

Review: Matrix calculus

Example 1. Matrix multiplication is not commutative!

$$\bullet \begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix} \cdot \begin{bmatrix} 1 & 2 \\ 0 & 1 \end{bmatrix} = \begin{bmatrix} 1 & 4 \\ 3 & 10 \end{bmatrix}$$

Multiplication (on the right) with that “almost identity matrix” is performing the column operation $C_2 + 2C_1 \Rightarrow C_2$ (i.e. 2 times the first column is added to the second column).

$$\bullet \begin{bmatrix} 1 & 2 \\ 0 & 1 \end{bmatrix} \cdot \begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix} = \begin{bmatrix} 7 & 10 \\ 3 & 4 \end{bmatrix}$$

Multiplication (on the left) with the same matrix is performing the row operation $R_1 + 2R_2 \Rightarrow R_1$.

First comment. This indicates a second interpretation of matrix multiplication: instead of taking linear combinations of columns of the first matrix, we can also take linear combinations of rows of the second matrix.

Second comment. The row operations we are doing during Gaussian elimination can be realized by multiplying (on the left) with “almost identity matrices”.

Example 2. $\begin{bmatrix} 1 & 2 & 3 \end{bmatrix} \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix} = \begin{bmatrix} 14 \end{bmatrix}$ whereas $\begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix} \begin{bmatrix} 1 & 2 & 3 \end{bmatrix} = \begin{bmatrix} 1 & 2 & 3 \\ 2 & 4 & 6 \\ 3 & 6 & 9 \end{bmatrix}$.

If you know about the dot product, do you see a connection with the first case?

Example 3. Suppose A is $m \times n$ and B is $p \times q$. When does AB make sense? In that case, what are the dimensions of AB ?

AB makes sense if $n = p$. In that case, AB is a $m \times q$ matrix.

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

On the RHS we have the **identity matrix**, usually denoted I or I_2 (since it's the 2×2 identity matrix here).

Hence, the two matrices on the left are inverses of each other: $\begin{bmatrix} 3 & 1 \\ 2 & 1 \end{bmatrix}^{-1} = \begin{bmatrix} 1 & -1 \\ -2 & 3 \end{bmatrix}$, $\begin{bmatrix} 1 & -1 \\ -2 & 3 \end{bmatrix}^{-1} = \begin{bmatrix} 3 & 1 \\ 2 & 1 \end{bmatrix}$.

Example 5. The following formula immediately gives us the inverse of a 2×2 matrix (if it exists). It is worth remembering!

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix}^{-1} = \frac{1}{ad - bc} \begin{bmatrix} d & -b \\ -c & a \end{bmatrix} \quad \text{provided that } ad - bc \neq 0$$

Let's check that! $\frac{1}{ad - bc} \begin{bmatrix} d & -b \\ -c & a \end{bmatrix} \begin{bmatrix} a & b \\ c & d \end{bmatrix} = \frac{1}{ad - bc} \begin{bmatrix} ad - bc & 0 \\ 0 & -cb + ad \end{bmatrix} = I_2$

In particular, a 2×2 matrix $\begin{bmatrix} a & b \\ c & d \end{bmatrix}$ is invertible $\iff ad - bc \neq 0$.

Recall that this is the **determinant**: $\det\left(\begin{bmatrix} a & b \\ c & d \end{bmatrix}\right) = ad - bc$.

In particular:

$$\det(A) = 0 \iff A \text{ is not invertible}$$

Similarly, for $n \times n$ matrices A :

A is invertible	(i.e. there is a matrix A^{-1} such that $AA^{-1} = I$)
$\iff \det(A) \neq 0$	
$\iff Ax = b$ has a unique solution	(namely, $x = A^{-1}b$)

Comment. Why is it not common to write $\frac{1}{A}$ instead of A^{-1} ?

The notation $\frac{1}{A}$ easily leads to ambiguities: for instance, should $\frac{B}{A}$ mean BA^{-1} or should it mean $A^{-1}B$?

[Of course, one could try to avoid this by notations like B/A which would more clearly mean BA^{-1} . It's just not common and doesn't have any real advantages.]

Example 6.

$$\begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ -4 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} = \begin{bmatrix} -7 & 2 & 3 \\ -16 & 5 & 6 \\ -25 & 8 & 9 \end{bmatrix}$$

Multiplication (on the right) with that "almost identity matrix" is performing the column operation $C_1 - 4C_2 \Rightarrow C_1$ (i.e. -4 times the second column is added to the first column).

$$\begin{bmatrix} 1 & 0 & 0 \\ -4 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{bmatrix} = \begin{bmatrix} 1 & 2 & 3 \\ 0 & -3 & -6 \\ 7 & 8 & 9 \end{bmatrix}$$

Multiplication (on the left) with the same matrix is performing the row operation $R_2 - 4R_1 \Rightarrow R_2$.

Comment (again). The row operations we are doing during Gaussian elimination can all be realized by multiplying (on the left) with "almost identity matrices".

These matrices are called **elementary matrices** (they are obtained by performing a single elementary row operation on an identity matrix).

Elementary matrices are **invertible** because elementary row operations are reversible:

$$\begin{bmatrix} 1 & 0 & 0 \\ 2 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}^{-1} = \begin{bmatrix} 1 & 0 & 0 \\ -2 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad \begin{bmatrix} 1 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 1 \end{bmatrix}^{-1} = \begin{bmatrix} 1 & & \\ & \frac{1}{2} & \\ & & 1 \end{bmatrix}, \quad \begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}^{-1} = \begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

Example 7. Let us do Gaussian elimination on $A = \begin{bmatrix} 2 & 1 \\ 4 & -6 \end{bmatrix}$ until we have an echelon form:

$$A = \begin{bmatrix} 2 & 1 \\ 4 & -6 \end{bmatrix} \xrightarrow{R_2 - 2R_1 \Rightarrow R_2} \begin{bmatrix} 2 & 1 \\ 0 & -8 \end{bmatrix}$$

As last class, the row operation can be encoded by multiplication with an “almost identity matrix” E :

$$\underbrace{\begin{bmatrix} 1 & 0 \\ -2 & 1 \end{bmatrix}}_E \underbrace{\begin{bmatrix} 2 & 1 \\ 4 & -6 \end{bmatrix}}_A = \underbrace{\begin{bmatrix} 2 & 1 \\ 0 & -8 \end{bmatrix}}_U$$

Since $\begin{bmatrix} 1 & 0 \\ -2 & 1 \end{bmatrix}^{-1} = \begin{bmatrix} 1 & 0 \\ 2 & 1 \end{bmatrix}$ (no calculation needed; this is the row operation $R_2 + 2R_1 \Rightarrow R_2$ which reverses our above operation), this means that

$$A = E^{-1}U = \begin{bmatrix} 1 & 0 \\ 2 & 1 \end{bmatrix} \begin{bmatrix} 2 & 1 \\ 0 & -8 \end{bmatrix}.$$

We factored A as the product of a lower and an upper triangular matrix!

$A = LU$ is known as the **LU decomposition** of A .
 L is lower triangular, U is upper triangular.

If A is $m \times n$, then L is an invertible lower triangular $m \times m$ matrix, and U is a usual **echelon form** of A . Every matrix A has a LU decomposition (after possibly swapping some rows of A first).

- The matrix U is just the echelon form of A produced during Gaussian elimination.
- The matrix L can be constructed, entry-by-entry, by simply recording the row operations used during Gaussian elimination. (No extra work needed!)

Example 8. Determine the LU decomposition of $A = \begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix}$.

Solution. $A = \begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix} \xrightarrow{R_2 - 3R_1 \Rightarrow R_2} \begin{bmatrix} 1 & 2 \\ 0 & -2 \end{bmatrix}$ translates into $\begin{bmatrix} 1 & 0 \\ -3 & 1 \end{bmatrix} \begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix} = \begin{bmatrix} 1 & 2 \\ 0 & -2 \end{bmatrix}$.

Since $\begin{bmatrix} 1 & 0 \\ -3 & 1 \end{bmatrix}^{-1} = \begin{bmatrix} 1 & 0 \\ 3 & 1 \end{bmatrix}$ (no calculation needed!), we therefore have $A = \begin{bmatrix} 1 & 0 \\ 3 & 1 \end{bmatrix} \begin{bmatrix} 1 & 2 \\ 0 & -2 \end{bmatrix}$.

Example 9. Determine the LU decomposition of $A = \begin{bmatrix} 1 & 1 & 2 & 1 \\ 3 & 2 & 7 & 2 \\ -2 & 6 & -3 & 1 \end{bmatrix}$.

Solution. We perform Gaussian elimination until we arrive at an echelon form:

$$\begin{bmatrix} 1 & 1 & 2 & 1 \\ 3 & 2 & 7 & 2 \\ -2 & 6 & -3 & 1 \end{bmatrix} \xrightarrow{\substack{R_2 - 3R_1 \Rightarrow R_2 \\ R_3 + 2R_1 \Rightarrow R_3}} \begin{bmatrix} 1 & 1 & 2 & 1 \\ 0 & -1 & 1 & -1 \\ 0 & 8 & 1 & 3 \end{bmatrix} \xrightarrow{R_3 + 8R_2 \Rightarrow R_3} \begin{bmatrix} 1 & 1 & 2 & 1 \\ 0 & -1 & 1 & -1 \\ 0 & 0 & 9 & -5 \end{bmatrix}$$

Observe that we can reverse both of these steps using the row operations $R_2 + 3R_1 \Rightarrow R_2$ and $R_3 - 2R_1 - 8R_2 \Rightarrow R_3$.

Encoding these in L , the corresponding LU decomposition of A is

$$A = LU = \begin{bmatrix} 1 & & & \\ 3 & 1 & & \\ -2 & -8 & 1 & \end{bmatrix} \begin{bmatrix} 1 & 1 & 2 & 1 \\ & -1 & 1 & -1 \\ & & 9 & -5 \end{bmatrix}.$$

Note that no further computation was required to obtain L . (The entries in the matrix L are precisely the (negative) coefficients in the original row operations.)

Comment. By contrast, combining the operations $R_2 - 3R_1 \Rightarrow R_2$ and $R_3 + 8R_2 \Rightarrow R_3$ requires computation.

That is because we change R_2 in the first step, and then use the changed R_2 in the second step. Indeed, note that

$$\begin{bmatrix} 1 & & \\ 0 & 1 & \\ 0 & 8 & 1 \end{bmatrix} \begin{bmatrix} 1 & & \\ -3 & 1 & \\ 2 & 0 & 1 \end{bmatrix} = \begin{bmatrix} 1 & & \\ -3 & 1 & \\ -22 & 8 & 1 \end{bmatrix},$$

so the combined operations are $R_2 - 3R_1 \Rightarrow R_2$ and $R_3 - 22R_1 + 8R_2 \Rightarrow R_3$ (you can also see that directly from the operations).

On the other hand, there was no such complication when combining the reversed operations:

Combining $R_3 - 8R_2 \Rightarrow R_3$ and $R_2 + 3R_1 \Rightarrow R_2$ simply results in $R_2 + 3R_1 \Rightarrow R_2$ and $R_3 - 2R_1 - 8R_2 \Rightarrow R_3$, as used above.

The difference is that, here, we change R_3 in the first step but then don't use the changed R_3 in the second step. In terms of matrix multiplication, we have

$$\begin{bmatrix} 1 & & \\ 3 & 1 & \\ -2 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & & \\ 0 & 1 & \\ 0 & -8 & 1 \end{bmatrix} = \begin{bmatrix} 1 & & \\ 3 & 1 & \\ -2 & -8 & 1 \end{bmatrix},$$

where, because of their special form, the product of the two lower triangular matrices is just "putting together" the entries (unlike in the non-reversed product).

Review. The RREF (row-reduced echelon form) of A is obtained from an echelon form by

- scaling the pivots to 1, and then
- eliminating the entries above the pivots.

A typical RREF has the shape

[* represents an entry that could be anything]

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

Example 10. Let's compute the RREF of the 3×4 matrix from Example 9.

Solution.

$$\begin{bmatrix} 1 & 1 & 2 & 1 \\ 3 & 2 & 7 & 2 \\ -2 & 6 & -3 & 1 \end{bmatrix} \xrightarrow{\substack{R_2 - 3R_1 \Rightarrow R_2 \\ R_3 + 2R_1 \Rightarrow R_3}} \begin{bmatrix} 1 & 1 & 2 & 1 \\ 0 & -1 & 1 & -1 \\ 0 & 8 & 1 & 3 \end{bmatrix} \xrightarrow{R_3 + 8R_2 \Rightarrow R_3} \begin{bmatrix} 1 & 1 & 2 & 1 \\ 0 & -1 & 1 & -1 \\ 0 & 0 & 9 & -5 \end{bmatrix}$$

$$\xrightarrow{\substack{-R_2 \Rightarrow R_2 \\ \frac{1}{9}R_3 \Rightarrow R_3}} \begin{bmatrix} 1 & 1 & 2 & 1 \\ 0 & 1 & -1 & 1 \\ 0 & 0 & 1 & -\frac{5}{9} \end{bmatrix} \xrightarrow{\substack{R_1 - 2R_3 \Rightarrow R_1 \\ R_2 + R_3 \Rightarrow R_2}} \begin{bmatrix} 1 & 1 & 0 & \frac{19}{9} \\ 0 & 1 & 0 & \frac{4}{9} \\ 0 & 0 & 1 & -\frac{5}{9} \end{bmatrix} \xrightarrow{R_1 - R_2 \Rightarrow R_1} \begin{bmatrix} 1 & 0 & 0 & \frac{5}{3} \\ 0 & 1 & 0 & \frac{4}{9} \\ 0 & 0 & 1 & -\frac{5}{9} \end{bmatrix}$$

Example 11. The RREF of $A = \begin{bmatrix} 2 & 1 \\ 4 & -6 \end{bmatrix}$ from earlier is the 2×2 identity matrix.

Comment. That's not surprising: A square matrix is invertible if and only if its RREF is the identity matrix. If that isn't obvious to you, think about how you invert a matrix using Gaussian elimination (reviewed next).

Review. Recall the Gauss–Jordan method of computing A^{-1} . Starting with the augmented matrix $[A \mid I]$, we do Gaussian elimination until we obtain the RREF, which will be of the form $[I \mid A^{-1}]$ so that we can read off A^{-1} .

Why does that work? By our discussion, the steps of Gaussian elimination can be expressed by multiplication (on the left) with a matrix B . Only looking at the first part of the augmented matrix, and since the RREF of an invertible matrix is I , we have $BA = I$, which means that we must have $B = A^{-1}$. The other part of the augmented matrix (which is I initially) gets multiplied with $B = A^{-1}$ as well, so that, in the end, it is $BI = A^{-1}$. That’s why we can read off A^{-1} !

For instance. To invert $\begin{bmatrix} 2 & 1 \\ 4 & -6 \end{bmatrix}$ using the Gauss–Jordan method, we would proceed as follows:

$$\left[\begin{array}{cc|cc} 2 & 1 & 1 & 0 \\ 4 & -6 & 0 & 1 \end{array} \right] \xrightarrow{R_2 - 2R_1 \Rightarrow R_2} \left[\begin{array}{cc|cc} 2 & 1 & 1 & 0 \\ 0 & -8 & -2 & 1 \end{array} \right] \xrightarrow{\begin{array}{l} \frac{1}{2}R_1 \Rightarrow R_1 \\ -\frac{1}{8}R_2 \Rightarrow R_2 \end{array}} \left[\begin{array}{cc|cc} 1 & \frac{1}{2} & \frac{1}{2} & 0 \\ 0 & 1 & \frac{1}{4} & -\frac{1}{8} \end{array} \right] \xrightarrow{R_1 - \frac{1}{2}R_2 \Rightarrow R_1} \left[\begin{array}{cc|cc} 1 & 0 & \frac{3}{8} & \frac{1}{16} \\ 0 & 1 & \frac{1}{4} & -\frac{1}{8} \end{array} \right]$$

We conclude that $\begin{bmatrix} 2 & 1 \\ 4 & -6 \end{bmatrix}^{-1} = \begin{bmatrix} \frac{3}{8} & \frac{1}{16} \\ \frac{1}{4} & -\frac{1}{8} \end{bmatrix}$.

Of course, for 2×2 matrices it is much simpler to use the formula $\begin{bmatrix} a & b \\ c & d \end{bmatrix}^{-1} = \frac{1}{ad-bc} \begin{bmatrix} d & -b \\ -c & a \end{bmatrix}$.

Review: Vector spaces, bases, dimension, null spaces

Review.

- Vectors are things that can be **added** and **scaled**.
- Hence, given vectors $\mathbf{v}_1, \dots, \mathbf{v}_n$, the most general we can do is form the **linear combination** $\lambda_1\mathbf{v}_1 + \dots + \lambda_n\mathbf{v}_n$. The set of all these linear combinations is the **span** of $\mathbf{v}_1, \dots, \mathbf{v}_n$, denoted by $\text{span}\{\mathbf{v}_1, \dots, \mathbf{v}_n\}$.

- Vector **spaces** are spans.

Equivalently. Vector spaces are sets of vectors so that the result of adding and scaling remains within that set.

Homework. Of course, the latter is a very informal statement. Revisit the formal definition, probably consisting of a list of axioms, and observe how that matches with the above (for instance, several of the axioms are concerned with addition and scaling satisfying the “expected” rules).

- Recall that vectors from a vector space V form a **basis** of V if and only if
 - the vectors span V , and
 - the vectors are (linearly) independent.

Equivalently. $\mathbf{v}_1, \dots, \mathbf{v}_n$ from V form a basis of V if and only if every vector in V can be expressed as a unique linear combination of $\mathbf{v}_1, \dots, \mathbf{v}_n$.

Just checking. Make sure that you can define precisely what it means for vectors $\mathbf{v}_1, \dots, \mathbf{v}_n$ to be independent.

- The **dimension** of a vector space V is the number of vectors in a basis for V .
No matter what basis one chooses for V , it always has the same number of vectors.

Example 12. \mathbb{R}^3 is the vector space of all vectors with 3 real entries.

\mathbb{R} itself refers to the set of real numbers. We will later also discuss \mathbb{C} , the set of complex numbers.

The **standard basis** of \mathbb{R}^3 is $\begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}$. The dimension of \mathbb{R}^3 is 3.

Review. The **null space** $\text{null}(A)$ of a matrix A consists of those vectors \mathbf{x} such that $A\mathbf{x} = \mathbf{0}$.

Make sure that you see why $\text{null}(A)$ is a vector space. [For instance, if you pick two vectors in $\text{null}(A)$ why is it that the sum of them is in $\text{null}(A)$ again?]

Example 13. What is $\text{null}(A)$ if the matrix A is invertible?

Solution. If A is invertible, then $A\mathbf{x} = \mathbf{0}$ has the unique solution $\mathbf{x} = A^{-1}\mathbf{0} = \mathbf{0}$.

Hence, $\text{null}(A) = \{\mathbf{0}\}$ which is the trivial vector space (consisting of only the null vector) and has dimension 0.

Example 14. Compute a basis for $\text{null}(A)$ where $A = \begin{bmatrix} -1 & 0 & 2 \\ 2 & -3 & 2 \\ 1 & 0 & -2 \end{bmatrix}$.

Solution. We perform row operations and obtain

$$\text{null}\left(\begin{bmatrix} -1 & 0 & 2 \\ 2 & -3 & 2 \\ 1 & 0 & -2 \end{bmatrix}\right) \stackrel{\substack{R_2 + 2R_1 \Rightarrow R_2 \\ R_3 + R_1 \Rightarrow R_3}}{=} \text{null}\left(\begin{bmatrix} -1 & 0 & 2 \\ 0 & -3 & 6 \\ 0 & 0 & 0 \end{bmatrix}\right) \stackrel{\substack{-R_1 \Rightarrow R_1 \\ -\frac{1}{3}R_2 \Rightarrow R_2}}{=} \text{null}\left(\begin{bmatrix} 1 & 0 & -2 \\ 0 & 1 & -2 \\ 0 & 0 & 0 \end{bmatrix}\right).$$

From the RREF, we can now read off the general solution to $A\mathbf{x} = \mathbf{0}$:

- x_1 and x_2 are pivot variables. [For each we have an equation expressing it in terms of the other variables; for instance, $x_1 - 2x_3 = 0$ tells us that $x_1 = 2x_3$.]
- x_3 is a free variable. [There is no equation forcing a value on x_3 .]
- Hence, without computation, we see that the general solution is $\begin{bmatrix} 2x_3 \\ 2x_3 \\ x_3 \end{bmatrix}$.

In other words, a basis is $\begin{bmatrix} 2 \\ 2 \\ 1 \end{bmatrix}$.

Comment. We are starting with the three equations $-x_1 + 2x_3 = 0$, $2x_1 - 3x_2 + 2x_3 = 0$, $x_1 - 2x_3 = 0$. Performing row operations on the matrix is the same as combining these equations (with the objective to form simpler equations by eliminating variables).

Example 15. Compute a basis for $\text{null}\left(\begin{bmatrix} 2 & 0 & 2 \\ 2 & 0 & 2 \\ 1 & 0 & 1 \end{bmatrix}\right)$.

Solution.

$$\text{null}\left(\begin{bmatrix} 2 & 0 & 2 \\ 2 & 0 & 2 \\ 1 & 0 & 1 \end{bmatrix}\right) \stackrel{\substack{R_2 - R_1 \Rightarrow R_2 \\ R_3 - \frac{1}{2}R_1 \Rightarrow R_3}}{=} \text{null}\left(\begin{bmatrix} 2 & 0 & 2 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}\right) \stackrel{\frac{1}{2}R_1 \Rightarrow R_1}{=} \text{null}\left(\begin{bmatrix} 1 & 0 & 1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}\right)$$

This time, x_2 and x_3 are free variables. The general solution is $\begin{bmatrix} -x_3 \\ x_2 \\ x_3 \end{bmatrix} = x_2 \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix} + x_3 \begin{bmatrix} -1 \\ 0 \\ 1 \end{bmatrix}$.

Hence, a basis is $\begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} -1 \\ 0 \\ 1 \end{bmatrix}$.

Review: Eigenvalues and eigenvectors

If $Ax = \lambda x$ (and $x \neq 0$), then x is an **eigenvector** of A with **eigenvalue** λ (just a number).

Note that for the equation $Ax = \lambda x$ to make sense, A needs to be a square matrix (i.e. $n \times n$).

Key observation:

$$\begin{aligned} Ax &= \lambda x \\ \iff Ax - \lambda x &= 0 \\ \iff (A - \lambda I)x &= 0 \end{aligned}$$

This homogeneous system has a nontrivial solution x if and only if $\det(A - \lambda I) = 0$.

To find eigenvectors and eigenvalues of A :

(a) First, find the eigenvalues λ by solving $\det(A - \lambda I) = 0$.

$\det(A - \lambda I)$ is a polynomial in λ , called the **characteristic polynomial** of A .

(b) Then, for each eigenvalue λ , find corresponding eigenvectors by solving $(A - \lambda I)x = 0$.

More precisely, we find a basis of eigenvectors for the λ -**eigenspace** $\text{null}(A - \lambda I)$.

Example 16. $A = \begin{bmatrix} 4 & 0 & 2 \\ 2 & 2 & 2 \\ 1 & 0 & 3 \end{bmatrix}$ has one eigenvector that is “easy” to see. Do you see it?

Solution. Note that $A \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix} = \begin{bmatrix} 2 \\ 2 \\ 1 \end{bmatrix} = 2 \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}$. Hence, $\begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}$ is a 2-eigenvector.

Just for contrast. Note that $A \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} = \begin{bmatrix} 2 \\ 2 \\ 3 \end{bmatrix} \neq \lambda \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}$. Hence, $\begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}$ is not an eigenvector.

Suppose that A is $n \times n$ and has independent eigenvectors x_1, \dots, x_n .

Then A can be **diagonalized** as $A = PDP^{-1}$, where

- the columns of P are the eigenvectors, and
- the diagonal matrix D has the eigenvalues on the diagonal.

Such a diagonalization is possible if and only if A has enough (independent) eigenvectors.

Comment. If you don't quite recall why these choices result in the diagonalization $A = PDP^{-1}$, note that the diagonalization is equivalent to $AP = PD$.

- Put the eigenvectors x_1, \dots, x_n as columns into a matrix P .

$$\begin{aligned} Ax_i = \lambda_i x_i \implies A \begin{bmatrix} | & & | \\ x_1 & \dots & x_n \\ | & & | \end{bmatrix} &= \begin{bmatrix} | & & | \\ \lambda_1 x_1 & \dots & \lambda_n x_n \\ | & & | \end{bmatrix} \\ &= \begin{bmatrix} | & & | \\ x_1 & \dots & x_n \\ | & & | \end{bmatrix} \begin{bmatrix} \lambda_1 & & \\ & \ddots & \\ & & \lambda_n \end{bmatrix} \end{aligned}$$

- In summary: $AP = PD$

Example 17. Let $A = \begin{bmatrix} 4 & 0 & 2 \\ 2 & 2 & 2 \\ 1 & 0 & 3 \end{bmatrix}$.

- Find the eigenvalues and bases for the eigenspaces of A .
- Diagonalize A . That is, determine matrices P and D such that $A = PDP^{-1}$.

Solution.

- By expanding by the second column, we find that the characteristic polynomial $\det(A - \lambda I)$ is

$$\begin{vmatrix} 4-\lambda & 0 & 2 \\ 2 & 2-\lambda & 2 \\ 1 & 0 & 3-\lambda \end{vmatrix} = (2-\lambda) \begin{vmatrix} 4-\lambda & 2 \\ 1 & 3-\lambda \end{vmatrix} = (2-\lambda)[(4-\lambda)(3-\lambda) - 2] = (2-\lambda)^2(5-\lambda).$$

Hence, the eigenvalues are $\lambda = 2$ (with multiplicity 2) and $\lambda = 5$.

Comment. At this point, we know that we will find one eigenvector for $\lambda = 5$ (more precisely, the 5-eigenspace definitely has dimension 1). On the other hand, the 2-eigenspace might have dimension 2 or 1. In order for A to be diagonalizable, the 2-eigenspace must have dimension 2. (Why?)

- The 5-eigenspace is $\text{null}\left(\begin{bmatrix} -1 & 0 & 2 \\ 2 & -3 & 2 \\ 1 & 0 & -2 \end{bmatrix}\right)$. Proceeding as in Example 14, we obtain

$$\text{null}\left(\begin{bmatrix} -1 & 0 & 2 \\ 2 & -3 & 2 \\ 1 & 0 & -2 \end{bmatrix}\right) \stackrel{\text{RREF}}{=} \text{null}\left(\begin{bmatrix} 1 & 0 & -2 \\ 0 & 1 & -2 \\ 0 & 0 & 0 \end{bmatrix}\right) = \text{span}\left\{\begin{bmatrix} 2 \\ 2 \\ 1 \end{bmatrix}\right\}.$$

In other words, the 5-eigenspace has basis $\begin{bmatrix} 2 \\ 2 \\ 1 \end{bmatrix}$.

- The 2-eigenspace is $\text{null}\left(\begin{bmatrix} 2 & 0 & 2 \\ 2 & 0 & 2 \\ 1 & 0 & 1 \end{bmatrix}\right)$. Proceeding as in Example 15, we obtain

$$\text{null}\left(\begin{bmatrix} 2 & 0 & 2 \\ 2 & 0 & 2 \\ 1 & 0 & 1 \end{bmatrix}\right) \stackrel{\text{RREF}}{=} \text{null}\left(\begin{bmatrix} 1 & 0 & 1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}\right) = \text{span}\left\{\begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} -1 \\ 0 \\ 1 \end{bmatrix}\right\}$$

In other words, the 2-eigenspace has basis $\begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} -1 \\ 0 \\ 1 \end{bmatrix}$.

