

DAVID LERNER

LINEAR ALGEBRA

-LECTURE NOTES-

Table of contents

1. Matrices and matrix algebra	Page 1
2. Systems of linear equations	Page 8
3. Elementary row operations and the corresponding matrices	Page 14
4. Elementary matrices, continued	Page 19
5. The homogeneous system $A\mathbf{x} = \mathbf{0}$	Page 25
6. The inhomogeneous system $A\mathbf{x} = \mathbf{y}$, $\mathbf{y} \neq \mathbf{0}$	Page 31
7. Square matrices, inverses, and related matters	Page 37
8. Determinants	Page 42
9. The derivative	Page 47
10. Subspaces of $\mathbb{R}^n$	Page 51
11. Linear dependence and independence	Page 55
12. Basis and dimension	Page 59
13. The rank-nullity (dimension) theorem	Page 63
14. Change of basis	Page 67
15. Matrices and linear transformations	Page 74
16. Computations with eigenvalues and eigenvectors	Page 81
17. Inner products	Page 88
18. Orthogonality and related matters	Page 93
19. Orthogonal projections and Gram-Schmidt	Page 100
20. Approximations - the method of least squares (I)	Page 106

21. Least squares approximations (II)	Page 111
---	----------

1 Matrices and matrix algebra

1.1 Examples of matrices

A *matrix* is a rectangular array of numbers and/or variables. For instance

$$A = \begin{pmatrix} 4 & -2 & 0 & -3 & 1 \\ 5 & 1.2 & -0.7 & x & 3 \\ \pi & -3 & 4 & 6 & 27 \end{pmatrix}$$

is a matrix with 3 rows and 5 columns (a 3×5 matrix). The 15 *entries* of the matrix are referenced by the row and column in which they sit: the (2,3) entry of A is -0.7 . We may also write $a_{23} = -0.7$, $a_{24} = x$, etc. We indicate the fact that A is 3×5 (this is read as "three by five") by writing $A_{3 \times 5}$. Matrices can also be enclosed in square brackets as well as large parentheses. That is, both

$$\begin{pmatrix} 2 & 4 \\ 1 & -6 \end{pmatrix} \text{ and } \begin{bmatrix} 2 & 4 \\ 1 & -6 \end{bmatrix}$$

are perfectly good ways to write this 2×2 matrix.

Real numbers are 1×1 matrices. A vector such as

$$\mathbf{v} = \begin{pmatrix} x \\ y \\ z \end{pmatrix}$$

is a 3×1 matrix. We will generally use upper case Latin letters as symbols for matrices, boldface lower case letters for vectors, and ordinary lower case letters for real numbers.

Definition: Real numbers, when used in matrix computations, are called *scalars*.

Matrices are ubiquitous in mathematics and the sciences. Some instances include:

- Systems of linear algebraic equations (the main subject matter of this course) are normally written as simple matrix equations of the form $A\mathbf{x} = \mathbf{y}$.

- The *derivative* of a function $f : R^3 \rightarrow R^2$ is a 2×3 matrix.
- First order systems of linear differential equations are written in matrix form.
- The symmetry groups of mathematics and physics, which we'll look at later, are groups of matrices.
- Quantum mechanics can be formulated using infinite-dimensional matrices.

1.2 Operations with matrices

- *Addition*: matrices of the same size can be added or subtracted by adding or subtracting the corresponding entries:

$$\begin{pmatrix} 2 & 1 \\ -3 & 4 \\ 7 & 0 \end{pmatrix} + \begin{pmatrix} 6 & -1.2 \\ \pi & x \\ 1 & -1 \end{pmatrix} = \begin{pmatrix} 8 & -0.2 \\ \pi - 3 & 4 + x \\ 8 & -1 \end{pmatrix}.$$

Definition: If the matrices A and B have the same size, then their *sum* is the matrix $A + B$ defined by

$$(A + B)_{ij} = a_{ij} + b_{ij}.$$

Their *difference* is the matrix $A - B$ defined by

$$(A - B)_{ij} = a_{ij} - b_{ij}$$

.

- **Definition:** A matrix A can be multiplied by a scalar c to obtain the matrix cA , where

$$(cA)_{ij} = ca_{ij}.$$

This is called *scalar multiplication*. We just multiply each entry of A by c . For example

$$-3 \begin{pmatrix} 1 & 2 \\ 3 & 4 \end{pmatrix} = \begin{pmatrix} -3 & -6 \\ -9 & -12 \end{pmatrix}$$

- **Definition:** The $m \times n$ matrix whose entries are all 0 is denoted 0_{mn} (or, more often, just by 0 if the dimensions are obvious from context). It's called the *zero* matrix.
- **Definition:** Two matrices A and B are *equal* if all their corresponding entries are equal:

$$A = B \iff a_{ij} = b_{ij} \text{ for all } i, j.$$

- **Definition:** If the number of columns of A equals the number of rows of B , then the *product* AB is defined by

$$(AB)_{ij} = \sum_{s=1}^k a_{is}b_{sj}.$$

Here k is the number of columns of A or rows of B .

Example:

$$\begin{pmatrix} 1 & 2 & 3 \\ -1 & 0 & 4 \end{pmatrix} \begin{pmatrix} -1 & 0 \\ 4 & 2 \\ 1 & 3 \end{pmatrix} = \begin{pmatrix} 1 \cdot -1 + 2 \cdot 4 + 3 \cdot 1 & 1 \cdot 0 + 2 \cdot 2 + 3 \cdot 3 \\ -1 \cdot -1 + 0 \cdot 4 + 4 \cdot 1 & -1 \cdot 0 + 0 \cdot 2 + 4 \cdot 3 \end{pmatrix} = \begin{pmatrix} 10 & 13 \\ 5 & 12 \end{pmatrix}$$

If AB is defined, then the number of rows of AB is the same as the number of rows of A , and the number of columns is the same as the number of columns of B :

$$A_{m \times n} B_{n \times p} = (AB)_{m \times p}.$$

Why define multiplication like this? The answer is that this is the definition that corresponds to what shows up in practice.

Example: Recall from calculus (**Exercise!**) that if a point (x, y) in the plane is rotated counterclockwise about the origin through an angle θ to obtain a new point (x', y') , then

$$x' = x \cos \theta - y \sin \theta$$

$$y' = x \sin \theta + y \cos \theta.$$

In matrix notation, this can be written

$$\begin{pmatrix} x' \\ y' \end{pmatrix} = \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix}.$$

If the new point (x', y') is now rotated through an additional angle ϕ to get (x'', y'') , then

$$\begin{aligned} \begin{pmatrix} x'' \\ y'' \end{pmatrix} &= \begin{pmatrix} \cos \phi & -\sin \phi \\ \sin \phi & \cos \phi \end{pmatrix} \begin{pmatrix} x' \\ y' \end{pmatrix} \\ &= \begin{pmatrix} \cos \phi & -\sin \phi \\ \sin \phi & \cos \phi \end{pmatrix} \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} \\ &= \begin{pmatrix} \cos \theta \cos \phi - \sin \theta \sin \phi & -(\cos \theta \sin \phi + \sin \theta \cos \phi) \\ \cos \theta \sin \phi + \sin \theta \cos \phi & \cos \theta \cos \phi - \sin \theta \sin \phi \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} \\ &= \begin{pmatrix} \cos(\theta + \phi) & -\sin(\theta + \phi) \\ \sin(\theta + \phi) & \cos(\theta + \phi) \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} \end{aligned}$$

This is obviously the correct answer, since it shows that the point has been rotated through the total angle of $\theta + \phi$. The correct answer is given by matrix multiplication as we've defined it, and not some other way.

- **Matrix multiplication is not commutative:** in English, $AB \neq BA$, for arbitrary matrices A and B . For instance, if A is 3×5 and B is 5×2 , then AB is 3×2 , but BA is not defined. Even if both matrices are square and of the same size, so that both AB and BA are defined and have the same size, the two products are not generally equal.

Exercise: Write down two 2×2 matrices and compute both products. Unless you've been very selective, the two products won't be equal. Can you think of cases in which they *are* equal?

Another example: If

$$A = \begin{pmatrix} 2 \\ 3 \end{pmatrix}, \text{ and } B = \begin{pmatrix} 1 & 2 \end{pmatrix},$$

then

$$AB = \begin{pmatrix} 2 & 4 \\ 3 & 6 \end{pmatrix}, \text{ while } BA = (8).$$

- **Two properties of matrix multiplication:**
 1. If AB and AC are defined, then $A(B + C) = AB + AC$.

2. If AB is defined, and c is a scalar, then $A(cB) = c(AB)$.

(Although we won't do it here, both these properties can be proven by showing that, in each equation, the (i, j) entry on the right hand side of the equation is equal to the (i, j) entry on the left.)

- **Definition:** The *transpose* of the matrix A , denoted A^t , is obtained from A by making the first row of A into the first column of A^t , the second row of A into the second column of A^t , and so on. Formally,

$$a_{ij}^t = a_{ji}.$$

So

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

Here's a standard consequence of the non-commutativity of matrix multiplication: If AB is defined, then $(AB)^t = B^t A^t$ (not $A^t B^t$ as you might expect).

Example: If

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

then

$$AB = \begin{pmatrix} 2 & 7 \\ -3 & 6 \end{pmatrix}, \text{ so } (AB)^t = \begin{pmatrix} 2 & -3 \\ 7 & 6 \end{pmatrix}.$$

And

$$B^t A^t = \begin{pmatrix} -1 & 4 \\ 2 & 3 \end{pmatrix} \begin{pmatrix} 2 & 3 \\ 1 & 0 \end{pmatrix} = \begin{pmatrix} 2 & -3 \\ 7 & 6 \end{pmatrix}$$

as advertised.

- **Definition:** A is *square* if it has the same number of rows and columns. An important instance is the *identity matrix* I_n , which has ones on the main diagonal and zeros elsewhere:

Example:

$$I_3 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

Often, we'll just write I without the subscript for an identity matrix, when the dimension is clear from the context. The identity matrices behave, in some sense, like the number 1. If A is $n \times m$, then $I_n A = A$, and $A I_m = A$.

- **Definition:** Suppose A and B are square matrices of the same dimension, and suppose that $AB = I = BA$. Then B is said to be the *inverse* of A , and we write this as $B = A^{-1}$. Similarly, $B^{-1} = A$. For instance, you can easily check that

$$\begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix} \begin{pmatrix} 1 & -1 \\ -1 & 2 \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix},$$

and so these two matrices are inverses of one another:

$$\begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix}^{-1} = \begin{pmatrix} 1 & -1 \\ -1 & 2 \end{pmatrix} \text{ and } \begin{pmatrix} 1 & -1 \\ -1 & 2 \end{pmatrix}^{-1} = \begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix}.$$

Example: Not every square matrix has an inverse. For instance

$$A = \begin{pmatrix} 3 & 1 \\ 3 & 1 \end{pmatrix}$$

has no inverse.

Exercise: Show that the matrix A in the above example has no inverse. Hint: Suppose that

$$B = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$$

is the inverse of A . Then we must have $BA = I$. Write this out and show that the equations for the entries of B are inconsistent.

Exercise: Which 1×1 matrices are invertible, and what are their inverses?

Exercise: Show that if

$$A = \begin{pmatrix} a & b \\ c & d \end{pmatrix}, \text{ and } ad - bc \neq 0, \text{ then } A^{-1} = \frac{1}{ad - bc} \begin{pmatrix} d & -b \\ -c & a \end{pmatrix}.$$

If $ad - bc = 0$, then the matrix is not invertible. You should probably memorize this formula.

2 Matrices and systems of linear equations

You have all seen systems of linear equations such as

$$3x + 4y = 5 \tag{1}$$

$$2x - y = 0. \tag{2}$$

This system can easily be solved: just multiply the 2nd equation by 4, and add the two resulting equations to get $11x = 5$ or $x = 5/11$. Substituting this into either equation gives $y = 10/11$. In this case, a solution exists (obviously) and is *unique* (there's just one solution, namely $(5/11, 10/11)$).

We can write this system as a matrix equation, that is in the form $A\mathbf{x} = \mathbf{y}$.

$$\begin{pmatrix} 3 & 4 \\ 2 & -1 \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} 5 \\ 0 \end{pmatrix}. \tag{3}$$

Here

$$\mathbf{x} = \begin{pmatrix} x \\ y \end{pmatrix}, \text{ and } \mathbf{y} = \begin{pmatrix} 5 \\ 0 \end{pmatrix}.$$

This works because if we multiply the two matrices on the left, we get the 2×1 matrix equation

$$\begin{pmatrix} 3x + 4y \\ 2x - y \end{pmatrix} = \begin{pmatrix} 5 \\ 0 \end{pmatrix}.$$

And the two matrices are equal if both their entries are equal, which gives us the two equations in (1).

Of course, rewriting the system in matrix form does not, by itself, simplify the way in which we solve it. The simplification results from the following observation: the variables x and y can be eliminated from the computation by simply writing down a matrix in which the coefficients of x are in the first column, the coefficients of y in the second, and the right hand side of the system is the third column:

$$\begin{pmatrix} 3 & 4 & 5 \\ 2 & -1 & 0 \end{pmatrix}. \tag{4}$$

We are using the columns as "place markers" instead of x , y and the $=$ sign. That is, the first column consists of the coefficients of x , the second has the coefficients of y , and the third has the numbers on the right hand side of (1).

Definition: The matrix in (2) is called the *augmented matrix* of the system, and can be written in matrix shorthand as $(A|\mathbf{y})$.

We can do exactly the same operations on this matrix as we did on the original system¹:

$$\begin{pmatrix} 3 & 4 & 5 \\ 8 & -4 & 0 \end{pmatrix} : \text{Multiply the 2nd eqn by 4}$$

$$\begin{pmatrix} 3 & 4 & 5 \\ 11 & 0 & 5 \end{pmatrix} : \text{Add the 1st eqn to the 2nd}$$

$$\begin{pmatrix} 3 & 4 & 5 \\ 1 & 0 & \frac{5}{11} \end{pmatrix} : \text{Divide the 2nd eqn by 11}$$

The second equation now reads $1 \cdot x + 0 \cdot y = 5/11$, and we've solved for x ; we can now substitute for x in the first equation to solve for y as above.

Even though the solution to the system of equations is unique, it can be solved in many different ways (all of which, clearly, must give the same answer). Here are two other ways to solve it, both using the augmented matrix. As before, start with

$$\begin{pmatrix} 3 & 4 & 5 \\ 2 & -1 & 0 \end{pmatrix},$$

$$\begin{pmatrix} 1 & 5 & 5 \\ 2 & -1 & 0 \end{pmatrix} : \text{Replace eqn 1 with eqn 1 - eqn 2}$$

$$\begin{pmatrix} 1 & 5 & 5 \\ 0 & -11 & -10 \end{pmatrix} : \text{Subtract 2 times eqn 1 from eqn 2}$$

$$\begin{pmatrix} 1 & 5 & 5 \\ 0 & 1 & \frac{10}{11} \end{pmatrix} : \text{Divide eqn 2 by 11 to get } y = 10/11$$

¹The purpose of this lecture is to remind you of the mechanics for solving simple linear systems. We'll give precise definitions and statements of the algorithms later.

Now the second equation tells us that $y = 10/11$, and we can substitute this into the first equation $x + 5y = 5$ to get $x = 5/11$. We could even take this one step further:

$$\begin{pmatrix} 1 & 0 & \frac{5}{11} \\ 0 & 1 & \frac{10}{11} \end{pmatrix} : \text{ We added } -5 \cdot \text{eqn 2 to eqn 1}$$

Now the complete solution can just be read off from the matrix. What we've done is to eliminate x from the second equation, (the 0 in position (2,1)) and y from the first (the 0 in position (1,2)).

Exercise: What's wrong with writing the final matrix as

$$\begin{pmatrix} 1 & 0 & 0.45 \\ 0 & 1 & 0.91 \end{pmatrix}?$$

Exercise: (Do this BEFORE continuing with the text!) The system we just looked at consisted of two linear equations in two unknowns. Each equation, by itself, is the equation of a line in the plane and so has infinitely many solutions. To solve both equations simultaneously, we need to find the points, if any, which lie on *both* lines. There are 3 possibilities: (a) there's just one (the usual case), (b) there is no solution (if the two lines are parallel and distinct), or (c) there are an infinite number of solutions (if the two lines coincide).

Given all this food for thought, what are the possibilities for 2 equations in 3 unknowns? That is, what geometric object does each equation represent, and what are the possibilities for solution(s)?

Let's throw another variable into the mix and consider two equations in three unknowns:

$$\begin{aligned} 2x - 4y + z &= 1 \\ 4x + y - z &= 3 \end{aligned} \tag{5}$$

Rather than solving this directly, we'll work with the augmented matrix for the system which is

$$\begin{pmatrix} 2 & -4 & 1 & 1 \\ 4 & 1 & -1 & 3 \end{pmatrix}.$$

We proceed in more or less the same manner as above - that is, we try to eliminate x from the second equation, and y from the first by doing simple operations on the matrix. Before we start, observe that each time we do such an "operation", we are, in effect, replacing the original system of equations by an equivalent system which has the same solutions. For instance, if we multiply equation 1 by the number 2, we get a "new" equation 1 which has exactly the same solutions as the original. This is also true if we replace, say, equation 2 with equation 2 plus some multiple of equation 1. (Can you see why?)

So, to business:

$$\begin{pmatrix} 1 & -2 & \frac{1}{2} & \frac{1}{2} \\ 4 & 1 & -1 & 3 \end{pmatrix} : \text{Mult eqn 1 by } 1/2$$

$$\begin{pmatrix} 1 & -2 & \frac{1}{2} & \frac{1}{2} \\ 0 & 9 & -3 & 1 \end{pmatrix} : \text{Mult eqn 1 by } -4 \text{ and add it to eqn 2}$$

$$\begin{pmatrix} 1 & -2 & \frac{1}{2} & \frac{1}{2} \\ 0 & 1 & -\frac{1}{3} & \frac{1}{9} \end{pmatrix} : \text{Mult eqn 2 by } 1/9 \tag{6}$$

$$\begin{pmatrix} 1 & 0 & -\frac{1}{6} & \frac{13}{18} \\ 0 & 1 & -\frac{1}{3} & \frac{1}{9} \end{pmatrix} : \text{Add } -2 \cdot \text{eqn 2 to eqn 1} \tag{7}$$

The matrix (4) is called an *echelon form* of the augmented matrix. The matrix (5) is called the *reduced echelon form*. (Precise definitions of these terms will be given in the next lecture.) Either one can be used to solve the system of equations. Working with the echelon form in (4), the two equations now read

$$\begin{aligned} x - 2y + z/2 &= 1/2 \\ y - z/3 &= 1/9. \end{aligned}$$

So $y = z/3 + 1/9$. Substituting this into the first equation gives

$$\begin{aligned} x &= 2y - z/2 + 1/2 \\ &= 2(z/3 + 1/9) - z/2 + 1/2 \\ &= z/6 + 13/18 \end{aligned}$$

Exercise: Verify that the reduced echelon matrix (5) gives exactly the same solutions. This is as it should be. All "equivalent" systems of equations have the same solutions.

We see that for *any* choice of z , we get a solution to (3). If we take $z = 0$, then the solution is $x = 13/18$, $y = 1/9$. But if $z = 1$, then $x = 8/9$, $y = 4/9$ is the solution. Similarly for any other choice of z which for this reason is called a *free variable*. If we write $z = t$, a more familiar expression for the solution is

$$\begin{pmatrix} x \\ y \\ z \end{pmatrix} = \begin{pmatrix} \frac{t}{6} + \frac{13}{18} \\ \frac{t}{3} + \frac{1}{9} \\ t \end{pmatrix} = t \begin{pmatrix} \frac{1}{6} \\ \frac{1}{3} \\ 1 \end{pmatrix} + \begin{pmatrix} \frac{13}{18} \\ \frac{1}{9} \\ 0 \end{pmatrix}. \quad (8)$$

This is of the form $\mathbf{r}(t) = t\mathbf{v} + \mathbf{a}$, and you will recognize it as the (vector) parametric form of a line in $\mathbb{R}^3$. This (with t a free variable) is called the *general solution* to the system (3). If we choose a particular value of t , say $t = 3\pi$, and substitute into (6), then we have a *particular solution*.

Exercises: Write down the augmented matrix and solve these. If there are free variables, write your answer in the form given in (6) above. Also, give a geometric interpretation of the solution set (e.g., the common intersection of three planes in $\mathbb{R}^3$.)

1.

$$\begin{aligned} 3x + 2y - 4z &= 3 \\ -x - 2y + 3z &= 4 \end{aligned}$$

2.

$$\begin{aligned} 2x - 4y &= 3 \\ 3x + 2y &= -1 \\ x - y &= 10 \end{aligned}$$

3.

$$x + y + 3z = 4$$

It is now time to put on our mathematician's hats and think about what we've just been doing:

- Can we formalize the algorithm we've been using to solve these equations?
- Can we show that the algorithm always works? That is, are we guaranteed to get all the solutions if we use the algorithm?

To begin with, let's write down the different "operations" we've been using on the systems of equations and on the corresponding augmented matrices:

1. We can multiply any equation by a *non-zero* real number (scalar). The corresponding matrix operation consists of multiplying a row of the matrix by a scalar.
2. We can replace any equation by the original equation plus a scalar multiple of another equation. Equivalently, we can replace any row of a matrix by that row plus a multiple of another row.
3. We can interchange two equations (or two rows of the augmented matrix); we haven't needed to do this yet, but sometimes it's necessary, as we'll see in a bit.

Definition: These three operations are called *elementary row operations*.

In the next lecture, we'll assemble the solution algorithm, and show that it can be reformulated in terms of matrix multiplication.

3 Elementary row operations and their corresponding matrices

As we'll see shortly, each of the 3 elementary row operations can be performed by multiplying the augmented matrix $(A|\mathbf{y})$ on the *left* by what we'll call an *elementary matrix*. Just so this doesn't come as a total shock, let's look at some simple matrix operations:

- Suppose EA is defined, and suppose the first row of E is $(1, 0, 0, \dots, 0)$. Then the first row of EA is *identical* to the first row of A .
- Similarly, if the i^{th} row of E is all zeros except for a 1 in the i^{th} slot, then the i^{th} row of the product EA is identical to the i^{th} row of A .
- It follows that if we want to *change* only row i of the matrix A , we should multiply A on the left by some matrix E with the following property:

Every row *except* row i should be the i^{th} row of the corresponding identity matrix.

The procedure that we illustrate below can (and is) used to reduce *any* matrix to echelon form (not just augmented matrices).

Example: Let

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

1. To multiply the first row of A by $1/3$, we can multiply A on the left by the elementary matrix

$$E_1 = \begin{pmatrix} \frac{1}{3} & 0 \\ 0 & 1 \end{pmatrix}.$$

The result is

$$E_1 A = \begin{pmatrix} 1 & \frac{4}{3} & \frac{5}{3} \\ 2 & -1 & 0 \end{pmatrix}.$$

You should check this on your own. Same with the remaining computations.

2. To add $-2 \times \text{row 1}$ to row 2 in the resulting matrix, multiply it by

$$E_2 = \begin{pmatrix} 1 & 0 \\ -2 & 1 \end{pmatrix}$$

to obtain

$$E_2 E_1 A = \begin{pmatrix} 1 & \frac{4}{3} & \frac{5}{3} \\ 0 & -\frac{11}{3} & -\frac{10}{3} \end{pmatrix}.$$

3. Now multiply row 2 of $E_2 E_1 A$ by $-3/11$ using the matrix

$$E_3 = \begin{pmatrix} 1 & 0 \\ 0 & -\frac{3}{11} \end{pmatrix},$$

yielding

$$E_3 E_2 E_1 A = \begin{pmatrix} 1 & \frac{4}{3} & \frac{5}{3} \\ 0 & 1 & \frac{10}{11} \end{pmatrix}.$$

4. Finally, we clean out the second column by adding $(-4/3) \times \text{row 2}$ to row 1. We multiply by

$$E_4 = \begin{pmatrix} 1 & -\frac{4}{3} \\ 0 & 1 \end{pmatrix}$$

obtaining

$$E_4 E_3 E_2 E_1 A = \begin{pmatrix} 1 & 0 & \frac{5}{11} \\ 0 & 1 & \frac{10}{11} \end{pmatrix}.$$

Of course we get the same result as before, so why bother? The answer is that we're in the process of developing an algorithm that will work in the general case. So it's about time to formally identify our goal in the general case. We begin with some definitions.

