

Math 18.06 - Linear algebra

All readings are from the textbook “Introduction to linear algebra”, 5–th edition, by Gilbert Strang.

All viewings are from Gilbert Strang’s videos, available online at:

<https://ocw.mit.edu/courses/mathematics/18-06-linear-algebra-spring-2010/video-lectures/>

Note that often, my lectures will correspond to parts of the corresponding video lectures.

Lecture 1 (February 17)

Reading: Sections 1.1–2.1

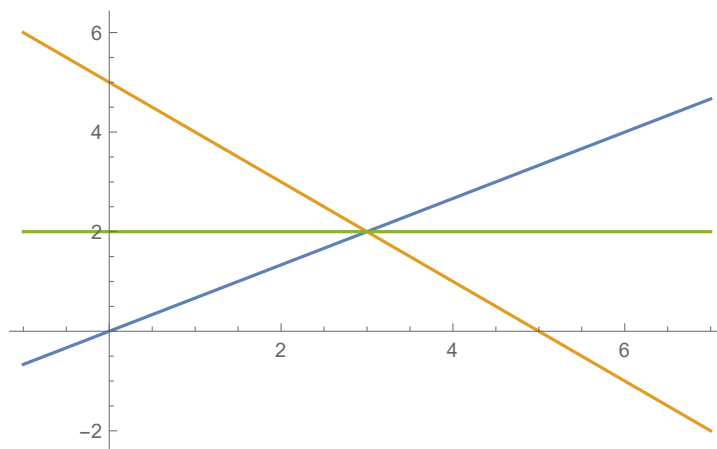
Viewing: Lecture 1

In 18.02, you learned how to solve systems of linear equations:

$$\begin{cases} 2x - 3y = 0 \\ x + y = 5 \\ 2y = 4 \end{cases} \quad (1)$$

The high-school way to solve this system is to take y out from the last equation (since it is the more manageable of the three) and get $y = 2$. Then we plug this answer into either the first or the second equation, and we get $x = 3$. So the solution is $(x, y) = (3, 2)$.

There's also a geometric way to think about the system of equations (1). Linear equations, such as $2x - 3y = 0$, $x + y = 5$ or $2y = 4$, are equations of lines in the xy -plane. Specifically, they correspond to the blue, orange and green lines in the picture below:



Asking for a solution to the system of equations (1) is the same thing as asking for a point (x, y) which lies on all three of these lines, or in other words, the intersection of these lines. It's easy to see from the picture above that this point is $(x, y) = (3, 2)$.

Another point of view on the system of equations (1) is to rewrite it in matrix form:

$$\begin{bmatrix} 2 & -3 \\ 1 & 1 \\ 0 & 2 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} 0 \\ 5 \\ 4 \end{bmatrix} \quad (2)$$

In this way, you solve for the quantity $\begin{bmatrix} x \\ y \end{bmatrix}$, which is thought of as the point (x, y) in the plane.

There are a number of ways to solve matrix equations like (2), the most general form of which is:

$$A\mathbf{v} = \mathbf{b} \quad (3)$$

where A is an $m \times n$ matrix, \mathbf{v} is an $n \times 1$ vector, and \mathbf{b} is a $m \times 1$ vector. In situations such as (3), usually A and \mathbf{b} are given as part of the problem, while \mathbf{v} is the unknown you have to solve for.

Before we delve into ways to solve (3), let us say a thing or two about vectors. We think of:

$$\mathbf{v} = \begin{bmatrix} v_1 \\ v_2 \\ \vdots \\ v_n \end{bmatrix} \quad (4)$$

as a point in n -dimensional space. For example, 2-dimensional space is a plane (such as the surface of the paper, or a blackboard), while 3-dimensional space, is... well, ordinary space. Consider the following operations:

- given two vectors $\mathbf{v} = \begin{bmatrix} v_1 \\ v_2 \\ \vdots \\ v_n \end{bmatrix}$ and $\mathbf{w} = \begin{bmatrix} w_1 \\ w_2 \\ \vdots \\ w_n \end{bmatrix}$, you may add them componentwise:

$$\mathbf{v} + \mathbf{w} = \begin{bmatrix} v_1 + w_1 \\ v_2 + w_2 \\ \vdots \\ v_n + w_n \end{bmatrix}$$

- given a vector \mathbf{v} and a number α (called scalar), we can multiply them together:

$$\alpha\mathbf{v} = \begin{bmatrix} \alpha v_1 \\ \alpha v_2 \\ \vdots \\ \alpha v_n \end{bmatrix}$$

Combining the two operations above, we will say a **linear combination** of two vectors \mathbf{v} and \mathbf{w} is any expression of the form:

$$\alpha\mathbf{v} + \beta\mathbf{w} = \begin{bmatrix} \alpha v_1 + \beta w_1 \\ \alpha v_2 + \beta w_2 \\ \vdots \\ \alpha v_n + \beta w_n \end{bmatrix}$$

The numbers α and β are called the **coefficients** of the linear combination. So we've seen the definition, but what does the concept of "linear combination" actually mean? Let's take an easy example:

$$\mathbf{v} = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} \quad \mathbf{w} = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \quad \Rightarrow \quad \text{linear combinations are } \alpha\mathbf{v} + \beta\mathbf{w} = \begin{bmatrix} \alpha \\ 0 \\ \beta \end{bmatrix}$$

As α and β run over all possible numbers, the right-hand side runs over all vectors in the xz -plane (which in turn is a subset of xyz -space), since the y -component of such vectors is 0. Given that \mathbf{v} is the unit vector on the x -axis and \mathbf{w} is the unit vector on the z -axis, the principle here is the following.

Fact 1. Almost ¹ any two $n \times 1$ vectors \mathbf{v} and \mathbf{w} trace out a 2-dimensional plane in n -dimensional space. Any vector in the aforementioned plane is a linear combination of \mathbf{v} and \mathbf{w} , and vice-versa.

Example 1. Let's draw the plane traced out by the vectors $\mathbf{v} = \begin{bmatrix} 2 \\ 0 \\ -1 \end{bmatrix}$ and $\mathbf{w} = \begin{bmatrix} 0 \\ 3 \\ 1 \end{bmatrix}$.

To do this, represent \mathbf{v} and \mathbf{w} as the points $(2, 0, -1)$ and $(0, 3, 1)$, respectively, in 3-dimensional space. Then trace out the lines passing through the origin and either the point \mathbf{v} or \mathbf{w} . Finally, trace out the 2-plane spanned by these two lines.

We can use the language of linear combinations to make sense of systems of equations such as (3). Let's look at the specific example (2). Because:

$$\begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} x \\ 0 \end{bmatrix} + \begin{bmatrix} 0 \\ y \end{bmatrix} = x \cdot \begin{bmatrix} 1 \\ 0 \end{bmatrix} + y \cdot \begin{bmatrix} 0 \\ 1 \end{bmatrix}$$

then formula (2) reads:

$$\begin{bmatrix} 2 & -3 \\ 1 & 1 \\ 0 & 2 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = x \cdot \begin{bmatrix} 2 & -3 \\ 1 & 1 \\ 0 & 2 \end{bmatrix} \begin{bmatrix} 1 \\ 0 \end{bmatrix} + y \cdot \begin{bmatrix} 2 & -3 \\ 1 & 1 \\ 0 & 2 \end{bmatrix} \begin{bmatrix} 0 \\ 1 \end{bmatrix} = x \cdot \begin{bmatrix} 2 \\ 1 \\ 0 \end{bmatrix} + y \cdot \begin{bmatrix} -3 \\ 1 \\ 2 \end{bmatrix} = \begin{bmatrix} 0 \\ 5 \\ 4 \end{bmatrix}$$

Therefore, solving for (x, y) is the same thing as finding a certain linear combination of the vectors $\begin{bmatrix} 2 \\ 1 \\ 0 \end{bmatrix}$ and $\begin{bmatrix} -3 \\ 1 \\ 2 \end{bmatrix}$ (which are nothing but the columns of the 3×2 matrix in the left-hand side of (2))

that produces the vector $\begin{bmatrix} 0 \\ 5 \\ 4 \end{bmatrix}$. Such a linear combination exists if and only if the latter vector lies in the plane spanned out by the former vectors, which is indeed the case in our example. In general:

Fact 2. The equation $A\mathbf{v} := \begin{bmatrix} a_{11} & \dots & a_{1n} \\ \vdots & \ddots & \vdots \\ a_{m1} & \dots & a_{mn} \end{bmatrix} \begin{bmatrix} v_1 \\ \vdots \\ v_n \end{bmatrix} = \begin{bmatrix} b_1 \\ \vdots \\ b_m \end{bmatrix} =: \mathbf{b}$ is equivalent to:

$$v_1 \begin{bmatrix} a_{11} \\ \vdots \\ a_{m1} \end{bmatrix} + \dots + v_n \begin{bmatrix} a_{1n} \\ \vdots \\ a_{mn} \end{bmatrix} = \begin{bmatrix} b_1 \\ \vdots \\ b_m \end{bmatrix} \quad (5)$$

In other words, solving the system of equations for \mathbf{v} precisely entails finding out which linear combinations of the columns of the matrix A are equal to the vector \mathbf{b} .

¹The word “almost” requires us to exclude the case when \mathbf{v} and \mathbf{w} are collinear, i.e. on the same line passing through the origin; we will discuss this issue further when introducing the notion of linear independence

There is another operation that one can perform on two vectors, the **dot product**:

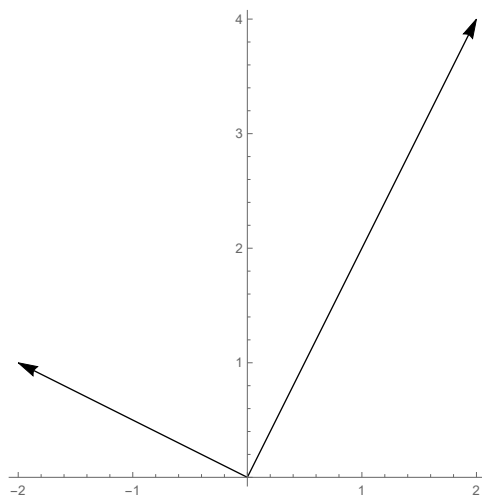
$$\text{if } \mathbf{v} = \begin{bmatrix} v_1 \\ \vdots \\ v_n \end{bmatrix} \text{ and } \mathbf{w} = \begin{bmatrix} w_1 \\ \vdots \\ w_n \end{bmatrix} \text{ then } \boxed{\mathbf{v} \cdot \mathbf{w} = v_1 w_1 + \cdots + v_n w_n}$$

VERY IMPORTANT: the dot product of two vectors is a number (a.k.a. scalar). There is no “reasonable” way to multiply two vectors in such a way as to obtain a vector (with the notable exception of the cross-product for vectors with $n = 3$ components, but that’s another story).

The dot product allows us to describe the angle θ between two vectors \mathbf{v} and \mathbf{w} . Recall that:

$$\mathbf{v} \cdot \mathbf{w} = 0 \quad \text{means that} \quad \theta = 90^\circ$$

For example, take $\mathbf{v} = \begin{bmatrix} -2 \\ 1 \end{bmatrix}$ and $\mathbf{w} = \begin{bmatrix} 2 \\ 4 \end{bmatrix}$:



The angle between these two vectors is 90° , and indeed $\mathbf{v} \cdot \mathbf{w} = (-2) \cdot 2 + 1 \cdot 4 = 0$. But the dot product describes arbitrary angles, not just right angles. The general formula is:

$$\boxed{\cos \theta = \frac{\mathbf{v} \cdot \mathbf{w}}{\|\mathbf{v}\| \|\mathbf{w}\|}} \quad (6)$$

where the **length** of a vector is given by:

$$\boxed{\|\mathbf{v}\| = \sqrt{\mathbf{v} \cdot \mathbf{v}} = \sqrt{v_1^2 + v_2^2 + \cdots + v_n^2}} \quad (7)$$

So if you want to compute the angle between two lines, apply formula (6) with \mathbf{v} and \mathbf{w} being the unit (i.e. length 1) vectors corresponding to those lines. Then the formula in question precisely says that the cosine of the angle is equal to the dot product of unit vectors.

Lecture 2 (February 19)

Reading: Sections 2.2–2.3

Viewing: Lecture 2

The basic method for solving systems of equations is **Gaussian elimination**. For instance:

$$\begin{cases} x - y + 2z = 1 \\ -2x + 2y - 3z = -1 \\ -3x - y + 2z = -3 \end{cases} \quad \text{or, equivalently,} \quad \begin{bmatrix} 1 & -1 & 2 \\ -2 & 2 & -3 \\ -3 & -1 & 2 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \end{bmatrix} = \begin{bmatrix} 1 \\ -1 \\ -3 \end{bmatrix} \quad (8)$$

The 3×3 matrix $A = \begin{bmatrix} 1 & -1 & 2 \\ -2 & 2 & -3 \\ -3 & -1 & 2 \end{bmatrix}$ is called the **coefficient matrix**, while $\mathbf{v} = \begin{bmatrix} x \\ y \\ z \end{bmatrix}$ is the unknown vector. Geometrically, the system (8) consists of three linear equations in xyz -space. Each of these equations determines a 2-dimensional plane in 3-dimensional space, so the system of equations computes the coordinates (x, y, z) of the intersection point of three planes.

Gaussian elimination is the following algorithm. Consider the **augmented matrix**:

$$\left[\begin{array}{ccc|c} \boxed{1} & -1 & 2 & 1 \\ \boxed{-2} & 2 & -3 & -1 \\ \boxed{-3} & -1 & 2 & -3 \end{array} \right] \quad (9)$$

obtained by tacking the right-hand side of (8) onto the right of the matrix A . The **pivot** of each row (marked by a box) refers to the leftmost non-zero entry on that row. The game is the following: by performing **row eliminations** (namely adding any multiple of row j to row i , for various $i > j$), your goal is to ensure that the pivot of the first row is strictly to the left of the pivot on the second row, which in turn is strictly to the left of the pivot on the third row etc.

The particular row eliminations you have to do are imposed upon you by the following algorithm. Start from the pivot on the first row. I claim that there's a single choice of row eliminations which will ensure that the entries below the pivot on the first row will be 0. This is because:

$$a \cdot 1 + (-2) = 0 \text{ requires } a = 2$$

$$b \cdot 1 + (-3) = 0 \text{ requires } b = 3$$

so we must add 2 times the first row to the second row and 3 times the first row to the third row, in order to achieve our stated goal of having the entries below the first pivot be 0.

$$\left[\begin{array}{ccc|c} 1 & -1 & 2 & 1 \\ -2 + 2 \cdot 1 & 2 + 2 \cdot (-1) & -3 + 2 \cdot 2 & -1 + 2 \cdot 1 \\ -3 + 3 \cdot 1 & -1 + 3 \cdot (-1) & 2 + 3 \cdot 2 & -3 + 3 \cdot 1 \end{array} \right] = \left[\begin{array}{ccc|c} \boxed{1} & -1 & 2 & 1 \\ 0 & 0 & \boxed{1} & 1 \\ 0 & \boxed{-4} & 8 & 0 \end{array} \right]$$

The next step is to find the the pivot of the second row, and repeat the process. But wait! In the example above, the pivot on the second row is to the right of the pivot on the third row. Whenever this happens, there's no row elimination which will put the pivots in the right order (namely the pivot on the second row to the left of the pivot on the third row). So we allow ourselves to make an

extra move: **row exchanges**. Specifically, we simply swap the second and third rows. We obtain:

$$\left[\begin{array}{ccc|c} \boxed{1} & -1 & 2 & 1 \\ 0 & \boxed{-4} & 8 & 0 \\ 0 & 0 & \boxed{1} & 1 \end{array} \right] \quad (10)$$

This augmented matrix is in **row echelon form**, which is the purpose of Gaussian elimination: all the pivots are non-zero, and each pivot is strictly to the right of the one on the row above.

Remark. We only allow the pivots to be in the original part of the matrix, not in the augmented part (i.e. not in the right-most column of (9)). If at any point during Gaussian elimination, some i -th row of the matrix does not have a pivot (which can only happen if that row has all zeroes except maybe on the last column) then we use a row exchange to swap the i -th row with the last row. Thus, at the end of Gaussian elimination, all rows without pivots (if any) will be at the bottom of the matrix.

Remark. There's a variant of this algorithm, called Gauss-Jordan elimination, which entails doing two additional steps. Firstly, multiply all the rows of the augmented matrix by appropriate constants, so that all the pivots are equal to 1:

$$\left[\begin{array}{ccc|c} \boxed{1} & -1 & 2 & 1 \\ 0 & \boxed{1} & -2 & 0 \\ 0 & 0 & \boxed{1} & 1 \end{array} \right]$$

Then, we do some further row operations so that all the entries above the pivots are 0 as well. In the case of the matrix at hand, we add 2 times the third row to the second row:

$$\left[\begin{array}{ccc|c} \boxed{1} & -1 & 2 & 1 \\ 0 & \boxed{1} & 0 & 2 \\ 0 & 0 & \boxed{1} & 1 \end{array} \right]$$

and then we add -2 times the third row and 1 times the second row to the first row:

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

The output of Gauss-Jordan elimination is said to be in **reduced row echelon form**, i.e. a matrix which is both in row echelon form, and satisfies the two extra conditions underlined above. We will revisit this notion in a few weeks' time.

The augmented matrix (10) corresponds to the following system of linear equations:

$$\begin{bmatrix} 1 & -1 & 2 \\ 0 & -4 & 8 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \end{bmatrix} = \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix} \quad \text{or, equivalently,} \quad \begin{cases} x - y + 2z = 1 \\ -4y + 8z = 0 \\ z = 1 \end{cases} \quad (11)$$

and the great thing about Gaussian elimination is that this system of equations is actually equivalent to the system (8) that we started from:

any solution (x, y, z) of (8) is also a solution of (11), and vice versa

So why did we go through the trouble of Gaussian (or Gauss-Jordan) elimination? This is because systems where the coefficient matrix is in (reduced) row echelon form are really easy to solve by hand, by the following procedure (for which the technical term is **back substitution**): the third equation in (11) gives you $z = 1$, which you can plug into the second equation to get $y = 2$, which you can plug into the first equation to get $x = 1$. So the solution is $(x, y, z) = (1, 2, 1)$.

Remark. *If the row echelon form of a matrix A has a full row of zeroes (i.e. no pivot on the last row), then the matrix A is called **singular**. In this case, the system (11) may not have any solutions or may have more solutions than expected. Otherwise, the matrix is called **non-singular**.*

So what is going on behind this method? Back substitution provided such a quick answer because:

$$\boxed{U\mathbf{v} = \mathbf{c}} \tag{12}$$

is a very easy system to solve if U is in row echelon form. Indeed, you solve for the last component of the vector \mathbf{v} first, then for the next-to-last one, then for the next-to-next-to-last one etc. If you need to solve a system of equations $A\mathbf{v} = \mathbf{b}$ for a general matrix A , then Gaussian elimination precisely provides a recipe for putting A in row echelon form, at the cost of performing the following row operations (for any $i > j$):

- Adding the j -th row times λ to the i -th row of A , which is achieved by multiplying the latter on the left with an **elimination matrix**:

$$A \rightsquigarrow E_{ij}^{(\lambda)} A \quad \text{where} \quad E_{ij}^{(\lambda)} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & \lambda & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

The entry λ is on the i -th row and j -th column (in the formula above, $i = 4$ and $j = 2$).²

- Switching the i -th and j -th rows of A is achieved by multiplying the latter on the left with a **permutation matrix**:

$$A \rightsquigarrow P_{ij} A \quad \text{where} \quad P_{ij} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

Explicitly, P_{ij} has coefficients 1 on positions (i, j) , (j, i) and (k, k) for all $k \neq i, j$, and all other coefficients of P_{ij} are 0 (in the formula above, $i = 3$ and $j = 4$).

²If you're doing Gauss-Jordan elimination, you will also need the elimination matrices $E_{ij}^{(\lambda)}$ for $i < j$.

- If you wish to put the matrix in reduced row echelon (as opposed to just row echelon) form, you will also need to multiply the i -th row of A by various constants λ , which is achieved by multiplying A on the left by a **diagonal** matrix:

$$A \rightsquigarrow D_i^{(\lambda)} A \quad \text{where} \quad D_i^{(\lambda)} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & \lambda & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

The entry λ is on the i -th row (in the formula above, $i = 2$).

So let's see what these matrices amount to in the example of (8). We will forget about the 3×4 augmented matrix $[A|\mathbf{b}]$ and just discuss Gaussian elimination on the 3×3 matrix:

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

The first thing we did was to add 2 times the first row to the second row:

$$E_{21}^{(2)} A = \begin{bmatrix} 1 & 0 & 0 \\ 2 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & -1 & 2 \\ -2 & 2 & -3 \\ -3 & -1 & 2 \end{bmatrix} = \begin{bmatrix} 1 & -1 & 2 \\ 0 & 0 & 1 \\ -3 & -1 & 2 \end{bmatrix}$$

The next thing we did was to add 3 times the first row to the third row:

$$E_{31}^{(3)} E_{21}^{(2)} A = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 3 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & -1 & 2 \\ 0 & 0 & 1 \\ -3 & -1 & 2 \end{bmatrix} = \begin{bmatrix} 1 & -1 & 2 \\ 0 & 0 & 1 \\ 0 & -4 & 8 \end{bmatrix}$$

Then we exchanged the second and third rows:

$$P_{23} E_{31}^{(3)} E_{21}^{(2)} A = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} 1 & -1 & 2 \\ 0 & 0 & 1 \\ 0 & -4 & 8 \end{bmatrix} = \begin{bmatrix} 1 & -1 & 2 \\ 0 & -4 & 8 \\ 0 & 0 & 1 \end{bmatrix} =: U \tag{13}$$

and we have achieved precisely the row echelon form matrix of (11).

Remark. Suppose we wish to go further, and do Gauss-Jordan elimination. Then we need to make all the pivots 1, and to this end we must multiply the second row by $-1/4$:

$$D_2^{(-\frac{1}{4})} U = \begin{bmatrix} 1 & 0 & 0 \\ 0 & -\frac{1}{4} & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & -1 & 2 \\ 0 & -4 & 8 \\ 0 & 0 & 1 \end{bmatrix} = \begin{bmatrix} 1 & -1 & 2 \\ 0 & 1 & -2 \\ 0 & 0 & 1 \end{bmatrix}$$

The pivots are the 1's on the diagonal. To put U in reduced row echelon form, we need to annihilate all the non-zero entries above the 1's. This is achieved by adding 2 times the third row to the second row, then adding -2 times the third row to the first row, then adding the second row to the first:

$$E_{12}^{(1)} E_{13}^{(-2)} E_{23}^{(2)} D_2^{(-\frac{1}{4})} U = \begin{bmatrix} 1 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & -2 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 2 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & -1 & 2 \\ 0 & 1 & -2 \\ 0 & 0 & 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} =: R$$

The matrix R is now in reduced row echelon form.

Very importantly, this procedure also allows us to convert the system of equations (8) into a system of the form (12). Specifically, take the equality (8) and multiply it on the left with the particular product of elimination/permutation/diagonal matrices which appears in (13). We obtain:

$$\underbrace{P_{23}E_{31}^{(3)}E_{21}^{(2)}}_U A \mathbf{v} = P_{23}E_{31}^{(3)}E_{21}^{(2)} \mathbf{b} \quad (14)$$

If we let:

$$\begin{aligned} \mathbf{c} := P_{23}E_{31}^{(3)}E_{21}^{(2)} \mathbf{b} &= \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 3 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 2 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 \\ -1 \\ -3 \end{bmatrix} = \\ &= \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 3 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 \\ 1 \\ -3 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix} = \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix} \end{aligned}$$

then formula (14) reads precisely (12), which establishes the following important principle:

Fact 3. *If:*

$$\boxed{A = LU \text{ and } \mathbf{b} = L\mathbf{c}}$$

for some matrix L (which in practice will be a product of elimination/permutation/diagonal matrices), then the system $A\mathbf{v} = \mathbf{b}$ is equivalent (i.e. has the same solution \mathbf{v}) as the system $U\mathbf{v} = \mathbf{c}$.

Since we have already seen that it is easy to solve a system of equations where the matrix U is in row echelon form, this gives a pretty general algorithm for solving general systems of equations.

Remark. *Gaussian (or Gauss-Jordan) elimination can be applied to any $m \times n$ matrix, producing a matrix in row (or reduced row) echelon form. The fact that in the first part of this lecture, we applied it to the augmented matrix $[A|\mathbf{b}]$ was simply because this is what you're supposed to do if you want Gaussian elimination to help you solve systems of equations. But as we have seen in the second part of this lecture, the method of Gaussian elimination can be applied to the matrix A itself.*

Lecture 3 (February 22)

Reading: Sections 2.4–2.5

Viewing: Lecture 3

Let's delve a little deeper into matrix multiplication. The key thing is the following:

$$(m \times n \text{ matrix})(n \times p \text{ matrix}) = (m \times p \text{ matrix}) \quad (15)$$

AB only makes sense if the number of columns of A equals the number of rows of B .

Assume A is an $m \times n$ matrix with entries a_{ij} and that B is an $n \times p$ matrix with entries b_{ij} . Then the product $C = AB$ is a $m \times p$ matrix, whose entry on the i -th row and j -th column is equal to:

$$c_{ij} = a_{i1}b_{1j} + a_{i2}b_{2j} + \cdots + a_{in}b_{nj} = \sum_{k=1}^n a_{ik}b_{kj} \quad (16)$$

Here's how to see the same computation by looking at the matrices themselves:

$$\underbrace{\begin{bmatrix} * & * & * & * & * \\ * & * & c_{ij} & * & * \\ * & * & * & * & * \end{bmatrix}}_{AB} = \underbrace{\begin{bmatrix} * & * & * & * \\ a_{i1} & a_{i2} & \cdots & a_{in} \\ * & * & * & * \end{bmatrix}}_A \underbrace{\begin{bmatrix} * & * & b_{1j} & * & * \\ * & * & b_{2j} & * & * \\ * & * & \cdots & * & * \\ * & * & b_{nj} & * & * \end{bmatrix}}_B$$

(in the example above, $m = 3$, $n = 4$, $p = 5$, $i = 2$, $j = 3$). The i -th row of the matrix A and the j -th column of the matrix B are vectors of the same size, namely n . The only distinction is that the former vector is written horizontally and the latter vector is written vertically. Formula (16) precisely says that the (i, j) -entry in the matrix AB is the **dot product** of these two vectors.

What if we wanted to look at the entire i -th row of the matrix AB ?

$$\underbrace{\begin{bmatrix} * & * & * & * & * \\ c_{i1} & c_{i2} & \cdots & \cdots & c_{ip} \\ * & * & * & * & * \end{bmatrix}}_{AB} = \underbrace{\begin{bmatrix} * & * & * & * \\ a_{i1} & a_{i2} & \cdots & a_{in} \\ * & * & * & * \end{bmatrix}}_A \underbrace{\begin{bmatrix} b_{11} & b_{12} & \cdots & \cdots & b_{1p} \\ b_{21} & b_{22} & \cdots & \cdots & b_{2p} \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ b_{n1} & b_{n2} & \cdots & \cdots & b_{np} \end{bmatrix}}_B$$

The entire i -th row of AB is a linear combination of the rows of B , according to the rule:

$$[c_{i1} \ c_{i2} \ \cdots \ \cdots \ c_{ip}] = a_{i1} [b_{11} \ b_{12} \ \cdots \ \cdots \ b_{1p}] + \cdots + a_{in} [b_{n1} \ b_{n2} \ \cdots \ \cdots \ b_{np}]$$

Similarly, we may ask about the entire j -th column of the matrix AB :

$$\underbrace{\begin{bmatrix} * & * & c_{1j} & * & * \\ * & * & \cdots & * & * \\ * & * & c_{mj} & * & * \end{bmatrix}}_{AB} = \underbrace{\begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1n} \\ \vdots & \vdots & \ddots & \vdots \\ a_{m1} & a_{m2} & \cdots & a_{mn} \end{bmatrix}}_A \underbrace{\begin{bmatrix} * & * & b_{1j} & * & * \\ * & * & b_{2j} & * & * \\ * & * & \cdots & * & * \\ * & * & b_{nj} & * & * \end{bmatrix}}_B$$

The entire j -th column of AB is a linear combination of the columns of A , according to the rule:

$$\begin{bmatrix} c_{1j} \\ \vdots \\ c_{mj} \end{bmatrix} = \begin{bmatrix} a_{11} \\ \vdots \\ a_{m1} \end{bmatrix} b_{1j} + \begin{bmatrix} a_{12} \\ \vdots \\ a_{m2} \end{bmatrix} b_{2j} + \cdots + \begin{bmatrix} a_{1n} \\ \vdots \\ a_{mn} \end{bmatrix} b_{nj}$$

The two computations we have just seen are particular instances of **block multiplication** of matrices. Explicitly, we say that matrices A and B are written in **block form**:

$$A = \left[\begin{array}{c|c} A_{11} & A_{12} \\ \hline A_{21} & A_{22} \end{array} \right] \quad \text{and} \quad B = \left[\begin{array}{c|c} B_{11} & B_{12} \\ \hline B_{21} & B_{22} \end{array} \right] \quad (17)$$

if A_{ij} is an $m_i \times n_j$ matrix, and B_{jk} is a $n_j \times p_k$ matrix for all $i, j, k \in \{1, 2\}$ and some natural numbers $m_1, m_2, n_1, n_2, p_1, p_2$. **Very important:** the sizes of the blocks in the block form of a matrix always match at the boundaries (for example, the number of rows of A_{11} is the same as the number of rows of A_{12} , and the number of columns of A_{11} is the same as the number of columns of A_{21}). Then the product matrix AB also has a block form:

$$AB = \left[\begin{array}{c|c} A_{11}B_{11} + A_{12}B_{21} & A_{11}B_{12} + A_{12}B_{22} \\ \hline A_{21}B_{11} + A_{22}B_{21} & A_{21}B_{12} + A_{22}B_{22} \end{array} \right] \quad (18)$$

where each of the four blocks in (18) is matrix multiplication in its own right (the sizes of these blocks are $m_i \times p_k$). Thus, you may reduce the multiplication of the bigger matrices A and B to the multiplication of the smaller matrices A_{ij} and B_{jk} . This allows you to recursively reduce matrix multiplication to the smallest blocks, namely 1×1 , which are simply individual matrix entries.

Example 2. *Block matrices such as (17) are the same thing as ordinary matrices, they are just represented (construed) differently. For example, consider the following matrices in block form:*

$$A = \left[\begin{array}{cc|c} 1 & -1 & 2 \\ -4 & 1 & 0 \end{array} \right] \quad \text{and} \quad B = \left[\begin{array}{ccc|c} 1 & -1 & 2 & -3 \\ -4 & 0 & 3 & -6 \\ 1 & 4 & 0 & 2 \end{array} \right] \quad (19)$$

You could either think of A and B as 2×3 and 3×4 matrices (respectively) by deleting the various horizontal and vertical lines, or you could think of them as block matrices in the form (17), with:

$$\begin{aligned} A_{11} &= [1 \quad -1] & A_{12} &= [2] \\ A_{21} &= [-4 \quad 1] & A_{22} &= [0] \\ B_{11} &= \begin{bmatrix} 1 & -1 & 2 \\ -4 & 0 & 3 \end{bmatrix} & B_{12} &= \begin{bmatrix} -3 \\ -6 \end{bmatrix} \\ B_{21} &= [1 \quad 4 \quad 0] & B_{22} &= [2] \end{aligned}$$

Then you may apply formula (18) to compute the product of matrices AB by summing products of their constituent blocks A_{ij} and B_{jk} . Doing so, one obtains:

$$AB = \left[\begin{array}{ccc|c} 7 & 7 & -1 & 7 \\ -8 & 4 & -5 & 6 \end{array} \right]$$

One could also obtain the same result for AB by multiplying the matrices (19) directly, just as you would multiply a 2×3 with a 3×4 matrix. Therefore, whether we think of matrices as being in block form or not, computations yield the same answer. So block form is just a matter of perspective.

Matrix multiplication satisfies many important rules, such as:

- associativity: $A(BC) = (AB)C$

- distributivity: $A(B + C) = AB + AC$ and $(A + B)C = AC + BC$

- unit element: $AI = IA = A$, where $I = \begin{bmatrix} 1 & 0 & \dots & 0 \\ 0 & 1 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 1 \end{bmatrix}$ is called the **identity** (or unit) matrix. People sometimes write I_n instead of I if they wish to specify the identity matrix of size $n \times n$.

Given a square matrix A , we can form its powers:

$$A^p = \underbrace{AA \dots A}_{p \text{ times}}$$

Just like with numbers, we set:

$$A^0 = I = \text{the identity matrix}$$

but can we talk about negative powers? This is the point of the following definition.

Definition 1. Given a square matrix A , its **inverse** is a matrix A^{-1} with the property that:

$$\boxed{AA^{-1} = A^{-1}A = I} \tag{20}$$

Not all square matrices have inverses (those which do are called **non-singular**³ and those which do not are called **singular**) but if an inverse exists, it is unique. For example, consider the matrix:

$$A = \begin{bmatrix} 1 & 5 \\ 2 & 10 \end{bmatrix}$$

If it had an inverse $A^{-1} = \begin{bmatrix} a & b \\ c & d \end{bmatrix}$, then we would have:

$$\begin{bmatrix} 1 & 5 \\ 2 & 10 \end{bmatrix} \begin{bmatrix} a & b \\ c & d \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \tag{21}$$

As we have seen, this implies that the columns of the identity matrix, namely:

$$\begin{bmatrix} 1 \\ 0 \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} 0 \\ 1 \end{bmatrix} \tag{22}$$

³There is no ambiguity between this notion of “(non)-singularity” and the one from the previous Lecture, as we will soon see that the two notions are equivalent for square matrices.

are both linear combinations of the columns of A , namely:

$$\begin{bmatrix} 1 \\ 2 \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} 5 \\ 10 \end{bmatrix} \quad (23)$$

But this is clearly impossible, as the vectors (23) are both on the line of slope 2 in the plane, and therefore so is any linear combination of them. Meanwhile, the vectors (22) are not confined to this line, so we conclude that formula (21) cannot hold for any a, b, c, d , and therefore A is singular.

Gauss-Jordan elimination provides an algorithm for computing the inverse of a matrix A , if such an inverse exists. Start by constructing the augmented matrix:

$$\left[A \mid I \right] \quad (24)$$

where A and I are both $n \times n$. If A is non-singular, then Gauss-Jordan elimination will output:

$$\left[I \mid B \right] \quad (25)$$

for some matrix B . Now how does Gauss-Jordan elimination get you from (24) to (25)? We have seen last time that it does this by multiplying (24) on the left with another matrix C :

$$C \left[A \mid I \right] = \left[I \mid B \right]$$

But from the rules of multiplying block matrices, this is equivalent to the two matrix equalities:

$$CA = I \quad \text{and} \quad CI = B$$

Therefore, $C = CI = B$, hence $BA = I$, hence B is the sought-for inverse of A . So the right-most block of (25), which is produced by Gauss-Jordan elimination, is precisely the inverse of A . If Gauss-Jordan elimination does not produce the identity matrix in the left-most block of (25), which would be the case if there were rows without pivots at the end of the elimination algorithm, then A would not have an inverse (i.e. A would be singular).

Being able to calculate inverses of matrices is a very helpful skill, since they allow you to solve square systems of equations (i.e. systems where the number of equations is equal to the number of unknowns). Specifically, if you multiply a system $A\mathbf{v} = \mathbf{b}$ on the left by A^{-1} , then you obtain:

$$\boxed{\mathbf{v} = A^{-1}\mathbf{b}} \quad (26)$$

which is precisely the same as solving for the unknown vector \mathbf{v} . A useful tool in computing matrices is the following formula:

$$\boxed{(AB)^{-1} = B^{-1}A^{-1}} \quad (27)$$

which says that the inverse of a product (of square matrices) is the product of the inverses, but taken in reverse order. This is very meaningful, since the order in which we multiply matrices is important (a fancy way of saying this is that matrix multiplication is not a commutative operation). Finally, a nice consequence of being able to define inverse matrices is that the expression:

$$A^n \quad (28)$$

now makes sense for any integer n , if A is a non-singular square matrix. If n is negative, then you just define (28) as the $-n$ th power of the inverse matrix A^{-1} . This is just the basic requirement that the usual formulas for powers:

$$A^m A^n = A^{m+n} \quad \text{and} \quad (A^m)^n = A^{mn} \quad (29)$$

continue to hold for matrices, as they do for numbers.

Lecture 4 (February 24)

Reading: Section 2.6

Viewing: Lecture 4

Let us return to Gaussian elimination, in the particular case of a square $n \times n$ matrix. For simplicity, let us assume that the algorithm doesn't require any rows to be exchanged (which is the case for most, but not all matrices). Then under this hypothesis, the algorithm goes as follows:

- first step: add a multiple λ_{21} of the first row to the second, then a multiple λ_{31} of the first row to the third, ... then a multiple λ_{n1} of the first row to the n -th row
- second step: add a multiple λ_{32} of the second row to the third, then a multiple λ_{42} of the second row to the fourth row, ... then a multiple λ_{n2} of the second row to the n -th row
- ...
- $(n-2)$ -th step: add a multiple $\lambda_{n-1,n-2}$ of the $(n-2)$ -th row to the $(n-1)$ -th row, then add a multiple $\lambda_{n,n-2}$ of the $(n-2)$ -th row to the n -th row
- $(n-1)$ -th step: add a multiple $\lambda_{n,n-1}$ of the $(n-1)$ -th row to the n -th row.

As we have seen in Lecture 2, each step in the algorithm above involves multiplying A on the left by an elimination matrix. As a formula, this states that:

$$\underbrace{E_{n,n-1}^{(\lambda_{n,n-1})}}_{(n-1)\text{-th step}} \underbrace{E_{n,n-2}^{(\lambda_{n,n-2})} E_{n-1,n-2}^{(\lambda_{n-1,n-2})}}_{(n-2)\text{-th step}} \underbrace{E_{n,n-3}^{(\lambda_{n,n-3})} E_{n-1,n-3}^{(\lambda_{n-1,n-3})} E_{n-2,n-3}^{(\lambda_{n-2,n-3})}}_{(n-3)\text{-th step}} \cdots \underbrace{\cdots E_{n2}^{(\lambda_{n2})} \cdots E_{42}^{(\lambda_{42})} E_{32}^{(\lambda_{32})}}_{\text{second step}} \underbrace{E_{n1}^{(\lambda_{n1})} \cdots E_{31}^{(\lambda_{31})} E_{21}^{(\lambda_{21})}}_{\text{first step}} A = U \quad (30)$$

where U is a matrix in row echelon form. Because we are in the square matrix case, U is also a square matrix, so row echelon form is the same thing as **upper triangular**:

$$\begin{bmatrix} u_{11} & u_{12} & u_{13} & u_{14} & u_{15} \\ 0 & u_{22} & u_{23} & u_{24} & u_{25} \\ 0 & 0 & u_{33} & u_{34} & u_{35} \\ 0 & 0 & 0 & u_{44} & u_{45} \\ 0 & 0 & 0 & 0 & u_{55} \end{bmatrix}$$

(the matrix above is how a 5×5 upper triangular matrix looks like). The opposite notion, where a square matrix has all its non-zero entries on or underneath the diagonal, is called **lower triangular**.

Fact 4. Each elimination matrix $E_{ij}^{(\lambda)}$ is a lower triangular matrix, and its inverse is given by:

$$\left(E_{ij}^{(\lambda)}\right)^{-1} = E_{ij}^{(-\lambda)} \quad (31)$$

Indeed, let's establish formula (31) in the case of 2×2 matrices (so $i = 2, j = 1$):

$$\begin{bmatrix} 1 & 0 \\ \lambda & 1 \end{bmatrix}^{-1} = \begin{bmatrix} 1 & 0 \\ -\lambda & 1 \end{bmatrix}$$

which is an easy computation to prove by hand. The proof of (31) for general $n \times n$ matrices is similar, and you can think about it, if you'd like.

Using formula (31), let's convert (30) into a formula for A . The first step to doing so is to multiply equation (30) on the left by the inverse of $E_{n,n-1}^{(\lambda_{n,n-1})}$, which according to (31) is precisely $E_{n,n-1}^{(-\lambda_{n,n-1})}$:

$$\begin{aligned} \left(E_{n,n-2}^{(\lambda_{n,n-2})} E_{n-1,n-2}^{(\lambda_{n-1,n-2})} \right) \left(E_{n,n-3}^{(\lambda_{n,n-3})} E_{n-1,n-3}^{(\lambda_{n-1,n-3})} E_{n-2,n-3}^{(\lambda_{n-2,n-3})} \right) \cdots \left(E_{n1}^{(\lambda_{n1})} \cdots E_{31}^{(\lambda_{31})} E_{21}^{(\lambda_{21})} \right) A = \\ = \left(E_{n,n-1}^{(\lambda_{n,n-1})} \right)^{-1} U = E_{n,n-1}^{(-\lambda_{n,n-1})} U \end{aligned}$$

Then multiply the equation above on the left by the inverse of $E_{n,n-2}^{(\lambda_{n,n-2})} E_{n-1,n-2}^{(\lambda_{n-1,n-2})}$, which according to (27) and (31) is precisely $E_{n-1,n-2}^{(-\lambda_{n-1,n-2})} E_{n,n-2}^{(-\lambda_{n,n-2})}$. We obtain:

$$\begin{aligned} \left(E_{n,n-3}^{(\lambda_{n,n-3})} E_{n-1,n-3}^{(\lambda_{n-1,n-3})} E_{n-2,n-3}^{(\lambda_{n-2,n-3})} \right) \cdots \left(E_{n1}^{(\lambda_{n1})} \cdots E_{31}^{(\lambda_{31})} E_{21}^{(\lambda_{21})} \right) A = \\ = \left(E_{n,n-2}^{(\lambda_{n,n-2})} E_{n-1,n-2}^{(\lambda_{n-1,n-2})} \right)^{-1} \left(E_{n,n-1}^{(-\lambda_{n,n-1})} \right) U = \left(E_{n-1,n-2}^{(-\lambda_{n-1,n-2})} E_{n,n-2}^{(-\lambda_{n,n-2})} \right) \left(E_{n,n-1}^{(-\lambda_{n,n-1})} \right) U \end{aligned}$$

We repeat this multiplication procedure until we get rid of all the matrices to the left of A , yielding:

$$A = \left(E_{21}^{(-\lambda_{21})} E_{31}^{(-\lambda_{31})} \cdots E_{n1}^{(-\lambda_{n1})} \right) \cdots \left(E_{n-1,n-2}^{(-\lambda_{n-1,n-2})} E_{n,n-2}^{(-\lambda_{n,n-2})} \right) \left(E_{n,n-1}^{(-\lambda_{n,n-1})} \right) U \quad (32)$$

Now here's the kicker: each $E_{ij}^{(\lambda)}$ is a lower triangular matrix, with all 1's on the diagonal. It is easy to see that a product of such lower triangular matrices is still lower triangular, for example:

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

Therefore, the product of elimination matrices in (32), namely:

$$L = \left(E_{21}^{(-\lambda_{21})} E_{31}^{(-\lambda_{31})} \cdots E_{n1}^{(-\lambda_{n1})} \right) \cdots \left(E_{n-1,n-2}^{(-\lambda_{n-1,n-2})} E_{n,n-2}^{(-\lambda_{n,n-2})} \right) \left(E_{n,n-1}^{(-\lambda_{n,n-1})} \right) \quad (33)$$

is still lower triangular, and moreover its diagonal entries are all 1. We conclude that:

$$\boxed{A = LU} \quad (34)$$

where L is a lower triangular matrix and U is an upper triangular matrix. Therefore, we summarize this whole discussion by saying that: **Gaussian elimination on a square matrix is equivalent to writing it as a product of a lower and an upper triangular matrix** (assuming there are no row exchanges). Moreover, if A is non-singular, then the matrices L and U in (34) are unique.