Comment. So, indeed, the 2-eigenspace has dimension 2. In particular, A is diagonalizable.

- A possible choice is $P = \begin{bmatrix} 2 & 0 & -1 \\ 2 & 1 & 0 \\ 1 & 0 & 1 \end{bmatrix}$, $D = \begin{bmatrix} 5 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 2 \end{bmatrix}$.

Comment. However, many other choices are possible and correct. For instance, the order of the eigenvalues in D doesn't matter (as long as the same order is used for P). Also, for P , the columns can be chosen to be any other set of eigenvectors.

Example 18. (extra practice) Diagonalize, if possible, the matrices

$$A = \begin{bmatrix} 3 & 4 & 1 \\ 0 & 2 & 0 \\ 1 & 4 & 3 \end{bmatrix}, \quad B = \begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \quad C = \begin{bmatrix} 1 & 0 & 1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}.$$

Solution. For instance, $A = PDP^{-1}$ with $P = \begin{bmatrix} 1 & -4 & -1 \\ 0 & 1 & 0 \\ 1 & 0 & 1 \end{bmatrix}$ and $D = \begin{bmatrix} 4 & & \\ & 2 & \\ & & 2 \end{bmatrix}$. B is not diagonalizable.

For instance, $C = PDP^{-1}$ with $P = \begin{bmatrix} 1 & 0 & -1 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$ and $D = \begin{bmatrix} 1 & & \\ & 0 & \\ & & 0 \end{bmatrix}$.

Review: Computing determinants using cofactor expansion

Review. Let A be an $n \times n$ matrix. The **determinant** of A , written as $\det(A)$ or $|A|$, is a number with the property that:

$$\begin{aligned} \det(A) \neq 0 &\iff A \text{ is invertible} \\ &\iff Ax = b \text{ has a (unique) solution } x \text{ (for all } b) \\ &\iff Ax = 0 \text{ is only solved by } x = 0 \end{aligned}$$

Example 19. $\begin{vmatrix} a & b \\ c & d \end{vmatrix} = ad - bc$

Example 20. Compute $\begin{vmatrix} 1 & 2 & 0 \\ 3 & -1 & 2 \\ 2 & 0 & 1 \end{vmatrix}$ by **cofactor expansion**.

Solution. We expand by the first row:

$$\begin{aligned} \begin{vmatrix} 1 & 2 & 0 \\ 3 & -1 & 2 \\ 2 & 0 & 1 \end{vmatrix} &= 1 \cdot \begin{vmatrix} + & & \\ & -1 & 2 \\ & 0 & 1 \end{vmatrix} - 2 \cdot \begin{vmatrix} & - & \\ 3 & & 2 \\ 2 & & 1 \end{vmatrix} + 0 \cdot \begin{vmatrix} & & + \\ 3 & -1 & \\ 2 & 0 & \end{vmatrix} \\ &\stackrel{\text{i.e.}}{=} 1 \cdot \begin{vmatrix} -1 & 2 \\ 0 & 1 \end{vmatrix} - 2 \cdot \begin{vmatrix} 3 & 2 \\ 2 & 1 \end{vmatrix} + 0 \cdot \begin{vmatrix} 3 & -1 \\ 2 & 0 \end{vmatrix} = 1 \cdot (-1) - 2 \cdot (-1) + 0 = 1 \end{aligned}$$

Each term in the cofactor expansion is ± 1 times an entry times a smaller determinant (row and column of entry deleted).

The ± 1 is assigned to each entry according to $\begin{bmatrix} + & - & + & \dots \\ - & + & - & \\ + & - & + & \\ \vdots & & & \ddots \end{bmatrix}$.

Solution. We expand by the second column:

$$\begin{aligned} \begin{vmatrix} 1 & 2 & 0 \\ 3 & -1 & 2 \\ 2 & 0 & 1 \end{vmatrix} &= -2 \cdot \begin{vmatrix} & - & \\ 3 & & 2 \\ 2 & & 1 \end{vmatrix} + (-1) \cdot \begin{vmatrix} 1 & & 0 \\ & + & \\ 2 & & 1 \end{vmatrix} - 0 \cdot \begin{vmatrix} 1 & & 0 \\ 3 & & 2 \\ & - & \end{vmatrix} \\ &= -2 \cdot (-1) + (-1) \cdot 1 - 0 = 1 \end{aligned}$$

Example 21. Compute $\begin{vmatrix} 1 & 0 & 3 & 4 \\ 0 & 2 & 1 & 5 \\ 0 & 0 & 2 & 1 \\ 2 & 0 & 8 & 5 \end{vmatrix}$.

Solution. We can expand by the second column:

$$\begin{vmatrix} 1 & 0 & 3 & 4 \\ 0 & 2 & 1 & 5 \\ 0 & 0 & 2 & 1 \\ 2 & 0 & 8 & 5 \end{vmatrix} = -0 \begin{vmatrix} 0 & 1 & 5 \\ 0 & 2 & 1 \\ 2 & 8 & 5 \end{vmatrix} + 2 \begin{vmatrix} 1 & 3 & 4 \\ 0 & 2 & 1 \\ 2 & 8 & 5 \end{vmatrix} - 0 \begin{vmatrix} 1 & 3 & 4 \\ 0 & 1 & 5 \\ 2 & 8 & 5 \end{vmatrix} + 0 \begin{vmatrix} 1 & 3 & 4 \\ 0 & 1 & 5 \\ 0 & 2 & 1 \end{vmatrix}$$

[Of course, you don't have to spell out the 3×3 matrices that get multiplied with 0.]

We can compute the remaining 3×3 matrix in any way we prefer. One option is to expand by the first column:

$$2 \begin{vmatrix} 1 & 3 & 4 \\ 0 & 2 & 1 \\ 2 & 8 & 5 \end{vmatrix} = 2 \left(+1 \begin{vmatrix} 2 & 1 \\ 8 & 5 \end{vmatrix} + 2 \begin{vmatrix} 3 & 4 \\ 2 & 1 \end{vmatrix} \right) = 2(1 \cdot 2 + 2 \cdot (-5)) = -16$$

Comment. For cofactor expansion, choosing to expand by the second column is the best choice because this column has more zeros than any other column or row.

The determinant of a triangular matrix is the product of the diagonal entries.

Why? Can you explain this (you can use the next example) using cofactor expansion?

Example 22. Compute $\begin{vmatrix} 1 & 0 & 3 & -1 \\ 0 & 3 & 1 & 5 \\ 0 & 0 & -2 & 1 \\ 0 & 0 & 0 & 5 \end{vmatrix}$.

Solution. Since the matrix is (upper) triangular, $\begin{vmatrix} 1 & 0 & 3 & -1 \\ 0 & 3 & 1 & 5 \\ 0 & 0 & -2 & 1 \\ 0 & 0 & 0 & 5 \end{vmatrix} = 1 \cdot 3 \cdot (-2) \cdot 5 = -30$.

Review.

- Effect of row (or column) operations on determinant.
- $\det(AB) = \det(A)\det(B)$
- In particular, the LU decomposition provides us with a way to compute determinants: If $A = LU$, then $\det(A) = \det(L)\det(U)$ and the latter determinants are just products of diagonal entries (because both L and U are triangular).

Comment. Unless a row swap is required, we can compute the LU decomposition of $A = LU$ using only row operations of the form $R_i + cR_j \Rightarrow R_i$ (those don't change the determinant!).

In that case, the matrix L will have 1's on the diagonal. In particular, $\det(L) = 1$.

Consequently, in that case, $\det(A) = \det(U)$.

Practical comment. For larger matrices, cofactor expansion is a terribly inefficient way of computing determinants. Instead, Gaussian elimination (i.e. LU decomposition) is much more efficient.

On the other hand, cofactor expansion is a good choice when working by hand with small matrices.

Example 23. (review) If $A = \begin{bmatrix} 1 & 4 \\ 2 & 5 \\ 3 & 6 \end{bmatrix}$, then its **transpose** is $A^T = \begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{bmatrix}$.

Recall that $(AB)^T = B^T A^T$. This reflects the fact that, in the column-centric versus the row-centric interpretation of matrix multiplication, the order of the matrices is reversed.

Comment. When working with complex numbers, the fundamental role is not played by the transpose but by the **conjugate transpose** instead (we'll see that in our discussion of orthogonality): $A^* = \overline{A^T}$.

For instance, if $A = \begin{bmatrix} 1-3i & 5i \\ 2+i & 3 \end{bmatrix}$, then $A^* = \begin{bmatrix} 1+3i & 2-i \\ -5i & 3 \end{bmatrix}$.

Orthogonality

The inner product and distances

Definition 24. The **inner product** (or **dot product**) of \mathbf{v} , \mathbf{w} in \mathbb{R}^n :

$$\mathbf{v} \cdot \mathbf{w} = \mathbf{v}^T \mathbf{w} = v_1 w_1 + \dots + v_n w_n.$$

Because we can think of this as a special case of the matrix product, it satisfies the basic rules like associativity and distributivity.

In addition: $\mathbf{v} \cdot \mathbf{w} = \mathbf{w} \cdot \mathbf{v}$.

Example 25. $\begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix} \cdot \begin{bmatrix} 2 \\ -1 \\ 4 \end{bmatrix} = 2 - 2 + 12 = 12$


Definition 26.

- The **norm** (or **length**) of a vector \mathbf{v} in \mathbb{R}^n is

$$\|\mathbf{v}\| = \sqrt{\mathbf{v} \cdot \mathbf{v}} = \sqrt{v_1^2 + \dots + v_n^2}.$$

- The **distance** between points \mathbf{v} and \mathbf{w} in \mathbb{R}^n is

$$\text{dist}(\mathbf{v}, \mathbf{w}) = \|\mathbf{v} - \mathbf{w}\|.$$



Example 27. For instance, in \mathbb{R}^2 , $\text{dist}\left(\begin{bmatrix} x_1 \\ y_1 \end{bmatrix}, \begin{bmatrix} x_2 \\ y_2 \end{bmatrix}\right) = \left\| \begin{bmatrix} x_1 - x_2 \\ y_1 - y_2 \end{bmatrix} \right\| = \sqrt{(x_1 - x_2)^2 + (y_1 - y_2)^2}$.

Example 28. Write $\|\mathbf{v} - \mathbf{w}\|^2$ as a dot product, and multiply it out.

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

Comment. This is a vector version of $(x - y)^2 = x^2 - 2xy + y^2$.

The reason we were careful and first wrote $-\mathbf{v} \cdot \mathbf{w} - \mathbf{w} \cdot \mathbf{v}$ before simplifying it to $-2\mathbf{v} \cdot \mathbf{w}$ is that we should not take rules such as $\mathbf{v} \cdot \mathbf{w} = \mathbf{w} \cdot \mathbf{v}$ for granted. For instance, for the cross product $\mathbf{v} \times \mathbf{w}$, that you may have seen in Calculus, we have $\mathbf{v} \times \mathbf{w} \neq \mathbf{w} \times \mathbf{v}$ (instead, $\mathbf{v} \times \mathbf{w} = -\mathbf{w} \times \mathbf{v}$).

Orthogonal vectors

Definition 29. v and w in \mathbb{R}^n are **orthogonal** if

$$v \cdot w = 0.$$

Why? How is this related to our understanding of right angles?


Pythagoras!

v and w are orthogonal

$$\begin{aligned} \iff \|v\|^2 + \|w\|^2 &= \underbrace{\|v - w\|^2}_{= \|v\|^2 - 2v \cdot w + \|w\|^2} \\ &\text{(by previous example)} \end{aligned}$$

$$\iff -2v \cdot w = 0$$

$$\iff v \cdot w = 0$$



Definition 30. We say that two subspaces V and W of \mathbb{R}^n are **orthogonal** if and only if every vector in V is orthogonal to every vector in W .

The **orthogonal complement** of V is the space V^\perp of all vectors that are orthogonal to V .

Exercise. Show that the orthogonal complement is indeed a vector space. Alternatively, this follows from our discussion in the next example which leads to Theorem 32. Namely, every space V can be written as $V = \text{col}(A)$ for a suitable matrix A (for instance, we can choose the columns of A to be basis vectors of V). It then follows that $V^\perp = \text{null}(A^T)$ (which is clearly a space).

Example 31. Determine a basis for the orthogonal complement of $V = \text{span}\left\{\begin{bmatrix} 1 \\ 2 \\ 1 \end{bmatrix}, \begin{bmatrix} 3 \\ 1 \\ 2 \end{bmatrix}\right\}$.

Solution. The orthogonal complement V^\perp consists of all vectors $\begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}$ that are orthogonal to $\begin{bmatrix} 1 \\ 2 \\ 1 \end{bmatrix}$ and $\begin{bmatrix} 3 \\ 1 \\ 2 \end{bmatrix}$.

Using the dot product, this means we must have $\begin{bmatrix} 1 \\ 2 \\ 1 \end{bmatrix} \cdot \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = 0$ as well as $\begin{bmatrix} 3 \\ 1 \\ 2 \end{bmatrix} \cdot \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = 0$.

Note that this is equivalent to the equations $1x_1 + 2x_2 + 1x_3 = 0$ and $3x_1 + 1x_2 + 2x_3 = 0$.

In matrix-vector form, these two equations combine to $\begin{bmatrix} 1 & 2 & 1 \\ 3 & 1 & 2 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$.

This is the same as saying that $\begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}$ has to be in $\text{null}\left(\begin{bmatrix} 1 & 2 & 1 \\ 3 & 1 & 2 \end{bmatrix}\right)$. This means that $V^\perp = \text{null}\left(\begin{bmatrix} 1 & 2 & 1 \\ 3 & 1 & 2 \end{bmatrix}\right)$.

[Note that we have done no computations up to this point! Instead, we have derived Theorem 32 below.]

We compute (fill in the work!) that $V^\perp = \text{null}\left(\begin{bmatrix} 1 & 2 & 1 \\ 3 & 1 & 2 \end{bmatrix}\right) \stackrel{\text{RREF}}{=} \text{null}\left(\begin{bmatrix} 1 & 0 & 3/5 \\ 0 & 1 & 1/5 \end{bmatrix}\right) = \text{span}\left\{\begin{bmatrix} -3/5 \\ -1/5 \\ 1 \end{bmatrix}\right\}$.

Check. $\begin{bmatrix} -3/5 \\ -1/5 \\ 1 \end{bmatrix}$ is indeed orthogonal to both $\begin{bmatrix} 1 \\ 2 \\ 1 \end{bmatrix}$ and $\begin{bmatrix} 3 \\ 1 \\ 2 \end{bmatrix}$.

Note. If $\begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}$ is orthogonal to both basis vectors $\begin{bmatrix} 1 \\ 2 \\ 1 \end{bmatrix}$ and $\begin{bmatrix} 3 \\ 1 \\ 2 \end{bmatrix}$, then it is orthogonal to every vector in V .

Indeed, vectors in V are of the form $v = a \begin{bmatrix} 1 \\ 2 \\ 1 \end{bmatrix} + b \begin{bmatrix} 3 \\ 1 \\ 2 \end{bmatrix}$ and we have $v \cdot \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = a \underbrace{\begin{bmatrix} 1 \\ 2 \\ 1 \end{bmatrix} \cdot \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}}_{=0} + b \underbrace{\begin{bmatrix} 3 \\ 1 \\ 2 \end{bmatrix} \cdot \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}}_{=0} = 0$.

Just to make sure. Why is it geometrically clear that the orthogonal complement of V is 1-dimensional?

The following theorem follows by the same reasoning that we used in the previous example.

In that example, we started with $V = \text{col}\left(\begin{bmatrix} 1 & 3 \\ 2 & 1 \\ 1 & 2 \end{bmatrix}\right)$ and found that $V^\perp = \text{null}\left(\begin{bmatrix} 1 & 2 & 1 \\ 3 & 1 & 2 \end{bmatrix}\right)$.

Theorem 32. If $V = \text{col}(A)$, then $V^\perp = \text{null}(A^T)$.
 In particular, if V is a subspace of \mathbb{R}^n with $\dim(V) = r$, then $\dim(V^\perp) = n - r$.

For short. $\text{col}(A)^\perp = \text{null}(A^T)$

Note that the second part can be written as $\dim(V) + \dim(V^\perp) = n$.

To see that this is true, suppose we choose the columns of A to be a basis of V . If V is a subspace of \mathbb{R}^n with $\dim(V) = r$, then A is a $r \times n$ matrix with r pivot columns. Correspondingly, A^T is a $n \times r$ matrix with r pivot rows. Since $n \geq r$ there are $n - r$ free variables when computing a basis for $\text{null}(A^T)$. Hence, $\dim(V^\perp) = n - r$.

Example 33. Suppose that V is spanned by 3 linearly independent vectors in \mathbb{R}^5 . Determine the dimension of V and its orthogonal complement V^\perp .

Solution. This means that $\dim V = 3$. By Theorem 32, we have $\dim V^\perp = 5 - 3 = 2$.

Example 34. Determine a basis for the orthogonal complement of (the span of) $\begin{bmatrix} 1 \\ 2 \\ 1 \end{bmatrix}$.

Solution. Here, $V = \text{span}\left\{\begin{bmatrix} 1 \\ 2 \\ 1 \end{bmatrix}\right\}$ and we are looking for the orthogonal complement V^\perp .

Since $V = \text{col}\left(\begin{bmatrix} 1 \\ 2 \\ 1 \end{bmatrix}\right)$, it follows from Theorem 32 that $V^\perp = \text{null}([1 \ 2 \ 1])$.

Computing a basis for $\text{null}([1 \ 2 \ 1])$ is easy since $[1 \ 2 \ 1]$ is already in RREF.

Note that the general solution to $[1 \ 2 \ 1]\mathbf{x} = 0$ is $\begin{bmatrix} -2s - t \\ s \\ t \end{bmatrix} = s \begin{bmatrix} -2 \\ 1 \\ 0 \end{bmatrix} + t \begin{bmatrix} -1 \\ 0 \\ 1 \end{bmatrix}$.

A basis for $V^\perp = \text{null}([1 \ 2 \ 1])$ therefore is $\begin{bmatrix} -2 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} -1 \\ 0 \\ 1 \end{bmatrix}$.

Check. We easily check (do it!) that both of these are indeed orthogonal to the original vector $\begin{bmatrix} 1 \\ 2 \\ 1 \end{bmatrix}$.

The fundamental theorem

Review. The four **fundamental subspaces** associated with a matrix A are

$$\text{col}(A), \quad \text{row}(A), \quad \text{null}(A), \quad \text{null}(A^T).$$

Note that $\text{row}(A) = \text{col}(A^T)$. (In particular, we usually write vectors in $\text{row}(A)$ as column vectors.)

Comment. $\text{null}(A^T)$ is called the **left null space** of A .

Why that name? Recall that, by definition \mathbf{x} is in $\text{null}(A) \iff A\mathbf{x} = \mathbf{0}$.

Likewise, \mathbf{x} is in $\text{null}(A^T) \iff A^T\mathbf{x} = \mathbf{0} \iff \mathbf{x}^T A = \mathbf{0}$.

[Recall that $(AB)^T = B^T A^T$. In particular, $(A^T \mathbf{x})^T = \mathbf{x}^T A$, which is what we used in the last equivalence.]

Review. The **rank** of a matrix is the number of pivots in its RREF.

Equivalently, as showcased in the next result, the rank is the dimension of either the column or the row space.

Theorem 35. (Fundamental Theorem of Linear Algebra, Part I)

Let A be an $m \times n$ matrix of **rank** r .

- $\dim \text{col}(A) = r$ (subspace of \mathbb{R}^m)
- $\dim \text{row}(A) = r$ (subspace of \mathbb{R}^n) $\text{row}(A) = \text{col}(A^T)$
- $\dim \text{null}(A) = n - r$ (subspace of \mathbb{R}^n)
- $\dim \text{null}(A^T) = m - r$ (subspace of \mathbb{R}^m)

Example 36. Let $A = \begin{bmatrix} 1 & 2 \\ 2 & 4 \\ 3 & 6 \end{bmatrix}$. Determine bases for all four fundamental subspaces.

Solution. Make sure that, for such a simple matrix, you can see all of these that at a glance!

$$\text{col}(A) = \text{span} \left\{ \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix} \right\}, \quad \text{row}(A) = \text{span} \left\{ \begin{bmatrix} 1 \\ 2 \end{bmatrix} \right\}, \quad \text{null}(A) = \text{span} \left\{ \begin{bmatrix} -2 \\ 1 \end{bmatrix} \right\}, \quad \text{null}(A^T) = \text{span} \left\{ \begin{bmatrix} -2 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} -3 \\ 0 \\ 1 \end{bmatrix} \right\}$$

Important observation. The basis vectors for $\text{row}(A)$ and $\text{null}(A)$ are orthogonal! $\begin{bmatrix} -2 \\ 1 \end{bmatrix} \cdot \begin{bmatrix} 1 \\ 2 \end{bmatrix} = 0$

The same is true for the basis vectors for $\text{col}(A)$ and $\text{null}(A^T)$: $\begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix} \cdot \begin{bmatrix} -2 \\ 1 \\ 0 \end{bmatrix} = 0$ and $\begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix} \cdot \begin{bmatrix} -3 \\ 0 \\ 1 \end{bmatrix} = 0$

Always. Vectors in $\text{null}(A)$ are orthogonal to vectors in $\text{row}(A)$. In short, $\text{null}(A)$ is orthogonal to $\text{row}(A)$.

Why? Suppose that \mathbf{x} is in $\text{null}(A)$. That is, $A\mathbf{x} = \mathbf{0}$. But think about what $A\mathbf{x} = \mathbf{0}$ means (row-product rule). It means that the inner product of every row with \mathbf{x} is zero. Which implies that \mathbf{x} is orthogonal to the row space.

Theorem 37. (Fundamental Theorem of Linear Algebra, Part II)

- $\text{null}(A)$ is orthogonal to $\text{row}(A)$. (both subspaces of \mathbb{R}^n)

Note that $\dim \text{null}(A) + \dim \text{row}(A) = n$. Hence, the two spaces are orthogonal complements.

- $\text{null}(A^T)$ is orthogonal to $\text{col}(A)$.

Again, the two spaces are orthogonal complements. (This is just the first part with A replaced by A^T .)

Example 38. Let $A = \begin{bmatrix} 1 & 2 & 1 & 4 \\ 2 & 4 & 0 & 2 \\ 3 & 6 & 0 & 3 \end{bmatrix}$. Check that $\text{null}(A)$ and $\text{row}(A)$ are orthogonal complements.

Solution.

$$\begin{bmatrix} 1 & 2 & 1 & 4 \\ 2 & 4 & 0 & 2 \\ 3 & 6 & 0 & 3 \end{bmatrix} \xrightarrow[\begin{smallmatrix} R_3 - 3R_1 \Rightarrow R_3 \\ \rightsquigarrow \end{smallmatrix}]{\begin{smallmatrix} R_2 - 2R_1 \Rightarrow R_2 \\ \rightsquigarrow \end{smallmatrix}} \begin{bmatrix} 1 & 2 & 1 & 4 \\ 0 & 0 & -2 & -6 \\ 0 & 0 & -3 & -9 \end{bmatrix} \xrightarrow[\rightsquigarrow]{R_3 - \frac{3}{2}R_2 \Rightarrow R_3} \begin{bmatrix} 1 & 2 & 1 & 4 \\ 0 & 0 & -2 & -6 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

$$\xrightarrow[\rightsquigarrow]{-\frac{1}{2}R_2 \Rightarrow R_2} \begin{bmatrix} 1 & 2 & 1 & 4 \\ 0 & 0 & 1 & 3 \\ 0 & 0 & 0 & 0 \end{bmatrix} \xrightarrow[\rightsquigarrow]{R_1 - R_2 \Rightarrow R_1} \begin{bmatrix} 1 & 2 & 0 & 1 \\ 0 & 0 & 1 & 3 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

$$\text{Hence, } \text{null}(A) = \text{span} \left\{ \begin{bmatrix} -2 \\ 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} -1 \\ 0 \\ -3 \\ 1 \end{bmatrix} \right\}, \text{row}(A) = \text{span} \left\{ \begin{bmatrix} 1 \\ 2 \\ 0 \\ 1 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 1 \\ 3 \end{bmatrix} \right\}.$$

$\text{null}(A)$ and $\text{row}(A)$ are indeed orthogonal, as certified by:

$$\begin{bmatrix} -2 \\ 1 \\ 0 \\ 0 \end{bmatrix} \cdot \begin{bmatrix} 1 \\ 2 \\ 0 \\ 1 \end{bmatrix} = 0, \quad \begin{bmatrix} -2 \\ 1 \\ 0 \\ 0 \end{bmatrix} \cdot \begin{bmatrix} 0 \\ 0 \\ 1 \\ 3 \end{bmatrix} = 0, \quad \begin{bmatrix} -1 \\ 0 \\ -3 \\ 1 \end{bmatrix} \cdot \begin{bmatrix} 1 \\ 2 \\ 0 \\ 1 \end{bmatrix} = 0, \quad \begin{bmatrix} -1 \\ 0 \\ -3 \\ 1 \end{bmatrix} \cdot \begin{bmatrix} 0 \\ 0 \\ 1 \\ 3 \end{bmatrix} = 0.$$

In fact, $\text{null}(A)$ and $\text{row}(A)$ are orthogonal complements because the dimensions add up to $2 + 2 = 4$.

In particular, $\begin{bmatrix} -2 \\ 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} -1 \\ 0 \\ -3 \\ 1 \end{bmatrix}, \begin{bmatrix} 1 \\ 2 \\ 0 \\ 1 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 1 \\ 3 \end{bmatrix}$ form a basis of all of \mathbb{R}^4 .

Example 39. (extra) Determine bases for all four fundamental subspaces of

$$A = \begin{bmatrix} 1 & 2 & 1 & 3 \\ 2 & 4 & 0 & 1 \\ 3 & 6 & 0 & 1 \end{bmatrix}.$$

Verify all parts of the Fundamental Theorem, especially that $\text{null}(A)$ and $\text{row}(A)$ (as well as $\text{null}(A^T)$ and $\text{col}(A)$) are orthogonal complements.

Partial solution. One can almost see that $\text{rank}(A) = 3$. Hence, the dimensions of the fundamental subspaces are ...

Consistency of a system of equations

Example 40. (warmup) $\begin{bmatrix} 1 & 2 \\ 3 & 1 \\ 0 & 5 \end{bmatrix} \cdot \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = x_1 \begin{bmatrix} 1 \\ 3 \\ 0 \end{bmatrix} + x_2 \begin{bmatrix} 2 \\ 1 \\ 5 \end{bmatrix}$

Note that this means that the system of equations $\begin{matrix} x_1 + 2x_2 = 1 \\ 3x_1 + x_2 = 1 \\ 5x_2 = 1 \end{matrix}$ can also be written as $\begin{bmatrix} 1 & 2 \\ 3 & 1 \\ 0 & 5 \end{bmatrix} \cdot \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}$.

[This was the motivation for introducing matrix-vector multiplication.]

In the same way, any system can be written as $A\mathbf{x} = \mathbf{b}$, where A is a matrix and \mathbf{b} a vector. In particular, this makes it obvious that:

$$A\mathbf{x} = \mathbf{b} \text{ is consistent} \iff \mathbf{b} \text{ is in } \text{col}(A)$$

Recall that, by the FTLA, $\text{col}(A)$ and $\text{null}(A^T)$ are orthogonal complements.

Theorem 41. $A\mathbf{x} = \mathbf{b}$ is consistent $\iff \mathbf{b}$ is orthogonal to $\text{null}(A^T)$

Proof. $A\mathbf{x} = \mathbf{b}$ is consistent $\iff \mathbf{b}$ is in $\text{col}(A) \xleftrightarrow{\text{FTLA}} \mathbf{b}$ is orthogonal to $\text{null}(A^T)$

Note. \mathbf{b} is orthogonal to $\text{null}(A^T)$ means that $\mathbf{y}^T \mathbf{b} = 0$ whenever $\mathbf{y}^T A = \mathbf{0}$. Why?!