Definition: The *leading entry* of a matrix row is the first non-zero entry in the row, starting from the left. A row without a leading entry is a row of zeros.

Definition: The matrix R is said to be in *echelon form* provided that

1. The leading entry of every non-zero row is a 1.

2. If the leading entry of row i is in position k , and the next row is not a row of zeros, then the leading entry of row $i + 1$ is in position $k + j$, where $j \geq 1$.
3. All zero rows are at the bottom of the matrix.

The following matrices are in echelon form:

$$\begin{pmatrix} 1 & * \\ 0 & 1 \end{pmatrix}, \begin{pmatrix} 1 & * & * \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{pmatrix}, \text{ and } \begin{pmatrix} 0 & 1 & * & * \\ 0 & 0 & 1 & * \\ 0 & 0 & 0 & 1 \end{pmatrix}.$$

Here the asterisks (*) stand for any number at all, including 0.

Definition: The matrix R is said to be in *reduced echelon form* if (a) R is in echelon form, and (b) each leading entry is the *only* non-zero entry in its column. The reduced echelon form of a matrix is also called the *Gauss-Jordan* form.

The following matrices are in reduced row echelon form:

$$\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \begin{pmatrix} 1 & * & 0 & * \\ 0 & 0 & 1 & * \\ 0 & 0 & 0 & 0 \end{pmatrix}, \text{ and } \begin{pmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}.$$

Exercise: Suppose A is 3×5 . What is the maximum number of leading 1's that can appear when it's been reduced to echelon form? Same questions for $A_{5 \times 3}$. Can you generalize your results to a statement for $A_{m \times n}$? (State it as a theorem.)

Once a matrix has been brought to echelon form, it can be put into reduced echelon form by cleaning out the non-zero entries in any column containing a leading 1. For example, if

$$R = \begin{pmatrix} 1 & 2 & -1 & 3 \\ 0 & 1 & 2 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix},$$

which is in echelon form, then it can be reduced to Gauss-Jordan form by adding (-2) row 2 to row 1, and then (-3) row 3 to row 1. Thus

$$\begin{pmatrix} 1 & -2 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 2 & -1 & 3 \\ 0 & 1 & 2 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} = \begin{pmatrix} 1 & 0 & -5 & 3 \\ 0 & 1 & 2 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}.$$

and

$$\begin{pmatrix} 1 & 0 & -3 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & -5 & 3 \\ 0 & 1 & 2 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} = \begin{pmatrix} 1 & 0 & -5 & 0 \\ 0 & 1 & 2 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}.$$

Note that column 3 cannot be "cleaned out" since there's no leading 1 there.

There is one more elementary row operation and corresponding elementary matrix we may need. Suppose we want to reduce the following matrix to Gauss-Jordan form

$$A = \begin{pmatrix} 2 & 2 & -1 \\ 0 & 0 & 3 \\ 1 & -1 & 2 \end{pmatrix}.$$

Multiplying row 1 by $1/2$, and then adding $-$ row 1 to row 3 leads to

$$E_2 E_1 A = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ -1 & 0 & 1 \end{pmatrix} \begin{pmatrix} \frac{1}{2} & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 2 & 2 & -1 \\ 0 & 0 & 3 \\ 1 & -1 & 2 \end{pmatrix} = \begin{pmatrix} 1 & 1 & \frac{-1}{2} \\ 0 & 0 & 3 \\ 0 & -2 & \frac{5}{2} \end{pmatrix}.$$

Now we can clearly do 2 more operations to get a leading 1 in the $(2,3)$ position, and another leading 1 in the $(3,2)$ position. But this won't be in echelon form (why not?) We need to interchange rows 2 and 3. This corresponds to changing the order of the equations, and evidently doesn't change the solutions. We can accomplish this by multiplying on the left with a matrix obtained from I by interchanging rows 2 and 3:

$$E_3 E_2 E_1 A = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix} \begin{pmatrix} 1 & 1 & \frac{-1}{2} \\ 0 & 0 & 3 \\ 0 & -2 & \frac{5}{2} \end{pmatrix} = \begin{pmatrix} 1 & 1 & \frac{-1}{2} \\ 0 & -2 & \frac{5}{2} \\ 0 & 0 & 3 \end{pmatrix}.$$

Exercise: Without doing any further computation, write down the Gauss-Jordan form for this matrix.

Exercise: Use elementary matrices to reduce

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

to Gauss-Jordan form. You should wind up with an expression of the form

$$E_k \cdots E_2 E_1 A = I.$$

What can you say about the matrix $B = E_k \cdots E_2 E_1$?

4 Elementary matrices, continued

We have identified 3 types of row operations and their corresponding elementary matrices. If you check the previous examples, you'll find that [these matrices are constructed by performing the given row operation on the identity matrix](#):

1. To multiply $\text{row}_j(A)$ by the scalar c use the matrix E obtained from I by multiplying j^{th} row of I by c .
2. To add $c\text{row}_j(A)$ to $\text{row}_k(A)$, use the identity matrix with its k^{th} row replaced by $(\dots, c, \dots, 1, \dots)$. Here c is in position j and the 1 is in position k . All other entries are 0
3. To interchange rows j and k , use the identity matrix with rows j and k interchanged.

4.1 Properties of elementary matrices

1. [Elementary matrices are always square](#). If the operation is to be performed on $A_{m \times n}$, then the elementary matrix E is $m \times m$. So the product EA has the same dimension as the original matrix A .
2. [Elementary matrices are invertible](#). If E is elementary, then E^{-1} is the matrix which undoes the operation that created E , and $E^{-1}EA = IA = A$; the matrix followed by its inverse does nothing to A : [Examples](#):

•

$$E = \begin{pmatrix} 1 & 0 \\ -2 & 1 \end{pmatrix}$$

adds $(-2)\text{row}_1(A)$ to $\text{row}_2(A)$. Its inverse is

$$E^{-1} = \begin{pmatrix} 1 & 0 \\ 2 & 1 \end{pmatrix},$$

which adds $(2)\text{row}_1(A)$ to $\text{row}_2(A)$.

- If E multiplies the second row by $\frac{1}{2}$, then

$$E^{-1} = \begin{pmatrix} 1 & 0 \\ 0 & 2 \end{pmatrix}.$$

- If E interchanges two rows, then $E = E^{-1}$.

Exercises:

1. If A is 3×4 , what is the elementary matrix that (a) subtracts $7\text{row}_3(A)$ from $\text{row}_2(A)$?, (b) interchanges the first and third rows? (c) multiplies $\text{row}_1(A)$ by 2?
2. What are the inverses of the matrices in exercise 1?
3. Do elementary matrices commute? That is, does it matter in which order they're multiplied? Give an example or two to illustrate your answer.

4.2 The algorithm for Gaussian elimination

We can now state the algorithm which will reduce any matrix first to row echelon form, and then, if needed to reduced echelon form:

1. Begin with the $(1, 1)$ entry. If it's some $a \neq 0$, divide through row 1 by a to get a 1 in the $(1, 1)$ position. If it *is* zero, then interchange row 1 with another row to get a nonzero $(1, 1)$ entry and proceed as above. If every entry in column 1 is zero, go to the top of column 2 and, by multiplication and permuting rows if necessary, get a 1 in the $(1, 2)$ slot. If column 2 won't work, then go to column 3, etc. If you can't arrange for a leading 1 somewhere in row 1, then your original matrix was the zero matrix, and it's already reduced.
2. You now have a leading 1 in some column. Use this leading 1 and operations of the type $(a)\text{row}_i(A) + \text{row}_k(A) \rightarrow \text{row}_k(A)$ to replace every entry in the column below the

location of the leading 1 by 0. In other words, the column will now look like

$$\begin{pmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{pmatrix}.$$

3. Now move one column to the right, and one row down and attempt to repeat the process, getting a leading 1 in this location. You may need to permute this row with a row *below* it. If it's not possible to get a non-zero entry in this position, move right one column and try again. At the end of this second procedure, your matrix might look like

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

where the second leading entry is in column 3. Notice that once a leading 1 has been installed in the $(1, 1)$ position, none of the subsequent row operations will change any of the elements in column 1. Similarly, for the matrix above, no subsequent row operations in our reduction process will change any of the entries in the first 3 columns.

4. The process continues until there are no more positions for leading entries – we either run out of rows or columns or both because the matrix has only a finite number of each. We have arrived at the row echelon form.

The three matrices below are all in row echelon form:

$$\begin{pmatrix} 1 & * & * & * & * \\ 0 & 0 & 1 & * & * \\ 0 & 0 & 0 & 1 & * \end{pmatrix}, \text{ or } \begin{pmatrix} 1 & * & * \\ 0 & 0 & 1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \text{ or } \begin{pmatrix} 1 & * & * \\ 0 & 1 & * \\ 0 & 0 & 1 \end{pmatrix} \quad (1)$$

Remark: The description of the algorithm doesn't involve elementary matrices. As a practical matter, it's much simpler to just do the row operation directly on A , instead of writing down

an elementary matrix and multiplying the matrices. But the fact that we *could* do this with the elementary matrices will turn out to be very useful theoretically.

Exercise: Find the echelon form for each of the following:

$$\begin{pmatrix} 1 & 2 \\ 3 & 4 \\ 5 & 6 \\ 7 & 8 \end{pmatrix}, \quad \begin{pmatrix} 0 & 4 \\ 7 & -2 \end{pmatrix}, \quad (3, 4), \quad \begin{pmatrix} 3 & 2 & -1 & 4 \\ 2 & -5 & 2 & 6 \end{pmatrix}$$

4.3 Observations

- The leading entries progress strictly downward, from left to right. We could just as easily have written an algorithm in which the leading entries progress downward as we move from right to left.
- The row echelon form of the matrix is *upper triangular*: any entry a_{ij} with $i > j$ satisfies $a_{ij} = 0$.
- To continue the reduction to Gauss-Jordan form, it is only necessary to use each leading 1 to clean out any remaining non-zero entries in its column. For the first matrix in (1) above, the Gauss-Jordan form will look like

$$\begin{pmatrix} 1 & * & 0 & 0 & * \\ 0 & 0 & 1 & 0 & * \\ 0 & 0 & 0 & 1 & * \end{pmatrix}$$

(Of course, cleaning out the columns may lead to changes in the entries labelled with *.)

4.4 Application to the solution(s) of $Ax = y$

Suppose that we have reduced the augmented matrix $(A:y)$ to either echelon or Gauss-Jordan form. Then

1. If there is a leading 1 anywhere in the last column, the system $A\mathbf{x} = \mathbf{y}$ is *inconsistent*. That is, there is no $\mathbf{x}$ which satisfies the system of equations. Why?
2. If there's no leading entry in the last column, then at least one solution exists. The question then becomes "How many solutions are there?" The answer to this question depends on the number of free variables:

Definition: Suppose the augmented matrix for the linear system $A\mathbf{x} = \mathbf{y}$ has been brought to echelon form. If there is a leading 1 in any column except the last, then the corresponding variable is called a *leading variable*. For instance, if there's a leading 1 in column 3, then x_3 is a leading variable.

Definition: Any variable which is not a leading variable is a *free variable*.

Example: Suppose the echelon form of $(A;\mathbf{y})$ is

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

Then the original matrix A is 2×3 , and if x_1, x_2 , and x_3 are the variables in the original equations, we see that x_1 and x_3 are leading variables, and x_2 is a free variable.

- If the system is consistent and there are no free variables, then the solution is unique — there's just one. Here's an example of this:

$$\begin{pmatrix} 1 & * & * & * \\ 0 & 1 & * & * \\ 0 & 0 & 1 & * \\ 0 & 0 & 0 & 0 \end{pmatrix}$$

- If the system is consistent and there are one or more free variables, then there are infinitely many solutions.

$$\begin{pmatrix} 1 & * & * & * \\ 0 & 0 & 1 & * \\ 0 & 0 & 0 & 0 \end{pmatrix}$$

Here x_2 is a free variable, and we get a different solution for each of the infinite number of ways we could choose x_2 .

- Just because there are free variables does not mean that the system is consistent.

$$\begin{pmatrix} 1 & * & * \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{pmatrix}$$

Here x_2 is a free variable, but the system is inconsistent because of the leading 1 in the last column. There are no solutions to this system.

5 Homogeneous systems

Definition A *homogeneous* (ho-mo-geen'-ius) system of linear algebraic equations is one in which all the numbers on the right hand side are equal to 0:

$$\begin{aligned}a_{11}x_1 + \dots + a_{1n}x_n &= 0 \\ &\vdots \\ a_{m1}x_1 + \dots + a_{mn}x_n &= 0\end{aligned}$$

In matrix form, this reads $A\mathbf{x} = \mathbf{0}$, where A is $m \times n$,

$$\mathbf{x} = \begin{pmatrix} x_1 \\ \vdots \\ x_n \end{pmatrix}_{n \times 1},$$

and $\mathbf{0}$ is $n \times 1$. The homogenous system $A\mathbf{x} = \mathbf{0}$ always has the solution $\mathbf{x} = \mathbf{0}$. It follows that any homogeneous system of equations is *always consistent*. Any non-zero solutions, if they exist, are said to be *non-trivial* solutions. These may or may not exist. We can find out by row reducing the corresponding augmented matrix $(A:\mathbf{0})$.

Example: Given the augmented matrix

$$(A:\mathbf{0}) = \begin{pmatrix} 1 & 2 & 0 & -1 & 0 \\ -2 & -3 & 4 & 5 & 0 \\ 2 & 4 & 0 & -2 & 0 \end{pmatrix},$$

row reduction leads quickly to the echelon form

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

Observe that nothing happened to the last column — row operations don't do anything to a column of zeros. In particular, doing a row operation on a system of homogeneous equations doesn't change the fact that it's homogeneous. For this reason, when working with homogeneous systems, we'll just use the matrix A . The echelon form of A is

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

Here, the leading variables are x_1 and x_2 , while x_3 and x_4 are free variables, since there are no leading entries in the third or fourth columns. Continuing along, we obtain the Gauss-Jordan form (You *are* working out all the details on your scratch paper as we go along, aren't you!?)

$$\begin{pmatrix} 1 & 0 & -8 & -7 \\ 0 & 1 & 4 & 3 \\ 0 & 0 & 0 & 0 \end{pmatrix}.$$

No further simplification is possible; any further row operations will destroy the Gauss-Jordan structure of the columns with leading entries. The resulting system of equations reads

$$\begin{aligned} x_1 - 8x_3 - 7x_4 &= 0 \\ x_2 + 4x_3 + 3x_4 &= 0, \end{aligned}$$

In principle, we're done in the sense that we have the solution in hand. However, it's customary to rewrite the solution in vector form so that its properties are more clearly displayed. First, we solve for the leading variables; everything else goes on the right hand side of the equations:

$$\begin{aligned} x_1 &= 8x_3 + 7x_4 \\ x_2 &= -4x_3 - 3x_4. \end{aligned}$$

Assigning any values we choose to the two free variables x_3 and x_4 gives us a solution to the original homogeneous system. This is, of course, where the variables are called "free". We can distinguish the free variables from the leading variables by denoting them as s , t , u , etc.

Thus, setting $x_3 = s$, $x_4 = t$, we rewrite the solution in the form

$$\begin{aligned}x_1 &= 8s + 7t \\x_2 &= -4s - 3t \\x_3 &= s \\x_4 &= t\end{aligned}$$

Better yet, the solution can also be written in matrix (vector) form as

$$\mathbf{x} = \begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{pmatrix} = s \begin{pmatrix} 8 \\ -4 \\ 1 \\ 0 \end{pmatrix} + t \begin{pmatrix} 7 \\ -3 \\ 0 \\ 1 \end{pmatrix} \quad (1)$$

We call (1) the *general solution* to the homogeneous equation. The notation is misleading, since the left hand side $\mathbf{x}$ looks like a single vector, while the right hand side clearly represents an infinite collection of objects with 2 degrees of freedom. We'll address this later in the lecture.

We won't do it here, but If we were to carry out the above procedure on a general homogeneous system $A_{m \times n} \mathbf{x} = \mathbf{0}$, we'd establish the following facts:

5.1 Properties of the homogenous system for A_{mn}

- The number of leading variables is $\leq \min(m, n)$.
- The number of non-zero equations in the echelon form of the system is equal to the number of leading entries.
- The number of free variables plus the number of leading variables $= n$, the number of columns of A .
- The homogenous system $A\mathbf{x} = \mathbf{0}$ has *non-trivial* solutions if and only if there are free variables.

- If there are more unknowns than equations, the homogeneous system *always* has non-trivial solutions. Why? This is one of the few cases in which we can tell something about the solutions without doing any work.
- A homogeneous system of equations is always consistent (i.e., always has at least one solution).

Exercise: What sort of geometric object does $\mathbf{x}_H$ represent?

There are two other fundamental properties:

1. **Theorem:** If $\mathbf{x}$ is a solution to $A\mathbf{x} = \mathbf{0}$, then so is $c\mathbf{x}$ for any real number c .

Proof: $\mathbf{x}$ is a solution means $A\mathbf{x} = \mathbf{0}$. But $A(c\mathbf{x}) = cA\mathbf{x} = c\mathbf{0} = \mathbf{0}$, so $c\mathbf{x}$ is also a solution.

2. **Theorem:** If $\mathbf{x}$ and $\mathbf{y}$ are two solutions to the homogeneous equation, then so is $\mathbf{x} + \mathbf{y}$.

Proof: $A(\mathbf{x} + \mathbf{y}) = A\mathbf{x} + A\mathbf{y} = \mathbf{0} + \mathbf{0} = \mathbf{0}$.

These two properties constitute the famous *principle of superposition* which holds for homogeneous systems (but NOT for inhomogeneous ones).

Definition: if $\mathbf{x}$ and $\mathbf{y}$ are two vectors and s and t two scalars, then $s\mathbf{x} + t\mathbf{y}$ is called a *linear combination* of $\mathbf{x}$ and $\mathbf{y}$.

Example: $3\mathbf{x} - 4\pi\mathbf{y}$ is a linear combination of $\mathbf{x}$ and $\mathbf{y}$.

We can restate the superposition principle as:

Superposition principle: if $\mathbf{x}$ and $\mathbf{y}$ are two solutions to the homogenous equation $A\mathbf{x} = \mathbf{0}$, then any linear combination of $\mathbf{x}$ and $\mathbf{y}$ is also a solution.

Remark: This is just a compact way of restating the two properties: If $\mathbf{x}$ and $\mathbf{y}$ are solutions, then by property 1, $s\mathbf{x}$ and $t\mathbf{y}$ are also solutions. And by property 2, their sum $s\mathbf{x} + t\mathbf{y}$ is a

solution. Conversely, if $s\mathbf{x} + t\mathbf{y}$ is a solution to the homogeneous equation for all s, t , then taking $t = 0$ gives property 1, and taking $s = t = 1$ gives property 2.

You have seen this principle at work in your calculus courses.

Example: Suppose $\phi(x, y)$ satisfies Laplace's equation

$$\frac{\partial^2 \phi}{\partial x^2} + \frac{\partial^2 \phi}{\partial y^2} = 0.$$

We write this as

$$\Delta \phi = 0, \quad \text{where } \Delta = \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2}.$$

The *differential operator* Δ has the same property as matrix multiplication, namely: if $\phi(x, y)$ and $\psi(x, y)$ are two differentiable functions, and s and t any two real numbers, then

$$\Delta(s\phi + t\psi) = s\Delta\phi + t\Delta\psi.$$

It follows that if ϕ and ψ are two solutions to Laplace's equation, then any linear combination of ϕ and ψ is also a solution. The principle of superposition also holds for solutions to the wave equation, Maxwell's equations in free space, and Schrödinger's equation in quantum mechanics.

Example: Start with "white" light (e.g., sunlight); it's a collection of electromagnetic waves which satisfy Maxwell's equations. Pass the light through a prism, obtaining red, orange, ..., violet light; these are also solutions to Maxwell's equations. The original solution (white light) is seen to be a superposition of many other solution, corresponding to the various different colors. The process can be reversed to obtain white light again by passing the different colors of the spectrum through an inverted prism.

Referring back to the example (see Eqn (1)), if we set

$$\mathbf{x} = \begin{pmatrix} 8 \\ -4 \\ 1 \\ 0 \end{pmatrix}, \quad \text{and } \mathbf{y} = \begin{pmatrix} 7 \\ -3 \\ 0 \\ 1 \end{pmatrix},$$

then the susperposition principle tells us that any linear combination of $\mathbf{x}$ and $\mathbf{y}$ is also a solution. In fact, these are *all* of the solutions to this system.

Definition: We write

$$\mathbf{x}_H = \{s\mathbf{x} + t\mathbf{y} : \forall \text{ real } s, t\}$$

and say that $\mathbf{x}_H$ is the *general solution* to the homogeneous system $A\mathbf{x} = \mathbf{0}$.

6 The Inhomogeneous system $A\mathbf{x} = \mathbf{y}$, $\mathbf{y} \neq \mathbf{0}$

Definition: The system $A\mathbf{x} = \mathbf{y}$ is *inhomogeneous* if it's not homogeneous.

Mathematicians love definitions like this! It means of course that the vector $\mathbf{y}$ is not the zero vector. And this means that at least one of the equations has a non-zero right hand side.

As an example, we can use the same system as in the previous lecture, except we'll change the right hand side to something non-zero:

$$\begin{aligned}x_1 + 2x_2 - x_4 &= 1 \\-2x_1 - 3x_2 + 4x_3 + 5x_4 &= 2 \\2x_1 + 4x_2 - 2x_4 &= 3\end{aligned}$$

Those of you with sharp eyes should be able to tell at a glance that this system is *inconsistent* — that is, there are *no* solutions. Why? We're going to proceed anyway because this is hardly an exceptional situation.

The augmented matrix is

$$(A:\mathbf{y}) = \begin{pmatrix} 1 & 2 & 0 & -1 & 1 \\ -2 & -3 & 4 & 5 & 2 \\ 2 & 4 & 0 & -2 & 3 \end{pmatrix}.$$

We can't discard the 5th column here since it's not zero. The row echelon form of the augmented matrix is

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

And the reduced echelon form is

$$\begin{pmatrix} 1 & 0 & -8 & -7 & 0 \\ 0 & 1 & 4 & 3 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix}.$$

The third equation, from either of these, now reads

$$0x_1 + 0x_2 + 0x_3 + 0x_4 = 1, \text{ or } 0 = 1.$$

This is false! How can we wind up with a false statement? The actual reasoning that led us here is this: *If* the original system has a solution, then performing elementary row operations gives us an equivalent system of equations which has the same solution. But this equivalent system of equations is *inconsistent*. It has no solutions; that is *no* choice of $x_1, \dots, x_4$ satisfies the equation. So the original system is also inconsistent.

In general: *If the echelon form of $(A;\mathbf{y})$ has a leading 1 in any position of the last column, the system of equations is inconsistent.*

Now it's not true that any inhomogenous system with the same matrix A is inconsistent. It depends completely on the particular $\mathbf{y}$ which sits on the right hand side. For instance, if

$$\mathbf{y} = \begin{pmatrix} 1 \\ 2 \\ 2 \end{pmatrix},$$

then (work this out!) the echelon form of $(A;\mathbf{y})$ is

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

and the reduced echelon form is

$$\begin{pmatrix} 1 & 0 & -8 & -7 & -7 \\ 0 & 1 & 4 & 3 & 4 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix}.$$

Since this is consistent, we have, as in the homogeneous case, the leading variables x_1 and x_2 , and the free variables x_3 and x_4 . Renaming the free variables by s and t , and writing out the equations solved for the leading variables gives us

$$\begin{aligned}x_1 &= 8s + 7t - 7 \\x_2 &= -4s - 3t + 4 \\x_3 &= s \\x_4 &= t\end{aligned}.$$

This looks like the solution to the homogeneous equation found in the previous section except for the additional scalars -7 and $+4$ in the first two equations. If we rewrite this using vector notation, we get

$$\mathbf{x} = \begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{pmatrix} = s \begin{pmatrix} 8 \\ -4 \\ 1 \\ 0 \end{pmatrix} + t \begin{pmatrix} 7 \\ -3 \\ 0 \\ 1 \end{pmatrix} + \begin{pmatrix} -7 \\ 4 \\ 0 \\ 0 \end{pmatrix}$$

Compare this with the general solution $\mathbf{x}_H$ to the homogenous equation found before. Once again, we have a 2-parameter family of solutions. We can get what is called a *particular* solution by making some specific choice for s and t . For example, taking $s = t = 0$, we get the particular solution

$$\mathbf{x}_p = \begin{pmatrix} -7 \\ 4 \\ 0 \\ 0 \end{pmatrix}.$$

We can get other particular solutions by making other choices. Observe that the *general solution to the inhomogeneous system* worked out here can be written in the form $\mathbf{x} = \mathbf{x}_H + \mathbf{x}_p$. In fact, this is true in general:

Theorem: Let $\mathbf{x}_p$ and $\mathbf{y}_p$ be two solutions to $A\mathbf{x} = \mathbf{y}$. Then their difference $\mathbf{x}_p - \mathbf{y}_p$ is a

solution to the homogeneous equation $A\mathbf{x} = \mathbf{0}$. The *general solution* to $A\mathbf{x} = \mathbf{y}$ can be written as $\mathbf{x}_p + \mathbf{x}_h$ where $\mathbf{x}_h$ denotes the general solution to the homogeneous system.