Let's do an example, namely Gaussian elimination of the matrix:

$$A = \begin{bmatrix} 2 & 4 & 1 \\ -4 & -5 & 0 \\ -2 & 5 & 6 \end{bmatrix}$$

First, let's add 2 times the first row to the second row, and we get:

$$E_{21}^{(2)} A = \begin{bmatrix} 1 & 0 & 0 \\ 2 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 2 & 4 & 1 \\ -4 & -5 & 0 \\ -2 & 5 & 6 \end{bmatrix} = \begin{bmatrix} 2 & 4 & 1 \\ 0 & 3 & 2 \\ -2 & 5 & 6 \end{bmatrix}$$

Then, let us add 1 times the first row to the third row, and we get:

$$E_{31}^{(1)} E_{21}^{(2)} A = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 1 & 0 & 1 \end{bmatrix} \begin{bmatrix} 2 & 4 & 1 \\ 0 & 3 & 2 \\ -2 & 5 & 6 \end{bmatrix} = \begin{bmatrix} 2 & 4 & 1 \\ 0 & 3 & 2 \\ 0 & 9 & 7 \end{bmatrix}$$

Finally, we add -3 times the second row to the third row, and we get:

$$E_{32}^{(-3)} E_{31}^{(1)} E_{21}^{(2)} A = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & -3 & 1 \end{bmatrix} \begin{bmatrix} 2 & 4 & 1 \\ 0 & 3 & 2 \\ 0 & 9 & 7 \end{bmatrix} = \begin{bmatrix} 2 & 4 & 1 \\ 0 & 3 & 2 \\ 0 & 0 & 1 \end{bmatrix} =: U \quad (35)$$

Since the matrix on the right is upper triangular, hence in row echelon form, we are done with Gaussian elimination! To write A in the form (34), we need to isolate A in the left-hand side of (35). To do so, we must successively multiply the equation above with the inverses of the elimination matrices $E_{32}^{(-3)}$, $E_{31}^{(1)}$ and $E_{21}^{(2)}$, in this order. To this end, we have:

$$\begin{aligned} E_{31}^{(1)} E_{21}^{(2)} A &= E_{32}^{(3)} U \\ E_{21}^{(2)} A &= E_{31}^{(-1)} E_{32}^{(3)} U \\ A &= E_{21}^{(-2)} E_{31}^{(-1)} E_{32}^{(3)} U \end{aligned}$$

If we set:

$$L = E_{21}^{(-2)} E_{31}^{(-1)} E_{32}^{(3)} = \begin{bmatrix} 1 & 0 & 0 \\ -2 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ -1 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 3 & 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ -2 & 1 & 0 \\ -1 & 3 & 1 \end{bmatrix} \quad (36)$$

then we conclude that $A = LU$ holds with U and L given by (35) and (36), respectively.

A **diagonal** matrix is a square matrix, all of whose non-zero entries are on the diagonal:

$$D = \begin{bmatrix} d_1 & 0 & \dots & 0 \\ 0 & d_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & d_n \end{bmatrix} \quad (37)$$

A diagonal matrix is invertible if and only if all the diagonal entries d_i 's are non-zero, and in this case its inverse is given by:

$$D^{-1} = \begin{bmatrix} \frac{1}{d_1} & 0 & \dots & 0 \\ 0 & \frac{1}{d_2} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \frac{1}{d_n} \end{bmatrix} \quad (38)$$

VERY IMPORTANT: If a matrix is not diagonal, then its inverse is **NOT** obtained by simply inverting all of its entries; for example you could think about the matrix all of whose entries are 1.

In the LU factorization (34), we note that L has 1's on the diagonal, while U is not required to. This is simply a matter of convention because if you'd like, you can write U as a diagonal matrix times an upper triangular matrix with 1's on the diagonal, for example:

$$\begin{bmatrix} 2 & 4 & 1 \\ 0 & 3 & 2 \\ 0 & 0 & 1 \end{bmatrix} = \begin{bmatrix} 2 & 0 & 0 \\ 0 & 3 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 2 & \frac{1}{2} \\ 0 & 1 & \frac{2}{3} \\ 0 & 0 & 1 \end{bmatrix} \quad (39)$$

Basically, this is done by writing $U = D\tilde{U}$, where D is the diagonal matrix with the same diagonal entries as the upper triangular matrix U . Then \tilde{U} is forced to equal $D^{-1}U$, with the inverse D^{-1} given by formula (38). In these circumstances, \tilde{U} will be upper triangular with 1's on the diagonal.

Definition 2. *The LU factorization of a matrix A consists of the unique matrices U and L (the former upper and the latter lower triangular with 1's on the diagonal) such that $A = LU$ holds.*

Similarly, the LDU factorization of a matrix A consists of the unique matrices U , D and L (where D is diagonal, while U and L are upper/lower triangular with 1's on the diagonal) such that $A = LDU$.

Note that if you have the LU factorization, you can easily get the LDU factorization, essentially by applying procedure (39). Let's end with a remark: suppose you do Gaussian elimination on a square matrix A , and you bring it in the form (32). Now you may be taken aback by the task of actually having to perform the multiplication (33) in order to compute L . However, a great fact about lower triangular matrices is the following formula:

$$\left(E_{21}^{(-\lambda_{21})} \dots E_{n1}^{(-\lambda_{n1})} \right) \dots \left(E_{n-1,n-2}^{(-\lambda_{n-1,n-2})} E_{n,n-2}^{(-\lambda_{n,n-2})} \right) \left(E_{n,n-1}^{(-\lambda_{n,n-1})} \right) = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ -\lambda_{21} & 1 & 0 & 0 & 0 \\ \vdots & \vdots & \ddots & 0 & 0 \\ -\lambda_{n-1,1} & -\lambda_{n-1,2} & \dots & 1 & 0 \\ -\lambda_{n1} & -\lambda_{n2} & \dots & -\lambda_{n,n-1} & 1 \end{bmatrix}$$

Note that this formula only holds because of the specific order in which the matrices $E_{ij}^{(-\lambda)}$ were multiplied in the left-hand side. If you had multiplied the matrices in (almost) any other order, you would not have obtained a result as nice as the right-hand side.

Lecture 5 (February 26)

Reading: Section 2.7

Viewing: Lecture 5

A **permutation matrix** is a square matrix which has a single 1 on each row and column, and 0 everywhere else. For example, the following is a 4×4 permutation matrix:

$$\begin{bmatrix} 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix} \quad (40)$$

For all $i \in \{1, \dots, n\}$, look at the single 1 on the i -th row, and let $\sigma(i) \in \{1, \dots, n\}$ be the column where that 1 lies. By the very definition of a permutation matrix, the numbers $\sigma(1), \dots, \sigma(n)$ are all different, hence they determine a permutation of the numbers $1, \dots, n$. For example, the permutation corresponding to the matrix (40) is $(2, 1, 4, 3)$. This is why permutation matrices are called this way, because they are in one-to-one correspondence with usual permutations. Therefore:

$$\text{the number of } n \times n \text{ permutation matrices is } n! = 1 \cdot 2 \cdot 3 \cdot \dots \cdot (n-1) \cdot n \quad (41)$$

The particular matrix P_{ij} we encountered in Lecture 2 (when discussing the matrix meaning of the various steps in Gaussian elimination) corresponds to the permutation of the numbers $1, \dots, n$ that switches i and j , but keeps all other numbers unchanged.

Suppose you do Gaussian elimination on a square matrix A , always taking care that the pivots are in the right place before doing elimination. For example, we may have:

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

Since the pivot on the first row is to the right of the pivots on the other two rows, let us remedy the situation by exchanging the roles of the first and third rows. As we have seen in Lecture 2, this is achieved by multiplying A on the left with the permutation matrix P_{13} :

$$P_{13}A = \begin{bmatrix} 1 & 2 & -1 \\ 1 & 1 & 2 \\ 0 & 0 & 3 \end{bmatrix}$$

Then we do elimination, by subtracting the first row from the second one:

$$E_{21}^{(-1)}P_{13}A = \begin{bmatrix} 1 & 2 & -1 \\ 0 & -1 & 3 \\ 0 & 0 & 3 \end{bmatrix} =: U$$

The right-hand side is now in row echelon form, i.e. upper triangular. We may move any product of E 's to the right-hand side by multiplying the equation above with their inverses, thus obtaining:

$$P_{13}A = E_{21}^{(1)}U$$

Since $E_{21}^{(1)}$ is a lower triangular matrix and P_{13} is a permutation matrix, this is a particular example of the following general formula, which holds for any non-singular matrix A :

$$\boxed{PA = LU} \tag{42}$$

where P is a permutation matrix, L is lower triangular and U is upper triangular. As we have seen in the example above, P is obtained from A by first rearranging its rows in such a way that the pivots are in the standard order (i.e. ordered left-to-right as we go from top-to-bottom). Then L and U are obtained by the usual Gaussian elimination we described in the previous lecture.

Remark. Look up to where the word “before” is underlined. If we replace this word by “after”, and do Gaussian elimination by first applying elimination matrices and then applying row exchanges, then the factorization we get is $A = LPU$, where P is a permutation matrix. This is a variant of the $PA = LU$ factorization described above.

Another feature of matrices is the operation of taking **transposes**. Specifically, if A is an $m \times n$ matrix, its transpose A^T is an $n \times m$ matrix obtained by switching the roles of rows and columns:

$$\text{if } A = \begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{bmatrix} \quad \text{then } A^T = \begin{bmatrix} 1 & 4 \\ 2 & 5 \\ 3 & 6 \end{bmatrix}$$

So the first row of A becomes the first column of A^T etc. In terms of entries, the ij entry of A is the ji entry of A^T , for all i and j . Transposes follow certain rules:

$$\boxed{(A + B)^T = A^T + B^T} \qquad \boxed{(AB)^T = B^T A^T} \tag{43}$$

$$\boxed{(A^T)^T = A} \qquad \boxed{(A^{-1})^T = (A^T)^{-1}} \tag{44}$$

(the last property requires A to be square and non-singular, or we couldn't talk about its inverse).

Transposes are related to a lot of things, for example, the dot product of vectors. Compare:

$$\begin{array}{l} \text{the dot product} \\ \text{the matrix multiplication} \end{array} \quad \begin{array}{l} \begin{bmatrix} a_1 \\ \vdots \\ a_n \end{bmatrix} \cdot \begin{bmatrix} b_1 \\ \vdots \\ b_n \end{bmatrix} = a_1 b_1 + \cdots + a_n b_n \\ \begin{bmatrix} a_1 \\ \vdots \\ a_n \end{bmatrix}^T \begin{bmatrix} b_1 \\ \vdots \\ b_n \end{bmatrix} = [a_1 \ \dots \ a_n] \begin{bmatrix} b_1 \\ \vdots \\ b_n \end{bmatrix} = [a_1 b_1 + \cdots + a_n b_n] \end{array}$$

In other words, given two vectors \mathbf{v} and \mathbf{w} written in column form, the only entry of the 1×1 matrix $\mathbf{v}^T \mathbf{w}$ is the same number as the dot product $\mathbf{v} \cdot \mathbf{w}$.

Another instance of transposes comes about when studying **symmetric** matrices, namely square matrices which are equal to their own transposes:

$$\boxed{S \text{ is called symmetric if } S = S^T} \tag{45}$$

For example, the following matrix is symmetric:

$$\begin{bmatrix} 2 & 5 & -1 \\ 5 & 4 & 8 \\ -1 & 8 & 0 \end{bmatrix}$$

because it is unchanged upon reflecting it across its diagonal. All diagonal matrices (37) are symmetric. More generally, here's a general recipe to obtain symmetric matrices:

$$\boxed{\text{for any matrix } A, \text{ the matrix } S = A^T A \text{ is symmetric}} \quad (46)$$

which holds even if the matrix A is not square. To prove this, note that (using (43) and (44)):

$$S^T = (A^T A)^T = A^T (A^T)^T = A^T A = S$$

Finally, let us note an important reason why symmetric matrices are important from a computational viewpoint. The LU factorization of a symmetric matrix takes a particularly nice form:

$$\boxed{S = LDL^T} \quad (47)$$

where L is a lower triangular matrix with 1's on the diagonal, and D is a diagonal matrix. So in other words, symmetric matrices are precisely those whose LDU factorization has the property that $U = L^T$. You can actually prove that (47) holds for any non-singular symmetric matrix S which admits a LDU factorization, just by using the uniqueness of the LDU factorization. Indeed:

$$S = LDU \xrightarrow{(43)} S^T = U^T D^T L^T$$

It is easy to see that U^T is lower triangular with 1's on the diagonal, L^T is upper triangular with 1's on the diagonal, and $D^T = D$ is diagonal. Since $S = S^T$, the uniqueness of the LDU factorization of $S = S^T$ implies that $U^T = L$ and $L^T = U$ (these two conditions are actually equivalent).

Lecture 6 (March 1)

Reading: Section 3.1

Viewing: Lectures 5-6

A key idea of linear algebra is to think of a line/plane/space as a **vector space**. Specifically, this means that points of n -dimensional space:

$$\mathbb{R}^n = \text{set of all } n\text{-tuples of real numbers} = \{(x_1, \dots, x_n), x_1, \dots, x_n \text{ are real numbers}\} \quad (48)$$

can be thought of as column vectors, according to the correspondence:

$$(x_1, \dots, x_n) \leftrightarrow \begin{bmatrix} x_1 \\ \vdots \\ x_n \end{bmatrix}$$

For example, \mathbb{R} is the line, \mathbb{R}^2 is the plane, \mathbb{R}^3 is usual space. But the key thing about thinking of these as sets of vectors is that vectors can be added, and vectors can be multiplied with scalars:

$$\begin{bmatrix} x_1 \\ \vdots \\ x_n \end{bmatrix} + \begin{bmatrix} y_1 \\ \vdots \\ y_n \end{bmatrix} = \begin{bmatrix} x_1 + y_1 \\ \vdots \\ x_n + y_n \end{bmatrix} \quad \text{and} \quad \alpha \begin{bmatrix} x_1 \\ \vdots \\ x_n \end{bmatrix} = \begin{bmatrix} \alpha x_1 \\ \vdots \\ \alpha x_n \end{bmatrix} \quad (49)$$

Definition 3. A **vector space** is a set V , together with a notion of how to add two elements in V , and a notion of how to multiply an element of V by a scalar. In other words:

$$\text{for any } \mathbf{v}, \mathbf{w} \in V, \text{ there must be a rule for defining } \mathbf{v} + \mathbf{w} \in V \quad (50)$$

$$\text{for any } \mathbf{v} \in V, \alpha \in \mathbb{R}, \text{ there must be a rule for defining } \alpha \mathbf{v} \in V \quad (51)$$

This may seem abstract, but I will now spell out the main examples of vector spaces we will encounter in this course. First of all, \mathbb{R}^n is a vector space, because of (49): the sum of vectors is a well-defined vector, and a vector times a scalar is a well-defined vector. More generally, we have:

Definition 4. If V is a vector space, then a subset $S \subset V$ is called a **subspace** if:

$$\text{for any } \mathbf{v}, \mathbf{w} \in S, \text{ then } \mathbf{v} + \mathbf{w} \in S \quad (52)$$

$$\text{for any } \mathbf{v} \in S, \alpha \in \mathbb{R}, \text{ then } \alpha \mathbf{v} \in S \quad (53)$$

where the notions “ $\mathbf{v} + \mathbf{w}$ ” and “ $\alpha \mathbf{v}$ ” are well-defined, because S is a subset of a vector space V .

By combining (52) and (53), we have the more general fact for any subspace S of a vector space:

$$\text{for all } \mathbf{v}_1, \dots, \mathbf{v}_k \in S \text{ and } \alpha_1, \dots, \alpha_k \in \mathbb{R}$$

then:

$$\boxed{\alpha_1 \mathbf{v}_1 + \dots + \alpha_k \mathbf{v}_k \text{ lies in } S} \quad (54)$$

For any choice of scalars $\alpha_1, \dots, \alpha_k$, the vector in (54) will be called a **linear combination** of the vectors $\mathbf{v}_1, \dots, \mathbf{v}_k$.

In this course, we will be interested in subspaces $S \subset \mathbb{R}^n$. There is a mind-boggling infinity of subsets of \mathbb{R}^n , but only a select few of them will turn out to be subspaces. In the case $n = 1$, there are only two subsets $S \subset \mathbb{R}$ which are subspaces: $S = \{0\}$ (the zero subspace) and $S = \mathbb{R}$ itself (the whole line). In the case $n = 2$, the subsets $S \subset \mathbb{R}^2$ which are subspaces are of three types:

- $S = \{0\}$ (the zero subspace)
- $S =$ any line passing through the origin
- $S = \mathbb{R}^2$ (the whole plane)

Indeed, if a subspace $S \subset \mathbb{R}^2$ contains any non-zero vector, then it contains the whole line passing through the origin and that vector, due to (53). And if S contains two non-zero vectors not on the same line, then it contains the entire plane, due to (54). Based on this pattern, you can probably guess that the subsets $S \subset \mathbb{R}^3$ which are subspaces are: the zero subspace, any line passing through the origin, any plane passing through the origin, and the whole space.

Now comes the point where we explain what all of this has to do with linear algebra, specifically with solving systems of equations like $A\mathbf{v} = \mathbf{b}$. We start from the following definition.

Definition 5. *The **column space** of an $m \times n$ matrix A is the subspace:*

$$C(A) \subset \mathbb{R}^m \tag{55}$$

spanned by the columns of A (the matrix A has n columns).

*Here and throughout our course, the subspace **spanned** by any vectors $\mathbf{v}_1, \dots, \mathbf{v}_k$ is defined to be the set of all linear combinations (54) of those vectors.*

Exercise: prove that the subspace spanned by any number of vectors in \mathbb{R}^m is a subspace, in the sense of Definition 4. While you ponder this question, here is an explicit example of column space:

$$\text{if } A = \begin{bmatrix} 1 & 2 \\ 1 & 4 \\ 1 & 6 \end{bmatrix} \text{ then } C(A) = \left\{ \alpha \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} + \beta \begin{bmatrix} 2 \\ 4 \\ 6 \end{bmatrix}, \text{ for all scalars } \alpha \text{ and } \beta \right\}$$

So in other words, the column space in the example above is the plane traced out by the vectors

$\begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}$ and $\begin{bmatrix} 2 \\ 4 \\ 6 \end{bmatrix}$ in space. However:

$$\text{if } B = \begin{bmatrix} 1 & 2 \\ 2 & 4 \\ 3 & 6 \end{bmatrix} \text{ then } C(B) = \left\{ \alpha \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix} + \beta \begin{bmatrix} 2 \\ 4 \\ 6 \end{bmatrix}, \text{ for all scalars } \alpha \text{ and } \beta \right\} =$$

$$= \left\{ \gamma \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix}, \text{ for all scalars } \gamma \right\} \quad (56)$$

so the column space of B is the line determined by the vector $\begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix}$. This is because the second

column $\begin{bmatrix} 2 \\ 4 \\ 6 \end{bmatrix}$ is already on the line spanned by the first column. Taking linear combinations of vectors that lie on the same line passing through the origin will still remain on the given line:

$$\alpha \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix} + \beta \begin{bmatrix} 2 \\ 4 \\ 6 \end{bmatrix} = \alpha \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix} + 2\beta \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix} = (\alpha + 2\beta) \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix}$$

for any scalars α and β .

Fact 5. *A system of equations $A\mathbf{v} = \mathbf{b}$ has solutions if and only if $\mathbf{b} \in C(A)$.*

The fact above is simply a restatement of Fact 2. It is simply saying that if there is a solution \mathbf{v} to the system $A\mathbf{v} = \mathbf{b}$, then (5) implies that \mathbf{b} is a linear combination of the columns of A . So the column space of a matrix precisely consists of those vectors \mathbf{b} which can appear as the right-hand sides of systems of linear equations $A\mathbf{v} = \mathbf{b}$.

For example, does the system:

$$\begin{bmatrix} 1 & 2 \\ 2 & 4 \\ 3 & 6 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} \quad (57)$$

have any solutions? The answer is no, because we have already seen in (56) that the column space of the matrix in question is the line spanned by the vector $\begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix}$. Since the vector $\begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}$ is not

on this line (as it is not a multiple of $\begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix}$), then Fact 5 implies that the system (57) has no solutions.

Lecture 7 (March 3)

Reading: Section 3.2

Viewing: Lectures 6-7

Consider the subspace $S \subset \mathbb{R}^3$ spanned by the vectors:

$$\begin{bmatrix} 1 \\ 1 \\ -2 \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} 5 \\ -2 \\ -3 \end{bmatrix} \quad (58)$$

These two vectors are not on the same line (passing through the origin) so we conclude that S is a plane. As such, S is cut out by a single equation in the plane, and it is not hard to guess what this equation is:

$$S = \left\{ \begin{bmatrix} x \\ y \\ z \end{bmatrix} \text{ such that } x + y + z = 0 \right\} \quad (59)$$

Indeed, both vectors in (58) have the property that the sum of their entries is 0. Moreover, any linear combination of those two vectors has the property that the sum of its entries is 0, and this allows us to conclude (59). A more linear algebra way to say this is to consider the 1×3 matrix:

$$A = [1 \quad 1 \quad 1]$$

(the entries of A are precisely the coefficients of x , y and z in the equation of the plane (59)) and to observe the fact that:

$$A \begin{bmatrix} 1 \\ 1 \\ -2 \end{bmatrix} = [0] \quad \text{and} \quad A \begin{bmatrix} 5 \\ -2 \\ -3 \end{bmatrix} = [0]$$

Moreover, taking appropriate linear combinations of these two equations allows us to conclude that:

$$A \left(\alpha \begin{bmatrix} 1 \\ 1 \\ -2 \end{bmatrix} + \beta \begin{bmatrix} 5 \\ -2 \\ -3 \end{bmatrix} \right) = [0]$$

for any scalars α and β . We can summarize the formula above by saying that the plane S is precisely what people call the nullspace of A , as in the following definition.

Definition 6. The **nullspace** of an $m \times n$ matrix A is the subspace:

$$N(A) \subset \mathbb{R}^n \quad (60)$$

consisting of those vectors $\mathbf{v} \in \mathbb{R}^n$ such that $A\mathbf{v} = \mathbf{0}$.

Remark. A little quirk concerning our notation: the right-hand side of the equation $A\mathbf{v} = \mathbf{0}$ should be an $m \times 1$ column vector, specifically the zero vector $\begin{bmatrix} 0 \\ \vdots \\ 0 \end{bmatrix}$. People usually abbreviate the zero vector by using the notation “0”, although it is still technically a vector.

The fact that the nullspace of the matrix A satisfies the requirements (52) and (53) of being a subspace follows from the easy to prove relations:

$$A\mathbf{v} = 0 \text{ and } A\mathbf{w} = 0 \Rightarrow A(\mathbf{v} + \mathbf{w}) = 0$$

$$A\mathbf{v} = 0 \text{ and } \alpha \in \mathbb{R} \Rightarrow A(\alpha\mathbf{v}) = 0$$

So computing the nullspace of a matrix A boils down to solving the equation $A\mathbf{v} = 0$. However, many matrices have the same nullspace. For example, if:

$$\boxed{A = LR} \tag{61}$$

where L is an invertible matrix (such as the product of elimination/permutation/diagonal matrices that you encounter in Gaussian and Gauss-Jordan elimination, hint-hint) then:

$$\boxed{N(A) = N(R)} \tag{62}$$

which is essentially saying that:

$$A\mathbf{v} = LR\mathbf{v} = 0 \quad \text{is equivalent to} \quad R\mathbf{v} = 0$$

(the equivalence is showed by multiplying on the left with L^{-1}). Gauss-Jordan elimination gives us a recipe for writing A in the form (61), where R is in reduced row echelon form, i.e.:

- the pivots are all 1
- the pivots go to the right as we read the rows from top to bottom
- all the entries above a pivot are zeroes

For example, the following matrix is in reduced row echelon form:

$$R = \begin{bmatrix} \boxed{1} & 3 & 0 & 7 \\ 0 & 0 & \boxed{1} & 2 \end{bmatrix}$$

(as usual, the pivots are in boxes). A column which has a pivot is called a (unsurprisingly) **pivot column** and all other columns are called **free columns**. In the example of R above, the first and third are pivot columns, and the second and fourth are free columns. Explicitly, we have:

$$R \begin{bmatrix} a \\ b \\ c \\ d \end{bmatrix} = 0 \quad \text{means that} \quad \begin{bmatrix} a + 3b + 7d \\ c + 2d \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix} \quad \text{i.e.} \quad \begin{cases} a = -3b - 7d \\ c = -2d \end{cases}$$

So here's the catch: for any choice of the numbers b and d , the equalities above tell you what a and c have to be in order for the vector:

$$\begin{bmatrix} a \\ b \\ c \\ d \end{bmatrix}$$

to lie in the nullspace of R . Thus, b and d are **free variables**, and they precisely correspond to the free columns of R . We conclude that the nullspace of R can be explicitly described as:

$$N(R) = \left\{ \begin{bmatrix} -3b - 7d \\ b \\ -2d \\ d \end{bmatrix} \text{ for all } b, d \text{ real numbers} \right\}$$

The set above is simply a plane in 4-dimensional space. If you want to describe this plane as being spanned by two vectors, then all you need to do is to assign concrete values to the free variables. A reasonable choice is to set $(b, d) = (1, 0)$ or $(0, 1)$, and this implies that:

$$N(R) \text{ is spanned by the vectors } \begin{bmatrix} -3 \\ 1 \\ 0 \\ 0 \end{bmatrix} \text{ and } \begin{bmatrix} -7 \\ 0 \\ -2 \\ 1 \end{bmatrix}$$

Therefore, in general, here is the algorithm for computing the nullspace of any $m \times n$ matrix A (or equivalently, for finding all solutions to the system of equations $A\mathbf{v} = 0$):

- Write A in reduced row echelon form, and call the resulting matrix R
- Identify the pivot columns and the free columns of R ; the pivot/free variables will be those x_i among x_1, \dots, x_n such that the i -th column of R is a pivot/free column
- The nullspace of A , which coincides with the nullspace of R , is spanned by those vectors (x_1, \dots, x_n) where a given free variable is 1, all other free variables are 0, and all the pivot variables are determined by the free variables and the equations:

$$R \begin{bmatrix} x_1 \\ \vdots \\ x_n \end{bmatrix} = 0$$

It is easy to solve for the pivot variables (in terms of the free variables) in the equation above by back-substitution, precisely because R is in reduced row echelon form. The aforementioned ease is one of the main reasons why people bother with reduced row echelon form in the first place.

Lecture 8 (March 5)

Reading: Section 3.3

Viewing: Lecture 8

We will now put together the ideas of the previous two lectures in order to find the complete (a.k.a. general) solution \mathbf{v} of the system:

$$A\mathbf{v} = \mathbf{b} \tag{63}$$

for any $m \times n$ matrix A . First of all, as we have seen in Fact 5, there are no solutions if the $m \times 1$ vector \mathbf{b} is not in the column space $C(A)$. So we henceforth assume that $\mathbf{b} \in C(A)$, which is just saying that there exists a particular solution:

$$A\mathbf{v}_{\text{particular}} = \mathbf{b} \tag{64}$$

This $n \times 1$ vector $\mathbf{v}_{\text{particular}}$ is a solution of (63), but it may not be the only one. In fact, for any other solution \mathbf{v} of the original system (63), we may subtract the two equations above and obtain:

$$A(\mathbf{v} - \mathbf{v}_{\text{particular}}) = 0$$

So the difference $\mathbf{v} - \mathbf{v}_{\text{particular}}$ lies in the nullspace of A . This immediately leads to the following.

Fact 6. *The general solution of the system (63) is given by:*

$$\boxed{\mathbf{v}_{\text{general}} = \mathbf{v}_{\text{particular}} + \mathbf{w}_{\text{general}}} \tag{65}$$

where:

- $\mathbf{v}_{\text{particular}}$ is a particular solution of (63), which exists if and only if $\mathbf{b} \in C(A)$
- $\mathbf{w}_{\text{general}}$ is any element of $N(A)$, i.e. the general solution to $A\mathbf{w}_{\text{general}} = 0$

The fact above is a general mathematical feature of linear equations. For example, it also holds in the theory of linear ordinary differential equations, as you can see in 18.03.

Reduced row echelon forms allow you to deal with both bullets in Fact 6 at the same time. Let me show you how to do this in an example. Consider the system:

$$\begin{bmatrix} 2 & 6 & 2 & 2 & -1 \\ 2 & 6 & 1 & -1 & 2 \\ 3 & 9 & -1 & -9 & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \end{bmatrix} = \begin{bmatrix} -7 \\ 6 \\ -6 \end{bmatrix} \tag{66}$$

Let us now apply Gauss-Jordan elimination to the augmented matrix of the system:

$$\left[\begin{array}{ccccc|c} 2 & 6 & 2 & 2 & -1 & -7 \\ 2 & 6 & 1 & -1 & 2 & 6 \\ 3 & 9 & -1 & -9 & 1 & -6 \end{array} \right] \rightsquigarrow \left[\begin{array}{ccccc|c} \boxed{1} & 3 & 0 & -2 & 0 & -3 \\ 0 & 0 & \boxed{1} & 3 & 0 & 2 \\ 0 & 0 & 0 & 0 & \boxed{1} & 5 \end{array} \right]$$

(I will leave you the joy of going through all the steps in Gauss-Jordan elimination above, as you did in Lecture 2). The matrix on the right is in reduced row echelon form (the pivots are boxed) and it corresponds to the system:

$$\begin{bmatrix} \boxed{1} & 3 & 0 & -2 & 0 \\ 0 & 0 & \boxed{1} & 3 & 0 \\ 0 & 0 & 0 & 0 & \boxed{1} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \end{bmatrix} = \begin{bmatrix} -3 \\ 2 \\ 5 \end{bmatrix} \quad (67)$$

By Fact 3, finding the general solution of the system (66) is equivalent to finding the solution of the system (67) (note: the right-hand sides of (66) and (67) are different, because the right-most column must also undergo Gauss-Jordan elimination together with the rest of the matrix) so let us solve the latter. To do so, let us look at the coefficient matrix:

$$R = \begin{bmatrix} \boxed{1} & 3 & 0 & -2 & 0 \\ 0 & 0 & \boxed{1} & 3 & 0 \\ 0 & 0 & 0 & 0 & \boxed{1} \end{bmatrix}$$

and identify the pivot columns (these are columns 1,3 and 5 in the example at hand) and the free columns (these are columns 2 and 4 in the example at hand). Therefore, the pivot variables are x_1, x_3, x_5 and the free variables are x_2, x_4 . Now we make the following observation:

For any choice of free variables x_2, x_4 there exists a unique choice of pivot variables x_1, x_3, x_5 which solve the system (67)

So we keep x_2, x_4 as arbitrary numbers, and solve for x_1, x_3, x_5 in (67) by back substitution:

$$\begin{cases} x_1 + 3x_2 - 2x_4 = -3 \\ x_3 + 3x_4 = 2 \\ x_5 = 5 \end{cases} \rightsquigarrow \begin{cases} x_1 = -3 - 3x_2 + 2x_4 \\ x_3 = 2 - 3x_4 \\ x_5 = 5 \end{cases}$$

So we conclude that the general solution to the system (67) (which is the same as the general solution to the system (66), by the very nature of Gauss-Jordan elimination) is:

$$\begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \end{bmatrix}_{\text{general}} = \begin{bmatrix} -3 - 3a + 2b \\ a \\ 2 - 3b \\ b \\ 5 \end{bmatrix} \quad (68)$$

(we relabeled $x_2 = a$ and $x_4 = b$ to emphasize the fact that they are arbitrary real numbers). So what if you didn't need the general solution of (66) and you were satisfied with a particular solution? One easy way to get a particular solution is to just plug in $a = b = 0$, thus obtaining:

$$\begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \end{bmatrix}_{\text{particular}} = \begin{bmatrix} -3 \\ 0 \\ 2 \\ 0 \\ 5 \end{bmatrix} \quad (69)$$

You could also have gotten this by setting all the free variables equal to 0 in (67) and solving for the pivot variables by using back substitution. Similarly, if you subtract the particular solution (69) from the general solution (68), you get:

$$\begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \end{bmatrix}_{\text{general}} - \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \end{bmatrix}_{\text{particular}} = \begin{bmatrix} -3a + 2b \\ a \\ -3b \\ b \\ 0 \end{bmatrix} \quad (70)$$

The right-hand side is just the general element of the nullspace of the matrix R .

Let us just mention one thing which could have gone wrong in the procedure above, and how to remedy it. It could have been that the reduced row echelon form of the coefficient matrix had a full row of zeroes. For example, suppose that instead of the system (67) we would have encountered:

$$\begin{bmatrix} \boxed{1} & 3 & 0 & -2 & 0 \\ 0 & 0 & \boxed{1} & 3 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \end{bmatrix} = \begin{bmatrix} -3 \\ 2 \\ 5 \end{bmatrix} \quad (71)$$

(when doing elimination, it is customary to put all the zero rows at the very bottom of the matrix, which we can achieve by suitable row exchanges). It is easy to see that the system above has no solutions, because equating the third rows of equality (71) to each other would force us to have $0 = 5$, which is impossible. However, if the system were instead:

$$\begin{bmatrix} \boxed{1} & 3 & 0 & -2 & 0 \\ 0 & 0 & \boxed{1} & 3 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \end{bmatrix} = \begin{bmatrix} -3 \\ 2 \\ 0 \end{bmatrix} \quad (72)$$

then we do have solutions. To find these solutions, just throw out the zero row at the bottom of the equality (72) and proceed as we did on the previous page. The upshot is that for any matrix R in reduced row echelon form:

a system $Rv = c$ has solutions if and only if the vector c has entry 0 corresponding to every full row of zeroes of R

Definition 7. The **rank** r of a matrix A is equal to the number of its pivot columns.

Here and throughout, the pivot columns of a matrix A refer to those columns where the pivots lie in the (reduced) row echelon form of A .

If the matrix A is $m \times n$, then the rank satisfies the inequality:

$$\boxed{r \leq \min(m, n)}$$

since the number of pivots cannot surpass either the number of rows or of columns. We say that:

- A has full row rank if $r = m$. In this case, the column space of A is as big as possible, i.e. $C(A) = \mathbb{R}^m$, and the system $A\mathbf{v} = \mathbf{b}$ has at least one solution \mathbf{v} for any \mathbf{b} .
- A has full column rank if $r = n$. In this case, the nullspace of A is as small as possible, i.e. $N(A) = \{0\}$, and the system $A\mathbf{v} = \mathbf{b}$ has at most one solution \mathbf{v} for any \mathbf{b} .

Lecture 9 (March 9)

Reading: Section 3.4

Viewing: Lecture 9

For vectors $\mathbf{v}_1, \dots, \mathbf{v}_n \in \mathbb{R}^m$, recall that the subspace V spanned by these vectors was explained in Definition 5. Specifically, vectors in the subspace $V \subset \mathbb{R}^m$ are arbitrary linear combinations:

$$\alpha_1 \mathbf{v}_1 + \dots + \alpha_n \mathbf{v}_n \tag{73}$$

for various scalars $\alpha_1, \dots, \alpha_n$. But what if one of the vectors $\mathbf{v}_1, \dots, \mathbf{v}_n$ was itself a linear combination of the other vectors? For example, we could have:

$$\mathbf{v}_i = \beta_1 \mathbf{v}_1 + \dots + \beta_{i-1} \mathbf{v}_{i-1} + \beta_{i+1} \mathbf{v}_{i+1} + \dots + \beta_n \mathbf{v}_n$$

for some $i \in \{1, \dots, n\}$. Then any linear combination of the form (73) can be written as a linear combination of the vectors $\mathbf{v}_1, \dots, \mathbf{v}_{i-1}, \mathbf{v}_{i+1}, \dots, \mathbf{v}_n$ only:

$$\begin{aligned} \alpha_1 \mathbf{v}_1 + \dots + \alpha_{i-1} \mathbf{v}_{i-1} + \alpha_i \underbrace{(\beta_1 \mathbf{v}_1 + \dots + \beta_{i-1} \mathbf{v}_{i-1} + \beta_{i+1} \mathbf{v}_{i+1} + \dots + \beta_n \mathbf{v}_n)}_{\mathbf{v}_i} + \alpha_{i+1} \mathbf{v}_{i+1} + \dots + \alpha_n \mathbf{v}_n &= \\ = (\alpha_1 + \alpha_i \beta_1) \mathbf{v}_1 + \dots + (\alpha_{i-1} + \alpha_i \beta_{i-1}) \mathbf{v}_{i-1} + (\alpha_{i+1} + \alpha_i \beta_{i+1}) \mathbf{v}_{i+1} + \dots + (\alpha_n + \alpha_i \beta_n) \mathbf{v}_n \end{aligned}$$

So we don't need the vector \mathbf{v}_i to span the subspace $V \subset \mathbb{R}^m$: this subspace is in fact spanned by the slightly smaller collection of vectors $\mathbf{v}_1, \dots, \mathbf{v}_{i-1}, \mathbf{v}_{i+1}, \dots, \mathbf{v}_n$.

Definition 8. Vectors $\mathbf{v}_1, \dots, \mathbf{v}_n \in \mathbb{R}^m$ are called **(linearly) independent** if none of them is a linear combination of the others. Equivalently, this means that any linear combination is non-zero:

$$\boxed{\alpha_1 \mathbf{v}_1 + \dots + \alpha_n \mathbf{v}_n \neq \mathbf{0}} \tag{74}$$

except for the zero combination (namely the one where $\alpha_1 = \dots = \alpha_n = 0$).

If the vectors $\mathbf{v}_1, \dots, \mathbf{v}_n \in \mathbb{R}^m$ are not independent, we call them **(linearly) dependent**, which means that there exists a linear combination:

$$\alpha_1 \mathbf{v}_1 + \dots + \alpha_n \mathbf{v}_n = \mathbf{0} \tag{75}$$

where not all the $\alpha_1, \dots, \alpha_n$ are 0. For example, if $\alpha_i \neq 0$, then we can rewrite equation (75) as:

$$\mathbf{v}_i = -\frac{\alpha_1}{\alpha_i} \mathbf{v}_1 - \dots - \frac{\alpha_{i-1}}{\alpha_i} \mathbf{v}_{i-1} - \frac{\alpha_{i+1}}{\alpha_i} \mathbf{v}_{i+1} - \dots - \frac{\alpha_n}{\alpha_i} \mathbf{v}_n$$

We conclude that a collection of vectors is linear dependent precisely when one of them can be written as a linear combination of the others.

For example, the two vectors:

$$\begin{bmatrix} 2 \\ 5 \\ 1 \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} 1 \\ -3 \\ -1 \end{bmatrix} \tag{76}$$

are linearly independent, because otherwise one of them would be a linear combination (i.e. a multiple) of the other, and this is visibly not true. Geometrically, we can see this because the two vectors above are not on the same line passing through the origin. However, the three vectors:

$$\begin{bmatrix} 2 \\ 5 \\ 1 \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} 1 \\ -3 \\ -1 \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} 4 \\ -1 \\ -1 \end{bmatrix} \quad (77)$$

are linearly dependent, because the following linear combination is 0:

$$1 \cdot \begin{bmatrix} 2 \\ 5 \\ 1 \end{bmatrix} + 2 \cdot \begin{bmatrix} 1 \\ -3 \\ -1 \end{bmatrix} - 1 \cdot \begin{bmatrix} 4 \\ -1 \\ -1 \end{bmatrix} = 0 \quad (78)$$

Alternatively, you could think of this as saying that one of the three vectors is a linear combination of the other two vectors, e.g.:

$$\begin{bmatrix} 4 \\ -1 \\ -1 \end{bmatrix} = 1 \cdot \begin{bmatrix} 2 \\ 5 \\ 1 \end{bmatrix} + 2 \cdot \begin{bmatrix} 1 \\ -3 \\ -1 \end{bmatrix} \quad (79)$$

Definition 9. A **basis** of a vector space V is a collection of independent vectors which span V .

For example, both the collection (76) and the collection (77) span the plane:

$$\left\{ (x, y, z) \text{ such that } 2x - 3y + 11z = 0 \right\} \subset \mathbb{R}^3 \quad (80)$$

But among these two collections of vectors, only (76) is a basis, because they are linearly independent. The collection of vectors (77) is not a basis, because each of the three vectors in (77) can be written as a linear combination of the other two (see (79)). So any of these three vectors is redundant in spanning the plane (80), and that is why we don't need all three of them to have a basis.

Remark. A basis of any vector space is not unique. In fact, any collection of two of the three vectors in (77) form a basis of the subspace (80).

Definition 10. The **dimension** of a vector space V is the number of vectors in a basis of V .