Example 42. Let $A = \begin{bmatrix} 1 & 2 \\ 3 & 1 \\ 0 & 5 \end{bmatrix}$. For which \mathbf{b} does $A\mathbf{x} = \mathbf{b}$ have a solution?

Solution. (old)

$$\left[\begin{array}{cc|c} 1 & 2 & b_1 \\ 3 & 1 & b_2 \\ 0 & 5 & b_3 \end{array} \right] \xrightarrow{R_2 - 3R_1 \Rightarrow R_2} \left[\begin{array}{cc|c} 1 & 2 & b_1 \\ 0 & -5 & -3b_1 + b_2 \\ 0 & 5 & b_3 \end{array} \right] \xrightarrow{R_3 + R_2 \Rightarrow R_3} \left[\begin{array}{cc|c} 1 & 2 & b_1 \\ 0 & -5 & -3b_1 + b_2 \\ 0 & 0 & -3b_1 + b_2 + b_3 \end{array} \right]$$

So, $A\mathbf{x} = \mathbf{b}$ is consistent if and only if $-3b_1 + b_2 + b_3 = 0$.

Solution. (new) We determine a basis for $\text{null}(A^T)$:

$$\left[\begin{array}{ccc} 1 & 3 & 0 \\ 2 & 1 & 5 \end{array} \right] \xrightarrow{R_2 - 2R_1 \Rightarrow R_2} \left[\begin{array}{ccc} 1 & 3 & 0 \\ 0 & -5 & 5 \end{array} \right] \xrightarrow{-\frac{1}{5}R_2 \Rightarrow R_2} \left[\begin{array}{ccc} 1 & 3 & 0 \\ 0 & 1 & -1 \end{array} \right] \xrightarrow{R_1 - 3R_2 \Rightarrow R_1} \left[\begin{array}{ccc} 1 & 0 & 3 \\ 0 & 1 & -1 \end{array} \right]$$

We read off from the RREF that $\text{null}(A^T)$ has basis $\begin{bmatrix} -3 \\ 1 \\ 1 \end{bmatrix}$.

\mathbf{b} has to be orthogonal to $\text{null}(A^T)$. That is, $\mathbf{b} \cdot \begin{bmatrix} -3 \\ 1 \\ 1 \end{bmatrix} = 0$. As above!

Comment. Below is how we can use Sage to (try and) solve $A\mathbf{x} = \mathbf{b}$ for $\mathbf{b} = \begin{bmatrix} 1 \\ 1 \\ 2 \end{bmatrix}$ and $\mathbf{b} = \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}$.

```
Sage] A = matrix([[1,2],[3,1],[0,5]])
```

```
Sage] A.solve_right(vector([1,1,2]))
```

$$\begin{pmatrix} 1 & 2 \\ \frac{1}{5} & \frac{2}{5} \end{pmatrix}$$

```
Sage] A.solve_right(vector([1,1,1]))
```

```
Traceback (most recent call last):
ValueError: matrix equation has no solutions
```


Least squares

Example 43. Not all linear systems have solutions.


In fact, for many applications, data needs to be fitted and there is no hope for a perfect match.

For instance, $Ax = b$ with

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

has no solution:

- $\begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}$ is not in $\text{col}(A)$ since $\begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} \cdot \begin{bmatrix} -3 \\ 1 \\ 1 \end{bmatrix} \neq 0$ (see previous example).
- Instead of giving up, we want the x which makes Ax and b as close as possible.
- Such x is characterized by the error $Ax - b$ being **orthogonal** to $\text{col}(A)$ (i.e. all possible Ax).



Definition 44. \hat{x} is a **least squares solution** of the system $Ax = b$ if \hat{x} is such that $A\hat{x} - b$ is as small as possible (i.e. minimal norm).

- If $Ax = b$ is consistent, then \hat{x} is just an ordinary solution. (in that case, $A\hat{x} - b = 0$)
- Interesting case: $Ax = b$ is inconsistent. (in particular, if the system is overdetermined)

The normal equations

The following result provides a straightforward recipe (thanks to the FTLA) to find least squares solutions for all systems $Ax = b$.

Theorem 45. \hat{x} is a least squares solution of $Ax = b$

$$\iff A^T A \hat{x} = A^T b \quad (\text{the normal equations})$$

Proof.

\hat{x} is a least squares solution of $Ax = b$

$\iff A\hat{x} - b$ is as small as possible

$\iff A\hat{x} - b$ is orthogonal to $\text{col}(A)$

$\stackrel{\text{FTLA}}{\iff} A\hat{x} - b$ is in $\text{null}(A^T)$

$\iff A^T(A\hat{x} - b) = 0$

$\iff A^T A \hat{x} = A^T b$

□

Example 46. Find the least squares solution to $A\mathbf{x} = \mathbf{b}$, where

$$A = \begin{bmatrix} 1 & 1 \\ -1 & 1 \\ 0 & 0 \end{bmatrix}, \quad \mathbf{b} = \begin{bmatrix} 2 \\ 1 \\ 1 \end{bmatrix}.$$

Solution. First, $A^T A = \begin{bmatrix} 1 & -1 & 0 \\ 1 & 1 & 0 \end{bmatrix} \begin{bmatrix} 1 & 1 \\ -1 & 1 \\ 0 & 0 \end{bmatrix} = \begin{bmatrix} 2 & 0 \\ 0 & 2 \end{bmatrix}$ and $A^T \mathbf{b} = \begin{bmatrix} 1 & -1 & 0 \\ 1 & 1 & 0 \end{bmatrix} \begin{bmatrix} 2 \\ 1 \\ 1 \end{bmatrix} = \begin{bmatrix} 1 \\ 3 \end{bmatrix}$.

Hence, the normal equations $A^T A \hat{\mathbf{x}} = A^T \mathbf{b}$ take the form $\begin{bmatrix} 2 & 0 \\ 0 & 2 \end{bmatrix} \hat{\mathbf{x}} = \begin{bmatrix} 1 \\ 3 \end{bmatrix}$.

Solving, we immediately find $\hat{\mathbf{x}} = \begin{bmatrix} 1/2 \\ 3/2 \end{bmatrix}$.

Check. Since $A\hat{\mathbf{x}} = \begin{bmatrix} 2 \\ 1 \\ 0 \end{bmatrix}$, the error is $A\hat{\mathbf{x}} - \mathbf{b} = \begin{bmatrix} 0 \\ 0 \\ -1 \end{bmatrix}$. Recall that the error must be orthogonal to $\text{col}(A)$!

This error is indeed orthogonal to $\text{col}(A)$ because $\begin{bmatrix} 0 \\ 0 \\ -1 \end{bmatrix} \cdot \begin{bmatrix} 1 \\ -1 \\ 0 \end{bmatrix} = 0$ and $\begin{bmatrix} 0 \\ 0 \\ -1 \end{bmatrix} \cdot \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix} = 0$.

Comment. Why are the normal equations so particularly simple (compare with example below for the typical case) here? Note how each entry of the product $A^T A$ is computed as the dot product of two columns of A (matrix products of a row of A^T times a column of A). That $A^T A$ is a diagonal matrix reflects the fact that the two columns of A are orthogonal to each other.

Example 47. Find the least squares solution to $A\mathbf{x} = \mathbf{b}$, where

$$A = \begin{bmatrix} 1 & 2 \\ 3 & 1 \\ 0 & 5 \end{bmatrix}, \quad \mathbf{b} = \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}.$$

Solution. First, $A^T A = \begin{bmatrix} 1 & 3 & 0 \\ 2 & 1 & 5 \end{bmatrix} \begin{bmatrix} 1 & 2 \\ 3 & 1 \\ 0 & 5 \end{bmatrix} = \begin{bmatrix} 10 & 5 \\ 5 & 30 \end{bmatrix}$ and $A^T \mathbf{b} = \begin{bmatrix} 1 & 3 & 0 \\ 2 & 1 & 5 \end{bmatrix} \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} = \begin{bmatrix} 4 \\ 8 \end{bmatrix}$.

Hence, the normal equations $A^T A \hat{\mathbf{x}} = A^T \mathbf{b}$ take the form $\begin{bmatrix} 10 & 5 \\ 5 & 30 \end{bmatrix} \hat{\mathbf{x}} = \begin{bmatrix} 4 \\ 8 \end{bmatrix}$.

Since $\begin{bmatrix} 10 & 5 \\ 5 & 30 \end{bmatrix}^{-1} = \frac{1}{275} \begin{bmatrix} 30 & -5 \\ -5 & 10 \end{bmatrix} = \frac{1}{55} \begin{bmatrix} 6 & -1 \\ -1 & 2 \end{bmatrix}$, we find $\hat{\mathbf{x}} = \frac{1}{55} \begin{bmatrix} 6 & -1 \\ -1 & 2 \end{bmatrix} \begin{bmatrix} 4 \\ 8 \end{bmatrix} = \frac{1}{55} \begin{bmatrix} 16 \\ 12 \end{bmatrix}$.

Check. Since $A\hat{\mathbf{x}} = \frac{1}{55} \begin{bmatrix} 40 \\ 60 \\ 60 \end{bmatrix}$, the error $A\hat{\mathbf{x}} - \mathbf{b} = \frac{1}{55} \begin{bmatrix} -15 \\ 5 \\ 5 \end{bmatrix} = \frac{1}{11} \begin{bmatrix} -3 \\ 1 \\ 1 \end{bmatrix}$ must be orthogonal to $\text{col}(A)$.

The error is indeed orthogonal to $\text{col}(A)$ because $\begin{bmatrix} 1 \\ 3 \\ 0 \end{bmatrix} \cdot \frac{1}{11} \begin{bmatrix} -3 \\ 1 \\ 1 \end{bmatrix} = 0$ and $\begin{bmatrix} 2 \\ 1 \\ 5 \end{bmatrix} \cdot \frac{1}{11} \begin{bmatrix} -3 \\ 1 \\ 1 \end{bmatrix} = 0$.

Any serious linear algebra problems are done by a machine. Let us see how to use the open-source computer algebra system **Sage** to do basic computations for us.

Sage is freely available at sagemath.org. Instead of installing it locally (it's huge!) we can conveniently use it in the cloud at cocalc.com from any browser. For short computations, like the one below, you can also just use the input field on our course website.

Sage is built as a **Python** library, so any Python code is valid. Here, we will just use it as a fancy calculator.

Let's revisit Example 38 and let Sage do the work for us:

```
Sage] A = matrix([[1,2,1,4],[2,4,0,2],[3,6,0,3]])
```

```
Sage] A.rref()
```

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

Similarly, if we wanted to compute a basis for $\text{null}(A^T)$, we can simply do:

```
Sage] A.transpose().rref()
```

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

Here are some other standard things we might be interested in (compare with Example 17):

```
Sage] A = matrix([[4,0,2],[2,2,2],[1,0,3]])
```

```
Sage] A.eigenvalues()
```

```
[5, 2, 2]
```

```
Sage] A.eigenvectors_right()
```

$$\left[\left(5, \left[\left(1, 1, \frac{1}{2} \right), 1 \right] \right), \left(2, [(1, 0, -1), (0, 1, 0)], 2 \right) \right]$$

```
Sage] A.eigenmatrix_right()
```

$$\left(\begin{bmatrix} 5 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 2 \end{bmatrix}, \begin{bmatrix} 1 & 1 & 0 \\ 1 & 0 & 1 \\ \frac{1}{2} & -1 & 0 \end{bmatrix} \right)$$

```
Sage] A.rank()
```

```
3
```

```
Sage] A.determinant()
```

```
20
```

```
Sage] A.inverse()
```

$$\begin{bmatrix} \frac{3}{10} & 0 & -\frac{1}{5} \\ -\frac{1}{5} & \frac{1}{2} & -\frac{1}{5} \\ -\frac{1}{10} & 0 & \frac{2}{5} \end{bmatrix}$$

Application: least squares lines

Given data points (x_i, y_i) , we wish to find optimal parameters a, b such that $y_i \approx a + bx_i$ for all i .

Example 48. Determine the line that “best fits” the data points $(2, 1), (5, 2), (7, 3), (8, 3)$.

Comment. Can you see that there is no line fitting the data perfectly? (Check out the last two points!)

Solution. We need to determine the values a, b for the best-fitting line $y = a + bx$.

If there was a line that fit the data perfectly, then:

$$\begin{aligned} a + 2b &= 1 && (2, 1) \\ a + 5b &= 2 && (5, 2) \\ a + 7b &= 3 && (7, 3) \\ a + 8b &= 3 && (8, 3) \end{aligned}$$

In matrix form, this is: $\underbrace{\begin{bmatrix} 1 & x_1 \\ 1 & x_2 \\ 1 & x_3 \\ 1 & x_4 \end{bmatrix}}_{\text{design matrix } X} \begin{bmatrix} a \\ b \end{bmatrix} = \underbrace{\begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ y_4 \end{bmatrix}}_{\text{observation vector } \mathbf{y}}$ (writing the points as (x_i, y_i))

Using our points, these equations become $\begin{bmatrix} 1 & 2 \\ 1 & 5 \\ 1 & 7 \\ 1 & 8 \end{bmatrix} \begin{bmatrix} a \\ b \end{bmatrix} = \begin{bmatrix} 1 \\ 2 \\ 3 \\ 3 \end{bmatrix}$. [This system is inconsistent (as expected).]

We compute a least squares solution.

$$X^T X = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 2 & 5 & 7 & 8 \end{bmatrix} \begin{bmatrix} 1 & 2 \\ 1 & 5 \\ 1 & 7 \\ 1 & 8 \end{bmatrix} = \begin{bmatrix} 4 & 22 \\ 22 & 142 \end{bmatrix}, \quad X^T \mathbf{y} = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 2 & 5 & 7 & 8 \end{bmatrix} \begin{bmatrix} 1 \\ 2 \\ 3 \\ 3 \end{bmatrix} = \begin{bmatrix} 9 \\ 57 \end{bmatrix}.$$

Solving the normal equations $\begin{bmatrix} 4 & 22 \\ 22 & 142 \end{bmatrix} \begin{bmatrix} a \\ b \end{bmatrix} = \begin{bmatrix} 9 \\ 57 \end{bmatrix}$, we find $\begin{bmatrix} a \\ b \end{bmatrix} = \begin{bmatrix} 2/7 \\ 5/14 \end{bmatrix}$.

Hence, the least squares line is $y = \frac{2}{7} + \frac{5}{14}x$.

The plot above shows our points together with this line. It does look like a very good fit!

Important comment. In what sense is this the line of “best fit”? By computing a least squares solution the way we do, we are minimizing the error $\mathbf{y} - X \begin{bmatrix} a \\ b \end{bmatrix}$. The components of that error are $y_i - (a + bx_i)$.

Hence, we see that we are minimizing the **residual sum of squares** $SS_{\text{res}} = \sum_i [y_i - (a + bx_i)]^2$.

Also see the discussion after the next example (where we swap the role of x and y) as well as the example at the beginning of next class (where we discuss making predictions and why minimizing SS_{res} corresponds to minimizing the error of those predictions).

Example 49. (again) Determine the least squares line for the points $(2, 1)$, $(5, 2)$, $(7, 3)$, $(8, 3)$.

Solution. Let's repeat the computation we did in the previous example. This time, we let Sage do the actual work for us:

```
Sage] X = matrix([[1,2],[1,5],[1,7],[1,8]]); y = vector([1,2,3,3])
```

```
Sage] (X.transpose()*X).solve_right(X.transpose()*y)
```

$$\left(\frac{2}{7}, \frac{5}{14}\right)$$

Here are some intermediate steps to help see what's going on (and that it matches our earlier work):

```
Sage] X.transpose()*X
```

$$\begin{bmatrix} 4 & 22 \\ 22 & 142 \end{bmatrix}$$


```
Sage] X.transpose()*y
```

$$(9, 57)$$


Let's plot the least squares line $y = \frac{2}{7} + \frac{5}{14}x$ in Sage to marvel at the good fit!

```
Sage] points = [[2,1],[5,2],[7,3],[8,3]]
```


```
Sage] scatter_plot(points)
```



```
Sage] scatter_plot(points) + plot(2/7+5/14*x,1,9)
```



Comment. As mentioned earlier, the least squares line minimizes the (sum of squares of the) vertical offsets:



<http://mathworld.wolfram.com/LeastSquaresFitting.html>

Comment. We get a (slightly) different “best fit” line if we change the role of x and y ! Can you explain that?


```
Sage] X = matrix([[1,1],[1,2],[1,3],[1,3]]); y = vector([2,5,7,8])
```

```
Sage] (X.transpose()*X).solve_right(X.transpose()*y)
```

$$\left(-\frac{7}{11}, \frac{30}{11}\right)$$

Note that $x = -\frac{7}{11} + \frac{30}{11}y$ is equivalent to $y = \frac{7}{30} + \frac{11}{30}x$.

```
Sage] scatter_plot([[2,1],[5,2],[7,3],[8,3]]) + plot(2/7+5/14*x,1,9) + plot(7/30+11/30*x,
1,9,color='red')
```



The explanation is that (see pictures at the beginning of this example) we are minimizing vertical offsets in one case and horizontal offsets in the other case.

In linear regression, the relationship between a dependent variable and one or more explanatory variables is modeled. If y is the dependent variable, with x the explanatory variable, then it is natural to minimize the error we make in “predicting y through x ” (vertical offsets). See example at the beginning of next class!

Example 50. A car rental company wants to predict the annual maintenance cost y (in 100USD/year) of a car using the age x (in years) of that car (as an explanatory variable). Based on the observations $(x, y) = (2, 1), (5, 2), (7, 3), (8, 3)$, predict the cost for a 4.5 year old car (using linear regression).

Solution. Once we compute the regression line $y = a + bx$ (we already did that: $y = \frac{2}{7} + \frac{5}{14}x$), our prediction is $\frac{2}{7} + \frac{5}{14} \cdot 4.5 = \frac{53}{28} \approx 1.89$, that is, 189 USD/year.

Application: multiple linear regression

In statistics, **linear regression** is an approach for modeling the relationship between a scalar dependent variable and one or more explanatory variables.

The case of one explanatory variable is called simple linear regression.

For more than one explanatory variable, the process is called multiple linear regression.

http://en.wikipedia.org/wiki/Linear_regression

The experimental data might be of the form (x_i, y_i, z_i) , where now the dependent variable z_i depends on two explanatory variables x_i, y_i (instead of just x_i).

Example 51. Set up a linear system to find values for the parameters a, b, c such that $z = a + bx + cy$ best fits some given points $(x_1, y_1, z_1), (x_2, y_2, z_2), \dots$

Solution. The equations $a + bx_i + cy_i = z_i$ translate into the system:

$$\underbrace{\begin{bmatrix} 1 & x_1 & y_1 \\ 1 & x_2 & y_2 \\ 1 & x_3 & y_3 \\ \vdots & \vdots & \vdots \end{bmatrix}}_{\text{design matrix } A} \underbrace{\begin{bmatrix} a \\ b \\ c \end{bmatrix}}_{\text{observation vector } \mathbf{z}} = \underbrace{\begin{bmatrix} z_1 \\ z_2 \\ z_3 \\ \vdots \end{bmatrix}}_{\text{observation vector } \mathbf{z}}$$

Of course, this is usually inconsistent. To find the best possible a, b, c we compute a least squares solution by solving $A^T A \begin{bmatrix} a \\ b \\ c \end{bmatrix} = A^T \mathbf{z}$.

Application: Fitting data to other curves

We can also fit the experimental data (x_i, y_i) using other curves.

Example 52. Set up a linear system to find values for the parameters a, b, c that result in the quadratic curve $y = a + bx + cx^2$ that best fits some given points $(x_1, y_1), (x_2, y_2), \dots$

Solution. $y_i \approx a + bx_i + cx_i^2$ with parameters a, b, c .

The equations $y_i = a + bx_i + cx_i^2$ in matrix form:

$$\underbrace{\begin{bmatrix} 1 & x_1 & x_1^2 \\ 1 & x_2 & x_2^2 \\ 1 & x_3 & x_3^2 \\ \vdots & \vdots & \vdots \end{bmatrix}}_{\text{design matrix } A} \underbrace{\begin{bmatrix} a \\ b \\ c \end{bmatrix}}_{\text{observation vector } \mathbf{y}} = \underbrace{\begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ \vdots \end{bmatrix}}_{\text{observation vector } \mathbf{y}}$$

Again, we determine values for a, b, c by computing a least squares solution to that system.

That is, we need to solve the system $A^T A \begin{bmatrix} a \\ b \\ c \end{bmatrix} = A^T \mathbf{y}$.

Example 53. (homework) Use Sage to find values for a, b, c that result in the quadratic curve $y = a + bx + cx^2$ that best fits the points $(0, 1), (1, 2), (2, 3), (3, -4), (4, -7), (5, -12)$.

Solution. We first input the points:

```
Sage] points = [[0,1],[1,2],[2,3],[3,-4],[4,-7],[5,-12]]
```

We set up the system described in the previous example, then determine a least-squares solution.

```
Sage] X = matrix([[1,0,0],[1,1,1],[1,2,4],[1,3,9],[1,4,16],[1,5,25]])
```


```
Sage] y = vector([1,2,3,-4,-7,-12])
```

```
Sage] (X.transpose()*X).solve_right(X.transpose()*y)
```

$$\left(\frac{3}{2}, \frac{179}{140}, -\frac{23}{28} \right)$$

Hence, the best fitting quadratic curve is $y = \frac{3}{2} + \frac{179}{140}x - \frac{23}{28}x^2$. Here's a plot:

```
Sage] scatter_plot(points) + plot(3/2+179/140*x-23/28*x^2,0,5,color='red')
```



Advanced comment. If you are comfortable with Python, you can avoid typing out X and \mathbf{y} :
[The plot command above now won't work anymore because we are overwriting x with numbers.]

```
Sage] X = matrix([[1,x,x^2] for x,y in points])
```

```
Sage] y = vector([y for x,y in points])
```


More on orthogonality

Example 54. (review) Find the least squares solution to $Ax = b$, where

$$A = \begin{bmatrix} 4 & 0 \\ 0 & 2 \\ 1 & 1 \end{bmatrix}, \quad b = \begin{bmatrix} 2 \\ 0 \\ 11 \end{bmatrix}.$$

Solution. First, $A^T A = \begin{bmatrix} 4 & 0 & 1 \\ 0 & 2 & 1 \end{bmatrix} \begin{bmatrix} 4 & 0 \\ 0 & 2 \\ 1 & 1 \end{bmatrix} = \begin{bmatrix} 17 & 1 \\ 1 & 5 \end{bmatrix}$ and $A^T b = \begin{bmatrix} 4 & 0 & 1 \\ 0 & 2 & 1 \end{bmatrix} \begin{bmatrix} 2 \\ 0 \\ 11 \end{bmatrix} = \begin{bmatrix} 19 \\ 11 \end{bmatrix}$.

Hence, the normal equations $A^T A \hat{x} = A^T b$ take the form $\begin{bmatrix} 17 & 1 \\ 1 & 5 \end{bmatrix} \hat{x} = \begin{bmatrix} 19 \\ 11 \end{bmatrix}$. Solving, we find $\hat{x} = \begin{bmatrix} 1 \\ 2 \end{bmatrix}$.

Check. The error $A\hat{x} - b = \begin{bmatrix} 2 \\ 4 \\ -8 \end{bmatrix}$ is indeed orthogonal to $\text{col}(A)$. Because $\begin{bmatrix} 2 \\ 4 \\ -8 \end{bmatrix} \cdot \begin{bmatrix} 4 \\ 0 \\ 1 \end{bmatrix} = 0$ and $\begin{bmatrix} 2 \\ 4 \\ -8 \end{bmatrix} \cdot \begin{bmatrix} 0 \\ 2 \\ 1 \end{bmatrix} = 0$.

Orthogonal projections

The **(orthogonal) projection** \hat{b} of a vector b onto a subspace W is the vector in W closest to b .

We can compute \hat{b} as follows:

- Write $W = \text{col}(A)$ for some matrix A .
- Then $\hat{b} = A\hat{x}$ where \hat{x} is a least squares solution to $Ax = b$. (i.e. \hat{x} solves $A^T A \hat{x} = A^T b$)

Why? Why is $A\hat{x}$ the projection of b onto $\text{col}(A)$?

Because, if \hat{x} is a least squares solution then $A\hat{x} - b$ is as small as possible (and any element in $\text{col}(A)$ is of the form Ax for some x).

Note. This is a recipe for computing any orthogonal projection! That's because every subspace W can be written as $\text{col}(A)$ for some choice of the matrix A (take, for instance, A so that its columns are a basis for W).

Assuming $A^T A$ is invertible (which, as discussed in the lemma below, is automatically the case if the columns of A are independent), we have $\hat{x} = (A^T A)^{-1} A^T b$ and hence:

(projection matrix) The projection \hat{b} of b onto $\text{col}(A)$ is (assuming cols of A are independent)

$$\hat{b} = \underbrace{A(A^T A)^{-1} A^T}_P b.$$

The matrix $P = A(A^T A)^{-1} A^T$ is the **projection matrix** for projecting onto $\text{col}(A)$.

Example 55.

(a) What is the orthogonal projection of $\begin{bmatrix} 2 \\ 0 \\ 11 \end{bmatrix}$ onto $W = \text{span} \left\{ \begin{bmatrix} 4 \\ 0 \\ 1 \end{bmatrix}, \begin{bmatrix} 0 \\ 2 \\ 1 \end{bmatrix} \right\}$?

(b) What is the matrix P for projecting onto $W = \text{span} \left\{ \begin{bmatrix} 4 \\ 0 \\ 1 \end{bmatrix}, \begin{bmatrix} 0 \\ 2 \\ 1 \end{bmatrix} \right\}$?

(c) **(once more)** Using P , what is the orthogonal projection of $\begin{bmatrix} 2 \\ 0 \\ 11 \end{bmatrix}$ onto W ?

(d) Using P , what is the orthogonal projection of $\begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}$ onto W ?

Solution.

- (a) In other words, what is the orthogonal projection of $\mathbf{b} = \begin{bmatrix} 2 \\ 0 \\ 11 \end{bmatrix}$ onto $\text{col}(A)$ with $A = \begin{bmatrix} 4 & 0 \\ 0 & 2 \\ 1 & 1 \end{bmatrix}$.

In Example 54, we found that the system $A\mathbf{x} = \mathbf{b}$ has the least squares solution $\mathbf{x} = \begin{bmatrix} 1 \\ 2 \end{bmatrix}$.

The projection $\hat{\mathbf{b}}$ of \mathbf{b} onto $\text{col}(A)$ thus is $A\hat{\mathbf{x}} = \begin{bmatrix} 4 & 0 \\ 0 & 2 \\ 1 & 1 \end{bmatrix} \begin{bmatrix} 1 \\ 2 \end{bmatrix} = 1 \begin{bmatrix} 4 \\ 0 \\ 1 \end{bmatrix} + 2 \begin{bmatrix} 0 \\ 2 \\ 1 \end{bmatrix} = \begin{bmatrix} 4 \\ 4 \\ 3 \end{bmatrix}$.

Check. The error $\hat{\mathbf{b}} - \mathbf{b} = \begin{bmatrix} 2 \\ 4 \\ -8 \end{bmatrix}$ needs to be orthogonal to $\text{col}(A)$. Indeed: $\begin{bmatrix} 2 \\ 4 \\ -8 \end{bmatrix} \cdot \begin{bmatrix} 4 \\ 0 \\ 1 \end{bmatrix} = 0$ and $\begin{bmatrix} 2 \\ 4 \\ -8 \end{bmatrix} \cdot \begin{bmatrix} 0 \\ 2 \\ 1 \end{bmatrix} = 0$.