Proof: Since $\mathbf{x}_p$ and $\mathbf{y}_p$ are solutions, we have $A(\mathbf{x}_p - \mathbf{y}_p) = A\mathbf{x}_p - A\mathbf{y}_p = \mathbf{y} - \mathbf{y} = \mathbf{0}$. So their difference solves the homogeneous equation. Conversely, given a particular solution $\mathbf{x}_p$, then the entire set $\mathbf{x}_p + \mathbf{x}_h$ consists of solutions to $A\mathbf{x} = \mathbf{y}$: if $\mathbf{z}$ belongs to $\mathbf{x}_h$, then $A(\mathbf{x}_p + \mathbf{z}) = A\mathbf{x}_p + A\mathbf{z} = \mathbf{y} + \mathbf{0} = \mathbf{y}$ and so $\mathbf{x}_p + \mathbf{z}$ is a solution to $A\mathbf{x} = \mathbf{y}$.

Going back to the example, suppose we write the general solution to $A\mathbf{x} = \mathbf{y}$ in the vector form

$$\mathbf{x} = s\mathbf{v}_1 + t\mathbf{v}_2 + \mathbf{x}_p,$$

where

$$\mathbf{v}_1 = \begin{pmatrix} 8 \\ -4 \\ 1 \\ 0 \end{pmatrix}, \quad \mathbf{v}_2 = \begin{pmatrix} 7 \\ -3 \\ 0 \\ 1 \end{pmatrix}, \quad \text{and } \mathbf{x}_p = \begin{pmatrix} -7 \\ 4 \\ 0 \\ 0 \end{pmatrix}$$

Now we can get another particular solution to the system by taking $s = 1, t = 1$. This gives

$$\mathbf{y}_p = \begin{pmatrix} 8 \\ -3 \\ 1 \\ 1 \end{pmatrix}.$$

We can rewrite the general solution as

$$\begin{aligned} \mathbf{x} &= (s - 1 + 1)\mathbf{v}_1 + (t - 1 + 1)\mathbf{v}_2 + \mathbf{x}_p \\ &= (s - 1)\mathbf{v}_1 + (t - 1)\mathbf{v}_2 + \mathbf{y}_p \\ &= \hat{s}\mathbf{v}_1 + \hat{t}\mathbf{v}_2 + \mathbf{y}_p \end{aligned}.$$

As $\hat{s}$ and $\hat{t}$ run over all possible pairs of real numbers we get exactly the same set of solutions as before. So the general solution can be written as $\mathbf{y}_p + \mathbf{x}_h$ as well as $\mathbf{x}_p + \mathbf{x}_h$! This is a bit confusing until you remember that these are *sets* of solutions, rather than single solutions; $(\hat{s}, \hat{t})$ and (s, t) are just different sets of coordinates. But running through either set of coordinates (or parameters) produces the same set.

Remarks

- Those of you taking a course in differential equations will encounter a similar situation: the general solution to a linear differential equation has the form $y = y_p + y_h$, where y_p is any particular solution to the DE, and y_h denotes the set of all solutions to the homogeneous DE.




Figure 1: The lower plane (the one passing through $\mathbf{0}$) represents $\mathbf{x}_H$. Given the particular solution $\mathbf{x}_p$ and a $\mathbf{z}$ in $\mathbf{x}_H$, we get another solution to the inhomogeneous equation. As $\mathbf{z}$ varies in $\mathbf{x}_H$, we get all the solutions to $A\mathbf{x} = \mathbf{y}$.

- We can visualize the general solutions to the homogeneous and inhomogeneous equations we've worked out in detail as follows. The set $\mathbf{x}_H$ is a 2-plane in $\mathbb{R}^4$ which goes through the origin since $\mathbf{x} = \mathbf{0}$ is a solution. The general solution to $A\mathbf{x} = \mathbf{y}$ is obtained by adding the vector $\mathbf{x}_p$ to every point in this 2-plane. Geometrically, this gives another 2-plane parallel to the first, but *not* containing the origin (since $\mathbf{x} = \mathbf{0}$ is not

a solution to $A\mathbf{x} = \mathbf{y}$ unless $\mathbf{y} = \mathbf{0}$). Now pick *any* point in this parallel 2-plane and add to it all the vectors in the 2-plane corresponding to $\mathbf{x}_h$. What do you get? You get the same parallel 2-plane! This is why $\mathbf{x}_p + \mathbf{x}_h = \mathbf{y}_p + \mathbf{x}_h$.

7 Square matrices, inverses and related matters

Square matrices are the only matrices that can have inverses, and for this reason, they are a bit special.

In a system of linear algebraic equations, if the number of equations equals the number of unknowns, then the associated coefficient matrix A is square. Suppose we row reduce A to its Gauss-Jordan form. There are two possible outcomes:

1. The Gauss-Jordan form for $A_{n \times n}$ is the $n \times n$ identity matrix I_n (commonly written as just I).
2. The Gauss-Jordan form for A has at least one row of zeros.

The second case is clear: The GJ form of $A_{n \times n}$ can have at most n leading entries. If the GJ form of A is not I , then the GJ form has $n - 1$ or fewer leading entries, and therefore has a row of zeros.

In the first case, we can show that A is invertible. To see this, remember that A is reduced to GJ form by multiplication on the left by a finite number of elementary matrices. If the GJ form is I , then we have an expression like

$$E_k E_{k-1} \dots E_2 E_1 A = I,$$

where E_i is the matrix corresponding to the i^{th} row operation used in the reduction. If we set $B = E_k E_{k-1} \dots E_2 E_1$, then clearly $BA = I$ and so $B = A^{-1}$. Furthermore, multiplying BA on the left by (note the order!!!) E_k^{-1} , then by E_{k-1}^{-1} , and continuing to E_1^{-1} , we *undo* all the row operations that brought A to GJ form, and we get back A :

$$\begin{aligned} (E_1^{-1} E_2^{-1} \dots E_{k-1}^{-1} E_k^{-1}) BA &= (E_1^{-1} E_2^{-1} \dots E_{k-1}^{-1} E_k^{-1}) I \text{ or} \\ (E_1^{-1} E_2^{-1} \dots E_{k-1}^{-1} E_k^{-1}) (E_k E_{k-1} \dots E_2 E_1) A &= (E_1^{-1} E_2^{-1} \dots E_{k-1}^{-1} E_k^{-1}) \\ IA &= (E_1^{-1} E_2^{-1} \dots E_{k-1}^{-1} E_k^{-1}) \\ A &= (E_1^{-1} E_2^{-1} \dots E_{k-1}^{-1} E_k^{-1}) \end{aligned}$$

We summarize this in a

Theorem: The following are equivalent (i.e., each of the statements below implies and is implied by any of the others)

- The square matrix A is invertible.
- The Gauss-Jordan or reduced echelon form of A is the identity matrix.
- A can be written as a product of elementary matrices

Example - (fill in the details on your scratch paper)

We start with

$$A = \begin{pmatrix} 2 & 1 \\ 1 & 2 \end{pmatrix}.$$

We multiply row 1 by $1/2$ using the matrix E_1 :

$$E_1 A = \begin{pmatrix} \frac{1}{2} & 0 \\ 0 & 1 \end{pmatrix} A = \begin{pmatrix} 1 & \frac{1}{2} \\ 1 & 2 \end{pmatrix}.$$

We now add $-(\text{row } 1)$ to row 2, using E_2 :

$$E_2 E_1 A = \begin{pmatrix} 1 & 0 \\ -1 & 1 \end{pmatrix} \begin{pmatrix} 1 & \frac{1}{2} \\ 1 & 2 \end{pmatrix} = \begin{pmatrix} 1 & \frac{1}{2} \\ 0 & \frac{3}{2} \end{pmatrix}.$$

Now multiply the second row by $\frac{2}{3}$:

$$E_3 E_2 E_1 A = \begin{pmatrix} 1 & 0 \\ 0 & \frac{2}{3} \end{pmatrix} \begin{pmatrix} 1 & \frac{1}{2} \\ 0 & \frac{3}{2} \end{pmatrix} = \begin{pmatrix} 1 & \frac{1}{2} \\ 0 & 1 \end{pmatrix}.$$

And finally, add $-\frac{1}{2}(\text{row } 2)$ to row 1:

$$E_4 E_3 E_2 E_1 A = \begin{pmatrix} 1 & -\frac{1}{2} \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & \frac{1}{2} \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}.$$

So

$$A^{-1} = E_4 E_3 E_2 E_1 = \frac{1}{3} \begin{pmatrix} 2 & -1 \\ -1 & 2 \end{pmatrix}.$$

Exercises:

- *Check* the last expression by multiplying the elementary matrices.
- Write A as the product of elementary matrices.
- The individual factors in the product of A^{-1} are not unique. They depend on how we do the row reduction. Find another factorization of A^{-1} . (Hint: Start things out a different way, for example by adding $-(\text{row } 2)$ to row 1.)

- Let

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

Express both A and A^{-1} as the products of elementary matrices.

7.1 Solutions to $A\mathbf{x} = \mathbf{y}$ when A is square

- If A is invertible, then the equation $A\mathbf{x} = \mathbf{y}$ has the *unique* solution $A^{-1}\mathbf{y}$ for any right hand side $\mathbf{y}$. For,

$$A\mathbf{x} = \mathbf{y} \iff A^{-1}A\mathbf{x} = A^{-1}\mathbf{y} \iff \mathbf{x} = A^{-1}\mathbf{y}.$$

In this case, the solution to the homogeneous equation is also unique - it's the trivial solution.

- If A is not invertible, then there is at least one free variable. So there are non-trivial solutions to $A\mathbf{x} = \mathbf{0}$. If $\mathbf{y} \neq \mathbf{0}$, then either $A\mathbf{x} = \mathbf{y}$ is inconsistent (the most likely case) or solutions to the system exist, but there are infinitely many.

Exercise: If the square matrix A is not invertible, why is it “likely” that the inhomogeneous equation is inconsistent? “Likely”, in this case, means that the system should be inconsistent for a $\mathbf{y}$ chosen at random.

7.2 An algorithm for constructing A^{-1}

The work we've just done leads us immediately to an algorithm for constructing the inverse of A . (You've probably seen this before, but now you'll know why it works!). It's based on the following observation: suppose $B_{n \times p}$ is another matrix with the same number of rows as $A_{n \times n}$, and $E_{n \times n}$ is an elementary matrix which can multiply A on the left. Then E can also multiply B on the left, and if we form the *partitioned* matrix

$$C = (A:B)_{n \times n+p},$$

Then, in what should be an obvious notation, we have

$$EC = (EA:EB)_{n \times n+p},$$

where EA is $n \times n$ and EB is $n \times p$. (**Exercise:** Check this for yourself with a simple example. Better yet, prove it in general.)

The algorithm consists of forming the partitioned matrix $C = (A:I)$, and doing the row operations that reduce A to Gauss-Jordan form on the larger matrix C . If A is invertible, we'll end up with

$$\begin{aligned} E_k \dots E_1(A:I) &= (E_k \dots E_1 A : E_k \dots E_1 I) \\ &= (I : A^{-1}) \end{aligned}.$$

In words: the same sequence of row operations that reduces A to I will convert I to A^{-1} . The advantage to doing things this way is that **you don't have to write down the elementary matrices**. They're working away in the background, as we know from the theory, but if all we want is A^{-1} , then we don't need them explicitly; we just do the row operations.

Example:

$$\text{Let } A = \begin{pmatrix} 1 & 2 & 3 \\ 1 & 0 & -1 \\ 2 & 3 & 1 \end{pmatrix}.$$

Then row reducing $(A:I)$, we get

$$(A:I) = \begin{pmatrix} 1 & 2 & 3 & 1 & 0 & 0 \\ 1 & 0 & -1 & 0 & 1 & 0 \\ 2 & 3 & 1 & 0 & 0 & 1 \end{pmatrix}$$

$$r1 \leftrightarrow r2 \quad \begin{pmatrix} 1 & 0 & -1 & 0 & 1 & 0 \\ 1 & 2 & 3 & 1 & 0 & 0 \\ 2 & 3 & 1 & 0 & 0 & 1 \end{pmatrix}$$

$$\text{do col 1} \quad \begin{pmatrix} 1 & 0 & -1 & 0 & 1 & 0 \\ 0 & 2 & 4 & 1 & -1 & 0 \\ 0 & 3 & 3 & 0 & -2 & 1 \end{pmatrix}$$

$$\text{do column 2} \quad \begin{pmatrix} 1 & 0 & -1 & 0 & 1 & 0 \\ 0 & 1 & 2 & \frac{1}{2} & -\frac{1}{2} & 0 \\ 0 & 0 & -3 & -\frac{3}{2} & -\frac{1}{2} & 1 \end{pmatrix}$$

$$\text{and column 3} \quad \begin{pmatrix} 1 & 0 & 0 & \frac{1}{2} & \frac{7}{6} & -\frac{1}{3} \\ 0 & 1 & 0 & -\frac{1}{2} & -\frac{5}{6} & \frac{2}{3} \\ 0 & 0 & 1 & \frac{1}{2} & \frac{1}{6} & -\frac{1}{3} \end{pmatrix}$$

So,

$$A^{-1} = \begin{pmatrix} \frac{1}{2} & \frac{7}{6} & -\frac{1}{3} \\ -\frac{1}{2} & -\frac{5}{6} & \frac{2}{3} \\ \frac{1}{2} & \frac{1}{6} & -\frac{1}{3} \end{pmatrix}.$$

Exercise: Write down a 2×2 matrix and do this yourself. Same with a 3×3 matrix.

8 Square matrices continued: Determinants

8.1 Introduction

Determinants give us important information about square matrices, and, as we'll see in the next lecture, are essential for the computation of eigenvalues. You have seen determinants in your precalculus courses. For a 2×2 matrix

$$A = \begin{pmatrix} a & b \\ c & d \end{pmatrix},$$

the formula reads

$$\det(A) = ad - bc.$$

For a 3×3 matrix

$$\begin{pmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{pmatrix},$$

life is more complicated. Here the formula reads

$$\det(A) = a_{11}a_{22}a_{33} + a_{13}a_{21}a_{32} + a_{12}a_{23}a_{31} - a_{12}a_{21}a_{33} - a_{11}a_{23}a_{32} - a_{13}a_{22}a_{31}.$$

Things get worse quickly as the dimension increases. For an $n \times n$ matrix A , the expression for $\det(A)$ has n factorial $= n! = 1 \cdot 2 \cdot \dots \cdot (n-1) \cdot n$ terms, each of which is a product of n matrix entries. Even on a computer, calculating the determinant of a 10×10 matrix using this sort of formula would be unnecessarily time-consuming, and doing a 1000×1000 matrix would take years!

Fortunately, as we'll see below, computing the determinant is easy if the matrix happens to be in echelon form. You just need to do a little bookkeeping on the side as you reduce the matrix to echelon form.

8.2 The definition of $\det(A)$

Let A be $n \times n$, and write $\mathbf{r}_1$ for the first row, $\mathbf{r}_2$ for the second row, etc.

The **determinant of A** is a real-valued function of the rows of A which we write as

$$\det(A) = \det(\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_n).$$

It is completely determined by the following four properties:

1. Multiplying a row by the constant c multiplies the determinant by c :

$$\det(\mathbf{r}_1, \mathbf{r}_2, \dots, c\mathbf{r}_i, \dots, \mathbf{r}_n) = c \det(\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_i, \dots, \mathbf{r}_n)$$

2. If row i is the sum of $\mathbf{r}_i$ and $\mathbf{y}_i$, then the determinant is the sum of the two corresponding determinants:

$$\det(\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_i + \mathbf{y}_i, \dots, \mathbf{r}_n) = \det(\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_i, \dots, \mathbf{r}_n) + \det(\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{y}_i, \dots, \mathbf{r}_n)$$

(These two properties are summarized by saying that the determinant is a *linear function* of each row.)

3. Interchanging any two rows of the matrix changes the sign of the determinant:

$$\det(\dots, \mathbf{r}_i, \dots, \mathbf{r}_j, \dots) = -\det(\dots, \mathbf{r}_j, \dots, \mathbf{r}_i, \dots)$$

4. The determinant of the $n \times n$ identity matrix is 1.

8.3 Some consequences of the definition

- If A has a row of zeros, then $\det(A) = 0$: Because if $A = (\dots, \mathbf{0}, \dots)$, then A also $= (\dots, c\mathbf{0}, \dots)$ for any c , and therefore, $\det(A) = c \det(A)$ for any c (property 1). This can only happen if $\det(A) = 0$.

- If $\mathbf{r}_i = \mathbf{r}_j$, $i \neq j$, then $\det(A) = 0$: Because then $\det(A) = \det(\dots, \mathbf{r}_i, \dots, \mathbf{r}_j, \dots) = -\det(\dots, \mathbf{r}_j, \dots, \mathbf{r}_i, \dots)$, by property 3, so $\det(A) = -\det(A)$ which means $\det(A) = 0$.
- If B is obtained from A by replacing row i with row $i + c(\text{row } j)$, then $\det(B) = \det(A)$:

$$\begin{aligned}
 \det(B) &= \det(\dots, \mathbf{r}_i + c\mathbf{r}_j, \dots) \\
 &= \det(\dots, \mathbf{r}_i, \dots) + \det(\dots, c\mathbf{r}_j, \dots) \\
 &= \det(A) + c \det(\dots, \mathbf{r}_j, \dots) \\
 &= \det(A) + 0
 \end{aligned}$$

The second determinant vanishes because both the i^{th} and j^{th} rows are equal to $\mathbf{r}_j$.

- **Theorem:** The determinant of an upper or lower triangular matrix with non-zero entries on the main diagonal is equal to the product of the entries on the main diagonal.

Proof: Suppose A is upper triangular. This means all the entries beneath the main diagonal are zero. This means we can clean out each column above the diagonal by using a row operation of the type just considered above. The end result is a matrix with the original non zero entries on the main diagonal and zeros elsewhere. Then repeated use of property 1 gives the result.

Remark: This is the property we use to compute determinants, because, as we know, row reduction leads to an upper triangular matrix.

Exercise: If A is an upper triangular matrix with one or more 0s on the main diagonal, then $\det(A) = 0$.

Examples

1. Let

$$A = \begin{pmatrix} 2 & 1 \\ 3 & -4 \end{pmatrix}.$$

Note that $\text{row } 1 = (2, 1) = 2(1, \frac{1}{2})$, so that

$$\begin{aligned}\det(A) &= 2 \det \begin{pmatrix} 1 & \frac{1}{2} \\ 3 & -4 \end{pmatrix} \\ &= 2 \det \begin{pmatrix} 1 & \frac{1}{2} \\ 0 & -\frac{11}{2} \end{pmatrix} \\ &= (2)(-\frac{11}{2}) \det \begin{pmatrix} 1 & \frac{1}{2} \\ 0 & 1 \end{pmatrix} \\ &= -11\end{aligned}$$

Exercise: Justify each of the above steps.

2. We can derive the formula for a 2×2 determinant in the same way: Let

$$A = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$$

Then

$$\begin{aligned}\det(A) &= a \det \begin{pmatrix} 1 & \frac{b}{a} \\ c & d \end{pmatrix} \\ &= \det \begin{pmatrix} 1 & \frac{b}{a} \\ 0 & d - \frac{bc}{a} \end{pmatrix} \\ &= a(d - \frac{bc}{a}) = ad - bc\end{aligned}$$

Exercises::

- Suppose $a = 0$ in the matrix A . Then we can't divide by a and the above computation won't work. Show that it's still true that $\det(A) = ad - bc$.
- Show that the three types of elementary matrices all have nonzero determinants.
- Suppose that $\text{row}_k(A)$ is a linear combination of rows i and j , where $i \neq j \neq k$: So $\mathbf{r}_k = a\mathbf{r}_i + b\mathbf{r}_j$. Show that $\det(A) = 0$.

There are two other important properties of the determinant, which we won't prove here (you can find the proofs in more advanced linear algebra texts):

- The determinant of A is the same as that of its transpose A^t .
- If A and B are square matrices of the same size, then

$$\det(AB) = \det(A) \det(B)$$

From the second of these, it follows that if A is invertible, then $\det(AA^{-1}) = \det(I) = 1 = \det(A) \det(A^{-1})$, so $\det(A^{-1}) = 1/\det(A)$.

Definition: If the (square) matrix A is invertible, then A is said to be *non-singular*. Otherwise, A is *singular*.

Exercises:

- Show that A is invertible $\iff \det(A) \neq 0$. (Hint: use the properties of determinants together with the theorem on GJ form and existence of the inverse.)
- A is singular $\iff$ the homogeneous equation $A\mathbf{x} = \mathbf{0}$ has nontrivial solutions. (Hint: If you don't want to do this directly, make an argument that this statement is logically equivalent to: A is non-singular $\iff$ the homogeneous equation has only the trivial solution.)
- Compute the determinants of the following matrices using the properties of the determinant; justify your work:

$$\begin{pmatrix} 1 & 2 & 3 \\ 1 & 0 & -1 \\ 2 & 3 & 1 \end{pmatrix}, \quad \begin{pmatrix} 1 & 2 & -3 & 0 \\ 2 & 6 & 0 & 1 \\ 1 & 4 & 3 & 1 \\ 2 & 4 & 6 & 8 \end{pmatrix}, \quad \text{and} \quad \begin{pmatrix} 1 & 0 & 0 \\ \pi & 4 & 0 \\ 3 & 7 & 5 \end{pmatrix}$$

9 The derivative as a linear transformation

9.1 Redefining the derivative

Matrices appear in many different contexts in mathematics, not just when we need to solve a system of linear equations. An important instance is linear approximation. Recall from your calculus course that a differentiable function f can be expanded about any point a in its domain using Taylor's theorem. We can write

$$f(x) = f(a) + f'(a)(x - a) + \frac{f''(c)}{2!}(x - a)^2,$$

where c is some point between x and a . The *remainder* term $\frac{f''(c)}{2!}(x - a)^2$ can be thought of as the “error” made by using the linear approximation to f at $x = a$,

$$f(x) \approx f(a) + f'(a)(x - a).$$

In fact, we can write Taylor's theorem in the more suggestive form

$$f(x) = f(a) + f'(a)(x - a) + \epsilon(x, a),$$

where the error function $\epsilon(x, a)$ has the important property

$$\lim_{x \rightarrow a} \frac{\epsilon(x, a)}{x - a} = 0.$$

(The existence of this limit is another way of saying that the error function “looks like” $(x - a)^2$.)

This observation gives us an alternative (and in fact, much better) definition of the derivative:

Definition: The real-valued function f is said to be *differentiable* at $x = a$ if there exists a number A and a function $\epsilon(x, a)$ such that

$$f(x) = f(a) + A(x - a) + \epsilon(x, a),$$

where

$$\lim_{x \rightarrow a} \frac{\epsilon(x, a)}{x - a} = 0.$$

Theorem: This is equivalent to the usual calculus definition.

Proof: If the new definition holds, then

$$\lim_{x \rightarrow a} \frac{f(x) - f(a)}{x - a} = A + \lim_{x \rightarrow a} \frac{\epsilon(x)}{x - a} = A + 0 = A,$$

and $A = f'(a)$ according to the standard definition. Conversely, if the standard definition of differentiability holds, then *define* ϵ to be the error made in the linear approximation:

$$\epsilon(x, a) = f(x) - f(a) - f'(a)(x - a).$$

Then

$$\lim_{x \rightarrow a} \frac{\epsilon(x, a)}{x - a} = \lim_{x \rightarrow a} \frac{f(x) - f(a)}{x - a} - f'(a) = f'(a) - f'(a) = 0,$$

so f can be written in the new form, with $A = f'(a)$.