Implicit in the Definition above is that all bases of a vector space (and there are infinitely many of them, see the Remark above) have the same number of vectors in them. The abstract notion of dimension introduced above precisely matches our usual, geometric, notion of dimension: a line has dimension 1 (it is spanned by a single vector) and a plane has dimension 2 (it is spanned by two vectors, for example the plane (80) with the basis (76)).

So, practically, how can we tell when a collection of vectors $\mathbf{v}_1, \dots, \mathbf{v}_n$ is linearly independent? And if they are not linearly independent, how can we remove some of these vectors, in such a way as to obtain a basis of the subspace they span? The way to answer both of these questions is to construct the $m \times n$ matrix whose columns are the vectors in question:

$$A = [\mathbf{v}_1 \mid \dots \mid \mathbf{v}_n] \quad (81)$$

For any vector $\mathbf{w} = \begin{bmatrix} \alpha_1 \\ \vdots \\ \alpha_n \end{bmatrix}$, we have seen that:

$$A\mathbf{w} = \alpha_1\mathbf{v}_1 + \cdots + \alpha_n\mathbf{v}_n \quad (82)$$

Therefore, the vectors $\mathbf{v}_1, \dots, \mathbf{v}_n$ are linearly independent if and only if $A\mathbf{w} = 0$ happens only when $\mathbf{w} = 0$, or in other words, if $N(A) = 0$. Conversely, the vectors are dependent precisely when there exists $\mathbf{w} \neq 0$ such that $A\mathbf{w} = 0$. So to find a linear combination (82) which is 0, one only needs to find a non-zero vector $\mathbf{w} \in N(A)$, and this can be done by putting A in reduced row echelon form.

More concretely, given a collection of vectors $\mathbf{v}_1, \dots, \mathbf{v}_n$, how do we find a subset of them which are linearly independent and thus give a basis of the vector subspace they span? Simple: put the given vectors into a matrix A (see (81)), put A in row echelon form, identify the pivot columns, and look at the corresponding columns of A . Those will be a basis. For example, consider the matrix:

$$A = \begin{bmatrix} 2 & 1 & 4 \\ 5 & -3 & -1 \\ 1 & -1 & -1 \end{bmatrix} \xrightarrow{\text{REF}} U = \begin{bmatrix} \boxed{2} & 1 & 4 \\ 0 & \boxed{-5.5} & -11 \\ 0 & 0 & 0 \end{bmatrix}$$

The pivot columns are the first and second ones, so the upshot is that a basis of $C(A)$ consists of the first two columns of A . This is precisely the basis we found in (76).

Fact 7. *The rank of the matrix A (i.e. the number of pivot columns) equals the dimension of $C(A)$.*

Above, we saw how to find a basis of the vector space spanned by a collection of vectors $\mathbf{v}_1, \dots, \mathbf{v}_n$. But there are other ways to present vector spaces, namely by equations. For example consider:

$$V = \left\{ 2x - 2y - 6z - 4t = 5x + y - 3z + 8t = -3x + 2y + 7z + 3t = 0 \right\} \subset \mathbb{R}^4$$

Since V is cut out by three equations in four unknowns, it is tempting to believe that V is one-dimensional. But not so fast, since there may be redundancy in these equations! Let's find a basis of V systematically. The way to do so is to think of V as the nullspace of the coefficient matrix:

$$V = N(A) \quad \text{where} \quad A = \begin{bmatrix} 2 & -2 & -6 & -4 \\ 5 & 1 & -3 & 8 \\ -3 & 2 & 7 & 3 \end{bmatrix}$$

To describe this nullspace, put the matrix A in reduced row echelon form:

$$A = \begin{bmatrix} 2 & -2 & -6 & -4 \\ 5 & 1 & -3 & 8 \\ -3 & 2 & 7 & 3 \end{bmatrix} \xrightarrow{\text{RREF}} R = \begin{bmatrix} \boxed{1} & 0 & -1 & 1 \\ 0 & \boxed{1} & 2 & 3 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

Since $N(A) = N(R)$, the nullspace of the matrix R in reduced row echelon form can be described in terms of pivot variables (x and y) and free variables (z and t):

$$R \begin{bmatrix} x \\ y \\ z \\ t \end{bmatrix} = 0 \quad \text{i.e.} \quad \begin{cases} x - z + t = 0 \\ y + 2z + 3t = 0 \end{cases}$$

Hence:

$$V = \left\{ \begin{bmatrix} z - t \\ -2z - 3t \\ z \\ t \end{bmatrix} \text{ for any } z, t \right\}$$

To find basis vectors of V , just set one of the free variables (in this case, either z or t) equal to 1 and the other free variable equal to 0. We conclude that a basis of V is given by:

$$\begin{bmatrix} 1 \\ -2 \\ 1 \\ 0 \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} -1 \\ -3 \\ 0 \\ 1 \end{bmatrix}$$

Therefore, V is two-dimensional, hence a plane.

Let us summarize the two ways of thinking about a vector subspace $V \subset \mathbb{R}^m$: either as the column space of a matrix A (when you are given a set of vectors which spans V), or as the nullspace of a matrix A (when you are given equations satisfied by V). In either case, analyzing V (i.e. finding a basis and computing its dimension) is systematically done by putting the matrix A in (reduced) row echelon form.

Remark. *If A is a square $n \times n$ matrix, the rank always satisfies $r \leq n$. We can only have equality, i.e. $r = n$, if the columns of A are linearly independent. This is equivalent to the only solution of $A\mathbf{v} = \mathbf{0}$ being the zero vector $\mathbf{v} = \mathbf{0}$, so the nullspace of A just consists of the zero vector. Moreover:*

$$\boxed{r = n} \quad \text{if and only if} \quad \boxed{A \text{ is invertible}} \quad (83)$$

Indeed, everything we said in this Remark can be translated as saying that (for A a square matrix) the system $A\mathbf{v} = \mathbf{b}$ has a single solution \mathbf{v} for any vector \mathbf{b} . This solution is precisely $\mathbf{v} = A^{-1}\mathbf{b}$.

Lecture 10 (March 10)

Reading: Section 3.5

Viewing: Lecture 10

To a $m \times n$ matrix A , we associated the following subspaces:

the column space $C(A) \subset \mathbb{R}^m$

the nullspace $N(A) \subset \mathbb{R}^n$

We will now consider two more subspaces corresponding to the matrix A :

the row space $C(A^T) \subset \mathbb{R}^n$

the left nullspace $N(A^T) \subset \mathbb{R}^m$

Collectively, the four subspaces above are called the **four fundamental subspaces** of A . The row space is just the subspace spanned by the rows of the matrix A (which are the same things as the columns of the transpose matrix A^T). Meanwhile, the left nullspace consists of $m \times 1$ vectors such that:

$$A^T \begin{bmatrix} v_1 \\ \vdots \\ v_m \end{bmatrix} = 0$$

If you transpose the equation above, then you obtain (using formula (43)):

$$[v_1 \ \dots \ v_m] A = 0$$

hence the left nullspace of A consists of $1 \times m$ vectors (in other words, row vectors) \mathbf{v} such that:

$$\boxed{\mathbf{v}A = 0} \tag{84}$$

This equation explains why we call it the “left” nullspace.

Fact 8. *If r is the rank of the matrix A (number of pivots), then the four subspaces have dimensions:*

$$\boxed{\dim C(A) = r}$$

$$\boxed{\dim C(A^T) = r}$$

$$\boxed{\dim N(A) = n - r}$$

$$\boxed{\dim N(A^T) = m - r}$$

Let’s try and justify the fact above. The fact that the column space has dimension equal to r is just a consequence of the fact that the pivot columns form a basis of $C(A)$. The fact that the nullspace has dimension $n - r$ holds because a choice of basis vectors of $N(A)$ is given by setting all the free variables equal to 0 except for one of them; the number of such basis vectors is the same as the number of free variables, namely n minus the number of pivot variables, i.e. $n - r$.

Let’s justify the fact that the row space of A also has dimension r . To see this, let us bring the matrix in reduced row echelon form R (row operations do not change the row space). Of the m rows of R , exactly r of them will be non-zero (the rows which have pivots) and $m - r$ of them will be zero (the rows which do not have pivots). Therefore, the dimension of the row space is exactly r , since pivot rows are linearly independent. This implies that $\dim C(A^T) = r$, which also implies that

$\dim N(A^T) = m - r$, just by repeating the argument in the previous paragraph for A replaced by A^T .

The reason why we introduced the row space and the left nullspace is that they provide complements for the nullspace and column space, respectively. To understand what this means, suppose you had two vector subspaces:

$$V, W \subset \mathbb{R}^n$$

of dimensions k and $n - k$, respectively. If V and W are sufficiently general, then their intersection only consists of the zero vector, in which case we call V and W **complementary** subspaces. This notion only applies when the dimensions of V and W add up to n , the dimension of ambient space.

Example 3. *Two different lines passing through the origin are complementary subspaces of \mathbb{R}^2 , and any plane and line which only intersect at the origin are complementary subspaces of \mathbb{R}^3 .*

However, two planes passing through the origin can never be complementary subspaces in \mathbb{R}^3 , because their intersection always contains a line (moreover, the dimensions add up to $2 + 2 = 4 \neq 3$).

Here is a fact we will understand more closely in upcoming weeks:

The row space and nullspace of A are complementary subspaces of \mathbb{R}^n

The column space and left nullspace of A are complementary subspaces of \mathbb{R}^m

for any $m \times n$ matrix A . But practically, how do we compute the four fundamental subspaces of a matrix? For example, how do we find a basis of each of them? For the column space and nullspace of A , we saw that the key role was played by its reduced row echelon form R :

- the column space $C(A)$ is spanned by the pivot columns of A (i.e. those columns where the reduced row echelon form R has its pivots)

- the nullspace $N(A)$ is spanned by the vectors $\mathbf{v} = \begin{bmatrix} x_1 \\ \vdots \\ x_n \end{bmatrix}$ where all the free variables except for one take the value 0, and the pivot variables take values prescribed by $R\mathbf{v} = 0$

So how do we compute the spaces $C(A^T)$ and $N(A^T)$? The direct way is to just put A^T in reduced row echelon form, and apply the previous two bullets. But quite often, the problem naturally deals with the reduced row echelon form of the matrix A itself:

$$A \overset{\text{RREF}}{\rightsquigarrow} R \tag{85}$$

By analogy with (62), we have:

$$\boxed{C(A^T) = C(R^T)} \tag{86}$$

which happens because row operations do not change the row space of a matrix. Therefore:

a basis for $C(A^T)$ is given by the pivot rows of the matrix R

What about the left nullspace $N(A^T)$ in terms of the reduced row echelon form (85)? As we have seen a few lectures ago, Gauss-Jordan elimination involves multiplying the matrix A on the left:

$$KA = R$$

by an invertible matrix K . By (84), a row vector \mathbf{v} lies in $N(A^T)$ precisely means:

$$0 = \mathbf{v}A = \underbrace{\mathbf{v}K^{-1}}_{\text{call this } \mathbf{w}} R \quad (87)$$

where \mathbf{v} is an $1 \times m$ vector. The reduced row echelon matrix R will have the form, say:

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

and $0 = \mathbf{w}R$ happens if and only if:

$$\mathbf{w} = [0 \ 0 \ 0 \ * \ *]$$

(namely the only non-zero entries of \mathbf{w} are those corresponding to the zero rows of R). So the vectors \mathbf{v} that can satisfy (87) are those of the form:

$$\mathbf{v} = \mathbf{w}K = [0 \ 0 \ 0 \ * \ *] K$$

We conclude that:

a basis for $N(A^T)$ is given by the rows of the square matrix K which correspond to the zero rows of the reduced row echelon form matrix R

Lecture 11 (March 12)

Reading: Section 4.1

Viewing: Lecture 14

Today we will start discussing orthogonality (i.e. perpendicularity). First of all, we know that two vectors $\mathbf{v}, \mathbf{w} \in \mathbb{R}^n$ are orthogonal precisely when their dot product is 0:

$$\mathbf{v} \perp \mathbf{w} \quad \Leftrightarrow \quad \mathbf{v} \cdot \mathbf{w} = \mathbf{v}^T \mathbf{w} = 0 \quad (88)$$

If two vectors \mathbf{v} and \mathbf{w} are orthogonal, then we have the Pythagorean theorem:

$$\|\mathbf{v} + \mathbf{w}\|^2 = \|\mathbf{v}\|^2 + \|\mathbf{w}\|^2$$

where we recall formula (7) for the length $\|\mathbf{v}\|$ of a vector (proof of the Pythagorean theorem: $\|\mathbf{v} + \mathbf{w}\|^2 = (\mathbf{v} + \mathbf{w}) \cdot (\mathbf{v} + \mathbf{w}) = \mathbf{v} \cdot \mathbf{v} + \mathbf{v} \cdot \mathbf{w} + \mathbf{w} \cdot \mathbf{v} + \mathbf{w} \cdot \mathbf{w} = \|\mathbf{v}\|^2 + \|\mathbf{w}\|^2$).

But now that we have the necessary language to talk about more than just individual vectors in \mathbb{R}^n , we can generalize the concept of orthogonality to entire subspaces:

Definition 11. Two subspaces $V, W \subset \mathbb{R}^n$ are called **orthogonal**, which will be denoted by $V \perp W$, if we have $\mathbf{v} \perp \mathbf{w}$ for any two vectors $\mathbf{v} \in V$ and $\mathbf{w} \in W$.

Note that two subspaces can be orthogonal only if $\dim V + \dim W \leq n$. Otherwise, the two subspaces would intersect in a non-zero vector $\mathbf{v} \in V \cap W$ (think about two planes passing through the origin in \mathbb{R}^3 : they always contain a whole line) and it would be impossible for $\mathbf{v} \perp \mathbf{v}$.

Let us now consider an $m \times n$ matrix A , with the four fundamental subspaces considered in the previous Lecture. We have the following facts:

$$\boxed{C(A) \perp N(A^T)} \quad (89)$$

$$\boxed{N(A) \perp C(A^T)} \quad (90)$$

Let's prove these formulas: for (89), we need to take an arbitrary vector $\mathbf{a} \in C(A)$ and an arbitrary vector $\mathbf{v} \in N(A^T)$ and show that they are perpendicular. By the very definition of the column space and the left nullspace, we have:

$$\mathbf{a} = A\mathbf{w} \text{ for some vector } \mathbf{w} \quad \text{and} \quad A^T \mathbf{v} = 0 \Leftrightarrow \mathbf{v}^T A = 0$$

Then we have:

$$\mathbf{v} \cdot \mathbf{a} = \mathbf{v}^T \mathbf{a} = \mathbf{v}^T A\mathbf{w} = 0\mathbf{w} = 0 \quad \Rightarrow \quad \boxed{\mathbf{a} \perp \mathbf{v}}$$

Similarly, let's prove (90): we need to take an arbitrary $\mathbf{w} \in N(A)$ and an arbitrary $\mathbf{b} \in C(A^T)$ and show that they are perpendicular. By the very definition of nullspace and row space, we have:

$$A\mathbf{w} = 0 \quad \text{and} \quad \mathbf{b} = A^T \mathbf{v} \text{ for some vector } \mathbf{v}$$

Then we have:

$$\mathbf{b}^T = \mathbf{v}^T A \quad \Rightarrow \quad \mathbf{b}^T \mathbf{w} = \mathbf{v}^T A\mathbf{w} = \mathbf{v}^T 0 = 0 \quad \Rightarrow \quad \boxed{\mathbf{b} \perp \mathbf{w}}$$

(alternatively, you can think of statement (90) as simply restating (89) for A^T instead of A). Moreover, we saw in Fact 8 that the spaces in (89) are of complementary dimension, and thus they have the maximal total dimension that orthogonal subspaces can have. Ditto for (90). This is a special case of the following notion:

Definition 12. Given a subspace $V \subset \mathbb{R}^n$, its **orthogonal complement** is the subspace V^\perp consisting of all vectors in \mathbb{R}^n perpendicular to V . In other words:

$$V^\perp = \left\{ \mathbf{w} \in \mathbb{R}^n \text{ such that } \mathbf{v} \perp \mathbf{w} \text{ for all } \mathbf{v} \in V \right\} \quad (91)$$

Note that (91) is just a set. To show that it is a subspace, you would need to prove that:

$$\mathbf{v} \perp \mathbf{w}_1 \text{ and } \mathbf{v} \perp \mathbf{w}_2 \text{ for all } \mathbf{v} \in V \quad \Rightarrow \quad \mathbf{v} \perp (\alpha_1 \mathbf{w}_1 + \alpha_2 \mathbf{w}_2) \text{ for all } \mathbf{v} \in V$$

for any scalars α_1, α_2 . The implication above can be easily proved by resorting to the dot product interpretation of perpendicularity. So combining Fact 8 with (89), (90) precisely says that:

the left nullspace is the orthogonal complement of the column space
the row space is the orthogonal complement of the nullspace

This is sometimes called the Fundamental Theorem of Linear Algebra. Let's see an example:

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

(I chose a matrix which is already in reduced row echelon form to save us the work of Gauss-Jordan elimination, which is a routine computation for us by now). Then the row space of A is:

$$C(A^T) = \text{spanned by } \begin{bmatrix} 1 \\ 0 \\ 3 \\ -1 \end{bmatrix} \text{ and } \begin{bmatrix} 0 \\ 1 \\ -2 \\ 1 \end{bmatrix} \quad (92)$$

(we write the rows vertically to keep with our convention that vectors should be written as columns). To compute a basis for the nullspace of A , you must first note that the pivot columns of the matrix are 1 and 2, while the free columns are 3 and 4. Basis vectors for the nullspace are given by setting all the free variables equal to 0, except for one which is set equal to 1:

$$N(A) = \text{spanned by } \begin{bmatrix} a \\ b \\ 1 \\ 0 \end{bmatrix} \text{ and } \begin{bmatrix} c \\ d \\ 0 \\ 1 \end{bmatrix}$$

To solve for a, b, c, d , you must use the property that the vectors above are annihilated by A :

$$A \begin{bmatrix} a \\ b \\ 1 \\ 0 \end{bmatrix} = A \begin{bmatrix} c \\ d \\ 0 \\ 1 \end{bmatrix} = 0 \quad \Rightarrow \quad \begin{cases} a + 3 = 0 \\ b - 2 = 0 \\ c - 1 = 0 \\ d + 1 = 0 \end{cases}$$

so we conclude that:

$$N(A) = \text{spanned by } \begin{bmatrix} -3 \\ 2 \\ 1 \\ 0 \end{bmatrix} \text{ and } \begin{bmatrix} 1 \\ -1 \\ 0 \\ 1 \end{bmatrix} \quad (93)$$

It is easy to see that the two basis vectors in (92) are orthogonal to each of the two basis vectors in (93), thus establishing the fact that the row space and nullspace are orthogonal subspaces.

The fundamental theorem of linear algebra says that the four fundamental subspaces of a matrix are pairwise orthogonal complements. Here's why we like this setup: if you have two complementary subspaces $V, W \subset \mathbb{R}^n$ (as defined in Lecture 10) then any vector $\mathbf{a} \in \mathbb{R}^n$ can be written as:

$$\mathbf{a} = \mathbf{v} + \mathbf{w} \quad \text{where } \mathbf{v} \in V \text{ and } \mathbf{w} \in W \quad (94)$$

Moreover, \mathbf{v} and \mathbf{w} with this property are unique, and they are called the **components** of the vector \mathbf{a} in the complementary subspaces V and W , respectively. So for example, letting V and W be the nullspace and row space of an $m \times n$ matrix, then we conclude that any vector in \mathbb{R}^n can be written uniquely as a linear combination of the rows of A plus a vector annihilated by A .

In the special case when V, W are complementary orthogonal subspaces of \mathbb{R}^n , i.e.:

$$V = W^\perp \quad \text{or, equivalently} \quad W = V^\perp$$

the decomposition (94) takes on a more geometric meaning. In this case, \mathbf{v} and \mathbf{w} are none other than the projections of the general vector \mathbf{a} onto the subspaces V and W :

$$\mathbf{v} = \text{proj}_V \mathbf{a} \quad \text{and} \quad \mathbf{w} = \text{proj}_W \mathbf{a}$$

We conclude that any vector is the sum of its projections onto complementary orthogonal subspaces.

Lecture 12 (March 15)

Reading: Section 4.2

Viewing: Lecture 15

We will now give a rigorous definition of projections.

Definition 13. Given a vector $\mathbf{b} \in \mathbb{R}^n$ and a linear subspace $V \subset \mathbb{R}^n$, the **(orthogonal) projection** of \mathbf{b} onto V is the unique vector:

$$\mathbf{p} = \text{proj}_V \mathbf{b} \in V$$

such that $(\mathbf{b} - \mathbf{p}) \perp V$. Put differently, \mathbf{p} is the closest vector in V to \mathbf{b} .

Intuitively, suppose you want to measure a quantity, which general theory tells you must be a vector \mathbf{p} in a subspace $V \subset \mathbb{R}^n$. However, your measurement will probably have an error, and you might get a value \mathbf{b} outside of V . Then your best guess for the actual value $\mathbf{p} \in V$ is to set it equal to the orthogonal projection of \mathbf{b} on V , since this is the closest vector in V to the measured value \mathbf{b} . That is also why the vector $\mathbf{e} = \mathbf{b} - \mathbf{p}$ is sometimes called the **error**.

Let us work out formulas for the projection of a vector $\mathbf{b} = \begin{bmatrix} b_1 \\ \vdots \\ b_n \end{bmatrix}$. The easiest case is when V is one dimensional, i.e. the line spanned by a vector $\mathbf{a} = \begin{bmatrix} a_1 \\ \vdots \\ a_n \end{bmatrix}$. Then the orthogonal projection \mathbf{p} must be a multiple of this vector, so we are looking for:

$$\mathbf{p} = \lambda \mathbf{a} \quad \text{such that} \quad (\mathbf{b} - \lambda \mathbf{a}) \perp \mathbf{a}$$

Using the dot product interpretation of perpendicularity, we need:

$$0 = (\mathbf{b} - \lambda \mathbf{a}) \cdot \mathbf{a} \quad \Rightarrow \quad \lambda = \frac{\mathbf{a} \cdot \mathbf{b}}{\mathbf{a} \cdot \mathbf{a}} = \frac{\mathbf{a}^T \mathbf{b}}{\mathbf{a}^T \mathbf{a}}$$

So the formula for the projection of \mathbf{b} onto the line spanned by a single vector \mathbf{a} is:

$$\text{proj}_{\mathbf{a}} \mathbf{b} = \mathbf{a} \underbrace{\frac{\mathbf{a}^T \mathbf{b}}{\mathbf{a}^T \mathbf{a}}}_{\text{this is a scalar}} \quad (95)$$

Here's a sanity check: assume $\mathbf{a} = \mathbf{e}_i$ is the unit vector on the i -th axis. Then formula (95) reads:

$$\text{proj}_{\mathbf{e}_i} \mathbf{b} = \mathbf{e}_i b_i \quad (96)$$

so the projection only remembers the i -th component of the vector \mathbf{b} . It is easy to see that the operation (96) can be presented as a matrix multiplying the vector \mathbf{b} , specifically:

$$P_i \mathbf{b} = \mathbf{e}_i b_i \quad \text{where} \quad P_i = \begin{bmatrix} \dots & \dots & 0 & \dots & \dots \\ \dots & \dots & 0 & \dots & \dots \\ 0 & 0 & 1 & 0 & 0 \\ \dots & \dots & 0 & \dots & \dots \\ \dots & \dots & 0 & \dots & \dots \end{bmatrix} \quad (97)$$

with the unique 1 being at the intersection of row i and column i , and 0's everywhere else. The projection onto an arbitrary vector \mathbf{a} can be described by a similar matrix, specifically:

$$P_{\mathbf{a}}\mathbf{b} = \mathbf{a}\frac{\mathbf{a}^T\mathbf{b}}{\mathbf{a}^T\mathbf{a}} \quad \text{where} \quad P_{\mathbf{a}} = \frac{\mathbf{a}\mathbf{a}^T}{\mathbf{a}^T\mathbf{a}} \quad (98)$$

Indeed, since \mathbf{a} is an $n \times 1$ vector, the numerator of $P_{\mathbf{a}}$ is an $n \times n$ matrix, and the denominator is a 1×1 matrix, i.e. a scalar. This all fits into the following principle:

Fact 9. *For any subspace $V \subset \mathbb{R}^n$, there exists an $n \times n$ matrix P_V such that:*

$$\boxed{\text{proj}_V\mathbf{b} = P_V\mathbf{b}} \quad (99)$$

for any vector $\mathbf{b} \in \mathbb{R}^n$.

So all orthogonal projections are given by left multiplication with a suitable matrix (unoriginally called a “projection matrix”). The task for us now is to compute the matrix corresponding to a given projection, and we will do it in the special case when:

$$V = C(A)$$

for an $n \times m$ matrix A . In fact, this does not represent any restriction on the subspace V , since any subspace can be written as the column space of a suitably chosen matrix: just take basis vectors of the subspace and put them next to each other in a matrix. So given $\mathbf{b} \in \mathbb{R}^n$, we are looking for:

$$\mathbf{p} \in C(A) \quad \text{such that} \quad (\mathbf{b} - \mathbf{p}) \perp C(A)$$

But $\mathbf{p} \in C(A)$ precisely means:

$$\mathbf{p} = A\mathbf{v} \quad (100)$$

for some $m \times 1$ vector \mathbf{v} . Moreover, the condition $(\mathbf{b} - \mathbf{p}) \perp C(A)$ is equivalent to:

$$A^T(\mathbf{b} - \mathbf{p}) = 0 \quad (101)$$

because we know from the previous Lecture that a vector being orthogonal to the column space of A is equivalent to it lying in the left nullspace of A . Combining (100) with (101) gives us:

$$A^T(\mathbf{b} - A\mathbf{v}) = 0 \quad \Rightarrow \quad A^T A\mathbf{v} = A^T\mathbf{b}$$

Therefore, we may solve for $\mathbf{v} = (A^T A)^{-1}A^T\mathbf{b}$, and by combining this with (100) we obtain the following formula for the projection of the vector \mathbf{b} onto the subspace $C(A)$:

$$\boxed{\text{proj}_{C(A)}\mathbf{b} = A(A^T A)^{-1}A^T\mathbf{b}} \quad (102)$$

Therefore, we conclude that the $n \times n$ projection matrix that Fact 9 predicted gives rise to the projection onto the subspace $C(A)$ is:

$$\boxed{P_{C(A)} = A(A^T A)^{-1}A^T} \quad (103)$$

If the case $m = 1$, i.e. A is a single $n \times 1$ vector, formula (103) is equivalent to (98). The $m \times m$ matrix $S = A^T A$ which appears in (103) is symmetric, but we need it to be invertible in order for the formulas presented above to hold. Luckily for us, we have the following fact:

$$\text{if the columns of } A \text{ are independent, then } S = A^T A \text{ is invertible} \quad (104)$$

When working with formulas such as (102) or (103), note that the equality:

$$(A^T A)^{-1} \neq A^{-1}(A^T)^{-1}$$

does not hold in general, because when A is a rectangular matrix, its inverse is not well-defined.

Let's do an example. Suppose we are in $n = 3$ dimensional space, and we want to compute the projection onto the plane V spanned by the vectors:

$$\begin{bmatrix} 2 \\ 0 \\ -1 \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} -1 \\ 1 \\ 0 \end{bmatrix} \quad (105)$$

First of all, let's make sure that these vectors are linearly independent (if one were a linear combination of the other, we could just throw it out). Then, we put the vectors together in a matrix:

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

As we saw in (103), an important role is played by the symmetric matrix:

$$S = A^T A = \begin{bmatrix} 2 & 0 & -1 \\ -1 & 1 & 0 \end{bmatrix} \begin{bmatrix} 2 & -1 \\ 0 & 1 \\ -1 & 0 \end{bmatrix} = \begin{bmatrix} 5 & -2 \\ -2 & 2 \end{bmatrix}$$

whose inverse (this matrix is invertible because the vectors (105) are independent) is given by:

$$S^{-1} = \frac{1}{6} \begin{bmatrix} 2 & 2 \\ 2 & 5 \end{bmatrix}$$

Then formula (103) for the projection matrix onto the subspace spanned by the vectors (105) is:

$$P_V = \begin{bmatrix} 2 & -1 \\ 0 & 1 \\ -1 & 0 \end{bmatrix} \frac{1}{6} \begin{bmatrix} 2 & 2 \\ 2 & 5 \end{bmatrix} \begin{bmatrix} 2 & 0 & -1 \\ -1 & 1 & 0 \end{bmatrix} = \frac{1}{6} \begin{bmatrix} 5 & -1 & -2 \\ -1 & 5 & -2 \\ -2 & -2 & 2 \end{bmatrix} \quad (106)$$

We conclude that the projection of any vector $\begin{bmatrix} x \\ y \\ z \end{bmatrix}$ onto the given 2-dimensional plane is:

$$P_V \begin{bmatrix} x \\ y \\ z \end{bmatrix} = \frac{1}{6} \begin{bmatrix} 5x - y - 2z \\ -x + 5y - 2z \\ -2x - 2y + 2z \end{bmatrix}$$

Remark. *Projection matrices are special in that they satisfy the equation:*

$$P_V P_V = P_V$$

for any subspace V . This corresponds to the geometric fact that the projection of a projection is the projection itself. Check that this equation holds for the specific example (106).

Lecture 13 (March 17)

Reading: Section 4.3

Viewing: Lecture 16

A big computational application of projections is **least squares** approximation, which we will now delve into. Mathematically, the problem involves a system of equations for given A and \mathbf{b} :

$$A\mathbf{v} = \mathbf{b} \tag{107}$$

which does not have any solutions, because the number of equations is greater than the number of variables (mathematically, think of A as an $m \times n$ matrix with $m > n$). Therefore, the task is to find a value of \mathbf{v} which makes $A\mathbf{v}$ as close as possible to \mathbf{b} . In other words, the **error** vector:

$$\mathbf{e} = \mathbf{b} - A\mathbf{v}$$

should be made as “small” as possible. Since we are dealing with vectors, “small” means that the length of the vector should be as small as possible. Therefore, the problem is to find \mathbf{v} such that:

$$\|\mathbf{e}\| = \|\mathbf{b} - A\mathbf{v}\| \tag{108}$$

is as small as possible, for given A and \mathbf{b} . But what does this mean geometrically? As \mathbf{v} ranges over all possible vectors, $A\mathbf{v}$ ranges over the column space $C(A)$. If \mathbf{b} were already in $C(A)$, then the equation (107) could be satisfied with no error. But in general, $\mathbf{b} \notin C(A)$, in which case:

$$\|\mathbf{b} - A\mathbf{v}\| \geq \left(\text{distance from } \mathbf{b} \text{ to the subspace } C(A) \right) = \|\mathbf{b} - \mathbf{p}\|$$

where \mathbf{p} is the orthogonal projection of \mathbf{b} onto the subspace $C(A)$. Therefore, finding \mathbf{v} which minimizes the quantity (108) actually involves two steps:

- find the orthogonal projection \mathbf{p} of \mathbf{b} onto the subspace $C(A)$
- solve the equation $A\mathbf{v} = \mathbf{p}$ (which can be done, since $\mathbf{p} \in C(A)$)

We already know the answer to the first question from (102):

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

The answer to the second question is also manifest, because a solution to the boxed equation is:

$$\mathbf{v} = (A^T A)^{-1} A^T \mathbf{b} \tag{109}$$

Recall the fact (104): in order for the matrix $A^T A$ to be invertible and for the formula above to make sense, we need the matrix A to have independent columns. This can always be arranged by removing some dependent columns from the matrix A : for example, if the last column of A is a linear combination of the other ones, then we can simply remove the last column of A and the last entry of \mathbf{v} , and the task of approximating equation (107) as well as possible remains the same.

Remark. You could also have arrived at formula (109) from calculus. Specifically, it's obvious that $\|\mathbf{b} - A\mathbf{v}\|$ is minimum when $\|\mathbf{b} - A\mathbf{v}\|^2$ is minimum. The usual formulas give us:

$$\|\mathbf{b} - A\mathbf{v}\|^2 = (\mathbf{b} - A\mathbf{v})^T(\mathbf{b} - A\mathbf{v}) = \mathbf{b}^T\mathbf{b} - 2(A\mathbf{v})^T\mathbf{b} + (A\mathbf{v})^T(A\mathbf{v}) = \mathbf{b}^T\mathbf{b} - 2\mathbf{v}^T A^T\mathbf{b} + \mathbf{v}^T A^T A\mathbf{v}$$

Now, the latter quantity is minimum when its derivative with respect to \mathbf{v} vanishes (there is a mathematical notion of taking derivative with respect to a vector of variables, and it behaves just like derivative with respect to a single variable), which explicitly states that:

$$0 = \frac{\partial(\mathbf{b}^T\mathbf{b} - 2\mathbf{v}^T A^T\mathbf{b} + \mathbf{v}^T A^T A\mathbf{v})}{\partial\mathbf{v}} = -2A^T\mathbf{b} + 2A^T A\mathbf{v}$$

Solving the equation above for \mathbf{v} gives precisely (109).

Let's do a concrete example. Find numbers x, y, z such that the vector:

$$x \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix} + y \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix} + z \begin{bmatrix} 1 \\ 2 \\ 1 \end{bmatrix} \quad \text{is as close to} \quad \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix} \quad \text{as possible} \quad (110)$$

A straightforward application of the analysis above would have us set $A = \begin{bmatrix} 1 & 0 & 1 \\ 1 & 1 & 2 \\ 0 & 1 & 1 \end{bmatrix}$ and apply formula (109). But the columns of A are not independent: in fact, the third column is the sum of the first two. This is reflected in the fact that equation (110) reads:

$$x \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix} + y \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix} + z \left(\begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix} + \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix} \right) = (x+z) \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix} + (y+z) \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix}$$

Since x, y, z can be chosen arbitrarily, the problem is equivalent to finding a and b such that:

$$\text{the vector} \quad a \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix} + b \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix} \quad \text{is as close to} \quad \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix} \quad \text{as possible}$$

In matrix form, this is translated as:

$$\underbrace{\begin{bmatrix} 1 & 0 \\ 1 & 1 \\ 0 & 1 \end{bmatrix}}_A \underbrace{\begin{bmatrix} a \\ b \end{bmatrix}}_v - \underbrace{\begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix}}_b \quad \text{is as small as possible} \quad (111)$$

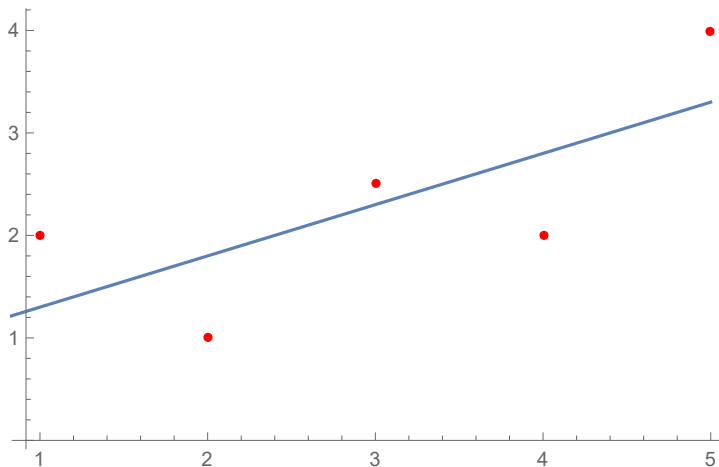
Now the columns are independent, so we can apply formula (109). We have:

$$A^T A = \begin{bmatrix} 1 & 1 & 0 \\ 0 & 1 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 1 & 1 \\ 0 & 1 \end{bmatrix} = \begin{bmatrix} 2 & 1 \\ 1 & 2 \end{bmatrix} \quad \Rightarrow \quad (A^T A)^{-1} = \frac{1}{3} \begin{bmatrix} 2 & -1 \\ -1 & 2 \end{bmatrix}$$

and so (109) implies that the difference (111) is minimized for:

$$\begin{bmatrix} a \\ b \end{bmatrix} = \frac{1}{3} \begin{bmatrix} 2 & -1 \\ -1 & 2 \end{bmatrix} \begin{bmatrix} 1 & 1 & 0 \\ 0 & 1 & 1 \end{bmatrix} \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix} = \frac{1}{3} \begin{bmatrix} 2 & 1 & -1 \\ -1 & 1 & 2 \end{bmatrix} \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix} = \begin{bmatrix} \frac{1}{3} \\ \frac{2}{3} \end{bmatrix}$$

We now turn to what is perhaps the standard application of least squares approximation, which will explain its name. Suppose you have a graph which seeks to represent a time-dependent quantity $y(t)$. You measure this quantity at times t_1, \dots, t_m and get values y_1, \dots, y_m . Mark the data points $(t_1, y_1), \dots, (t_m, y_m)$ on a graph, as follows (say $m = 5$ data points in the picture below):



The goal is to “fit” a straight line L between these data points as closely as possible. Mathematically, the word “closely” means that the following condition must be satisfied:

the sum of squares of the vertical distances from the data points to L is minimum

which explains the terminology “least squares”. So how do we do find this best fit line? To choose a line is the same thing as to choose two numbers a and b , so that the equation of the line is $y(t) = a + bt$. Then the sum of squares that we need to minimize is the quantity:

$$(y_1 - a - bt_1)^2 + \dots + (y_m - a - bt_m)^2 \quad (112)$$

where t_1, \dots, t_m and y_1, \dots, y_m are given to you by your experiment. How to convert this into a linear algebra problem? If the line were to pass precisely through our data points, this would require the following system of equations to be satisfied (where a and b are the unknowns):

$$\begin{cases} a + bt_1 = y_1 \\ \dots \\ a + bt_m = y_m \end{cases} \quad \text{or, equivalently} \quad \begin{bmatrix} 1 & t_1 \\ \vdots & \vdots \\ 1 & t_m \end{bmatrix} \begin{bmatrix} a \\ b \end{bmatrix} = \begin{bmatrix} y_1 \\ \vdots \\ y_m \end{bmatrix}$$

In practice, we will not be able to solve this system on the nose, but the condition that the quantity (112) is minimized is exactly the condition that the square of (108) is minimized, for:

$$A = \begin{bmatrix} 1 & t_1 \\ \vdots & \vdots \\ 1 & t_m \end{bmatrix} \quad \text{and} \quad \mathbf{b} = \begin{bmatrix} y_1 \\ \vdots \\ y_m \end{bmatrix}$$

We have:

$$A^T A = \begin{bmatrix} 1 & \dots & 1 \\ t_1 & \dots & t_m \end{bmatrix} \begin{bmatrix} 1 & t_1 \\ \vdots & \vdots \\ 1 & t_m \end{bmatrix} = \begin{bmatrix} m & t_1 + \dots + t_m \\ t_1 + \dots + t_m & t_1^2 + \dots + t_m^2 \end{bmatrix} \Rightarrow$$

$$\Rightarrow (A^T A)^{-1} = \frac{1}{m(t_1^2 + \dots + t_m^2) - (t_1 + \dots + t_m)^2} \begin{bmatrix} t_1^2 + \dots + t_m^2 & -t_1 - \dots - t_m \\ -t_1 - \dots - t_m & m \end{bmatrix}$$

So formula (109) tells us that the choice of a and b which gives the best straight line fit is:

$$\begin{bmatrix} a \\ b \end{bmatrix} = \frac{1}{m(t_1^2 + \dots + t_m^2) - (t_1 + \dots + t_m)^2} \begin{bmatrix} t_1^2 + \dots + t_m^2 & -t_1 - \dots - t_m \\ -t_1 - \dots - t_m & m \end{bmatrix} \begin{bmatrix} 1 & \dots & 1 \\ t_1 & \dots & t_m \end{bmatrix} \begin{bmatrix} y_1 \\ \vdots \\ y_m \end{bmatrix}$$

and therefore:

$$\begin{bmatrix} a \\ b \end{bmatrix} = \frac{1}{m(t_1^2 + \dots + t_m^2) - (t_1 + \dots + t_m)^2} \begin{bmatrix} \sum_{i=1}^n t_i(t_i - t_1) & \dots & \sum_{i=1}^n t_i(t_i - t_m) \\ \sum_{i=1}^n (t_1 - t_i) & \dots & \sum_{i=1}^n (t_m - t_i) \end{bmatrix} \begin{bmatrix} y_1 \\ \vdots \\ y_m \end{bmatrix}$$

The fact that the fraction in the formula above is non-zero is a consequence of the Cauchy-Schwarz inequality, which states that $m(t_1^2 + \dots + t_m^2) > (t_1 + \dots + t_m)^2$ if t_1, \dots, t_m are all distinct numbers.

Lecture 14 (March 19)

Reading: Section 4.4

Viewing: Lecture 17

Question: when are two planes V and W orthogonal? Answer: when any two vectors $\mathbf{v} \in V$ and $\mathbf{w} \in W$ are orthogonal. Of course, it is hard to check this property for any vectors, but it suffices to check it for all vectors $\mathbf{v}_1, \dots, \mathbf{v}_k$ and $\mathbf{w}_1, \dots, \mathbf{w}_l$ which form bases of V and W , respectively:

$$V \perp W \Leftrightarrow \mathbf{v}_i^T \mathbf{w}_j = 0 \text{ for all } i \in \{1, \dots, k\}, j \in \{1, \dots, l\} \quad (113)$$

Since there are k basis vectors of V and l basis vectors of W , there are $k \cdot l$ conditions to check above. You can replace this with a single condition, if you put the \mathbf{v} 's together in a matrix A and the \mathbf{w} 's together in a matrix B :

$$A = [\mathbf{v}_1 \mid \dots \mid \mathbf{v}_k] \quad B = [\mathbf{w}_1 \mid \dots \mid \mathbf{w}_l]$$

Then $V = C(A)$ and $W = C(B)$, and condition (113) reads:

$$V \perp W \Leftrightarrow A^T B = 0 \quad (114)$$

The following is a related notion, which applies to bases of a single vector space V .

Definition 14. A collection of non-zero vectors $\mathbf{q}_1, \dots, \mathbf{q}_n$ are called **orthogonal** if:

$$\mathbf{q}_i \perp \mathbf{q}_j \text{ or, equivalently } \mathbf{q}_i^T \mathbf{q}_j = 0 \text{ for all } 1 \leq i \neq j \leq n \quad (115)$$

It is called **orthonormal** if the vectors are both orthogonal and have length 1:

$$\|\mathbf{q}_i\| = 1 \text{ or, equivalently } \mathbf{q}_i^T \mathbf{q}_i = 1 \text{ for all } 1 \leq i \leq n \quad (116)$$

Orthogonal vectors are always linearly independent (so a basis of the vector space they span).

