrix could require orders of magnitude more memory than the original matrix. For some matrices, storing – much less computing – the inverse is impossible.

Despite its disadvantages, the matrix inverse is still a powerful construct. A multiplicative inverse for matrices is necessary for many algebraic manipulations, and the inverse can be used to simply or prove many matrix equations. Just remember to think critically about the need for a matrix inverse before calculating one.

4.6 Inverses of Elementary Matrices

We conclude with another interesting property of elementary matrices. We said before that left multiplication by an elementary matrix performs an elementary row operation (the same ERO that was used to construct the elementary matrix). Left multiplication by the inverse of an elementary matrix “undoes” the operation of the elementary matrix. For example, the elementary matrix \mathbf{E}_{3R_2} scales the second row by two. The inverse $\mathbf{E}_{3R_2}^{-1}$ would scale the second row by $1/2$, undoing the scaling by two. Similarly, $\mathbf{E}_{R_2 \leftrightarrow R_3}$ swaps rows two and three, and $\mathbf{E}_{R_2 \leftrightarrow R_3}^{-1}$ swaps them back. The proof of this property is straightforward.

Theorem. *If the elementary matrix \mathbf{E}_r performs the elementary row operation r , then left multiplication by the inverse \mathbf{E}_r^{-1} undoes this operation.*

Proof.

$$\mathbf{E}_r^{-1}(\mathbf{E}_r \mathbf{A}) = (\mathbf{E}_r^{-1} \mathbf{E}_r) \mathbf{A} = (\mathbf{I}) \mathbf{A} = \mathbf{A}$$

□

4.7 Rank

Consider the linear system

$$\begin{pmatrix} 1 & 0 & 3 \\ 0 & 2 & -4 \\ -2 & 0 & -6 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} = \begin{pmatrix} 2 \\ -2 \\ -4 \end{pmatrix}$$

and the row echelon form of the associated augmented matrix

$$\begin{pmatrix} 1 & 0 & 3 & 2 \\ 0 & 2 & -4 & -2 \\ -2 & 0 & -6 & -4 \end{pmatrix} \xrightarrow{\frac{1}{2}R_2 \rightarrow R_3 + 2R_1} \begin{pmatrix} 1 & 0 & 3 & 2 \\ 0 & 1 & -2 & -1 \\ 0 & 0 & 0 & 0 \end{pmatrix}$$

Notice that the last row is all zeros. We have no information about the last entry (x_3). However, this does not mean we cannot solve the

linear system. Since x_3 is unknown, let us assign its value the symbol α . Then, by back substitution

$$\begin{aligned} x_3 &= \alpha \\ x_2 - 2x_3 &= -1 \Rightarrow x_2 = 2\alpha - 1 \\ x_1 + 3x_3 &= 2 \Rightarrow x_1 = 2 - 3\alpha \end{aligned}$$

The above linear system has not one solution, but infinitely many. There is a solution for every value of the parameter α , so we say the system has a *parameterized solution*.

Parameterized solutions are necessary any time row echelon reduction of a matrix leads to one or more rows with all zero entries. The number of nonzero rows in the row echelon form of a matrix is the matrix's *rank*. The rank of a matrix can be calculated by counting the number of nonzero rows after the matrix is transformed into row echelon form by Gaussian elimination. In general, if a matrix with n columns has rank n , it is possible to find a unique solution to the system $\mathbf{Ax} = \mathbf{y}$. If $\text{rank}(\mathbf{A}) < n$, there may be infinitely many solutions. These solutions require that we specify $n - \text{rank}(\mathbf{A})$ parameters.

We denote the rank of a matrix \mathbf{A} as $\text{rank}(\mathbf{A})$.