Example: Let $f(x) = 4 + 2x - x^2$, and let $a = 1$. Then we can get a “linear approximation” by taking *any* number, say 43, and using it for A , writing $f(x) = f(1) + 43(x - 1) +$ “error term”, where by definition, the error term is what’s left: that is,

$$f(x) - f(1) - 43(x - 1) = 4 + 2x - x^2 - 5 - 43(x - 1) = 42 - 41x - x^2.$$

But you can see that if we were to define

$$\epsilon(x) = 42 - 41x - x^2 \quad (= 42(1 - x) + x(1 - x)),$$

then

$$\lim_{x \rightarrow 1} \frac{\epsilon}{x - 1} = -42 - 1 = -43,$$

which, you will notice, is not 0. The error term, instead of being purely quadratic in $x - 1$ (as required by the definition of differentiability), has a linear term: Using Taylor’s theorem to expand $\epsilon(x)$ about $x = 1$, we get (**exercise**)

$$\epsilon = 42 - 41x - x^2 = -43(x - 1) - (x - 1)^2$$

The *only* choice for the linear approximation in which the error term is purely quadratic is $f(x) \approx f(1) + f'(1)(x - 1)$.

Exercise: Interpret this geometrically in terms of the slope of various lines passing through the point $(1, f(1))$.

9.2 Generalization to higher dimensions

Our new definition of derivative is the one which generalizes to higher dimensions. We start with an

Example: Consider a function from R^2 to R^2 , say

$$\mathbf{f}(\mathbf{x}) = \mathbf{f} \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} u(x, y) \\ v(x, y) \end{pmatrix} = \begin{pmatrix} 2 + x + 4y + 4x^2 + 5xy - y^2 \\ 1 - x + 2y - 2x^2 + 3xy + y^2 \end{pmatrix}$$

By inspection, as it were, we can separate the right hand side into three parts. We have

$$\mathbf{f}(\mathbf{0}) = \begin{pmatrix} 2 \\ 1 \end{pmatrix}$$

and the linear part of $\mathbf{f}$ is the vector

$$\begin{pmatrix} x + 4y \\ -x + 2y \end{pmatrix},$$

which can be written in matrix form as

$$A\mathbf{x} = \begin{pmatrix} 1 & 4 \\ -1 & 2 \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix}.$$

By analogy with the one-dimensional case, we might guess that

$$\mathbf{f}(\mathbf{x}) = \mathbf{f}(\mathbf{0}) + A\mathbf{x} + \text{an error term of order 2 in } x, y.$$

where A is the matrix

$$A = \begin{pmatrix} \frac{\partial u}{\partial x} & \frac{\partial u}{\partial y} \\ \frac{\partial v}{\partial x} & \frac{\partial v}{\partial y} \end{pmatrix} (0, 0).$$

And this suggests the following

Definition: A function $f : R^n \rightarrow R^m$ is said to be *differentiable* at the point $\mathbf{x} = \mathbf{a} \in R^n$ if there exists an $m \times n$ matrix A and a function $\epsilon(\mathbf{x}, \mathbf{a})$ such that

$$\mathbf{f}(\mathbf{x}) = \mathbf{f}(\mathbf{a}) + A(\mathbf{x} - \mathbf{a}) + \epsilon(\mathbf{x}, \mathbf{a}),$$

where

$$\lim_{\|\mathbf{x}-\mathbf{a}\| \rightarrow 0} \frac{\epsilon}{\|\mathbf{x}-\mathbf{a}\|} = \mathbf{0}.$$

The matrix A is called the *derivative of $\mathbf{f}$ at $\mathbf{x} = \mathbf{a}$* , and is denoted by $D\mathbf{f}(\mathbf{a})$.

Generalizing the one-dimensional case, it can be shown that if

$$\mathbf{f}(\mathbf{x}) = \begin{pmatrix} u_1(\mathbf{x}) \\ \vdots \\ u_m(\mathbf{x}) \end{pmatrix},$$

is differentiable at $\mathbf{x} = \mathbf{a}$, then the derivative of $\mathbf{f}$ is given by the $m \times n$ matrix of partial derivatives

$$D\mathbf{f}(\mathbf{a}) = \begin{pmatrix} \frac{\partial u_1}{\partial x_1} & \cdots & \frac{\partial u_1}{\partial x_n} \\ \vdots & \vdots & \vdots \\ \frac{\partial u_m}{\partial x_1} & \cdots & \frac{\partial u_m}{\partial x_n} \end{pmatrix}_{m \times n} (\mathbf{a}).$$

Conversely, if all the indicated partial derivatives exist and are continuous at $\mathbf{x} = \mathbf{a}$, then the approximation

$$\mathbf{f}(x) \approx \mathbf{f}(\mathbf{a}) + D\mathbf{f}(\mathbf{a})(\mathbf{x} - \mathbf{a})$$

is accurate to the second order in $\mathbf{x} - \mathbf{a}$.

Exercise: Find the derivative of the function $\mathbf{f} : \mathbb{R}^2 \rightarrow \mathbb{R}^3$ at $\mathbf{a} = (1, 2)^t$, where

$$\mathbf{f}(\mathbf{x}) = \begin{pmatrix} (x + y)^3 \\ x^2 y^3 \\ y/x \end{pmatrix}$$

10 Subspaces

Definitions:

- A *linear combination of the vectors* $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_m$ is any vector of the form $c_1\mathbf{v}_1 + c_2\mathbf{v}_2 + \dots + c_m\mathbf{v}_m$, where $c_1, \dots, c_m \in \mathbb{R}$.
- A subset V of $\mathbb{R}^n$ is a *subspace* if, whenever $\mathbf{v}_1, \mathbf{v}_2 \in V$, and c_1 , and c_2 are any real numbers, the linear combination $c_1\mathbf{v}_1 + c_2\mathbf{v}_2 \in V$.

Remark: Suppose that V is a subspace, and that $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_m$ all belong to V . Then $c_1\mathbf{x}_1 + c_2\mathbf{x}_2 \in V$. Therefore, $(c_1\mathbf{x}_1 + c_2\mathbf{x}_2) + c_3\mathbf{x}_3 \in V$. Similarly, $(c_1\mathbf{x}_1 + \dots + c_{m-1}\mathbf{x}_{m-1}) + c_m\mathbf{x}_m \in V$. We say that the subspace V is *closed under linear combinations*.

Examples:

- The set of all solutions to the homogeneous equation $A\mathbf{x} = \mathbf{0}$ is a subspace of $\mathbb{R}^n$ if A is $m \times n$.

Proof: Suppose $\mathbf{x}_1$ and $\mathbf{x}_2$ are solutions; we need to show that $c_1\mathbf{x}_1 + c_2\mathbf{x}_2$ is also a solution. Because $\mathbf{x}_1$ is a solution, $A\mathbf{x}_1 = \mathbf{0}$. Similarly, $A\mathbf{x}_2 = \mathbf{0}$. Then for any scalars c_1, c_2 , $A(c_1\mathbf{x}_1 + c_2\mathbf{x}_2) = c_1A\mathbf{x}_1 + c_2A\mathbf{x}_2 = c_1\mathbf{0} + c_2\mathbf{0} = \mathbf{0}$. So $c_1\mathbf{x}_1 + c_2\mathbf{x}_2$ is also a solution. The set of solutions is closed under linear combinations and so it's a subspace.

Definition: This important subspace is called the *null space of A* , and is denoted $\text{Null}(A)$

- The set V of all vectors in $\mathbb{R}^3$ which are orthogonal (perpendicular) to the vector

$$\mathbf{v} = \begin{pmatrix} 2 \\ -1 \\ 3 \end{pmatrix}$$

is a subspace of $\mathbb{R}^3$.

Proof: $\mathbf{x}$ is orthogonal to $\mathbf{v}$ ($\mathbf{x} \perp \mathbf{v}$) means that $\mathbf{x} \cdot \mathbf{v} = 0$. So suppose that $\mathbf{x}_1$ and $\mathbf{x}_2$ are orthogonal to $\mathbf{v}$. Then, using the properties of the dot product, for any constants c_1, c_2 , we have

$$(c_1\mathbf{x}_1 + c_2\mathbf{x}_2) \cdot \mathbf{v} = c_1(\mathbf{x}_1 \cdot \mathbf{v}) + c_2(\mathbf{x}_2 \cdot \mathbf{v}) = c_1\mathbf{0} + c_2\mathbf{0} = \mathbf{0}.$$

And therefore $(c_1\mathbf{x}_1 + c_2\mathbf{x}_2) \perp \mathbf{v}$, so we have a subspace.

- The set consisting of the single vector $\mathbf{0}$ is a subspace of $\mathbb{R}^n$ for any n : any linear combination of elements of this set is a multiple of $\mathbf{0}$, and hence equal to $\mathbf{0}$ which is in the set.
- $\mathbb{R}^n$ is a subspace of itself since any linear combination of vectors in the set is again in the set.
- Take any finite or infinite set $S \subset \mathbb{R}^n$

Definition: The *span of S* is the set of all finite linear combinations of elements of S :

$$\text{span}(S) = \left\{ \mathbf{x} : \mathbf{x} = \sum_{i=1}^n c_i \mathbf{v}_i, \text{ where } \mathbf{v}_i \in S, \text{ and } n < \infty \right\}$$

Exercise: Show that $\text{span}(S)$ is a subspace of $\mathbb{R}^n$.

Definition; If $V = \text{span}(S)$, then the vectors in S are said to *span the subspace V* . (So the word “span” is used in 2 ways.)

Example: Referring back to the section on solutions to the homogeneous equation, we had an example for which the general solution to $A\mathbf{x} = \mathbf{0}$ took the form

$$\mathbf{x}_H = \{s\mathbf{x}_1 + t\mathbf{x}_2, \quad t, s \in \mathbb{R}\}$$

So $\mathbf{x}_H = \text{span}(\mathbf{v}_1, \mathbf{v}_2)$. And, of course, $\mathbf{x}_H = \text{Null}(A)$ is just the null space of the matrix A . (We will not use the obscure notation $\mathbf{x}_H$ for this subspace any longer.)

How can you tell if a particular vector belongs to $\text{span}(S)$? You have to show that you can (or cannot) write it as a linear combination of vectors in S .

Example:

$$\text{Is } \mathbf{v} = \begin{pmatrix} 1 \\ 2 \\ 3 \end{pmatrix}$$

in the span of

$$\left\{ \begin{pmatrix} 1 \\ 0 \\ 1 \end{pmatrix}, \begin{pmatrix} 2 \\ -1 \\ 2 \end{pmatrix} \right\} = \{\mathbf{x}_1, \mathbf{x}_2\}?$$

Answer: It is if there exist numbers c_1 and c_2 such that $\mathbf{v} = c_1\mathbf{x}_1 + c_2\mathbf{x}_2$. Writing this out gives a system of linear equations:

$$\mathbf{v} = \begin{pmatrix} 1 \\ 2 \\ 3 \end{pmatrix} = c_1 \begin{pmatrix} 1 \\ 0 \\ 1 \end{pmatrix} + c_2 \begin{pmatrix} 2 \\ -1 \\ 2 \end{pmatrix}.$$

In matrix form, this reads

$$\begin{pmatrix} 1 & 2 \\ 0 & -1 \\ 1 & 2 \end{pmatrix} \begin{pmatrix} c_1 \\ c_2 \end{pmatrix} = \begin{pmatrix} 1 \\ 2 \\ 3 \end{pmatrix}$$

As you can (and should!) verify, this system is inconsistent. No such c_1, c_2 exist. So $\mathbf{v}$ is not in the span of these two vectors.

- The set of all solutions to the inhomogeneous system $A\mathbf{x} = \mathbf{y}$, $\mathbf{y} \neq \mathbf{0}$ is *not* a subspace. To see this, suppose that $\mathbf{x}_1$ and $\mathbf{x}_2$ are two solutions. We'll have a subspace if any linear combination of these two vectors is again a solution. So we compute

$$\begin{aligned} A(c_1\mathbf{x}_1 + c_2\mathbf{x}_2) &= c_1A\mathbf{x}_1 + c_2A\mathbf{x}_2 \\ &= c_1\mathbf{y} + c_2\mathbf{y} \\ &= (c_1 + c_2)\mathbf{y}, \end{aligned}$$

Since for general c_1, c_2 the right hand side is *not* equal to $\mathbf{y}$, this is not a subspace.

NOTE: To show that V is or is not a subspace does not, as a general rule, require any prodigious intellectual effort. Just assume that $\mathbf{x}_1, \mathbf{x}_2 \in V$, and see if $c_1\mathbf{x}_1 + c_2\mathbf{x}_2 \in V$

for arbitrary scalars c_1, c_2 . If so, it's a subspace, otherwise no. The scalars must be arbitrary, and $\mathbf{x}_1, \mathbf{x}_2$ must be arbitrary elements of V . (So you can't pick two of your favorite vectors and two of your favorite scalars for this proof - that's why we always use "generic" elements like $\mathbf{x}_1$, and c_1 .)

- **Definition:** Suppose A is $m \times n$. The m rows of A form a subset of $\mathbb{R}^n$; the span of these vectors is called the *row space of the matrix*. Similarly, the n columns of A form a set of vectors in $\mathbb{R}^m$, and the space they span is called the *column space of the matrix A* .

Example: For the matrix

$$A = \begin{pmatrix} 1 & 0 & -1 & 2 \\ 3 & 4 & 6 & -1 \\ 2 & 5 & -9 & 7 \end{pmatrix},$$

the row space of A is $\text{span}\{(1, 0, -1, 2)^t, (3, 4, 6, -1)^t, (2, 5, -9, 7)^t\}$, and the column space is

$$\text{span}\left\{\begin{pmatrix} 1 \\ 3 \\ 2 \end{pmatrix}, \begin{pmatrix} 0 \\ 4 \\ 5 \end{pmatrix}, \begin{pmatrix} -1 \\ 6 \\ -9 \end{pmatrix}, \begin{pmatrix} 2 \\ -1 \\ 7 \end{pmatrix}\right\}$$

Exercises:

- A plane through $\mathbf{0}$ in $\mathbb{R}^3$ is a subspace of $\mathbb{R}^3$. A plane which does not contain the origin is not a subspace. (Hint: what are the equations for these planes?)
- When is a line in $\mathbb{R}^2$ a subspace of $\mathbb{R}^2$?

²In many texts, vectors are written as row vectors for typographical reasons (it takes up less space). But for computations the vectors should always be written as columns, which is why the symbols for the transpose appear here

11 Linear dependence and independence

Definition: A finite set $S = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_m\}$ of vectors in $\mathbb{R}^n$ is said to be *linearly dependent* if there exist scalars (real numbers) $c_1, c_2, \dots, c_m$, not all of which are 0, such that $c_1\mathbf{x}_1 + c_2\mathbf{x}_2 + \dots + c_m\mathbf{x}_m = \mathbf{0}$.

Examples:

1. The vectors

$$\mathbf{x}_1 = \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}, \mathbf{x}_2 = \begin{pmatrix} 1 \\ -1 \\ 2 \end{pmatrix}, \text{ and } \mathbf{x}_3 = \begin{pmatrix} 3 \\ 1 \\ 4 \end{pmatrix}$$

are linearly dependent because $2\mathbf{x}_1 + \mathbf{x}_2 - \mathbf{x}_3 = \mathbf{0}$.

2. Any set containing the vector $\mathbf{0}$ is linearly dependent, because for any $c \neq 0$, $c\mathbf{0} = \mathbf{0}$.
3. In the definition, we require that not all of the scalars $c_1, \dots, c_n$ are 0. The reason for this is that otherwise, any set of vectors would be linearly dependent.
4. If a set of non-zero vectors is linearly dependent, then one of them can be written as a linear combination of the others: (We just do this for 3 vectors, but it is true for any number). Suppose $c_1\mathbf{x}_1 + c_2\mathbf{x}_2 + c_3\mathbf{x}_3 = \mathbf{0}$, where at least one of the c 's is not zero. If, say, $c_2 \neq 0$, then we can solve for $\mathbf{x}_2$:

$$\mathbf{x}_2 = (-1/c_2)(c_1\mathbf{x}_1 + c_3\mathbf{x}_3).$$

And similarly if some other coefficient is not zero.

5. In principle, it is an easy matter to determine whether a set S is linearly dependent: We write down a system of linear algebraic equations and see if there are solutions. For instance, suppose

$$S = \left\{ \begin{pmatrix} 1 \\ 2 \\ 1 \end{pmatrix}, \begin{pmatrix} 1 \\ 0 \\ -1 \end{pmatrix}, \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} \right\} = \{\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3\}.$$

By the definition, S is linearly dependent $\iff$ we can find scalars c_1, c_2 , and c_3 , not all 0, such that

$$c_1\mathbf{x}_1 + c_2\mathbf{x}_2 + c_3\mathbf{x}_3 = \mathbf{0}.$$

We write this equation out in matrix form:

$$\begin{pmatrix} 1 & 1 & 1 \\ 2 & 0 & 1 \\ 1 & -1 & 1 \end{pmatrix} \begin{pmatrix} c_1 \\ c_2 \\ c_3 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}$$

Evidently, the set S is linearly dependent if and only if there is a non-trivial solution to this homogeneous equation. Row reduction of the matrix leads quickly to

$$\begin{pmatrix} 1 & 1 & 1 \\ 0 & 1 & \frac{1}{2} \\ 0 & 0 & 1 \end{pmatrix}.$$

This matrix is non-singular, so the only solution to the homogeneous equation is the trivial one with $c_1 = c_2 = c_3 = 0$. So the vectors are *not* linearly dependent.

Definition: the set S is *linearly independent* if it's not linearly dependent.

What could be clearer? The set S is not linearly dependent if, whenever some linear combination of the elements of S adds up to $\mathbf{0}$, it turns out that $c_1, c_2, \dots$ are *all* zero. In the last example above, we assumed that $c_1\mathbf{x}_1 + c_2\mathbf{x}_2 + c_3\mathbf{x}_3 = \mathbf{0}$ and were led to the conclusion that all the coefficients must be 0. So this set is linearly independent.

The "test" for linear independence is the same as that for linear dependence. We set up a homogeneous system of equations, and find out whether or not it has non-trivial solutions

Exercises:

1. A set S consisting of two different vectors $\mathbf{u}$ and $\mathbf{v}$ is linearly dependent $\iff$ one of the two is a nonzero multiple of the other. (Don't forget the possibility that one of

the vectors could be $\mathbf{0}$). If neither vector is $\mathbf{0}$, the vectors are linearly dependent if they are parallel. What is the geometric condition for three nonzero vectors in $\mathbb{R}^3$ to be linearly dependent?

2. Find two linearly independent vectors belonging to the null space of the matrix

$$A = \begin{pmatrix} 3 & 2 & -1 & 4 \\ 1 & 0 & 2 & 3 \\ -2 & -2 & 3 & -1 \end{pmatrix}.$$

3. Are the columns of A (above) linearly independent in $\mathbb{R}^3$? Why? Are the rows of A linearly independent in $\mathbb{R}^3$? Why?

11.1 Elementary row operations

We can show that elementary row operations performed on a matrix A don't change the row space. We just give the proof for one of the operations; the other two are left as exercises.

Suppose that, in the matrix A , $\text{row}_i(A)$ is replaced by $\text{row}_i(A) + c \cdot \text{row}_j(A)$. Call the resulting matrix B . If $\mathbf{x}$ belongs to the row space of A , then

$$\mathbf{x} = c_1 \text{row}_1(A) + \dots + c_i \text{row}_i(A) + \dots + c_j \text{row}_j(A) + c_m \text{row}_m(A).$$

Now add and subtract $c \cdot c_i \cdot \text{row}_j(A)$ to get

$$\begin{aligned} \mathbf{x} &= c_1 \text{row}_1(A) + \dots + c_i \text{row}_i(A) + c \cdot c_i \text{row}_j(A) + \dots + (c_j - c_i \cdot c) \text{row}_j(A) + c_m \text{row}_m(A) \\ &= c_1 \text{row}_1(B) + \dots + c_i \text{row}_i(B) + \dots + (c_j - c_i \cdot c) \text{row}_j(B) + \dots + c_m \text{row}_m(B). \end{aligned}$$

This shows that $\mathbf{x}$ can also be written as a linear combination of the rows of B . So any element in the row space of A is contained in the row space of B .

Exercise: Show the converse - that any element in the row space of B is contained in the row space of A .

Definition: Two *sets X and Y are equal* if $X \subseteq Y$ and $Y \subseteq X$.

This is what we've just shown for the two row spaces.

Exercises:

1. Show that the other two elementary row operations don't change the row space of A .
2. **Show that when we multiply any matrix A by another matrix B on the left, the rows of the product BA are linear combinations of the rows of A .
3. **Show that when we multiply A on the right by B , that the columns of AB are linear combinations of the columns of A .

12 Basis and dimension of subspaces

12.1 The concept of basis

It follows from what we've said above that if $S = \{\mathbf{e}_1, \dots, \mathbf{e}_m\}$ spans the subspace V ³ but is linearly dependent, we can express one of the elements in S as a linear combination of the others. By relabeling if necessary, we suppose that $\mathbf{e}_m$ can be written as a linear combination of the others. Then

$$\text{span}(S) = \text{span}(\mathbf{e}_1, \dots, \mathbf{e}_{m-1}). \text{ Why?}$$

If the remaining $m - 1$ vectors are still linearly dependent, we can repeat the process, writing one of them as a linear combination of the remaining $m - 2$, relabeling, and then

$$\text{span}(S) = \text{span}(\mathbf{e}_1, \dots, \mathbf{e}_{m-2}).$$

We can continue this until we arrive finally at a "minimal" spanning set, say $\{\mathbf{e}_1, \dots, \mathbf{e}_k\}$. Such a set will be called a basis for V :

Definition: The set $B = \{\mathbf{e}_1, \dots, \mathbf{e}_k\}$ is a *basis* for the subspace V if

- $\text{span}(B) = V$.
- B is linearly independent.

Remark: In definitions like that given above, we really should put "iff" (if and only if) instead of just "if", and that's the way you should read it. More precisely, if B is a basis, then B spans V and is linearly independent. Conversely, if B spans V and is linearly independent, then B is a basis.

Examples:

³We use the word *span* in two ways: if $V = \text{span}S$, then we say that S *spans* the subspace V .

- In $\mathbb{R}^3$, the set

$$B = \left\{ \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} \right\} = \{\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3\}$$

is a basis..

Why? (a) Any vector

$$\mathbf{v} = \begin{pmatrix} a \\ b \\ c \end{pmatrix}$$

in $\mathbb{R}^3$ can be written as $\mathbf{v} = a\mathbf{e}_1 + b\mathbf{e}_2 + c\mathbf{e}_3$, so B spans $\mathbb{R}^3$. And (b): if $c_1\mathbf{e}_1 + c_2\mathbf{e}_2 + c_3\mathbf{e}_3 = \mathbf{0}$, then

$$\begin{pmatrix} c_1 \\ c_2 \\ c_3 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix},$$

which means that $c_1 = c_2 = c_3 = 0$, so the set is linearly independent.

Definition: The set $\{\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3\}$ is called the *standard basis* for $\mathbb{R}^3$.

- **Exercise:** Any 4 vectors in $\mathbb{R}^3$ are linearly dependent and therefore do *not* form a basis. You should be able to supply the argument, which amounts to showing that a certain homogeneous system of equations has a nontrivial solution.
- **Exercise:** No 2 vectors can span $\mathbb{R}^3$. Why not?
- If a set B is a basis for $\mathbb{R}^3$, then it contains exactly 3 elements. This is a consequence of the previous two statements.

Exercise: Prove that any basis for $\mathbb{R}^n$ has precisely n elements.

- **Example:** Find a basis for the null space of the matrix

$$A = \begin{pmatrix} 1 & 0 & 0 & 3 & 2 \\ 0 & 1 & 0 & 1 & -1 \\ 0 & 0 & 1 & 2 & 3 \end{pmatrix}.$$

Solution: Since A is already in Gauss-Jordan form, we can just write down the general solution to the homogeneous equation. These are the elements of the null space of A .

We have, setting $x_4 = s$, and $x_5 = t$,

$$\begin{aligned}x_1 &= -3s - 2t \\x_2 &= -s + t \\x_3 &= -2s - 3t, \\x_4 &= s \\x_5 &= t\end{aligned}$$

so the general solution to $A\mathbf{x} = \mathbf{0}$ is given by $\mathbf{K}_A = \{s\mathbf{v}_1 + t\mathbf{v}_2 \mid s, t \in \mathbb{R}\}$, where

$$\mathbf{v}_1 = \begin{pmatrix} -3 \\ -1 \\ -2 \\ 1 \\ 0 \end{pmatrix}, \text{ and } \mathbf{v}_2 = \begin{pmatrix} -2 \\ 1 \\ -3 \\ 0 \\ 1 \end{pmatrix}.$$

It is obvious by inspection of the last two entries in each that the set $B = \{\mathbf{v}_1, \mathbf{v}_2\}$ is linearly independent. Furthermore, by construction, the set B spans the null space. So B is a basis.