The standard basis $\mathbf{e}_1, \dots, \mathbf{e}_n$ of \mathbb{R}^n (where $\mathbf{e}_i = (0, \dots, 0, 1, 0, \dots, 0)$ with the 1 on the i -th spot) is orthonormal. In fact, the notion of orthonormality seeks to generalize this behavior to other bases.

There are n^2 conditions (115)–(116), so there are n^2 properties to check in order to show that a collection of n vectors are orthogonal. As before, this can be packaged into a single matrix equality, if we consider the $m \times n$ matrix (assume $\mathbf{q}_1, \dots, \mathbf{q}_n \in \mathbb{R}^m$):

$$Q = [\mathbf{q}_1 \mid \dots \mid \mathbf{q}_n]$$

Then the orthogonality property (115) precisely states that:

$$Q^T Q = \begin{bmatrix} d_1 & 0 & \dots & 0 \\ 0 & d_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & d_n \end{bmatrix} \quad (117)$$

where $d_i = \|\mathbf{q}_i\|^2$. Therefore, the orthonormality properties (115)–(116) precisely state that:

$$\boxed{Q^T Q = I_n} \quad (118)$$

In the particular case when $m = n$, we have:

Definition 15. A square matrix Q is called **orthogonal** if (118) holds, i.e. if $Q^T = Q^{-1}$.

Very important: Definition 15 (as well as the term “orthogonal matrix”) only applies to square matrices. This is because while we can state property (118) for a non-square matrix, it does not make sense to take the inverse of a non-square matrix.

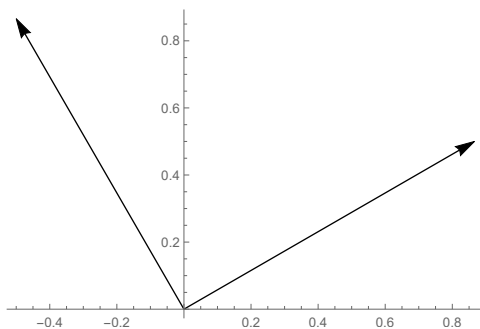
Examples of orthogonal matrices include all permutation matrices:

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

and the rotation matrices:

$$Q = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix} \quad (119)$$

The columns of the matrix Q of (119) are the usual unit vectors in \mathbb{R}^2 , rotated by an angle θ :



and they are orthonormal because $(\sin \theta)^2 + (\cos \theta)^2 = 1$. Both examples above share a feature, which is actually common of all orthogonal matrices: they preserve dot products, lengths and perpendicularity of vectors:

$$\text{if } Q \text{ is an orthogonal matrix, then } \boxed{(Q\mathbf{v})^T(Q\mathbf{w}) = \mathbf{v}^T\mathbf{w}} \quad (120)$$

for all vectors \mathbf{v} and \mathbf{w} of the same size. The reason for this is the chain of equalities:

$$(Q\mathbf{v})^T(Q\mathbf{w}) = \mathbf{v}^T Q^T Q \mathbf{w} = \mathbf{v}^T I \mathbf{w} = \mathbf{v}^T \mathbf{w}$$

Therefore, as a consequence of (120), we have:

$$\text{if } \mathbf{v} \perp \mathbf{w} \text{ then } Q\mathbf{v} \perp Q\mathbf{w} \quad (121)$$

and:

$$\|Q\mathbf{v}\| = \|\mathbf{v}\| \quad (122)$$

for all vectors \mathbf{v} and \mathbf{w} of the same size.

The discussion above was for square matrices, but now let’s consider general ones. More precisely, let’s return to the projection formula (103), which is generally cumbersome because we need to

invert the matrix $A^T A$. In the particular case when $A = Q$ has orthonormal columns (i.e. we have an orthonormal basis of the vector space we are projecting upon), the projection formula is nicer:

$$P_{C(Q)} = QQ^T \tag{123}$$

because as we have seen in (118):

$$Q \text{ has orthonormal columns} \iff Q^T Q = I$$

Similarly with the discussion above, in the case when $A = Q$ has orthonormal columns, the solution to the least squares problem (109) simplifies to:

$$\mathbf{v} = Q^T \mathbf{b} \tag{124}$$

To summarize: if you're projecting onto a subspace with general basis vectors $\mathbf{v}_1, \dots, \mathbf{v}_n$, then you need to apply the more complicated formulas (103), (109) for $A = [\mathbf{v}_1 | \dots | \mathbf{v}_n]$. But if you know that your subspace has a basis consisting of orthonormal vectors $\mathbf{q}_1, \dots, \mathbf{q}_n$, then you may apply the simpler formulas (123), (124) for $Q = [\mathbf{q}_1 | \dots | \mathbf{q}_n]$. It pays to have orthonormal bases of subspaces.

Let's do an example of computing the projection matrix using formula (123) instead of the original formula (103). Suppose you wanted a formula for the projection matrix onto the plane:

$$V = \{x + y + z = 0\} \subset \mathbb{R}^3 \tag{125}$$

The "old" way of doing this was to pick a basis of V (i.e. any two vectors in V which are not multiples of each other), form the 3×2 matrix A whose columns are the given vectors, and apply formula (103). The "new" way is to start not from an arbitrary basis, but from an orthonormal basis of V . We'll learn how to construct such bases systematically in the next lecture, but the first step is to pick two orthogonal (i.e. perpendicular) vectors in V . For example, the vector:

$$\mathbf{v}_1 = \begin{bmatrix} 1 \\ -1 \\ 0 \end{bmatrix}$$

lies in V , so it can be the first vector in our basis. Now we need to pick a second vector \mathbf{v}_2 which is at the same time orthogonal to \mathbf{v}_1 (i.e. the first two coordinates of \mathbf{v}_2 must be equal to each other) and also lie in V (i.e. the coordinates of \mathbf{v}_2 must sum up to 0). A natural choice then is:

$$\mathbf{v}_2 = \begin{bmatrix} 1 \\ 1 \\ -2 \end{bmatrix}$$

To turn an orthogonal basis into an orthonormal basis, we need to rescale the vectors \mathbf{v}_1 and \mathbf{v}_2 so that they have length 1. This is achieved by dividing \mathbf{v}_1 and \mathbf{v}_2 by their lengths, which are:

$$\|\mathbf{v}_1\| = \sqrt{2} \quad \text{and} \quad \|\mathbf{v}_2\| = \sqrt{6}$$

Therefore, we conclude that the vectors:

$$\mathbf{q}_1 = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ -1 \\ 0 \end{bmatrix} \quad \text{and} \quad \mathbf{q}_2 = \frac{1}{\sqrt{6}} \begin{bmatrix} 1 \\ 1 \\ -2 \end{bmatrix} \tag{126}$$

form an orthonormal basis of V . Then we form a matrix whose columns are the orthonormal basis:

$$Q = [\mathbf{q}_1 \mid \mathbf{q}_2] = \begin{bmatrix} \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{6}} \\ \frac{-1}{\sqrt{2}} & \frac{1}{\sqrt{6}} \\ \frac{0}{\sqrt{2}} & \frac{-2}{\sqrt{6}} \end{bmatrix}$$

and formula (123) tells us that the projection matrix we're looking for takes the form:

$$P_V = QQ^T = \begin{bmatrix} \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{6}} \\ \frac{-1}{\sqrt{2}} & \frac{1}{\sqrt{6}} \\ \frac{0}{\sqrt{2}} & \frac{-2}{\sqrt{6}} \end{bmatrix} \begin{bmatrix} \frac{1}{\sqrt{2}} & \frac{-1}{\sqrt{2}} & 0 \\ \frac{1}{\sqrt{6}} & \frac{1}{\sqrt{6}} & \frac{-2}{\sqrt{6}} \end{bmatrix} = \begin{bmatrix} \frac{2}{3} & -\frac{1}{3} & -\frac{1}{3} \\ -\frac{1}{3} & \frac{2}{3} & -\frac{1}{3} \\ -\frac{1}{3} & -\frac{1}{3} & \frac{2}{3} \end{bmatrix}$$

Lecture 15 (March 24)

Reading: Section 4.4

Viewing: Lecture 17

We finished the previous Lecture with an example that basically asked you to “pick an orthonormal basis” of a given vector space. We’ll now see how to explicitly construct such a basis, and even more so, we’ll give an explicit algorithm to construct an orthonormal basis from any given basis. This is called the **Gram-Schmidt process** and here’s how it goes: you start from a general collection of linearly independent vectors $\mathbf{v}_1, \dots, \mathbf{v}_n$ (in whatever vector space you are working with) and the goal is to “make” them orthonormal. You are allowed to modify the individual vectors \mathbf{v}_i by rescaling them and adding linear combinations of the other vectors:

$$\mathbf{v}_i \rightsquigarrow \alpha_1 \mathbf{v}_1 + \dots + \alpha_i \mathbf{v}_i + \dots + \alpha_n \mathbf{v}_n \quad \text{with } \alpha_i \neq 0$$

since these operations do not change the vector space spanned by $\mathbf{v}_1, \dots, \mathbf{v}_n$. Start by renormalizing \mathbf{v}_1 so that it has length 1:

$$\boxed{\text{Step 1 :}} \quad \mathbf{v}_1 \rightsquigarrow \mathbf{q}_1 = \frac{\mathbf{v}_1}{\|\mathbf{v}_1\|}$$

Then modify \mathbf{v}_2 so that it is perpendicular to \mathbf{q}_1 , and then renormalize it to have length 1:

$$\boxed{\text{Step 2 :}} \quad \mathbf{v}_2 \rightsquigarrow \mathbf{w}_2 = \mathbf{v}_2 - \text{proj}_{\mathbf{q}_1} \mathbf{v}_2 \quad \text{and} \quad \mathbf{w}_2 \rightsquigarrow \mathbf{q}_2 = \frac{\mathbf{w}_2}{\|\mathbf{w}_2\|}$$

Then modify \mathbf{v}_3 so that it is perpendicular to \mathbf{q}_1 and \mathbf{q}_2 , and then renormalize it to have length 1:

$$\boxed{\text{Step 3 :}} \quad \mathbf{v}_3 \rightsquigarrow \mathbf{w}_3 = \mathbf{v}_3 - \text{proj}_{\mathbf{q}_1} \mathbf{v}_3 - \text{proj}_{\mathbf{q}_2} \mathbf{v}_3 \quad \text{and} \quad \mathbf{w}_3 \rightsquigarrow \mathbf{q}_3 = \frac{\mathbf{w}_3}{\|\mathbf{w}_3\|}$$

... so far and so forth, until the last step where we modify \mathbf{v}_n so that it is perpendicular to $\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_{n-1}$ and then renormalize it to have length 1:

$$\boxed{\text{Step } n :} \quad \mathbf{v}_n \rightsquigarrow \mathbf{w}_n = \mathbf{v}_n - \text{proj}_{\mathbf{q}_1} \mathbf{v}_n - \dots - \text{proj}_{\mathbf{q}_{n-1}} \mathbf{v}_n \quad \text{and} \quad \mathbf{w}_n \rightsquigarrow \mathbf{q}_n = \frac{\mathbf{w}_n}{\|\mathbf{w}_n\|}$$

At the end of this procedure, all the vectors $\mathbf{q}_1, \dots, \mathbf{q}_n$ are orthogonal to each other by construction, and they all have length 1. This is because at each step, we subtract from every \mathbf{v}_i its projection onto the already constructed vectors $\mathbf{q}_1, \dots, \mathbf{q}_{i-1}$. What remains will be orthogonal to these vectors. The various projections that one must calculate in this process, namely $\text{proj}_{\mathbf{q}_j} \mathbf{v}_i$ can be computed using formula (95) (which is even simpler in our case, since the denominator of (95) is equal to 1).

Example 4. Let us construct an orthonormal basis of the vector space V of (125). First pick any basis, i.e. any two vectors \mathbf{v}_1 and \mathbf{v}_2 in the vector space V (being in the vector space V just means that the coordinates of these vectors should sum up to 0) which are not multiples of each other, say:

$$\mathbf{v}_1 = \begin{bmatrix} 1 \\ 2 \\ -3 \end{bmatrix} \quad \text{and} \quad \mathbf{v}_2 = \begin{bmatrix} 0 \\ 1 \\ -1 \end{bmatrix}$$

Step 1 is to rescale \mathbf{v}_1 in order for it to have length 1; the way to do so is to divide \mathbf{v}_1 by its length, which is $\|\mathbf{v}_1\| = \sqrt{14}$:

$$\mathbf{q}_1 = \frac{1}{\sqrt{14}} \begin{bmatrix} 1 \\ 2 \\ -3 \end{bmatrix}$$

Then we need to modify \mathbf{v}_2 in such a way as to be perpendicular to \mathbf{q}_1 . As we have seen in Step 2 above, the way to do so is to subtract from \mathbf{v}_2 its projection onto \mathbf{q}_1 :

$$\mathbf{w}_2 = \mathbf{v}_2 - \text{proj}_{\mathbf{q}_1} \mathbf{v}_2 \stackrel{(95)}{=} \begin{bmatrix} 0 \\ 1 \\ -1 \end{bmatrix} - \frac{1}{\sqrt{14}} \begin{bmatrix} 1 \\ 2 \\ -3 \end{bmatrix} \frac{1}{\sqrt{14}} [1 \ 2 \ -3] \begin{bmatrix} 0 \\ 1 \\ -1 \end{bmatrix}$$

(the reason why the denominator of (95) doesn't show up in the formula above is that $\|\mathbf{q}_1\| = 1$ by construction). Explicit computation shows that the vector above is:

$$\mathbf{w}_2 = \frac{1}{14} \begin{bmatrix} -5 \\ 4 \\ 1 \end{bmatrix}$$

The last thing we need to do is to rescale \mathbf{w}_2 in order for it to have length 1, which we do by dividing it by its length $\|\mathbf{w}_2\| = \sqrt{3/14}$:

$$\mathbf{q}_2 = \frac{1}{\sqrt{42}} \begin{bmatrix} -5 \\ 4 \\ 1 \end{bmatrix}$$

The vectors \mathbf{q}_1 and \mathbf{q}_2 computed above determine an orthonormal basis of the vector space V : it is easy to see that they both lie in V , that they both have length 1, and that they are perpendicular.

The example above shows two important things about the Gram-Schmidt process. Firstly, it can produce some square roots, so don't be scared of getting your hands a bit dirty with the computation, because the final answer might turn out to be reasonably nice. And secondly, there exist many possible choices of orthonormal bases of a vector space, e.g. just compare the vectors we guessed in (126) with those produced by the example above.

Let's see what the Gram-Schmidt process tells us about the relation between the matrix:

$$A = [\mathbf{v}_1 \mid \dots \mid \mathbf{v}_n] \quad \text{and} \quad Q = [\mathbf{q}_1 \mid \dots \mid \mathbf{q}_n]$$

where $\mathbf{v}_1, \dots, \mathbf{v}_n$ are the input vectors and $\mathbf{q}_1, \dots, \mathbf{q}_n$ are the output vectors of Gram-Schmidt. At each step, we are modifying the i -th column of a matrix by subtracting from it a linear combination of the previous columns, and then multiplying it by a scalar. This is doing the same thing for columns as row operations were doing for rows. Therefore, it shouldn't surprise you that:

- Adding the j -th column times λ to the i -th column of A is achieved by multiplying the latter on the right with an elimination matrix:

$$A \rightsquigarrow AE_{ji}^{(\lambda)} \quad \text{where} \quad E_{ji}^{(\lambda)} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & \lambda & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

The entry λ is on the j -th row and i -th column (in the formula above, $i = 4$ and $j = 2$).

- Multiplying the i -th column of A by the scalar λ is achieved by multiplying A on the right with a diagonal matrix:

$$A \rightsquigarrow AD_i^{(\lambda)} \quad \text{where} \quad D_i^{(\lambda)} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & \lambda & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

The entry λ is on i -th row (in the formula above, $i = 2$).

Compare these with the operations you learned in Lecture 2 when discussing Gaussian elimination and LU factorizations: the upshot is that while row operations correspond to multiplication on the left, column operations correspond to multiplication on the right. Therefore, we conclude that the connection between the input and the output of Gram-Schmidt is:

$$Q = A \underbrace{D_1^{(\mu_1)}}_{\text{first step}} \underbrace{E_{12}^{(\lambda_{12})} D_2^{(\mu_2)}}_{\text{second step}} \underbrace{E_{13}^{(\lambda_{13})} E_{23}^{(\lambda_{23})} D_3^{(\mu_3)}}_{\text{third step}} \dots \underbrace{E_{1n}^{(\lambda_{1n})} \dots E_{n-1,n}^{(\lambda_{n-1,n})} D_n^{(\mu_n)}}_{\text{last step}} \quad (127)$$

where the scalars μ_i and λ_{ij} are determined by the orthonormality requirement at every stage of Gram-Schmidt. By moving all the E 's and D 's in the other side of the equation, we obtain:

$$A = Q D_n^{(\frac{1}{\mu_n})} E_{n-1,n}^{(-\lambda_{n-1,n})} \dots E_{1n}^{(-\lambda_{1n})} \dots D_3^{(\frac{1}{\mu_3})} E_{23}^{(-\lambda_{23})} E_{13}^{(-\lambda_{13})} D_2^{(\frac{1}{\mu_2})} E_{12}^{(-\lambda_{12})} D_1^{(\frac{1}{\mu_1})}$$

If we multiply the D and E matrices together, we will obtain an upper triangular matrix R . Hence:

Fact 10. Any matrix A with linearly independent columns can be uniquely written as:

$$\boxed{A = QR} \quad (128)$$

where the columns of Q are orthonormal and R is an upper triangular square matrix:

$$R = \begin{bmatrix} * & * & * & * \\ 0 & * & * & * \\ 0 & 0 & * & * \\ 0 & 0 & 0 & * \end{bmatrix}$$

Note that this result also makes sense for non-square matrices A . As we have seen in (129), computing the QR factorization of a matrix is achieved by performing Gram-Schmidt on its columns.

MIDTERM 1 (March 26)

Lecture 16 (March 29)

Reading: Section 8.1

Viewing: Lecture 30

We will now introduce the ultimate level of abstraction in our course, which will finally explain the meaning of matrix multiplication. Recall n -dimensional space, which as a set is given by:

$$\mathbb{R}^n = \left\{ \begin{bmatrix} x_1 \\ \vdots \\ x_n \end{bmatrix} \text{ where } x_1, \dots, x_n \in \mathbb{R} \right\}$$

Definition 16. A *linear transformation* is a function:

$$\phi : \mathbb{R}^n \rightarrow \mathbb{R}^m \quad (129)$$

which respects the vector space structures of \mathbb{R}^n and \mathbb{R}^m , namely addition and scalar multiplication:

$$\phi(\mathbf{v} + \mathbf{v}') = \phi(\mathbf{v}) + \phi(\mathbf{v}') \quad \text{and} \quad \phi(\alpha \cdot \mathbf{v}) = \alpha \cdot \phi(\mathbf{v}) \quad (130)$$

for any vectors $\mathbf{v}, \mathbf{v}' \in \mathbb{R}^n$ and any scalar $\alpha \in \mathbb{R}$.

Let us consider a linear transformation ϕ as in (129), and express it in terms of the basis:

$$\mathbf{e}_j = \begin{bmatrix} 0 \\ \vdots \\ 1 \\ \vdots \\ 0 \end{bmatrix}$$

(whose only entry 1 is on the j -th row). The value $\phi(\mathbf{e}_j)$ will be a vector in \mathbb{R}^m , and therefore a linear combination of the basis vectors \mathbf{e}_i . This means that there exist constants $a_{ij} \in \mathbb{R}$ such that:

$$\phi(\mathbf{e}_j) = \sum_{i=1}^m a_{ij} \mathbf{e}_i \quad \forall j \in \{1, \dots, n\} \quad (131)$$

But once you know the numbers a_{ij} , the linear transformation ϕ is uniquely determined! Indeed,

since any vector $\mathbf{v} = \begin{bmatrix} v_1 \\ \vdots \\ v_n \end{bmatrix} \in \mathbb{R}^n$ is a linear combination of the \mathbf{e}_i 's, formulas (130) and (131) imply:

$$\phi(\mathbf{v}) = \phi(v_1 \mathbf{e}_1 + \dots + v_n \mathbf{e}_n) = \sum_{j=1}^n v_j \phi(\mathbf{e}_j) = \sum_{i=1}^m \sum_{j=1}^n a_{ij} v_j \mathbf{e}_i$$

which completely determines $\phi(\mathbf{v})$ for any \mathbf{v} . In column matrix form, this relation reads:

$$\phi \begin{bmatrix} v_1 \\ \vdots \\ v_n \end{bmatrix} = \begin{bmatrix} a_{11}v_1 + \dots + a_{1n}v_n \\ \vdots \\ a_{m1}v_1 + \dots + a_{mn}v_n \end{bmatrix} \quad (132)$$

Does this look familiar? Indeed, ϕ acts on column vectors precisely like the matrix:

$$A = \begin{bmatrix} a_{11} & \cdots & a_{1n} \\ \vdots & \ddots & \vdots \\ a_{m1} & \cdots & a_{mn} \end{bmatrix} \quad (133)$$

multiplies vectors. We conclude the following:

Fact 11. *There is a one-to-one correspondence between linear transformations and matrices, in that the matrix (133) represents the linear transformation (132). In symbols, this means:*

$$\boxed{\phi(\mathbf{v}) = A\mathbf{v}} \quad (134)$$

for all vectors \mathbf{v} .

Let us discuss the 2×2 case in detail, by showing which geometrically relevant linear transformations $\phi : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ correspond to which matrices:

- $\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$ is reflection in the line $\{x = y\}$
- $\begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$ is projection onto the x -axis
- $\begin{bmatrix} \lambda & 0 \\ 0 & \mu \end{bmatrix}$ is a scaling by a factor of λ in the x direction and by a factor of μ in the y direction
- $\begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix}$ is counter-clockwise rotation by the angle θ around the origin
- $\begin{bmatrix} 1 & c \\ 0 & 1 \end{bmatrix}$ is a shearing parallel to the x axis

As you can see from these examples, linear transformations send lines to lines and geometric shapes of a certain kind (triangles, quadrilaterals, ellipses) to shapes of the same kind. Another big class of linear transformations are projections, and formula (102) basically states that the function:

$$\text{proj}_V : \mathbb{R}^n \rightarrow \mathbb{R}^n$$

which projects onto a given m -dimensional subspace V corresponds to the matrix:

$$A(A^T A)^{-1} A^T$$

where the columns of the $n \times m$ matrix A are given by any basis of V .

Note that composition of linear transformations corresponds to multiplication of matrices. In symbols, if the linear transformations ϕ_1, ϕ_2 are given by matrices A_1, A_2 , respectively, then:

$$\boxed{\text{the transformation } \phi_1 \circ \phi_2 \text{ corresponds to the matrix } A_1 A_2} \quad (135)$$

For example, rotation by $\frac{\pi}{2}$ degrees followed by a horizontal shearing is given by the matrix:

$$\begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix} = \begin{bmatrix} 1 & -1 \\ 1 & 0 \end{bmatrix}$$

Take now $m = n$. If we have a function $\phi : \mathbb{R}^n \rightarrow \mathbb{R}^n$, we can talk about its **inverse** $\phi^{-1} : \mathbb{R}^n \rightarrow \mathbb{R}^n$: the latter is connected to the former by the condition that if $\phi(\mathbf{v}) = \mathbf{w}$, then $\phi^{-1}(\mathbf{w}) = \mathbf{v}$. It may not surprise you that if the linear transformation ϕ corresponds to a square matrix A , then:

$$\boxed{\text{the transformation } \phi^{-1} \text{ corresponds to the matrix } A^{-1}} \quad (136)$$

The function ϕ has an inverse precisely if and only if the matrix A is invertible.

The nullspace and column space of a matrix A also have an interpretation in terms of the corresponding linear transformation $\phi : \mathbb{R}^n \rightarrow \mathbb{R}^m$:

$$N(A) = \text{Ker } \phi \quad \text{and} \quad C(A) = \text{Im } \phi \quad (137)$$

where the notions of **kernel** and **image** of a linear transformation ϕ are defined by:

$$\begin{aligned} \text{Ker } \phi &= \left\{ \mathbf{v} \in \mathbb{R}^n \text{ such that } \phi(\mathbf{v}) = \mathbf{0} \right\} \\ \text{Im } \phi &= \left\{ \mathbf{w} \in \mathbb{R}^m \text{ such that } \mathbf{w} = \phi(\mathbf{v}) \text{ for some } \mathbf{v} \in \mathbb{R}^n \right\} \end{aligned}$$

Lecture 17 (March 31)

Reading: Section 8.2

Viewing: Lecture 31

In our last lecture, we said that a linear transformation $\phi : \mathbb{R}^n \rightarrow \mathbb{R}^m$ corresponds to an $m \times n$ matrix A . However, this is not completely precise: while a linear transformation ϕ is an intrinsic geometric notion (namely a rule that sends points to points and lines to lines) the matrix A is simply a representation of it. For example, think of elements of \mathbb{R}^2 as points in the plane, and consider the linear transformation:

$$\phi : \mathbb{R}^2 \rightarrow \mathbb{R}^2$$

given by the following rule for all points p in the plane:

$$\phi(p) = (p \text{ rotated counterclockwise by } 30^\circ)$$

The matrix $A = \frac{1}{2} \begin{bmatrix} \sqrt{3} & -1 \\ 1 & \sqrt{3} \end{bmatrix}$ is a representation of this rotation in the usual (x, y) system of coordinates, i.e. in the usual basis $\mathbf{e}_1 = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$, $\mathbf{e}_2 = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$ of \mathbb{R}^2 . Indeed, all that this is saying is that:

$$\phi(x \cdot \mathbf{e}_1 + y \cdot \mathbf{e}_2) = x \cdot \mathbf{a}_1 + y \cdot \mathbf{a}_2 = \left(\frac{x\sqrt{3}}{2} - \frac{y}{2} \right) \cdot \mathbf{e}_1 + \left(\frac{x}{2} + \frac{y\sqrt{3}}{2} \right) \cdot \mathbf{e}_2 \quad (138)$$

where \mathbf{a}_1 and \mathbf{a}_2 are the columns of A . But now suppose you wanted to express this same geometric transformation in a different reference frame, i.e. changing the basis from $\mathbf{e}_1, \mathbf{e}_2$ to:

$$\mathbf{v}_1 = 2\mathbf{e}_1, \quad \mathbf{v}_2 = \mathbf{e}_1 - \mathbf{e}_2.$$

When you change basis from \mathbf{e} 's to \mathbf{v} 's, you can always change back, i.e. express the \mathbf{e} 's in terms of \mathbf{v} 's. In the case above, we have (solving for \mathbf{e} 's in terms of \mathbf{v} 's):

$$\mathbf{e}_1 = \frac{\mathbf{v}_1}{2}, \quad \mathbf{e}_2 = \frac{\mathbf{v}_1}{2} - \mathbf{v}_2$$

We may plug these formulas into (138) and we get:

$$\phi(x \cdot \mathbf{v}_1 + y \cdot \mathbf{v}_2) = \phi(x \cdot 2\mathbf{e}_1 + y \cdot (\mathbf{e}_1 - \mathbf{e}_2)) = \phi((2x + y) \cdot \mathbf{e}_1 - y \cdot \mathbf{e}_2) \stackrel{(138)}{=} \quad (139)$$

$$\begin{aligned} & \left[\frac{(2x + y)\sqrt{3}}{2} + \frac{y}{2} \right] \mathbf{e}_1 + \left[\frac{2x + y}{2} - \frac{y\sqrt{3}}{2} \right] \mathbf{e}_2 = \left[\frac{(2x + y)\sqrt{3}}{2} + \frac{y}{2} \right] \frac{\mathbf{v}_1}{2} + \left[\frac{2x + y}{2} - \frac{y\sqrt{3}}{2} \right] (\frac{\mathbf{v}_1}{2} - \mathbf{v}_2) \\ & = \left[x \cdot \frac{\sqrt{3} + 1}{2} + y \cdot \frac{1}{2} \right] \cdot \mathbf{v}_1 + \left[x \cdot (-1) + y \cdot \frac{\sqrt{3} - 1}{2} \right] \cdot \mathbf{v}_2 \end{aligned}$$

So in terms of the new basis $\mathbf{v}_1, \mathbf{v}_2$, the linear transformation ϕ is represented by the matrix:

$$B = \begin{bmatrix} \frac{\sqrt{3}+1}{2} & \frac{1}{2} \\ -1 & \frac{\sqrt{3}-1}{2} \end{bmatrix}$$

The connection between this matrix and the original matrix A is precisely:

$$\boxed{B = V^{-1}AV} \quad (140)$$

where $V = \begin{bmatrix} 2 & 1 \\ 0 & -1 \end{bmatrix}$ is the matrix whose columns are the vectors $\mathbf{v}_1, \mathbf{v}_2$ written in the $\mathbf{e}_1, \mathbf{e}_2$ basis, i.e.:

$$V\mathbf{e}_1 = \mathbf{v}_1 \quad \text{and} \quad V\mathbf{e}_2 = \mathbf{v}_2 \quad (141)$$

In other words, V is the matrix whose columns are the vectors \mathbf{v}_1 and \mathbf{v}_2 . We call V the:

change of basis matrix from the basis $\mathbf{v}_1, \mathbf{v}_2$ to the basis $\mathbf{e}_1, \mathbf{e}_2$

Remark. You might be a bit surprised that we call (141) the “change of basis matrix from the basis $\mathbf{v}_1, \mathbf{v}_2$ to the basis $\mathbf{e}_1, \mathbf{e}_2$ ” and not vice-versa. This is a terminology issue, but let me try to justify it. Take a random vector, say $\mathbf{v} = 2\mathbf{v}_1 + 3\mathbf{v}_2$. In the basis $\mathbf{v}_1, \mathbf{v}_2$:

$$\mathbf{v} \text{ is represented by } \begin{bmatrix} 2 \\ 3 \end{bmatrix}$$

But in the standard basis $\mathbf{e}_1, \mathbf{e}_2$, since $\mathbf{v}_1 = V\mathbf{e}_1$ and $\mathbf{v}_2 = V\mathbf{e}_2$:

$$\mathbf{v} \text{ is represented by } V \begin{bmatrix} 2 \\ 3 \end{bmatrix}$$

To summarize: if an arbitrary vector $\mathbf{v} \in \mathbb{R}^2$ is represented by a column in the basis $\mathbf{v}_1, \mathbf{v}_2$, then it is represented by V times that column in the basis $\mathbf{e}_1, \mathbf{e}_2$. Hence the “...from...to...” terminology.

So we have seen that the matrices A and B connected by (140) describe the same linear transformation, albeit in different bases. Let us explain the general principle for any linear transformation:

$$\phi : \mathbb{R}^n \rightarrow \mathbb{R}^n$$

which is represented by the $n \times n$ matrix $A = \begin{bmatrix} a_{11} & \dots & a_{1n} \\ \vdots & \ddots & \vdots \\ a_{n1} & \dots & a_{nn} \end{bmatrix}$ in the standard basis:

$$\phi(x_1\mathbf{e}_1 + \dots + x_n\mathbf{e}_n) = (a_{11}x_1 + \dots + a_{1n}x_n)\mathbf{e}_1 + \dots + (a_{n1}x_1 + \dots + a_{nn}x_n)\mathbf{e}_n \quad (142)$$

If you wish to express the linear transformation ϕ in a different basis $\mathbf{v}_1, \dots, \mathbf{v}_n$, then you need to consider the **similar** matrix:

$$B = V^{-1}AV = \begin{bmatrix} b_{11} & \dots & b_{1n} \\ \vdots & \ddots & \vdots \\ b_{n1} & \dots & b_{nn} \end{bmatrix}$$

where $V = [\mathbf{v}_1 \mid \dots \mid \mathbf{v}_n]$ is the change of basis matrix from the basis $\mathbf{v}_1, \dots, \mathbf{v}_n$ to the basis $\mathbf{e}_1, \dots, \mathbf{e}_n$. In terms of the $\mathbf{v}_1, \dots, \mathbf{v}_n$ basis, formula (142) may be rewritten as:

$$\phi(x_1\mathbf{v}_1 + \dots + x_n\mathbf{v}_n) = (b_{11}x_1 + \dots + b_{1n}x_n)\mathbf{v}_1 + \dots + (b_{n1}x_1 + \dots + b_{nn}x_n)\mathbf{v}_n \quad (143)$$

Example 5. The projection onto the vector $\begin{bmatrix} 1 \\ 1 \end{bmatrix}$ in \mathbb{R}^2 is represented by the matrix:

$$A = \begin{bmatrix} 1 \\ 1 \end{bmatrix} \left(\begin{bmatrix} 1 & 1 \end{bmatrix} \begin{bmatrix} 1 \\ 1 \end{bmatrix} \right)^{-1} \begin{bmatrix} 1 & 1 \end{bmatrix} = \begin{bmatrix} \frac{1}{2} & \frac{1}{2} \\ \frac{1}{2} & \frac{1}{2} \end{bmatrix}$$

in the usual $\mathbf{e}_1, \mathbf{e}_2$ basis of the plane. But suppose you wanted to represent the same linear transformation in the basis $\mathbf{v}_1 = 2\mathbf{e}_1 + 2\mathbf{e}_2$, $\mathbf{v}_2 = 3\mathbf{e}_1 - 3\mathbf{e}_2$ instead. To do so, consider the change of basis matrix:

$$V = \begin{bmatrix} 2 & 3 \\ 2 & -3 \end{bmatrix}$$

Then formula (140) says that the given projection linear transformation is represented by the matrix:

$$B = V^{-1}AV = \begin{bmatrix} \frac{1}{4} & \frac{1}{4} \\ \frac{1}{6} & -\frac{1}{6} \end{bmatrix} \begin{bmatrix} \frac{1}{2} & \frac{1}{2} \\ \frac{1}{2} & \frac{1}{2} \end{bmatrix} \begin{bmatrix} 2 & 3 \\ 2 & -3 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$$

in the new $\mathbf{v}_1, \mathbf{v}_2$ basis. It's not a coincidence that the answer is such a simple matrix. Indeed, all that the formula for B is saying is that:

$$\text{proj}_{\begin{bmatrix} 1 \\ 1 \end{bmatrix}}(x \cdot \mathbf{v}_1 + y \cdot \mathbf{v}_2) = x \cdot \mathbf{v}_1$$

Geometrically, this says that projection sends \mathbf{v}_1 to itself (i.e. leaves it unchanged, which makes sense given that \mathbf{v}_1 is on the line being projected upon) and sends \mathbf{v}_2 to 0 (which makes sense given that $\mathbf{v}_1 \perp \mathbf{v}_2$).

To summarize: expressing one and the same linear transformation in different bases (i.e. reference frames) is achieved by multiplying the corresponding matrix on the left and on the right by appropriate change of basis matrices. The key notion here is that of the change of basis matrix:

$$V = [\mathbf{v}_1 \mid \dots \mid \mathbf{v}_n] \quad \Leftrightarrow \quad V\mathbf{e}_i = \mathbf{v}_i \quad \forall i \in \{1, \dots, n\}$$

for the change of basis matrix from the basis $\mathbf{v}_1, \dots, \mathbf{v}_n$ to the basis $\mathbf{e}_1, \dots, \mathbf{e}_n$.

Fact 12. All change of basis matrices are invertible.

Indeed, if V changes from the basis $\mathbf{v}_1, \dots, \mathbf{v}_n$ to $\mathbf{e}_1, \dots, \mathbf{e}_n$, then its inverse should change back:

$$V^{-1}\mathbf{v}_i = \mathbf{e}_i \quad \forall i \in \{1, \dots, n\}$$

from the basis $\mathbf{e}_1, \dots, \mathbf{e}_n$ to $\mathbf{v}_1, \dots, \mathbf{v}_n$. So in particular the inverse should exist.

What if we wanted to change between two general bases, say from $\mathbf{v}_1, \dots, \mathbf{v}_n$ to $\mathbf{w}_1, \dots, \mathbf{w}_n$? What is the change of basis matrix which achieves that? We can think about this as a two-step process:

$$\{\mathbf{v}_1, \dots, \mathbf{v}_n\} \xrightarrow{\text{change basis using matrix } V} \{\mathbf{e}_1, \dots, \mathbf{e}_n\} \xrightarrow{\text{change basis using matrix } W^{-1}} \{\mathbf{w}_1, \dots, \mathbf{w}_n\}$$

So if we had a vector \mathbf{a} expressed in the bases $\mathbf{v}_1, \dots, \mathbf{v}_n$ and $\mathbf{w}_1, \dots, \mathbf{w}_n$ and $\mathbf{e}_1, \dots, \mathbf{e}_n$ as:

$$\mathbf{a} = x_1\mathbf{v}_1 + \dots + x_n\mathbf{v}_n = y_1\mathbf{w}_1 + \dots + y_n\mathbf{w}_n = z_1\mathbf{e}_1 + \dots + z_n\mathbf{e}_n$$

then the Remark above says that the vectors of coefficients $\mathbf{x} = \begin{bmatrix} x_1 \\ \vdots \\ x_n \end{bmatrix}$ and $\mathbf{y} = \begin{bmatrix} y_1 \\ \vdots \\ y_n \end{bmatrix}$ are connected

to the vector of coefficients $\mathbf{z} = \begin{bmatrix} z_1 \\ \vdots \\ z_n \end{bmatrix}$ by:

$$\mathbf{z} = V\mathbf{x} \quad \text{and} \quad \mathbf{z} = W\mathbf{y}$$

Setting the two formulas equal to each other implies that \mathbf{x} and \mathbf{y} are connected by:

$$\boxed{\mathbf{y} = W^{-1}V\mathbf{x}}$$

which is the way we change from the \mathbf{v} basis to the \mathbf{w} basis. In other words, the entries of the matrix $W^{-1}V$ give the coordinates of the vectors $\mathbf{v}_1, \dots, \mathbf{v}_n$ when expressed in the basis $\mathbf{w}_1, \dots, \mathbf{w}_n$:

$$W^{-1}V = \begin{bmatrix} c_{11} & \dots & c_{1n} \\ \vdots & \ddots & \vdots \\ c_{n1} & \dots & c_{nn} \end{bmatrix} \quad \text{where} \quad \mathbf{v}_j = \sum_{i=1}^n c_{ij}\mathbf{w}_i \quad \forall j \in \{1, \dots, n\}$$

Lecture 18 (April 2)

Reading: Section 8.2

Viewing: Lecture 31

In the previous Lecture, we always worked in the same basis on the domain (source) of the linear transformation ϕ as on the codomain (target). But now we will see what happens if we allow different bases on the domain and codomain. So let us consider a general linear transformation:

$$\phi : \mathbb{R}^n \rightarrow \mathbb{R}^m$$

which is represented by the $m \times n$ matrix $A = \begin{bmatrix} a_{11} & \cdots & a_{1n} \\ \vdots & \ddots & \vdots \\ a_{m1} & \cdots & a_{mn} \end{bmatrix}$ in the standard bases:

$$\phi(x_1\mathbf{e}_1 + \cdots + x_n\mathbf{e}_n) = (a_{11}x_1 + \cdots + a_{1n}x_n)\mathbf{e}_1 + \cdots + (a_{m1}x_1 + \cdots + a_{mn}x_n)\mathbf{e}_m \quad (144)$$

Since \mathbb{R}^m and \mathbb{R}^n are (in general) different vector spaces, it will often make sense to change the basis in the input \mathbb{R}^n independently of the change of basis in the output \mathbb{R}^m . So let us fix bases:

$$\mathbf{v}_1, \dots, \mathbf{v}_n \text{ of } \mathbb{R}^n \quad \text{and} \quad \mathbf{w}_1, \dots, \mathbf{w}_m \text{ of } \mathbb{R}^m$$

If you wish to express the linear transformation ϕ in these new bases, consider the matrices:

$$V = [\mathbf{v}_1 \mid \cdots \mid \mathbf{v}_n] \quad \text{and} \quad W = [\mathbf{w}_1 \mid \cdots \mid \mathbf{w}_m]$$

As we have seen in the previous lecture, V is the change of basis matrix from the basis $\mathbf{v}_1, \dots, \mathbf{v}_n$ to the basis $\mathbf{e}_1, \dots, \mathbf{e}_n$, and W is the change of basis matrix from the basis $\mathbf{w}_1, \dots, \mathbf{w}_m$ to the basis $\mathbf{e}_1, \dots, \mathbf{e}_m$. Therefore, W^{-1} is the change of basis matrix from the basis $\mathbf{e}_1, \dots, \mathbf{e}_m$ to the basis $\mathbf{w}_1, \dots, \mathbf{w}_m$. We conclude that the linear transformation ϕ is represented by the matrix:

$$\boxed{B = W^{-1}AV} \quad (145)$$

in the new bases $\mathbf{v}_1, \dots, \mathbf{v}_n$ and $\mathbf{w}_1, \dots, \mathbf{w}_m$. Explicitly, this means that if:

$$B = \begin{bmatrix} b_{11} & \cdots & b_{1n} \\ \vdots & \ddots & \vdots \\ b_{m1} & \cdots & b_{mn} \end{bmatrix}$$

then formula (144) is equivalent to:

$$\phi(x_1\mathbf{v}_1 + \cdots + x_n\mathbf{v}_n) = (b_{11}x_1 + \cdots + b_{1n}x_n)\mathbf{w}_1 + \cdots + (b_{m1}x_1 + \cdots + b_{mn}x_n)\mathbf{w}_m \quad (146)$$

Formula (145) is called the **change of basis** formula.

The ability to change bases gives us a lot of freedom in studying matrices. Basically, you can always change bases in order to simplify any given matrix, and how far you go depends on the problem at hand. For example, in the square $n \times n$ case, almost any matrix is similar to a diagonal matrix:

$$A = V \begin{bmatrix} d_1 & 0 & \cdots & 0 \\ 0 & d_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & d_n \end{bmatrix} V^{-1} \quad (147)$$

In future weeks, we will see the meaning of the numbers d_1, \dots, d_n (the eigenvalues of A) and of the columns of the matrix V (the eigenvectors of A). We will also learn how to eliminate the word “almost” in the previous paragraph, by studying Jordan normal forms.