- (b) Note that $W = \text{col}(A)$ for $A = \begin{bmatrix} 4 & 0 \\ 0 & 2 \\ 1 & 1 \end{bmatrix}$ and that $A^T A = \begin{bmatrix} 17 & 1 \\ 1 & 5 \end{bmatrix}$. Thus $(A^T A)^{-1} = \frac{1}{84} \begin{bmatrix} 5 & -1 \\ -1 & 17 \end{bmatrix}$.

$$P = A(A^T A)^{-1} A^T = \frac{1}{84} \begin{bmatrix} 4 & 0 \\ 0 & 2 \\ 1 & 1 \end{bmatrix} \begin{bmatrix} 5 & -1 \\ -1 & 17 \end{bmatrix} \begin{bmatrix} 4 & 0 & 1 \\ 0 & 2 & 1 \end{bmatrix} = \frac{1}{21} \begin{bmatrix} 20 & -2 & 4 \\ -2 & 17 & 8 \\ 4 & 8 & 5 \end{bmatrix}$$

- (c) The orthogonal projection of $\begin{bmatrix} 2 \\ 0 \\ 11 \end{bmatrix}$ onto W is $P \begin{bmatrix} 2 \\ 0 \\ 11 \end{bmatrix} = \frac{1}{21} \begin{bmatrix} 20 & -2 & 4 \\ -2 & 17 & 8 \\ 4 & 8 & 5 \end{bmatrix} \begin{bmatrix} 2 \\ 0 \\ 11 \end{bmatrix} = \frac{1}{21} \begin{bmatrix} 84 \\ 84 \\ 63 \end{bmatrix} = \begin{bmatrix} 4 \\ 4 \\ 3 \end{bmatrix}$.

Note. Of course, that agrees with what our computations in the first part. Note that computing P is more work than what we did in in the first part. However, after having computed P once, we can easily project many vectors onto W .

- (d) The orthogonal projection of $\begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}$ onto W is $P \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} = \frac{1}{21} \begin{bmatrix} 20 & -2 & 4 \\ -2 & 17 & 8 \\ 4 & 8 & 5 \end{bmatrix} \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} = \frac{1}{21} \begin{bmatrix} 20 \\ -2 \\ 4 \end{bmatrix}$.

Check. The error $\begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} - \frac{1}{21} \begin{bmatrix} 20 \\ -2 \\ 4 \end{bmatrix} = \frac{1}{21} \begin{bmatrix} 1 \\ 2 \\ -4 \end{bmatrix}$ is indeed orthogonal to both $\begin{bmatrix} 4 \\ 0 \\ 1 \end{bmatrix}$ and $\begin{bmatrix} 0 \\ 2 \\ 1 \end{bmatrix}$.

Example 56. (extra)

- (a) What is the matrix P for projecting onto $W = \text{span} \left\{ \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}, \begin{bmatrix} 1 \\ -1 \\ 1 \end{bmatrix} \right\}$?

- (b) Using the projection matrix, project $\begin{bmatrix} 2 \\ 3 \\ 3 \end{bmatrix}$ onto $W = \text{span} \left\{ \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}, \begin{bmatrix} 1 \\ -1 \\ 1 \end{bmatrix} \right\}$.

Solution.

- (a) Choosing $A = \begin{bmatrix} 1 & 1 \\ 1 & -1 \\ 1 & 1 \end{bmatrix}$, the projection matrix P is $A(A^T A)^{-1} A^T = \begin{bmatrix} 1 & 1 \\ 1 & -1 \\ 1 & 1 \end{bmatrix} \begin{bmatrix} 3 & 1 \\ 1 & 3 \end{bmatrix}^{-1} \begin{bmatrix} 1 & 1 & 1 \\ 1 & -1 & 1 \end{bmatrix}$
 $= \begin{bmatrix} 1 & 1 \\ 1 & -1 \\ 1 & 1 \end{bmatrix} \frac{1}{8} \begin{bmatrix} 3 & -1 \\ -1 & 3 \end{bmatrix} \begin{bmatrix} 1 & 1 & 1 \\ 1 & -1 & 1 \end{bmatrix} = \frac{1}{8} \begin{bmatrix} 1 & 1 \\ 1 & -1 \\ 1 & 1 \end{bmatrix} \begin{bmatrix} 2 & 4 & 2 \\ 2 & -4 & 2 \end{bmatrix} = \frac{1}{2} \begin{bmatrix} 1 & 0 & 1 \\ 0 & 2 & 0 \\ 1 & 0 & 1 \end{bmatrix}$.

Comment. We can choose A in any way such that its columns are a basis for W . The final projection matrix will always be the same.

- (b) The projection is $\frac{1}{2} \begin{bmatrix} 1 & 0 & 1 \\ 0 & 2 & 0 \\ 1 & 0 & 1 \end{bmatrix} \begin{bmatrix} 2 \\ 3 \\ 3 \end{bmatrix} = \frac{1}{2} \begin{bmatrix} 5 \\ 6 \\ 5 \end{bmatrix}$.

Check. The error $\begin{bmatrix} 2 \\ 3 \\ 3 \end{bmatrix} - \frac{1}{2} \begin{bmatrix} 5 \\ 6 \\ 5 \end{bmatrix} = \begin{bmatrix} -1/2 \\ 0 \\ 1/2 \end{bmatrix}$ is indeed orthogonal to W .

Example 57. (extra)

(a) What is the orthogonal projection of $\begin{bmatrix} 1 \\ 2 \\ 2 \end{bmatrix}$ onto $\text{span}\left\{\begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix}, \begin{bmatrix} 1 \\ 1 \\ -1 \end{bmatrix}\right\}$?

(b) What is the orthogonal projection of $\begin{bmatrix} 1 \\ 2 \\ 2 \end{bmatrix}$ onto $\text{span}\left\{\begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix}\right\}$?

Solution. (final answer only) The projections are $\left(\frac{11}{6}, \frac{1}{3}, \frac{7}{6}\right)^T$ and $\left(\frac{3}{2}, 0, \frac{3}{2}\right)^T$.

Lemma 58. If the columns of a matrix A are independent, then $A^T A$ is invertible.

Proof. Assume $A^T A$ is not invertible, so that $A^T A \mathbf{x} = \mathbf{0}$ for some $\mathbf{x} \neq \mathbf{0}$. Multiply both sides with \mathbf{x}^T to get

$$\mathbf{x}^T A^T A \mathbf{x} = (A \mathbf{x})^T A \mathbf{x} = \|A \mathbf{x}\|^2 = 0,$$

which implies that $A \mathbf{x} = \mathbf{0}$. Since the columns of A are independent, this shows that $\mathbf{x} = \mathbf{0}$. A contradiction! \square

Example 59. If P is a projection matrix, then what is P^2 ?

For instance. For P as in Example 56, $P^2 = \frac{1}{4} \begin{bmatrix} 1 & 0 & 1 \\ 0 & 2 & 0 \\ 1 & 0 & 1 \end{bmatrix}^2 = \frac{1}{2} \begin{bmatrix} 1 & 0 & 1 \\ 0 & 2 & 0 \\ 1 & 0 & 1 \end{bmatrix} = P$.

Solution. Can you see why it is always true that $P^2 = P$?

[Recall that P projects a vector onto a space W (actually, $W = \text{col}(P)$). Hence P^2 takes a vector \mathbf{b} , projects it onto W to get $\hat{\mathbf{b}}$, and then projects $\hat{\mathbf{b}}$ onto W again. But the projection of $\hat{\mathbf{b}}$ onto W is just $\hat{\mathbf{b}}$ (why?!), so that P^2 always has the exact same effect as P . Therefore, $P^2 = P$.]

Example 60. True or false? If P is the matrix for projecting onto W , then $W = \text{col}(P)$.

Solution. True!

Why? The columns of P are the projections of the standard basis vectors and hence in W . On the other hand, for any vector \mathbf{w} in W , we have $P \mathbf{w} = \mathbf{w}$ so that \mathbf{w} is a combination of the columns of P .

[This may take several readings to digest but do read (or ask) until it makes sense!]

In particular. $\text{rank}(P) = \dim W$ (because, for any matrix, $\text{rank}(A) = \dim \text{col}(A)$)

Review. The **projection matrix** for projecting onto $\text{col}(A)$ is $P = A(A^T A)^{-1} A^T$.

Projecting onto 1-dimensional spaces

When we project onto a 1-dimensional space $\text{span}\{w\}$, we usually just say that we are projecting onto w .

The (orthogonal) projection of v onto w is $\frac{w \cdot v}{\|w\|^2} w$.

Why? Replace b with v and A with w in our general projection matrix formula to get $w(w^T w)^{-1} w^T v$, which equals $\frac{w \cdot v}{\|w\|^2} w$ (note that $w^T v = w \cdot v$ and $w^T w = \|w\|^2$ are scalars).

Comment. If you have taken Calculus 3, you have seen that formula before. Most likely, you were deriving it using angles at that time. Namely, the dot product has the following connection to angles:

$v \cdot w = \|v\| \|w\| \cos\theta$ where $\theta \in [0, \pi]$ is the angle between v and w

Why? You can derive this by repeating what we did, right after Definition 29 to show that v and w are orthogonal if and only if $v \cdot w = 0$. Just replace Pythagoras with the law of cosines ($c^2 = a^2 + b^2 - 2ab \cos\theta$ holds in any triangle!).


Two obvious cases. Observe that the cases $\theta = 0$ and $\theta = 90^\circ$ are clearly true.

We will not discuss angles much further in this class. Just in case it is helpful, here is the typical argument given in Calculus 3 to determine the projection $\text{proj}_w v$ of v onto w :

From the sketch, we see that “error” = $v - \text{proj}_w v$ and that this error is orthogonal to w .

Basic trigonometry tells us that the length of $\text{proj}_w v$ is $\|v\| \cos\theta$. Hence:

$$\begin{aligned} \text{proj}_w v &= \underbrace{\|v\| \cos\theta}_{\text{length}} \underbrace{\frac{w}{\|w\|}}_{\text{direction}} \\ &= \frac{\|v\| \|w\| \cos\theta}{\|w\|} \frac{w}{\|w\|} = \left(\frac{v \cdot w}{\|w\|^2} \right) w \end{aligned}$$



Orthogonal bases

Review. Vectors v_1, \dots, v_n are a **basis** for V .

$\iff V = \text{span}\{v_1, \dots, v_n\}$ and v_1, \dots, v_n are linearly independent.

\iff Any vector w in V can be written as $w = c_1 v_1 + \dots + c_n v_n$ in a unique way.

The latter is the practical reason why we care so much about bases!

V could be some abstract vector space (of polynomials or Fourier series), meaning that vectors are abstract objects and not just our usual column vectors. However, as soon as we pick a basis of V , then we can represent every (abstract) vector w by the (usual) column vector $(c_1, c_2, \dots, c_n)^T$.

This means all of our results can be used, too, when working with these abstract spaces!

Definition 61. A basis v_1, \dots, v_n of a vector space V is an **orthogonal basis** if the vectors are (pairwise) orthogonal. If, in addition, the basis vectors have length 1, then this is called an **orthonormal basis**.

Example 62. The standard basis $\begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}$ is an orthonormal basis for \mathbb{R}^3 .

Example 63. Are the vectors $\begin{bmatrix} 1 \\ -1 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}$ an orthogonal basis for \mathbb{R}^3 ? Is it orthonormal?

Solution. $\begin{bmatrix} 1 \\ -1 \\ 0 \end{bmatrix} \cdot \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix} = 0, \begin{bmatrix} 1 \\ -1 \\ 0 \end{bmatrix} \cdot \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} = 0, \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix} \cdot \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} = 0.$

So, this is an orthogonal basis.

On the other hand, the vectors do not all have length 1, so that this basis is not orthonormal.

Note. Orthogonal vectors are always linearly independent (see next class). Here, this certifies that the three vectors are linearly independent (and hence a basis for \mathbb{R}^3).

Normalize the vectors to produce an orthonormal basis.

Solution.

$\begin{bmatrix} 1 \\ -1 \\ 0 \end{bmatrix}$ has length $\sqrt{\begin{bmatrix} 1 \\ -1 \\ 0 \end{bmatrix} \cdot \begin{bmatrix} 1 \\ -1 \\ 0 \end{bmatrix}} = \sqrt{2} \implies$ normalized: $\frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ -1 \\ 0 \end{bmatrix}$

$\begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}$ has length $\sqrt{\begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix} \cdot \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}} = \sqrt{2} \implies$ normalized: $\frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}$

$\begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}$ has length $\sqrt{\begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \cdot \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}} = 1 \implies$ is already normalized: $\begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}$

The resulting orthonormal basis is $\frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ -1 \\ 0 \end{bmatrix}, \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}.$

Theorem 64. Suppose that $\mathbf{v}_1, \dots, \mathbf{v}_n$ are nonzero and pairwise orthogonal. Then $\mathbf{v}_1, \dots, \mathbf{v}_n$ are linearly independent.

Proof. Suppose that $c_1\mathbf{v}_1 + \dots + c_n\mathbf{v}_n = \mathbf{0}$. In order to show that $\mathbf{v}_1, \dots, \mathbf{v}_n$ are independent, we need to show that $c_1 = c_2 = \dots = c_n = 0$.

Take the dot product of \mathbf{v}_1 with both sides:

$$\begin{aligned} 0 &= \mathbf{v}_1 \cdot (c_1\mathbf{v}_1 + \dots + c_n\mathbf{v}_n) \\ &= c_1\mathbf{v}_1 \cdot \mathbf{v}_1 + c_2\mathbf{v}_1 \cdot \mathbf{v}_2 + \dots + c_n\mathbf{v}_1 \cdot \mathbf{v}_n \\ &= c_1\mathbf{v}_1 \cdot \mathbf{v}_1 = c_1\|\mathbf{v}_1\|^2 \end{aligned}$$

But $\|\mathbf{v}_1\| \neq 0$ and hence $c_1 = 0$. Likewise, we find $c_2 = 0, \dots, c_n = 0$. Hence, the vectors are independent. \square

Comment. Note that this result is intuitively obvious: if the vectors were linearly dependent, then one of them could be written as a linear combination of the others. However, all these other vectors (and hence any combination of them) are orthogonal to it.

Orthogonal projections if we have an orthogonal basis

Lemma 65. (orthogonal projection if we have an orthogonal basis)

If $\mathbf{v}_1, \dots, \mathbf{v}_n$ are orthogonal, then the orthogonal projection of \mathbf{w} onto $\text{span}\{\mathbf{v}_1, \dots, \mathbf{v}_n\}$ is

$$\hat{\mathbf{w}} = \underbrace{\frac{\mathbf{w} \cdot \mathbf{v}_1}{\mathbf{v}_1 \cdot \mathbf{v}_1}}_{\text{proj of } \mathbf{w} \text{ onto } \mathbf{v}_1} \mathbf{v}_1 + \dots + \underbrace{\frac{\mathbf{w} \cdot \mathbf{v}_n}{\mathbf{v}_n \cdot \mathbf{v}_n}}_{\text{proj of } \mathbf{w} \text{ onto } \mathbf{v}_n} \mathbf{v}_n.$$

Proof. It suffices to show that the error $\mathbf{w} - \hat{\mathbf{w}}$ is orthogonal to each \mathbf{v}_i . Indeed:

$$(\mathbf{w} - \hat{\mathbf{w}}) \cdot \mathbf{v}_i = \left(\mathbf{w} - \frac{\mathbf{w} \cdot \mathbf{v}_1}{\mathbf{v}_1 \cdot \mathbf{v}_1} \mathbf{v}_1 - \dots - \frac{\mathbf{w} \cdot \mathbf{v}_n}{\mathbf{v}_n \cdot \mathbf{v}_n} \mathbf{v}_n \right) \cdot \mathbf{v}_i = \mathbf{w} \cdot \mathbf{v}_i - \frac{\mathbf{w} \cdot \mathbf{v}_i}{\mathbf{v}_i \cdot \mathbf{v}_i} \mathbf{v}_i \cdot \mathbf{v}_i = 0.$$

Alternatively, can you deduce the formula (say, in the case of an orthonormal basis) from our earlier formula for the projection matrix? \square

Important consequence. If $\mathbf{v}_1, \dots, \mathbf{v}_n$ is an orthogonal basis of V , and \mathbf{w} is in V , then

$$\mathbf{w} = c_1 \mathbf{v}_1 + \dots + c_n \mathbf{v}_n \quad \text{with} \quad c_j = \frac{\mathbf{w} \cdot \mathbf{v}_j}{\mathbf{v}_j \cdot \mathbf{v}_j}.$$

If the $\mathbf{v}_1, \dots, \mathbf{v}_n$ are a basis, but not orthogonal, then we have to solve a system of equations to find the c_i . That is a lot more work than simply computing a few dot products.

Note. In other words, \mathbf{w} decomposes as the sum of its projections onto each basis vector.

Note. If $\mathbf{v}_1, \dots, \mathbf{v}_n$ are orthonormal, then the denominators are all 1.

Example 66. What is the projection of $\begin{bmatrix} 3 \\ 7 \\ 4 \end{bmatrix}$ onto $W = \text{span}\{\mathbf{v}_1, \mathbf{v}_2\}$ with $\mathbf{v}_1 = \begin{bmatrix} 1 \\ -1 \\ 0 \end{bmatrix}$, $\mathbf{v}_2 = \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}$?

Comment. We know how to do this using least squares. (Do it for practice!)

However, realizing that \mathbf{v}_1 and \mathbf{v}_2 are orthogonal makes things easier.

[Actually, here, it is obvious what the projection is going to be if we realized that W is the x - y -plane.]

Solution. (using orthogonality) Because \mathbf{v}_1 and \mathbf{v}_2 are orthogonal, the projection is

$$\underbrace{\begin{bmatrix} 3 \\ 7 \\ 4 \end{bmatrix} \cdot \begin{bmatrix} 1 \\ -1 \\ 0 \end{bmatrix}}_{\text{projection onto } \mathbf{v}_1} \begin{bmatrix} 1 \\ -1 \\ 0 \end{bmatrix} + \underbrace{\begin{bmatrix} 3 \\ 7 \\ 4 \end{bmatrix} \cdot \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}}_{\text{projection onto } \mathbf{v}_2} \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix} = \frac{-4}{2} \begin{bmatrix} 1 \\ -1 \\ 0 \end{bmatrix} + \frac{10}{2} \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix} = \begin{bmatrix} 3 \\ 7 \\ 0 \end{bmatrix}.$$

Important note. Note that, at this point, we can easily extend $\begin{bmatrix} 1 \\ -1 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}$ to an orthogonal basis of \mathbb{R}^3 :

That is because the error $\begin{bmatrix} 3 \\ 7 \\ 4 \end{bmatrix} - \begin{bmatrix} 3 \\ 7 \\ 0 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 4 \end{bmatrix}$ is orthogonal to both of the existing basis vectors.

Therefore $\begin{bmatrix} 1 \\ -1 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 4 \end{bmatrix}$ is an orthogonal basis of \mathbb{R}^3 .

This observation underlies the Gram-Schmidt process, which we will discuss next class.

Example 67. Express $\begin{bmatrix} 3 \\ 7 \\ 4 \end{bmatrix}$ in terms of the basis $\begin{bmatrix} 1 \\ -1 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}$.

Solution. Because $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3$ is an orthogonal basis of \mathbb{R}^3 , we get (much as in the previous example):

$$\begin{aligned} \begin{bmatrix} 3 \\ 7 \\ 4 \end{bmatrix} &= c_1 \begin{bmatrix} 1 \\ -1 \\ 0 \end{bmatrix} + c_2 \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix} + c_3 \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \\ &= \underbrace{\begin{bmatrix} 3 \\ 7 \\ 4 \end{bmatrix} \cdot \begin{bmatrix} 1 \\ -1 \\ 0 \end{bmatrix}}_{\text{projection of } \mathbf{x} \text{ onto } \mathbf{v}_1} \begin{bmatrix} 1 \\ -1 \\ 0 \end{bmatrix} + \underbrace{\begin{bmatrix} 3 \\ 7 \\ 4 \end{bmatrix} \cdot \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}}_{\text{projection of } \mathbf{x} \text{ onto } \mathbf{v}_2} \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix} + \underbrace{\begin{bmatrix} 3 \\ 7 \\ 4 \end{bmatrix} \cdot \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}}_{\text{projection of } \mathbf{x} \text{ onto } \mathbf{v}_3} \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \\ &= \frac{-4}{2} \begin{bmatrix} 1 \\ -1 \\ 0 \end{bmatrix} + \frac{10}{2} \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix} + \frac{4}{1} \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \end{aligned}$$

Because we spelled out all the details this looks more involved than it is. We only computed 6 dot products!

Alternative. We could have solved $\begin{bmatrix} 1 & 1 & 0 \\ -1 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} c_1 \\ c_2 \\ c_3 \end{bmatrix} = \begin{bmatrix} 3 \\ 7 \\ 4 \end{bmatrix}$ to also find $\begin{bmatrix} c_1 \\ c_2 \\ c_3 \end{bmatrix} = \begin{bmatrix} -2 \\ 5 \\ 4 \end{bmatrix}$.

The numbers are particularly easy here but in general, to find this solution, we have to go through the entire process of Gaussian elimination. On the other hand, if we have an orthogonal basis, the former approach requires less work, because it is just computing a few dot products.

Example 68. Express $\begin{bmatrix} 3 \\ 7 \\ 4 \end{bmatrix}$ in terms of the basis $\begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix}, \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix}$.

Solution. This is not an orthogonal basis, so we cannot proceed as in the previous example.

To write $\begin{bmatrix} 3 \\ 7 \\ 4 \end{bmatrix} = c_1 \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix} + c_2 \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix} + c_3 \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix}$, we need to solve $\begin{bmatrix} 1 & 0 & 1 \\ 1 & 1 & 0 \\ 0 & 1 & 1 \end{bmatrix} \begin{bmatrix} c_1 \\ c_2 \\ c_3 \end{bmatrix} = \begin{bmatrix} 3 \\ 7 \\ 4 \end{bmatrix}$.

Solving that system (do it!), we find $\begin{bmatrix} c_1 \\ c_2 \\ c_3 \end{bmatrix} = \begin{bmatrix} 3 \\ 4 \\ 0 \end{bmatrix}$.

Review. If v_1, \dots, v_n are orthogonal, the orthogonal projection of w onto $\text{span}\{v_1, \dots, v_n\}$ is

$$\hat{w} = \frac{w \cdot v_1}{v_1 \cdot v_1} v_1 + \dots + \frac{w \cdot v_n}{v_n \cdot v_n} v_n.$$

Example 69.

(a) Project $\begin{bmatrix} 3 \\ 2 \\ 1 \end{bmatrix}$ onto $W = \text{span}\left\{ \begin{bmatrix} 1 \\ 2 \\ 1 \end{bmatrix}, \begin{bmatrix} 2 \\ -1 \\ 0 \end{bmatrix} \right\}$.

(b) Express $\begin{bmatrix} 3 \\ 2 \\ 1 \end{bmatrix}$ in terms of the basis $\begin{bmatrix} 1 \\ 2 \\ 1 \end{bmatrix}, \begin{bmatrix} 2 \\ -1 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 2 \\ -5 \end{bmatrix}$.

Solution.

(a) We note that the vectors $\begin{bmatrix} 1 \\ 2 \\ 1 \end{bmatrix}, \begin{bmatrix} 2 \\ -1 \\ 0 \end{bmatrix}$ are orthogonal to each other.

Therefore, the projection can be computed as $\frac{\begin{bmatrix} 3 \\ 2 \\ 1 \end{bmatrix} \cdot \begin{bmatrix} 1 \\ 2 \\ 1 \end{bmatrix}}{\begin{bmatrix} 1 \\ 2 \\ 1 \end{bmatrix} \cdot \begin{bmatrix} 1 \\ 2 \\ 1 \end{bmatrix}} \begin{bmatrix} 1 \\ 2 \\ 1 \end{bmatrix} + \frac{\begin{bmatrix} 3 \\ 2 \\ 1 \end{bmatrix} \cdot \begin{bmatrix} 2 \\ -1 \\ 0 \end{bmatrix}}{\begin{bmatrix} 2 \\ -1 \\ 0 \end{bmatrix} \cdot \begin{bmatrix} 2 \\ -1 \\ 0 \end{bmatrix}} \begin{bmatrix} 2 \\ -1 \\ 0 \end{bmatrix} = \frac{8}{6} \begin{bmatrix} 1 \\ 2 \\ 1 \end{bmatrix} + \frac{4}{5} \begin{bmatrix} 2 \\ -1 \\ 0 \end{bmatrix}$.

Comment. If we didn't have an orthogonal basis for $W = \text{col}\left(\begin{bmatrix} 1 & 2 \\ 2 & -1 \\ 1 & 0 \end{bmatrix}\right)$, then we would have to solve the least squares problem $\begin{bmatrix} 1 & 2 \\ 2 & -1 \\ 1 & 0 \end{bmatrix} x = \begin{bmatrix} 3 \\ 2 \\ 1 \end{bmatrix}$ instead to get the same final result (with more work).

(b) Note that this basis is orthogonal! Therefore, we can compute $\begin{bmatrix} 3 \\ 2 \\ 1 \end{bmatrix} = \frac{8}{6} \begin{bmatrix} 1 \\ 2 \\ 1 \end{bmatrix} + \frac{4}{5} \begin{bmatrix} 2 \\ -1 \\ 0 \end{bmatrix} + \frac{5}{30} \begin{bmatrix} 1 \\ 2 \\ -5 \end{bmatrix}$.

(We proceed exactly as in the previous part to compute each coefficient as a quotient of dot products.)

Gram-Schmidt

(Gram-Schmidt orthogonalization)

Given a basis w_1, w_2, \dots for W , we produce an orthogonal basis q_1, q_2, \dots for W as follows:

- $q_1 = w_1$
- $q_2 = w_2 - \left(\begin{matrix} \text{projection of} \\ w_2 \text{ onto } q_1 \end{matrix} \right)$
- $q_3 = w_3 - \left(\begin{matrix} \text{projection of} \\ w_3 \text{ onto } q_1 \end{matrix} \right) - \left(\begin{matrix} \text{projection of} \\ w_3 \text{ onto } q_2 \end{matrix} \right)$
- $q_4 = \dots$

Note. Since q_1, q_2 are orthogonal, $\left(\begin{matrix} \text{projection of} \\ w_3 \text{ onto } \text{span}\{q_1, q_2\} \end{matrix} \right) = \left(\begin{matrix} \text{projection of} \\ w_3 \text{ onto } q_1 \end{matrix} \right) + \left(\begin{matrix} \text{projection of} \\ w_3 \text{ onto } q_2 \end{matrix} \right)$.

Important comment. When working numerically on a computer it actually saves time to compute an orthonormal basis q_1, q_2, \dots by the same approach but always normalizing each q_i along the way. The reason this saves time is that now the projections onto q_i only require a single dot product (instead of two). This is called **Gram-Schmidt orthonormalization**. When working by hand, it is usually simpler to wait until the end to normalize (so as to avoid working with square roots).

Note. When normalizing, the orthonormal basis q_1, q_2, \dots is the unique one (up to \pm signs) with the property that $\text{span}\{q_1, q_2, \dots, q_k\} = \text{span}\{w_1, w_2, \dots, w_k\}$ for all $k = 1, 2, \dots$

Example 70. Using Gram–Schmidt, find an orthogonal basis for $W = \text{span}\left\{\begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}, \begin{bmatrix} 1 \\ -1 \\ 1 \end{bmatrix}\right\}$.

Solution. We already have the basis $\mathbf{w}_1 = \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}$, $\mathbf{w}_2 = \begin{bmatrix} 1 \\ -1 \\ 1 \end{bmatrix}$ for W . However, that basis is not orthogonal.