12.2 Dimension

As we've seen above, any basis for $\mathbb{R}^n$ has precisely n elements. Although we're not going to prove it here, the same property holds for any subspace of $\mathbb{R}^n$: the number of elements in any basis for the subspace is the same. Given this, we make the following

Definition: Let $V \neq \{\mathbf{0}\}$ be a subspace of $\mathbb{R}^n$ for some n . The *dimension* of V , written $\dim(V)$, is the number of elements in any basis of V .

Examples:

- $\dim(\mathbb{R}^n) = n$. Why?

- For the matrix A above, the dimension of the null space of A is 2.
- The subspace $V = \{\mathbf{0}\}$ is a bit peculiar: it doesn't have a basis according to our definition, since any subset of V is linearly independent. We extend the definition of dimension to this case by *defining* $\dim(V) = 0$.

Exercises:

1. Show that the dimension of the null space of any matrix R in reduced echelon form is equal to the number of free variables in the echelon form.
2. Show that the dimension of the set

$$\{(x, y, z) \text{ such that } 2x - 3y + z = 0\}$$

is two.

13 The rank-nullity (dimension) theorem

13.1 Rank and nullity of a matrix

Definition: The *nullity* of the matrix A is the dimension of the null space of A , and is denoted by $N(A)$.

Examples: The nullity of I is 0. The nullity of the 3×5 matrix considered above is 2. The nullity of $0_{m \times n}$ is n .

Definition: The *rank* of the matrix A is the dimension of the row space of A , and is denoted $R(A)$

Examples: The rank of $I_{n \times n}$ is n ; the rank of $0_{m \times n}$ is 0. The rank of the 3×5 matrix considered above is 3.

Definition: The matrix B is said to be *row equivalent* to A if B can be obtained from A by a finite sequence of elementary row operations. In pure matrix terms, this means precisely that

$$B = E_k E_{k-1} \cdots E_2 E_1 A,$$

where $E_1, \dots, E_k$ are elementary row matrices. We can now establish two important results:

Theorem: If B is row equivalent to A , then $\text{Null}(B) = \text{Null}(A)$.

Proof: Suppose $\mathbf{x} \in \text{Null}(A)$. Then $A\mathbf{x} = \mathbf{0}$. Since $B = E_k \cdots E_1 A$, it follows that $B\mathbf{x} = E_k \cdots E_1 A\mathbf{x} = E_k \cdots E_1 \mathbf{0} = \mathbf{0}$, so $\mathbf{x} \in \text{Null}(B)$, and therefore that $\text{Null}(A) \subseteq \text{Null}(B)$. Conversely, if $\mathbf{x} \in \text{Null}(B)$, then $B\mathbf{x} = \mathbf{0}$. But $B = CA$, where C is invertible, being the product of elementary matrices. Thus $B\mathbf{x} = CA\mathbf{x} = \mathbf{0}$. Multiplying by C^{-1} gives $A\mathbf{x} = C^{-1}\mathbf{0} = \mathbf{0}$, so $\mathbf{x} \in \text{Null}(A)$, and $\text{Null}(B) \subseteq \text{Null}(A)$. So the two sets are equal, as advertised.

Theorem: If B is row equivalent to A , then the row space of B is identical to that of A

Proof: Suppose first that $B = EA$, where E is the matrix of some elementary row operation. If E interchanges rows, then the span of the new set of rows is the same as the span of the old set, so the row space doesn't change. Similarly if E multiplies one row by a nonzero scalar. Finally, if E is the operation corresponding to $\text{row}_i(A) \rightarrow \text{row}_i(A) + c \cdot \text{row}_j(A)$, then the span of the rows of the new matrix is the same as the span of the rows of A (why?). Since the theorem is true for any single row operation, it's true for any finite number of them, which completes the proof.

Summarizing these results: **Row operations do not change either the row space or the null space of A .**

Corollary 1: If R is the Gauss-Jordan form of A , then R has the same null space and row space as A .

Corollary 2: If B is row equivalent to A , then $R(B) = R(A)$, and $N(B) = N(A)$.

Exercise: $R(A)$ is equal to the number of leading 1's in the echelon form of A .

The following result may be somewhat surprising:

Theorem: The number of linearly independent rows of the matrix A is equal to the number of linearly independent columns of A . In particular, the rank of A is also equal to the number of linearly independent columns.

Proof (sketch): As an example, suppose that columns i , j , and k are linearly independent, with

$$\text{col}_i(A) = 2\text{col}_j(A) - 3\text{col}_k(A).$$

You should be able to convince yourself that doing any row operation on the matrix A doesn't affect this equation. Even though the row operation changes the entries of the various columns, it changes them all in the same way, and this equation continues to hold. The *span* of the columns can, and generally will change under row operations (why?), but this doesn't affect the result.

The actual proof would consist of the following steps: (1) identify a maximal linearly independent set of columns of A , (2) argue that this set remains linearly independent if row operations are done on A . (3) Then it follows that the number of linearly independent columns in the reduced echelon form of A is the same as the number of linearly independent columns in A . The number of linearly independent columns of A is then just the number of leading entries in the reduced echelon form of A which is, as we know, the same as the rank of A .

13.2 The rank-nullity theorem

This is also known as the dimension theorem, and version 1 (we'll see another later in the course) goes as follows:

Theorem: Let A be $m \times n$. Then

$$n = N(A) + R(A),$$

where n is the number of columns of A .

Let's assume, for the moment, that this is true. What good is it? Answer: You can read off both the rank and the nullity from the echelon form of the matrix A . Suppose A can be row-reduced to

$$\begin{pmatrix} 1 & * & * & * & * \\ 0 & 0 & 1 & * & * \\ 0 & 0 & 0 & 1 & * \end{pmatrix}.$$

Then it's clear (why?) that the dimension of the row space is 3, or equivalently, that the dimension of the column space is 3. Since there are 5 columns altogether, the dimension theorem says that $n = 5 = 3 + N(A)$, so $N(A) = 2$. We can therefore expect to find two linearly independent solutions to the homogeneous equation $A\mathbf{x} = \mathbf{0}$.

Alternatively, inspection of the echelon form of A reveals that there are precisely 2 free variables, x_2 and x_5 . So we know that $N(A) = 2$ (why?), and therefore, $\text{rank}(A) = 5 - 2 = 3$.

Proof of the theorem: This is, at this point, almost trivial. We have shown above that the rank of A is the same as the rank of the Gauss-Jordan form of A which is clearly equal to the number of leading entries in the Gauss-Jordan form. We also know that the dimension of the null space is equal to the number of free variables in the reduced echelon (GJ) form of A . And we know further that the number of free variables plus the number of leading entries is exactly the number of columns. So

$$n = N(A) + R(A),$$

as claimed.

14 Change of basis

When we first set up a problem in mathematics, we normally use the most familiar coordinates. In $\mathbb{R}^3$, this means using the Cartesian coordinates x , y , and z . In vector terms, this is equivalent to using what we've called the standard basis in $\mathbb{R}^3$; that is, we write

$$\begin{pmatrix} x \\ y \\ z \end{pmatrix} = x \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} + y \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix} + z \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} = x\mathbf{e}_1 + y\mathbf{e}_2 + z\mathbf{e}_3,$$

where $\{\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3\}$ is the standard basis.

But, as you know, for any particular problem, there is often another coordinate system that simplifies the problem. For example, to study the motion of a planet around the sun, we put the sun at the origin, and use polar or spherical coordinates. This happens in linear algebra as well.

Example: Let's look at a simple system of two first order linear differential equations

$$\begin{aligned} x_1' &= 3x_1 + x_2 \\ x_2' &= x_1 + 3x_2 \end{aligned} \quad (1)$$

Here, we seek functions $x_1(t)$, and $x_2(t)$ such that *both* equations hold simultaneously. Now there's no problem solving a single differential equation like

$$x' = 3x.$$

In fact, we can see by inspection that $x(t) = ce^{3t}$ is a solution for any scalar c . The difficulty with the system (1) is that x_1 and x_2 are "coupled", and the two equations must be solved simultaneously. There are a number of straightforward ways to solve this system which you'll learn when you take a course in differential equations, and we won't worry about that here.

But there's also a sneaky way to solve (1) by changing coordinates. We'll do this at the end of the lecture. First, we need to see what happens in general when we change the basis.

For simplicity, we're just going to work in $\mathbb{R}^2$; generalization to higher dimensions is (really!) straightforward.

Suppose we have a basis $\{\mathbf{e}_1, \mathbf{e}_2\}$ for $\mathbb{R}^2$. It doesn't have to be the standard basis. Then, by the definition of basis, any vector $\mathbf{v} \in \mathbb{R}^2$ can be written as a linear combination of $\mathbf{e}_1$ and $\mathbf{e}_2$. That is, there exist scalars c_1, c_2 such that $\mathbf{v} = c_1\mathbf{e}_1 + c_2\mathbf{e}_2$.

Definition: The numbers c_1 and c_2 are called the *coordinates* of $\mathbf{v}$ in the basis $\{\mathbf{e}_1, \mathbf{e}_2\}$. And

$$\mathbf{v}_e = \begin{pmatrix} c_1 \\ c_2 \end{pmatrix}$$

is called the *coordinate vector* of $\mathbf{v}$ in the basis $\{\mathbf{e}_1, \mathbf{e}_2\}$.

Theorem: The coordinates of the vector $\mathbf{v}$ are *unique*.

Proof: Suppose there are two sets of coordinates for $\mathbf{v}$. That is, suppose that $\mathbf{v} = c_1\mathbf{e}_1 + c_2\mathbf{e}_2$, and also that $\mathbf{v} = d_1\mathbf{e}_1 + d_2\mathbf{e}_2$. Subtracting the two expressions for $\mathbf{v}$ gives

$$\mathbf{0} = (c_1 - d_1)\mathbf{e}_1 + (c_2 - d_2)\mathbf{e}_2.$$

But $\{\mathbf{e}_1, \mathbf{e}_2\}$ is linearly independent, so the coefficients in this expression must vanish: $c_1 - d_1 = c_2 - d_2 = 0$. That is, $c_1 = d_1$ and $c_2 = d_2$, and the coordinates are unique, as claimed.

Example: Let us use the basis

$$\{\mathbf{e}_1, \mathbf{e}_2\} = \left\{ \begin{pmatrix} 1 \\ 2 \end{pmatrix}, \begin{pmatrix} -2 \\ 3 \end{pmatrix} \right\},$$

and suppose

$$\mathbf{v} = \begin{pmatrix} 3 \\ 5 \end{pmatrix}.$$

Then we can find the coordinate vector $\mathbf{v}_e$ in this basis in the usual way, by solving a system of linear equations. We are looking for numbers c_1 and c_2 (the coordinates of $\mathbf{v}$ in this basis) such that

$$c_1 \begin{pmatrix} 1 \\ 2 \end{pmatrix} + c_2 \begin{pmatrix} -2 \\ 3 \end{pmatrix} = \begin{pmatrix} 3 \\ 5 \end{pmatrix}.$$

In matrix form, this reads

$$A\mathbf{v}_e = \mathbf{v},$$

where

$$A = \begin{pmatrix} 1 & -2 \\ 2 & 3 \end{pmatrix}, \quad \mathbf{v} = \begin{pmatrix} 3 \\ 5 \end{pmatrix}, \quad \text{and } \mathbf{v}_e = \begin{pmatrix} c_1 \\ c_2 \end{pmatrix}.$$

We solve for $\mathbf{v}_e$ by multiplying both sides by A^{-1} :

$$\mathbf{v}_e = A^{-1}\mathbf{v} = (1/7) \begin{pmatrix} 3 & 2 \\ -2 & 1 \end{pmatrix} \begin{pmatrix} 3 \\ 5 \end{pmatrix} = (1/7) \begin{pmatrix} 19 \\ -1 \end{pmatrix} = \begin{pmatrix} 19/7 \\ -1/7 \end{pmatrix}$$

Exercise: Find the coordinates of the vector $\mathbf{v} = (-2, 4)^t$ in this basis.

14.1 Notation

In this section, we'll develop a compact notation for the above computation that is easy to remember. Start with an arbitrary basis $\{\mathbf{e}_1, \mathbf{e}_2\}$ and an arbitrary vector $\mathbf{v}$. We know that

$$\mathbf{v} = c_1\mathbf{e}_1 + c_2\mathbf{e}_2,$$

where

$$\begin{pmatrix} c_1 \\ c_2 \end{pmatrix} = \mathbf{v}_e$$

is the coordinate vector. We see that the expression for $\mathbf{v}$ is a linear combination of two column vectors. And we know that such a thing can be obtained by writing down a certain matrix product:

If we define the 2×2 matrix $E = (\mathbf{e}_1, \mathbf{e}_2)$ then the expression for $\mathbf{v}$ can be simply written as

$$\mathbf{v} = E \cdot \mathbf{v}_e.$$

Now suppose that $\{\mathbf{f}_1, \mathbf{f}_2\}$ is another basis for $\mathbb{R}^2$. Then the same vector $\mathbf{v}$ can also be written uniquely as a linear combination of these vectors. Of course it will have *different*

coordinates, and a different coordinate vector $\mathbf{v}_f$. In matrix form, we'll have

$$\mathbf{v} = F \cdot \mathbf{v}_f.$$

Exercise: Let $\{\mathbf{f}_1, \mathbf{f}_2\}$ be given by

$$\left\{ \begin{pmatrix} 1 \\ 1 \end{pmatrix}, \begin{pmatrix} 1 \\ -1 \end{pmatrix} \right\}.$$

If

$$\mathbf{v} = \begin{pmatrix} 3 \\ 5 \end{pmatrix},$$

(same vector as above) find $\mathbf{v}_f$ and verify that $\mathbf{v} = F \cdot \mathbf{v}_f = E \cdot \mathbf{v}_e$.

Remark: This works just the same in $\mathbb{R}^n$, where $E = (\mathbf{e}_1 \cdots \mathbf{e}_n)$ is $n \times n$, and $\mathbf{v}_e$ is $n \times 1$.

Continuing along with our examples, since E is a basis, the vectors $\mathbf{f}_1$ and $\mathbf{f}_2$ can each be written as linear combinations of $\mathbf{e}_1$ and $\mathbf{e}_2$. So there exist scalars a, b, c, d such that

$$\begin{aligned} \mathbf{f}_1 &= \begin{pmatrix} 1 \\ 1 \end{pmatrix} = a \begin{pmatrix} 1 \\ 2 \end{pmatrix} + b \begin{pmatrix} -2 \\ 3 \end{pmatrix} \\ \mathbf{f}_2 &= \begin{pmatrix} 1 \\ -1 \end{pmatrix} = c \begin{pmatrix} 1 \\ 2 \end{pmatrix} + d \begin{pmatrix} 1 \\ 2 \end{pmatrix} \end{aligned}$$

We won't worry now about the precise values of a, b, c, d , since you can easily solve for them.

Definition: The *change of basis matrix* from E to F is

$$P = \begin{pmatrix} a & c \\ b & d \end{pmatrix}.$$

Note that this is the transpose of what you might think it should be; this is because we're doing column operations, and it's the first column of P which takes linear combinations of the columns of E and replaces the first column of E with the first column of F , and so on. In matrix form, we have

$$F = E \cdot P$$

and, of course, $E = F \cdot P^{-1}$.

Exercise: Find a, b, c, d and the change of basis matrix from E to F .

Given the change of basis matrix, we can figure out everything else we need to know.

- Suppose $\mathbf{v}$ has the known coordinates $\mathbf{v}_e$ in the basis E , and $F = E \cdot P$. Then

$$\mathbf{v} = E \cdot \mathbf{v}_e = F \cdot P^{-1} \mathbf{v}_e = F \cdot \mathbf{v}_f.$$

Remember that the coordinate vector is unique. This means that

$$\mathbf{v}_f = P^{-1} \mathbf{v}_e.$$

If P changes the basis from E to F , then P^{-1} changes the coordinates from $\mathbf{v}_e$ to $\mathbf{v}_f$ ⁴.

Compare this with the example at the end of the first section.

- For *any* nonsingular matrix P , the following holds:

$$\mathbf{v} = E \cdot \mathbf{v}_e = E \cdot P \cdot P^{-1} \cdot \mathbf{v}_e = G \cdot \mathbf{v}_g,$$

where P is the change of basis matrix from E to G : $G = E \cdot P$, and $P^{-1} \cdot \mathbf{v}_e = \mathbf{v}_g$ are the coordinates of the vector $\mathbf{v}$ in this basis.

- This notation is consistent with the standard basis as well. Since

$$\mathbf{e}_1 = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \text{ and } \mathbf{e}_2 = \begin{pmatrix} 0 \\ 1 \end{pmatrix},$$

we have $E = I_2$, and $\mathbf{v} = I_2 \cdot \mathbf{v}$

Remark: When we change from the standard basis to the basis $\{\mathbf{e}_1, \mathbf{e}_2\}$, the corresponding matrices are I (for the standard basis) and E . So according to what's just been shown, the change of basis matrix will be the matrix P which satisfies

$$E = I \cdot P.$$

In other words, the change of basis matrix in this case is just the matrix E .

⁴**Warning:** Some texts use P^{-1} instead of P for the change of basis matrix. This is a convention, but you need to check.

First example, cont'd We can write the system of differential equations in matrix form as

$$\dot{\mathbf{v}} = \begin{pmatrix} 1 & 3 \\ 3 & 1 \end{pmatrix} \mathbf{v} = A\mathbf{v}, \quad ((2))$$

where the dot indicates d/dt . We change from the standard basis to F via the matrix

$$F = \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}.$$

Then, according to what we've just worked out, we'll have

$$\mathbf{v}_f = F^{-1}\mathbf{v}, \text{ and taking derivatives, } \dot{\mathbf{v}}_f = F^{-1}\dot{\mathbf{v}}.$$

So using $\mathbf{v} = F\mathbf{v}_f$ and substituting into (2), we find

$$F\dot{\mathbf{v}}_f = AF\mathbf{v}_f, \text{ or } \dot{\mathbf{v}}_f = F^{-1}AF\mathbf{v}_f.$$

Now an easy computation shows that

$$F^{-1}AF = \begin{pmatrix} 4 & 0 \\ 0 & -2 \end{pmatrix},$$

and in the new coordinates, we have the system

$$\begin{aligned} \dot{v}_{f1} &= 4v_{f1} \\ \dot{v}_{f2} &= -2v_{f2} \end{aligned}$$

In the new coordinates, the system is now *decoupled* and easily solved to give

$$\begin{aligned} v_{f1} &= c_1 e^{4t} \\ v_{f2} &= c_2 e^{-2t}, \end{aligned}$$

where c_1, c_2 are arbitrary constants of integration. We can now transform back to the original (standard) basis to get the solution in the original coordinates:

$$\mathbf{v} = F\mathbf{v}_f = \begin{pmatrix} v_1 \\ v_2 \end{pmatrix} = \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \begin{pmatrix} c_1 e^{4t} \\ c_2 e^{-2t} \end{pmatrix} = \begin{pmatrix} c_1 e^{4t} + c_2 e^{-2t} \\ c_1 e^{4t} - c_2 e^{-2t} \end{pmatrix}.$$

A reasonable question at this point is "How does one come up with this new basis F ? It clearly was not chosen at random. The answer has to do with the eigenvalues and eigenvectors of the coefficient matrix of the differential equation, namely the matrix

$$A = \begin{pmatrix} 1 & 3 \\ 3 & 1 \end{pmatrix}.$$

All of which brings us to the subject of the next lecture.

15 Matrices and Linear transformations

We have been thinking of matrices in connection with solutions to linear systems of equations like $A\mathbf{x} = \mathbf{b}$. It is time to broaden our horizons a bit and start thinking of matrices as *functions*. In particular, if A is $m \times n$, we can use A to define a function $\mathbf{f}_A$ from $\mathbb{R}^n$ to $\mathbb{R}^m$ which sends $\mathbf{v} \in \mathbb{R}^n$ to $A\mathbf{v} \in \mathbb{R}^m$. That is, $\mathbf{f}_A(\mathbf{v}) = A\mathbf{v}$.

Example: Let

$$A_{2 \times 3} = \begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{pmatrix}.$$

If

$$\mathbf{v} = \begin{pmatrix} x \\ y \\ z \end{pmatrix} \in \mathbb{R}^3,$$

then

$$\mathbf{f}_A(\mathbf{v}) = A\mathbf{v} = \begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{pmatrix} \begin{pmatrix} x \\ y \\ z \end{pmatrix} = \begin{pmatrix} x + 2y + 3z \\ 4x + 5y + 6z \end{pmatrix}$$

sends the vector $\mathbf{v} \in \mathbb{R}^3$ to $A\mathbf{v} \in \mathbb{R}^2$.

Definition: A function $\mathbf{f} : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is said to be *linear* if

- $\mathbf{f}(\mathbf{v}_1 + \mathbf{v}_2) = \mathbf{f}(\mathbf{v}_1) + \mathbf{f}(\mathbf{v}_2)$, and
- $\mathbf{f}(c\mathbf{v}) = c\mathbf{f}(\mathbf{v})$ for all $\mathbf{v}_1, \mathbf{v}_2 \in \mathbb{R}^n$ and for all scalars c .

A linear function $\mathbf{f}$ is also known as a *linear transformation*.

Examples:

- Define $f : \mathbb{R}^3 \rightarrow \mathbb{R}$ by

$$f \begin{pmatrix} x \\ y \\ z \end{pmatrix} = 3x - 2y + z.$$

Then f is linear because for any

$$\mathbf{v}_1 = \begin{pmatrix} x_1 \\ y_1 \\ z_1 \end{pmatrix}, \quad \text{and } \mathbf{v}_2 = \begin{pmatrix} x_2 \\ y_2 \\ z_2 \end{pmatrix},$$

we have

$$f(\mathbf{v}_1 + \mathbf{v}_2) = f \begin{pmatrix} x_1 + x_2 \\ y_1 + y_2 \\ z_1 + z_2 \end{pmatrix} = 3(x_1 + x_2) - 2(y_1 + y_2) + (z_1 + z_2).$$

And the right hand side can be rewritten as $(3x_1 - 2y_1 + z_1) + (3x_2 - 2y_2 + z_2)$, which is the same as $f(\mathbf{v}_1) + f(\mathbf{v}_2)$. So the first property holds. So does the second, since $f(c\mathbf{v}) = 3cx - 2cy + cz = c(3x - 2y + z) = cf(\mathbf{v})$.

- Notice that the function f is actually f_A for the right A : if $A_{1 \times 3} = (3, -2, 1)$, then $f(\mathbf{v}) = A\mathbf{v}$.
- If $A_{m \times n}$ is a matrix, then $\mathbf{f}_A : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is a linear transformation because $\mathbf{f}_A(\mathbf{v}_1 + \mathbf{v}_2) = A(\mathbf{v}_1 + \mathbf{v}_2) = A\mathbf{v}_1 + A\mathbf{v}_2 = \mathbf{f}_A(\mathbf{v}_1) + \mathbf{f}_A(\mathbf{v}_2)$. And $A(c\mathbf{v}) = cA\mathbf{v} \Rightarrow \mathbf{f}_A(c\mathbf{v}) = c\mathbf{f}_A(\mathbf{v})$. (These are two fundamental properties of matrix multiplication.)
- Although we don't give the proof, it can be shown that *any* linear transformation can be written as $\mathbf{f}_A$ for a suitable matrix A .
- The derivative (see Lecture 9) is a linear transformation. $D\mathbf{f}(\mathbf{a})$ is the linear approximation to $\mathbf{f}(\mathbf{x}) - \mathbf{f}(\mathbf{a})$.
- There are many other examples of linear transformations; some of the most interesting ones do *not* go from $\mathbb{R}^n$ to $\mathbb{R}^m$:

1. If f and g are differentiable functions, then

$$\frac{d}{dx}(f + g) = \frac{df}{dx} + \frac{dg}{dx}, \text{ and } \frac{d}{dx}(cf) = c \frac{df}{dx}.$$

Thus the function $\mathcal{D}(f) = df/dx$ is linear.

2. If f is continuous, then we can define

$$If(x) = \int_0^x f(s) \, ds,$$

and I is linear, by well-known properties of the integral.

3. The Laplace operator, Δ , defined before, is linear.
4. Let y be twice continuously differentiable and define

$$L(y) = y'' - 2y' - 3y.$$

Then L is linear, as you can (and should!) verify.