Example 6. Consider the linear transformation $\phi : \mathbb{R}^3 \rightarrow \mathbb{R}^2$ represented by the matrix:

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

In the standard bases, this linear transformation sends a vector:

$$x_1\mathbf{e}_1 + x_2\mathbf{e}_2 + x_3\mathbf{e}_3 \quad \text{to} \quad (x_1 + x_2 + 2x_3)\mathbf{e}_1 + (2x_1 + 3x_2 + 4x_3)\mathbf{e}_2$$

But can we change bases so that ϕ is represented by a “nicer” matrix than A ? Say we wanted to find bases $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3$ of \mathbb{R}^3 and $\mathbf{w}_1, \mathbf{w}_2$ of \mathbb{R}^2 such that the same linear transformation ϕ is represented by the simple matrix:

$$B = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix}$$

in the new bases. In other words, the same linear transformation ϕ should send a vector:

$$x_1\mathbf{v}_1 + x_2\mathbf{v}_2 + x_3\mathbf{v}_3 \quad \text{to} \quad x_1\mathbf{w}_1 + x_2\mathbf{w}_2$$

Fortunately, the change of basis formula (145) tells us exactly what we need to do. We need to find invertible matrices V and W such that:

$$\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} = W^{-1} \begin{bmatrix} 1 & 1 & 2 \\ 2 & 3 & 4 \end{bmatrix} V$$

or in other words:

$$\begin{bmatrix} 1 & 1 & 2 \\ 2 & 3 & 4 \end{bmatrix} = W \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} V^{-1} \tag{148}$$

We don’t need to just guess V and W , since we already have a system in place for simplifying a matrix A : Gauss-Jordan elimination. Running this algorithm on the original matrix A gives us:

$$E_{12}^{(-1)} E_{21}^{(-2)} \begin{bmatrix} 1 & 1 & 2 \\ 2 & 3 & 4 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 2 \\ 0 & 1 & 0 \end{bmatrix}$$

By moving the elimination matrices to the right, we get:

$$\begin{bmatrix} 1 & 1 & 2 \\ 2 & 3 & 4 \end{bmatrix} = E_{21}^{(2)} E_{12}^{(1)} \begin{bmatrix} 1 & 0 & 2 \\ 0 & 1 & 0 \end{bmatrix} = \underbrace{\begin{bmatrix} 1 & 1 \\ 2 & 3 \end{bmatrix}}_{\text{call this matrix } W} \begin{bmatrix} 1 & 0 & 2 \\ 0 & 1 & 0 \end{bmatrix}$$

The right-most matrix isn’t quite as simple as B , and we cannot make it any simpler using row operations. But we can do column operations, which we know corresponds to multiplication on the right with elimination/diagonal/permutation matrices. In the case at hand, this is quite easy to do:

$$\begin{bmatrix} 1 & 0 & 2 \\ 0 & 1 & 0 \end{bmatrix} E_{13}^{(-2)} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix}$$

By moving the elimination matrix to the right, we get:

$$\begin{bmatrix} 1 & 0 & 2 \\ 0 & 1 & 0 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} E_{13}^{(2)} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} \underbrace{\begin{bmatrix} 1 & 0 & 2 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}}_{\text{call this matrix } V^{-1}}$$

Therefore, we have achieved the formula $A = WBV^{-1}$ using the matrices:

$$V = \begin{bmatrix} 1 & 0 & -2 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad \left(\text{this is equivalent to } V^{-1} = \begin{bmatrix} 1 & 0 & 2 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \right) \quad \text{and} \quad W = \begin{bmatrix} 1 & 1 \\ 2 & 3 \end{bmatrix}$$

Hence the desired bases have $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3$ and $\mathbf{w}_1, \mathbf{w}_2$ equal to the columns of V and W , respectively.

Lecture 19 (April 5)

Reading: Section 5.1

Viewing: Lecture 18

Consider any linear transformation $\phi : \mathbb{R}^n \rightarrow \mathbb{R}^n$, represented by an $n \times n$ square matrix A .

Definition 17. The *determinant* of the square matrix A is the number:

$$\boxed{\det A} \tag{149}$$

defined as the factor by which the linear transformation ϕ scales the n -dimensional volumes of regions in \mathbb{R}^n (the determinant can be negative if ϕ switches left-handedness to right-handedness).

Let's take for instance the linear transformations of \mathbb{R}^2 considered immediately after Fact 11. Two-dimensional volume just means area. Any reflection preserves the magnitude of areas but switches handedness (the mirror images of left hands are right hands), hence:

$$\det \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} = -1 \tag{150}$$

Projections flatten regions $R \subset \mathbb{R}^2$ to line segments, which have area 0, so therefore:

$$\det \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} = 0 \tag{151}$$

Scaling transformations increase areas by the product of the scaling factors, and therefore:

$$\det \begin{bmatrix} \lambda & 0 \\ 0 & \mu \end{bmatrix} = \lambda\mu \tag{152}$$

Finally, rotations and shearings preserve areas, so we have:

$$\det \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix} = 1$$
$$\det \begin{bmatrix} 1 & c \\ 0 & 1 \end{bmatrix} = 1 \tag{153}$$

In general, the determinant of a 2×2 matrix is given by the following formula:

$$\boxed{\det \begin{bmatrix} a & b \\ c & d \end{bmatrix} = ad - bc} \tag{154}$$

and indeed, the determinant plays an important role in the formula for the inverse of a 2×2 matrix:

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix}^{-1} = \frac{1}{ad - bc} \begin{bmatrix} d & -b \\ -c & a \end{bmatrix} \tag{155}$$

But before we work out general means for computing determinants, let us spell out two key facts.

Fact 13. For any $n \times n$ matrices A and B , we have:

$$\det AB = (\det A)(\det B) \tag{156}$$

Let's give a quick proof of this fact: if A and B correspond to linear transformations ϕ and ψ , respectively, then we have seen in (135) that AB corresponds to the composition $\phi \circ \psi$ of the two transformations. Then (156) is simply saying that the rate at which $\phi \circ \psi$ scales volumes is equal to the product of rates at which ϕ and ψ individually scale volumes. Indeed, if you take any region $R \subset \mathbb{R}^n$ and assume that:

$$\begin{cases} \psi \text{ takes } R \text{ to some region } R' \\ \phi \text{ takes } R' \text{ to some region } R'' \end{cases} \Rightarrow \phi \circ \psi \text{ takes } R \text{ to the region } R''$$

With this in mind, we have:

$$\text{LHS of (156)} = \frac{\text{vol } R''}{\text{vol } R} = \frac{\text{vol } R''}{\text{vol } R'} \cdot \frac{\text{vol } R'}{\text{vol } R} = \text{RHS of (156)}$$

which proves (156). Caveat: it is not true that $\det(A + B) = \det A + \det B$, in general!

Fact 14. Row operations have the following effect on determinants:

- adding a multiple of one row to another leaves the determinant unchanged
- exchanging two rows multiplies the determinant by -1
- multiplying a single row by λ has the effect of multiplying the determinant by λ

The analogous properties still hold if the word “row” is replaced by “column” everywhere.

Fact 14 actually follows from Fact 13, as we will now proceed to show. Recall from Lecture 2 that performing the three kinds of row operations above on a matrix A is precisely achieved by multiplying A on the left by elimination/permutation/diagonal matrices. Then the statements in the bullets of Fact 14 become equivalent to:

$\det E_{ij}^{(\lambda)} A = \det A$	which happens because	$\det E_{ij}^{(\lambda)} = 1$
$\det P_{ij} A = -\det A$	which happens because	$\det P_{ij} = -1$
$\det D_i^{(\lambda)} A = \lambda \cdot \det A$	which happens because	$\det D_i^{(\lambda)} = \lambda$

The fact that P_{ij} , $D_i^{(\lambda)}$, $E_{ij}^{(\lambda)}$ have determinants -1 , λ , 1 respectively, is simply a natural generalization of (150), (152), (153). Fact 14 has the following important consequence:

If A is a square matrix, then $\det A = \pm$ product of pivots of REF(A) (157)

which means that a computationally efficient way to compute the determinant of a matrix is to just put it in row echelon form and multiply its pivots. Indeed, Fact 14 implies that:

$$\det A = \pm \det U \quad (158)$$

where U is the row echelon form of A , and the sign \pm is simply (-1) raised to the number of row exchanges required by Gaussian elimination. A square matrix in row echelon form is the same thing as upper triangular, and its pivots are just its diagonal entries. Therefore, to get from (158) to (157) we need to invoke the following fact, which simplifies a lot of computations:

the determinant of a triangular/diagonal matrix is the product of its diagonal entries

The fact above is true for both upper and lower triangular matrices, as well as for diagonal ones:

$$\det \begin{bmatrix} d_1 & 0 & \dots & 0 \\ * & d_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ * & * & \dots & d_n \end{bmatrix} = \det \begin{bmatrix} d_1 & 0 & \dots & 0 \\ 0 & d_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & d_n \end{bmatrix} = \det \begin{bmatrix} d_1 & * & \dots & * \\ 0 & d_2 & \dots & * \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & d_n \end{bmatrix} = d_1 d_2 \dots d_n \quad (159)$$

where the $*$'s denote arbitrary numbers. Let's prove (159) in the case of an upper triangular matrix (the other two cases featured in the formula above are analogous). We have:

$$\begin{bmatrix} d_1 & * & \dots & * \\ 0 & d_2 & \dots & * \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & d_n \end{bmatrix} = D_1^{(d_1)} \dots D_n^{(d_n)} \begin{bmatrix} 1 & \# & \dots & \# \\ 0 & 1 & \dots & \# \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 1 \end{bmatrix} \quad (160)$$

where the $*$'s and $\#$'s denote arbitrary numbers. By Gauss-Jordan elimination, the right-most matrix in (160) can be reduced to the identity matrix I_n using only the first bullet in Fact 14. As:

$$\det I_n = 1$$

(the identity linear transformation does not change volumes) then applying Fact 13 to (160) yields:

$$\det \begin{bmatrix} d_1 & * & \dots & * \\ 0 & d_2 & \dots & * \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & d_n \end{bmatrix} = \det \left(D_1^{(d_1)} \dots D_n^{(d_n)} \right) = \left(\det D_1^{(d_1)} \right) \dots \left(\det D_n^{(d_n)} \right) = d_1 \dots d_n$$

as we needed to show in (159). This completes the proof of (157).

Remark. Another way to restate formula (157) is to remember that any square matrix A has a factorization of the following form for a suitable permutation matrix P :

$$PA = LDU$$

where L and U are lower/upper triangular with all 1's on the diagonal, and D is a diagonal matrix. Applying determinant and Fact 13 to the formula above gives us;

$$(\det P)(\det A) = (\det L)(\det D)(\det U)$$

and then (157) follows from the fact that $\det P = (-1)$ raised to the number of row exchanges, while $\det L = \det U = 1$ and $\det D =$ product of pivots.

If a square matrix A has a full row of zeroes (or a full column of zeroes, for that matter), then its determinant is 0: this is simply a consequence of (157) and the fact that one of the pivots would be 0 if we had a full row of zeroes. But we also have the following more general property:

$$\text{If the rows/columns of } A \text{ are linearly dependent, then } \det A = 0 \quad (161)$$

Indeed, if the rows were dependent, then one of them (say row i) would be a linear combination of the other ones. But then by successively subtracting multiples of the other rows from row i , which we know does not change the determinant, we would arrive at a matrix with a full row of zeroes. Since such a matrix would have a zero pivot, it has determinant 0, thus establishing (161).

Expanding on the discussion above, we know exactly when a matrix has linearly dependent rows or columns: precisely when it is singular, i.e. not invertible. Therefore, (161) can be restated as:

Fact 15. *A square matrix A is non-singular (invertible) if and only if $\det A \neq 0$.*

Indeed, if the matrix A had an inverse A^{-1} , then applying determinant and Fact 13 to the formula $AA^{-1} = I$ would imply that:

$$(\det A)(\det A^{-1}) = \det I = 1 \quad \Rightarrow \quad \det A^{-1} = \frac{1}{\det A} \quad (162)$$

This immediately implies that the determinant of an invertible matrix cannot be 0. Finally, we give one more fact for the road:

Fact 16. *For any square matrix A , we have $\boxed{\det A = \det A^T}$.*

If Q is an orthogonal matrix (see Definition 15), i.e. $Q^T = Q^{-1}$, then (162) and Fact 16 imply that its determinant can only be 1 or -1 .

Lecture 20 (April 7)

Reading: Section 5.2

Viewing: Lecture 19

Formula (157) computes the determinant of a square matrix from its pivots, which is just what a computer would do. But there is a different way, namely the so-called “big formula” for the determinant, which is very important for understanding the theory and properties of the determinant.

Let’s start from the following observation: determinants are linear functions of each row separately. As a formula, this means that for any numbers x_{ij} and $a_1, \dots, a_n, b_1, \dots, b_n$, we have:

$$\det \begin{bmatrix} x_{11} & x_{12} & \dots & x_{1n} \\ \vdots & \vdots & \ddots & \vdots \\ a_1 + b_1 & a_2 + b_2 & \dots & a_n + b_n \\ \vdots & \vdots & \ddots & \vdots \\ x_{n1} & x_{n2} & \dots & x_{nn} \end{bmatrix} = \det \begin{bmatrix} x_{11} & x_{12} & \dots & x_{1n} \\ \vdots & \vdots & \ddots & \vdots \\ a_1 & a_2 & \dots & a_n \\ \vdots & \vdots & \ddots & \vdots \\ x_{n1} & x_{n2} & \dots & x_{nn} \end{bmatrix} + \det \begin{bmatrix} x_{11} & x_{12} & \dots & x_{1n} \\ \vdots & \vdots & \ddots & \vdots \\ b_1 & b_2 & \dots & b_n \\ \vdots & \vdots & \ddots & \vdots \\ x_{n1} & x_{n2} & \dots & x_{nn} \end{bmatrix} \quad (163)$$

Very importantly, the matrix on the left is **NOT** the sum of the matrices on the right. Instead, all three matrices have the same elements on all rows except for the i -th one (that’s the one where the a ’s and b ’s are in the formula), and i could be anything in $\{1, \dots, n\}$. Moreover, we have:

$$\det \begin{bmatrix} x_{11} & x_{12} & \dots & x_{1n} \\ \vdots & \vdots & \ddots & \vdots \\ \lambda a_1 & \lambda a_2 & \dots & \lambda a_n \\ \vdots & \vdots & \ddots & \vdots \\ x_{n1} & x_{n2} & \dots & x_{nn} \end{bmatrix} = \lambda \cdot \det \begin{bmatrix} x_{11} & x_{12} & \dots & x_{1n} \\ \vdots & \vdots & \ddots & \vdots \\ a_1 & a_2 & \dots & a_n \\ \vdots & \vdots & \ddots & \vdots \\ x_{n1} & x_{n2} & \dots & x_{nn} \end{bmatrix} \quad (164)$$

Remark. Properties (163) and (164) also hold if we replace rows by columns (so the matrices in the left and right-hand sides would only differ in a single column) which makes sense given the fact that determinant is unchanged by transposing matrices.

Here’s an argument for (163). You can perform row exchanges in order to make $i = n$, i.e. the three matrices only differ on the last row. Perform Gaussian elimination on the three matrices in (163). It’s not hard to see that the first $n - 1$ rows will look the same in all three matrices, so they will all have the same pivots p_1, \dots, p_{n-1} . When you finally get to do Gaussian elimination on the n -th row, denote the pivots on the n -th rows of the matrices in the right-hand side of (163) by α and β , respectively. Then a little thought shows that the pivot on the n -th row of the matrix in the left-hand side of (163) has to be $\alpha + \beta$. Then formula (157) implies that (163) boils down to:

$$\pm p_1 \dots p_{n-1} (\alpha + \beta) = \pm p_1 \dots p_{n-1} \alpha \pm p_1 \dots p_{n-1} \beta$$

which is definitely a true statement. Formula (164) is proved by a similar argument.

By applying formula (164) to all the rows of an $n \times n$ matrix, we get:

$$\boxed{\det(\lambda A) = \lambda^n \cdot \det A} \quad (165)$$

which you could also get by applying (156) to the matrix $B = \begin{bmatrix} \lambda & 0 & \dots & 0 \\ 0 & \lambda & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \lambda \end{bmatrix}$.

Let's return to computing determinants. Since $\det(A + B) \neq \det A + \det B$ in general, we must resort to formula (163) in order to break up a general matrix into simpler, more manageable pieces. Let's take the 2×2 case for illustration. By applying (163) to the first row, we have:

$$\det \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} = \det \begin{bmatrix} a_{11} & 0 \\ a_{21} & a_{22} \end{bmatrix} + \det \begin{bmatrix} 0 & a_{12} \\ a_{21} & a_{22} \end{bmatrix}$$

Then apply (163) to the second rows of each of the matrices in the right-hand side:

$$\det \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} = \det \begin{bmatrix} a_{11} & 0 \\ a_{21} & 0 \end{bmatrix} + \det \begin{bmatrix} a_{11} & 0 \\ 0 & a_{22} \end{bmatrix} + \det \begin{bmatrix} 0 & a_{12} \\ a_{21} & 0 \end{bmatrix} + \det \begin{bmatrix} 0 & a_{12} \\ 0 & a_{22} \end{bmatrix}$$

The first and fourth determinants in the right-hand side vanish because of (161), since they have full 0 columns. Therefore, we are left with:

$$\det \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} = \det \begin{bmatrix} a_{11} & 0 \\ 0 & a_{22} \end{bmatrix} + \det \begin{bmatrix} 0 & a_{12} \\ a_{21} & 0 \end{bmatrix} = a_{11}a_{22} - a_{12}a_{21}$$

Indeed, the first underlined matrix is already in row echelon form, so its determinant is the product of the two pivots. Meanwhile, the second underlined matrix would be in row echelon form as soon as we exchanged the two rows, which is responsible for the fact that its determinant has a -1 sign in front. Thus, we have proved formula (154).

In general, the procedure of successively applying (163) to the rows of the matrix leaves us with a sum of many determinants, all of which have a single non-zero entry on each row. Of all those determinants, the ones which have a full zero column will vanish, which means that the only surviving determinants are the ones where the non-zero elements are all on different rows, and also all on different columns. For example, in the 3×3 case, we have:

$$\begin{aligned} \det \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{bmatrix} &= \det \begin{bmatrix} a_{11} & 0 & 0 \\ 0 & a_{22} & 0 \\ 0 & 0 & a_{33} \end{bmatrix} + \det \begin{bmatrix} 0 & a_{12} & 0 \\ 0 & 0 & a_{23} \\ a_{31} & 0 & 0 \end{bmatrix} + \det \begin{bmatrix} 0 & 0 & a_{13} \\ a_{21} & 0 & 0 \\ 0 & a_{32} & 0 \end{bmatrix} \\ &+ \det \begin{bmatrix} a_{11} & 0 & 0 \\ 0 & 0 & a_{23} \\ 0 & a_{32} & 0 \end{bmatrix} + \det \begin{bmatrix} 0 & a_{12} & 0 \\ a_{21} & 0 & 0 \\ 0 & 0 & a_{33} \end{bmatrix} + \det \begin{bmatrix} 0 & 0 & a_{13} \\ 0 & a_{22} & 0 \\ a_{31} & 0 & 0 \end{bmatrix} \end{aligned}$$

There are as many surviving determinants as possible permutations of the rows, and indeed, all of these determinants can be transformed into diagonal matrices by an appropriate row exchange. This leads to the following formula:

$$\det \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{bmatrix} = a_{11}a_{22}a_{33} + a_{12}a_{23}a_{31} + a_{13}a_{21}a_{32} - a_{11}a_{23}a_{32} - a_{12}a_{21}a_{33} - a_{13}a_{22}a_{31} \quad (166)$$

In general, the procedure gives a formula for the determinant as a sum over all permutations:

$$\{\sigma(1), \dots, \sigma(n)\} \quad \text{of} \quad \{1, \dots, n\}$$

and we have the following **big formula for the determinant** :

$$\det \begin{bmatrix} a_{11} & a_{12} & \dots & a_{1n} \\ a_{21} & a_{22} & \dots & a_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ a_{n1} & a_{n2} & \dots & a_{nn} \end{bmatrix} = \sum_{\{\sigma(1), \dots, \sigma(n)\}}^{\text{permutations}} (-1)^{\text{sgn } \sigma} a_{1\sigma(1)} a_{2\sigma(2)} \dots a_{n\sigma(n)} \quad (167)$$

The number $\text{sgn } \sigma$ is called the **signature** of the permutation σ : it is the minimum number of successive (e.g. the i -th and $i + 1$ -th) row exchanges that you need to apply to the permutation matrix corresponding to σ , in order to turn it into a diagonal matrix.

Formula (167) is nice and explicit, but awful in practice. The number of summands for an $n \times n$ matrix is equal to the number of permutations of $\{1, \dots, n\}$, which is $n! = 1 \cdot 2 \cdot \dots \cdot n$. But there are matrices, particularly those which have a lot of zeroes, for which this formula is quite manageable. For example, if the i -th row only has a single non-zero entry (say on the j -th column) then we get:

$$\begin{aligned} \det \begin{bmatrix} \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots \\ a_{i-1,1} & \dots & a_{i-1,j-1} & a_{i-1,j} & a_{i-1,j+1} & \dots & a_{i-1,n} \\ 0 & \dots & 0 & a_{ij} & 0 & \dots & 0 \\ a_{i+1,1} & \dots & a_{i+1,j-1} & a_{i+1,j} & a_{i+1,j+1} & \dots & a_{i+1,n} \\ \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots \end{bmatrix} = \\ = \sum_{\substack{\{\sigma(1), \dots, \sigma(i-1), \sigma(i+1), \dots, \sigma(n)\} \\ \text{of } \{1, \dots, j-1, j+1, \dots, n\}}}^{\text{permutations}} (-1)^{\text{signature}} a_{1\sigma(1)} \dots a_{i-1, \sigma(i-1)} a_{ij} a_{i+1, \sigma(i+1)} \dots a_{n\sigma(n)} \\ = a_{ij} \cdot (-1)^{i+j} \det \underbrace{\begin{bmatrix} \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ a_{i-1,1} & \dots & a_{i-1,j-1} & a_{i-1,j+1} & \dots & a_{i-1,n} \\ a_{i+1,1} & \dots & a_{i+1,j-1} & a_{i+1,j+1} & \dots & a_{i+1,n} \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \end{bmatrix}}_{\text{cofactor } C_{ij}} \end{aligned} \quad (168)$$

This is because the only terms which give a non-zero contribution to (167) are those with $\sigma(i) = j$. So computing an $n \times n$ determinant whose only non-zero entry on row i is on column j boils down to computing the determinant of the $(n - 1) \times (n - 1)$ matrix obtained by removing row i and column j .

Formula (168) leads to a general way of computing determinants, known as **cofactor expansion**.

Explicitly, for a matrix $A = \begin{bmatrix} a_{11} & \dots & a_{1n} \\ \vdots & \ddots & \vdots \\ a_{n1} & \dots & a_{nn} \end{bmatrix}$ and any row i , we have:

$$\det A \stackrel{(163)}{=} \sum_{j=1}^n \det \begin{bmatrix} \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots \\ a_{i-1,1} & \dots & a_{i-1,j-1} & a_{i-1,j} & a_{i-1,j+1} & \dots & a_{i-1,n} \\ 0 & \dots & 0 & a_{ij} & 0 & \dots & 0 \\ a_{i+1,1} & \dots & a_{i+1,j-1} & a_{i+1,j} & a_{i+1,j+1} & \dots & a_{i+1,n} \\ \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots \end{bmatrix}$$

We may invoke formula (168) to compute the right-hand side, and we obtain:

$$\boxed{\det A = a_{i1} \cdot C_{i1} + \cdots + a_{in} \cdot C_{in}} \quad (169)$$

where C_{ij} is called the (i, j) **cofactor**:

$$\boxed{C_{ij} = (-1)^{i+j} \det M_{ij}} \quad (170)$$

where M_{ij} is the matrix obtained by removing row i and column j from A .

Remark. *The cofactor expansion can also be done along columns instead of rows:*

$$\boxed{\det A = a_{1j} \cdot C_{1j} + \cdots + a_{nj} \cdot C_{nj}} \quad (171)$$

for all $j \in \{1, \dots, n\}$.

Let's do an example, by computing the determinant of the matrix:

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

This matrix has a lot of zeroes, so it's perfect for cofactor expansion. To keep your work as light as possible, pick a row with a lot of zeroes. For example, cofactor expansion for the third row gives:

$$\det A = 2 \cdot (-1)^{3+1} \det \begin{bmatrix} 0 & 3 & -1 \\ 4 & 3 & 0 \\ 2 & 0 & 1 \end{bmatrix} + 1 \cdot (-1)^{3+4} \det \begin{bmatrix} 7 & 0 & 3 \\ 0 & 4 & 3 \\ 0 & 2 & 0 \end{bmatrix}$$

We can compute the two determinants in the right-hand side by cofactor expansion for the first and third rows, respectively:

$$\begin{aligned} \det A &= 2 \cdot (-1)^{3+1} \left(3 \cdot (-1)^{1+2} \det \begin{bmatrix} 4 & 0 \\ 2 & 1 \end{bmatrix} + (-1) \cdot (-1)^{1+3} \det \begin{bmatrix} 4 & 3 \\ 2 & 0 \end{bmatrix} \right) + \\ &+ 1 \cdot (-1)^{3+4} \cdot 2 \cdot (-1)^{3+2} \det \begin{bmatrix} 7 & 3 \\ 0 & 3 \end{bmatrix} = 2((-3) \cdot 4 + (-1) \cdot (-6)) + 2 \cdot 21 = 30 \end{aligned}$$

Just to do a sanity check, let's also compute the determinant of A by cofactor expansion along a column (say the second column) and hope that we get the same answer:

$$\det A = 4 \cdot (-1)^{2+2} \det \begin{bmatrix} 7 & 3 & -1 \\ 2 & 0 & 1 \\ 0 & 0 & 1 \end{bmatrix} + 2 \cdot (-1)^{2+4} \det \begin{bmatrix} 7 & 3 & -1 \\ 0 & 3 & 0 \\ 2 & 0 & 1 \end{bmatrix}$$

Now let's compute the two determinants in the right-hand side by cofactor expansion for the second and third column, respectively:

$$\begin{aligned} \det A &= 4 \cdot (-1)^{2+2} \left(3 \cdot (-1)^{1+2} \det \begin{bmatrix} 2 & 1 \\ 0 & 1 \end{bmatrix} \right) + 2 \cdot (-1)^{2+4} \left((-1) \cdot (-1)^{1+3} \det \begin{bmatrix} 0 & 3 \\ 2 & 0 \end{bmatrix} + \right. \\ &\left. + 1 \cdot (-1)^{3+3} \det \begin{bmatrix} 7 & 3 \\ 0 & 3 \end{bmatrix} \right) = 4 \cdot (-3) \cdot 2 + 2((-1) \cdot (-6) + 1 \cdot 21) = 30 \end{aligned}$$

Lecture 21 (April 9)

Reading: Section 5.3

Viewing: Lecture 20

Remember systems of linear equations like $A\mathbf{v} = \mathbf{b}$? We learned how to solve these through Gaussian elimination, but when A is a square $n \times n$ matrix, there is another way called Cramer's rule. In this setup, you have as many equations as unknowns, and if A is invertible then:

$$A\mathbf{v} = \mathbf{b} \quad \text{has the \underline{unique} solution} \quad \mathbf{v} = A^{-1}\mathbf{b} \quad (172)$$

So the task has become to compute the inverse matrix A^{-1} . But if you look back at formula (169), the inverse matrix is almost staring back at you! Explicitly, consider the matrix:

$$X \quad \text{whose entries are} \quad x_{ij} = C_{ji}$$

and C_{ij} are the cofactors (170) (note the indices above!). Then formula (169) precisely says that:

$$\det A = a_{i1}x_{1i} + \cdots + a_{in}x_{ni}$$

for all $i \in \{1, \dots, n\}$. Together with the fact that for all $i \neq j$:

$$0 = a_{i1}x_{1j} + \cdots + a_{in}x_{nj}$$

(proof: the right-hand side of the formula above is precisely the right-hand side of (169) for the matrix A' obtained from A by copying the i -row into of the j -th row; since A' has two equal rows, we have $\det A' = 0$) we conclude that:

$$\begin{bmatrix} \det A & 0 & \cdots & 0 \\ 0 & \det A & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \det A \end{bmatrix} = AX$$

and therefore:

$$\boxed{A^{-1} = \frac{X}{\det A}} \quad (173)$$

In other words, the (i, j) entry of A^{-1} is given by:

$$\boxed{A_{ij}^{-1} = \frac{C_{ji}}{\det A}} \quad (174)$$

Note that you need to put the (j, i) cofactor in the formula above, not the (i, j) cofactor!

So now let's write a formula for the solution of the equation:

$$A \begin{bmatrix} v_1 \\ \vdots \\ v_n \end{bmatrix} = \begin{bmatrix} b_1 \\ \vdots \\ b_n \end{bmatrix} \quad (175)$$

By (172), we must have:

$$v_i = A_{i1}^{-1} \cdot b_1 + \cdots + A_{in}^{-1} \cdot b_n$$

By formula (174), we therefore have:

$$v_i = \frac{C_{1i} \cdot b_1 + \cdots + C_{ni} \cdot b_n}{\det A}$$

If you stare the right-hand side of the formula above, you will notice that it's precisely the column cofactor expansion (171) of the matrix:

$$B_i \text{ obtained by replacing the } i\text{-th column of } A \text{ by } \mathbf{b} \quad (176)$$

Therefore, we get **Cramer's rule**, which says that the solution to the system (175) has:

$$\boxed{v_1 = \frac{\det B_1}{\det A}, \dots, v_n = \frac{\det B_n}{\det A}} \quad (177)$$

This formula is particularly good when solving small systems of equations, like 3×3 ones:

$$\begin{bmatrix} 1 & -2 & 0 \\ 1 & 0 & -3 \\ 2 & 0 & -5 \end{bmatrix} \begin{bmatrix} v_1 \\ v_2 \\ v_3 \end{bmatrix} = \begin{bmatrix} 0 \\ 4 \\ 1 \end{bmatrix}$$

By (177), the solution is:

$$v_1 = \frac{\det \begin{bmatrix} 0 & -2 & 0 \\ 4 & 0 & -3 \\ 1 & 0 & -5 \end{bmatrix}}{\det \begin{bmatrix} 1 & -2 & 0 \\ 1 & 0 & -3 \\ 2 & 0 & -5 \end{bmatrix}} \quad v_2 = \frac{\det \begin{bmatrix} 1 & 0 & 0 \\ 1 & 4 & -3 \\ 2 & 1 & -5 \end{bmatrix}}{\det \begin{bmatrix} 1 & -2 & 0 \\ 1 & 0 & -3 \\ 2 & 0 & -5 \end{bmatrix}} \quad v_3 = \frac{\det \begin{bmatrix} 1 & -2 & 0 \\ 1 & 0 & 4 \\ 2 & 0 & 1 \end{bmatrix}}{\det \begin{bmatrix} 1 & -2 & 0 \\ 1 & 0 & -3 \\ 2 & 0 & -5 \end{bmatrix}}$$

The 3×3 determinants above can be computed either by cofactor expansion (169), (171) (preferably along a row or column with many zeroes) or even by applying the brute force formula (166). In any case, we obtain:

$$v_1 = -17 \quad v_2 = -\frac{17}{2} \quad v_3 = -7$$

As a final application of determinants, let us discuss about cross products. Remember when I said that it does not make sense to multiply vectors and get a vector, bar a rather particular exception for three dimensional vectors? Well, here is that exception. If you have vectors:

$$\mathbf{v} = \begin{bmatrix} v_1 \\ v_2 \\ v_3 \end{bmatrix} \quad \text{and} \quad \mathbf{w} = \begin{bmatrix} w_1 \\ w_2 \\ w_3 \end{bmatrix}$$

in three-dimensional space, their **cross product** is defined as the following vector in three dimensional space:

$$\mathbf{v} \times \mathbf{w} = \det \begin{bmatrix} \mathbf{i} & v_1 & w_1 \\ \mathbf{j} & v_2 & w_2 \\ \mathbf{k} & v_3 & w_3 \end{bmatrix} \quad (178)$$

where $\mathbf{i} = \mathbf{e}_1, \mathbf{j} = \mathbf{e}_2, \mathbf{k} = \mathbf{e}_3$ are the standard unit vectors on the axes of three-dimensional space. You can compute the determinant in (178) by cofactor expansion on the first column, and obtain:

$$\mathbf{v} \times \mathbf{w} = \mathbf{i} \cdot (v_2 w_3 - v_3 w_2) + \mathbf{j} \cdot (v_3 w_1 - v_1 w_3) + \mathbf{k} \cdot (v_1 w_2 - v_2 w_1) = \begin{bmatrix} v_2 w_3 - v_3 w_2 \\ v_3 w_1 - v_1 w_3 \\ v_1 w_2 - v_2 w_1 \end{bmatrix} \quad (179)$$

The determinant definition of the cross product has a few advantages: it is easy to memorize, and it manifestly explains why the cross product is 0 if the vectors \mathbf{v} and \mathbf{w} are on the same line (because then columns 2 and 3 of the determinant (178) would be multiples of each other). Moreover, because switching two columns flips the sign of the determinant, we conclude that:

$$\mathbf{v} \times \mathbf{w} = -(\mathbf{w} \times \mathbf{v})$$

Moreover, if we have a third vector \mathbf{u} , then we could dot it with the cross product $\mathbf{v} \times \mathbf{w}$ and get a number. This number is actually quite nice, namely:

$$\mathbf{u} \cdot (\mathbf{v} \times \mathbf{w}) = \det \begin{bmatrix} u_1 & v_1 & w_1 \\ u_2 & v_2 & w_2 \\ u_3 & v_3 & w_3 \end{bmatrix} \quad (180)$$

Geometrically, the determinant in the right-hand side of (180) is the volume of the parallelepiped whose edges are given by the vectors \mathbf{u} , \mathbf{v} , \mathbf{w} . If \mathbf{u} lies in the plane spanned by \mathbf{v} and \mathbf{w} , then the volume of this parallelepiped is 0. This is manifestly apparent from (180), since a matrix where one of the columns is a linear combination of the others has determinant 0.

Lecture 22 (April 12)

Reading: Section 6.1

Viewing: Lecture 21

Just like in the last few classes, we're sticking with square matrices, so let A be an $n \times n$ matrix. The goal is to make the matrix "as simple as possible". What this means precisely might be the subject of some debate, but some of the simplest square matrices out there are diagonal ones:

$$\text{diag}_{d_1, \dots, d_n} = \begin{bmatrix} d_1 & 0 & \dots & 0 \\ 0 & d_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & d_n \end{bmatrix} \quad (181)$$

Indeed, the linear transformation $\mathbb{R}^n \rightarrow \mathbb{R}^n$ that corresponds to $\text{diag}_{d_1, \dots, d_n}$ simply scales along the unit vectors $\mathbf{e}_1, \dots, \mathbf{e}_n$ of the coordinate axes (specifically, it scales by a factor of d_i along the i -th coordinate axis of \mathbb{R}^n). As we have seen in (159), we have $\det(\text{diag}_{d_1, \dots, d_n}) = d_1 \dots d_n$, which reflects the fact that our linear transformation scales volumes by a factor equal to $d_1 \dots d_n$.

Definition 18. A matrix is called **diagonalizable** if it is similar to a diagonal matrix, i.e.:

$$A = V \cdot \text{diag}_{d_1, \dots, d_n} \cdot V^{-1} \quad (182)$$

for some invertible $n \times n$ matrix V , and some numbers d_1, \dots, d_n .

The property (182) means that the linear transformation $\mathbb{R}^n \rightarrow \mathbb{R}^n$ that corresponds to A scales by a factor of d_i in the direction of the vector \mathbf{v}_i , where \mathbf{v}_i is the i -th column of V . Therefore, a diagonalizable matrix is diagonal in the basis $\mathbf{v}_1, \dots, \mathbf{v}_n$ of \mathbb{R}^n .

Remark. Almost all $n \times n$ matrices are diagonalizable, and next week we will see what one can say about those matrices which are not.

Suppose you knew that the matrix A is diagonalizable, but you didn't know the specific numbers d_i and the matrix V . You could reconstruct them from the notion of eigenvalues and eigenvectors:

Definition 19. Given a $n \times n$ matrix A , a non-zero vector $\mathbf{v} \in \mathbb{R}^n$ is called an **eigenvector** if:

$$\boxed{A\mathbf{v} = \lambda\mathbf{v}} \quad (183)$$

for some number λ called an **eigenvalue**.

Indeed, \mathbf{v} is an eigenvector of A with corresponding eigenvalue λ if the linear transformation corresponding to A scales by a factor of λ in the direction of the vector \mathbf{v} . This story is very closely connected with (182): if the matrix A is equal to $V \text{diag}_{d_1, \dots, d_n} V^{-1}$, then all the d_i 's are eigenvalues of A and the corresponding eigenvectors are the columns of V .

So how many eigenvalues/eigenvectors are there? It's easy to see that if \mathbf{v} is an eigenvector, then so is any multiple of \mathbf{v} . Moreover, if \mathbf{v} and \mathbf{w} are both eigenvectors corresponding to the same eigenvalue

λ , then $\mathbf{v} + \mathbf{w}$ is also an eigenvector corresponding to λ . Therefore, a better question would be “how many linearly independent eigenvectors are there”? Well, from the equation (183) we obtain:

$$(A - \lambda I)\mathbf{v} = 0$$

Since \mathbf{v} is supposed to be non-zero, this implies that the matrix $A - \lambda I$ should be singular (i.e. not invertible). Therefore, Fact 15 implies that:

$$\boxed{\det(A - \lambda I) = 0} \tag{184}$$

for any eigenvalue λ of A . Conversely, any number λ for which (184) holds is an eigenvalue of A (try to think of an argument). This leads us to the following:

Definition 20. *The characteristic polynomial of A is:*

$$p(\lambda) = \det(A - \lambda I) \tag{185}$$

as a function of the variable λ . The roots (i.e. those numbers λ for which $p(\lambda) = 0$) of the characteristic polynomial are the eigenvalues of A .

Although not very practical, you could in theory compute the characteristic polynomial by applying the big formula (167) for the determinant:

$$p(\lambda) = \det \begin{bmatrix} a_{11} - \lambda & a_{12} & \dots & a_{1n} \\ a_{21} & a_{22} - \lambda & \dots & a_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ a_{n1} & a_{n2} & \dots & a_{nn} - \lambda \end{bmatrix} = \sum_{\{\sigma(1), \dots, \sigma(n)\}}^{\text{permutations}} (-1)^{\text{sgn } \sigma} a'_{1\sigma(1)} a'_{2\sigma(2)} \dots a'_{n\sigma(n)}$$

where $a'_{ij} = a_{ij}$ if $i \neq j$ and $a'_{ii} = a_{ii} - \lambda$. Since every summand in the right-hand side is a product of n numbers a'_{ij} for various i and j , no summand can contain more than n factors of λ . Moreover, the only summand which contains exactly n factors of λ is $a'_{11} a'_{22} \dots a'_{nn} = (a_{11} - \lambda)(a_{22} - \lambda) \dots (a_{nn} - \lambda)$. Therefore, the characteristic polynomial has degree n in λ and its top degree coefficient is $(-1)^n$:

$$p(\lambda) = (-1)^n \lambda^n + \alpha_{n-1} \lambda^{n-1} + \dots + \alpha_1 \lambda + \alpha_0$$

The coefficients $\alpha_{n-1}, \dots, \alpha_0$ depend on the entries of the matrix A . For example:

$$\alpha_{n-1} = (-1)^{n-1} \cdot \text{tr } A \tag{186}$$

where:

$$\boxed{\text{tr } A = a_{11} + a_{22} + \dots + a_{n-1, n-1} + a_{nn}} \tag{187}$$

is called the **trace** of the matrix A . Clearly:

$$\alpha_0 = \det A \tag{188}$$

is just the determinant of A (this is just saying that the value of the characteristic polynomial at $\lambda = 0$ is the determinant of A , which is obvious given the definition (185)). Because of Vieta's formulas, which express the sum/product of the roots of a polynomial in terms of its coefficients, we conclude the following:

Fact 17. *The sum of the eigenvalues of a matrix A is equal to the trace of A .*

The product of the eigenvalues of a matrix A is the determinant of A .

As an exercise, compute the eigenvalues of the matrix:

$$A = \begin{bmatrix} 2 & 1 \\ 5 & 2 \end{bmatrix} \quad (189)$$

First let us form the characteristic polynomial:

$$p(\lambda) = \det(A - \lambda I) = \det \begin{bmatrix} 2 - \lambda & 1 \\ 5 & 2 - \lambda \end{bmatrix} = (2 - \lambda)(2 - \lambda) - 5 \cdot 1 = \lambda^2 - 4\lambda - 1$$

The eigenvalues of A are the roots of this polynomial, which you can get by the quadratic formula:

$$\lambda = \frac{4 \pm \sqrt{4^2 + 4}}{2} = 2 \pm \sqrt{5} \quad \Rightarrow \quad \begin{cases} \lambda_1 = 2 + \sqrt{5} \\ \lambda_2 = 2 - \sqrt{5} \end{cases}$$

In general, because the eigenvalues of a 2×2 matrix are the solutions of a quadratic equation, you should expect to encounter square roots in their formulas. As we have seen in the example above, in the 2×2 case, one can deduce the formula for the characteristic polynomial from the facts that:

- it is quadratic and its leading coefficient is 1
- its linear term is minus the trace of the matrix (186)
- its constant term is the determinant of a matrix (188)

So the characteristic polynomial of the matrix $\begin{bmatrix} a & b \\ c & d \end{bmatrix}$ is:

$$p(\lambda) = \lambda^2 - \lambda(a + d) + (ad - bc) \quad (190)$$

For matrices of size 3×3 and higher, the characteristic polynomial is of order 3 and higher, so it is a bit more complicated to write it down and a lot more complicated to find its roots.

Remark. *In general, it is **NOT** true that the eigenvalues of $A + B$ are the sums of the eigenvalues of A and B , nor are the eigenvalues of AB the products of the eigenvalues of A and B .*

Now suppose that you know λ is an eigenvalue of a square matrix A ; how do you find an eigenvector \mathbf{v} such that $A\mathbf{v} = \lambda\mathbf{v}$? You already have the tools for this: just recast the eigenvector relation as $(A - \lambda I)\mathbf{v} = \mathbf{0}$, so all you need to do is to find $\mathbf{v} \neq \mathbf{0}$ in the nullspace of the singular matrix $A - \lambda I$.