We can construct an orthogonal basis $\mathbf{q}_1, \mathbf{q}_2$ for W as follows:

- $\mathbf{q}_1 = \mathbf{w}_1 = \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}$
- $\mathbf{q}_2 = \mathbf{w}_2 - \left(\text{projection of } \mathbf{w}_2 \text{ onto } \mathbf{q}_1\right) = \begin{bmatrix} 1 \\ -1 \\ 1 \end{bmatrix} - \frac{1}{3} \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} = \frac{1}{3} \begin{bmatrix} 2 \\ -4 \\ 2 \end{bmatrix}$

Note. \mathbf{q}_2 is the error of the projection of \mathbf{w}_2 onto \mathbf{q}_1 . This guarantees that it is orthogonal to \mathbf{q}_1 . On the other hand, since \mathbf{q}_2 is a combination of \mathbf{w}_2 and \mathbf{q}_1 , we know that \mathbf{q}_2 actually is in W .

We have thus found the orthogonal basis $\begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}, \frac{2}{3} \begin{bmatrix} 1 \\ -2 \\ 1 \end{bmatrix}$ for W (if we like, we can, of course, drop that $\frac{2}{3}$).

Important comment. By normalizing, we get an orthonormal basis for W : $\frac{1}{\sqrt{3}} \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}, \frac{1}{\sqrt{6}} \begin{bmatrix} 1 \\ -2 \\ 1 \end{bmatrix}$.

Practical comment. When implementing Gram–Schmidt on a computer, it is beneficial (slightly less work) to normalize each \mathbf{q}_i during the Gram–Schmidt process. This typically introduces square roots, which is why normalizing at the end is usually preferable when working by hand.

Comment. There are, of course, many orthogonal bases $\mathbf{q}_1, \mathbf{q}_2$ for W . Up to the length of the vectors, ours is the unique one with the property that $\text{span}\{\mathbf{q}_1\} = \text{span}\{\mathbf{w}_1\}$ and $\text{span}\{\mathbf{q}_1, \mathbf{q}_2\} = \text{span}\{\mathbf{w}_1, \mathbf{w}_2\}$.

A matrix Q has orthonormal columns $\iff Q^T Q = I$

Why? Let $\mathbf{q}_1, \mathbf{q}_2, \dots$ be the columns of Q . By the way matrix multiplication works, the entries of $Q^T Q$ are dot products of these columns:

$$\begin{bmatrix} - & \mathbf{q}_1^T & - \\ - & \mathbf{q}_2^T & - \\ & \vdots & \end{bmatrix} \begin{bmatrix} | & | & \dots \\ \mathbf{q}_1 & \mathbf{q}_2 & \dots \\ | & | & \dots \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & \ddots \end{bmatrix}$$

Hence, $Q^T Q = I$ if and only if $\mathbf{q}_i^T \mathbf{q}_j = 0$ (that is, the columns are orthogonal), for $i \neq j$, and $\mathbf{q}_i^T \mathbf{q}_i = 1$ (that is, the columns are normalized).

Example 71. $Q = \begin{bmatrix} 1/\sqrt{3} & 1/\sqrt{6} \\ 1/\sqrt{3} & -2/\sqrt{6} \\ 1/\sqrt{3} & 1/\sqrt{6} \end{bmatrix}$ obtained from Example 70 satisfies $Q^T Q = I$.

The QR decomposition

Just like the LU decomposition encodes the steps of Gaussian elimination, the QR decomposition encodes the steps of Gram–Schmidt.

(QR decomposition) Every $m \times n$ matrix A of rank n can be decomposed as $A = QR$, where

- Q has orthonormal columns, $(m \times n)$
- R is upper triangular and invertible. $(n \times n)$

How to find Q and R ?

- Gram–Schmidt orthonormalization on (columns of) A , to get (columns of) Q
- $R = Q^T A$
Why? If $A = QR$, then $Q^T A = Q^T QR$ which simplifies to $R = Q^T A$ (since $Q^T Q = I$).

The decomposition $A = QR$ is unique if we require the diagonal entries of R to be positive (and this is exactly what happens when applying Gram–Schmidt).

Practical comment. Actually, no extra work is needed for computing R . All of its entries have been computed during Gram–Schmidt.

Variations. We can also arrange things so that Q is an $m \times m$ **orthogonal** matrix (this means Q is square and has orthonormal columns) and R a $m \times n$ upper triangular matrix. This is a tiny bit more work (and not required for many applications): we need to complement “our” Q with additional orthonormal columns and add corresponding zero rows to R . For square matrices this makes no difference.

Example 72. Determine the QR decomposition of $A = \begin{bmatrix} 1 & 1 \\ 1 & -1 \\ 1 & 1 \end{bmatrix}$.

Solution. The first step is Gram–Schmidt orthonormalization on the columns of A . We then use the resulting orthonormal vectors (they need to be normalized!) as the columns of Q .

We already did Gram–Schmidt in Example 70: from that work, we have $Q = \begin{bmatrix} 1/\sqrt{3} & 1/\sqrt{6} \\ 1/\sqrt{3} & -2/\sqrt{6} \\ 1/\sqrt{3} & 1/\sqrt{6} \end{bmatrix}$.

Hence, $R = Q^T A = \begin{bmatrix} 1/\sqrt{3} & 1/\sqrt{3} & 1/\sqrt{3} \\ 1/\sqrt{6} & -2/\sqrt{6} & 1/\sqrt{6} \end{bmatrix} \begin{bmatrix} 1 & 1 \\ 1 & -1 \\ 1 & 1 \end{bmatrix} = \begin{bmatrix} \sqrt{3} & 1/\sqrt{3} \\ 0 & 4/\sqrt{6} \end{bmatrix}$.

Comment. The entries of R have actually all been computed during Gram–Schmidt, so that, if we pay attention, we could immediately write down R (no extra work required). Looking back at Example 70, can you see this?

Check. Indeed, $QR = \begin{bmatrix} 1/\sqrt{3} & 1/\sqrt{6} \\ 1/\sqrt{3} & -2/\sqrt{6} \\ 1/\sqrt{3} & 1/\sqrt{6} \end{bmatrix} \begin{bmatrix} \sqrt{3} & 1/\sqrt{3} \\ 0 & 4/\sqrt{6} \end{bmatrix} = \begin{bmatrix} 1 & 1 \\ 1 & -1 \\ 1 & 1 \end{bmatrix}$ equals A .

Example 73. Using Gram–Schmidt, find an orthogonal basis for $W = \text{span} \left\{ \begin{bmatrix} 0 \\ 3 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 2 \\ 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 1 \\ 1 \\ 1 \end{bmatrix} \right\}$.

Solution. We begin with the (not orthogonal) basis $w_1 = \begin{bmatrix} 0 \\ 3 \\ 0 \\ 0 \end{bmatrix}$, $w_2 = \begin{bmatrix} 2 \\ 1 \\ 0 \\ 0 \end{bmatrix}$, $w_3 = \begin{bmatrix} 1 \\ 1 \\ 1 \\ 1 \end{bmatrix}$.

We then construct an orthogonal basis q_1, q_2, q_3 :

- $q_1 = w_1 = \begin{bmatrix} 0 \\ 3 \\ 0 \\ 0 \end{bmatrix}$
- $q_2 = w_2 - \left(\begin{array}{l} \text{projection of} \\ w_2 \text{ onto } q_1 \end{array} \right) = \begin{bmatrix} 2 \\ 1 \\ 0 \\ 0 \end{bmatrix} - \frac{3}{9} \begin{bmatrix} 0 \\ 3 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} 2 \\ 0 \\ 0 \\ 0 \end{bmatrix}$
- $q_3 = w_3 - \left(\begin{array}{l} \text{projection of } w_3 \\ \text{onto } \text{span}\{q_1, q_2\} \end{array} \right) = w_3 - \left(\begin{array}{l} \text{projection of} \\ w_3 \text{ onto } q_1 \end{array} \right) - \left(\begin{array}{l} \text{projection of} \\ w_3 \text{ onto } q_2 \end{array} \right)$
 $= \begin{bmatrix} 1 \\ 1 \\ 1 \\ 1 \end{bmatrix} - \frac{3}{9} \begin{bmatrix} 0 \\ 3 \\ 0 \\ 0 \end{bmatrix} - \frac{2}{4} \begin{bmatrix} 2 \\ 0 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 1 \\ 1 \end{bmatrix}$

Make sure you understand how q_3 was designed to be orthogonal to both q_1 and q_2 !

Also note that breaking up the projection onto $\text{span}\{q_1, q_2\}$ into the projections onto q_1 and q_2 is only possible because q_1 and q_2 are orthogonal.

Hence, $\begin{bmatrix} 0 \\ 3 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 2 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 1 \\ 1 \end{bmatrix}$ is an orthogonal basis of W .

Important. Normalizing, we obtain an orthonormal basis: $\begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \frac{1}{\sqrt{2}} \begin{bmatrix} 0 \\ 0 \\ 1 \\ 1 \end{bmatrix}$.

Example 74. Determine the QR decomposition of $A = \begin{bmatrix} 0 & 2 & 1 \\ 3 & 1 & 1 \\ 0 & 0 & 1 \\ 0 & 0 & 1 \end{bmatrix}$.

Solution. The first step is Gram–Schmidt orthonormalization on the columns of A . We then use the resulting orthonormal vectors as the columns of Q .

We already did Gram–Schmidt in Example 73: from that work, we have $Q = \begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1/\sqrt{2} \\ 0 & 0 & 1/\sqrt{2} \end{bmatrix}$.

Hence, $R = Q^T A = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 1/\sqrt{2} & 1/\sqrt{2} \end{bmatrix} \begin{bmatrix} 0 & 2 & 1 \\ 3 & 1 & 1 \\ 0 & 0 & 1 \\ 0 & 0 & 1 \end{bmatrix} = \begin{bmatrix} 3 & 1 & 1 \\ 0 & 2 & 1 \\ 0 & 0 & \sqrt{2} \end{bmatrix}$.

Comment. As commented earlier, the entries of R have actually all been computed during Gram–Schmidt, so that, if we pay attention, we could immediately write down R (no extra work required). Looking back at Example 73, can you see this?

Letting Sage do the work for us.

```
Sage] A = matrix(QQbar, [[0,2,1],[3,1,1],[0,0,1],[0,0,1]])
```

```
Sage] A.QR(full=false)
```

$$\left(\begin{bmatrix} 0 & 1 & & 0 \\ 1 & 0 & & 0 \\ 0 & 0 & 0.7071067811865475? & \\ 0 & 0 & 0.7071067811865475? & \end{bmatrix}, \begin{bmatrix} 3 & 1 & & 1 \\ 0 & 2 & & 1 \\ 0 & 0 & 1.414213562373095? & \end{bmatrix} \right)$$

Comment. Can you figure out what happens if you omit the `full=false`? Check out the comment under **Variations** in the statement of the QR decomposition. On the other hand, the `QQbar` is telling Sage to compute with algebraic numbers (instead of just rational numbers); if omitted, it would complain that square roots are not available

Example 75. (extra) Determine the QR decomposition of $A = \begin{bmatrix} 1 & 2 & 4 \\ 0 & 0 & -5 \\ 0 & 3 & 6 \end{bmatrix}$.

Solution. We first apply Gram–Schmidt orthonormalization to the columns of A . For a variation, like a computer, we normalize after each step (rather than normalize at the end):

- $\mathbf{b}_1 = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}$, so that $\mathbf{q}_1 = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}$.
- $\mathbf{b}_2 = \begin{bmatrix} 2 \\ 0 \\ 3 \end{bmatrix} - \left(\begin{bmatrix} 2 \\ 0 \\ 3 \end{bmatrix} \cdot \mathbf{q}_1 \right) \mathbf{q}_1 = \begin{bmatrix} 0 \\ 0 \\ 3 \end{bmatrix}$, so that $\mathbf{q}_2 = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}$.
- $\mathbf{b}_3 = \begin{bmatrix} 4 \\ -5 \\ 6 \end{bmatrix} - \left(\begin{bmatrix} 4 \\ -5 \\ 6 \end{bmatrix} \cdot \mathbf{q}_1 \right) \mathbf{q}_1 - \left(\begin{bmatrix} 4 \\ -5 \\ 6 \end{bmatrix} \cdot \mathbf{q}_2 \right) \mathbf{q}_2 = \begin{bmatrix} 0 \\ -5 \\ 0 \end{bmatrix}$, so that $\mathbf{q}_3 = \begin{bmatrix} 0 \\ -1 \\ 0 \end{bmatrix}$.

Therefore, $Q = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & -1 \\ 0 & 1 & 0 \end{bmatrix}$. Finally, $R = Q^T A = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & -1 & 0 \end{bmatrix} \begin{bmatrix} 1 & 2 & 4 \\ 0 & 0 & -5 \\ 0 & 3 & 6 \end{bmatrix} = \begin{bmatrix} 1 & 2 & 4 \\ 0 & 3 & 6 \\ 0 & 0 & 5 \end{bmatrix}$.

In conclusion, we have found the QR decomposition: $\underbrace{\begin{bmatrix} 1 & 2 & 4 \\ 0 & 0 & -5 \\ 0 & 3 & 6 \end{bmatrix}}_A = \underbrace{\begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & -1 & 0 \end{bmatrix}}_Q \underbrace{\begin{bmatrix} 1 & 2 & 4 \\ 0 & 3 & 6 \\ 0 & 0 & 5 \end{bmatrix}}_R$

Comment. As noted before, we actually could write down R without any additional computation. Indeed, realize that the second column of R , that is $[2, 3, 0]^T$ means that

$$\text{2nd col of } A = 2\mathbf{q}_1 + 3\mathbf{q}_2.$$

Which we already knew from our computation of \mathbf{q}_2 ! Also, by construction, we know that the second column of A is a linear combination of \mathbf{q}_1 and \mathbf{q}_2 only, and that \mathbf{q}_3 enters the story later on. This corresponds to the fact that R is always upper triangular.

Letting Sage do the work for us.

```
Sage] A = matrix(QQbar, [[1,2,4], [0,0,-5], [0,3,6]])
```

```
Sage] A.QR()
```

$$\left(\begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & -1 \\ 0 & 1 & 0 \end{bmatrix}, \begin{bmatrix} 1 & 2 & 4 \\ 0 & 3 & 6 \\ 0 & 0 & 5 \end{bmatrix} \right)$$

Comment. The `QQbar` is telling Sage to compute with algebraic numbers (instead of just rational numbers); in general, if omitted, it would complain that square roots are not available (because the matrices Q and R typically involve square roots). Here, we are lucky that square roots didn't creep in.

Example 76. (extra) Find the QR decomposition of $A = \begin{bmatrix} 1 & 1 & 2 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{bmatrix}$.

Solution. (final answer only) $A = QR$ with $Q = \begin{bmatrix} \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} & 0 \\ 0 & 0 & 1 \\ \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} & 0 \end{bmatrix}$ and $R = \begin{bmatrix} \sqrt{2} & \frac{1}{\sqrt{2}} & \sqrt{2} \\ 0 & \frac{1}{\sqrt{2}} & \sqrt{2} \\ 0 & 0 & 1 \end{bmatrix}$.

Example 77. One practical application of the QR decomposition is solving systems of linear equations.

$$\begin{aligned} Ax = b &\iff QRx = b && \text{(now, multiply with } Q^T \text{ from the left)} \\ &\implies Rx = Q^T b \end{aligned}$$

The last system is triangular and can be solved by back-substitution.

A couple of comments are in order:

- If A is $n \times n$ and invertible, then the “ \implies ” is actually a “ \iff ”.
- The equation $Rx = Q^T b$ is always consistent! (Recall that R is invertible.)
Indeed, if A is not $n \times n$ or not invertible, then $Rx = Q^T b$ gives the least squares solutions!
Why? $A^T A \hat{x} = A^T b \iff \underbrace{(QR)^T Q R \hat{x}}_{=R^T Q^T Q R} = (QR)^T b \iff R^T R \hat{x} = R^T Q^T b \iff R \hat{x} = Q^T b$
[For the last step we need that R is invertible, which is always the case when A is $m \times n$ of rank n .]

- So, how does the QR way of solving linear systems compare to our beloved Gaussian elimination (LU)?
It turns out that QR is a little slower than LU but makes up for it in “numerical stability”.

What does that mean? When computing numerically, we use floating point arithmetic and approximate each number by an expression of the form $0.1234 \cdot 10^{-16}$. A certain (fixed) number of bits is used to store the part 0.1234 (here, 4 decimal places of accuracy) as well as the exponent -16 .

Now, here is something terrible that can happen in numerical computations: mathematically, the quantities x and $(x+1) - 1$ are exactly the same. However, numerically, they might not. Take, for instance, $x = 0.1234 \cdot 10^{-6}$. Then, to an accuracy of 4 decimal places, $x+1 = 0.1000 \cdot 10^1$, so that $(x+1) - 1 = 0.0000$. But $x \neq 0$. We completely lost all the information about x .

To be numerically stable, an algorithm must avoid issues like that.

\hat{x} is a least squares solution of $Ax = b$
 $\iff R\hat{x} = Q^T b$ (where $A = QR$)

Review. A matrix A has orthonormal columns $\iff A^T A = I$.

Example 78. Suppose Q has orthonormal columns. What is the projection matrix P for orthogonally projecting onto $\text{col}(Q)$?

Solution. Recall that, to project onto $\text{col}(A)$, the projection matrix is $P = A(A^T A)^{-1} A^T$.

Since $Q^T Q = I$, to project onto $\text{col}(Q)$, the projection matrix is $P = Q Q^T$.

Comment. A familiar special case is when we project onto a unit vector q : in that case, the projection of b onto q is $(q \cdot b)q = q(q^T b) = (qq^T)b$, so the projection matrix here is qq^T .

Comment. In particular, if Q is not square, then $Q^T Q = I$ but $Q Q^T \neq I$. In some sense, $Q Q^T$ still “tries” to be as close to the identity as possible: since it is the matrix projecting onto $\text{col}(Q)$ it does act like the identity for vectors in $\text{col}(Q)$. (Vectors not in $\text{col}(Q)$ are sent to their projection, that is, the closest to themselves while restricted to $\text{col}(Q)$.)

Example 79. Suppose A is invertible. What is the projection matrix P for orthogonally projecting onto $\text{col}(A)$?

Solution. If A is an invertible $n \times n$ matrix, then $\text{col}(A) = \mathbb{R}^n$ (because the n columns of A are linearly independent and hence form a basis for \mathbb{R}^n).

Since $\text{col}(A)$ is the entire space we are not really projecting at all: every vector is sent to itself.

In particular, the projection matrix is $P = I$.

Definition 80. An **orthogonal matrix** is a square matrix with orthonormal columns.

[This is not a typo (but a confusing convention): the columns need to be orthonormal, not just orthogonal.]

An $n \times n$ matrix Q is orthogonal $\iff Q^T Q = I$

In other words, $Q^{-1} = Q^T$.

Review. Recall the following properties of determinants:

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

Comment. In fancy language, this means that the determinant is a group homomorphism between the group of (invertible) $n \times n$ matrices and (nonzero) complex numbers. Note that, on the left hand, we are multiplying matrices while, on the right hand, we are multiplying numbers. The key point is that it doesn't matter which multiplication we do: the two multiplications are compatible.

- $\det(A^{-1}) = \frac{1}{\det(A)}$

Comment. Can you derive this from the previous property?

- $\det(A^T) = \det(A)$

Comment. We are familiar with this in the context of cofactor expansion: it doesn't matter whether we expand by a column or by a row.

Example 81. What can we say about $\det(Q)$ if Q is orthogonal?

Solution. Write $d = \det(Q)$. Since $Q^{-1} = Q^T$, we have $\frac{1}{d} = d$ (recall that $\det(Q^{-1}) = 1 / \det(Q)$ and $\det(Q^T) = \det(Q)$) or, equivalently, $d^2 = 1$. Hence, $d = \pm 1$.

Both of these are possible as the examples $Q = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$ and $Q = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}$ illustrate.

Preview: The spectral theorem

Example 82. (review) In Example 17, we diagonalized $A = \begin{bmatrix} 4 & 0 & 2 \\ 2 & 2 & 2 \\ 1 & 0 & 3 \end{bmatrix}$ as $A = PDP^{-1}$.

We found that one choice for P and D is $P = \begin{bmatrix} 2 & 0 & -1 \\ 2 & 1 & 0 \\ 1 & 0 & 1 \end{bmatrix}$, $D = \begin{bmatrix} 5 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 2 \end{bmatrix}$.

Spell out what that tells us about A !

Solution. The diagonal entries 5, 2, 2 of D are the eigenvalues of A .

The columns of P are corresponding eigenvectors of A .

- $\begin{bmatrix} 2 \\ 2 \\ 1 \end{bmatrix}$ is a 5-eigenvector of A (that is, $A \begin{bmatrix} 2 \\ 2 \\ 1 \end{bmatrix} = 5 \begin{bmatrix} 2 \\ 2 \\ 1 \end{bmatrix}$).
- The 2-eigenspace of A is 2-dimensional. A basis is $\begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} -1 \\ 0 \\ 1 \end{bmatrix}$.

Example 83. Diagonalize the symmetric matrix $A = \begin{bmatrix} 8 & -6 & 2 \\ -6 & 7 & -4 \\ 2 & -4 & 3 \end{bmatrix}$ as $A = PDP^{-1}$.

Review. Recall that a matrix A is symmetric if $A^T = A$.

Solution. We let Sage do the work for us:

Sage] `A = matrix([[8,-6,2],[-6,7,-4],[2,-4,3]])`

Sage] `A.eigenmatrix_right()`

$$\left(\begin{bmatrix} 15 & 0 & 0 \\ 0 & 3 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \begin{bmatrix} 1 & 1 & 1 \\ -1 & \frac{1}{2} & 2 \\ \frac{1}{2} & -1 & 2 \end{bmatrix} \right)$$

This output shows that A is diagonalizable as $A = PDP^{-1}$ with $D = \begin{bmatrix} 15 & 0 & 0 \\ 0 & 3 & 0 \\ 0 & 0 & 0 \end{bmatrix}$ and $P = \begin{bmatrix} 1 & 1 & 1 \\ -1 & \frac{1}{2} & 2 \\ \frac{1}{2} & -1 & 2 \end{bmatrix}$.

Just to make sure. This means that the eigenvalues of A are 15, 3, 0 (the diagonal entries of D).

Moreover, we have that $\begin{bmatrix} 1 \\ -1 \\ \frac{1}{2} \end{bmatrix}$ is a 15-eigenvector, $\begin{bmatrix} 1 \\ \frac{1}{2} \\ -1 \end{bmatrix}$ is a 3-eigenvector, and $\begin{bmatrix} 1 \\ 2 \\ 2 \end{bmatrix}$ is a 0-eigenvector.

Important observation. Note that the eigenspaces of A are orthogonal to each other here.

The spectral theorem says that this is true for all symmetric matrices A .

Example 84. Diagonalize the symmetric matrix $A = \begin{bmatrix} 8 & -6 & 2 \\ -6 & 7 & -4 \\ 2 & -4 & 3 \end{bmatrix}$ as $A = PDP^T$.

Solution. By the previous example, we can diagonalize A as $\tilde{P}D\tilde{P}^{-1}$ with $\tilde{P} = \begin{bmatrix} 2 & 2 & 1 \\ -2 & 1 & 2 \\ 1 & -2 & 2 \end{bmatrix}$ and $D = \begin{bmatrix} 15 & 0 & 0 \\ 0 & 3 & 0 \\ 0 & 0 & 0 \end{bmatrix}$.

(To avoid fractions, we just scaled the first two columns of \tilde{P} , which are eigenvectors.)

Note that the columns of \tilde{P} are orthogonal (this is due the spectral theorem). If we normalize them (they all have length $\sqrt{2^2 + 2^2 + 1} = 3$), then we obtain the orthogonal matrix $P = \frac{1}{3} \begin{bmatrix} 2 & 2 & 1 \\ -2 & 1 & 2 \\ 1 & -2 & 2 \end{bmatrix}$.

Since $P^{-1} = P^T$, we now have $A = PDP^T$.

Example 85.

- (a) Determine the eigenspaces of the symmetric matrix $A = \begin{bmatrix} 1 & 2 \\ 2 & 4 \end{bmatrix}$.
- (b) Diagonalize A as $A = PDP^T$.

Solution.

- (a) The characteristic polynomial is $\begin{vmatrix} 1-\lambda & 2 \\ 2 & 4-\lambda \end{vmatrix} = \lambda^2 - 5\lambda = \lambda(\lambda - 5)$, and so A has eigenvalues 5, 0.

The 5-eigenspace is $\text{null}\left(\begin{bmatrix} -4 & 2 \\ 2 & -1 \end{bmatrix}\right)$ has basis $\begin{bmatrix} 1 \\ 2 \end{bmatrix}$.

The 0-eigenspace is $\text{null}\left(\begin{bmatrix} 1 & 2 \\ 2 & 4 \end{bmatrix}\right)$ has basis $\begin{bmatrix} -2 \\ 1 \end{bmatrix}$.

Important observation. The 5-eigenvector $\begin{bmatrix} 1 \\ 2 \end{bmatrix}$ and the 0-eigenvector $\begin{bmatrix} -2 \\ 1 \end{bmatrix}$ are orthogonal!

- (b) Note that a usual diagonalization is of the form $A = PDP^{-1}$.

We need to choose P so that $P^{-1} = P^T$, which means that P must be **orthogonal** (meaning orthonormal columns). [Choosing such a P is only possible if the eigenspaces of A are orthogonal.]

Hence, we normalize the two eigenvectors to $\frac{1}{\sqrt{5}}\begin{bmatrix} 1 \\ 2 \end{bmatrix}$ and $\frac{1}{\sqrt{5}}\begin{bmatrix} -2 \\ 1 \end{bmatrix}$.

With $P = \frac{1}{\sqrt{5}}\begin{bmatrix} 1 & -2 \\ 2 & 1 \end{bmatrix}$ and $D = \begin{bmatrix} 5 & 0 \\ 0 & 0 \end{bmatrix}$, we then have $A = PDP^T$.

Review: Diagonalizability

Example 86. (review) If A is a 2×2 matrix with $\det(A) = -8$ and eigenvalue 4. What is the second eigenvalue?

Solution. Recall that $\det(A)$ is the product of the eigenvalues (see below). Hence, the second eigenvalue is -2 .

$\det(A)$ is the product of the eigenvalues of A .

Why? Recall how we determine the eigenvalues $\lambda_1, \lambda_2, \dots, \lambda_n$ of an $n \times n$ matrix A . We compute the characteristic polynomial $\det(A - \lambda I)$ and determine the λ_i as the roots of that polynomial.

That means that we have the factorization $\det(A - \lambda I) = (\lambda_1 - \lambda)(\lambda_2 - \lambda) \cdots (\lambda_n - \lambda)$. Now, set $\lambda = 0$ to conclude that $\det(A) = \lambda_1 \lambda_2 \cdots \lambda_n$.

Lemma 87. A matrix A is diagonalizable if and only if, for every eigenvalue λ that is k times repeated, the λ -eigenspace of A has dimension k .

In short, an $n \times n$ matrix A is diagonalizable if and only if there exists a basis of \mathbb{R}^n consisting of eigenvectors of A (i.e. "there are enough eigenvectors").

The next two examples illustrate that not all matrices are diagonalizable and that, even if a real matrix is diagonalizable, the eigenvalues and eigenvectors might be complex.

Example 88. What are the eigenvalues and eigenvectors of $A = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}$? Is A diagonalizable?

Solution. The characteristic polynomial is $\det\left(\begin{bmatrix} -\lambda & 1 \\ 0 & -\lambda \end{bmatrix}\right) = \lambda^2$, which has $\lambda = 0$ as a double root.