Linear transformations acting on functions, like the above, are generally known as *linear operators*. They're a bit more complicated than matrix multiplication operators, but they have the same essential property of linearity.

Exercises:

1. Give an example of a function from $\mathbb{R}^2$ to itself which is not linear.
2. Identify all the linear transformations from $\mathbb{R}$ to $\mathbb{R}$.
3. If $\mathbf{f} : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is linear then

$$\text{Ker}(\mathbf{f}) := \{\mathbf{v} \in \mathbb{R}^n \text{ such that } \mathbf{f}(\mathbf{v}) = \mathbf{0}\}$$

is a subspace of $\mathbb{R}^n$, called the *kernel* of $\mathbf{f}$.

4. If $\mathbf{f} : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is linear, then

$$\text{Range}(\mathbf{f}) = \{\mathbf{y} \in \mathbb{R}^m \text{ such that } \mathbf{y} = \mathbf{f}(\mathbf{v}) \text{ for some } \mathbf{v}\}$$

is a subspace of $\mathbb{R}^m$ called the *range* of $\mathbf{f}$.

Everything we've been doing regarding the solution of linear systems of equations can be recast in the framework of linear transformations. In particular, if $\mathbf{f}_A$ is multiplication by some matrix A , then the range of $\mathbf{f}_A$ is just the set of all $\mathbf{y}$ such that the linear system $A\mathbf{v} = \mathbf{y}$ has a solution (i.e., it's the column space of A). And the kernel of $\mathbf{f}_A$ is the set of all solutions to the homogeneous equation $A\mathbf{v} = \mathbf{0}$.

15.1 The rank-nullity theorem - version 2

Recall that for $A_{m \times n}$, we have $n = N(A) + R(A)$. Now think of A as the linear transformation $\mathbf{f}_A : \mathbb{R}^n \rightarrow \mathbb{R}^m$. The *domain* of $\mathbf{f}_A$ is $\mathbb{R}^n$; $\text{Ker}(\mathbf{f}_A)$ is the null space of A , and $\text{Range}(\mathbf{f}_A)$ is the column space of A . We can therefore restate the rank-nullity theorem as the

Dimension theorem: Let $\mathbf{f}_A : \mathbb{R}^n \rightarrow \mathbb{R}^m$. Then

$$\dim(\text{domain}(\mathbf{f}_A)) = \dim(\text{Range}(\mathbf{f}_A)) + \dim(\text{Null}(\mathbf{f}_A)).$$

15.2 Choosing a useful basis for A

We now want to study square matrices, regarding an $n \times n$ matrix A as a linear transformation from $\mathbb{R}^n$ to itself. We'll just write $A\mathbf{v}$ for $\mathbf{f}_A(\mathbf{v})$ to simplify the notation, and to keep things really simple, we'll just talk about 2×2 matrices – all the problems that exist in higher dimensions are present in $\mathbb{R}^2$.

There are several questions that present themselves:

- Can we visualize the linear transformation $\mathbf{x} \rightarrow A\mathbf{x}$? One thing we *can't* do in general is draw a graph! Why not?
- Connected with the first question is: can we choose a better coordinate system in which to view the problem?

The answer is not an unequivocal "yes" to either of these, but we can generally do some useful things.

To pick up at the end of the last lecture, note that when we write $\mathbf{f}_A(\mathbf{v}) = \mathbf{y} = A\mathbf{v}$, we are actually using the coordinate vector of $\mathbf{v}$ in the standard basis. Suppose we change to some other basis $\{\mathbf{e}_1, \mathbf{e}_2\}$ using the invertible matrix E . Then we can rewrite the equation in the new coordinates and basis:

We have $\mathbf{v} = E\mathbf{v}_e$, and $\mathbf{y} = E\mathbf{y}_e$, so

$$\begin{aligned}\mathbf{y} &= A\mathbf{v} \\ E\mathbf{y}_e &= AE\mathbf{v}_e, \text{ and} \\ \mathbf{y}_e &= E^{-1}AE\mathbf{v}_e\end{aligned}$$

That is, the matrix equation $\mathbf{y} = A\mathbf{v}$ is given in the new basis by the equation

$$\mathbf{y}_e = E^{-1}AE\mathbf{v}_e.$$

Definition: The matrix $E^{-1}AE$ will be denoted by A_e and called the *matrix of the linear transformation in the basis E* .

We can now restate the second question: Can we find a nonsingular matrix E so that $E^{-1}AE$ is particularly useful?

Definition: The matrix A is *diagonal* if the only nonzero entries lie on the main diagonal. That is, $a_{ij} = 0$ if $i \neq j$.

Example:

$$A = \begin{pmatrix} 4 & 0 \\ 0 & -3 \end{pmatrix}$$

is diagonal. This is useful because we can (partially) visualize the linear transformation corresponding to multiplication by A : a vector $\mathbf{v}$ lying along the first coordinate axis is mapped to $4\mathbf{v}$, a multiple of itself. A vector $\mathbf{w}$ lying along the second coordinate axis is also mapped to a multiple of itself: $A\mathbf{w} = -3\mathbf{w}$. It's length is tripled, and its direction is

reversed. An arbitrary vector $(a, b)^t$ is a linear combination of the basis vectors, and it's mapped to $(4a, -3b)^t$.

It turns out that we can find vectors like $\mathbf{v}$ and $\mathbf{w}$, which are mapped to multiples of themselves, *without* first finding the matrix E . This is the subject of the next few sections.

15.3 Eigenvalues and eigenvectors

Definitions: If a vector $\mathbf{v} \neq \mathbf{0}$ satisfies the equation $A\mathbf{v} = \lambda\mathbf{v}$, for some real number λ , then λ is said to be an *eigenvalue of the matrix A* , and $\mathbf{v}$ is said to be an *eigenvector of A corresponding to λ* .

Example: If

$$A = \begin{pmatrix} 2 & 3 \\ 3 & 2 \end{pmatrix}, \text{ and } \mathbf{v} = \begin{pmatrix} 1 \\ 1 \end{pmatrix},$$

then

$$A\mathbf{v} = \begin{pmatrix} 5 \\ 5 \end{pmatrix} = 5\mathbf{v}.$$

So $\lambda = 5$ is an eigenvalue of A , and $\mathbf{v}$ an eigenvector corresponding to this eigenvalue.

Remark: Note that the definition of eigenvector *requires* that $\mathbf{v} \neq \mathbf{0}$. The reason for this is that if $\mathbf{v} = \mathbf{0}$ were allowed, then any number λ would be an eigenvalue since the statement $A\mathbf{0} = \lambda\mathbf{0}$ holds for any λ . On the other hand, we *can* have $\lambda = 0$, and $\mathbf{v} \neq \mathbf{0}$. See the exercise below.

Exercises:

1. Show that

$$\begin{pmatrix} 1 \\ -1 \end{pmatrix}$$

is also an eigenvector of the matrix A above. What's the eigenvalue?

2. Eigenvectors are not unique. Show that if $\mathbf{v}$ is an eigenvector for A , then so is $c\mathbf{v}$, for any real number $c \neq 0$.
3. Suppose λ is an eigenvalue of A .

Definition:

$$E_\lambda = \{\mathbf{v} \in \mathbb{R}^n \text{ such that } A\mathbf{v} = \lambda\mathbf{v}\}$$

is called the *eigenspace of A corresponding to the eigenvalue λ* .

Show that E_λ is a subspace of $\mathbb{R}^n$. (N.b: the definition of E_λ does not require $\mathbf{v}$ to be an eigenvector of A , so $\mathbf{v} = \mathbf{0}$ is allowed; otherwise, it wouldn't be a subspace.)

4. $E_0 = \text{Ker}(\mathbf{f}_A)$ is just the null space of the matrix A .

Example: The matrix

$$A = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} = \begin{pmatrix} \cos(\pi/2) & -\sin(\pi/2) \\ \sin(\pi/2) & \cos(\pi/2) \end{pmatrix}$$

represents a counterclockwise rotation through the angle $\pi/2$. Apart from $\mathbf{0}$, there is no vector which is mapped by A to a multiple of itself. So not every matrix has eigenvectors.

Exercise: What are the eigenvalues of this matrix?

16 Computations with eigenvalues and eigenvectors

How do we find the eigenvalues and eigenvectors of a matrix A ?

Suppose $\mathbf{v} \neq \mathbf{0}$ is an eigenvector. Then for some $\lambda \in \mathbb{R}$, $A\mathbf{v} = \lambda\mathbf{v}$. Then

$$\begin{aligned} A\mathbf{v} - \lambda\mathbf{v} &= \mathbf{0}, \text{ or, equivalently} \\ (A - \lambda I)\mathbf{v} &= \mathbf{0}. \end{aligned}$$

So $\mathbf{v}$ is a nontrivial solution to the homogeneous system of equations determined by the square matrix $A - \lambda I$. This can only happen if $\det(A - \lambda I) = 0$. On the other hand, *if* λ is a real number such that $\det(A - \lambda I) = 0$, this means exactly that there's a nontrivial solution to $(A - \lambda I)\mathbf{v} = \mathbf{0}$. So λ is an eigenvalue, and $\mathbf{v} \neq \mathbf{0}$ is an eigenvector. Summarizing, we have the

Theorem: λ is an eigenvalue of A if and only if $\det(A - \lambda I) = 0$.

For a 2×2 matrix

$$A = \begin{pmatrix} a & b \\ c & d \end{pmatrix},$$

we compute

$$\det(A - \lambda I) = \det \begin{pmatrix} a - \lambda & b \\ c & d - \lambda \end{pmatrix} = \lambda^2 - (a + d)\lambda + (ad - bc).$$

This polynomial of degree 2 is called the *characteristic polynomial* of the matrix A , and is denoted by $p_A(\lambda)$. By the above theorem, the eigenvalues of A are just the roots of the characteristic polynomial. The equation for the roots, $p_A(\lambda) = 0$, is called the *characteristic equation of A* .

Example: If

$$A = \begin{pmatrix} 1 & 3 \\ 3 & 1 \end{pmatrix}.$$

Then

$$A - \lambda I = \begin{pmatrix} 1 - \lambda & 3 \\ 3 & 1 - \lambda \end{pmatrix}, \text{ and } p_A(\lambda) = (1 - \lambda)^2 - 9 = \lambda^2 - 2\lambda - 8.$$

This factors as $p_A(\lambda) = (\lambda - 4)(\lambda + 2)$, so there are two eigenvalues: $\lambda_1 = 4$, and $\lambda_2 = -2$.

We should be able to find an eigenvector for each of these eigenvalues. To do so, we must find a nontrivial solution to the corresponding homogeneous equation

$(A - \lambda I)\mathbf{v} = \mathbf{0}$. For $\lambda_1 = 4$, we have the homogeneous system

$$\begin{pmatrix} 1-4 & 3 \\ 3 & 1-4 \end{pmatrix} \mathbf{v} = \begin{pmatrix} -3 & 3 \\ 3 & -3 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}.$$

This leads to the two equations $-3x_1 + 3x_2 = 0$, and $3x_1 - 3x_2 = 0$. Notice that the first equation is a multiple of the second, so there's really only one equation to solve.

Exercise: What property of the matrix $A - \lambda I$ guarantees that one of these equations will be a multiple of the other?

The general solution to the homogeneous system can be written as

$$\begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = c \begin{pmatrix} 1 \\ 1 \end{pmatrix}, \text{ where } c \text{ is arbitrary.}$$

This one-dimensional subspace of $\mathbb{R}^2$ is what we called E_4 in the last section.

We get an eigenvector by choosing any nonzero element of E_4 . Taking $c = 1$ gives the eigenvector

$$\mathbf{v}_1 = \begin{pmatrix} 1 \\ 1 \end{pmatrix}$$

Exercises:

1. Find the subspace E_{-2} and show that

$$\mathbf{v}_2 = \begin{pmatrix} 1 \\ -1 \end{pmatrix}$$

is an eigenvector corresponding to $\lambda_2 = -2$.

2. Find the eigenvalues and corresponding eigenvectors of the matrix

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

3. Same question for the matrix

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

16.1 Some observations

What are the possibilities for the characteristic polynomial p_A ? It's of degree 2, so there are 3 cases:

1. The two roots are real and distinct: $\lambda_1 \neq \lambda_2$, $\lambda_1, \lambda_2 \in \mathbb{R}$. We just worked out an example of this.
2. The roots are complex conjugates of one another: $\lambda_1 = a + ib$, $\lambda_2 = a - ib$.

Example:

$$A = \begin{pmatrix} 2 & 3 \\ -3 & 2 \end{pmatrix}.$$

Here, $p_A(\lambda) = \lambda^2 - 4\lambda + 13 = 0$ has the two roots $\lambda_{\pm} = 2 \pm 3i$. Now there's certainly no real vector $\mathbf{v}$ with the property that $A\mathbf{v} = (2 + 3i)\mathbf{v}$, so there are no eigenvectors in the usual sense. But there are *complex* eigenvectors corresponding to the complex eigenvalues. For example, if

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

$p_A(\lambda) = \lambda^2 + 1$ has the complex eigenvalues $\lambda_{\pm} = \pm i$. You can easily check that $A\mathbf{v} = i\mathbf{v}$, where

$$\mathbf{v} = \begin{pmatrix} i \\ 1 \end{pmatrix}.$$

We won't worry about complex eigenvectors in this course.

3. $p_A(\lambda)$ has a repeated root. An example is

$$A = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} = I_2.$$

Here $p_A(\lambda) = (1 - \lambda)^2$ and $\lambda = 1$ is the only eigenvalue. The matrix $A - \lambda I$ is the zero matrix. So there are no restrictions on the components of the eigenvectors. Any nonzero vector in $\mathbb{R}^2$ is an eigenvector corresponding to this eigenvalue.

But for

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

as you saw in the exercise above, we also have $p_A(\lambda) = (1 - \lambda)^2$. In this case, though, there is just a one-dimensional eigenspace.

16.2 Diagonalizable matrices

Example: In the preceding lecture, we showed that, for the matrix

$$A = \begin{pmatrix} 1 & 3 \\ 3 & 1 \end{pmatrix},$$

if we change the basis using

$$E = (\mathbf{e}_1 : \mathbf{e}_2) = \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix},$$

then, in this new basis, we have

$$A_E = E^{-1}AE = \begin{pmatrix} 4 & 0 \\ 0 & -2 \end{pmatrix}.$$

The matrix A_E is called a *diagonal* matrix; the only non-zero entries lie on the main diagonal.

Definition: Let A be $n \times n$. We say that A is *diagonalizable* if there exists a basis $\{\mathbf{e}_1, \dots, \mathbf{e}_n\}$ of $\mathbb{R}^n$, with corresponding change of basis matrix $E = (\mathbf{e}_1 : \dots : \mathbf{e}_n)$ such that

$$A_E = E^{-1}AE$$

is diagonal.

In the example, our matrix E has the form $E = (\mathbf{e}_1 : \mathbf{e}_2)$, where the two columns are two eigenvectors of A corresponding to the eigenvalues $\lambda = 4$, and $\lambda = 2$. In fact, this is the general recipe:

Theorem: The matrix A is diagonalizable $\iff$ there is a basis for $\mathbb{R}^n$ consisting of eigenvectors of A .

Proof: Suppose $\{\mathbf{e}_1, \dots, \mathbf{e}_n\}$ is a basis for $\mathbb{R}^n$ with the property that $A\mathbf{e}_j = \lambda_j\mathbf{e}_j$, $1 \leq j \leq n$. Form the matrix $E = (\mathbf{e}_1 : \mathbf{e}_2 : \dots : \mathbf{e}_n)$. We have

$$\begin{aligned} AE &= (A\mathbf{e}_1 : A\mathbf{e}_2 : \dots : A\mathbf{e}_n) \\ &= (\lambda_1\mathbf{e}_1 : \lambda_2\mathbf{e}_2 : \dots : \lambda_n\mathbf{e}_n) \\ &= ED, \end{aligned}$$

where $D = \text{Diag}(\lambda_1, \lambda_2, \dots, \lambda_n)$. Evidently, $A_E = D$ and A is diagonalizable. Conversely, if A is diagonalizable, then the columns of the matrix which diagonalizes A are the required basis of eigenvectors.

So, in $\mathbb{R}^2$, a matrix A can be diagonalized $\iff$ we can find two linearly independent eigenvectors. (To *diagonalize* a matrix A means to find a matrix E such that $E^{-1}AE$ is diagonal.)

Examples:

- Diagonalize the matrix

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

Solution: From the previous exercise set, we have $\lambda_1 = 3$, $\lambda_2 = -2$ with corresponding eigenvectors

$$\mathbf{v}_1 = \begin{pmatrix} 1 \\ 1 \end{pmatrix}, \quad \mathbf{v}_2 = \begin{pmatrix} -2 \\ 3 \end{pmatrix}.$$

We form the matrix

$$E = (\mathbf{v}_1 : \mathbf{v}_2) = \begin{pmatrix} 1 & -2 \\ 1 & 3 \end{pmatrix}, \text{ with } E^{-1} = (1/5) \begin{pmatrix} 3 & 2 \\ -1 & 1 \end{pmatrix},$$

and check that $E^{-1}AE = \text{Diag}(3, -2)$. Of course, we don't really need to check: the result is guaranteed by the theorem above!

- The matrix

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

has only the one-dimensional eigenspace spanned by the eigenvector

$$\begin{pmatrix} 1 \\ 0 \end{pmatrix}.$$

There is no basis of $\mathbb{R}^2$ consisting of eigenvectors of A , so *this matrix cannot be diagonalized*.

This can only happen in the case of repeated or complex roots because of the following

Theorem: If λ_1 and λ_2 are distinct eigenvalues of A , with corresponding eigenvectors $\mathbf{v}_1$, $\mathbf{v}_2$, then $\{\mathbf{v}_1, \mathbf{v}_2\}$ are linearly independent.

Proof: Suppose $c_1\mathbf{v}_1 + c_2\mathbf{v}_2 = \mathbf{0}$, where one of the coefficients, say c_1 is nonzero. Then $\mathbf{v}_1 = \alpha\mathbf{v}_2$, for some $\alpha \neq 0$. (If $\alpha = 0$, then $\mathbf{v}_1 = \mathbf{0}$ and $\mathbf{v}_1$ by definition is not an eigenvector.) Multiplying both sides on the left by A gives

$$A\mathbf{v}_1 = \lambda_1\mathbf{v}_1 = \alpha A\mathbf{v}_2 = \alpha\lambda_2\mathbf{v}_2.$$

On the other hand, multiplying the same equation by λ_1 and then subtracting the two equations gives

$$\mathbf{0} = \alpha(\lambda_2 - \lambda_1)\mathbf{v}_2$$

which is impossible, since neither α nor $(\lambda_1 - \lambda_2)$ nor $\mathbf{v}_2 = \mathbf{0}$.

It follows that if $A_{2 \times 2}$ has two distinct eigenvalues, then it has two linearly independent eigenvectors and can be diagonalized. In a similar way, if $A_{n \times n}$ has n distinct eigenvalues, it is diagonalizable.

Exercises:

1. Find the eigenvalues and eigenvectors of the matrix

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

Form the matrix E and verify that $E^{-1}AE$ is diagonal.

2. List the two reasons a matrix may fail to be diagonalizable. Give examples of both cases.
3. An arbitrary 2×2 symmetric matrix ($A = A^t$) has the form

$$A = \begin{pmatrix} a & b \\ b & c \end{pmatrix},$$

where a, b, c can be any real numbers. Show that A always has real eigenvalues. When are the two eigenvalues equal?

17 Inner products

Up until now, we have only examined the properties of vectors and matrices in $\mathbb{R}^n$. But normally, when we think of $\mathbb{R}^n$, we're really thinking of n-dimensional Euclidean space - that is, $\mathbb{R}^n$ together with the dot product. Once we have the dot product, or more generally an "inner product" on $\mathbb{R}^n$, we can talk about angles, lengths, distances, etc.

Definition: An *inner product* on $\mathbb{R}^n$ is a function

$$(\cdot, \cdot) : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$$

with the following properties:

1. It is *bilinear*, meaning it's linear in each argument:

- $(c_1\mathbf{x}_1 + c_2\mathbf{x}_2, \mathbf{y}) = c_1(\mathbf{x}_1, \mathbf{y}) + c_2(\mathbf{x}_2, \mathbf{y})$, and
- $(\mathbf{x}, c_1\mathbf{y}_1 + c_2\mathbf{y}_2) = c_1(\mathbf{x}, \mathbf{y}_1) + c_2(\mathbf{x}, \mathbf{y}_2)$.

2. It is *symmetric*: $(\mathbf{x}, \mathbf{y}) = (\mathbf{y}, \mathbf{x})$, $\forall \mathbf{x}, \mathbf{y} \in \mathbb{R}^n$.

3. It is *non-degenerate*: If $(\mathbf{x}, \mathbf{y}) = 0, \forall \mathbf{y} \in \mathbb{R}^n$, then $\mathbf{x} = \mathbf{0}$.

The inner product is said to be *positive definite* if, in addition

4. $(\mathbf{x}, \mathbf{x}) > 0$ whenever $\mathbf{x} \neq \mathbf{0}$.

Examples of inner products

- The *dot product* in $\mathbb{R}^n$ given in the standard basis by

$$(\mathbf{x}, \mathbf{y}) = \mathbf{x} \cdot \mathbf{y} = x_1y_1 + x_2y_2 + \cdots + x_ny_n$$

The dot product is positive definite - all four of the properties above hold (**exercise**).

$\mathbb{R}^n$ with the dot product as an inner product is called *n-dimensional Euclidean space*, and is denoted $\mathbb{E}^n$.

- In $\mathbb{R}^4$, with coordinates t, x, y, z , we can define

$$(\mathbf{v}_1, \mathbf{v}_2) = t_1 t_2 - x_1 x_2 - y_1 y_2 - z_1 z_2.$$

This is an inner product too. But for $\mathbf{x} = (1, 1, 0, 0)^t$, we have $(\mathbf{x}, \mathbf{x}) = 0$, so it's *not* positive definite. $\mathbb{R}^4$ with this inner product is called *Minkowski space*. It is the spacetime of special relativity (invented by Einstein in 1905, and made into a nice geometric space by Minkowski several year later). It is denoted $\mathbb{M}^4$, and if time permits, we'll look more closely at this space later in the course.

- Let G be an $n \times n$ symmetric matrix ($G = G^t$), with $\det(G) \neq 0$. Define

$$(\mathbf{x}, \mathbf{y})_G = \mathbf{x}^t G \mathbf{y}.$$

It is not difficult to verify that this satisfies the properties in the definition. For example, if $(\mathbf{x}, \mathbf{y})_G = \mathbf{x}^t G \mathbf{y} = 0 \forall \mathbf{y}$, then $\mathbf{x}^t G = \mathbf{0}$, because if we write $\mathbf{x}^t G$ as the row vector $(a_1, a_2, \dots, a_n)$, then $\mathbf{x}^t G \mathbf{e}_1 = 0 \Rightarrow a_1 = 0$, $\mathbf{x}^t G \mathbf{e}_2 = 0 \Rightarrow a_2 = 0$, etc. So all the components of $\mathbf{x}^t G$ are 0 and hence $\mathbf{x}^t G = \mathbf{0}$. Now taking transposes, we find that $G^t \mathbf{x} = G \mathbf{x} = \mathbf{0}$. Since G is nonsingular by definition, this means that $\mathbf{x} = \mathbf{0}$, (otherwise the homogeneous system $G \mathbf{x} = \mathbf{0}$ would have non-trivial solutions and G would be singular) and the inner product is non-degenerate.

In fact, any inner product on $\mathbb{R}^n$ can be written in this form for a suitable matrix G :

- * $\mathbf{x} \bullet \mathbf{y} = \mathbf{x}^t G \mathbf{y}$ with $G = I$. For instance, if

$$\mathbf{x} = \begin{pmatrix} 3 \\ 2 \\ 1 \end{pmatrix}, \text{ and } \mathbf{y} = \begin{pmatrix} -1 \\ 2 \\ 4 \end{pmatrix},$$

then

$$\mathbf{x} \bullet \mathbf{y} = \mathbf{x}^t I \mathbf{y} = \mathbf{x}^t \mathbf{y} = (3, 2, 1) \begin{pmatrix} -1 \\ 2 \\ 4 \end{pmatrix} = -3 + 4 + 4 = 5$$

- * The Minkowski inner product has the form $\mathbf{x}^t G \mathbf{y}$ with $G = \text{Diag}(1, -1, -1, -1)$

Remark: If we replace $\mathbf{y}$ by $\mathbf{x}$ in all of the above, we get what's called a *quadratic form*, which is a function of just the single vector variable $\mathbf{x}$. Its general form is $\mathbf{x}^t G \mathbf{x}$. It's no longer linear in $\mathbf{x}$, but quadratic (hence the name).