Matrices has both a *row rank* (the number of nonzero rows in row-reduced echelon form) and a *column rank* (the number of nonzero columns in a column-reduced echelon form). Thus the concept of rank also applies to nonsquare matrices. However, the row and column ranks are always equivalent, even if the matrix is not square:

Theorem. *The row rank of a matrix equals the column rank of the matrix, i.e. $\text{rank}(\mathbf{A}) = \text{rank}(\mathbf{A}^T)$.*

Proof. We will defer the proof of this theorem until our discussion of the Fundamental Theorem of Linear Algebra. \square

The equivalence of the row and column ranks implies an upper bound on the rank of nonsquare matrices.

Theorem. *The rank of a matrix is less than or equal to the smallest dimension of the matrix, i.e. $\text{rank}(\mathbf{A}) \leq \min(\dim \mathbf{A})$.*

Proof. The row rank of \mathbf{A} is the number of nonzero rows in the row-reduced \mathbf{A} , so the rank of \mathbf{A} must be less than the number of rows in \mathbf{A} . Since the row rank is also equal to the column rank, there must also be $\text{rank}(\mathbf{A})$ nonzero columns in the column-reduced \mathbf{A} . So the rank of \mathbf{A} must never be larger than either the number of rows or number of columns in \mathbf{A} . \square

A matrix that has the maximum possible rank (rank n for an $n \times n$ square matrix or rank $\min(m, n)$ for an $m \times n$ rectangular matrix), we say the matrix is *full rank*. A matrix that is not full rank is *rank deficient*.

Linear Independence

The notion of rank is deeply tied to the concept of *linear independence*. A vector \mathbf{x}_i is linearly dependent on a set of n vectors $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n$ if there exists coefficients c_1, c_2, \dots, c_n such that

$$\mathbf{x}_i = c_1\mathbf{x}_1 + c_2\mathbf{x}_2 + \dots + c_n\mathbf{x}_n$$

A set of vectors are *linearly dependent* if one of the vectors can be expressed as a linear combination of some of the others. This is analogous to saying there exists a set of coefficients c_1, \dots, c_n , not all equal to zero, such that

$$c_1\mathbf{x}_1 + c_2\mathbf{x}_2 + \dots + c_n\mathbf{x}_n = \mathbf{0}$$

If a matrix with n rows has rank $k < n$, then $n - k$ of the rows are linearly dependent on the other k rows. Going back to our previous example, the matrix

$$\begin{pmatrix} 1 & 0 & 3 \\ 0 & 2 & -4 \\ -2 & 0 & -6 \end{pmatrix} \rightarrow \begin{pmatrix} 1 & 0 & 3 \\ 0 & 1 & -2 \\ 0 & 0 & 0 \end{pmatrix}$$

has rank 2 since there are two nonzero rows in the row-reduced matrix. Therefore, one of the rows must be linearly dependent on the other rows. Indeed, we see that the last row $(-2 \ 0 \ -6)$ is -2 times the first row $(1 \ 0 \ 3)$. During row reduction by Gaussian elimination, any linearly dependent rows will be completely zeroed out, revealing the rank of the matrix.

Rank and linear dependence tell us about the information content of a coefficient matrix. If some of the rows of the coefficient matrix are linearly dependent, then matrix is rank deficient and no unique solution exists. These matrices are also information deficient – we do not have one independent expression for each variable. Without a separate piece of information for each variable, we cannot unique map between the input \mathbf{x} and the output \mathbf{y} . However, if we introduce a separate parameter for each zeroed row, we are artificially providing the missing information. We can find a new solution every time we specify values for the parameters.

Homogeneous Systems ($\mathbf{Ax} = \mathbf{0}$)

A linear systems of equations is *homogeneous* if and only if the right hand side vector (\mathbf{y}) is equal to the zero vector ($\mathbf{0}$). Homogeneous systems always have at least one solution, $\mathbf{x} = \mathbf{0}$, since $\mathbf{A}\mathbf{0} = \mathbf{0}$. The zero solution to a homogeneous system is called the *trivial solution*.