However, the 0-eigenspace $\text{null}(A) = \text{span}\left\{\begin{bmatrix} 1 \\ 0 \end{bmatrix}\right\}$ is only 1-dimensional.

As a consequence, A is not diagonalizable.

Example 89. What are the eigenvalues and eigenvectors of $A = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}$? Is A diagonalizable?

Solution. The characteristic polynomial is $\det\left(\begin{bmatrix} -\lambda & -1 \\ 1 & -\lambda \end{bmatrix}\right) = \lambda^2 + 1 = (\lambda - i)(\lambda + i)$.

Hence, the eigenvalues are $\pm i$.

The i -eigenspace $\text{null}\left(\begin{bmatrix} -i & -1 \\ 1 & -i \end{bmatrix}\right)$ has basis $\begin{bmatrix} i \\ 1 \end{bmatrix}$.

The $-i$ -eigenspace $\text{null}\left(\begin{bmatrix} i & -1 \\ 1 & i \end{bmatrix}\right)$ has basis $\begin{bmatrix} -i \\ 1 \end{bmatrix}$.

Thus, A has the diagonalization $A = PDP^{-1}$ with $D = \begin{bmatrix} i & \\ & -i \end{bmatrix}$ and $P = \begin{bmatrix} i & -i \\ 1 & 1 \end{bmatrix}$.

The spectral theorem

Recall that a matrix A is symmetric if and only if $A^T = A$.

Theorem 90. (spectral theorem, long version) Suppose A is a symmetric matrix.

- A is always diagonalizable.
- All eigenvalues of A are real.
- The eigenspaces of A are orthogonal.

Proof. We will prove (parts of) the spectral theorem later on. For now, we just appreciate that the spectral theorem guarantees all these nice things to happen for symmetric matrices (for any specific A we know how to determine whether A is diagonalizable and what its eigenspaces are).

Comment. The eigenspaces of A being orthogonal means that eigenvectors for different eigenvalues are always orthogonal.

Important consequence. In the diagonalization $A = PDP^{-1}$, we can choose P to be orthogonal (in which case $P^{-1} = P^T$). In that case, the diagonalization takes the special form $A = PDP^T$, where P is orthogonal and D is diagonal.

(spectral theorem, compact version) A symmetric matrix A can always be diagonalized as $A = PDP^T$, where P is orthogonal and D is diagonal (and both are real).

How? We proceed as in the diagonalization $A = PDP^{-1}$. For a symmetric matrix A , we can arrange P to be orthogonal, by normalizing its columns. If there is a repeated eigenvalue, then we also need to make sure to pick an orthonormal basis for the corresponding eigenspace (for instance, using Gram–Schmidt).

Advanced comment. A matrix such that $A^T A = A A^T$ is called **normal**. For normal matrices, the (complex!) eigenspaces are again orthogonal to each other. However, normal matrices which are not symmetric will always have complex eigenvalues. (In that case, the orthogonal matrix P gets replaced with a unitary matrix, the complex version of orthogonal matrices, and the P^T becomes the conjugate transpose $P^* = \bar{P}^T$.)

Example 91.

- (a) Determine the eigenspaces of the symmetric matrix $A = \begin{bmatrix} 1 & 3 \\ 3 & 1 \end{bmatrix}$.
- (b) Diagonalize A as $A = PDP^T$.

Solution.

- (a) The characteristic polynomial is $\begin{vmatrix} 1-\lambda & 3 \\ 3 & 1-\lambda \end{vmatrix} = (\lambda-4)(\lambda+2)$, and so A has eigenvalues $4, -2$.

The 4 -eigenspace is $\text{null}\left(\begin{bmatrix} -3 & 3 \\ 3 & -3 \end{bmatrix}\right)$ has basis $\begin{bmatrix} 1 \\ 1 \end{bmatrix}$.

The -2 -eigenspace is $\text{null}\left(\begin{bmatrix} 3 & 3 \\ 3 & 3 \end{bmatrix}\right)$ has basis $\begin{bmatrix} -1 \\ 1 \end{bmatrix}$.

Important observation. The 4 -eigenvector $\begin{bmatrix} 1 \\ 1 \end{bmatrix}$ and the -2 -eigenvector $\begin{bmatrix} -1 \\ 1 \end{bmatrix}$ are orthogonal!

Review. The product of all eigenvalues $-2 \cdot 4 = -8$ equals the determinant $\det(A) = 1 - 9 = -8$.

- (b) Note that a usual diagonalization is of the form $A = PDP^{-1}$.

We need to choose P so that $P^{-1} = P^T$, which means that P must be **orthogonal** (meaning orthonormal columns). [Choosing such a P is only possible if the eigenspaces of A are orthogonal.]

Hence, we normalize the two eigenvectors to $\frac{1}{\sqrt{2}}\begin{bmatrix} 1 \\ 1 \end{bmatrix}$ and $\frac{1}{\sqrt{2}}\begin{bmatrix} -1 \\ 1 \end{bmatrix}$.

With $P = \frac{1}{\sqrt{2}}\begin{bmatrix} 1 & -1 \\ 1 & 1 \end{bmatrix}$ and $D = \begin{bmatrix} 4 & 0 \\ 0 & -2 \end{bmatrix}$, we then have $A = PDP^T$.

Example 92. (again, simplified) Diagonalize the symmetric matrix $A = \begin{bmatrix} 1 & 3 \\ 3 & 1 \end{bmatrix}$ as $A = PDP^T$.

Solution. See Example 91 for a solution that illustrates how to diagonalize any symmetric matrix. For a simplified solution, note that we can see right away that $\begin{bmatrix} 1 \\ 1 \end{bmatrix}$ is a 4-eigenvector (since the row sums are equal!).

Because the eigenspaces are orthogonal (since A is symmetric!), $\begin{bmatrix} -1 \\ 1 \end{bmatrix}$ must also be an eigenvector.

Indeed, $\begin{bmatrix} 1 & 3 \\ 3 & 1 \end{bmatrix} \begin{bmatrix} -1 \\ 1 \end{bmatrix} = \begin{bmatrix} 2 \\ -2 \end{bmatrix}$ shows that the corresponding eigenvalue is -2 .

We normalize the two eigenvectors and use them as the columns of P , so that $P = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & -1 \\ 1 & 1 \end{bmatrix}$ is an orthogonal matrix ($P^{-1} = P^T$). With $D = \begin{bmatrix} 4 & 0 \\ 0 & -2 \end{bmatrix}$ we then have $A = PDP^T$.

Example 93. Let A be a symmetric 2×2 matrix with 7-eigenvector $\begin{bmatrix} 2 \\ 5 \end{bmatrix}$ and $\det(A) = -21$. Determine the second eigenvalue and a corresponding eigenvector.

Further, diagonalize A as $A = PDP^T$.

Solution. A has $-\frac{21}{7} = -3$ -eigenvector $\begin{bmatrix} -5 \\ 2 \end{bmatrix}$.

Hence, $A = PDP^T$ with $D = \begin{bmatrix} 7 & \\ & -3 \end{bmatrix}$ and $P = \frac{1}{\sqrt{29}} \begin{bmatrix} 2 & -5 \\ 5 & 2 \end{bmatrix}$.

Comment. Recall that, because A is symmetric, the eigenvector must be orthogonal to $\begin{bmatrix} 2 \\ 5 \end{bmatrix}$.

[In general, $\begin{bmatrix} a \\ b \end{bmatrix}$ and $\begin{bmatrix} -b \\ a \end{bmatrix}$ are orthogonal.]

Let us prove the following important part of the spectral theorem.

Theorem 94. If A is symmetric, then the eigenspaces of A are orthogonal.

Proof. To prove the claim we need to show that, if \mathbf{v} and \mathbf{w} are eigenvectors of A with different eigenvalues (say λ and μ), then $\mathbf{v} \cdot \mathbf{w} = 0$. Suppose therefore that $A\mathbf{v} = \lambda\mathbf{v}$ and $A\mathbf{w} = \mu\mathbf{w}$ with $\lambda \neq \mu$.

First, we observe that, for any matrix A and vectors \mathbf{v} , \mathbf{w} , we have the following:

$$(A\mathbf{v}) \cdot \mathbf{w} = (A\mathbf{v})^T \mathbf{w} = (\mathbf{v}^T A^T) \mathbf{w} = \mathbf{v}^T (A^T \mathbf{w}) = \mathbf{v} \cdot (A^T \mathbf{w})$$

If A is symmetric, we therefore have

$$(A\mathbf{v}) \cdot \mathbf{w} = \mathbf{v} \cdot (A\mathbf{w}).$$

We now use that $A\mathbf{v} = \lambda\mathbf{v}$ and $A\mathbf{w} = \mu\mathbf{w}$ to conclude from the above that $\lambda\mathbf{v} \cdot \mathbf{w} = \mu\mathbf{v} \cdot \mathbf{w}$.

However, since $\lambda \neq \mu$, this is only possible if $\mathbf{v} \cdot \mathbf{w} = 0$. □

Powers of matrices

Example 95. (warmup) Consider $A = \begin{bmatrix} -2 & 0 \\ 0 & 3 \end{bmatrix}$.

- What are the eigenspaces?
- What are A^{-1} and A^{100} ? What is A^n ?

Solution.

- $\begin{bmatrix} 1 \\ 0 \end{bmatrix}$ is a -2 -eigenvector, and $\begin{bmatrix} 0 \\ 1 \end{bmatrix}$ is a 3 -eigenvector. In other words, the -2 -eigenspace is $\text{span}\left\{\begin{bmatrix} 1 \\ 0 \end{bmatrix}\right\}$ and the 3 -eigenspace is $\text{span}\left\{\begin{bmatrix} 0 \\ 1 \end{bmatrix}\right\}$.
- $A^{-1} = \begin{bmatrix} -1/2 & 0 \\ 0 & 1/3 \end{bmatrix}$ and $A^{100} = \begin{bmatrix} (-2)^{100} & 0 \\ 0 & 3^{100} \end{bmatrix} = \begin{bmatrix} 2^{100} & 0 \\ 0 & 3^{100} \end{bmatrix}$. In general, $A^n = \begin{bmatrix} (-2)^n & 0 \\ 0 & 3^n \end{bmatrix}$.

Comment. Algebraically, the map $v \mapsto Av$ looks very simple. However, notice that it is not so easy to say what happens to, say, $v = \begin{bmatrix} 3 \\ 4 \end{bmatrix}$ geometrically. That is because two things are happening: part of the vector v is scaled by -2 , the other part is scaled by 3 .

Example 96. If A has λ -eigenvector v , then what can we say about A^2 ?

Solution. A^2 has λ^2 -eigenvector v .

[Indeed, $A^2v = A(Av) = A(\lambda v) = \lambda Av = \lambda^2v$. This is even easier in words: multiplying v with A has the effect of scaling it by λ ; hence, multiplying it with A^2 scales it by λ^2 .]

Important comment. Similarly, A^{100} has λ^{100} -eigenvector v .

Example 97. If a matrix A can be diagonalized as $A = PDP^{-1}$, what can we say about A^n ?

Solution. First, note that $A^2 = (PDP^{-1})(PDP^{-1}) = PD^2P^{-1}$. Likewise, $A^n = PD^nP^{-1}$.

[The point being that D^n is trivial to compute because D is diagonal.]

In particular. $A^{-1} = PD^{-1}P^{-1}$

Important comment. In the previous example, we observed that, if A has λ -eigenvector v , then A^n has λ^n -eigenvector v . Note that this is also expressed in $A^n = PD^nP^{-1}$, because the latter is a diagonalization of A^n . The diagonalization shows that A^n and A have the same eigenvectors (since we can use the same matrix P) and that the eigenvalues of A^n are the n -th powers of the eigenvalues of A (which are the entries of the diagonal matrix D).

(computing matrix powers) If A is a square matrix with diagonalization $A = PDP^{-1}$, then

$$A^n = PD^nP^{-1}.$$

Example 98. Let $A = \begin{bmatrix} 6 & 1 \\ 4 & 9 \end{bmatrix}$. Compute A^n .

Solution. First, we diagonalize: $A = PDP^{-1}$ with $P = \begin{bmatrix} 1 & -1 \\ 4 & 1 \end{bmatrix}$ and $D = \begin{bmatrix} 10 & \\ & 5 \end{bmatrix}$. (Fill in the details!)

$$A^n = PD^nP^{-1} = \begin{bmatrix} 1 & -1 \\ 4 & 1 \end{bmatrix} \begin{bmatrix} 10^n & \\ & 5^n \end{bmatrix} \frac{1}{5} \begin{bmatrix} 1 & 1 \\ -4 & 1 \end{bmatrix} = \frac{1}{5} \begin{bmatrix} 1 & -1 \\ 4 & 1 \end{bmatrix} \begin{bmatrix} 10^n & 10^n \\ -4 \cdot 5^n & 1 \cdot 5^n \end{bmatrix} = \frac{1}{5} \begin{bmatrix} 10^n + 4 \cdot 5^n & 10^n - 5^n \\ 4 \cdot 10^n - 4 \cdot 5^n & 4 \cdot 10^n + 5^n \end{bmatrix}$$

Check. Verify the cases $n = 0$ ($A^0 = I$) and $n = 1$.

Example 99. (extra) Let $A = \begin{bmatrix} 4 & 0 & 2 \\ 2 & 2 & 2 \\ 1 & 0 & 3 \end{bmatrix}$. Determine A^n .

Solution. We first repeat our work from Example 17 to find a diagonalization of A :

By expanding by the second column, we find that the characteristic polynomial $\det(A - \lambda I)$ is

$$\begin{vmatrix} 4-\lambda & 0 & 2 \\ 2 & 2-\lambda & 2 \\ 1 & 0 & 3-\lambda \end{vmatrix} = (2-\lambda) \begin{vmatrix} 4-\lambda & 2 \\ 1 & 3-\lambda \end{vmatrix} = (2-\lambda)[(4-\lambda)(3-\lambda) - 2] = (2-\lambda)^2(5-\lambda).$$

Hence, the eigenvalues are $\lambda = 2$ (with multiplicity 2) and $\lambda = 5$.

- $\lambda = 5$: $\text{null}\left(\begin{bmatrix} -1 & 0 & 2 \\ 2 & -3 & 2 \\ 1 & 0 & -2 \end{bmatrix}\right) \stackrel{\text{RREF}}{=} \text{null}\left(\begin{bmatrix} 1 & 0 & -2 \\ 0 & 1 & -2 \\ 0 & 0 & 0 \end{bmatrix}\right) = \text{span}\left\{\begin{bmatrix} 2 \\ 2 \\ 1 \end{bmatrix}\right\}$

- $\lambda = 2$: $\text{null}\left(\begin{bmatrix} 2 & 0 & 2 \\ 2 & 0 & 2 \\ 1 & 0 & 1 \end{bmatrix}\right) \stackrel{\text{RREF}}{=} \text{null}\left(\begin{bmatrix} 1 & 0 & 1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}\right) = \text{span}\left\{\begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} -1 \\ 0 \\ 1 \end{bmatrix}\right\}$

We therefore have the diagonalization $A = PDP^{-1}$ with $P = \begin{bmatrix} 2 & 0 & -1 \\ 2 & 1 & 0 \\ 1 & 0 & 1 \end{bmatrix}$, $D = \begin{bmatrix} 5 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 2 \end{bmatrix}$.

[Keep in mind that other choices for P and D exist.]

With some labor (do it!), we find $P^{-1} = \frac{1}{3} \begin{bmatrix} 1 & 0 & 1 \\ -2 & 3 & -2 \\ -1 & 0 & 2 \end{bmatrix}$.

It follows that

$$\begin{aligned} A^n &= PD^nP^{-1} \\ &= \begin{bmatrix} 2 & 0 & -1 \\ 2 & 1 & 0 \\ 1 & 0 & 1 \end{bmatrix} \begin{bmatrix} 5^n & 0 & 0 \\ 0 & 2^n & 0 \\ 0 & 0 & 2^n \end{bmatrix} \frac{1}{3} \begin{bmatrix} 1 & 0 & 1 \\ -2 & 3 & -2 \\ -1 & 0 & 2 \end{bmatrix} \\ &= \frac{1}{3} \begin{bmatrix} 2 \cdot 5^n & 0 & -2^n \\ 2 \cdot 5^n & 2^n & 0 \\ 5^n & 0 & 2^n \end{bmatrix} \begin{bmatrix} 1 & 0 & 1 \\ -2 & 3 & -2 \\ -1 & 0 & 2 \end{bmatrix} \\ &= \frac{1}{3} \begin{bmatrix} 2 \cdot 5^n + 2^n & 0 & 2 \cdot 5^n - 2 \cdot 2^n \\ 2 \cdot 5^n - 2 \cdot 2^n & 3 \cdot 2^n & 2 \cdot 5^n - 2 \cdot 2^n \\ 5^n - 2^n & 0 & 5^n + 2 \cdot 2^n \end{bmatrix}. \end{aligned}$$


Check. Notice that it is particularly easy to verify the cases $n = 0$ ($A^0 = I$) and $n = 1$.

Application: Markov chains

Example 100. Consider a fixed population of people with or without active immunization against some disease (like tetanus). Suppose that, each year, 40% of those unprotected get vaccinated while 10% of those with immunization lose their protection.

What is the immunization rate in the long run? (The long term equilibrium.)

Solution.



x_t : proportion of population unprotected at time t (in years)

y_t : proportion of population immune at time t

[Note that $x_t + y_t = 1$.]

$$\begin{bmatrix} x_{t+1} \\ y_{t+1} \end{bmatrix} = \begin{bmatrix} 0.6x_t + 0.1y_t \\ 0.4x_t + 0.9y_t \end{bmatrix} = \begin{bmatrix} 0.6 & 0.1 \\ 0.4 & 0.9 \end{bmatrix} \begin{bmatrix} x_t \\ y_t \end{bmatrix}$$

The matrix $M = \begin{bmatrix} 0.6 & 0.1 \\ 0.4 & 0.9 \end{bmatrix}$ is the **transition matrix** of this dynamical system, because it describes the transition from time t to time $t + 1$. This particular one is a **Markov matrix** (or stochastic matrix): its columns add to 1 and it has no negative entries.

Powers of the transition matrix. Note that $\begin{bmatrix} x_n \\ y_n \end{bmatrix} = M^n \begin{bmatrix} x_0 \\ y_0 \end{bmatrix}$. In other words, M^n describes the transition over n years. In particular, the powers of M are the key to understanding what happens in this model over time.

Equilibrium. $\begin{bmatrix} x_\infty \\ y_\infty \end{bmatrix}$ is an equilibrium if $\begin{bmatrix} x_\infty \\ y_\infty \end{bmatrix} = \begin{bmatrix} 0.6 & 0.1 \\ 0.4 & 0.9 \end{bmatrix} \begin{bmatrix} x_\infty \\ y_\infty \end{bmatrix}$. In other words, $\begin{bmatrix} x_\infty \\ y_\infty \end{bmatrix}$ is an eigenvector with eigenvalue 1.

The 1-eigenspace is $\text{null}\left(\begin{bmatrix} -0.4 & 0.1 \\ 0.4 & -0.1 \end{bmatrix}\right)$, which has basis $\begin{bmatrix} 1 \\ 4 \end{bmatrix}$.

Since $x_\infty + y_\infty = 1$, we conclude that $\begin{bmatrix} x_\infty \\ y_\infty \end{bmatrix} = \frac{1}{1+4} \begin{bmatrix} 1 \\ 4 \end{bmatrix} = \begin{bmatrix} 1/5 \\ 4/5 \end{bmatrix}$.

Hence, the immunization rate in the long term equilibrium is $4/5 = 80\%$.

[Ponder about why this is a reasonable value!]

Comment. What's the other eigenvalue of the transition matrix? No need to compute the characteristic polynomial: we can easily see that it is $0.5 = 0.6 \cdot 0.9 - 0.1 \cdot 0.4$ because the product of the eigenvalues equals the determinant!

The 0.5-eigenspace is spanned by $\begin{bmatrix} -1 \\ 1 \end{bmatrix}$.

Comment. Will the immunization rate always stabilize and approach the long term equilibrium? Yes! This is a consequence of the other eigenvalue of the transition matrix satisfying $|0.5| < 1$. If we start in state $\begin{bmatrix} x_0 \\ y_0 \end{bmatrix} = a \begin{bmatrix} 1 \\ 4 \end{bmatrix} + b \begin{bmatrix} -1 \\ 1 \end{bmatrix}$, then $\begin{bmatrix} x_n \\ y_n \end{bmatrix} = \begin{bmatrix} 0.6 & 0.1 \\ 0.4 & 0.9 \end{bmatrix}^n \begin{bmatrix} x_0 \\ y_0 \end{bmatrix} = 1^n \cdot a \begin{bmatrix} 1 \\ 4 \end{bmatrix} + (0.5)^n \cdot b \begin{bmatrix} -1 \\ 1 \end{bmatrix} \xrightarrow{\text{as } n \rightarrow \infty} a \begin{bmatrix} 1 \\ 4 \end{bmatrix}$.

Random comment. A rule of thumb is that a tetanus vaccination begins to wear off after about 10 years (somewhat in line with the 0.1 transition proportion in this example). However, the tetanus immunization rate in the United States appears to be considerable less than the 80% we found in this (awfully simplistic) example.

<https://www.cdc.gov/mmwr/preview/mmwrhtml/mm5940a3.htm>

Example 101. Compute M^n for $M = \begin{bmatrix} 0.6 & 0.1 \\ 0.4 & 0.9 \end{bmatrix}$.


Solution. In Example 98, we computed that $A = \begin{bmatrix} 6 & 1 \\ 4 & 9 \end{bmatrix}$ had powers $A^n = \frac{1}{5} \begin{bmatrix} 10^n + 4 \cdot 5^n & 10^n - 5^n \\ 4 \cdot 10^n - 4 \cdot 5^n & 4 \cdot 10^n + 5^n \end{bmatrix}$.

Since $M = \frac{1}{10}A$, this implies that $M^n = \frac{1}{10^n}A^n = \frac{1}{5} \begin{bmatrix} 1 + 4 \cdot 0.5^n & 1 - 0.5^n \\ 4 - 4 \cdot 0.5^n & 4 + 0.5^n \end{bmatrix}$.

Note that $M^n \rightarrow \frac{1}{5} \begin{bmatrix} 1 & 1 \\ 4 & 4 \end{bmatrix}$ as $n \rightarrow \infty$. This reflects the fact that $\frac{1}{5} \begin{bmatrix} 1 \\ 4 \end{bmatrix}$ is the long term equilibrium.

Example 102. (extra) Consider a fixed population of people with or without a job. Suppose that, each year, 50% of those unemployed find a job while 10% of those employed lose their job. What is the unemployment rate in the long term equilibrium?

Solution. Let x_t and y_t be the proportions of those employed and unemployed. Proceeding, as in the previous example, the transition matrix is $M = \begin{bmatrix} 0.9 & 0.5 \\ 0.1 & 0.5 \end{bmatrix}$.



The 1-eigenspace of M , that is $\text{null}\left(\begin{bmatrix} -0.1 & 0.5 \\ 0.1 & -0.5 \end{bmatrix}\right)$, has basis $\begin{bmatrix} 5 \\ 1 \end{bmatrix}$. The corresponding equilibrium is $\frac{1}{5+1}\begin{bmatrix} 5 \\ 1 \end{bmatrix}$. In particular, the unemployment rate in the long term equilibrium is $1/6 \approx 16.7\%$.

Example 103. Which of the following are true for all square matrices A ?

- Is it true that A^T has the same eigenvalues as A ?
- Is it true that A^T has the same eigenspaces as A ?
- Is it true that A^T has the same characteristic polynomial as A ?

Solution. True. False. True.

First, note that the characteristic polynomial $\det(A - \lambda I)$ is the same as $\det(A^T - \lambda I)$. [Make sure you can fill in the details of why this is the case!] Hence, the eigenvalues (which are the roots of the characteristic polynomial) are also the same for A and A^T .

On the other hand, A^T and A in general have very different eigenspaces. Take, for instance, the matrix $A = \begin{bmatrix} 0.6 & 0.1 \\ 0.4 & 0.9 \end{bmatrix}$ from Example 100. Then both A and A^T have eigenvalues $\lambda = 0.5, 1$.

However, the 1-eigenspace of A is spanned by $\begin{bmatrix} 1 \\ 4 \end{bmatrix}$, while the 1-eigenspace of A^T is spanned by $\begin{bmatrix} 1 \\ 1 \end{bmatrix}$.

Example 104. Show that a Markov matrix A (so that the columns of A sum to 1) always has eigenvalue 1.

Solution. This follows because the transpose A^T always has $[1 \ 1 \ \dots \ 1]^T$ as a 1-eigenvector (by virtue of the rows of A^T summing to 1). [Make sure that makes sense!]

By the previous example, A must also have eigenvalue 1 (but we have no idea what a 1-eigenvector is until we compute it).


Application: PageRank

Example 105. Suppose the internet consists of only the three webpages A, B, C .

We wish to rank these webpages in order of “importance”.

The idea. Instead of analyzing each webpage (which would be a lot of work!) we will try to only use the information how the pages are linked to each other. The idea being that an “important” page should be linked to from many other pages.

A and B have a link to each other. Also, A links to C and C links to B . If you keep randomly clicking from one webpage to the next, what proportion of the time will you be at each page?




The idea. We will assign ranking to the pages according to how frequently such a random surfer would visit these pages.

Comment. Before we start computing, stop for a moment, and think about how you would rank the webpages.

Solution. Let a_t be the probability that we will be on page A at time t . Likewise, b_t, c_t are the probabilities that we will be on page B or C .

The transition from one state to the next now works exactly as in the previous example. We get the following transition matrix:

$$\begin{bmatrix} a_{t+1} \\ b_{t+1} \\ c_{t+1} \end{bmatrix} = \begin{bmatrix} 0 \cdot a_t + 1 \cdot b_t + 0 \cdot c_t \\ \frac{1}{2} \cdot a_t + 0 \cdot b_t + 1 \cdot c_t \\ \frac{1}{2} \cdot a_t + 0 \cdot b_t + 0 \cdot c_t \end{bmatrix} = \begin{bmatrix} 0 & 1 & 0 \\ \frac{1}{2} & 0 & 1 \\ \frac{1}{2} & 0 & 0 \end{bmatrix} \begin{bmatrix} a_t \\ b_t \\ c_t \end{bmatrix}$$



To find the equilibrium state, we again determine an appropriate 1-eigenvector.

The 1-eigenspace is $\text{null}\left(\begin{bmatrix} -1 & 1 & 0 \\ \frac{1}{2} & -1 & 1 \\ \frac{1}{2} & 0 & -1 \end{bmatrix}\right)$ which has basis $\begin{bmatrix} 2 \\ 2 \\ 1 \end{bmatrix}$.

The corresponding equilibrium state is $\frac{1}{5} \begin{bmatrix} 2 \\ 2 \\ 1 \end{bmatrix}$. In this context, this is also known as the **PageRank vector**.

In other words, after browsing randomly for a long time, there is (about) a $\frac{2}{5} = 40\%$ chance to be at page A , a $\frac{2}{5} = 40\%$ chance to be at page B , and a $\frac{1}{5} = 20\%$ chance to be at page C .


We therefore rank A and B highest (tied), and C lowest.

Just checking. Maybe we were expecting B to be ranked above A , because B is the only page that has two incoming links. However, if we are at page B , then our next click will be to page A , which is why A and B receive equal ranking.

This method of ranking is the famous **PageRank** algorithm (underlying Google’s search algorithm).

By the way, the algorithm is named, not after ranking web“pages”, but after Larry Page (who founded Google in 1998 together with Sergey Brin).


Example 106. Suppose the internet consists of only the four webpages A, B, C, D which link to each other as indicated in the diagram.



Rank these webpages by computing their PageRank vector.