Exercise:** Show that under a change of basis matrix E , the matrix G of the inner product becomes $G_E = E^t G E$. For instance, if $G = I$, so that $\mathbf{x} \bullet \mathbf{y} = \mathbf{x}^t I \mathbf{y}$, and

$$E = \begin{pmatrix} 1 & 3 \\ 3 & 1 \end{pmatrix}, \text{ then } \mathbf{x} \bullet \mathbf{y} = \mathbf{x}_E^t G_E \mathbf{y}_E, \text{ with } G_E = \begin{pmatrix} 10 & 4 \\ 4 & 10 \end{pmatrix}.$$

This is different from the way in which an ordinary matrix (which can be viewed as a linear transformation) behaves. Thus the matrix representing an inner product is a different object from that representing a linear transformation.

17.1 Euclidean space

We now restrict attention to Euclidean space $\mathbb{E}^n$. We'll always be using the dot product, whether we write it as $\mathbf{x} \bullet \mathbf{y}$ or $(\mathbf{x}, \mathbf{y})$.

Definition: The *norm* of the vector $\mathbf{x}$ is defined by

$$\|\mathbf{x}\| = \sqrt{\mathbf{x} \bullet \mathbf{x}}.$$

In the standard coordinates, this is equal to

$$\|\mathbf{x}\| = \left(\sum_{i=1}^n x_i^2 \right)^{1/2}.$$

Example:

$$\text{If } \mathbf{x} = \begin{pmatrix} -2 \\ 4 \\ 1 \end{pmatrix}, \text{ then } \|\mathbf{x}\| = \sqrt{(-2)^2 + 4^2 + 1^2} = \sqrt{21}$$

Proposition:

- $||\mathbf{x}|| > 0$ if $\mathbf{x} \neq \mathbf{0}$.
- $||c\mathbf{x}|| = |c| ||\mathbf{x}||$, $\forall c \in \mathbb{R}$.

Proof: **Exercise**

As you know, $||\mathbf{x}||$ is the distance from the origin $\mathbf{0}$ to the point $\mathbf{x}$. Or it's the length of the vector $\mathbf{x}$. (Same thing.) The next few properties all follow from the **law of cosines**:

For a triangle with sides a, b , and c , and angles opposite these sides of A, B , and C ,

$$c^2 = a^2 + b^2 - 2ab \cos(C).$$

This reduces to Pythagoras' theorem if C is a right angle, of course. In the present context, we imagine two vectors $\mathbf{x}$ and $\mathbf{y}$ with their "tails" located at $\mathbf{0}$. The vector going from the tip of $\mathbf{x}$ to the tip of $\mathbf{y}$ is $\mathbf{x} - \mathbf{y}$. If θ is the angle between $\mathbf{x}$ and $\mathbf{y}$, then the law of cosines reads

$$||\mathbf{x} - \mathbf{y}||^2 = ||\mathbf{x}||^2 + ||\mathbf{y}||^2 - 2||\mathbf{x}|| ||\mathbf{y}|| \cos \theta. \quad (1)$$

On the other hand, from the definition of the norm, we have

$$\begin{aligned} ||\mathbf{x} - \mathbf{y}||^2 &= (\mathbf{x} - \mathbf{y}) \bullet (\mathbf{x} - \mathbf{y}) \\ &= \mathbf{x} \bullet \mathbf{x} - \mathbf{x} \bullet \mathbf{y} - \mathbf{y} \bullet \mathbf{x} + \mathbf{y} \bullet \mathbf{y} \text{ or} \\ ||\mathbf{x} - \mathbf{y}||^2 &= ||\mathbf{x}||^2 + ||\mathbf{y}||^2 - 2\mathbf{x} \bullet \mathbf{y} \end{aligned} \quad (2)$$

Comparing (1) and (2), we conclude that

$$\mathbf{x} \bullet \mathbf{y} = \cos \theta ||\mathbf{x}|| ||\mathbf{y}||, \text{ or } \cos \theta = \frac{\mathbf{x} \bullet \mathbf{y}}{||\mathbf{x}|| ||\mathbf{y}||} \quad (3)$$

Since $|\cos \theta| \leq 1$, taking absolute values we get

$$\text{Theorem: } \frac{|\mathbf{x} \bullet \mathbf{y}|}{||\mathbf{x}|| ||\mathbf{y}||} \leq 1, \text{ or } |\mathbf{x} \bullet \mathbf{y}| \leq ||\mathbf{x}|| ||\mathbf{y}|| \quad (4)$$

The inequality (4) is known as the *Cauchy-Schwarz* inequality.

Exercises:

1. Find the angle θ between the two vectors $\mathbf{v} = (1, 0, -1)^t$ and $(2, 1, 3)^t$.
2. When does $|\mathbf{x} \cdot \mathbf{y}| = \|\mathbf{x}\| \|\mathbf{y}\|$? What is θ when $\mathbf{x} \cdot \mathbf{y} = 0$?

Using the Cauchy-Schwarz inequality, we (i.e., you) can prove the *triangle inequality*:

Theorem: For all $\mathbf{x}, \mathbf{y}$, $\|\mathbf{x} + \mathbf{y}\| \leq \|\mathbf{x}\| + \|\mathbf{y}\|$.

Proof: **Exercise** (Expand the dot product $\|\mathbf{x} + \mathbf{y}\|^2 = (\mathbf{x} + \mathbf{y}) \cdot (\mathbf{x} + \mathbf{y})$, use the Cauchy-Schwarz inequality, and take the square root.)

Exercise: The triangle inequality as it's usually encountered in geometry courses states that, in $\triangle ABC$, the distance from A to B is $\leq$ the distance from A to C plus the distance from C to B . Is this the same thing?

18 Orthogonality and related notions

18.1 Orthogonality

Definition: Two vectors $\mathbf{x}$ and $\mathbf{y}$ are said to be *orthogonal* if $\mathbf{x} \bullet \mathbf{y} = 0$. (This is the fancy version of "perpendicular".)

Examples: The two vectors

$$\begin{pmatrix} 1 \\ -1 \\ 0 \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} 2 \\ 2 \\ 4 \end{pmatrix}$$

are orthogonal, since their dot product is $(2)(1) + (2)(-1) + (4)(0) = 0$. The standard basis vectors $\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3 \in \mathbb{R}^3$ are *mutually orthogonal*. That is $\mathbf{e}_i \bullet \mathbf{e}_j = 0$ whenever $i \neq j$. The vector $\mathbf{0}$ is orthogonal to everything.

Definition: A *unit vector* is a vector of length 1. If its length is 1, then the square of its length is also 1. So $\mathbf{v}$ is a unit vector if $\mathbf{v} \bullet \mathbf{v} = 1$.

If $\mathbf{w}$ is an arbitrary nonzero vector, then a *unit vector in the direction of $\mathbf{w}$* is obtained by multiplying $\mathbf{w}$ by $\|\mathbf{w}\|^{-1}$: $\hat{\mathbf{w}} = (1/\|\mathbf{w}\|)\mathbf{w}$ is a unit vector in the direction of $\mathbf{w}$. The caret mark over the vector will always be used to indicate a unit vector.

Examples: The standard basis vectors are all unit vectors. If

$$\mathbf{w} = \begin{pmatrix} 1 \\ 2 \\ 3 \end{pmatrix},$$

then a unit vector in the direction of $\mathbf{w}$ is

$$\hat{\mathbf{w}} = \frac{1}{\|\mathbf{w}\|} \mathbf{w} = \frac{1}{\sqrt{14}} \begin{pmatrix} 1 \\ 2 \\ 3 \end{pmatrix}.$$

The process of replacing a vector $\mathbf{w}$ by a unit vector in its direction is called *normalizing* the vector.

For an arbitrary nonzero vector in $\mathbb{R}^3$

$$\begin{pmatrix} x \\ y \\ z \end{pmatrix},$$

the corresponding unit vector is

$$\frac{1}{\sqrt{x^2 + y^2 + z^2}} \begin{pmatrix} x \\ y \\ z \end{pmatrix}$$

In physics and engineering courses, this particular vector is often denoted by $\hat{\mathbf{r}}$. For instance, the gravitational force on a particle of mass m sitting at $(x, y, z)^t$ due to a particle of mass M sitting at the origin is

$$\mathbf{F} = \frac{-GMm}{r^2} \hat{\mathbf{r}},$$

where $r^2 = x^2 + y^2 + z^2$.

18.2 Orthonormal bases

Although we know that any set of n linearly independent vectors in $\mathbb{R}^n$ can be used as a basis, there is a particularly nice collection of bases that we can use in Euclidean space.

Definition: A basis $\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n\}$ of $\mathbb{R}^n$ is said to be *orthonormal* if

1. $\mathbf{v}_i \bullet \mathbf{v}_j = 0$, whenever $i \neq j$. That is, they are mutually orthogonal, and
2. $\mathbf{v}_i \bullet \mathbf{v}_i = 1$ for all i . They are all unit vectors.

Examples: The standard basis is orthonormal. The basis

$$\left\{ \begin{pmatrix} 1 \\ 1 \end{pmatrix}, \begin{pmatrix} 1 \\ -1 \end{pmatrix} \right\}$$

is orthogonal, but not orthonormal. We can normalize these vectors to get the orthonormal basis

$$\left\{ \begin{pmatrix} 1/\sqrt{2} \\ 1/\sqrt{2} \end{pmatrix}, \begin{pmatrix} 1/\sqrt{2} \\ -1/\sqrt{2} \end{pmatrix} \right\}$$

You may recall that it's quite tedious to compute the coordinates of a vector $\mathbf{w}$ in an arbitrary basis. The advantage of using an orthonormal basis is

Theorem: Let $\{\mathbf{v}_1, \dots, \mathbf{v}_n\}$ be an orthonormal basis in $\mathbb{E}^n$. Let $\mathbf{w} \in \mathbb{E}^n$. Then

$$\mathbf{w} = (\mathbf{w} \bullet \mathbf{v}_1) \mathbf{v}_1 + (\mathbf{w} \bullet \mathbf{v}_2) \mathbf{v}_2 + \dots + (\mathbf{w} \bullet \mathbf{v}_n) \mathbf{v}_n.$$

That is, the i^{th} coordinate of $\mathbf{w}$ in this basis is given by $\mathbf{w} \bullet \mathbf{v}_i$, the dot product of $\mathbf{w}$ with the i^{th} basis vector.

Proof: Since we have a basis, there are unique numbers $c_1, \dots, c_n$ such that

$$\mathbf{w} = c_1 \mathbf{v}_1 + c_2 \mathbf{v}_2 + \dots + c_n \mathbf{v}_n.$$

Take the dot product of both sides of the equation with $\mathbf{v}_1$: using the linearity of the dot product, we get

$$\mathbf{v}_1 \bullet \mathbf{w} = c_1 (\mathbf{v}_1 \bullet \mathbf{v}_1) + c_2 (\mathbf{v}_1 \bullet \mathbf{v}_2) + \dots + c_n (\mathbf{v}_1 \bullet \mathbf{v}_n).$$

Since the basis is orthonormal, all these dot products vanish except for the first, and we have $(\mathbf{v}_1 \bullet \mathbf{w}) = c_1 (\mathbf{v}_1 \bullet \mathbf{v}_1) = c_1$. An identical argument holds for the general $\mathbf{v}_i$.

Example: Find the coordinates of the vector

$$\mathbf{w} = \begin{pmatrix} 2 \\ -3 \end{pmatrix}$$

in the basis

$$\{\mathbf{v}_1, \mathbf{v}_2\} = \left\{ \begin{pmatrix} 1/\sqrt{2} \\ 1/\sqrt{2} \end{pmatrix}, \begin{pmatrix} 1/\sqrt{2} \\ -1/\sqrt{2} \end{pmatrix} \right\}.$$

Solution: $\mathbf{w} \cdot \mathbf{v}_1 = 2/\sqrt{2} - 3/\sqrt{2} = -1/\sqrt{2}$, and $\mathbf{w} \cdot \mathbf{v}_2 = 2/\sqrt{2} + 3/\sqrt{2} = 5/\sqrt{2}$. So the coordinates of $\mathbf{w}$ in this basis are

$$\frac{1}{\sqrt{2}} \begin{pmatrix} -1 \\ 5 \end{pmatrix}.$$

Exercises:

1. Let

$$\{\mathbf{e}_1(\theta), \mathbf{e}_2(\theta)\} = \left\{ \begin{pmatrix} \cos \theta \\ \sin \theta \end{pmatrix}, \begin{pmatrix} -\sin \theta \\ \cos \theta \end{pmatrix} \right\}.$$

What's the relation between $\{\mathbf{e}_1(\theta), \mathbf{e}_2(\theta)\}$ and $\{\mathbf{i}, \mathbf{j}\} = \{\mathbf{e}_1(0), \mathbf{e}_2(0)\}$?

2. Let

$$\mathbf{v} = \begin{pmatrix} 2 \\ -3 \end{pmatrix}.$$


Find the coordinates of $\mathbf{v}$ in the basis $\{\mathbf{e}_1(\theta), \mathbf{e}_2(\theta)\}$

- By writing $\mathbf{v} = c_1 \mathbf{e}_1(\theta) + c_2 \mathbf{e}_2(\theta)$ and solving for c_1, c_2 .
- By using the theorem above.

18.3 Orthogonal projections

It is frequently useful to decompose a given vector $\mathbf{v}$ as $\mathbf{v} = \mathbf{v}_{||} + \mathbf{v}_{\perp}$, where $\mathbf{v}_{||}$ is parallel to a vector $\mathbf{w}$, and $\mathbf{v}_{\perp}$ is orthogonal to $\mathbf{w}$.

Example: Suppose a mass m is at the end of a rigid, massless rod (a pendulum, approximately), and the rod makes an angle θ with the vertical. The force acting on the pendulum is the gravitational force $-mg\mathbf{e}_2$. Since the pendulum is rigid, the force directed along the rod's direction doesn't do anything (i.e., doesn't cause the pendulum to move). Only the force orthogonal to the rod produces motion. The magnitude of the force parallel to the pendulum is $mg \cos \theta$, and the orthogonal force has magnitude $mg \sin \theta$. If the pendulum



The pendulum bob makes an angle θ with the vertical. The magnitude of the force (gravity) acting on the bob is mg .

The component of the force acting in the direction of motion of the pendulum has magnitude $mg \sin(\theta)$.

has length l , Newton's law ($\mathbf{F} = m\mathbf{a}$) reads

$$ml\ddot{\theta} = -mg \sin \theta,$$

or

$$\ddot{\theta} + \frac{g}{l} \sin \theta = 0.$$

This is the differential equation for the motion of the pendulum. For small angles, we have, approximately, $\sin \theta \approx \theta$, and the equation can be linearized to give

$$\ddot{\theta} + \omega^2 \theta = 0, \text{ where } \omega = \sqrt{\frac{g}{l}}.$$


18.4 Algorithm for the decomposition

We want to write $\mathbf{v} = \mathbf{v}_{\parallel} + \mathbf{v}_{\perp}$, where $\mathbf{v}_{\parallel}$ is in the direction of $\mathbf{w}$. See the figure. Suppose θ is the angle between $\mathbf{w}$ and $\mathbf{v}$. We assume for the moment that $0 \leq \theta \leq \pi/2$. Then

$$\|\mathbf{v}_{\parallel}\| = \|\mathbf{v}\| \cos \theta = \|\mathbf{v}\| \left(\frac{\mathbf{v} \cdot \mathbf{w}}{\|\mathbf{v}\| \|\mathbf{w}\|} \right) = \frac{\mathbf{v} \cdot \mathbf{w}}{\|\mathbf{w}\|},$$

or

$$\|\mathbf{v}_{\parallel}\| = \mathbf{v} \cdot \text{a unit vector in the direction of } \mathbf{w}$$



If the angle between $\mathbf{v}$ and $\mathbf{w}$ is θ , then the magnitude of the projection of $\mathbf{v}$ onto $\mathbf{w}$ is $||\mathbf{v}|| \cos(\theta)$.

And $\mathbf{v}_{||}$ is this number times a unit vector in the direction of $\mathbf{w}$:

$$\mathbf{v}_{||} = \frac{\mathbf{v} \cdot \mathbf{w}}{||\mathbf{w}||} \frac{\mathbf{w}}{||\mathbf{w}||} = \left(\frac{\mathbf{v} \cdot \mathbf{w}}{\mathbf{w} \cdot \mathbf{w}} \right) \mathbf{w}.$$

In other words, if $\widehat{\mathbf{w}} = (1/||\mathbf{w}||)\mathbf{w}$, then $\mathbf{v}_{||} = (\mathbf{v} \cdot \widehat{\mathbf{w}})\widehat{\mathbf{w}}$.

The vector $\mathbf{v}_{||}$ is called the *orthogonal projection* of $\mathbf{v}$ onto $\mathbf{w}$. The nonzero vector $\mathbf{w}$ also determines a 1-dimensional subspace, denoted W , consisting of all multiples of $\mathbf{w}$, and $\mathbf{v}_{||}$ is also known as *the orthogonal projection onto the subspace W* .

Since $\mathbf{v} = \mathbf{v}_{||} + \mathbf{v}_{\perp}$, we have

$$\mathbf{v}_{\perp} = \mathbf{v} - \mathbf{v}_{||}.$$

Example: Let

$$\mathbf{v} = \begin{pmatrix} 1 \\ -1 \\ 2 \end{pmatrix}, \text{ and } \mathbf{w} = \begin{pmatrix} 1 \\ 0 \\ 1 \end{pmatrix}.$$

Then $||\mathbf{w}|| = \sqrt{2}$, so $\widehat{\mathbf{w}} = (1/\sqrt{2})\mathbf{w}$, and

$$(\mathbf{v} \cdot \widehat{\mathbf{w}})\widehat{\mathbf{w}} = \begin{pmatrix} 3/2 \\ 0 \\ 3/2 \end{pmatrix}.$$

Then

$$\mathbf{v}_{\perp} = \mathbf{v} - \mathbf{v}_{||} = \begin{pmatrix} 1 \\ -1 \\ 2 \end{pmatrix} - \begin{pmatrix} 3/2 \\ 0 \\ 3/2 \end{pmatrix} = \begin{pmatrix} -1/2 \\ -1 \\ 1/2 \end{pmatrix}.$$

and you can easily check that $\mathbf{v}_{\parallel} \cdot \mathbf{v}_{\perp} = 0$.

Remark: Suppose that, in the above, $\pi/2 < \theta \leq \pi$, so the angle is not acute. In this case, $\cos \theta$ is negative, and $\cos \theta \|\mathbf{v}\|$ is not the length of $\mathbf{v}_{\parallel}$ (since it's negative, it can't be a length). It has to be interpreted as a signed length, since the correct projection points in the opposite direction from that of $\mathbf{w}$. In other words, the formula is correct, no matter what the value of θ .

Exercise: This refers to the pendulum figure. Suppose the mass is located at $(x, y) \in \mathbb{R}^2$. Find the unit vector parallel to the direction of the rod, say $\hat{\mathbf{r}}$, and a unit vector orthogonal to $\hat{\mathbf{r}}$, say $\hat{\theta}$, obtained by rotating $\hat{\mathbf{r}}$ counterclockwise through an angle $\pi/2$. Express these orthonormal vectors in terms of the angle θ . And show that $\mathbf{F} \cdot \hat{\theta} = -mg \sin \theta$ as claimed above.

19 Orthogonal projections and Gram-Schmidt

19.1 Orthogonal matrices

Suppose we take an orthonormal (o.n.) basis $\{\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_n\}$ of $\mathbb{R}^n$ and form the $n \times n$ matrix $E = (\mathbf{e}_1 \vdots \dots \vdots \mathbf{e}_n)$. Then

$$E^t E = \begin{pmatrix} \mathbf{e}_1^t \\ \mathbf{e}_2^t \\ \vdots \\ \mathbf{e}_n^t \end{pmatrix} (\mathbf{e}_1 \vdots \dots \vdots \mathbf{e}_n) = I_n,$$

because

$$(E^t E)_{ij} = \mathbf{e}_i^t \mathbf{e}_j = \mathbf{e}_i \bullet \mathbf{e}_j = \delta_{ij},$$

where δ_{ij} are the components of the identity matrix:

$$\delta_{ij} = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{if } i \neq j \end{cases}$$

Since $E^t E = I$, this means that $E^t = E^{-1}$.

Definition: A square matrix E such that $E^t = E^{-1}$ is called an *orthogonal matrix*.

Example:

$$\{\mathbf{e}_1, \mathbf{e}_2\} = \left\{ \begin{pmatrix} 1/\sqrt{2} \\ 1/\sqrt{2} \end{pmatrix}, \begin{pmatrix} 1/\sqrt{2} \\ -1/\sqrt{2} \end{pmatrix} \right\}$$

is an o.n. basis for $\mathbb{R}^2$. The corresponding matrix

$$E = (1/\sqrt{2}) \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$$

is easily verified to be orthogonal. Of course the identity matrix is also orthogonal. As a converse to the above, if E is an orthogonal matrix, the columns of E form an o.n. basis of $\mathbb{R}^n$.

Exercises:

- If E is orthogonal, so is E^t , so the rows of E also form an o.n. basis.
- If E and F are orthogonal and of the same dimension, then EF is orthogonal.
- Let

$$\{\mathbf{e}_1(\theta), \mathbf{e}_2(\theta)\} = \left\{ \begin{pmatrix} \cos \theta \\ \sin \theta \end{pmatrix}, \begin{pmatrix} -\sin \theta \\ \cos \theta \end{pmatrix} \right\}.$$

Let $R(\theta) = (\mathbf{e}_1(\theta) : \mathbf{e}_2(\theta))$. Show that $R(\theta)R(\tau) = R(\theta + \tau)$.

- If E and F are the two orthogonal matrices corresponding to two o.n. bases, then $F = EP$, where P is the change of basis matrix from E to F . Show that P is also orthogonal.

19.2 Construction of orthonormal bases

It is not obvious that any subspace V of $\mathbb{R}^n$ has an orthonormal basis, but it's true. Here we give an algorithm for constructing such a basis, starting from an arbitrary basis. This is called the *Gram-Schmidt* procedure. We'll do it first for a 2-dimensional subspace of $\mathbb{R}^3$, and then do it in general at the end:

Let V be a 2-dimensional subspace of $\mathbb{R}^3$, and let $\{\mathbf{f}_1, \mathbf{f}_2\}$ be a basis for V . We want to construct an o.n. basis $\{\mathbf{e}_1, \mathbf{e}_2\}$ for V .

- The first step is easy. We define $\mathbf{e}_1 = \frac{1}{\|\mathbf{f}_1\|}\mathbf{f}_1$ by normalizing $\mathbf{f}_1$.
- We now need a vector orthogonal to $\mathbf{e}_1$ which lies in the plane spanned by $\mathbf{f}_1$ and $\mathbf{f}_2$. We get this by decomposing $\mathbf{f}_2$ into vectors which are parallel to and orthogonal to $\mathbf{e}_1$: we have $\mathbf{f}_{2\parallel} = (\mathbf{f}_2 \cdot \mathbf{e}_1)\mathbf{e}_1$, and $\mathbf{f}_{2\perp} = \mathbf{f}_2 - \mathbf{f}_{2\parallel}$.
- We now normalize this to get $\mathbf{e}_2 = (1/\|\mathbf{f}_{2\perp}\|)\mathbf{f}_{2\perp}$.

- Since $\mathbf{f}_{2\perp}$ is orthogonal to $\mathbf{e}_1$, so is $\mathbf{e}_2$. Moreover

$$\mathbf{f}_{2\perp} = \mathbf{f}_2 - \left(\frac{\mathbf{f}_2 \cdot \mathbf{f}_1}{\|\mathbf{f}_1\|^2} \right) \mathbf{f}_1,$$

so $\mathbf{f}_{2\perp}$ and hence $\mathbf{e}_2$ are linear combinations of $\mathbf{f}_1$ and $\mathbf{f}_2$. Therefore, $\mathbf{e}_1$ and $\mathbf{e}_2$ span the same space and give an orthonormal basis for V .