Note the difference between $\mathbf{x}_1, \dots, \mathbf{x}_n$, a set of n vectors; and x_1, \dots, x_n , a set of n scalars that form the elements of a vector \mathbf{x} .

Some homogeneous systems have a nontrivial solution, i.e. a solution $\mathbf{x} \neq \mathbf{0}$. If a homogeneous system has a nontrivial solution, then it has infinitely many solutions, a result we state as follows

Theorem. *Any linear combination of nontrivial solutions to a homogeneous linear system is also a solution.*

Proof. Suppose we had two solutions, \mathbf{x} and \mathbf{x}' to the homogeneous system $\mathbf{A}\mathbf{x} = \mathbf{0}$. Then

$$\begin{aligned}\mathbf{A}(k\mathbf{x} + k'\mathbf{x}') &= \mathbf{A}(k\mathbf{x}) + \mathbf{A}(k'\mathbf{x}') \\ &= k(\mathbf{A}\mathbf{x}) + k'(\mathbf{A}\mathbf{x}') \\ &= k(\mathbf{0}) + k'(\mathbf{0}) \\ &= \mathbf{0}\end{aligned}$$

Since there are infinitely many scalars k and k' , we can generate infinitely many solutions to the homogeneous system $\mathbf{A}\mathbf{x} = \mathbf{0}$. \square

There is a connection between the solvability of nonhomogeneous systems $\mathbf{A}\mathbf{x} = \mathbf{y}$ and the corresponding homogeneous system $\mathbf{A}\mathbf{x} = \mathbf{0}$. If there exists at least one solution to $\mathbf{A}\mathbf{x} = \mathbf{y}$ and a nontrivial solution to $\mathbf{A}\mathbf{x} = \mathbf{0}$, then there are infinitely many solutions to $\mathbf{A}\mathbf{x} = \mathbf{y}$. To see why, let \mathbf{x}_{nh} be the solution to the nonhomogeneous system ($\mathbf{A}\mathbf{x}_{\text{nh}} = \mathbf{y}$) and \mathbf{x}_{h} be a nontrivial solution to the homogeneous system $\mathbf{A}\mathbf{x}_{\text{h}} = \mathbf{0}$. Then any of the infinite linear combinations $\mathbf{x}_{\text{nh}} + k\mathbf{x}_{\text{h}}$ is also a solution to $\mathbf{A}\mathbf{x} = \mathbf{y}$ since

$$\mathbf{A}(\mathbf{x}_{\text{nh}} + k\mathbf{x}_{\text{h}}) = \mathbf{A}\mathbf{x}_{\text{nh}} + k\mathbf{A}\mathbf{x}_{\text{h}} = \mathbf{y} + k\mathbf{0} = \mathbf{y}$$

General Solvability

For any linear system of equations, we can use the rank of the coefficient matrix and the augmented matrix to determine the existence and number of solutions. The relationship between solvability and rank is captured by the Rouché-Capelli theorem:

Theorem. *A linear system $\mathbf{A}\mathbf{x} = \mathbf{y}$ where $\mathbf{x} \in \mathbb{R}^n$ has a solution if and only if the rank of the coefficient matrix equals the rank of the augmented matrix, i.e. $\text{rank}(\mathbf{A}) = \text{rank}([\mathbf{A} \ \mathbf{y}])$. Furthermore, the solution is unique if $\text{rank}(\mathbf{A}) = n$; otherwise there are infinitely many solutions.*

Proof. We will sketch several portions of this proof to give intuition about the theorem. A more rigorous proof is beyond the scope of this class.

This proof is equivalent to showing that $\mathbf{A}\mathbf{x} = \mathbf{0}$ satisfies our definition of linear systems: $f(k_1x_1 + k_2x_2) = k_1f(x_1) + k_2f(x_2)$.