Solution. Recall that we model a random surfer, who randomly clicks on links. Let a_t be the probability that such a surfer will be on page A at time t . Likewise, b_t, c_t, d_t are the probabilities that the surfer will be on page B, C or D .

The transition probabilities are indicated in the diagram to the right. As in the previous example, we obtain the following transition behaviour:



$$\begin{bmatrix} a_{t+1} \\ b_{t+1} \\ c_{t+1} \\ d_{t+1} \end{bmatrix} = \begin{bmatrix} 0 \cdot a_t + \frac{1}{2} \cdot b_t + 1 \cdot c_t + 0 \cdot d_t \\ \frac{1}{3} \cdot a_t + 0 \cdot b_t + 0 \cdot c_t + 0 \cdot d_t \\ \frac{1}{3} \cdot a_t + 0 \cdot b_t + 0 \cdot c_t + 1 \cdot d_t \\ \frac{1}{3} \cdot a_t + \frac{1}{2} \cdot b_t + 0 \cdot c_t + 0 \cdot d_t \end{bmatrix} = \underbrace{\begin{bmatrix} 0 & \frac{1}{2} & 1 & 0 \\ \frac{1}{3} & 0 & 0 & 0 \\ \frac{1}{3} & 0 & 0 & 1 \\ \frac{1}{3} & \frac{1}{2} & 0 & 0 \end{bmatrix}}_{=M} \begin{bmatrix} a_t \\ b_t \\ c_t \\ d_t \end{bmatrix}$$

To find the equilibrium state, we determine an appropriate 1-eigenvector of the transition matrix M .

The 1-eigenspace is $\text{null}(M - 1 \cdot I) = \text{null}\left(\begin{bmatrix} -1 & \frac{1}{2} & 1 & 0 \\ \frac{1}{3} & -1 & 0 & 0 \\ \frac{1}{3} & 0 & -1 & 1 \\ \frac{1}{3} & \frac{1}{2} & 0 & -1 \end{bmatrix}\right)$.

To compute a basis, we perform Gaussian elimination:

$$\begin{bmatrix} -1 & \frac{1}{2} & 1 & 0 \\ \frac{1}{3} & -1 & 0 & 0 \\ \frac{1}{3} & 0 & -1 & 1 \\ \frac{1}{3} & \frac{1}{2} & 0 & -1 \end{bmatrix} \xrightarrow{\text{RREF}} \begin{bmatrix} 1 & 0 & 0 & -2 \\ 0 & 1 & 0 & -\frac{2}{3} \\ 0 & 0 & 1 & -\frac{5}{3} \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

We conclude that the 1-eigenspace has basis $\begin{bmatrix} 2 \\ \frac{2}{3} \\ \frac{5}{3} \\ 1 \end{bmatrix}$. (Note that its entries add up to $2 + \frac{2}{3} + \frac{5}{3} + 1 = \frac{16}{3}$.)

The corresponding equilibrium state is $\frac{3}{16} \begin{bmatrix} 2 \\ \frac{2}{3} \\ \frac{5}{3} \\ 1 \end{bmatrix} \approx \begin{bmatrix} 0.375 \\ 0.125 \\ 0.313 \\ 0.188 \end{bmatrix}$. This is the **PageRank vector**.

[For instance, after browsing randomly for a long time, there is (about) a 12.5% chance to be at page B .] Correspondingly, we rank the pages as $A > C > D > B$.

The real internet. [Google is getting more secretive about this kind of data, so the numbers are estimates from a while ago.]

- Google reports (2016) doing “trillions” of searches per year. [2 trillion means 63,000 searches per second.]
- Google’s search index contains almost 50 billion pages (2016). [Estimated to exceed 100,000,000 gigabytes.]
- More than 1,000,000,000 websites (i.e. hostnames; about 75% not active)

[The “average” user apparently only visits about 100 websites per month; wikipedia.org is one website, consisting of many webpages (more than 2,000,000).]

Gory details. (extra) There's nothing interesting about the Gaussian elimination above. Here are the full details:

$$\begin{array}{c}
 \begin{bmatrix} -1 & \frac{1}{2} & 1 & 0 \\ \frac{1}{3} & -1 & 0 & 0 \\ \frac{1}{3} & 0 & -1 & 1 \\ \frac{1}{3} & \frac{1}{2} & 0 & -1 \end{bmatrix} \xrightarrow{\substack{R_2 + \frac{1}{3}R_1 \Rightarrow R_2 \\ R_3 + \frac{1}{3}R_1 \Rightarrow R_3 \\ R_4 + \frac{1}{3}R_1 \Rightarrow R_4}} \begin{bmatrix} -1 & \frac{1}{2} & 1 & 0 \\ 0 & -\frac{5}{6} & \frac{1}{3} & 0 \\ 0 & \frac{1}{6} & -\frac{2}{3} & 1 \\ 0 & \frac{2}{3} & \frac{1}{3} & -1 \end{bmatrix} \xrightarrow{\substack{R_3 + \frac{1}{5}R_2 \Rightarrow R_3 \\ R_4 + \frac{1}{5}R_2 \Rightarrow R_4}} \begin{bmatrix} -1 & \frac{1}{2} & 1 & 0 \\ 0 & -\frac{5}{6} & \frac{1}{3} & 0 \\ 0 & 0 & -\frac{1}{3} & 1 \\ 0 & 0 & \frac{1}{3} & -1 \end{bmatrix} \\
 \\
 \begin{bmatrix} -1 & \frac{1}{2} & 1 & 0 \\ 0 & -\frac{5}{6} & \frac{1}{3} & 0 \\ 0 & 0 & -\frac{1}{3} & 1 \\ 0 & 0 & 0 & 0 \end{bmatrix} \xrightarrow{R_4 + R_3 \Rightarrow R_4} \begin{bmatrix} -1 & \frac{1}{2} & 1 & 0 \\ 0 & -\frac{5}{6} & \frac{1}{3} & 0 \\ 0 & 0 & -\frac{1}{3} & 1 \\ 0 & 0 & 0 & 0 \end{bmatrix} \xrightarrow{\substack{-1R_1 \Rightarrow R_1 \\ -\frac{6}{5}R_2 \Rightarrow R_2 \\ -\frac{6}{5}R_3 \Rightarrow R_3}} \begin{bmatrix} 1 & -\frac{1}{2} & -1 & 0 \\ 0 & 1 & -\frac{2}{5} & 0 \\ 0 & 0 & 1 & -\frac{5}{3} \\ 0 & 0 & 0 & 0 \end{bmatrix} \xrightarrow{\substack{R_1 + R_3 \Rightarrow R_1 \\ R_2 + \frac{2}{5}R_3 \Rightarrow R_2}} \begin{bmatrix} 1 & -\frac{1}{2} & 0 & -\frac{5}{3} \\ 0 & 1 & 0 & -\frac{2}{3} \\ 0 & 0 & 1 & -\frac{5}{3} \\ 0 & 0 & 0 & 0 \end{bmatrix} \xrightarrow{R_1 + \frac{1}{2}R_2 \Rightarrow R_1} \begin{bmatrix} 1 & 0 & 0 & -2 \\ 0 & 1 & 0 & -\frac{2}{3} \\ 0 & 0 & 1 & -\frac{5}{3} \\ 0 & 0 & 0 & 0 \end{bmatrix}
 \end{array}$$

Practical comment. The transition matrix we would get for the entire internet indexed by Google is prohibitively large (a 50 billion by 50 billion matrix). While gigantic in size, it is a very **sparse matrix**, meaning that almost all of its entries are zero (each column has 50 billion entries but only a handful are nonzero, namely those corresponding to a link to another webpage). This is typical for many applications in linear algebra: we often deal with big but sparse matrices.

Another practical comment. It's not an issue in our simple example, but what if our random surfer gets stuck on a webpage without links? Or, similarly, gets stuck in a loop of links? To deal with these, it is customary to include "teleportation". That is, each time, one of two things happens: with probability p (typically, something like $p = 0.85$) our surfer clicks a link as before; otherwise, with probability $1 - p$, he is teleported to some unrelated other page. Further, if the surfer comes to a page without links, he would teleport away.

A final practical comment. In practical situations, the system might be too large for finding the equilibrium vector by elimination, as we did above. An alternative to elimination is the power method: it is based on the idea that the equilibrium vector is what we expect in the long-term. We can approximate this "long-term" behaviour by simulating a few transitions. For instance, in our example, if we start with the state $[1/4 \ 1/4 \ 1/4 \ 1/4]^T$, which corresponds to equal chances of being on each webpage, then the next state (that is, after one random click) is

$$M \begin{bmatrix} 1/4 \\ 1/4 \\ 1/4 \\ 1/4 \end{bmatrix} = \begin{bmatrix} 0 & \frac{1}{2} & 1 & 0 \\ \frac{1}{3} & 0 & 0 & 0 \\ \frac{1}{3} & 0 & 0 & 1 \\ \frac{1}{3} & \frac{1}{2} & 0 & 0 \end{bmatrix} \begin{bmatrix} 1/4 \\ 1/4 \\ 1/4 \\ 1/4 \end{bmatrix} = \begin{bmatrix} 3/8 \\ 1/12 \\ 1/3 \\ 5/24 \end{bmatrix} = \begin{bmatrix} 0.375 \\ 0.083 \\ 0.333 \\ 0.208 \end{bmatrix}.$$

Note that the ranking of the webpages is already A, C, D, B if we stop right here.

The state after that (that is, after two random clicks) is $M^2 \begin{bmatrix} 1/4 \\ 1/4 \\ 1/4 \\ 1/4 \end{bmatrix} = \begin{bmatrix} 0.375 \\ 0.125 \\ 0.333 \\ 0.167 \end{bmatrix}$, and $M^3 \begin{bmatrix} 1/4 \\ 1/4 \\ 1/4 \\ 1/4 \end{bmatrix} = \begin{bmatrix} 0.396 \\ 0.125 \\ 0.292 \\ 0.188 \end{bmatrix}$.

Observe how we are (overall) approaching the equilibrium vector $\begin{bmatrix} 0.375 \\ 0.125 \\ 0.313 \\ 0.188 \end{bmatrix}$.

Iterating like this is guaranteed to converge to a 1-eigenvector under mild technical assumptions on the transition matrix (for instance, that all its entries be positive; in that case, the other eigenvalues λ satisfy $|\lambda| < 1$ so that their contributions go to zero exponentially, as in Example 100).

Application: Fibonacci numbers

The numbers 0, 1, 1, 2, 3, 5, 8, 13, 21, 34, ... are called **Fibonacci numbers**.

They are defined by the recursion $F_{n+1} = F_n + F_{n-1}$ and $F_0 = 0, F_1 = 1$.

How fast are they growing? Have a look at ratios of Fibonacci numbers:

$$\frac{2}{1} = 2, \frac{3}{2} = 1.5, \frac{5}{3} \approx 1.667, \frac{8}{5} = 1.6, \frac{13}{8} = 1.625, \frac{21}{13} \approx 1.615, \frac{34}{21} \approx 1.619, \dots$$

These ratios approach the **golden ratio** $\varphi = \frac{1+\sqrt{5}}{2} = 1.618\dots$

In other words, it appears that $\lim_{n \rightarrow \infty} \frac{F_{n+1}}{F_n} = \frac{1+\sqrt{5}}{2}$. This indeed follows from Theorem 108 below.

The crucial insight is the following simple observation:


$$F_{n+2} = F_{n+1} + F_n \quad \text{is equivalent to} \quad \begin{bmatrix} F_{n+2} \\ F_{n+1} \end{bmatrix} = \begin{bmatrix} 1 & 1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} F_{n+1} \\ F_n \end{bmatrix}.$$

In particular, $\begin{bmatrix} F_{n+1} \\ F_n \end{bmatrix} = \begin{bmatrix} 1 & 1 \\ 1 & 0 \end{bmatrix}^n \begin{bmatrix} F_1 \\ F_0 \end{bmatrix}$.

Comment. Recurrence equations are discrete analogs of differential equations. We will later see the same idea applied when we reduce the order of a differential equation by introducing additional equations.

Example 107. We model rabbit reproduction as follows.

Each month, every pair of adult rabbits produces one pair of baby rabbit as offspring. Meanwhile, it takes baby rabbits one month to mature to adults.



Comment. In this simplified model, rabbits always come in male/female pairs and no rabbits die. Though these features might make it sound fairly useless, the model may have some merit when describing populations under ideal conditions (unlimited resources) and over short time (no deaths).

Historical comment. The question how many rabbits there are after one year, when starting out with a pair of baby rabbits is famously included in the 1202 textbook of the Italian mathematician Leonardo of Pisa, known as Fibonacci.

Describe the transition from one month to the next.

Solution. Let a_t be the number of adult rabbit pairs after t months. Likewise, b_t is the number of baby rabbit pairs. Then the transition from one month to the next is described by

$$\begin{bmatrix} a_{t+1} \\ b_{t+1} \end{bmatrix} = \begin{bmatrix} a_t + b_t \\ a_t \end{bmatrix} = \begin{bmatrix} 1 & 1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} a_t \\ b_t \end{bmatrix}.$$

That's precisely the transition for the Fibonacci numbers!

It follows that Fibonacci numbers count the number of rabbits in this model.

Comment. Note that the setup is very much as for Markov chains. Here, however, the outgoing values do not add to 100% for each state. Consequently, we cannot expect an equilibrium (and, indeed, the number of rabbits increases without bound).

Everything we observe here about Fibonacci numbers holds for other sequences that satisfy similar recursion equations.

Theorem 108. (Binet's formula) $F_n = \frac{1}{\sqrt{5}} \left[\left(\frac{1+\sqrt{5}}{2} \right)^n - \left(\frac{1-\sqrt{5}}{2} \right)^n \right]$

Proof.

- We already observed that the recurrence $F_{n+2} = F_{n+1} + F_n$ translates into $\begin{bmatrix} F_{n+2} \\ F_{n+1} \end{bmatrix} = \begin{bmatrix} 1 & 1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} F_{n+1} \\ F_n \end{bmatrix}$ and, thus, $\begin{bmatrix} F_{n+1} \\ F_n \end{bmatrix} = \begin{bmatrix} 1 & 1 \\ 1 & 0 \end{bmatrix}^n \begin{bmatrix} F_1 \\ F_0 \end{bmatrix}$.

- We therefore diagonalize $M = \begin{bmatrix} 1 & 1 \\ 1 & 0 \end{bmatrix}$ as $M = PDP^{-1}$ with

$$D = \begin{bmatrix} \lambda_1 & \\ & \lambda_2 \end{bmatrix}, \quad P = \begin{bmatrix} \lambda_1 & \lambda_2 \\ 1 & 1 \end{bmatrix}, \quad \lambda_1 = \frac{1+\sqrt{5}}{2} \approx 1.618, \quad \lambda_2 = \frac{1-\sqrt{5}}{2} \approx -0.618.$$

Comment. λ_1 is the golden ratio!

- It follows that:

$$\begin{aligned} \begin{bmatrix} F_{n+1} \\ F_n \end{bmatrix} &= M^n \begin{bmatrix} F_1 \\ F_0 \end{bmatrix} = PD^n P^{-1} \begin{bmatrix} 1 \\ 0 \end{bmatrix} \\ &= \begin{bmatrix} \lambda_1 & \lambda_2 \\ 1 & 1 \end{bmatrix} \begin{bmatrix} \lambda_1^n & \\ & \lambda_2^n \end{bmatrix} \frac{1}{\lambda_1 - \lambda_2} \begin{bmatrix} 1 & -\lambda_2 \\ -1 & \lambda_1 \end{bmatrix} \begin{bmatrix} 1 \\ 0 \end{bmatrix} \\ &= \begin{bmatrix} \lambda_1^{n+1} & \lambda_2^{n+1} \\ \lambda_1^n & \lambda_2^n \end{bmatrix} \frac{1}{\sqrt{5}} \begin{bmatrix} 1 \\ -1 \end{bmatrix} \\ &= \frac{1}{\sqrt{5}} \begin{bmatrix} \lambda_1^{n+1} - \lambda_2^{n+1} \\ \lambda_1^n - \lambda_2^n \end{bmatrix} \end{aligned}$$

- Hence, $F_n = \frac{1}{\sqrt{5}} (\lambda_1^n - \lambda_2^n)$, which is the claimed formula. □

Comment. For large n , $F_n \approx \frac{1}{\sqrt{5}} \lambda_1^n$ (because λ_2^n becomes very small). In fact, $F_n = \text{round}\left(\frac{1}{\sqrt{5}} \left(\frac{1+\sqrt{5}}{2}\right)^n\right)$.

Back to the quotient of Fibonacci numbers. In particular, because λ_1^n dominates λ_2^n , it is now transparent that the ratios $\frac{F_{n+1}}{F_n}$ approach $\lambda_1 = \frac{1+\sqrt{5}}{2} \approx 1.618$. To be precise, note that

$$\frac{F_{n+1}}{F_n} = \frac{\frac{1}{\sqrt{5}}(\lambda_1^{n+1} - \lambda_2^{n+1})}{\frac{1}{\sqrt{5}}(\lambda_1^n - \lambda_2^n)} = \frac{\lambda_1^{n+1} - \lambda_2^{n+1}}{\lambda_1^n - \lambda_2^n} = \frac{\lambda_1 - \lambda_2 \left(\frac{\lambda_2}{\lambda_1}\right)^n}{1 - \left(\frac{\lambda_2}{\lambda_1}\right)^n} \xrightarrow{n \rightarrow \infty} \frac{\lambda_1 - 0}{1 - 0} = \lambda_1.$$

Comment. It follows from $\lambda_2 < 0$ that the ratios $\frac{F_{n+1}}{F_n}$ approach λ_1 in the alternating fashion that we observed numerically earlier. Can you see that?

Note that, given any Fibonacci-like recursion, we can apply our linear algebra skills in the same fashion. The next example illustrates how this is set up.

Example 109. Suppose the sequence a_n satisfies $a_{n+3} = 3a_{n+2} - 2a_{n+1} + 7a_n$. Write down a matrix-vector version of this recursion.

Solution.
$$\begin{bmatrix} a_{n+3} \\ a_{n+2} \\ a_{n+1} \end{bmatrix} = \begin{bmatrix} 3 & -2 & 7 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} a_{n+2} \\ a_{n+1} \\ a_n \end{bmatrix}$$

Important. If we write $\mathbf{a}_n = \begin{bmatrix} a_{n+2} \\ a_{n+1} \\ a_n \end{bmatrix}$, then this is simply $\mathbf{a}_{n+1} = M\mathbf{a}_n$ with $M = \begin{bmatrix} 3 & -2 & 7 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix}$.

In particular, it follows that $\mathbf{a}_n = M^n \mathbf{a}_0$.

If we compute M^n , then this produces an explicit formula for a_n (the third entry of \mathbf{a}_n). This formula is called a **Binet-like formula** (in the case of the Fibonacci numbers, this is precisely the classical Binet formula).

Review. Fibonacci numbers, Binet formula

Example 110. Consider the sequence a_n defined by $a_{n+2} = 2a_{n+1} + 3a_n$ and $a_0 = -1, a_1 = 5$.

- (a) Determine the first few terms of the sequence.
- (b) Write down a matrix-vector version of the recursion.
- (c) Find a Binet-like formula for a_n .
- (d) Determine $\lim_{n \rightarrow \infty} \frac{a_{n+1}}{a_n}$.

Solution.

(a) $-1, 5, 7, 29, 79, 245, 727, 2189, 6559, \dots$

(b) The recursion can be translated to $\begin{bmatrix} a_{n+2} \\ a_{n+1} \end{bmatrix} = \begin{bmatrix} 2 & 3 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} a_{n+1} \\ a_n \end{bmatrix}$.

(c) **(solution using matrix powers)** Thus, $\begin{bmatrix} a_{n+1} \\ a_n \end{bmatrix} = \begin{bmatrix} 2 & 3 \\ 1 & 0 \end{bmatrix}^n \begin{bmatrix} a_1 \\ a_0 \end{bmatrix}$.

After some work (do it!), we diagonalize $\begin{bmatrix} 2 & 3 \\ 1 & 0 \end{bmatrix} = PDP^{-1}$ with $D = \begin{bmatrix} 3 & 0 \\ 0 & -1 \end{bmatrix}$ and $P = \begin{bmatrix} 3 & -1 \\ 1 & 1 \end{bmatrix}$.

Therefore, $\begin{bmatrix} a_{n+1} \\ a_n \end{bmatrix} = PD^nP^{-1} \begin{bmatrix} a_1 \\ a_0 \end{bmatrix} = \begin{bmatrix} 3 & -1 \\ 1 & 1 \end{bmatrix} \begin{bmatrix} 3^n & 0 \\ 0 & (-1)^n \end{bmatrix} \frac{1}{4} \begin{bmatrix} 1 & 1 \\ -1 & 3 \end{bmatrix} \begin{bmatrix} 5 \\ -1 \end{bmatrix} = \begin{bmatrix} 3^{n+1} - 2(-1)^{n+1} \\ 3^n - 2(-1)^n \end{bmatrix}$
 $= \begin{bmatrix} 3^{n+1} & (-1)^{n+1} \\ 3^n & (-1)^n \end{bmatrix} = \begin{bmatrix} 3 & -1 \\ 1 & 1 \end{bmatrix}^n \begin{bmatrix} 5 \\ -1 \end{bmatrix}$

In particular, $a_n = 3^n - 2(-1)^n$.

(simplified solution) The eigenvalues of $\begin{bmatrix} 2 & 3 \\ 1 & 0 \end{bmatrix}$ are 3 and -1 .

Looking back at our work above, we can see that a_n therefore must have a formula of the form $a_n = C_1 \cdot 3^n + C_2 \cdot (-1)^n$ for some unknown constants C_1, C_2 which we still need to figure out

Using the two initial conditions, we get two equations:

$$(a_0 =) C_1 + C_2 = -1, (a_1 =) 3C_1 - C_2 = 5.$$

Solving, we find $C_1 = 1$ and $C_2 = -2$ so that, in conclusion, $a_n = 3^n - 2(-1)^n$.

(d) It follows from the Binet-like formula that $\lim_{n \rightarrow \infty} \frac{a_{n+1}}{a_n} = 3$ (the eigenvalue of largest absolute value).

Important comment. Right after computing the eigenvalues, we knew that this limit would be 3, except in the special (degenerate) case of $C_1 = 0$.

Definition 111. A sequence a_n satisfying a recursion of the form

$$a_{n+d} = r_1 a_{n+d-1} + r_2 a_{n+d-2} + \dots + r_d a_n$$

is called **C-finite** (or, **constant recursive**) of order d .

For instance. For the Fibonacci numbers, $d = 2$ and $r_1 = r_2 = 1$.

In matrix-vector form.

$$\begin{bmatrix} a_{n+d} \\ a_{n+d-1} \\ \vdots \\ a_{n+1} \end{bmatrix} = \underbrace{\begin{bmatrix} r_1 & r_2 & \dots & r_{d-1} & r_d \\ 1 & & & & 0 \\ & 1 & & & 0 \\ & & \ddots & & \vdots \\ & & & 1 & 0 \end{bmatrix}}_M \begin{bmatrix} a_{n+d-1} \\ a_{n+d-2} \\ \vdots \\ a_n \end{bmatrix}$$

By the same reasoning as for Fibonacci numbers, C -finite sequences have a Binet-like formula:

Theorem 112. (generalized Binet formula) Suppose the recursion matrix M has distinct eigenvalues $\lambda_1, \dots, \lambda_d$. Then

$$a_n = C_1\lambda_1^n + C_2\lambda_2^n + \dots + C_d\lambda_d^n$$

for certain numbers C_1, \dots, C_d .

For instance. For the Fibonacci numbers, $\lambda_1 = \frac{1+\sqrt{5}}{2}$, $\lambda_2 = \frac{1-\sqrt{5}}{2}$, and $C_1 = \frac{1}{\sqrt{5}}$, $C_2 = -\frac{1}{\sqrt{5}}$.

Comment. A little more care is needed in the case that eigenvalues are repeated.

Corollary 113. Under the assumptions of the previous theorem, if λ_1 is the eigenvalue with the largest absolute value and $\lambda_1 > 0$, as well as $\alpha_1 \neq 0$, then $\lim_{n \rightarrow \infty} \frac{a_{n+1}}{a_n} = \lambda_1$.

Proof. This follows from $a_n = C_1\lambda_1^n + C_2\lambda_2^n + \dots + C_d\lambda_d^n$ because, for large n , the term $C_1\lambda_1^n$ dominates the others. Indeed, we have

$$\frac{a_{n+1}}{a_n} = \frac{C_1\lambda_1^{n+1} + C_2\lambda_2^{n+1} + \dots + C_d\lambda_d^{n+1}}{C_1\lambda_1^n + C_2\lambda_2^n + \dots + C_d\lambda_d^n} = \frac{C_1\lambda_1 + C_2\lambda_2\left(\frac{\lambda_2}{\lambda_1}\right)^n + \dots + C_d\lambda_d\left(\frac{\lambda_d}{\lambda_1}\right)^n}{C_1 + C_2\left(\frac{\lambda_2}{\lambda_1}\right)^n + \dots + C_d\left(\frac{\lambda_d}{\lambda_1}\right)^n} \xrightarrow{n \rightarrow \infty} \frac{C_1\lambda_1}{C_1} = \lambda_1.$$

□

Example 114. Consider the sequence a_n defined by $a_{n+3} = 4a_{n+2} - a_{n+1} - 6a_n$ and $a_0 = 0$, $a_1 = -2$, $a_2 = 2$.

- Determine the first few terms of the sequence.
- Find a Binet-like formula for a_n .
- Determine $\lim_{n \rightarrow \infty} \frac{a_{n+1}}{a_n}$.

Solution.

- $0, -2, 2, 10, 50, 178, 602, 1930, 6050, \dots$

Note that this sequence is C -finite of order 3.

- The recursion can be translated to
$$\begin{bmatrix} a_{n+3} \\ a_{n+2} \\ a_{n+1} \end{bmatrix} = \begin{bmatrix} 4 & -1 & -6 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} a_{n+2} \\ a_{n+1} \\ a_n \end{bmatrix}.$$

Expanding by the 2nd row:
$$\begin{vmatrix} 4-\lambda & -1 & -6 \\ 1 & -\lambda & 0 \\ 0 & 1 & -\lambda \end{vmatrix} = -1 \cdot \begin{vmatrix} -1 & -6 \\ 1 & -\lambda \end{vmatrix} - \lambda \cdot \begin{vmatrix} 4-\lambda & -6 \\ 0 & -\lambda \end{vmatrix} = -\lambda^3 + 4\lambda^2 - \lambda - 6$$

The eigenvalues of the transition matrix are the roots of this polynomial: $\lambda = -1, 2, 3$

[You will not be asked to find roots of cubic polynomials by hand.]

Hence, $a_n = C_1 \cdot (-1)^n + C_2 \cdot 2^n + C_3 \cdot 3^n$ and we only need to figure out the two unknowns C_1, C_2, C_3 .

Using the three initial conditions, we get three equations:

$$(a_0 =) C_1 + C_2 + C_3 = 0, (a_1 =) -C_1 + 2C_2 + 3C_3 = -2, (a_2 =) C_1 + 4C_2 + 9C_3 = 2.$$

Solving, we find $C_1 = 1$, $C_2 = -2$ and $C_3 = 1$ so that, in conclusion, $a_n = (-1)^n - 2 \cdot 2^n + 3^n$.

Comment. Do you see how we might have found the characteristic polynomial directly from the recursion?

- It follows from the Binet-like formula that $\lim_{n \rightarrow \infty} \frac{a_{n+1}}{a_n} = 3$ (the eigenvalue of largest absolute value).

Important comment. Right after computing the eigenvalues, we knew that this limit would be 3, except in the special (degenerate) case of $C_3 = 0$.