For example, in the case of the matrix (189), let us look for an eigenvector \mathbf{v} corresponding to the eigenvalue $\lambda = 2 + \sqrt{5}$. To this end, let us construct the matrix:

$$A - \lambda I = \begin{bmatrix} -\sqrt{5} & 1 \\ 5 & -\sqrt{5} \end{bmatrix}$$

and we will simply pick a vector $\mathbf{v} \in N(A - \lambda I)$. Remember that, in order to compute the nullspace of a matrix, the way to go is to put it in reduced row echelon form:

$$\begin{bmatrix} -\sqrt{5} & 1 \\ 5 & -\sqrt{5} \end{bmatrix} \rightsquigarrow \begin{bmatrix} -\sqrt{5} & 1 \\ 0 & 0 \end{bmatrix} \rightsquigarrow \begin{bmatrix} 1 & -\frac{1}{\sqrt{5}} \\ 0 & 0 \end{bmatrix}$$

Therefore, an eigenvector $\mathbf{v} = \begin{bmatrix} v_1 \\ v_2 \end{bmatrix}$ can be obtained by solving the equation:

$$\begin{bmatrix} 1 & -\frac{1}{\sqrt{5}} \\ 0 & 0 \end{bmatrix} \begin{bmatrix} v_1 \\ v_2 \end{bmatrix} = 0 \quad \Leftrightarrow \quad v_1 - \frac{v_2}{\sqrt{5}} = 0$$

So one particular choice of eigenvector can be obtained by setting the free variable v_2 equal to 1, and solving for the pivot variable from the equation above: $v_1 = \frac{1}{\sqrt{5}}$. Therefore, an eigenvector of the matrix A corresponding to the eigenvalue $\lambda = 2 + \sqrt{5}$ is:

$$\mathbf{v} = \begin{bmatrix} \frac{1}{\sqrt{5}} \\ 1 \end{bmatrix}$$

Any multiple of \mathbf{v} would be an equally good eigenvector.

Lecture 23 (April 14)

Reading: Section 6.2

Viewing: Lecture 21

In the previous class, we saw how to compute the eigenvalues and eigenvectors of an $n \times n$ matrix A . Specifically, the eigenvalues are the roots of the characteristic polynomial:

$$p(\lambda) = \det(A - \lambda I) = 0$$

Since the characteristic polynomial has degree n , it has at most n roots, therefore there are at most n eigenvalues. But there could be fewer than n eigenvalues, for example the matrix:

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

has characteristic polynomial

$$p(\lambda) = (\lambda - 5)^2$$

which has $\lambda = 5$ as a double root. In this case, we say that $\lambda = 5$ has **algebraic multiplicity 2**, or that the matrix A has repeated (or equal) eigenvalues. In other words, we still think of A as having two eigenvalues, but they are both equal to 5. With this in mind, we have:

Fact 18. *An $n \times n$ matrix has exactly n eigenvalues, although some could be repeated, and some could be complex numbers (more on complex numbers later).*

So what can we say about eigenvectors? First of all, eigenvectors corresponding to different eigenvalues are always linearly independent. For example, if:

$$A\mathbf{v}_1 = \lambda_1\mathbf{v}_1 \quad \text{and} \quad A\mathbf{v}_2 = \lambda_2\mathbf{v}_2$$

for $\lambda_1 \neq \lambda_2$, then there can be no linear relation between \mathbf{v}_1 and \mathbf{v}_2 (argument: if there were such a relation, then we would have $\mathbf{v}_1 = c\mathbf{v}_2$ for some scalar c , and applying A would give us $\lambda_1\mathbf{v}_1 = A\mathbf{v}_1 = cA\mathbf{v}_2 = c\lambda_2\mathbf{v}_2 = \lambda_2\mathbf{v}_1$, which would force $\lambda_1 = \lambda_2$). So the best case is when the n eigenvalues of the matrix are all distinct numbers, because then there would exist n linearly independent eigenvectors, one for each eigenvalue. These eigenvectors would form a basis of \mathbb{R}^n .

Fact 19. *An $n \times n$ matrix A with n distinct eigenvalues $\lambda_1, \dots, \lambda_n$ is diagonalizable. Specifically:*

$$A = V \begin{bmatrix} \lambda_1 & 0 & \dots & 0 \\ 0 & \lambda_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \lambda_n \end{bmatrix} V^{-1} \tag{191}$$

where V is the matrix whose columns are $\mathbf{v}_1, \dots, \mathbf{v}_n$, the eigenvectors corresponding to $\lambda_1, \dots, \lambda_n$.

A couple of remarks: if the eigenvalues are all distinct, then the eigenvectors are really only defined up to constant multiple. So the matrix V in (191) is not unique, but it depends on a choice of constant multiples of the eigenvectors. Secondly, the fact that the eigenvectors are linearly independent

(which was discussed right before Fact 19) is crucial, because otherwise the matrix V would have rank $< n$ and would not be invertible and the right-hand side of formula (191) would not make sense.

In practice, how to diagonalize (i.e. to write as in (191)) a matrix A ? As explained in Fact 19, the way to do so is to find the eigenvalues and eigenvectors of A . Say we're looking at the matrix:

$$A = \begin{bmatrix} 1 & 1 \\ 1 & 0 \end{bmatrix} \quad (192)$$

To find its eigenvalues, construct the characteristic polynomial:

$$p(\lambda) = \lambda^2 - \lambda \cdot \text{tr } A + \det A = \lambda^2 - \lambda - 1$$

(here we're applying the shortcut formula (190), which only holds in the 2×2 case). Its roots are:

$$\lambda_1 = \frac{1 + \sqrt{5}}{2} \quad \text{and} \quad \lambda_2 = \frac{1 - \sqrt{5}}{2}$$

Let us now compute eigenvectors corresponding to these two eigenvalues:

$$\begin{aligned} A\mathbf{v}_1 = \frac{1 + \sqrt{5}}{2}\mathbf{v}_1 &\Rightarrow \left(A - \frac{1 + \sqrt{5}}{2}I \right) \mathbf{v}_1 = 0 \Rightarrow \mathbf{v}_1 \in N \left(A - \frac{1 + \sqrt{5}}{2}I \right) \\ A\mathbf{v}_2 = \frac{1 - \sqrt{5}}{2}\mathbf{v}_2 &\Rightarrow \left(A - \frac{1 - \sqrt{5}}{2}I \right) \mathbf{v}_2 = 0 \Rightarrow \mathbf{v}_2 \in N \left(A - \frac{1 - \sqrt{5}}{2}I \right) \end{aligned}$$

To work out the nullspace of $A - \frac{1 + \sqrt{5}}{2}I$, put it in row echelon form:

$$A - \frac{1 + \sqrt{5}}{2}I = \begin{bmatrix} \frac{1 - \sqrt{5}}{2} & 1 \\ 1 & \frac{-1 - \sqrt{5}}{2} \end{bmatrix} \rightsquigarrow \begin{bmatrix} \frac{1 - \sqrt{5}}{2} & 1 \\ 0 & 0 \end{bmatrix}$$

So the eigenvector $\mathbf{v}_1 = \begin{bmatrix} x \\ y \end{bmatrix}$ needs to satisfy:

$$\begin{bmatrix} \frac{1 - \sqrt{5}}{2} & 1 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = 0 \Leftrightarrow y + \frac{1 - \sqrt{5}}{2}x = 0$$

and therefore we may take $\mathbf{v}_1 = \begin{bmatrix} 1 \\ \frac{-1 + \sqrt{5}}{2} \end{bmatrix}$. Similarly, $\mathbf{v}_2 = \begin{bmatrix} 1 \\ \frac{-1 - \sqrt{5}}{2} \end{bmatrix}$, so we have:

$$A = V \begin{bmatrix} \frac{1 + \sqrt{5}}{2} & 0 \\ 0 & \frac{1 - \sqrt{5}}{2} \end{bmatrix} V^{-1} \quad \text{where} \quad V = \begin{bmatrix} 1 & 1 \\ \frac{-1 + \sqrt{5}}{2} & \frac{-1 - \sqrt{5}}{2} \end{bmatrix} \quad (193)$$

So we've done all this work, but you might ask: what are applications of diagonalization? Well, here's a classic one. Let's say you want to obtain an explicit formula for the Lucas numbers (close cousins of the Fibonacci numbers considered in the textbook), which are defined by the conditions:

$$L_0 = 2, \quad L_1 = 1, \quad \text{and the recursive relation} \quad L_{n+1} = L_n + L_{n-1} \quad (194)$$

for all $n \geq 1$. One way would be to consider the vectors:

$$\mathbf{a}_n = \begin{bmatrix} L_{n+1} \\ L_n \end{bmatrix}$$

and to observe that the conditions (194) can be interpreted as the matrix identities:

$$\mathbf{a}_0 = \begin{bmatrix} 1 \\ 2 \end{bmatrix} \quad \text{and} \quad \mathbf{a}_n = \begin{bmatrix} 1 & 1 \\ 1 & 0 \end{bmatrix} \mathbf{a}_{n-1}$$

for all $n \geq 1$. You can iterate the latter formula and obtain:

$$\mathbf{a}_n = A\mathbf{a}_{n-1} = A^2\mathbf{a}_{n-2} = \cdots = A^n\mathbf{a}_0 = A^n \begin{bmatrix} 1 \\ 2 \end{bmatrix} \quad (195)$$

where A is the matrix (192). So computing a formula for the Lucas numbers boils down to having an effective method for explicitly computing the matrix powers A^n for all natural numbers n . This is quite straightforward if you know how to diagonalize A . Specifically, because A can be written in the form (193), we have:

$$\begin{aligned} A^2 &= V \begin{bmatrix} \frac{1+\sqrt{5}}{2} & 0 \\ 0 & \frac{1-\sqrt{5}}{2} \end{bmatrix} \underbrace{V^{-1}V}_I \begin{bmatrix} \frac{1+\sqrt{5}}{2} & 0 \\ 0 & \frac{1-\sqrt{5}}{2} \end{bmatrix} V^{-1} = V \begin{bmatrix} \left(\frac{1+\sqrt{5}}{2}\right)^2 & 0 \\ 0 & \left(\frac{1-\sqrt{5}}{2}\right)^2 \end{bmatrix} V^{-1} \\ A^3 &= V \begin{bmatrix} \frac{1+\sqrt{5}}{2} & 0 \\ 0 & \frac{1-\sqrt{5}}{2} \end{bmatrix} \underbrace{V^{-1}V}_I \begin{bmatrix} \frac{1+\sqrt{5}}{2} & 0 \\ 0 & \frac{1-\sqrt{5}}{2} \end{bmatrix} \underbrace{V^{-1}V}_I \begin{bmatrix} \frac{1+\sqrt{5}}{2} & 0 \\ 0 & \frac{1-\sqrt{5}}{2} \end{bmatrix} V^{-1} = V \begin{bmatrix} \left(\frac{1+\sqrt{5}}{2}\right)^3 & 0 \\ 0 & \left(\frac{1-\sqrt{5}}{2}\right)^3 \end{bmatrix} V^{-1} \\ \dots & \\ A^n &= V \begin{bmatrix} \left(\frac{1+\sqrt{5}}{2}\right)^n & 0 \\ 0 & \left(\frac{1-\sqrt{5}}{2}\right)^n \end{bmatrix} V^{-1} \end{aligned}$$

for all integers n . We could plug these formulas in (195) and obtain:

$$\begin{aligned} \begin{bmatrix} L_{n+1} \\ L_n \end{bmatrix} &= \mathbf{a}_n = V \begin{bmatrix} \left(\frac{1+\sqrt{5}}{2}\right)^n & 0 \\ 0 & \left(\frac{1-\sqrt{5}}{2}\right)^n \end{bmatrix} V^{-1} \begin{bmatrix} 1 \\ 2 \end{bmatrix} = \\ &= \begin{bmatrix} 1 & 1 \\ \frac{-1+\sqrt{5}}{2} & \frac{-1-\sqrt{5}}{2} \end{bmatrix} \begin{bmatrix} \left(\frac{1+\sqrt{5}}{2}\right)^n & 0 \\ 0 & \left(\frac{1-\sqrt{5}}{2}\right)^n \end{bmatrix} \frac{1}{-\sqrt{5}} \begin{bmatrix} \frac{-1-\sqrt{5}}{2} & -1 \\ \frac{1-\sqrt{5}}{2} & 1 \end{bmatrix} \begin{bmatrix} 1 \\ 2 \end{bmatrix} \end{aligned}$$

I'll let you have all the fun in multiplying the matrices on the bottom row above. When you do so, you will finally obtain the formula for the Lucas numbers:

$$L_n = \left(\frac{1+\sqrt{5}}{2}\right)^n + \left(\frac{1-\sqrt{5}}{2}\right)^n \quad (196)$$

The example we have just seen generalizes to the following principle:

$$\text{if } A = V \begin{bmatrix} \lambda_1 & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & \lambda_n \end{bmatrix} V^{-1} \quad \text{then} \quad A^k = V \begin{bmatrix} \lambda_1^k & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & \lambda_n^k \end{bmatrix} V^{-1} \quad (197)$$

This is computationally very important, because if you want to compute the k -th power of an $n \times n$ matrix efficiently for some really large number k , it is faster to just diagonalize A and apply formula (197). Also, this formula allows us to describe the behavior of A^k as $k \rightarrow \infty$: its growth rate is dominated by the size of the k -th power of the largest eigenvalue of A .

Remark. The Lucas numbers L_n grow like $\left(\frac{1+\sqrt{5}}{2}\right)^n$ as $n \rightarrow \infty$. This makes sense, since the other eigenvalue $\frac{1-\sqrt{5}}{2}$ lies between -1 and 1 , and so its n -th powers will converge to 0 as $n \rightarrow \infty$.

Geometrically, suppose you want to compute the value of A^k applied to any vector \mathbf{v} ? Just decompose this vector as a linear combination of the eigenvectors:

$$\mathbf{v} = c_1 \mathbf{v}_1 + \cdots + c_n \mathbf{v}_n \quad \Rightarrow \quad A^k \mathbf{v} = c_1 A^k \mathbf{v}_1 + \cdots + c_n A^k \mathbf{v}_n$$

Since $A \mathbf{v}_i = \lambda_i \mathbf{v}_i \Rightarrow A^k \mathbf{v}_i = \lambda_i^k \mathbf{v}_i$ for all i (on account of $\mathbf{v}_1, \dots, \mathbf{v}_n$ being eigenvectors), we get:

$$\boxed{A^k \mathbf{v} = c_1 \lambda_1^k \mathbf{v}_1 + \cdots + c_n \lambda_n^k \mathbf{v}_n} \quad (198)$$

We conclude that the powers of A scale by powers of the eigenvalues $\lambda_1, \dots, \lambda_n$ in the direction of the eigenvectors $\mathbf{v}_1, \dots, \mathbf{v}_n$, and this gives an effective way to compute A^k times any vector.

Finally, let us mention an easy, but fundamental thing: **the eigenvalues of a diagonal matrix are the entries on the diagonal.** This is because the roots of the polynomial:

$$\det \left(\begin{bmatrix} d_1 & 0 & \cdots & 0 \\ 0 & d_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & d_n \end{bmatrix} - \lambda \cdot I \right) = \det \begin{bmatrix} d_1 - \lambda & 0 & \cdots & 0 \\ 0 & d_2 - \lambda & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & d_n - \lambda \end{bmatrix} = (d_1 - \lambda) \cdots (d_n - \lambda)$$

are precisely d_1, \dots, d_n . Applying this to (191), we see that the matrices A and $\text{diag}_{\lambda_1, \dots, \lambda_n}$ have the same eigenvalues. This is true for any pair of similar matrices:

$$\text{if } A = VB V^{-1} \text{ then } A \text{ and } B \text{ have the same eigenvalues} \quad (199)$$

for any invertible matrix V . However, A and B will in general have different eigenvectors, and these two bases of eigenvectors will be connected by the change of basis matrix V in formula (199):

$$V \mathbf{w} \text{ is an eigenvector of } A \quad \Leftrightarrow \quad \mathbf{w} \text{ is an eigenvector of } B$$

Lecture 24 (April 16)

Reading: Section 8.3

Viewing: Lecture 28 (after minute 12:00)

The previous class was all about the situation when an $n \times n$ matrix A has distinct eigenvalues. But what if the eigenvalues are repeated, i.e. what if the characteristic polynomial takes the form:

$$p(\lambda) = (d_1 - \lambda)^{r_1} (d_2 - \lambda)^{r_2} \dots (d_s - \lambda)^{r_s} \quad (200)$$

for certain distinct numbers d_1, \dots, d_s , and for certain powers r_1, \dots, r_s ? We think of the eigenvalues of A as being the following sequence of numbers:

$$\underbrace{d_1, \dots, d_1}_{r_1 \text{ times}}, \underbrace{d_2, \dots, d_2}_{r_2 \text{ times}}, \dots, \underbrace{d_s, \dots, d_s}_{r_s \text{ times}} \quad (201)$$

The number r_i is called the **algebraic multiplicity** of the eigenvalue d_i , and it represents the number of times it appears among the roots of the characteristic polynomial. Because the degree of the characteristic polynomial is n , the sum of the algebraic multiplicities of all eigenvalues is:

$$r_1 + \dots + r_s = n \quad (202)$$

What about eigenvectors for each eigenvalue d_i ? The set of eigenvectors for any given eigenvalue is a vector subspace of \mathbb{R}^n , namely:

$$\left\{ \text{eigenvectors of } A \text{ corresponding to eigenvalue } \lambda \right\} = N(A - \lambda I)$$

The **geometric multiplicity** of any eigenvalue d_i is defined as the dimension of the corresponding space of eigenvectors, namely $\dim N(A - d_i I)$. It never exceeds the algebraic multiplicity:

$$\boxed{\text{geometric multiplicity} \leq \text{algebraic multiplicity}} \quad (203)$$

For example, consider the matrices below:

$$A = \begin{bmatrix} d & 0 \\ 0 & d \end{bmatrix} \quad \text{and} \quad B = \begin{bmatrix} d & 1 \\ 0 & d \end{bmatrix} \quad (204)$$

Both of them have characteristic polynomial $(\lambda - d)^2$, so d is the only eigenvalue, with algebraic multiplicity 2. But the geometric multiplicity differs among the two matrices. In the case of the matrix A , any vector is an eigenvector because $A\mathbf{v} = d\mathbf{v}$ for all \mathbf{v} , thus implying that the eigenvalue d has geometric multiplicity 2. But for the matrix B , we have that:

$$\mathbf{v} = \begin{bmatrix} x \\ y \end{bmatrix} \text{ is an eigenvector} \Leftrightarrow \begin{bmatrix} d & 1 \\ 0 & d \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = d \begin{bmatrix} x \\ y \end{bmatrix} \Leftrightarrow y = 0$$

which implies that the subspace of eigenvectors is one-dimensional, hence the eigenvalue d has geometric multiplicity 1. This is strictly smaller than its algebraic multiplicity, which is 2.

Fact 20. *A matrix is diagonalizable if and only if all of its eigenvalues have geometric multiplicity equal to their algebraic multiplicity.*

Indeed, if the geometric multiplicities of all eigenvalues are equal to the corresponding algebraic multiplicities, then (202) implies that the sum of all geometric multiplicities is equal to n . This implies that the subspaces of eigenvectors have dimensions which add up to n , hence we can pick a basis of \mathbb{R}^n consisting only of eigenvectors. This allows us to define the invertible matrix V whose columns are the n chosen eigenvectors, and formula (191) holds.

But what can we say when some of the eigenvalues have geometric multiplicity strictly smaller than their algebraic multiplicity? In this case, the matrix A is no longer similar to a diagonal matrix, as in (191). Instead, it will be similar to a matrix in **Jordan normal form**:

$$A = V \left[\begin{array}{c|c|c|c} J_1 & 0 & \dots & 0 \\ \hline 0 & J_2 & \dots & 0 \\ \hline \vdots & \vdots & \ddots & \vdots \\ \hline 0 & 0 & \dots & J_t \end{array} \right] V^{-1} \quad (205)$$

where each J_i is a **Jordan block**, i.e. a matrix of the form:

$$\begin{bmatrix} \lambda & 1 & 0 & 0 & \dots & 0 \\ 0 & \lambda & 1 & 0 & \dots & 0 \\ 0 & 0 & \lambda & 1 & \dots & 0 \\ 0 & 0 & 0 & \lambda & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & 0 & \dots & \lambda \end{bmatrix} \quad (206)$$

where λ is among the eigenvalues of A . The total number of d_i 's that appear on the diagonal of the block matrix in (205) is equal to r_i , the algebraic multiplicity of the eigenvalue d_i .

There are two things one needs to know in order to set up the Jordan normal form (205): the specific sizes of the Jordan blocks J_1, \dots, J_t and the matrix V . There is a general prescription for figuring out the sizes of the Jordan blocks, but it goes beyond the scope of our course. As for the matrix V , we know from the diagonalizable case that its columns have something to do with the nullspaces $N(A - d_1 I), \dots, N(A - d_s I)$. However, in the situation at hand (when the geometric multiplicities are smaller than the algebraic multiplicities), the dimensions of these nullspaces do not add up to n , so they will not contain n linearly independent vectors. So instead, look at:

$$N\left((A - d_1 I)^{r_1}\right), \dots, N\left((A - d_s I)^{r_s}\right) \quad (207)$$

(where r_1, \dots, r_s are the algebraic multiplicities of the eigenvalues) and we claim that these subspaces are both linearly independent, and span the entire \mathbb{R}^n . Then the subspaces (207) will contain n linearly independent vectors, and these will be the columns of the matrix V of (205).

Let's present this in more detail in the case when the algebraic multiplicities are all either 1 or 2. Consider the matrix:

$$A = \begin{bmatrix} 10 & -3 & 7 \\ 27 & -10 & 18 \\ 5 & -3 & 3 \end{bmatrix}$$

and we want to compute its Jordan normal form: i.e. a matrix V and Jordan blocks J_1, \dots, J_t such

that formula (205) holds. First of all, let's compute the characteristic polynomial:

$$p(\lambda) = \det(A - \lambda I) = \det \begin{bmatrix} 10 - \lambda & -3 & 7 \\ 27 & -10 - \lambda & 18 \\ 5 & -3 & 3 - \lambda \end{bmatrix}$$

Let us compute the determinant by cofactor expansion along the last row:

$$\begin{aligned} p(\lambda) &= 5 \det \begin{bmatrix} -3 & 7 \\ -10 - \lambda & 18 \end{bmatrix} - (-3) \det \begin{bmatrix} 10 - \lambda & 7 \\ 27 & 18 \end{bmatrix} + (3 - \lambda) \det \begin{bmatrix} 10 - \lambda & -3 \\ 27 & -10 - \lambda \end{bmatrix} = \\ &= 5(-54 + 70 + 7\lambda) + 3(180 - 18\lambda - 189) + (3 - \lambda)(\lambda^2 - 100 + 81) = -\lambda^3 + 3\lambda^2 - 4 \end{aligned}$$

It's not easy to find the roots of a general cubic polynomial. But when you're out of ideas, just try plugging in some simple numbers for λ to see if $p(\lambda) = 0$. For example, trying $\lambda = -1$ shows that $p(-1) = 0$, which implies that $p(\lambda)$ is divisible by $\lambda + 1$. Then polynomial long division shows that:

$$p(\lambda) = (\lambda + 1)(-\lambda^2 + 4\lambda - 4) = -(\lambda + 1)(\lambda - 2)^2$$

So the eigenvalues are $d_1 = -1$ with algebraic multiplicity one, and $d_2 = 2$ with algebraic multiplicity two. Now for the subspaces of eigenvectors. We have:

$$A + I = \begin{bmatrix} 11 & -3 & 7 \\ 27 & -9 & 18 \\ 5 & -3 & 4 \end{bmatrix} \xrightarrow{\text{RREF}} \begin{bmatrix} 1 & 0 & \frac{1}{2} \\ 0 & 1 & -\frac{1}{2} \\ 0 & 0 & 0 \end{bmatrix} \quad \text{hence} \quad N(A + I) = N \left(\begin{bmatrix} 1 & 0 & \frac{1}{2} \\ 0 & 1 & -\frac{1}{2} \\ 0 & 0 & 0 \end{bmatrix} \right)$$

The matrix on the right has pivot columns 1 and 2, and free column 3, so its null-space is one dimensional and spanned by the vector:

$$\begin{bmatrix} x \\ y \\ 1 \end{bmatrix} \quad \text{such that} \quad \begin{bmatrix} 1 & 0 & \frac{1}{2} \\ 0 & 1 & -\frac{1}{2} \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = 0 \quad \Rightarrow \quad x = -\frac{1}{2} \quad \text{and} \quad y = \frac{1}{2}$$

Therefore, the eigenvalue $d_1 = -1$ also has geometric multiplicity one, and an eigenvector is:

$$\mathbf{v}_1 = \begin{bmatrix} -\frac{1}{2} \\ \frac{1}{2} \\ 1 \end{bmatrix}$$

As for the other eigenvalue, we have:

$$A - 2I = \begin{bmatrix} 8 & -3 & 7 \\ 27 & -12 & 18 \\ 5 & -3 & 1 \end{bmatrix} \xrightarrow{\text{RREF}} \begin{bmatrix} 1 & 0 & 2 \\ 0 & 1 & 3 \\ 0 & 0 & 0 \end{bmatrix} \quad \text{hence} \quad N(A - 2I) = N \left(\begin{bmatrix} 1 & 0 & 2 \\ 0 & 1 & 3 \\ 0 & 0 & 0 \end{bmatrix} \right)$$

The matrix on the right has pivot columns 1 and 2, and free column 3, so its null-space is one dimensional and spanned by the vector:

$$\begin{bmatrix} x \\ y \\ 1 \end{bmatrix} \quad \text{such that} \quad \begin{bmatrix} 1 & 0 & 2 \\ 0 & 1 & 3 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = 0 \quad \Rightarrow \quad x = -2 \quad \text{and} \quad y = -3$$

Therefore, the eigenvalue $d_2 = 2$ has geometric multiplicity one, and an eigenvector is:

$$\mathbf{v}_2 = \begin{bmatrix} -2 \\ -3 \\ 1 \end{bmatrix}$$

So we are in the situation when geometric multiplicity is strictly smaller than algebraic multiplicity, hence the subspaces of eigenvectors are not enough to produce a basis of \mathbb{R}^3 . In this case, (207) tells us to replace the nullspace of $A - 2I$ by the nullspace of $(A - 2I)^2$, namely:

$$(A - 2I)^2 = \begin{bmatrix} 18 & -9 & 9 \\ -18 & 9 & -9 \\ -36 & 18 & -18 \end{bmatrix} \xrightarrow{\text{RREF}} \begin{bmatrix} 1 & -\frac{1}{2} & \frac{1}{2} \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \quad \text{hence} \quad N((A - 2I)^2) = N\left(\begin{bmatrix} 1 & -\frac{1}{2} & \frac{1}{2} \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}\right)$$

The matrix on the right has pivot column 1 and free columns 2 and 3, so its nullspace is two dimensional and consists of vectors:

$$\begin{bmatrix} x \\ y \\ z \end{bmatrix} \quad \text{such that} \quad 2x - y + z = 0$$

One of these vectors is \mathbf{v}_2 , but we can take any other such vector, let's say $\mathbf{v}_3 = \begin{bmatrix} a \\ 2a \\ 0 \end{bmatrix}$ (there will be a benefit in making the last entry 0, as we will soon see) for any a . Therefore, the matrix:

$$V = [\mathbf{v}_1 \mid \mathbf{v}_2 \mid \mathbf{v}_3] = \begin{bmatrix} -\frac{1}{2} & -2 & a \\ \frac{1}{2} & -3 & 2a \\ 1 & 1 & 0 \end{bmatrix}$$

is invertible. The constant $a \neq 0$ is yours to choose freely, but if you want to have the equality:

$$A = V \left[\begin{array}{c|cc} -1 & 0 & 0 \\ \hline 0 & 2 & 1 \\ 0 & 0 & 2 \end{array} \right] V^{-1}$$

then you need to choose $a = -1$ (if you were to use any other value for a , then you would just need to replace the entry 1 in the 2×2 Jordan block above by the entry $-a$, which is not too bad).

Remark. As I mentioned, I did not give you a complete recipe for how to set up the Jordan blocks in (205). But as you were able to see from the example above, if the Jordan blocks corresponding to an eigenvalue λ are located on columns i, \dots, j , then the columns i, \dots, j of the matrix V need to be filled with a basis of $N((A - \lambda I)^r)$, where r is the algebraic multiplicity of the eigenvalue λ .

Lecture 25 (April 21)

Reading: Section 9.1

Viewing: none

Consider the problem of diagonalizing the matrix:

$$A = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix} \quad (208)$$

As we have seen in Fact 19, the recipe for doing so is to find the eigenvalues and eigenvectors of A first. Now, the eigenvalues are the roots of the characteristic polynomial:

$$p(\lambda) = \det \begin{bmatrix} -\lambda & -1 \\ 1 & -\lambda \end{bmatrix} = \lambda^2 + 1 \quad (209)$$

You may be tempted to believe that this polynomial has no roots, because after all, no real number has the property that its square is -1 . And your guess would seem to be justified: the linear transformation corresponding to the matrix (208) is rotation by 90 degrees of the plane. If the matrix had an eigenvector, then rotation by 90 degrees would have to scale along the direction of that eigenvector, which does not seem to make any sense geometrically.

However, in math, a paradox is rarely the end of the story, but more often it is the catalyst for new discoveries. In the example at hand, the way to fix the fact that the polynomial (209) has no solutions is to add these solutions by hand. What this means is that we are no longer satisfied with real numbers, and instead we want to consider the **imaginary number**:

$$i \text{ defined such that } \boxed{i^2 = -1} \quad (210)$$

which is often abbreviated as saying $i = \sqrt{-1}$. You should read this as saying that i is defined to be the square root of -1 , and it certainly is not a real number. But it is a very useful tool for defining **complex numbers**, which are by definition all expressions of the form:

$$z = a + bi \quad (211)$$

for any real numbers a and b . You may ask: why don't we also throw in various powers of i into the expression (211) while we're at it? The reason is that doing so would be redundant, since the powers of i can themselves be expressed in the form (211):

$$i^2 = -1 \quad i^3 = -i \quad i^4 = 1 \quad i^5 = i \quad i^6 = -1 \quad \dots \quad (212)$$

Remark. *I hope you won't be put off by the feeling that complex numbers are "contrived". In the past, people have argued for hundreds of years whether negative integers make any sense, and ultimately we accepted them. And if you think about it, what is the meaning of " -2 " other than "the thing to which you add 2 to get 0" and what is the meaning of " $\sqrt{3}$ " other than "the thing which you square to get 3". Similarly, " i " is defined as "the thing which you square to get -1 ".*

Complex numbers can be added and multiplied using the usual algebra operations:

$$(a + ib) + (c + id) = (a + c) + i(b + d) \quad (213)$$

$$(a + ib)(c + id) = ac + adi + bci + bdi^2 = (ac - bd) + (ad + bc)i \quad (214)$$

so the sum and product of complex numbers is also a complex number. But complex numbers have other kinds of operations associated to them. If $z = a + bi$ is a complex number, then:

$$\operatorname{Re} z = a \text{ is called the } \mathbf{real\ part} \text{ of } z \quad (215)$$

$$\operatorname{Im} z = b \text{ is called the } \mathbf{imaginary\ part} \text{ of } z \quad (216)$$

$$\bar{z} = a - bi \text{ is called the } \mathbf{conjugate} \text{ of } z \quad (217)$$

Therefore, the conjugate of a complex number is that complex number with the same real part, but the opposite imaginary part. Moreover, define the **absolute value** of a complex number as:

$$|z| = \sqrt{a^2 + b^2} \quad (218)$$

If this looks familiar, then it should! Since a complex number (211) is just a pair of real numbers, it encodes the same information as a point in the plane, or a vector in \mathbb{R}^2 . Then the absolute value of the complex number is just the length of the corresponding vector. Let us prove that:

$$\boxed{z\bar{z} = |z|^2} \quad (219)$$

Indeed, this is because:

$$(a + bi)(a - bi) = a^2 - abi + abi - b^2i^2 = a^2 + b^2$$

Formula (219) is important because it allows us to write:

$$\frac{1}{z} = \frac{\bar{z}}{|z|^2} \quad \text{i.e.} \quad \frac{1}{a + bi} = \frac{a - bi}{a^2 + b^2} \quad (220)$$

and this is great because it allows us to divide complex numbers:

$$\frac{2 + i}{4 + 3i} = \frac{(2 + i)(4 - 3i)}{4^2 + 3^2} = \frac{8 - 6i + 4i - 3i^2}{25} = \frac{11}{25} - \frac{2}{25}i$$

Example 7. We now have the tools to diagonalize the matrix (208). Its characteristic polynomial has two roots, which are complex numbers:

$$p(\lambda) = \lambda^2 + 1 = (\lambda - i)(\lambda + i) \quad \text{has roots} \quad \begin{cases} \lambda_1 = i \\ \lambda_2 = -i \end{cases}$$

so the eigenvalues are i and $-i$. Now we need to find eigenvectors corresponding to these eigenvalues:

$$\begin{aligned} A\mathbf{v}_1 = i\mathbf{v}_1 &\Rightarrow (A - iI)\mathbf{v}_1 = 0 &\Rightarrow \mathbf{v}_1 \in N(A - iI) \\ A\mathbf{v}_2 = -i\mathbf{v}_2 &\Rightarrow (A + iI)\mathbf{v}_2 = 0 &\Rightarrow \mathbf{v}_2 \in N(A + iI) \end{aligned}$$

To compute the nullspaces in question, we perform our good friend Gaussian elimination:

$$A - iI = \begin{bmatrix} -i & -1 \\ 1 & -i \end{bmatrix} \xrightarrow{\text{REF}} \begin{bmatrix} -i & -1 \\ 0 & 0 \end{bmatrix}$$

just by adding $(-i)$ times the first row to the second row. Therefore, $\mathbf{v}_1 = \begin{bmatrix} x \\ y \end{bmatrix}$ lies in $N(A - iI)$ if:

$$\begin{bmatrix} -i & -1 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = 0 \quad \Rightarrow \quad -ix - y = 0$$

so a choice of eigenvector for the eigenvalue i is $\mathbf{v}_1 = \begin{bmatrix} 1 \\ -i \end{bmatrix}$. Similarly, one finds that a choice of eigenvector for the eigenvalue $-i$ is $\mathbf{v}_2 = \begin{bmatrix} 1 \\ i \end{bmatrix}$. Therefore, Fact 19 implies that:

$$\begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix} = V \begin{bmatrix} i & 0 \\ 0 & -i \end{bmatrix} V^{-1} \quad \text{where} \quad V = \begin{bmatrix} 1 & 1 \\ -i & i \end{bmatrix}$$

and this concludes the diagonalization of the matrix A of (208).

Any real number can also be interpreted as a complex number, with imaginary part 0. With this in mind, we note that any quadratic polynomial has two complex roots, specifically:

$$ax^2 + bx + c = 0 \quad \text{has solutions} \quad \frac{-b \pm \sqrt{b^2 - 4ac}}{2a} \quad (221)$$

If $b^2 - 4ac \geq 0$, the solutions are real numbers. However, if $b^2 - 4ac < 0$, the square root in formula (221) will be an imaginary number (i.e. a multiple of i) and the solutions are complex non-real numbers. What's more, it turns out that complex numbers allow us to find roots of arbitrary polynomials. To this end, the **fundamental theorem of algebra** states that:

$$\text{any polynomial of degree } n \text{ has exactly } n \text{ complex roots} \quad (222)$$

where the word “exactly” should be taken to mean “counted with multiplicities”, as in Lecture 23.

Let's discuss one more feature of complex numbers, which is very useful in computations: the polar form. As we have seen, a complex number can be represented as a point in the plane via **Cartesian coordinates**:

$$z = a + ib \quad \rightsquigarrow \quad (a, b)$$

However, we can also represent the same point via **polar coordinates**:

$$z = a + ib \quad \rightsquigarrow \quad (r, \theta)$$

where $r = \sqrt{a^2 + b^2}$ is the absolute value of z (i.e. the size of a circle around the origin where the point is located) and $\theta \in [0, 2\pi)$ is the angle between the point in question and the horizontal line:

$$\cos \theta = \frac{a}{\sqrt{a^2 + b^2}} \quad \Rightarrow \quad \theta = \arccos \left(\frac{a}{\sqrt{a^2 + b^2}} \right)$$

The angle θ is called the **argument** of z , and is only defined up to adding arbitrary integer multiples of 2π (because θ and $\theta + 2\pi$ determine the same angle, there is an inherent ambiguity in the notion of “argument”). Since $(\cos \theta)^2 + (\sin \theta)^2 = 1$, we have $\sin \theta = \frac{b}{\sqrt{a^2 + b^2}}$ and therefore:

$$\boxed{z = a + bi = r(\cos \theta + i \sin \theta)} \quad (223)$$

The above is called the **polar form** of z . Recall the Taylor series expansions of $\cos \theta$ and $\sin \theta$:

$$\begin{aligned}\cos \theta &= 1 - \frac{\theta^2}{2!} + \frac{\theta^4}{4!} - \dots \\ \sin \theta &= \theta - \frac{\theta^3}{3!} + \frac{\theta^5}{5!} - \dots\end{aligned}$$

which implies that:

$$\cos \theta + i \sin \theta = 1 + i\theta - \frac{\theta^2}{2!} - \frac{i\theta^3}{3!} + \frac{\theta^4}{4!} + \frac{i\theta^5}{5!} - \dots$$

The right-hand side is the Taylor series of $e^{i\theta}$ (see (212)), so we obtain the following identity:

$$\boxed{e^{i\theta} = \cos \theta + i \sin \theta} \quad (224)$$

The polar form (223) of a complex number can therefore be written in a nicer form as:

$$\boxed{z = a + bi = re^{i\theta}} \quad (225)$$

Since $|e^{i\theta}| = 1$, the polar form of a complex number can be said to separate its absolute value part (namely r) from its angular part (namely $e^{i\theta}$). One of the good things about the polar form is that, as opposed from the Cartesian form, it's handy for multiplying numbers and raising to powers:

$$z^n = r^n e^{in\theta} \quad (226)$$

and:

$$zz' = rr' e^{i(\theta+\theta')} \quad (227)$$

if $z = re^{i\theta}$ and $z' = r'e^{i\theta'}$. In other words, we have:

$$|zz'| = |z||z'|$$

$$\arg(zz') = \arg(z) + \arg(z')$$

(the sum of arguments formula only holds up to adding integer multiples of 2π). Formula (226) allows us to describe the roots of unity, i.e. solutions of:

$$z^n = 1 \quad (228)$$

If z is a real number, the only solution to this equation is $z = 1$ (and also $z = -1$ if n is even), but the equation in question has more solutions in the world of complex numbers. Specifically, as a consequence of the fundamental theorem of algebra (222), we expect the equation (228) to have n complex solutions. If $z = re^{i\theta}$, then combining formulas (226) and (228) gives us:

$$r^n e^{in\theta} = 1$$

Since r is a positive real number, we must have $r = 1$. Since θ is an angle, the equation above implies $n\theta$ to be an even integral multiple of 2π . Therefore, we conclude that the solutions of the equation (228), namely the **n -th roots of unity**, are:

$$1, e^{\frac{2\pi i}{n}}, e^{\frac{4\pi i}{n}}, \dots, e^{\frac{2(n-1)\pi i}{n}} \quad (229)$$

There's no need to go beyond the $(n-1)$ -th multiple of $\frac{2\pi}{n}$, since then we will just recover the same numbers over again (this is simply because θ and $\theta + 2\pi$ represent the same angle).

Lecture 26 (April 23)

Reading: Section 6.4–6.5

Viewing: Lecture 25

Let us now consider eigenvalues and eigenvectors of symmetric matrices.

Fact 21. *A $n \times n$ symmetric matrix S has n real eigenvalues and n orthonormal eigenvectors.*

You can find a proof of the fact that the eigenvalues are real in the textbook, the bottom of page 339. For example, in the 2×2 case, a symmetric matrix:

$$S = \begin{bmatrix} a & b \\ b & c \end{bmatrix} \quad (230)$$

has characteristic polynomial:

$$p(\lambda) = \lambda^2 - \lambda(a + c) + ac - b^2$$

whose roots are:

$$\frac{a + c \pm \sqrt{(a + c)^2 - 4ac + 4b^2}}{2} = \frac{a + c \pm \sqrt{(a - c)^2 + 4b^2}}{2} \quad (231)$$

It is clear that these are real, because $(a - c)^2 + 4b^2 \geq 0$. As for the fact that the eigenvectors of a symmetric matrix S are orthogonal, suppose $\lambda \neq \lambda'$ are distinct eigenvalues with corresponding eigenvectors \mathbf{v}, \mathbf{v}' . Then:

$$\lambda \cdot \mathbf{v}^T \mathbf{v}' = (S\mathbf{v})^T \mathbf{v}' = \mathbf{v}^T S^T \mathbf{v}' = \mathbf{v}^T S \mathbf{v}' = \lambda' \cdot \mathbf{v}^T \mathbf{v}'$$

which implies $\mathbf{v}^T \mathbf{v}' = 0$. Once you know that the eigenvectors are orthogonal, they can be made orthonormal by rescaling them appropriately (since any multiple of an eigenvector is an eigenvector). As a consequence of the fact above, we may put the orthonormal eigenvectors $\mathbf{q}_1, \dots, \mathbf{q}_n$ in a matrix:

$$Q = [\mathbf{q}_1 \mid \dots \mid \mathbf{q}_n]$$

which is, by definition, an orthogonal matrix. Then (191) implies that we may write:

$$S = Q \begin{bmatrix} \lambda_1 & 0 & \dots & 0 \\ 0 & \lambda_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \lambda_n \end{bmatrix} Q^{-1} = Q \begin{bmatrix} \lambda_1 & 0 & \dots & 0 \\ 0 & \lambda_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \lambda_n \end{bmatrix} Q^T \quad (232)$$

The last equality follows from the fact that $Q^{-1} = Q^T$, a general feature of orthogonal matrices.