1. **Homogeneous systems.** For a homogeneous system $\mathbf{Ax} = \mathbf{0}$, we know that $\text{rank}(\mathbf{A}) = \text{rank}([\mathbf{A} \ \mathbf{0}])$. (Since the rank of \mathbf{A} is equal to the number of nonzero columns, adding another column of zeros will never change the rank.) Thus, we know that homogeneous systems are always solvable, at least by the trivial solution $\mathbf{x} = \mathbf{0}$. If $\text{rank}(\mathbf{A}) = n$, then the trivial solution is unique and is the only solution. If $\text{rank}(\mathbf{A}) < n$, there are infinitely many parameterized solutions.
2. **Full rank, nonhomogeneous systems.** For a nonhomogeneous system ($\mathbf{Ax} = \mathbf{y}$, $\mathbf{y} \neq \mathbf{0}$), we expect a unique solution if and only if adding the column \mathbf{y} to the coefficient matrix doesn't change the rank. For this to be true, \mathbf{y} must be linearly dependent on the other columns in \mathbf{A} ; otherwise, adding a new linearly independent column would increase the rank. If \mathbf{y} is linearly dependent on the n columns of \mathbf{A} , it must be true that there exists weights c_1, c_2, \dots, c_n such that

$$c_1 \mathbf{A}(:, 1) + c_2 \mathbf{A}(:, 2) + \dots + c_n \mathbf{A}(:, n) = \mathbf{y}$$

based on the definition of linear dependence. But the above expression can be rewritten in matrix form as

$$\begin{pmatrix} \mathbf{A}(:, 1) & \mathbf{A}(:, 2) & \dots & \mathbf{A}(:, n) \end{pmatrix} \begin{pmatrix} c_1 \\ c_2 \\ \vdots \\ c_n \end{pmatrix} = \mathbf{Ac} = \mathbf{y}$$

which shows that the system has a unique solution $\mathbf{x} = \mathbf{c}$.

3. **Rank deficient, nonhomogeneous systems.** Let $\text{rank}(\mathbf{A}) = k < n$. Then the row-reduced form of \mathbf{A} will have k rows that resemble the identity matrix and $n - k$ rows of all zeros:

$$\begin{pmatrix} a_{11} & a_{12} & \dots & a_{1k} & \dots & a_{1n} \\ a_{21} & a_{22} & \dots & a_{2k} & \dots & a_{2n} \\ \vdots & \vdots & & \vdots & & \vdots \\ a_{k1} & a_{k2} & \dots & a_{kk} & \dots & a_{kn} \\ a_{k+1,1} & a_{k+1,2} & \dots & a_{k+1,k} & \dots & a_{k+1,n} \\ \vdots & \vdots & & \vdots & & \vdots \\ a_{n1} & a_{n2} & \dots & a_{nk} & \dots & a_{nn} \end{pmatrix} \xrightarrow{\text{EROs}} \begin{pmatrix} 1 & a'_{12} & \dots & a'_{1k} & \dots & a'_{1n} \\ 0 & 1 & \dots & a'_{2k} & \dots & a'_{2n} \\ \vdots & \vdots & & \vdots & & \vdots \\ 0 & 0 & \dots & 1 & \dots & a'_{kn} \\ 0 & 0 & \dots & 0 & \dots & 0 \\ \vdots & \vdots & & \vdots & & \vdots \\ 0 & 0 & \dots & 0 & \dots & 0 \end{pmatrix}$$

Now imagine we performed the same row reduction on the augmented matrix (\mathbf{Ay}). We would still end up with $n - k$ rows with

Take time to understand the connection between the solvability of $\mathbf{Ax} = \mathbf{y}$ and the ability to express \mathbf{y} as a linear combination of the columns in \mathbf{A} . This relationship is the basis of Part III of the course.