Example: Let V be the subspace of $\mathbb{R}^3$ spanned by

$$\{\mathbf{v}_1, \mathbf{v}_2\} = \left\{ \begin{pmatrix} 2 \\ 1 \\ 1 \end{pmatrix}, \begin{pmatrix} 1 \\ 2 \\ 0 \end{pmatrix} \right\}.$$

Then $\|\mathbf{v}_1\| = \sqrt{6}$, so

$$\mathbf{e}_1 = (1/\sqrt{6}) \begin{pmatrix} 2 \\ 1 \\ 1 \end{pmatrix}.$$

And

$$\mathbf{v}_{2\perp} = \mathbf{v}_2 - (\mathbf{v}_2 \cdot \mathbf{e}_1) \mathbf{e}_1 = \begin{pmatrix} 1 \\ 2 \\ 0 \end{pmatrix} - (2/3) \begin{pmatrix} 2 \\ 1 \\ 1 \end{pmatrix} = (1/3) \begin{pmatrix} -1 \\ 4 \\ -2 \end{pmatrix}.$$

Normalizing, we find

$$\mathbf{e}_2 = (1/\sqrt{21}) \begin{pmatrix} -1 \\ 4 \\ -2 \end{pmatrix}.$$

Exercise: Let $E_{3 \times 2} = \{\mathbf{e}_1, \mathbf{e}_2\}$, where the columns are the orthonormal basis vectors found above. What is $E^t E$? What is $E E^t$? Is E an orthogonal matrix? Why or why not?

Exercise: Find an orthonormal basis for the null space of the 1×3 matrix $A = (1, -2, 4)$.

Exercise: Let $\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_k\}$ be a set of (non-zero) orthogonal vectors. Prove that the set is linearly independent. (Hint: suppose that some linear combination is zero and show that all the coefficients must vanish.) Did you really need this hint?

19.3 Orthogonal projection onto a subspace V

Suppose $V \subseteq \mathbb{R}^n$ is a subspace, and suppose that $\{\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_k\}$ is an *orthonormal* basis for V . For any vector $\mathbf{x} \in \mathbb{R}^n$, we define

$$\Pi_V(\mathbf{x}) = \sum_{i=1}^k (\mathbf{x} \cdot \mathbf{e}_i) \mathbf{e}_i.$$

$\Pi_V(\mathbf{x})$ is called the *orthogonal projection* of $\mathbf{x}$ onto V . This is the natural generalization to higher dimensions of the projection of $\mathbf{x}$ onto a one-dimensional space considered before. Notice what we're doing: we're projecting $\mathbf{x}$ onto each of the 1-dimensional spaces determined by the basis vectors and then adding them all up.

Example: Let V be the column space of the matrix

$$A = \begin{pmatrix} 2 & 1 \\ 1 & 2 \\ 1 & 0 \end{pmatrix}.$$

As we found above, an orthonormal basis for V is given by

$$\{\mathbf{e}_1, \mathbf{e}_2\} = \left\{ (1/\sqrt{6}) \begin{pmatrix} 2 \\ 1 \\ 1 \end{pmatrix}, (1/\sqrt{21}) \begin{pmatrix} -1 \\ 4 \\ -2 \end{pmatrix} \right\}.$$

So if $\mathbf{x} = (1, 2, 3)^t$,

$$\begin{aligned} \Pi_V(\mathbf{x}) &= (\mathbf{x} \cdot \mathbf{e}_1) \mathbf{e}_1 + (\mathbf{x} \cdot \mathbf{e}_2) \mathbf{e}_2 \\ &= (7/\sqrt{6}) \mathbf{e}_1 + (1/\sqrt{21}) \mathbf{e}_2 \\ &= (7/6) \begin{pmatrix} 2 \\ 1 \\ 1 \end{pmatrix} + (1/21) \begin{pmatrix} -1 \\ 4 \\ 2 \end{pmatrix} \end{aligned}$$

Exercises:

- Show that the function $\Pi_V : \mathbb{R}^n \rightarrow V$ is a linear transformation.

- (Extra credit): Normally we don't define geometric objects by using a basis. When we do, as in the case of $\Pi_V(\mathbf{x})$, we need to show that the concept is *well-defined*. In this case, we need to show that $\Pi_V(\mathbf{x})$ is the same, no matter which orthonormal basis in V is used.

1. Suppose that $\{\mathbf{e}_1, \dots, \mathbf{e}_k\}$ and $\{\hat{\mathbf{e}}_1, \dots, \hat{\mathbf{e}}_k\}$ are two bases for V . Then $\hat{\mathbf{e}}_j = \sum_{i=1}^k P_{ij} \mathbf{e}_i$ for some $k \times k$ matrix P . Show that P is an orthogonal matrix.
2. Use this result to show that $\sum (\mathbf{x} \bullet \mathbf{e}_i) \mathbf{e}_i = \sum (\mathbf{x} \bullet \hat{\mathbf{e}}_j) \hat{\mathbf{e}}_j$, so that $\Pi_V(\mathbf{x})$ is independent of the basis.

19.4 Orthogonal complements

Definition: $V^\perp = \{\mathbf{x} \in \mathbb{R}^n \text{ such that } \mathbf{x} \bullet \mathbf{v} = 0, \text{ for all } \mathbf{v} \in V\}$ is called the *orthogonal complement of V* in $\mathbb{R}^n$.

Exercise: $V^\perp$ is a subspace of $\mathbb{R}^n$.

Example: Let

$$V = \text{span} \left\{ \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} \right\}.$$

Then

$$V^\perp = \left\{ \mathbf{v} \in \mathbb{R}^3 \text{ such that } \mathbf{v} \bullet \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} = 0 \right\} = \left\{ \begin{pmatrix} x \\ y \\ z \end{pmatrix} \text{ such that } x + y + z = 0 \right\}$$

This is the same as the null space of the matrix $A = (1, 1, 1)$. (Isn't it?). So writing $s = y$, $t = z$, we have

$$V^\perp = \left\{ \begin{pmatrix} -s-t \\ s \\ t \end{pmatrix} = s \begin{pmatrix} -1 \\ 1 \\ 0 \end{pmatrix} + t \begin{pmatrix} -1 \\ 0 \\ 1 \end{pmatrix}, s, t \in \mathbb{R} \right\}.$$

A basis for $V^\perp$ is clearly given by the two indicated vectors; of course, it's not orthonormal, but we could remedy that if we wanted.

Exercises:

1. Let $\{\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_k\}$ be a basis for W . Show that $v \in W^\perp \iff \mathbf{v} \bullet \mathbf{w}_i = 0, \forall i$.
2. Let

$$W = \text{span} \left\{ \begin{pmatrix} 1 \\ 2 \\ 1 \end{pmatrix}, \begin{pmatrix} 1 \\ -1 \\ 2 \end{pmatrix} \right\}$$

Find a basis for $W^\perp$. Hint: Use the result of exercise 1 to get a system of two equations in two unknowns and solve it.

19.5 Gram-Schmidt - the general algorithm

Let V be a subspace of $\mathbb{R}^n$, and $\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_m\}$ an arbitrary basis for V . We construct an orthonormal basis out of this as follows:

1. $\mathbf{e}_1 = \widehat{\mathbf{v}}_1$ (recall that this means we normalize $\mathbf{v}_1$ so that it has length 1. Let W_1 be the subspace $\text{span}\{\mathbf{e}_1\}$.
2. Take $\mathbf{f}_2 = \mathbf{v}_2 - \Pi_{W_1}(\mathbf{v}_2)$; then let $\mathbf{e}_2 = \widehat{\mathbf{f}}_2$. Let $W_2 = \text{span}\{\mathbf{e}_1, \mathbf{e}_2\}$.
3. Now assuming that W_k has been constructed, we define, recursively
 $\mathbf{f}_{k+1} = \mathbf{v}_{k+1} - \Pi_{W_k}(\mathbf{v}_{k+1})$, $\mathbf{e}_{k+1} = \widehat{\mathbf{f}}_{k+1}$, and $W_{k+1} = \text{span}\{\mathbf{e}_1, \dots, \mathbf{e}_{k+1}\}$.
4. Continue until W_m has been defined. Then $\{\mathbf{e}_1, \dots, \mathbf{e}_m\}$ is an orthonormal set in V , hence linearly independent, and thus a basis, since there are m vectors in the set.

20 Approximations - the method of least squares (1)

In many applications, we have to consider the following problem:

Suppose that for some $\mathbf{y}$, the equation $A\mathbf{x} = \mathbf{y}$ has no solutions. It could be that this is an important problem and we can't just forget about it. We could try to find an *approximate* solution. But which one? Suppose we choose an $\mathbf{x}$ at random. Then $A\mathbf{x} \neq \mathbf{y}$. In choosing this $\mathbf{x}$, we'll make an error $\mathbf{e} = A\mathbf{x} - \mathbf{y}$. A reasonable choice (not the only one) is to seek an $\mathbf{x}$ with the property that $\|A\mathbf{x} - \mathbf{y}\|$, the magnitude of the error, is as small as possible. (If this error is 0, then we have an exact solution, so it seems like a reasonable thing to try and minimize it.) Since this is a bit abstract, we can look at a familiar example:

Example: Suppose we have a bunch of data in the form of ordered pairs:

$\{(x_1, y_1), (x_2, y_2), \dots, (x_n, y_n)\}$. These data might come from an experiment; for instance, x_i might be the current through some device and y_i might be the temperature of the device while the given current flows through it. The n data points correspond to n different experimental observations.

The problem is to "fit" a straight line to this data. Another way to put this is: find the linear model that "best" predicts y , given x . Clearly, this is a problem which has no exact solution unless all the data points are collinear - there's no single line which goes through all the points. So how do we choose? The problem is to find m and b such that $y = mx + b$ is, in some sense, the best possible fit. The first thing to do is to convince ourselves that this problem is a special case of finding an approximate solution to $A\mathbf{x} = \mathbf{y}$:

Suppose we fix m and b . If the resulting line were a perfect fit, we'd have

$$\begin{aligned} y_1 &= mx_1 + b \\ y_2 &= mx_2 + b \\ &\vdots \\ y_n &= mx_n + b. \end{aligned}$$

Put

$$\mathbf{y} = \begin{pmatrix} y_1 \\ y_2 \\ \dots \\ y_n \end{pmatrix}, \quad A = \begin{pmatrix} x_1 & 1 \\ x_2 & 1 \\ \dots & \dots \\ x_n & 1 \end{pmatrix}, \quad \text{and } \mathbf{x} = \begin{pmatrix} m \\ b \end{pmatrix}.$$

Then the linear system above takes the form $\mathbf{y} = A\mathbf{x}$, where A and $\mathbf{y}$ are known, and the problem is that there is no solution $\mathbf{x} = (m, b)^t$.

20.1 The method of least squares

We can visualize the problem geometrically. Think of the matrix A as defining a linear function $\mathbf{f}_A : \mathbb{R}^n \rightarrow \mathbb{R}^m$. The range of $\mathbf{f}_A$ is a subspace of $\mathbb{R}^m$, and the source of our problem is that $\mathbf{y} \notin \text{Range}(\mathbf{f}_A)$. If we pick an arbitrary point $A\mathbf{x} \in \text{Range}(\mathbf{f}_A)$, then the error we've made is $\mathbf{e} = A\mathbf{x} - \mathbf{y}$. We want to choose $A\mathbf{x}$ so that $\|\mathbf{e}\|$ is as small as possible.

Exercise: This could clearly be handled as a calculus problem. How?

Instead of using calculus, we can do something simpler. We decompose the error as $\mathbf{e} = \mathbf{e}_{\parallel} + \mathbf{e}_{\perp}$, where $\mathbf{e}_{\parallel} \in \text{Range}(\mathbf{f}_A)$ and $\mathbf{e}_{\perp} \in \text{Range}(\mathbf{f}_A)^{\perp}$. See the figure on the next page.

Then $\|\mathbf{e}\|^2 = \|\mathbf{e}_{\parallel}\|^2 + \|\mathbf{e}_{\perp}\|^2$ (by Pythagoras' theorem!). Changing our choice of $A\mathbf{x}$ does not change $\mathbf{e}_{\perp}$, so the only variable at our disposal is $\mathbf{e}_{\parallel}$. We can make this $\mathbf{0}$ by choosing $A\mathbf{x}$ so that $\Pi(\mathbf{y}) = A\mathbf{x}$, where Π is the orthogonal projection of $\mathbb{R}^m$ onto the range of $\mathbf{f}_A$. And this is the answer to our question. Instead of solving $A\mathbf{x} = \mathbf{y}$, which is impossible, we solve for $\mathbf{x}$ in the equation $A\mathbf{x} = \Pi(\mathbf{y})$, which is guaranteed to have a solution. So we have




Figure 2: The plane is the range of $\mathbf{f}_A$. To minimize $\|\mathbf{e}\|$, we make $\mathbf{e}_{\parallel} = \mathbf{0}$ by choosing $\tilde{\mathbf{x}}$ so that $A\tilde{\mathbf{x}} = \Pi_V(\mathbf{y})$. So $A\tilde{\mathbf{x}}$ is the unlabeled vector from $\mathbf{0}$ to the foot of $\mathbf{e}_{\perp}$.

minimized the squared length of the error $\mathbf{e}$, thus the name *least squares* approximation. We collect this information in a

Definition: The vector $\tilde{\mathbf{x}}$ is said to be a *least squares solution* to $A\mathbf{x} = \mathbf{y}$ if the error vector $\mathbf{e} = A\tilde{\mathbf{x}} - \mathbf{y}$ is orthogonal to the range of $\mathbf{f}_A$.

Example (cont'd.): Note: We're writing this down to demonstrate that we could, if we had to, find the least squares solution by solving $A\mathbf{x} = \Pi(\mathbf{y})$ directly. But this is *not* what's done in practice, as we'll see in the next lecture. In particular, this is not an efficient way to proceed.

That having been said, let's use what we now know to find the line which best fits the data points. (This line is called the least squares *regression line*, and you've probably encountered it before.) We have to project $\mathbf{y}$ into the range of $\mathbf{f}_A$, where

$$A = \begin{pmatrix} x_1 & 1 \\ x_2 & 1 \\ \cdots & \cdots \\ x_n & 1 \end{pmatrix}.$$

To do this, we need an orthonormal basis for the range of $\mathbf{f}_A$, which is the same as the column space of the matrix A . We apply the Gram-Schmidt process to the columns of A , starting with the easy one:

$$\mathbf{e}_1 = \frac{1}{\sqrt{n}} \begin{pmatrix} 1 \\ 1 \\ \dots \\ 1 \end{pmatrix}.$$

If we write $\mathbf{v}$ for the first column of A , we now need to compute

$$\mathbf{v}_\perp = \mathbf{v} - (\mathbf{v} \cdot \mathbf{e}_1) \mathbf{e}_1$$

A routine computation (exercise!) gives

$$\mathbf{v}_\perp = \begin{pmatrix} x_1 - \bar{x} \\ x_2 - \bar{x} \\ \dots \\ x_n - \bar{x} \end{pmatrix}, \text{ where } \bar{x} = \frac{1}{n} \sum_{i=1}^n x_i$$

is the *mean* or average value of the x -measurements. Then

$$\mathbf{e}_2 = \frac{1}{\sigma} \begin{pmatrix} x_1 - \bar{x} \\ x_2 - \bar{x} \\ \dots \\ x_n - \bar{x} \end{pmatrix}, \text{ where } \sigma^2 = \sum_{i=1}^n (x_i - \bar{x})^2$$

is the *variance* of the x -measurements. Its square root, σ , is called the *standard deviation* of the measurements.

We can now compute

$$\begin{aligned} \Pi(\mathbf{y}) &= (\mathbf{y} \cdot \mathbf{e}_1) \mathbf{e}_1 + (\mathbf{y} \cdot \mathbf{e}_2) \mathbf{e}_2 \\ &= \text{routine computation here} \dots \\ &= \bar{y} \begin{pmatrix} 1 \\ 1 \\ \dots \\ 1 \end{pmatrix} + \frac{1}{\sigma^2} \left\{ \sum_{i=1}^n x_i y_i - n \bar{x} \bar{y} \right\} \begin{pmatrix} x_1 - \bar{x} \\ x_2 - \bar{x} \\ \dots \\ x_n - \bar{x} \end{pmatrix}. \end{aligned}$$

For simplicity, let

$$\alpha = \frac{1}{\sigma^2} \left\{ \sum_{i=1}^n x_i y_i - n \bar{x} \bar{y} \right\}.$$

Then the system of equations $A\mathbf{x} = \Pi(\mathbf{y})$ reads

$$\begin{aligned} mx_1 + b &= \alpha x_1 + \bar{y} - \alpha \bar{x} \\ mx_2 + b &= \alpha x_2 + \bar{y} - \alpha \bar{x} \\ &\dots \\ mx_n + b &= \alpha x_n + \bar{y} - \alpha \bar{x}, \end{aligned}$$

and we *know* (why?) that the augmented matrix for this system has rank 2. So we can solve for m and b just using the first two equations, assuming $x_1 \neq x_2$ so these two are not multiples of one another. Subtracting the second from the first gives

$$m(x_1 - x_2) = \alpha(x_1 - x_2), \text{ or } m = \alpha.$$

Now substituting α for m in either equation gives

$$b = \bar{y} - \alpha \bar{x}.$$

These are the formulas your graphing calculator uses to compute the slope and y -intercept of the regression line.

This is also about the simplest possible least squares computation we can imagine, and it's much too complicated to be of any practical use. Fortunately, there's a much easier way to do the computation, which is the subject of the next lecture.

21 Least squares approximation - II

21.1 The transpose of A

In the next section we'll develop an equation, known as the *normal* equation, which is much easier to solve than $A\mathbf{x} = \Pi(\mathbf{y})$, and which also gives the correct $\mathbf{x}$. Of course, we need a bit of background first.

The transpose of a matrix, which we haven't made much use of until now, begins to play a more important role once the dot product has been introduced. If A is an $m \times n$ matrix, then as you know, it can be regarded as a linear transformation from $\mathbb{R}^n$ to $\mathbb{R}^m$. Its transpose, A^t then gives a linear transformation from $\mathbb{R}^m$ to $\mathbb{R}^n$, since it's $n \times m$. Note that there is no implication here that $A^t = A^{-1}$ – the matrices needn't be square, and even if they are, they need not be invertible. But A and A^t are related by the dot product:

Theorem: $\mathbf{x} \bullet A^t \mathbf{y} = A \mathbf{x} \bullet \mathbf{y}$

Proof: (Notice that the dot product on the left is in $\mathbb{R}^n$, while the one on the right is in $\mathbb{R}^m$.)

The proof is a "straightforward" computation:

$$\begin{aligned} A \mathbf{x} \bullet \mathbf{y} &= \sum_{j=1}^m (A \mathbf{x})_j y_j \\ &= \sum_{j=1}^m \left(\sum_{i=1}^n A_{ji} x_i \right) y_j \\ &\quad \text{now we reverse the order of summation to get} \\ &= \sum_{i=1}^n x_i \left(\sum_{j=1}^m A_{ji} y_j \right) \\ &\quad \text{and since } A_{ji} = A_{ij}^t, \text{ we get} \\ &= \sum_{i=1}^n x_i (A^t \mathbf{y})_i \\ &= \mathbf{x} \bullet A^t \mathbf{y} \end{aligned}$$

What this says in plain English: we can "move" A from one side of the dot product to the other by replacing it with A^t . So for instance, if $A \mathbf{x} \bullet \mathbf{y} = 0$, then $\mathbf{x} \bullet A^t \mathbf{y} = 0$, and conversely. In fact, pushing this a bit, we get an important

Theorem: $\text{Ker}(A^t) = (\text{Range}(A))^\perp$.

Proof: Let $\mathbf{y} \in (\text{Range}(A))^\perp$. This means that for all $\mathbf{x} \in \mathbb{R}^n$, $A\mathbf{x} \cdot \mathbf{y} = 0$. But by the previous theorem, this means that $\mathbf{x} \cdot A^t\mathbf{y} = 0$ for all $\mathbf{x} \in \mathbb{R}^n$. But any vector in $\mathbb{R}^n$ which is orthogonal to everything must be the zero vector. So $A^t\mathbf{y} = \mathbf{0}$ and $\mathbf{y} \in \text{Ker}(A^t)$. Conversely, if $\mathbf{y} \in \text{Ker}(A^t)$, then for any $\mathbf{x} \in \mathbb{R}^n$, $\mathbf{x} \cdot A^t\mathbf{y} = 0$. And again by the theorem, this means that $A\mathbf{x} \cdot \mathbf{y} = 0$ for all such $\mathbf{x}$, which means that $\mathbf{y} \perp \text{Range}(A)$.

We have shown that $(\text{Range}(A))^\perp \subseteq \text{Ker}(A^t)$, and conversely, that $\text{Ker}(A^t) \subseteq (\text{Range}(A))^\perp$. So the two sets are equal.

21.2 Least squares approximations – the Normal equation

Now we're ready to take up the least squares problem again. Recall that the problem is to solve $A\mathbf{x} = \Pi(\mathbf{y})$, where $\mathbf{y}$ has been projected orthogonally on to the range of A . The problem with solving this, as you'll recall, is that finding the projection Π is tedious. And now we'll see that it's not necessary.

We write $\mathbf{y} = \Pi(\mathbf{y}) + \mathbf{y}_\perp$, where $\mathbf{y}_\perp$ is orthogonal to the range of A . Now suppose that $\mathbf{x}$ is a solution to the least squares problem $A\mathbf{x} = \Pi(\mathbf{y})$. Multiply this equation by A^t to get $A^t A\mathbf{x} = A^t \Pi(\mathbf{y})$. So $\mathbf{x}$ is certainly also a solution to this. But now we notice that, in consequence of the previous theorem,

$$A^t\mathbf{y} = A^t(\Pi(\mathbf{y}) + \mathbf{y}_\perp) = A^t\Pi(\mathbf{y}),$$

since $A^t\mathbf{y}_\perp = \mathbf{0}$. (It's orthogonal to the range, so the theorem says it's in $\text{Ker}(A^t)$.)

So $\mathbf{x}$ is also a solution to the *normal equation*

$$A^t A\mathbf{x} = A^t\mathbf{y}.$$

Conversely, if $\mathbf{x}$ is a solution to the normal equation, then

$$A^t(A\mathbf{x} - \mathbf{y}) = \mathbf{0},$$

and by the previous theorem, this means that $A\mathbf{x} - \mathbf{y}$ is orthogonal to the range of A . But $A\mathbf{x} - \mathbf{y}$ is the error made using an approximate solution, and this shows that the error vector is orthogonal to the range of A – this is our definition of the least squares solution!

The reason for all this fooling around is simple: we can compute $A^t\mathbf{y}$ by doing a simple matrix multiplication. We don't need to find an orthonormal basis for the range of A to compute Π . We summarize the results:

Theorem: $\tilde{\mathbf{x}}$ is a least-squares solution to $A\mathbf{x} = \mathbf{y} \iff \tilde{\mathbf{x}}$ is a solution to the normal equation $A^tA\mathbf{x} = A^t\mathbf{y}$.

Example: Find the least squares regression line through the 4 points $(1, 2), (2, 3), (-1, 1), (0, 1)$.

Solution: We've already set up this problem in the last lecture. We have

$$A = \begin{pmatrix} 1 & 1 \\ 2 & 1 \\ -1 & 1 \\ 0 & 1 \end{pmatrix}, \quad \mathbf{y} = \begin{pmatrix} 2 \\ 3 \\ 1 \\ 1 \end{pmatrix}, \quad \text{and } \mathbf{x} = \begin{pmatrix} m \\ b \end{pmatrix}.$$

We compute

$$A^tA = \begin{pmatrix} 6 & 2 \\ 2 & 4 \end{pmatrix}, \quad A^t\mathbf{y} = \begin{pmatrix} 7 \\ 7 \end{pmatrix}$$

And the solution to the normal equation is

$$\mathbf{x} = (A^tA)^{-1}A^t\mathbf{y} = (1/20) \begin{pmatrix} 4 & -2 \\ -2 & 6 \end{pmatrix} \begin{pmatrix} 7 \\ 7 \end{pmatrix} = \begin{pmatrix} 7/10 \\ 7/5 \end{pmatrix}.$$

So the regression line has the equation $y = (7/10)x + 7/5$.

Exercises:

1. For these problems, think of the row space as the column space of A^t . Show that $\mathbf{v}$ is in the row space of $A \iff \mathbf{v} = A^t\mathbf{y}$ for some $\mathbf{y}$. This means that the row space of A is the range of $\mathbf{f}_{A^t}$ (analogous to the fact that the column space of A is the range of $\mathbf{f}_A$).

2. Show that the null space of A is the orthogonal complement of the row space. (Hint: use the above theorem with A^t instead of A .)