Geometrically, equality (232) can be interpreted as the fact that the linear transformation corresponding to a symmetric matrix can always be factored as:

$$\left(\text{rotation} \right) \left(\text{scaling} \right) \left(\text{inverse rotation} \right) \quad (233)$$

Indeed, diagonal matrices correspond to scaling in the direction of the coordinate axes, while orthogonal matrices correspond to rotations in n -dimensional space (as an exercise, prove that any

2×2 orthogonal matrix must be literally equal to one of the rotation matrices from Lecture 16). Explicitly, formula (232) implies that we can always express a symmetric matrix as:

$$\boxed{S = \lambda_1 \mathbf{q}_1 \mathbf{q}_1^T + \cdots + \lambda_n \mathbf{q}_n \mathbf{q}_n^T} \quad (234)$$

(you can prove this by applying both sides of identity (234) to an arbitrary eigenvector \mathbf{q}_i , and then extending by linear combinations to arbitrary vectors).

Remark. For an arbitrary (not necessarily symmetric) matrix A , the eigenvalues can be complex numbers. But whenever this happens, the following principle holds:

The complex eigenvalues of any real matrix come in conjugate pairs (235)

For example, in a 2×2 example, the characteristic polynomial:

$$p(\lambda) = \lambda^2 - \lambda \cdot \text{Tr} + \det$$

has roots:

$$\frac{\text{Tr} \pm \sqrt{\text{Tr}^2 - 4 \cdot \det}}{2}$$

If $\text{Tr}^2 < 4 \cdot \det$, the roots are complex numbers, and they are manifestly conjugates of each other. Moreover, the eigenvectors corresponding to conjugate eigenvalues are also conjugate:

$$A\mathbf{v} = \lambda\mathbf{v} \quad \Rightarrow \quad A\bar{\mathbf{v}} = \bar{\lambda}\bar{\mathbf{v}}$$

While in general, the eigenvalues and pivots of a square matrix are not directly related to each other (other than the fact that the eigenvalues and the pivots both have the same product, i.e. the determinant), for a symmetric matrix we have:

there are as many positive eigenvalues as positive pivots (236)

The analogous statement is true if “positive” were replaced by “negative”.

Definition 21. A symmetric matrix with all eigenvalues/pivots > 0 is called **positive definite**. A symmetric matrix with all eigenvalues/pivots ≥ 0 is called **positive semidefinite**.

In general, if S is a positive definite (respectively semidefinite) matrix, then the “energy”⁴:

$$\boxed{\mathbf{v}^T S \mathbf{v} > 0} \quad (\text{respectively } \mathbf{v}^T S \mathbf{v} \geq 0) \quad (237)$$

for any vector $\mathbf{v} \neq 0$. Indeed, just decompose \mathbf{v} as a linear combination of the eigenvectors of S :

$$\mathbf{v} = c_1 \mathbf{q}_1 + \cdots + c_n \mathbf{q}_n$$

and then (237) becomes (because the \mathbf{q}_i ’s are orthonormal):

$$\mathbf{v}^T S \mathbf{v} = \sum_{1 \leq i, j \leq n} c_i c_j \mathbf{q}_i^T S \mathbf{q}_j = \sum_{1 \leq i, j \leq n} c_i c_j \lambda_j \mathbf{q}_i^T \mathbf{q}_j = \sum_{i=1}^n c_i^2 \lambda_i$$

⁴The name comes from physics, in the case when \mathbf{v} is the vector of momenta of a system of particles

which is positive (respectively non-negative) if $\lambda_1, \dots, \lambda_n > 0$ (respectively $\lambda_1, \dots, \lambda_n \geq 0$). Note: condition (237) is actually equivalent to the matrix S being positive definite (respectively semidefinite). This gives a simple argument for why S and T being positive definite/semidefinite matrices implies that $\alpha S + \beta T$ is positive definite/semidefinite for any positive numbers α, β .

Remark. We already know that a big source of symmetric matrices are $S = A^T A$ for an arbitrary matrix A . This matrix is positive semidefinite, because:

$$\mathbf{v}^T S \mathbf{v} = \mathbf{v}^T A^T A \mathbf{v} = (A\mathbf{v})^T (A\mathbf{v}) = \|A\mathbf{v}\|^2 \geq 0 \quad (238)$$

Moreover, $S = A^T A$ is positive definite precisely if $A\mathbf{v} \neq 0$ for all $\mathbf{v} \neq 0$, which happens if and only if A has linearly independent columns (i.e. has zero nullspace). Conversely, any positive definite matrix is of the form $A^T A$ for some matrix A with linearly independent columns.

Let's work out in detail the case of a 2×2 symmetric matrix S , given by (230). Since the eigenvalues of S are given by (231), in order for them to both be positive, we would need $a + c > 0$ and:

$$a + c > \sqrt{(a - c)^2 + 4b^2} \quad \Rightarrow \quad (a + c)^2 > (a - c)^2 + 4b^2 \quad \Rightarrow \quad ac > b^2$$

This is actually equivalent to the stronger property $a, c > 0$ and $ac > b^2$. So we conclude that:

$$\mathbf{A} \quad 2 \times 2 \text{ symmetric matrix is positive definite precisely when } \text{Tr } S, \det S > 0 \quad (239)$$

As for the energy in the case when S is given by (230), we have:

$$\begin{bmatrix} x & y \end{bmatrix} \begin{bmatrix} a & b \\ b & c \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = ax^2 + 2bxy + cy^2 \quad (240)$$

It is a well-known fact in algebra that the expression above is positive for all x and y if $a, c > 0$ and $ac > b^2$, precisely as predicted by (239). In fact, you may recognize the equation “(240) = 1” as the equation of a conic in the plane: if S is positive definite, then this conic is an ellipse, while otherwise it's a hyperbola (the intermediary case, when one of the eigenvalues is 0, is a parabola).

Let's do all of this in a numerical example:

$$S = \begin{bmatrix} 5 & 4 \\ 4 & 5 \end{bmatrix}$$

Its characteristic polynomial is $\lambda^2 - 10\lambda + 9$, whose roots are 9 and 1. The corresponding eigenvectors are $\mathbf{v}_1 = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ 1 \end{bmatrix}$ and $\mathbf{v}_2 = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ -1 \end{bmatrix}$, hence we conclude that:

$$S = QDQ^T$$

where:

$$D = \begin{bmatrix} 9 & 0 \\ 0 & 1 \end{bmatrix} \quad \text{and} \quad Q = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}$$

The ellipses corresponding to S and D are:

$$5x^2 + 8xy + 5y^2 = 1 \quad \text{and} \quad 9x^2 + y^2 = 1$$

The latter ellipse has its axes lined up with the coordinate axes. The former has its axes lined up with the eigenvectors \mathbf{v}_1 and \mathbf{v}_2 , which is consistent with the fact that going from D to S amounts to rotating the coordinate axes from $\mathbf{e}_1, \mathbf{e}_2$ to $\mathbf{v}_1, \mathbf{v}_2$, i.e. applying the linear transformation Q .

Lecture 27 (*April 26*)

Reading: Section 7.1–7.2

Viewing: Lecture 29

We will now study the **Singular Value Decomposition (SVD)** of a matrix, which is very useful in computations. Think of image processing: in general, this deals with a large rectangular table of pixels, and the color of every pixel is an integer. So the image is encoded by an $m \times n$ matrix A of integers, where m and n are two very large numbers which keep track of how many pixels we have vertically and horizontally. If the colors of all the pixels are random, then there's not much you can do to simplify the matrix A . But in real life, nearby pixels tend to have similar colors, and patterns emerge that allow you to conclude that the matrix A has a special form. An extreme example is where your image consists only of three colors separated by vertical lines, such as a flag:

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

Remembering the whole matrix is both memory and computational power-consuming, so you'd be more efficient by remembering instead the fact that A can be presented as the product:

$$A = \begin{bmatrix} 1 \\ 1 \\ 1 \\ 1 \\ 1 \\ 1 \\ 1 \\ 1 \end{bmatrix} \cdot [1 \quad 1 \quad 2 \quad 2 \quad 2 \quad 2 \quad 2 \quad 3 \quad 3 \quad 3]$$

of a column vector and a row vector. In general, the goal of SVD is to write any rectangular matrix as a sum of such products (let r be the rank of A):

$$A = \mathbf{u}_1 \sigma_1 \mathbf{v}_1^T + \cdots + \mathbf{u}_r \sigma_r \mathbf{v}_r^T \tag{241}$$

where each \mathbf{u}_i is a column vector, each \mathbf{v}_i^T is a row vector, and:

the positive numbers $\sigma_1, \dots, \sigma_r > 0$ **are called singular values**

Let us now explain how to obtain the decomposition (241).

Fact 22. *The $m \times m$ matrix AA^T and the $n \times n$ matrix $A^T A$ are both symmetric and positive semidefinite. These two matrices have the same set of positive eigenvalues; call them $\sigma_1^2, \dots, \sigma_r^2$.*

Proof: assume that $\lambda \neq 0$ is an eigenvalue of AA^T . Then:

$$AA^T \mathbf{v} = \lambda \mathbf{v} \quad \Rightarrow \quad A^T A(A^T \mathbf{v}) = \lambda(A^T \mathbf{v})$$

which implies that λ is also an eigenvalue of $A^T A$ corresponding to the eigenvector $A^T \mathbf{v}$ (the reason you need $\lambda \neq 0$ is in order to infer that $A^T \mathbf{v} \neq 0$, since eigenvectors must be non-zero by definition).

Since symmetric matrices (such as AA^T and $A^T A$) have orthonormal eigenvectors, we can choose:

$$\begin{aligned} \text{left singular vectors } & \mathbf{u}_1, \dots, \mathbf{u}_r, \dots, \mathbf{u}_m \in \mathbb{R}^m \\ \text{right singular vectors } & \mathbf{v}_1, \dots, \mathbf{v}_r, \dots, \mathbf{v}_n \in \mathbb{R}^n \end{aligned}$$

such that the \mathbf{u}_i 's are orthonormal and the \mathbf{v}_i 's are orthonormal, and:

$$AA^T \mathbf{u}_i = \sigma_i^2 \mathbf{u}_i \quad \text{and} \quad A^T A \mathbf{v}_i = \sigma_i^2 \mathbf{v}_i \quad (242)$$

(the formulas above even hold for $i > r$, if we let $\sigma_{r+1}, \sigma_{r+2}, \dots = 0$). But we actually want the left and right singular vectors to also be connected, via the following formulas:

$$\boxed{A \mathbf{v}_i = \sigma_i \mathbf{u}_i} \quad \text{and} \quad \boxed{A^T \mathbf{u}_i = \sigma_i \mathbf{v}_i} \quad (243)$$

which also hold for all i , even $i > r$. The numbers $\sigma_1, \dots, \sigma_r$ are called **singular values**. It is easy to see that (243) implies both (242) and (241): the former implication is an easy exercise, while the latter follows by multiplying both sides of (241) with any given vector \mathbf{v}_i and using that $\mathbf{v}_j^T \mathbf{v}_i = 0$ if $j \neq i$. Note that the singular value decomposition (241) can be written in compact form as:

$$\boxed{A = U \Sigma V^T} \quad (244)$$

where:

$$U = [\mathbf{u}_1 \mid \dots \mid \mathbf{u}_m] \quad \Sigma = \begin{bmatrix} \sigma_1 & 0 & 0 & 0 & \dots \\ 0 & \ddots & 0 & 0 & \dots \\ 0 & 0 & \sigma_r & 0 & \dots \\ 0 & 0 & 0 & 0 & \dots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix} \quad V = [\mathbf{v}_1 \mid \dots \mid \mathbf{v}_n]$$

Remark. From (243), it follows that $\mathbf{u}_1, \dots, \mathbf{u}_r$ form a basis of the column space of A , and $\mathbf{v}_1, \dots, \mathbf{v}_r$ form a basis of the row space. The orthogonality of the four subspaces implies that $\mathbf{u}_{r+1}, \dots, \mathbf{u}_m$ form a basis of the left nullspace of A , while $\mathbf{v}_{r+1}, \dots, \mathbf{v}_n$ form a basis of the nullspace.

Since V is an orthogonal matrix, we have $V^T = V^{-1}$, so (244) implies:

$$\boxed{AV = U \Sigma} \quad (245)$$

which is just another way to write the first identity in (243). As for the second identity in (243), we can obtain it as follows. First transpose relation (244):

$$A^T = V \Sigma^T U^T \quad (246)$$

Since U is an orthogonal matrix, we have $U^T = U^{-1}$, hence:

$$A^T U = V \Sigma^T \quad (247)$$

Therefore, the SVD of the transpose of A is obtained by transposing the SVD of A .

Let's do an example of SVD, which will also show us that the procedure works, for the matrix:

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

We first need to compute the singular values, which will be the square roots of the eigenvalues of:

$$A^T A = \begin{bmatrix} 3 & 2 & 2 \\ 2 & 3 & -2 \end{bmatrix} \begin{bmatrix} 3 & 2 \\ 2 & 3 \\ 2 & -2 \end{bmatrix} = \begin{bmatrix} 17 & 8 \\ 8 & 17 \end{bmatrix}$$

The characteristic polynomial of this matrix is:

$$p(\lambda) = \det \begin{bmatrix} 17 - \lambda & 8 \\ 8 & 17 - \lambda \end{bmatrix} = (17 - \lambda)^2 - 8^2$$

and its roots will be precisely those numbers for which $17 - \lambda = \pm 8$, so:

$$\begin{aligned} 17 - \sigma_1^2 = -8 &\Rightarrow \sigma_1^2 = 25 \Rightarrow \sigma_1 = 5 \\ 17 - \sigma_2^2 = 8 &\Rightarrow \sigma_2^2 = 9 \Rightarrow \sigma_2 = 3 \end{aligned}$$

(it is customary to order your singular values $\sigma_1, \dots, \sigma_r$ from the greatest to the smallest). Next, you need to find the right singular vectors, which will be eigenvectors of the matrix $A^T A$:

$$\begin{aligned} \mathbf{v}_1 &\in N \left(\begin{bmatrix} 17 - \sigma_1^2 & 8 \\ 8 & 17 - \sigma_1^2 \end{bmatrix} \right) = N \left(\begin{bmatrix} -8 & 8 \\ 8 & -8 \end{bmatrix} \right) \Rightarrow \mathbf{v}_1 = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ 1 \end{bmatrix} \\ \mathbf{v}_2 &\in N \left(\begin{bmatrix} 17 - \sigma_2^2 & 8 \\ 8 & 17 - \sigma_2^2 \end{bmatrix} \right) = N \left(\begin{bmatrix} 8 & 8 \\ 8 & 8 \end{bmatrix} \right) \Rightarrow \mathbf{v}_2 = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ -1 \end{bmatrix} \end{aligned}$$

Note that \mathbf{v}_1 and \mathbf{v}_2 are orthonormal (indeed, recall that eigenvectors of symmetric matrices corresponding to different eigenvalues are always orthogonal, and we have intentionally scaled them so as to have length 1). Now let's compute the left singular vectors. Formula (243) forces us to take:

$$\begin{aligned} \mathbf{u}_1 &= \frac{1}{5} A \mathbf{v}_1 = \frac{1}{5} \begin{bmatrix} 3 & 2 \\ 2 & 3 \\ 2 & -2 \end{bmatrix} \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ 1 \end{bmatrix} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix} \\ \mathbf{u}_2 &= \frac{1}{3} A \mathbf{v}_2 = \frac{1}{3} \begin{bmatrix} 3 & 2 \\ 2 & 3 \\ 2 & -2 \end{bmatrix} \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ -1 \end{bmatrix} = \frac{1}{3\sqrt{2}} \begin{bmatrix} 1 \\ -1 \\ 4 \end{bmatrix} \end{aligned}$$

The vectors \mathbf{u}_1 and \mathbf{u}_2 are also orthonormal (this follows from the fact that \mathbf{v}_1 and \mathbf{v}_2 are) but we need to complete them to an orthonormal basis of \mathbb{R}^3 . This is achieved by Gram-Schmidt. Start with any vector that is not a linear combination of \mathbf{u}_1 and \mathbf{u}_2 , say:

$$\mathbf{a}_3 = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}$$

and construct the vector:

$$\mathbf{w}_3 = \mathbf{a}_3 - \text{proj}_{\mathbf{u}_1} \mathbf{a}_3 - \text{proj}_{\mathbf{u}_2} \mathbf{a}_3 = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} - \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix} \cdot 0 - \frac{1}{3\sqrt{2}} \begin{bmatrix} 1 \\ -1 \\ 4 \end{bmatrix} \cdot \frac{4}{3\sqrt{2}} = \frac{1}{9} \begin{bmatrix} -2 \\ 2 \\ 1 \end{bmatrix}$$

The vector \mathbf{w}_3 thus constructed is orthogonal to both \mathbf{u}_1 and \mathbf{u}_2 , so we need to renormalize it:

$$\mathbf{u}_3 = \frac{\mathbf{w}_3}{\|\mathbf{w}_3\|} = \frac{1}{3} \begin{bmatrix} -2 \\ 2 \\ 1 \end{bmatrix}$$

With this in mind, we conclude the following SVD for the matrix A :

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

So in general, what can we say about a rank 1 matrix? The answer is that it will be of the form:

$$A = \mathbf{u}\sigma\mathbf{v}^T \tag{248}$$

where \mathbf{u} and \mathbf{v} are length 1 vectors, and σ is a non-zero number. In other words, any column of a rank 1 matrix is a multiple of \mathbf{u} and any row is a multiple of \mathbf{v} . Obviously, you can rescale \mathbf{u} and \mathbf{v} such that they have length 1 (at the cost of changing the number σ) and you can flip the sign on σ (at the cost of replacing \mathbf{v} by $-\mathbf{v}$) so as to ensure $\sigma > 0$. With this in mind, the singular value decomposition (241) can be thought of as:

SVD writes any matrix as a sum of rank 1 matrices (249)

where the columns (respectively rows) of the rank 1 matrices involved are orthonormal. This is good to keep in mind, since it can save you a lot of time when doing SVD on a matrix where it's intuitively pretty obvious how to break it up as a sum of rank 1 matrices. For example:

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

If you started doing the usual algorithm (compute the eigenvalues and eigenvectors of the 3×3 matrix $A^T A$) you might find yourself doing a lot of work. You can save a lot of time by observing that A decomposes as:

$$A = \begin{bmatrix} 0 & 5 & 0 \\ 0 & 0 & 0 \end{bmatrix} + \begin{bmatrix} 0 & 0 & 0 \\ -3 & 0 & 0 \end{bmatrix}$$

Each of the two matrices above has rank 1, because all of their columns are multiples of a single vector. Even more explicitly, we have:

$$A = \underbrace{\begin{bmatrix} 1 \\ 0 \end{bmatrix}}_{\mathbf{u}_1} \cdot \underbrace{5}_{\sigma_1} \cdot \underbrace{\begin{bmatrix} 0 & 1 & 0 \end{bmatrix}}_{\mathbf{v}_1} + \underbrace{\begin{bmatrix} 0 \\ -1 \end{bmatrix}}_{\mathbf{u}_2} \cdot \underbrace{3}_{\sigma_2} \cdot \underbrace{\begin{bmatrix} 1 & 0 & 0 \end{bmatrix}}_{\mathbf{v}_2}$$

However, when using this shortcut, you need to make sure that the vectors $\mathbf{u}_1, \mathbf{u}_2, \dots$ you construct are orthonormal (and also that the vectors $\mathbf{v}_1, \mathbf{v}_2, \dots$ are orthonormal).

Lecture 28 (April 28)

Reading: Section 7.4
Viewing: Lecture 33

In practice, the singular value decomposition (241) of a matrix gives us a computationally manageable way to compute how A acts on arbitrary vectors:

$$\mathbf{v} = c_1 \mathbf{v}_1 + \cdots + c_n \mathbf{v}_n \quad (250)$$

(since the \mathbf{v}_i are orthonormal, we have $c_i = \mathbf{v}_i^T \mathbf{v}$). Then we have:

$$A\mathbf{v} = \sigma_1 c_1 \mathbf{u}_1 + \cdots + \sigma_r c_r \mathbf{u}_r \quad (251)$$

It is customary to order your singular values $\sigma_1 \geq \cdots \geq \sigma_r > 0$. In this case, you can compare (250) with (251) to get:

$$\frac{\|A\mathbf{v}\|}{\|\mathbf{v}\|} = \sqrt{\frac{\sigma_1^2 c_1^2 + \cdots + \sigma_r^2 c_r^2}{c_1^2 + \cdots + c_r^2 + \cdots + c_n^2}} \leq \sigma_1$$

Moreover, we get equality in the inequality above when \mathbf{v} is a multiple of \mathbf{v}_1 . Therefore:

$$\sigma_1 \text{ is the maximum of } \frac{\|A\mathbf{v}\|}{\|\mathbf{v}\|}, \text{ achieved for } \mathbf{v} = \mathbf{v}_1 \quad (252)$$

which is an intrinsic characterization of the largest singular value. Specifically, we are saying that the largest singular value is equal to the **norm** of the matrix:

$$\|A\| := \max_{\mathbf{w} \in \mathbb{R}^n \setminus \{0\}} \frac{\|A\mathbf{w}\|}{\|\mathbf{w}\|} \in \mathbb{R}_+ \quad (253)$$

The norm of a matrix satisfies the usual “triangle inequality” satisfied by the norms of vectors:

$$\|A + B\| \leq \|A\| + \|B\| \quad (254)$$

(proof: for any vector \mathbf{w} , we have $\|(A + B)\mathbf{w}\| \leq \|A\mathbf{w}\| + \|B\mathbf{w}\| \leq \|A\| \cdot \|\mathbf{w}\| + \|B\| \cdot \|\mathbf{w}\| = (\|A\| + \|B\|) \cdot \|\mathbf{w}\|$, which implies (254)) as well as a similar inequality when multiplying matrices:

$$\|AB\| \leq \|A\| \cdot \|B\| \quad (255)$$

(proof: for any vector \mathbf{w} , we have $\|AB\mathbf{w}\| \leq \|A\| \cdot \|B\mathbf{w}\| \leq \|A\| \cdot \|B\| \cdot \|\mathbf{w}\|$, which implies (255)). You can obtain norm-like descriptions of the other singular values of a matrix, as follows:

$$\sigma_2 \text{ is the maximum of } \frac{\|A\mathbf{v}\|}{\|\mathbf{v}\|} \text{ among those } \mathbf{v} \perp \mathbf{v}_1, \text{ achieved for } \mathbf{v} = \mathbf{v}_2 \quad (256)$$

$$\sigma_3 \text{ is the maximum of } \frac{\|A\mathbf{v}\|}{\|\mathbf{v}\|} \text{ among those } \mathbf{v} \perp \mathbf{v}_1, \mathbf{v}_2, \text{ achieved for } \mathbf{v} = \mathbf{v}_3 \quad (257)$$

etc. In fact, this recursive procedure defines the singular values σ_i and the right singular vectors \mathbf{v}_i (from which the left singular vectors can be deduced from (243)): first σ_1 and \mathbf{v}_1 are defined by (252), then σ_2 and \mathbf{v}_2 are defined by (256), then σ_3 and \mathbf{v}_3 are defined by (257) etc.

Geometrically, the SVD in the form (244) means that any matrix A can be factored as:

$$\left(\text{rotation}\right) \left(\text{scaling}\right) \left(\text{another rotation}\right) \quad (258)$$

Compare this with (233) for symmetric matrices: the novelty in (258) is that the two rotations, which correspond to the orthogonal matrices U and V^T respectively, are now completely different linear transformations acting on two different vector spaces. This is natural, since A corresponds to a linear transformation from \mathbb{R}^n to \mathbb{R}^m . Meanwhile, the two rotations in (233) are inverses to each other, which is natural because the symmetric matrix therein had the same domain and codomain.

You can find a picture of the factorization (258) on page 392 of the textbook: the two rotations do not change any given shape, they just rotate it around space. But the scaling in the middle does change shapes by dilations across various axes: circles become ellipses, squares become rectangles etc. In the case of 2×2 matrices, the factorization (258) explicitly says that:

$$\underbrace{\begin{bmatrix} a & b \\ c & d \end{bmatrix}}_A = \underbrace{\begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix}}_U \underbrace{\begin{bmatrix} \sigma_1 & 0 \\ 0 & \sigma_2 \end{bmatrix}}_\Sigma \underbrace{\begin{bmatrix} \cos \phi & \sin \phi \\ -\sin \phi & \cos \phi \end{bmatrix}}_{V^T}$$

Now let us recall formula (246), which explains that the transpose of a matrix A has the transposed singular value decomposition. So what can we say about the “inverse” of a matrix? The reason why I put the word “inverse” in quotes is that we are in the rectangular matrix case, where a proper inverse most likely does not exist. However, if it did, then from formula (243) we would expect:

$$\text{“}A^{-1}\text{”}(\mathbf{u}_i) = \frac{\mathbf{v}_i}{\sigma_i} \quad (259)$$

The goal of the following is to define a notion of “inverse” which satisfies the equation above.

Definition 22. The *pseudo-inverse* of an $m \times n$ matrix $A = U\Sigma V^T$ is the $n \times m$ matrix:

$$A^+ = V\Sigma^+U^T \quad (260)$$

$$\text{where } \Sigma^+ = \begin{bmatrix} \frac{1}{\sigma_1} & 0 & 0 & 0 & \dots \\ 0 & \ddots & 0 & 0 & \dots \\ 0 & 0 & \frac{1}{\sigma_r} & 0 & \dots \\ 0 & 0 & 0 & 0 & \dots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix}.$$

The pseudo-inverse coincides with the inverse in the cases when the latter exists, i.e. $m = n = r$. In general, the pseudo-inverse satisfies the following inverse-like properties:

$$A^+A = V\Sigma^+U^T U\Sigma V^T = V(\Sigma^+\Sigma)V^T = V \left[\begin{array}{c|c} I_r & 0 \\ \hline 0 & 0 \end{array} \right] V^T$$

$$AA^+ = U\Sigma V^T V\Sigma^+U^T = U(\Sigma\Sigma^+)U^T = U \left[\begin{array}{c|c} I_r & 0 \\ \hline 0 & 0 \end{array} \right] U^T$$

If the block matrix in the right-hand sides were just equal to the identity matrix I , then A and A^+ would be bona fide inverses. However, as it is in the formulas above, the block matrix is simply the projection matrix onto the first r coordinate vectors. And when we conjugate this block matrix by the orthogonal matrices V and U , respectively, we conclude that:

$$A^+A = \text{projection onto } C(A^T) \quad (261)$$

$$AA^+ = \text{projection onto } C(A) \quad (262)$$

This happens because the row space (respectively the column space) of the matrix A is spanned by the first r vectors of the matrix V (respectively U), as in the Remark in the previous Lecture.

As an example, let's compute the pseudo-inverse of the matrix:

$$A = \begin{bmatrix} 9 & 12 \\ 12 & 16 \end{bmatrix}$$

This matrix has rank 1, so its SVD is computed just like at the end of the previous lecture:

$$A = \begin{bmatrix} 3 \\ 4 \end{bmatrix} [3 \quad 4] = \begin{bmatrix} \frac{3}{5} \\ \frac{4}{5} \end{bmatrix} \cdot 25 \cdot \begin{bmatrix} \frac{3}{5} & \frac{4}{5} \end{bmatrix} \Rightarrow A^+ = \begin{bmatrix} \frac{3}{5} \\ \frac{4}{5} \end{bmatrix} \cdot \frac{1}{25} \cdot \begin{bmatrix} \frac{3}{5} & \frac{4}{5} \end{bmatrix} = \frac{1}{625} \begin{bmatrix} 9 & 12 \\ 12 & 16 \end{bmatrix}$$

Pseudo-inverses have a nice application to least squares: remember that the closest approximation to the system $A\mathbf{v} = \mathbf{b}$ is $\mathbf{v} = (A^T A)^{-1} A^T \mathbf{b}$, but we always made the caveat that the columns of A need to be independent in order for $A^T A$ to be invertible. If A has dependent columns, while we cannot write $(A^T A)^{-1}$, we can use the pseudo-inverse to give us the following solution:

$$\text{the closest } A\mathbf{v} \text{ is to } \mathbf{b} \text{ is achieved for } \boxed{\mathbf{v}^+ = A^+ \mathbf{b}} \quad (263)$$

Such a \mathbf{v} is not unique, basically because the fact that A has dependent columns implies that its nullspace is non-trivial (so to any particular solution \mathbf{v} , you could add any vector in the nullspace of A , and you would still get a solution). But to see that the particular solution \mathbf{v}^+ in (263) does indeed satisfy $A\mathbf{v}^+ = \text{proj}_{C(A)} \mathbf{b}$, just apply A to both sides of the equation in the box. You get:

$$A\mathbf{v}^+ = (AA^+) \mathbf{b} = \text{proj}_{C(A)} \mathbf{b}$$

as a consequence of (262). So $A\mathbf{v}^+$ is precisely the projection of \mathbf{b} onto the column space of A , as required by the geometry of least squares approximation.

Let's now consider one last application of SVD's in the case of square matrices (i.e. $m = n$):

Definition 23. The *polar decomposition* of an $n \times n$ matrix A is:

$$A = QS \quad (264)$$

where Q is orthogonal, and S is positive semidefinite.

The name stems from the fact that in the 1×1 case, any number a has a polar form (225): the number $e^{i\theta}$ has absolute value 1 so it behaves like an orthogonal 1×1 matrix, and the number $r \geq 0$ is a positive semidefinite 1×1 matrix. In general, you can obtain the polar decomposition of a matrix from its singular value decomposition, as follows:

$$A = U\Sigma V^T = \underbrace{(UV^T)}_Q \cdot \underbrace{(V\Sigma V^T)}_S \quad (265)$$

Indeed, the product of orthogonal matrices is orthogonal (it is key here that $m = n$, i.e. this all applies to square matrices), which means that $Q := UV^T$ is orthogonal. Moreover, $S := V\Sigma V^T$ is the usual formula for the diagonalization of a symmetric matrix, as we have seen in (232). The fact

that the eigenvalues of S , i.e. the diagonal entries of Σ , are non-negative is precisely equivalent with saying that S is positive semidefinite. The eigenvalues of S are the singular values of A .

MIDTERM 2 (*April 30*)

Lecture 29 (May 3)

Reading: Section 12.1

Viewing: none

We will now start discussing major applications of linear algebra: probability and statistics. First of these is **probability**, which studies the likelihood of future events. Suppose we're in a situation with n possible outcomes x_1, \dots, x_n (think of these as real numbers), which arise with probabilities p_1, \dots, p_n , respectively. Then:

$$p_1, \dots, p_n > 0 \quad \text{and} \quad p_1 + \dots + p_n = 1 \quad (266)$$

Definition 24. The **mean** is the sum of the possible outcomes, weighted by their probabilities:

$$\mu = p_1 x_1 + \dots + p_n x_n \quad (267)$$

and the **variance** is:

$$\Sigma = p_1(x_1 - \mu)^2 + \dots + p_n(x_n - \mu)^2 \quad (268)$$

The square root of the variance, namely $\sigma = \sqrt{\Sigma}$, is called the **standard deviation**.

The mean is sometimes called the **expected value**, which is the plain English significance of this notion. The variance is a positive quantity which measures how far away from the mean our outcomes are. Clearly, the variance is 0 only if there exists a single outcome with probability 1.

The situation described above is when there are finitely many (or discrete) outcomes. But we could also study a quantity which takes infinitely many (or continuous) values, for example temperature. In this case, the probability is encoded not in a finite set of numbers, but in a function:

$$p(x) : \mathbb{R} \longrightarrow \mathbb{R}_{\geq 0}$$

called **probability distribution**. In this context, you shouldn't measure the probability of your quantity being exactly equal to x (for example, it makes little sense to measure the probability that temperature will be exactly 62.746 degrees) but you measure instead the probability that the quantity lies in a certain interval $[a, b]$ (for example, you measure the probability that temperature falls between 62 and 63 degrees):

$$\text{Prob}(x \in [a, b]) = \int_a^b p(x) dx \quad (269)$$

The reason for the integral is that it is the continuous analogue of a sum: if you subdivide the interval $[a, b]$ into many (say N very large) small intervals, then the integral in (269) is infinitesimally close to the sum of the probabilities that x lies in the N small intervals. In this continuous context, condition (266) is replaced by:

$$\int_{-\infty}^{\infty} p(x) dx = 1 \quad (270)$$

while the mean (267) and the variance (268) must be defined by integrals:

$$\mu = \int_{-\infty}^{\infty} p(x) x dx \quad (271)$$

$$\Sigma = \int_{-\infty}^{\infty} p(x)(x - \mu)^2 dx \quad (272)$$

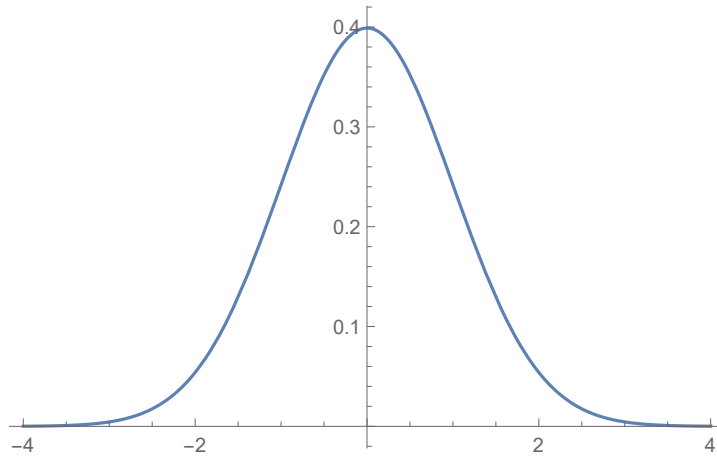
The most basic example of probability distribution is the uniform one, where the variable can lie in an interval (say $[0, c]$) with constant probability:

$$p(x) = \begin{cases} \frac{1}{c} & \text{if } x \in [0; c] \\ 0 & \text{otherwise} \end{cases}$$

However, this probability distribution does not really show up in nature. Real-life probabilities do not like the sharp cut-offs that the function above experiences at the endpoints of the interval $[0, c]$. Perhaps even more importantly, real-life probability distributions tend to be higher around the mean and lower away from the mean. The standard example of such a behavior is the **normal (or Gaussian) distribution**:

$$p(x) = \frac{e^{-\frac{x^2}{2}}}{\sqrt{2\pi}} \quad (273)$$

whose graph (sometimes called a “bell curve”) is plotted below. As always, the probability that the variable x lies in a given interval $[a, b]$ is equal to the integral of $p(x)$ (or the area under the graph) from $x = a$ to $x = b$.



The reason for dividing by the factor $\sqrt{2\pi}$ is because you want the total probability to be equal to 1, i.e. (270). The mean and variance of the normal distribution (273) are given by:

$$\mu = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} x e^{-\frac{x^2}{2}} dx = 0$$

$$\Sigma = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} x^2 e^{-\frac{x^2}{2}} dx = 1$$

More generally, the normal probability distribution with arbitrary mean $\mu \in \mathbb{R}$ and arbitrary variance $\Sigma > 0$ is given by the function:

$$p(x) = \frac{e^{-\frac{(x-\mu)^2}{2\Sigma}}}{\sqrt{2\pi\Sigma}} \quad (274)$$

Normal distributions are extremely important, because of the **central limit theorem**: if you perform an experiment $N \rightarrow \infty$ times, the mean of your samples tends to look as if it came from a normal distribution. For example, if you flip a coin N times, the probability of getting k heads is close to the probability that $x \in [k, k + 1]$ for the normal distribution (274) with mean $\mu = \frac{N}{2}$ and variance $\Sigma = \frac{N}{4}$. In the error analysis of the probability distribution (274), the probability that:

x is at most σ away from the mean is $\sim 68\%$

x is at most 2σ away from the mean is $\sim 95\%$

x is at most 3σ away from the mean is $\sim 99.7\%$

x is at most 4σ away from the mean is $\sim 99.99\%$

x is at most 5σ away from the mean is $\sim 99.9999\%$

You can actually see this in the graph of the bell curve above: about 99.7% of the area under the graph of the bell curve is between $x = -3$ and $x = 3$ etc.

Lecture 30 (May 5)

Reading: Section 12.2

Viewing: none

The true power of linear algebra in probability comes about when we run more than one experiment at once, and we want to compare the outcomes, like in the example of performing many coin tosses in the central limit theorem. The abstract concept that governs this is the following.

Definition 25. A *random variable* is a quantity X that takes values in \mathbb{R} , which is either:

- discrete, so takes finitely many values x_1, \dots, x_n with probabilities p_1, \dots, p_n , or
- continuous, associated to a probability distribution $p(x)$.

In either case, the mean (or expected value) is denoted by:

$$E[X] = \mu \quad (275)$$

and is defined by either (267) or (271). Then the variance is:

$$E[(X - \mu)^2] = \Sigma \quad (276)$$

and is given by either (268) or (272).

Definition 26. Given two random variables X and Y , their **covariance** is the quantity:

$$\Sigma_{XY} = E[(X - E[X])(Y - E[Y])] \quad (277)$$

Note that Σ_{XX} is just the **variance** of the random variable X itself.

Let's start with the case of discrete probabilities, so X and Y both take a finite number of values with certain probabilities. The whole point of covariance is to measure the extent to which X and Y are correlated with each other. So it does not suffice to simply know the probabilities of the various values of X and Y separately, but instead you need to know the **joint probabilities**:

$$p_{ij} \text{ is the probability that } (X, Y) = (x_i, y_j) \quad (278)$$

where x_1, x_2, \dots, x_k are the possible values of X and y_1, y_2, \dots, y_k are the possible values of Y .

Example 8. Let's assume that X and Y are coin flips, and their possible values are heads (H) and tails (T). If the two coin flips are independent, then the joint probabilities are:

$$p_{HH} = p_{TT} = p_{HT} = p_{TH} = \frac{1}{4} \quad (279)$$

But if the random variables are two coin flips which are forced to always land on the same side (think something involving magnets), then the joint probabilities are:

$$p_{HH} = p_{TT} = \frac{1}{2} \quad \text{and} \quad p_{HT} = p_{TH} = 0 \quad (280)$$

The means of two random variables X and Y are:

$$\mu = \sum_{i,j} p_{ij} x_i$$

$$\nu = \sum_{i,j} p_{ij} y_j$$

and therefore the covariance of the two random variables is:

$$\Sigma_{XY} = \sum_{i,j} p_{ij} (x_i - \mu)(y_j - \nu) \quad (281)$$

If X and Y are continuous random variables, their **joint probability distribution** is a function:

$$p(x, y) : \mathbb{R}^2 \rightarrow \mathbb{R}_{\geq 0} \quad (282)$$

which encodes probability in the sense that:

$$\text{Prob}(X \in [a, b] \text{ and } Y \in [c, d]) = \int_a^b \int_c^d p(x, y) dy dx \quad (283)$$

With this in mind, the covariance of the random variables X and Y is:

$$\Sigma_{XY} = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} p(x, y) (x - \mu)(y - \nu) dx dy \quad (284)$$

where μ and ν are the means of X and Y , defined by (271).

Remark. *The notion (282) is necessary because the random variables X and Y may be dependent on each other: mathematically, this means that where X lands on the real line might impact where Y lands. For example, if $X = Y$, then X and Y cannot be at different points of the real line, so in this case the function $p(x, y)$ is concentrated on the diagonal in \mathbb{R}^2 .*

An important special case is when two variables X and Y are independent of each other. In the discrete case, this means that the probabilities (278) are given by $p_{ij} = a_i b_j$, where a_1, a_2, \dots and b_1, b_2, \dots are the probabilities of X and Y separately. In the continuous case, the fact that X and Y are independent means that the joint probability distribution (282) is given by:

$$p(x, y) = a(x)b(y)$$

where a and b are the probability distributions of X and Y , respectively. In this case, the position of X on the real line has no impact on the position of Y and vice versa. The covariance (283) is:

$$\Sigma_{XY} = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} a(x)b(y)(x - \mu)(y - \nu) dx dy = \left(\int_{-\infty}^{\infty} a(x)(x - \mu) dx \right) \left(\int_{-\infty}^{\infty} b(y)(y - \nu) dy \right) = 0$$

(the last equality is a consequence of (270) and (271)). We conclude that the covariance of two independent random variables is 0. In general, we have the following principle:

- the more positive the covariance, the more correlated the random variables
- the more negative the covariance, the more anti-correlated the random variables
- if the covariance is 0, the random variables are called uncorrelated

For any random variables X and Y (discrete or continuous), the Cauchy-Schwartz inequality states:

$$|\Sigma_{XY}| \leq \sqrt{\Sigma_{XX}\Sigma_{YY}} \quad (285)$$

In other words, the covariance is bounded by the geometric mean of the variances. Therefore, if the variances of X and Y are fixed, the maximum covariance happens when $X = Y$ (the two random variables are as correlated as possible) and the minimum covariance happens when $X = -Y$ (the two random variables are as anti-correlated as possible).

Definition 27. *The covariance matrix of two random variables X and Y is:*

$$K = \begin{bmatrix} \Sigma_{XX} & \Sigma_{XY} \\ \Sigma_{YX} & \Sigma_{YY} \end{bmatrix} \quad (286)$$

Since $\Sigma_{XY} = \Sigma_{YX}$, the covariance matrix is symmetric. But it is actually positive semidefinite, because its trace is positive (as Σ_{XX} and Σ_{YY} are the variances of X and Y) and its determinant is ≥ 0 according to (285). To further analyze the covariance matrix, we have two sub-cases:

- *either* K is positive definite,
- *or* $\det K = 0$. When this happens, X and Y are called perfectly correlated.

Let's compute the covariance matrix of discrete random variables in terms of the probabilities (278):

$$K = \sum_{i,j=1}^k p_{ij} \begin{bmatrix} (x_i - \mu)^2 & (x_i - \mu)(y_j - \nu) \\ (x_i - \mu)(y_j - \nu) & (y_j - \nu)^2 \end{bmatrix} = \sum_{i,j=1}^k p_{ij} \underbrace{\begin{bmatrix} x_i - \mu \\ y_j - \nu \end{bmatrix} \begin{bmatrix} x_i - \mu & y_j - \nu \end{bmatrix}} \quad (287)$$

The underbraced term is a positive semidefinite symmetric matrix, since it is of the form $\mathbf{v}\mathbf{v}^T$. The p_{ij} 's are non-negative numbers, and it is easy to show that a non-negative linear combination of positive semidefinite matrices is positive semidefinite (for example, using the energy criterion (237)). This gives another proof of the fact that the covariance matrix is positive semidefinite.