zeros in the first n columns (the columns of \mathbf{A}):

$$\left(\begin{array}{cccccc} a_{11} & a_{12} & \cdots & a_{1k} & \cdots & a_{1n} & y_1 \\ a_{21} & a_{22} & \cdots & a_{2k} & \cdots & a_{2n} & y_2 \\ & \vdots & & & & & \vdots \\ a_{k1} & a_{k2} & \cdots & a_{kk} & \cdots & a_{kn} & y_k \\ a_{k+1,1} & a_{k+1,2} & \cdots & a_{k+1,k} & \cdots & a_{k+1,n} & y_{k+1} \\ & \vdots & & & & & \vdots \\ a_{n1} & a_{n2} & \cdots & a_{nk} & \cdots & a_{nn} & y_n \end{array} \right) \xrightarrow{\text{EROs}} \left(\begin{array}{cccccc} 1 & a'_{12} & \cdots & a'_{1k} & \cdots & a'_{1n} & y'_1 \\ 0 & 1 & \cdots & a'_{2k} & \cdots & a'_{2n} & y'_2 \\ & \vdots & & & & & \vdots \\ 0 & 0 & \cdots & 1 & \cdots & a'_{kn} & y'_k \\ 0 & 0 & \cdots & 0 & \cdots & 0 & y'_{k+1} \\ & \vdots & & & & & \vdots \\ 0 & 0 & \cdots & 0 & \cdots & 0 & y'_n \end{array} \right)$$

We know that if $y'_{k+1}, \dots, y'_n = 0$, we can solve this system by designating $n - k$ parameters for the variables x_{k+1}, \dots, x_n for which we have no information. However, notice what happens if any of the values y'_{k+1}, \dots, y'_n are nonzero. Then we have an expression of the form $0 = y'_i \neq 0$, which is nonsensical. Therefore, the only way we can solve this system is by requiring that $y'_{k+1}, \dots, y'_n = 0$. This is exactly the requirement that the rank of the augmented matrix equal k , the rank of the matrix \mathbf{A} by itself. If any of the y'_{k+1}, \dots, y'_n are nonzero, then the augmented matrix has one fewer row of zeros, so the rank of the augmented matrix would be greater than the rank of the original matrix. There are two ways to interpret this result. First, we require that the right hand side \mathbf{y} doesn't "mess up" our system by introducing a nonsensical expression. Second, if a row i in the matrix \mathbf{A} is linearly dependent on the other rows in \mathbf{A} , the corresponding values y_i must have the same dependency on the other values in \mathbf{y} . If so, when the row i is zeroed out during row reduction, the value y_i will also be zeroed out, avoiding any inconsistency.

□

4.8 Rank and Matrix Inverses

For a general nonhomogeneous system $\mathbf{Ax} = \mathbf{y}$ with $\mathbf{A} \in \mathbb{R}^{n \times n}$, we know that a unique solution only exists if $\text{rank}(\mathbf{A}) = \text{rank}([\mathbf{A} \ \mathbf{y}]) = n$. If $\text{rank}(\mathbf{A}) = n$, we know that \mathbf{A} can be transformed into reduced row form without generating any rows with all zero entries. We also know that if an inverse \mathbf{A}^{-1} exists for \mathbf{A} , we can use the inverse to uniquely solve for $\mathbf{x} = \mathbf{A}^{-1}\mathbf{y}$. Putting these facts together, we can now definitely state necessary and sufficient conditions for matrix inversion:

An $n \times n$ matrix \mathbf{A} has an inverse if and only if $\text{rank}(\mathbf{A}) = n$.

4.9 *Summary*

We've shown in this chapter the tight connections between matrix inversion, solvability, and the rank of a matrix. We will use these concepts many times to understand the solution of linear systems. However, we've also argued that despite their theoretical importance, these concepts have limited practical utility for solving linear systems. For example, computing the rank of a matrix requires transforming the matrix into reduced echelon form. This requires the same computations as solving a linear system involving the matrix, so one would rarely check the rank of a coefficient matrix before attempting to solve a linear system. Instead, we will see rank emerge as a useful tool only when considering matrices by themselves in Part III of this course.