Example 9. *Let's compute the covariance matrix in the case when X and Y are two coin tosses, as in Example 8. We will assign X and Y the value -1 if the coin falls on tails and 1 if the coin*

falls on heads. Clearly, both coin tosses have mean $\mu = \nu = 0$.

In the case of independent coin tosses (probabilities as in (279)) we have:

$$K = \frac{1}{4} \underbrace{\begin{bmatrix} 1 \\ 1 \end{bmatrix} \begin{bmatrix} 1 & 1 \end{bmatrix}}_{\text{outcome is HH}} + \frac{1}{4} \underbrace{\begin{bmatrix} -1 \\ -1 \end{bmatrix} \begin{bmatrix} -1 & -1 \end{bmatrix}}_{\text{outcome is TT}} + \frac{1}{4} \underbrace{\begin{bmatrix} 1 \\ -1 \end{bmatrix} \begin{bmatrix} 1 & -1 \end{bmatrix}}_{\text{outcome is HT}} + \frac{1}{4} \underbrace{\begin{bmatrix} -1 \\ 1 \end{bmatrix} \begin{bmatrix} -1 & 1 \end{bmatrix}}_{\text{outcome is TH}} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$$

The covariance matrix is diagonal because the random variables are independent, hence uncorrelated.

In the case of identical coin tosses (probabilities as in (280)) we have:

$$K = \frac{1}{2} \underbrace{\begin{bmatrix} 1 \\ 1 \end{bmatrix} \begin{bmatrix} 1 & 1 \end{bmatrix}}_{\text{outcome is HH}} + \frac{1}{2} \underbrace{\begin{bmatrix} -1 \\ -1 \end{bmatrix} \begin{bmatrix} -1 & -1 \end{bmatrix}}_{\text{outcome is TT}} + 0 \underbrace{\begin{bmatrix} 1 \\ -1 \end{bmatrix} \begin{bmatrix} 1 & -1 \end{bmatrix}}_{\text{outcome is HT}} + 0 \underbrace{\begin{bmatrix} -1 \\ 1 \end{bmatrix} \begin{bmatrix} -1 & 1 \end{bmatrix}}_{\text{outcome is TH}} = \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix}$$

The covariance matrix has determinant 0 because identical random variables are perfectly correlated.

Lecture 31 (May 10)

Reading: Section 12.2

Viewing: none

Let us generalize the previous setup to an arbitrary number of random variables.

Definition 28. Given n random variables X_1, \dots, X_n , their covariance matrix is:

$$K = \begin{bmatrix} \Sigma_{X_1 X_1} & \cdots & \Sigma_{X_1 X_n} \\ \vdots & \ddots & \vdots \\ \Sigma_{X_n X_1} & \cdots & \Sigma_{X_n X_n} \end{bmatrix} \quad (288)$$

The covariance matrix is a useful tool for determining how (un)correlated a collection of random variables is, for example if $X_1 = \text{age}$, $X_2 = \text{height}$, $X_3 = \text{weight}$ among a population. In fact, it is often convenient to put all of the random variables in a vector, called a **random vector**:

$$\mathbf{X} = \begin{bmatrix} X_1 \\ \vdots \\ X_n \end{bmatrix} \quad (289)$$

Let p_1, p_2, \dots to be the probability that \mathbf{X} takes a certain number of vector values $\mathbf{x}_1, \mathbf{x}_2, \dots$, where $\mathbf{x}_i = \begin{bmatrix} x_{1,i} \\ \vdots \\ x_{n,i} \end{bmatrix}$. In the setting of two random variables in (278), the probabilities in question

were denoted by $p_{ij} = \text{Prob} \left(\begin{bmatrix} X \\ Y \end{bmatrix} = \begin{bmatrix} x_i \\ y_j \end{bmatrix} \right)$. Then the mean of the random vector (289) is:

$$\boldsymbol{\mu} = \sum_i p_i \mathbf{x}_i$$

(note that $\boldsymbol{\mu}$ is a vector) and the covariance matrix is:

$$\begin{aligned} K &= \sum_i p_i \begin{bmatrix} (x_{1,i} - \mu_1)^2 & \cdots & (x_{1,i} - \mu_1)(x_{n,i} - \mu_n) \\ \vdots & \ddots & \vdots \\ (x_{1,i} - \mu_1)(x_{n,i} - \mu_n) & \cdots & (x_{n,i} - \mu_n)^2 \end{bmatrix} = \\ &= \sum_i p_i \begin{bmatrix} x_{1,i} - \mu_1 \\ \vdots \\ x_{n,i} - \mu_n \end{bmatrix} \begin{bmatrix} x_{1,i} - \mu_1 \\ \vdots \\ x_{n,i} - \mu_n \end{bmatrix}^T = \sum_i p_i (\mathbf{x}_i - \boldsymbol{\mu})(\mathbf{x}_i - \boldsymbol{\mu})^T = E[(\mathbf{X} - \boldsymbol{\mu})(\mathbf{X} - \boldsymbol{\mu})^T] \end{aligned} \quad (290)$$

Apart from the fact that vector and matrix notation gives the neat formula (290) for the covariance matrix, the equality above proves that the covariance matrix of any number of random variables is positive semidefinite. Moreover, any linear combination of the random variables X_1, \dots, X_n :

$$X = \mathbf{c}^T \mathbf{X} \quad \text{for some} \quad \mathbf{c} = \begin{bmatrix} c_1 \\ \vdots \\ c_n \end{bmatrix}$$

is a random variable (for example, in the age–height–weight example, X could be $2 \cdot \text{age} - 3 \cdot \text{height} + 7 \cdot \text{weight}$). The variance of X can be recovered from the covariance matrix K as:

$$\begin{aligned} E[(X - \mu)^2] &= E[(X - \mu)(X - \mu)^T] = E[(\mathbf{c}^T \mathbf{X} - \mathbf{c}^T \boldsymbol{\mu})(\mathbf{c}^T \mathbf{X} - \mathbf{c}^T \boldsymbol{\mu})^T] = \\ &= E[(\mathbf{c}^T \mathbf{X} - \mathbf{c}^T \boldsymbol{\mu})(\mathbf{X}^T \mathbf{c} - \boldsymbol{\mu}^T \mathbf{c})] = E[\mathbf{c}^T (\mathbf{X} - \boldsymbol{\mu})(\mathbf{X}^T - \boldsymbol{\mu}^T) \mathbf{c}] = \mathbf{c}^T K \mathbf{c} \end{aligned} \quad (291)$$

which is precisely the energy of the vector \mathbf{c} with respect to the positive semidefinite covariance matrix K . The matrix K is positive definite unless the energy is 0 for a non-zero vector \mathbf{c} . Hence:

the covariance matrix is positive definite, unless

the random variables X_1, \dots, X_n are linearly dependent

where “linearly dependent” means that there exists some linear combination of X_1, \dots, X_n which is constant (since the constant random variables are precisely those with 0 variance).

Example 10. *Here’s a quick application of this matrix formalism. Suppose you have random variables X and Y , with means μ and ν , and variances Σ_{XX} and Σ_{YY} . What are the mean and variance of the random variable $Z = X + Y$? You can compute these by writing:*

$$Z = \begin{bmatrix} 1 & 1 \end{bmatrix} \begin{bmatrix} X \\ Y \end{bmatrix}$$

The mean of Z is just:

$$\rho = \begin{bmatrix} 1 & 1 \end{bmatrix} \begin{bmatrix} \mu \\ \nu \end{bmatrix} = \mu + \nu$$

which is perhaps not so surprising. However, the variance of Z is given by formula (291):

$$\Sigma_{ZZ} = \begin{bmatrix} 1 & 1 \end{bmatrix} \begin{bmatrix} \Sigma_{XX} & \Sigma_{XY} \\ \Sigma_{XY} & \Sigma_{YY} \end{bmatrix} \begin{bmatrix} 1 \\ 1 \end{bmatrix} = \Sigma_{XX} + \Sigma_{YY} + 2\Sigma_{XY}$$

Thus, the variance of a sum of random variables involves not only the sum of the individual variances, but also a cross-term involving the covariance of the two variables.

In this language, we can ask what is the meaning of diagonalizing the covariance matrix:

$$K = QDQ^T \quad (292)$$

where D is a diagonal matrix with non-negative eigenvalues, and Q has orthonormal columns. This procedure is called **principal component analysis (PCA)** in probability and statistics. Mathematically, this entails defining the following random vector:

$$\begin{bmatrix} Y_1 \\ \vdots \\ Y_n \end{bmatrix} = \mathbf{Y} := Q^T \mathbf{X} \quad \text{whose mean will be} \quad \boldsymbol{\nu} = Q^T \boldsymbol{\mu} \quad (293)$$

This simply means that the entries Y_1, \dots, Y_n of \mathbf{Y} are all random variables, given by certain linear combinations of the random variables that make up \mathbf{X} . In this case, the covariance matrix corresponding to the random vector \mathbf{Y} is:

$$E[(\mathbf{Y} - \boldsymbol{\nu})(\mathbf{Y} - \boldsymbol{\nu})^T] = E[Q^T(\mathbf{X} - \boldsymbol{\mu})(\mathbf{X} - \boldsymbol{\mu})^T Q] = Q^T K Q = D = \begin{bmatrix} \Sigma_1 & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & \Sigma_n \end{bmatrix}$$

Since D is a diagonal matrix, this implies that the random variables Y_1, \dots, Y_n are mutually uncorrelated (since their pairwise covariances are 0), and that Σ_i is the variance of Y_i . Therefore, the diagonalization (292) corresponds to finding n linear combinations of the random variables X_1, \dots, X_n , denoted by Y_1, \dots, Y_n above, which are pairwise uncorrelated random variables.

Lecture 32 (May 12)

Reading: Section 12.3

Viewing: none

Last time, we discussed random vectors and covariance matrices in the discrete case, but you probably won't be surprised to hear that the formalism can be extrapolated to continuous probability distributions. Generalizing (282), a function:

$$p(x_1, \dots, x_n) : \mathbb{R}^n \rightarrow \mathbb{R}_{\geq 0} \quad (294)$$

is called the **joint probability distribution** of a random vector \mathbf{X} such as (289). The covariance matrix of this random vector is given by (290), if one replaces the sum by an integral:

$$K = E[(\mathbf{X} - \boldsymbol{\mu})(\mathbf{X} - \boldsymbol{\mu})^T] = \int_{-\infty}^{\infty} \dots \int_{-\infty}^{\infty} p(x_1, \dots, x_n) \begin{bmatrix} x_1 - \mu_1 \\ \vdots \\ x_n - \mu_n \end{bmatrix} \begin{bmatrix} x_1 - \mu_1 \\ \vdots \\ x_n - \mu_n \end{bmatrix}^T dx_1 \dots dx_n \quad (295)$$

where:

$$\boldsymbol{\mu} = \begin{bmatrix} \mu_1 \\ \vdots \\ \mu_n \end{bmatrix} = \int_{-\infty}^{\infty} \dots \int_{-\infty}^{\infty} p(x_1, \dots, x_n) \begin{bmatrix} x_1 \\ \vdots \\ x_n \end{bmatrix} dx_1 \dots dx_n \quad (296)$$

A rather simple case is a random vector \mathbf{Y} consisting of n independent normal distributions Y_1, \dots, Y_n , i.e. consider the following joint probability distribution:

$$p(y_1, \dots, y_n) = \frac{e^{-\frac{(y_1 - \nu_1)^2}{2\Sigma_1} - \dots - \frac{(y_n - \nu_n)^2}{2\Sigma_n}}}{\sqrt{(2\pi)^n \Sigma_1 \dots \Sigma_n}} \quad (297)$$

(we denote them by y 's instead of x 's to avoid confusion later on). As a consequence of the identity:

$$-\frac{(y_1 - \nu_1)^2}{2\Sigma_1} - \dots - \frac{(y_n - \nu_n)^2}{2\Sigma_n} = -\frac{1}{2} \begin{bmatrix} y_1 - \nu_1 \\ \vdots \\ y_n - \nu_n \end{bmatrix}^T D^{-1} \begin{bmatrix} y_1 - \nu_1 \\ \vdots \\ y_n - \nu_n \end{bmatrix}, \quad \text{where } D = \begin{bmatrix} \Sigma_1 & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & \Sigma_n \end{bmatrix}$$

we may write formula (295) as:

$$D = \int_{-\infty}^{\infty} \dots \int_{-\infty}^{\infty} \frac{e^{-\frac{1}{2}(\mathbf{y} - \boldsymbol{\nu})^T D^{-1}(\mathbf{y} - \boldsymbol{\nu})}}{\sqrt{(2\pi)^n \det D}} (\mathbf{y} - \boldsymbol{\nu})(\mathbf{y} - \boldsymbol{\nu})^T d\mathbf{y} \quad (298)$$

where $\mathbf{y} = \begin{bmatrix} y_1 \\ \vdots \\ y_n \end{bmatrix}$ and $d\mathbf{y} = dy_1 \dots dy_n$. Indeed, this is because the off-diagonal terms of the covariance matrix are 0 due to independence, while the diagonal terms are the variances of Y_1, \dots, Y_n .

Now let's consider a general vector \mathbf{X} of normal distributions X_1, \dots, X_n . The setup here is that we consider an arbitrary $\boldsymbol{\mu} \in \mathbb{R}^n$ and an arbitrary positive definite matrix S , and let:

$$p(x_1, \dots, x_n) = \frac{e^{-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})^T S^{-1}(\mathbf{x} - \boldsymbol{\mu})}}{\sqrt{(2\pi)^n \det S}} \quad (299)$$

Remark. The reason why we require S to be positive definite is that the quantity:

$$E = \frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})^T S^{-1}(\mathbf{x} - \boldsymbol{\mu})$$

can be interpreted as the kinetic energy of some physical system of particles with velocities x_1, \dots, x_n (the fact that we subtract the average $\boldsymbol{\mu}$ means that we choose our coordinates so that the center of mass of the system stands still). The expression e^{-E} is called the Boltzmann factor in physics. Its presence means that velocity vectors (x_1, \dots, x_n) with lower energy arise with higher probability and those with higher energy arise with lower probability.

In the situation at hand, the covariance matrix (295) can be written as:

$$K = \int_{-\infty}^{\infty} \dots \int_{-\infty}^{\infty} \frac{e^{-\frac{1}{2}(\mathbf{x}-\boldsymbol{\mu})^T S^{-1}(\mathbf{x}-\boldsymbol{\mu})}}{\sqrt{(2\pi)^n \det S}} (\mathbf{x} - \boldsymbol{\mu})(\mathbf{x} - \boldsymbol{\mu})^T d\mathbf{x} \quad (300)$$

To summarize: the entries of the random vector $\mathbf{X} = \begin{bmatrix} X_1 \\ \vdots \\ X_n \end{bmatrix}$ whose covariance matrix is K are no

longer independent. However, we can find appropriate linear combinations which are independent, by diagonalizing the matrices S and K . This is PCA for continuous probabilities, and it goes something like this. First diagonalize the positive definite matrix S that appears in (299):

$$S = QDQ^T$$

where Q is orthogonal and D is diagonal (with positive diagonal entries). Then if we perform the substitution $\mathbf{Y} = Q^T \mathbf{X} \Leftrightarrow \mathbf{X} = Q\mathbf{Y}$ (note that the mean of \mathbf{Y} is $\boldsymbol{\nu} = Q^T \boldsymbol{\mu}$), then (300) becomes:

$$\begin{aligned} K &= \int_{-\infty}^{\infty} \dots \int_{-\infty}^{\infty} \frac{e^{-\frac{1}{2}(\mathbf{y}-\boldsymbol{\nu})^T D^{-1}(\mathbf{y}-\boldsymbol{\nu})}}{\sqrt{(2\pi)^n \det S}} Q(\mathbf{y} - \boldsymbol{\nu})(\mathbf{y} - \boldsymbol{\nu})^T Q^T d\mathbf{y} = \\ &= Q \left[\int_{-\infty}^{\infty} \dots \int_{-\infty}^{\infty} \frac{e^{-\frac{1}{2}(\mathbf{y}-\boldsymbol{\nu})^T D^{-1}(\mathbf{y}-\boldsymbol{\nu})}}{\sqrt{(2\pi)^n \det D}} (\mathbf{y} - \boldsymbol{\nu})(\mathbf{y} - \boldsymbol{\nu})^T d\mathbf{y} \right] Q^T \stackrel{(298)}{=} QDQ^T = S \end{aligned}$$

(we are using $\det S = \det D$ and the fact that the determinant of the orthogonal matrix Q is ± 1).

Since the covariance matrix of the random vector $\mathbf{Y} = \begin{bmatrix} Y_1 \\ \vdots \\ Y_n \end{bmatrix}$ is diagonal, the random variables

Y_1, \dots, Y_n are uncorrelated normal distributions. Hence under orthogonal linear transformations Q , random vectors transform as $\mathbf{X} = Q\mathbf{Y}$, and their covariance matrix transforms as $K = QDQ^T$.

The principles of least squares also find their place in probability. Suppose A is a matrix, that b_1, \dots, b_m are random variables with variance $\Sigma_1, \dots, \Sigma_m$, and you are trying to minimize the error in the least squares approximation to the linear system:

$$A\mathbf{v} \sim \mathbf{b} = \begin{bmatrix} b_1 \\ \vdots \\ b_m \end{bmatrix} \quad (301)$$

The appropriate notion of error in the experiment is:

$$\varepsilon = \sum_{i=1}^n \frac{(\mathbf{b} - A\mathbf{v})_i^2}{\Sigma_i} \quad (302)$$

Indeed, if all the variances $\Sigma_1, \dots, \Sigma_n$ were equal, then the error would be a multiple of the square of the length of the vector $\mathbf{b} - A\mathbf{v}$, so we would be solving the ordinary least squares problem. But if the Σ_i 's are different, this means that we weigh the error by the inverse of the variance: those measurements b_i with smaller (respectively larger) variance are more (respectively less) reliable, hence they will weight more (respectively less) in the error. Still, the problem of minimizing (302) is actually equivalent with the ordinary least squares problem for the related system:

$$A'\mathbf{v} \sim \mathbf{b}', \quad \text{where } A' = D^{-\frac{1}{2}}A, \quad \mathbf{b}' = D^{-\frac{1}{2}}\mathbf{b} \quad \text{and} \quad D = \begin{bmatrix} \Sigma_1 & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & \Sigma_n \end{bmatrix}$$

Therefore, the least squares solution reads:

$$\mathbf{v} = (A'^T A')^{-1} A'^T \mathbf{b}' = (A^T D^{-1} A)^{-1} A^T D^{-1} \mathbf{b} \quad (303)$$

In terms of the system (301), this is known as **weighted least squares**. The most general scenario is when the random variables b_1, \dots, b_n are not independent, and they are related by the covariance matrix K . In this case, the appropriate analogue of the solution (303) is:

$$\mathbf{v} = (A^T K^{-1} A)^{-1} A^T K^{-1} \mathbf{b} \quad (304)$$

and it is known as **generalized least squares**. It is precisely the choice of \mathbf{v} which minimizes the error, whose appropriate definition is:

$$\varepsilon = (\mathbf{b} - A\mathbf{v})^T K^{-1} (\mathbf{b} - A\mathbf{v})$$

Another question you might ask is what is the covariance matrix of the solution (304) in terms of the covariance matrix $K = E[\mathbf{b}\mathbf{b}^T]$. Assume that the random vector \mathbf{b} has zero mean, to keep notations simple (it is reasonable to expect that measuring errors will have zero mean). Then:

$$\begin{aligned} E[\mathbf{v}\mathbf{v}^T] &= (A^T K^{-1} A)^{-1} A^T K^{-1} \underbrace{E[\mathbf{b}\mathbf{b}^T]}_K K^{-1,T} A (A^T K^{-1,T} A)^{-1} = \\ &= (A^T K^{-1} A)^{-1} A^T K^{-1,T} A (A^T K^{-1,T} A)^{-1} = (A^T K^{-1} A)^{-1} \end{aligned}$$

I hope this has been a convincing introduction (although we have just scratched the surface) of how linear algebra is useful in probability.

Lecture 33 (May 14)

Reading: none

Viewing: none

Statistics is the sibling of probability, but they differ slightly in their goals. While probability anticipates the likelihood of future events, statistics interprets data from past events. Let's say we are performing an experiment that measures a certain quantity. We run the experiment n times and we get the following values, called **samples**:

$$x_1, \dots, x_n \quad (305)$$

The collection of samples is often called the **data set**.

Definition 29. The *mean* of the data set (305) is:

$$\mu = \frac{1}{n}(x_1 + \dots + x_n) \quad (306)$$

The *variance* of the data set (305) is:

$$\Sigma = \frac{1}{n-1} \left[(x_1 - \mu)^2 + \dots + (x_n - \mu)^2 \right] \quad (307)$$

The *standard deviation* is the square root of the variance: $\sigma = \sqrt{\Sigma}$.

As $n \rightarrow \infty$, the quantities (306) and (307) tend to the mean and variance of the underlying probability distribution of whatever we're measuring. For example, if you're measuring a quantity which naturally obeys a normal distribution (e.g. the numeric grades in a large class), then the graph of the function “ $f(x)$ = number of students who got grade x ” will look like a bell curve.

Remark. You might wonder why we're using the denominator $n - 1$ in (307) instead of n , which might seem much more appropriate from a probability point of view. The discrepancy, namely $\frac{n}{n-1}$ is called **Bessel's correction**, and we will explain it from the point of view of linear algebra later on. It is not so significant when n is large, but it makes a difference when n is small.

Explicitly, the mean is just measuring the average of the samples in question, while the variance is measuring how far the samples are from their mean. These notions can be presented in the language of linear algebra, by introducing the vectors:

$$\mathbf{o} = \begin{bmatrix} 1 \\ \vdots \\ 1 \end{bmatrix} \quad \text{and} \quad \mathbf{x} = \begin{bmatrix} x_1 \\ \vdots \\ x_n \end{bmatrix}$$

Then the mean is equal to:

$$\mu = \frac{\mathbf{o}^T \mathbf{x}}{\mathbf{o}^T \mathbf{o}} \quad (308)$$

and thus $\mathbf{o}\mu$ can be interpreted as the projection of the vector of samples \mathbf{x} onto the vector \mathbf{o} . Meanwhile, the variance is:

$$\Sigma = \frac{\|\mathbf{x} - \mathbf{o}\mu\|^2}{n-1} = \frac{\|P\mathbf{x}\|^2}{n-1} \quad (309)$$

where $P = I - \frac{\mathbf{o}\mathbf{o}^T}{\mathbf{o}^T\mathbf{o}}$ is the projection matrix onto the orthogonal complement to \mathbf{o} . This orthogonal complement is an $n - 1$ dimensional vector space, and indeed we have:

$$P = \begin{bmatrix} \frac{n-1}{n} & \cdots & -\frac{1}{n} \\ \vdots & \ddots & \vdots \\ -\frac{1}{n} & \cdots & \frac{n-1}{n} \end{bmatrix} \quad (310)$$

It is elementary to show that the diagonalization of P is:

$$P = QDQ^T, \quad \text{where } Q \text{ is orthogonal and } D = \begin{bmatrix} 1 & \cdots & 0 & 0 \\ \vdots & \ddots & \vdots & \vdots \\ 0 & \cdots & 1 & 0 \\ 0 & \cdots & 0 & 0 \end{bmatrix} \quad (311)$$

Therefore, if the n entries of the sample vector \mathbf{x} are drawn from a normal probability distribution with mean 0 and variance 1, then so will the entries of the vector $\mathbf{y} = Q^T\mathbf{x}$ (this is because multiplying by orthogonal matrices does not change length). But we have:

$$\|P\mathbf{x}\|^2 = \|QDQ^T\mathbf{x}\|^2 = \|D\mathbf{y}\|^2 = y_1^2 + \cdots + y_{n-1}^2 \quad (312)$$

and so the variance of (312) is $n - 1$. We have just shown that the variance of the quantity $\|P\mathbf{x}\|^2$ is $\frac{n-1}{n}$ times that of the quantity $\|\mathbf{x}\|^2$, and Bessel's correction is simply a way to correct for this.

Now assume we have two data sets: x_1, \dots, x_n and y_1, \dots, y_n (say age and height) and we want to see how correlated they are. As in the case of probability, the appropriate notion is the **covariance**:

$$\Sigma_{xy} = \frac{1}{n-1} \left[(x_1 - \mu)(y_1 - \nu) + \cdots + (x_n - \mu)(y_n - \nu) \right] \quad (313)$$

where μ and ν are the means of the samples x_1, \dots, x_n and y_1, \dots, y_n , respectively. Note that Σ_{xx} is just the variance of the sample set x_1, \dots, x_n , so the reason why we use the denominator $n - 1$ instead of n in (313) is still Bessel's correction, as explained above. In terms of vectors:

$$\mathbf{x} = \begin{bmatrix} x_1 \\ \vdots \\ x_n \end{bmatrix} \quad \text{and} \quad \mathbf{y} = \begin{bmatrix} y_1 \\ \vdots \\ y_n \end{bmatrix}$$

the covariance takes the form:

$$\Sigma_{xy} = \frac{\mathbf{x}^T P \mathbf{y}}{n-1} \quad (314)$$

since $P^2 = P$, $P^T = P$ (being a projection matrix) and $P\mathbf{x} = \begin{bmatrix} x_1 - \mu \\ \vdots \\ x_n - \mu \end{bmatrix}$, $P\mathbf{y} = \begin{bmatrix} y_1 - \nu \\ \vdots \\ y_n - \nu \end{bmatrix}$. Also:

$$|\Sigma_{xy}| \leq \sqrt{\Sigma_{xx}\Sigma_{yy}} \quad (315)$$

just like in the case of probability. The covariance is greatest (or lowest) when x and y are the most correlated (anti-correlated). If the two data sets are independent, then you would expect the

covariance to be very close to 0, and actually tend to 0 as you take more and more samples $n \rightarrow \infty$.

Finally, assume we have m data sets, so let's put them in an $n \times m$ matrix:

$$\mathbf{A} = \begin{bmatrix} x_1 & y_1 & z_1 & \dots \\ \vdots & \vdots & \vdots & \ddots \\ x_n & y_n & z_n & \dots \end{bmatrix}$$

If the x 's, y 's, z 's ... have means $\mu, \nu, \xi \dots$, then:

$$P\mathbf{A} = \begin{bmatrix} x_1 - \mu & y_1 - \nu & z_1 - \xi & \dots \\ \vdots & \vdots & \vdots & \ddots \\ x_n - \mu & y_n - \nu & z_n - \xi & \dots \end{bmatrix} \quad (316)$$

The **covariance matrix** of the data sets is the $m \times m$ matrix defined by:

$$K = \begin{bmatrix} \Sigma_{xx} & \Sigma_{xy} & \Sigma_{xz} & \dots \\ \Sigma_{yx} & \Sigma_{yy} & \Sigma_{yz} & \dots \\ \Sigma_{zx} & \Sigma_{zy} & \Sigma_{zz} & \dots \\ \vdots & \vdots & \vdots & \ddots \end{bmatrix} \quad (317)$$

and is given by the formula:

$$K = \frac{\mathbf{A}^T P \mathbf{A}}{n - 1} \quad (318)$$

To see this, just note for instance that the (1, 2) entry of the matrix identity (318) is precisely (314).

Principal component analysis (PCA) involves diagonalizing the covariance matrix. If:

$$K = QDQ^T$$

with Q orthogonal and D diagonal (with positive diagonal entries) then (318) reads:

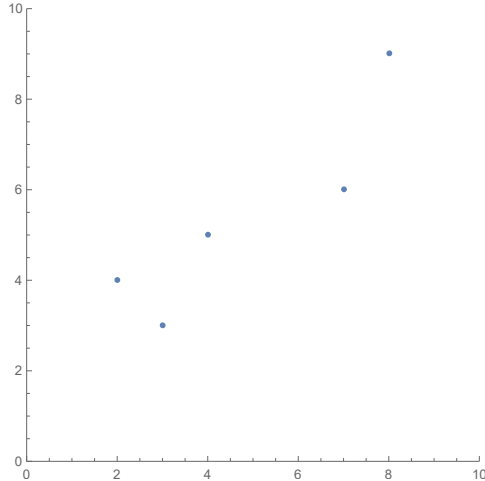
$$QDQ^T = \frac{\mathbf{A}^T P \mathbf{A}}{n - 1} \Rightarrow D = \frac{(\mathbf{A}Q)^T P (\mathbf{A}Q)}{n - 1}$$

This means that the data set $\mathbf{B} = \mathbf{A}Q$ (whose columns are linear combinations of the constituent data sets of \mathbf{A}) has diagonal covariance. Therefore, the constituent data sets of \mathbf{B} are mutually uncorrelated. Moreover, if the i -th diagonal entry of D is the largest (respectively smallest) then the i -th constituent data set of \mathbf{B} has the largest (respectively smallest) variance. Therefore, this approach allows us to isolate those linear combinations in a collection of data sets which produce the most variance from those which produce the least variance. For example, consider the plot:

of two data sets: the entries of $\mathbf{x} = \begin{bmatrix} x_1 \\ \vdots \\ x_5 \end{bmatrix}$ are on the horizontal axis and the entries of $\mathbf{y} = \begin{bmatrix} x_1 \\ \vdots \\ x_5 \end{bmatrix}$

are on the vertical axis. Explicitly, they are given by

$$\mathbf{x} = \begin{bmatrix} 2 \\ 3 \\ 4 \\ 7 \\ 8 \end{bmatrix} \quad \text{and} \quad \mathbf{y} = \begin{bmatrix} 4 \\ 3 \\ 5 \\ 6 \\ 9 \end{bmatrix}$$



and suppose we wish to compute their covariance matrix, and to apply PCA to them (i.e. to find linear combinations of the data sets which are mutually uncorrelated, and to compute the variance of those). Start by constructing the matrix:

$$\mathbf{A} = [\mathbf{x} \mid \mathbf{y}] = \begin{bmatrix} 2 & 4 \\ 3 & 3 \\ 4 & 5 \\ 7 & 6 \\ 8 & 9 \end{bmatrix}$$

Then compute:

$$P\mathbf{A} = \frac{1}{5} \begin{bmatrix} -14 & -7 \\ -9 & -12 \\ -4 & -2 \\ 11 & 3 \\ 16 & 18 \end{bmatrix} \quad (319)$$

Then the covariance matrix is given by formula (318):

$$K = \frac{1}{20} \begin{bmatrix} 134 & 107 \\ 107 & 106 \end{bmatrix}$$

You could perform PCA by diagonalizing the matrix K , or you could achieve the same result by instead computing the SVD of the matrix $P\mathbf{A}$. Let's opt for the latter:

$$P\mathbf{A} = U\Sigma V^T \quad \text{where} \quad \Sigma = \begin{bmatrix} \sqrt{24 + \frac{1}{5}\sqrt{11645}} & 0 \\ 0 & \sqrt{24 - \frac{1}{5}\sqrt{11645}} \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \end{bmatrix}$$

and U and V are orthogonal matrices. Then the covariance matrix is given by:

$$K = \frac{V\Sigma^T U^T U\Sigma V^T}{5-1} = V \begin{bmatrix} 6 + \frac{1}{20}\sqrt{11645} & 0 \\ 0 & 6 - \frac{1}{20}\sqrt{11645} \end{bmatrix} V^T \quad (320)$$

The meaning of this is the following: the entries of the 2×2 orthogonal matrix V teach you how to construct mutually independent combinations of the data sets \mathbf{x} and \mathbf{y} . Explicitly, if:

$$V = \begin{bmatrix} v_{11} & v_{12} \\ v_{21} & v_{22} \end{bmatrix}$$

then the data sets:

$$[\mathbf{x}' \mid \mathbf{y}'] = [\mathbf{x} \mid \mathbf{y}] V \quad \Leftrightarrow \quad \begin{cases} \mathbf{x}' = v_{11}\mathbf{x} + v_{21}\mathbf{y} \\ \mathbf{y}' = v_{12}\mathbf{x} + v_{22}\mathbf{y} \end{cases}$$

are mutually uncorrelated. Moreover, the variances of \mathbf{x}' and \mathbf{y}' are the two numbers on the diagonal of (320), and numerically they are roughly equal to 11.39 and 0.6. This implies that the linear combination \mathbf{y}' of our two data sets has very little variance, and so it is roughly constant. Meanwhile, the linear combination \mathbf{x}' is very “noisy”, since it has rather high variance.

Lecture 34 (May 17)

Reading: Section 10.5

Viewing: none

We will now study Fourier series, which is a very useful application of linear algebra in analysis, and from there in applied mathematics and engineering. But here's the twist: Fourier analysis is not concerned with finite-dimensional vector spaces, but with infinite-dimensional ones.

Definition 30. Let V be the vector space of differentiable functions $f : \mathbb{R} \rightarrow \mathbb{R}$ with period 2π , i.e.:

$$f(x + 2\pi) = f(x) \quad \forall x \in \mathbb{R} \quad (321)$$

Such functions are called **periodic**.

The fact that we only look at periodic functions might seem like a restriction, but it's suitable for many applications (think signals, waves). Also, the fact that the period is 2π rather than any other number is not a meaningful restriction, but more of a notational convention. The reason for this is that if $f(x)$ is a function with period L , then $f(\frac{2\pi x}{L})$ is a function with period 2π .

The fact that the set of f 's as in (321) is a vector space is easy to see, basically because any linear combination of differentiable periodic functions is differentiable and periodic (this entails the fact that all of these functions have the same period, but we have already assumed that this is the case). To convince you that it is not a finite-dimensional vector space, consider the functions:

$$1, \cos x, \cos 2x, \cos 3x, \dots, \sin x, \sin 2x, \sin 3x, \dots \quad (322)$$

We claim that all of the infinitely many functions in (322) are linearly independent: there is no linear combination of these functions which is equal to the zero function⁵. Therefore, the space V is an entirely new beast from the vector spaces \mathbb{R}^n that we have studied in most of our course. However, it does share a lot of features with finite-dimensional vector spaces, such as the existence of a dot product like operation, namely the **inner product** of functions:

$$(f, g) = \int_{-\pi}^{\pi} f(x)g(x)dx \quad (323)$$

⁶ In other words, the inner product takes two functions and produces a number, just like the dot product took two vectors and produced a number. And just like the length of a vector is its dot product with itself, we have the **norm** of a function:

$$\|f\|^2 = (f, f) = \int_{-\pi}^{\pi} f(x)^2 dx \quad (324)$$

which is a positive real number (unless $f = 0$, in which case the norm is 0). Here's one important fact: the functions (322) are all mutually orthogonal with respect to the inner product (323):

$$(\sin kx, \sin lx) = 0, \quad \text{unless } k = l, \text{ in which case } (\sin kx, \sin kx) = \pi \quad (325)$$

$$(\cos kx, \cos lx) = 0, \quad \text{unless } k = l, \text{ in which case } (\cos kx, \cos kx) = \pi \quad (326)$$

$$(\sin kx, \cos lx) = 0, \quad \text{for all } k, l \quad (327)$$

⁵Note the emphasis on the word "linear", since there are non-linear combinations of the above functions which are equal to the zero function, for example $(\sin x)^2 + (\cos x)^2 - 1 = 0$. However, when studying vector spaces, all we are concerned with are linear combinations

⁶The fact that the integral goes from $-\pi$ to π is not important. You could have taken it over any interval of length 2π and the value of the integral would have been unchanged, due to the periodicity of the functions f and g

The formulas above also hold if k or l are 0 (except that $\sin 0 = 0$ and $\cos 0 = 1$ imply that $(0, 0) = 0$ and $(1, 1) = 2\pi$). So (322) is a collection of linearly independent, mutually orthogonal elements of the vector space V . If this were a finite-dimensional vector space, we could hope that the collection (322) is also a basis. Well, something pretty close to this actually happens:

Fact 23. *Any function $f \in V$ can be written as a linear combination of the functions (322), but the linear combination has to be infinite. In other words, we can write any function $f \in V$ as:*

$$f(x) = a_0 + a_1 \cos x + a_2 \cos 2x + \cdots + b_1 \sin x + b_2 \sin 2x + \cdots \quad (328)$$

The numbers $a_0, a_1, a_2, \dots, b_1, b_2, \dots$ are called the **Fourier coefficients** of $f(x)$, and the sum (328) is called the **Fourier series** of $f(x)$. It is a theorem that the Fourier series converges.

In signal processing (and music) the functions $\cos kx, \sin kx$ are called **harmonics**, so the Fourier series is basically just a way to decompose an arbitrary periodic signal $f(x)$ into harmonics. As you would expect, in real-life examples the coefficients a_n, b_n tend to 0 as $n \rightarrow \infty$, so only the first few harmonics of a signal dominate it.

Linear algebra also gives us a very practical way to compute the Fourier coefficients of any function. Because the functions (322) are orthogonal, you can compute these coefficients by computing the inner product of f with any one of these functions:

$$(f, 1) = a_0(1, 1) \quad \Rightarrow \quad a_0 = \frac{1}{2\pi} \int_{-\pi}^{\pi} f(x) dx \quad (329)$$

$$(f, \cos kx) = a_k(\cos kx, \cos kx) \quad \Rightarrow \quad a_k = \frac{1}{\pi} \int_{-\pi}^{\pi} f(x)(\cos kx) dx \quad (330)$$

$$(f, \sin kx) = b_k(\sin kx, \sin kx) \quad \Rightarrow \quad b_k = \frac{1}{\pi} \int_{-\pi}^{\pi} f(x)(\sin kx) dx \quad (331)$$

Formulas (329), (330), (331) show that if you have a machine that computes various integrals of f (and there are many ways for getting approximate answers) this immediately gives a recipe for computing the Fourier series of f . Here's an example of this setup. Consider the function:

$$f(x) = x \quad \text{for all } x \in [-\pi, \pi] \quad (332)$$

then extended to all real numbers x so that it has period 2π (note that this function has a discontinuity at $x = \pi$; all the facts in this class still hold in this slightly more general setup). To compute its Fourier series, we need to compute the following integrals:

$$\begin{aligned} (f, \cos kx) &= \int_{-\pi}^{\pi} x(\cos kx) dx = 0 \\ (f, \sin kx) &= \int_{-\pi}^{\pi} x(\sin kx) dx = (-1)^{k-1} \frac{2\pi}{k} \end{aligned}$$

which can be done e.g. by integration by parts. Therefore (329)–(331) imply the Fourier series:

$$f(x) = \frac{2 \sin x}{1} - \frac{2 \sin 2x}{2} + \frac{2 \sin 3x}{3} - \frac{2 \sin 4x}{4} + \cdots \quad (333)$$

The fact that the Fourier series only contains sines in this example should not be surprising; after all, the particular function $f(x)$ in (332) is odd. But then let us compute the norm squared of the function (332). On one hand, we have:

$$(f, f) = \int_{-\pi}^{\pi} x^2 dx = \frac{2\pi^3}{3}$$

On the other hand, the norm squared of the right-hand side of (333) is easy to compute, because the functions $\sin kx$ have norm π and pairwise zero inner product:

$$(f, f) = \frac{4\pi}{1^2} + \frac{4\pi}{2^2} + \frac{4\pi}{3^2} + \frac{4\pi}{4^2} + \dots$$

By equating the two expressions above, we arrive at the interesting identity:

$$\frac{\pi^2}{6} = 1 + \frac{1}{2^2} + \frac{1}{3^2} + \frac{1}{4^2} + \dots$$

This (and others like it) is a neat identity involving π , which gives an easy way to approximate the powers of π . Based on the formula above, you could even approximate π using pen and paper (assuming you remember the algorithm for taking square roots by hand).

Let's now discuss a variant of Fourier series, wherein we replace the vector space V from Definition 30 by the vector space of differentiable functions with complex (instead of real) values:

$$f : \mathbb{R} \rightarrow \mathbb{C}$$

with period 2π . Because complex numbers are of the form $a + bi$ where a, b are real numbers, such complex-valued functions are of the form $a(x) + b(x)i$ where $a(x), b(x)$ are real-valued functions. Much of the theory concerning Fourier series carries through in the present setup, but in a sense it's simpler and more elegant. This is because instead of working with the particular real-valued functions (322), you can work with the complex-valued functions:

$$e^{ikx} \quad \text{for all integers } k \tag{334}$$

In fact, the functions in (322) are merely the real/imaginary parts of the functions in (334). In the complex-valued setting, the definition of inner product and norm from (323) should be changed to:

$$(f, g) = \int_{-\pi}^{\pi} f(x)\overline{g(x)}dx \tag{335}$$

where \bar{z} denotes the complex conjugate of the complex number z . The reason for this is that you want the norm of any function to be a non-negative real number:

$$\|f\|^2 = \int_{-\pi}^{\pi} |f(x)|^2 dx \geq 0 \tag{336}$$

and you would not have been able to achieve this without the conjugate in (335) (what I'm saying is that if z is an arbitrary complex number, z^2 is also a complex number, but $z\bar{z} = |z|^2$ is a non-negative real number). It is then easy to compute the inner products of the functions (334):

$$(e^{ikx}, e^{ilx}) = \int_{-\pi}^{\pi} e^{ikx}\overline{e^{ilx}}dx = \int_{-\pi}^{\pi} e^{i(k-l)x}dx = \begin{cases} \frac{e^{i(k-l)x}}{i(k-l)} \Big|_{-\pi}^{\pi} = 0 & \text{if } k \neq l \\ 2\pi & \text{if } k = l \end{cases}$$

So the functions (334) are pairwise orthogonal. The complex-valued analogue of Fact 23 is that any differentiable complex-valued periodic function has a Fourier series of the form:

$$f(x) = \sum_{k \in \mathbb{Z}} c_k e^{ikx} \quad (337)$$

where the Fourier coefficients are given by:

$$c_k = \frac{1}{2\pi} \int_{-\pi}^{\pi} f(x) e^{-ikx} dx \quad (338)$$

The advantage of complex-valued Fourier series over real-valued ones is clear: the Fourier series (337) is a sum of only one type of function, as opposed from two types of functions in (328). Moreover, you only need to remember the single formula (338) for the Fourier coefficients in the complex-valued case, instead of formulas (329), (330) and (331) in the real-valued case. Finally, since any real-valued function can also be thought of as a complex valued function, it still has a Fourier series of the form (337) (it's just that you will have $\overline{c_k} = c_{-k}$, in order for the imaginary parts of various complex exponentials to cancel each other out).

Lecture 35 (*May 19*)

Review and more examples on the previous class.