Example 115. (extra) Consider the sequence a_n defined by $a_{n+2} = 2a_{n+1} + 4a_n$ and $a_0 = 0$, $a_1 = 1$. Determine $\lim_{n \rightarrow \infty} \frac{a_{n+1}}{a_n}$.

Solution. The recursion can be translated to $\begin{bmatrix} a_{n+2} \\ a_{n+1} \end{bmatrix} = \begin{bmatrix} 2 & 4 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} a_{n+1} \\ a_n \end{bmatrix}$.

The eigenvalues of $\begin{bmatrix} 2 & 4 \\ 1 & 0 \end{bmatrix}$ are $1 \pm \sqrt{5}$. Hence, $a_n = C_1(1 + \sqrt{5})^n + C_2(1 - \sqrt{5})^n$ for certain numbers C_1, C_2 .

[Note that we cannot have $C_1 = 0$, because then $a_n = C_2(1 - \sqrt{5})^n$ so that $a_0 = 0$ would imply $C_2 = 0$.]

Therefore, $\lim_{n \rightarrow \infty} \frac{a_{n+1}}{a_n} = 1 + \sqrt{5} \approx 3.23607$.

Comment. With just a little more work, we find the Binet formula $a_n = \frac{(1 + \sqrt{5})^n - (1 - \sqrt{5})^n}{2\sqrt{5}}$.

First few terms of sequence. 0, 1, 2, 8, 24, 80, 256, 832, ...

These are actually related to Fibonacci numbers. Indeed, $a_n = 2^{n-1}F_n$. Can you prove this directly from the recursions? Alternatively, this follows from the Binet formulas.

Example 116. Consider the sequence a_n defined by $a_{n+2} = 2a_{n+1} + 5a_n$ and $a_0 = 0, a_1 = 1$.

- (a) Determine the first few terms of the sequence.
- (b) Find a Binet-like formula for a_n .
- (c) Determine $\lim_{n \rightarrow \infty} \frac{a_{n+1}}{a_n}$.

Solution.

(a) 0, 1, 2, 9, 28, 101, 342, 1189, 4088, ...

(b) The recursion can be translated to $\begin{bmatrix} a_{n+2} \\ a_{n+1} \end{bmatrix} = \begin{bmatrix} 2 & 5 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} a_{n+1} \\ a_n \end{bmatrix}$.

The eigenvalues of $\begin{bmatrix} 2 & 5 \\ 1 & 0 \end{bmatrix}$ are $1 \pm \sqrt{6}$.

Hence, $a_n = C_1(1 + \sqrt{6})^n + C_2(1 - \sqrt{6})^n$ and we only need to figure out the values of C_1 and C_2 .

Using the two initial conditions, we get two equations:

$$(a_0 = 0) \quad C_1 + C_2 = 0, \quad (a_1 = 1) \quad C_1(1 + \sqrt{6}) + C_2(1 - \sqrt{6}) = 1.$$

Solving, we find $C_1 = \frac{1}{2\sqrt{6}}$ and $C_2 = -\frac{1}{2\sqrt{6}}$ so that, in conclusion, $a_n = \frac{(1 + \sqrt{6})^n - (1 - \sqrt{6})^n}{2\sqrt{6}}$.

Comment. Alternatively, we could have proceeded as we did previously in the case of the Fibonacci numbers: starting with the recursion matrix M , we compute its diagonalization $M = PDP^{-1}$. Multiplying out $PD^nP^{-1} \begin{bmatrix} a_1 \\ a_0 \end{bmatrix}$, we obtain the Binet-like formula for a_n . However, this is more work than what we did.

(c) It follows from the Binet-like formula that $\lim_{n \rightarrow \infty} \frac{a_{n+1}}{a_n} = 1 + \sqrt{6} \approx 3.44949$.

Comment. Actually, we don't need the Binet-like formula for this conclusion. Just the eigenvalues and the observation that C_1 cannot be 0 are enough. [We cannot have $C_1 = 0$, because then $a_n = C_2(1 - \sqrt{6})^n$ so that $a_0 = 0$ would imply $C_2 = 0$.]

Another brief look at projections (and reflections)

(projections) Suppose that M is the projection matrix for projecting onto a subspace W .

- The 1-eigenspace of M is W .
- The 0-eigenspace of M is W^\perp .

In particular, M is symmetric.

Why? By definition, the 1-eigenspace of M consists of those vectors that get projected to themselves. But those are precisely the vectors in W (recall that projecting a vector v onto W means producing the vector in W that is closest to v). Can you likewise spell out the situation for the 0-eigenspace?

Note that M cannot have further eigenvalues (because the dimensions of W and W^\perp already add up to the dimension of the space that we are working in).

Because the eigenvalues of M are real and the eigenspaces are orthogonal, the matrix M has a diagonalization of the form $M = PDP^T$ (make sure you can explain why!) which implies that M is symmetric (that's because $M^T = (PDP^T)^T = (P^T)^T D^T P^T = PDP^T = M$).

Example 117. Let A be the matrix for orthogonally projecting onto $W = \text{span}\left\{\begin{bmatrix} 4 \\ 0 \\ 1 \end{bmatrix}, \begin{bmatrix} 0 \\ 2 \\ 1 \end{bmatrix}\right\}$.

- (a) Diagonalize A (without first computing A) as $A = PDP^{-1}$.
 (b) Diagonalize A as $A = PDP^T$.

Comment. This gives us yet another way to get our hands on projection matrices: we can directly write down the matrices P, D for the diagonalization $A = PDP^T$. The main point here is that the diagonalization of a A nicely reveals all the information about the projection.

[Can you see that this is not really a “new” way of computing projection matrices? In particular, note that, if Q is the matrix P with the third column omitted, then $A = QQ^T$.]

Solution.

- (a) The eigenvalues of A are $1, 1, 0$. The 1-eigenspace of A is W (2-dimensional), and the 0-eigenspace is W^\perp (1-dimensional).

We already have a basis for W . On the other hand, $W^\perp = \text{null}\left(\begin{bmatrix} 4 & 0 & 1 \\ 0 & 2 & 1 \end{bmatrix}\right)$ has basis $\begin{bmatrix} -1/4 \\ -1/2 \\ 1 \end{bmatrix}$.

We therefore choose $D = \begin{bmatrix} 1 & & \\ & 1 & \\ & & 0 \end{bmatrix}$ and $P = \begin{bmatrix} 4 & 0 & -1/4 \\ 0 & 2 & -1/2 \\ 1 & 1 & 1 \end{bmatrix}$.

- (b) In order to achieve a diagonalization PDP^T we need to choose P to be orthogonal (which we can do here because the eigenspaces are orthogonal).

Applying Gram–Schmidt to the basis $w_1 = \begin{bmatrix} 4 \\ 0 \\ 1 \end{bmatrix}$, $w_2 = \begin{bmatrix} 0 \\ 2 \\ 1 \end{bmatrix}$ (of the 1-eigenspace), we construct the

orthogonal basis $q_1 = w_1 = \begin{bmatrix} 4 \\ 0 \\ 1 \end{bmatrix}$, $q_2 = w_2 - \frac{w_2 \cdot q_1}{q_1 \cdot q_1} q_1 = \begin{bmatrix} 0 \\ 2 \\ 1 \end{bmatrix} - \frac{1}{17} \begin{bmatrix} 4 \\ 0 \\ 1 \end{bmatrix} = \frac{2}{17} \begin{bmatrix} -2 \\ 17 \\ 8 \end{bmatrix}$.

Next, we normalize the vectors $\begin{bmatrix} 4 \\ 0 \\ 1 \end{bmatrix}$, $\frac{1}{17} \begin{bmatrix} -4 \\ 34 \\ 16 \end{bmatrix}$, $\begin{bmatrix} -1/4 \\ -1/2 \\ 1 \end{bmatrix}$ to $\frac{1}{\sqrt{17}} \begin{bmatrix} 4 \\ 0 \\ 1 \end{bmatrix}$, $\frac{1}{\sqrt{357}} \begin{bmatrix} -2 \\ 17 \\ 8 \end{bmatrix}$, $\frac{1}{\sqrt{21}} \begin{bmatrix} -1 \\ -2 \\ 4 \end{bmatrix}$.

We therefore choose $D = \begin{bmatrix} 1 & & \\ & 1 & \\ & & 0 \end{bmatrix}$ and $P = \begin{bmatrix} 4/\sqrt{17} & -2/\sqrt{357} & -1/\sqrt{21} \\ 0 & 17/\sqrt{357} & -2/\sqrt{21} \\ 1/\sqrt{17} & 8/\sqrt{357} & 4/\sqrt{21} \end{bmatrix}$.

By the way. Multiplying out $A = PDP^T$, we can find that $A = \frac{1}{21} \begin{bmatrix} 20 & -2 & 4 \\ -2 & 17 & 8 \\ 4 & 8 & 5 \end{bmatrix}$ as in Example 55.

Example 118. Let A be the matrix for orthogonally projecting onto $W = \text{span}\left\{\begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}, \begin{bmatrix} -1 \\ 0 \\ 1 \end{bmatrix}\right\}$.

- (a) Diagonalize A (without first computing A) as $A = PDP^T$.
 (b) Is A invertible, orthogonal, symmetric?

Solution.

- (a) The eigenvalues of A are $1, 1, 0$. The 1-eigenspace of A is W (2-dimensional), and the 0-eigenspace is W^\perp (1-dimensional). Note that we are lucky and already have an orthogonal basis for W . On the other

hand, $W^\perp = \text{null}\left(\begin{bmatrix} 1 & 1 & 1 \\ -1 & 0 & 1 \end{bmatrix}\right)$ has basis $\begin{bmatrix} 1 \\ -2 \\ 1 \end{bmatrix}$.

We therefore choose $D = \begin{bmatrix} 1 & & \\ & 1 & \\ & & 0 \end{bmatrix}$ and, after normalizing columns, $P = \begin{bmatrix} 1/\sqrt{3} & -1/\sqrt{2} & 1/\sqrt{6} \\ 1/\sqrt{3} & 0 & -2/\sqrt{6} \\ 1/\sqrt{3} & 1/\sqrt{2} & 1/\sqrt{6} \end{bmatrix}$.

- (b) A is not invertible (because 0 is an eigenvalue) and therefore also cannot be orthogonal. Like any projection matrix, A is symmetric.

By the way. Multiplying out $A = PDP^T$, we can find that $A = \frac{1}{6} \begin{bmatrix} 1 & -2 & 1 \\ -2 & 4 & -2 \\ 1 & -2 & 1 \end{bmatrix}$.

(reflections) Suppose that M is the matrix for reflecting through the plane W in $\mathbb{3}$ -space.

- The 1 -eigenspace of M is W . (dimension 2)
- The -1 -eigenspace of M is W^\perp . (dimension 1)

In particular, M is symmetric.

Why? By definition, the 1 -eigenspace of M consists of those vectors that get reflected to themselves. But those are precisely the vectors in the plane W (only vectors on the plane are unchanged by the reflection). On the other hand, the -1 -eigenspace consists of those vectors v that get reflected to $-v$ (the exact opposite direction). These are precisely the vectors orthogonal to the plane.

As in the case of projection matrices, because the eigenvalues are real and the eigenspaces are orthogonal, the reflection matrices are symmetric.

Comment. In this context, the line W^\perp is often called the **normal line** of the plane W .

Example 119. Let A be the matrix for reflecting through the plane $W = \text{span}\left\{\begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}, \begin{bmatrix} -1 \\ 0 \\ 1 \end{bmatrix}\right\}$.

- (a) Diagonalize A (without first computing A) as $A = PDP^T$.
- (b) Is A invertible, orthogonal, symmetric?

Solution.

- (a) The eigenvalues of A are $1, 1, -1$. The 1 -eigenspace of A is W , and the -1 -eigenspace is W^\perp .
In order to achieve a diagonalization PDP^T we need to choose P to be orthogonal (which we can do here because the eigenspaces are orthogonal).

As in the previous example, $W^\perp = \text{span}\left\{\begin{bmatrix} 1 \\ -2 \\ 1 \end{bmatrix}\right\}$.

We therefore choose $D = \begin{bmatrix} 1 & & \\ & 1 & \\ & & -1 \end{bmatrix}$ and, after normalizing columns, $P = \begin{bmatrix} 1/\sqrt{3} & -1/\sqrt{2} & 1/\sqrt{6} \\ 1/\sqrt{3} & 0 & -2/\sqrt{6} \\ 1/\sqrt{3} & 1/\sqrt{2} & 1/\sqrt{6} \end{bmatrix}$.

- (b) A is invertible (because 0 is not an eigenvalue).
Like any reflection matrix, A is symmetric.
Finally, note that $A^2 = I$ (reflecting twice isn't doing anything), so that $A^{-1} = A$. It follows that A is orthogonal, because $A^{-1} = A = A^T$.

By the way. Multiplying out $A = PDP^T$, we can find that $A = \frac{1}{3} \begin{bmatrix} 2 & 2 & -1 \\ 2 & -1 & 2 \\ -1 & 2 & 2 \end{bmatrix}$.

Comment. Similarly, a $n \times n$ matrix corresponds to a reflection (through a hyperplane) if and only if it has a $(n - 1)$ -dimensional 1 -eigenspace and a 1 -dimensional -1 -eigenspace and these two spaces are orthogonal.

An alternative way of computing reflection matrices. Realize that, if n is the vector orthogonal to the plane (i.e. n is the normal vector of the plane), then reflecting v means sending it to $v - 2(\text{projection of } v \text{ onto } n)$.

We already observed that $n = \begin{bmatrix} 1 \\ -2 \\ 1 \end{bmatrix}$.

Hence, the reflection of v is $v - 2(\text{projection of } v \text{ onto } n) = v - 2\frac{n \cdot v}{n \cdot n}n = v - 2\frac{nn^T v}{n^T n} = \left(I - 2\frac{nn^T}{n^T n}\right)v$.

Accordingly, the reflection matrix is $A = I - 2\frac{nn^T}{n^T n} = \begin{bmatrix} 1 & & \\ & 1 & \\ & & 1 \end{bmatrix} - \frac{2}{6} \begin{bmatrix} 1 & -2 & 1 \\ -2 & 4 & -2 \\ 1 & -2 & 1 \end{bmatrix} = \frac{1}{3} \begin{bmatrix} 2 & 2 & -1 \\ 2 & -1 & 2 \\ -1 & 2 & 2 \end{bmatrix}$.

Comment. In other words, we got A from subtracting 2 times the projection matrix onto n from I .

Application: Linear differential equations

Example 120. (warmup) Solve the differential equation (DE) $y' = 2$.

Solution. From calculus, we know that the solutions are of the form $y(t) = 2t + C$.

Comment. To get a unique solution, we need to specify additional information, like an initial condition.

Example 121. (warmup) Solve the initial value problem (IVP) $y' = 2$, $y(0) = 1$.

Solution. This has the unique solution $y(t) = 2t + 1$.

Example 122. Which functions $y(t)$ satisfy the differential equation $y' = y$?

Solution. $y(t) = e^t$ and, more generally, $y(t) = Ce^t$. (And nothing else.)

(exponential function) e^t is the unique solution to $y' = y$, $y(0) = 1$.

From here, it follows that $e^t = 1 + t + \frac{t^2}{2!} + \frac{t^3}{3!} + \dots$

The latter is the Taylor series for e^t at $t = 0$ that we have seen in Calculus II.

Important note. We can actually construct this infinite sum directly from $y' = y$ and $y(0) = 1$.

Indeed, observe how each term, when differentiated, produces the term before it. For instance, $\frac{d}{dt} \frac{t^3}{3!} = \frac{t^2}{2!}$.

Example 123. Show that the differential equation $y' = 3y$ is solved by $y(t) = Ce^{3t}$.

Solution. Indeed, if $y(t) = Ce^{3t}$, then $y'(t) = 3Ce^{3t} = 3y(t)$.

Comment. It is important to realize that we can always easily check whether a function solves a differential equation. This means that (although you might be unfamiliar with the techniques for solving) you can use computer algebra systems like Sage to solve differential equations without trust issues.

Example 124. Solve the differential equation $y' = ay$ with initial condition $y(0) = y_0$.

Solution. As in the previous example, the general solution to $y' = ay$ is $y(t) = Ce^{at}$.

Since $y(0) = Ce^0 = C = y_0$, we conclude that the unique solution to the IVP is $y(t) = e^{at}y_0$.

Comment. It looks silly to write $e^{at}y_0$ instead of y_0e^{at} here, but we will soon replace the number a with a matrix A , and in that case only $e^{At}y_0$ makes sense.

Example 125. Our goal is to solve (systems of) differential equations like:

$$\begin{aligned} y_1' &= 2y_1 & y_1(0) &= 1 \\ y_2' &= -y_1 + 3y_2 + y_3 & y_2(0) &= 0 \\ y_3' &= -y_1 + y_2 + 3y_3 & y_3(0) &= 2 \end{aligned}$$

In matrix form, this becomes

$$\mathbf{y}' = \begin{bmatrix} 2 & 0 & 0 \\ -1 & 3 & 1 \\ -1 & 1 & 3 \end{bmatrix} \mathbf{y}, \quad \mathbf{y}(0) = \begin{bmatrix} 1 \\ 0 \\ 2 \end{bmatrix}.$$

The key idea will be to solve $\mathbf{y}' = A\mathbf{y}$ by introducing e^{At} .

Theorem 126. The solution to $\mathbf{y}' = A\mathbf{y}$, $\mathbf{y}(0) = \mathbf{y}_0$ is $\mathbf{y}(t) = e^{At}\mathbf{y}_0$.

Recall from Example 124 that the solution to $y' = ay$, $y(0) = y_0$ is $y(t) = e^{at}y_0$. Here, however, At is a matrix and so we need to make sense of the matrix exponential. Next time, we will define e^A by the familiar Taylor series for e^x .

Definition 127. Let A be $n \times n$. The **matrix exponential** is

$$e^A = I + A + \frac{1}{2!}A^2 + \frac{1}{3!}A^3 + \dots$$

Why? As a consequence of this definition (which is the motivation for that definition in the first place),

$$\begin{aligned} \frac{d}{dt}e^{At} &= \frac{d}{dt}\left[I + At + \frac{1}{2!}A^2t^2 + \frac{1}{3!}A^3t^3 + \dots\right] \\ &= 0 + A + A^2t + \frac{1}{2!}A^3t^2 + \dots = Ae^{At}. \end{aligned}$$

Therefore, $\mathbf{y}(t) = e^{At}\mathbf{y}_0$ indeed solves the initial value problem $\mathbf{y}' = A\mathbf{y}$, $\mathbf{y}(0) = \mathbf{y}_0$.

How to actually compute e^A ? Well, this Taylor series involves the powers A^n of A . How would you compute, say, A^{100} ? The answer is diagonalization!

Theorem 128. Suppose $A = PDP^{-1}$. Then, $e^A = Pe^DP^{-1}$.

Why? Recall that, if $A = PDP^{-1}$, then $A^n = PD^nP^{-1}$.

$$\begin{aligned} e^A &= I + A + \frac{1}{2!}A^2 + \frac{1}{3!}A^3 + \dots \\ &= I + PDP^{-1} + \frac{1}{2!}PD^2P^{-1} + \frac{1}{3!}PD^3P^{-1} + \dots \\ &= P\left(I + D + \frac{1}{2!}D^2 + \frac{1}{3!}D^3 + \dots\right)P^{-1} = Pe^DP^{-1} \end{aligned}$$

Comment. By the same argument, if $A = PDP^{-1}$, then $f(A) = Pf(D)P^{-1}$ for every “nice” function f . Here, “nice” means that f has a convergent Taylor series $f(x) = \sum_{n \geq 0} a_n x^n$.

More explicitly, if $A = P \operatorname{diag}(\lambda_1, \dots, \lambda_n) P^{-1}$, then $f(A) = P \operatorname{diag}(f(\lambda_1), \dots, f(\lambda_n)) P^{-1}$.

Example 129. If $A = \begin{bmatrix} 2 & 0 \\ 0 & 5 \end{bmatrix}$, then $A^{100} = \begin{bmatrix} 2^{100} & 0 \\ 0 & 5^{100} \end{bmatrix}$.

Example 130. If $A = \begin{bmatrix} 2 & 0 \\ 0 & 5 \end{bmatrix}$, then $e^A = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} + \begin{bmatrix} 2 & 0 \\ 0 & 5 \end{bmatrix} + \frac{1}{2!} \begin{bmatrix} 2^2 & 0 \\ 0 & 5^2 \end{bmatrix} + \dots = \begin{bmatrix} e^2 & 0 \\ 0 & e^5 \end{bmatrix}$.

Clearly, this works to obtain e^D for every diagonal matrix D .

In particular, for $At = \begin{bmatrix} 2t & 0 \\ 0 & 5t \end{bmatrix}$, $e^{At} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} + \begin{bmatrix} 2t & 0 \\ 0 & 5t \end{bmatrix} + \frac{1}{2!} \begin{bmatrix} (2t)^2 & 0 \\ 0 & (5t)^2 \end{bmatrix} + \dots = \begin{bmatrix} e^{2t} & 0 \\ 0 & e^{5t} \end{bmatrix}$.

Example 131. (homework) Diagonalize $A = \begin{bmatrix} 2 & 0 & 0 \\ -1 & 3 & 1 \\ -1 & 1 & 3 \end{bmatrix}$.

Solution. (final solution only) $A = PDP^{-1}$ with $P = \begin{bmatrix} 1 & 1 & 0 \\ 1 & 0 & 1 \\ 0 & 1 & 1 \end{bmatrix}$ and $D = \begin{bmatrix} 2 & & \\ & 2 & \\ & & 4 \end{bmatrix}$.

Example 132. Solve the initial value problem

$$\mathbf{y}' = \begin{bmatrix} 2 & 0 & 0 \\ -1 & 3 & 1 \\ -1 & 1 & 3 \end{bmatrix} \mathbf{y}, \quad \mathbf{y}(0) = \begin{bmatrix} 1 \\ 2 \\ 1 \end{bmatrix}.$$

Solution. Recall that the solution to $\mathbf{y}' = A\mathbf{y}$, $\mathbf{y}(0) = \mathbf{y}_0$ is $\mathbf{y} = e^{At}\mathbf{y}_0$.

- First, we diagonalize:

For $A = \begin{bmatrix} 2 & 0 & 0 \\ -1 & 3 & 1 \\ -1 & 1 & 3 \end{bmatrix}$, $A = PDP^{-1}$ with $P = \begin{bmatrix} 1 & 1 & 0 \\ 1 & 0 & 1 \\ 0 & 1 & 1 \end{bmatrix}$ and $D = \begin{bmatrix} 2 & & \\ & 2 & \\ & & 4 \end{bmatrix}$. (That's homework!)

- We can then compute the solution $\mathbf{y}(t) = e^{At}\mathbf{y}_0$:

$$\begin{aligned} \mathbf{y}(t) = e^{At}\mathbf{y}_0 &= Pe^{Dt}P^{-1}\mathbf{y}_0 \\ &= \begin{bmatrix} 1 & 1 & 0 \\ 1 & 0 & 1 \\ 0 & 1 & 1 \end{bmatrix} \begin{bmatrix} e^{2t} & & \\ & e^{2t} & \\ & & e^{4t} \end{bmatrix} \begin{bmatrix} 1 & 1 & 0 \\ 1 & 0 & 1 \\ 0 & 1 & 1 \end{bmatrix}^{-1} \begin{bmatrix} 1 \\ 2 \\ 1 \end{bmatrix} \\ &= \begin{bmatrix} 1 & 1 & 0 \\ 1 & 0 & 1 \\ 0 & 1 & 1 \end{bmatrix} \begin{bmatrix} e^{2t} & & \\ & e^{2t} & \\ & & e^{4t} \end{bmatrix} \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix} \\ &= \begin{bmatrix} 1 & 1 & 0 \\ 1 & 0 & 1 \\ 0 & 1 & 1 \end{bmatrix} \begin{bmatrix} e^{2t} \\ 0 \\ e^{4t} \end{bmatrix} = \begin{bmatrix} e^{2t} \\ e^{2t} + e^{4t} \\ e^{4t} \end{bmatrix} \end{aligned}$$

Comment. It is not necessary to compute $\begin{bmatrix} 1 & 1 & 0 \\ 1 & 0 & 1 \\ 0 & 1 & 1 \end{bmatrix}^{-1}$ (of course, you could do it, but that's more work).

Instead, recall that $A^{-1}\mathbf{b}$ is the unique solution to $A\mathbf{x} = \mathbf{b}$. Here, solving $\begin{bmatrix} 1 & 1 & 0 \\ 1 & 0 & 1 \\ 0 & 1 & 1 \end{bmatrix} \mathbf{x} = \begin{bmatrix} 1 \\ 2 \\ 1 \end{bmatrix}$, we find $\mathbf{x} = \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix}$.

Check. $\mathbf{y} = \begin{bmatrix} e^{2t} \\ e^{2t} + e^{4t} \\ e^{4t} \end{bmatrix}$ indeed solves the original problem:

$$\mathbf{y}' = \begin{bmatrix} 2e^{2t} \\ 2e^{2t} + 4e^{4t} \\ 4e^{4t} \end{bmatrix} \stackrel{\checkmark}{=} \begin{bmatrix} 2 & 0 & 0 \\ -1 & 3 & 1 \\ -1 & 1 & 3 \end{bmatrix} \begin{bmatrix} e^{2t} \\ e^{2t} + e^{4t} \\ e^{4t} \end{bmatrix}, \quad \mathbf{y}(0) = \begin{bmatrix} 1 \\ 1+1 \\ 1 \end{bmatrix} \stackrel{\checkmark}{=} \begin{bmatrix} 1 \\ 2 \\ 1 \end{bmatrix}$$

Example 133. We only discuss linear differential equations (DEs). Non-linear DEs include $y' = y^2 + 1$ or the second-order equation $y'' = \sin(ty') + y$.

The order of a DE indicates the highest occurring derivative.

Note, however, that $y'' = \sin(t)y' + y$ is a linear DE, because y and its derivatives occur linearly.

We will see here how to solve those linear DEs which have constant coefficients. That is, the coefficients of y are constants, as opposed to functions (like $\sin(t)$) depending on t .

Review.

- The solution to $y' = Ay$, $y(0) = y_0$ is $y(t) = e^{At}y_0$.
Why? Because $y'(t) = Ae^{At}y_0 = Ay(t)$ and $y(0) = e^{0A}y_0 = y_0$.
- If we have the diagonalization $A = PDP^{-1}$, then $e^A = Pe^DP^{-1}$ (and $e^{At} = Pe^{Dt}P^{-1}$).
- If $A = \begin{bmatrix} 2 & 0 \\ 0 & 5 \end{bmatrix}$, then $e^A = \begin{bmatrix} e^2 & 0 \\ 0 & e^5 \end{bmatrix}$ and $e^{At} = \begin{bmatrix} e^{2t} & 0 \\ 0 & e^{5t} \end{bmatrix}$.

Example 134. Solve the initial value problem $y' = \begin{bmatrix} 0 & -2 \\ -1 & 1 \end{bmatrix}y$, $y(0) = \begin{bmatrix} 3 \\ 0 \end{bmatrix}$.

Solution.

- $A = \begin{bmatrix} 0 & -2 \\ -1 & 1 \end{bmatrix}$ has characteristic polynomial $-\lambda(1-\lambda) - 2 = (\lambda+1)(\lambda-2)$.
Hence, the eigenvalues of A are $-1, 2$.
The -1 -eigenspace $\text{null}\left(\begin{bmatrix} 1 & -2 \\ -1 & 2 \end{bmatrix}\right)$ has basis $\begin{bmatrix} 2 \\ 1 \end{bmatrix}$.
The 2 -eigenspace $\text{null}\left(\begin{bmatrix} -2 & -2 \\ -1 & -1 \end{bmatrix}\right)$ has basis $\begin{bmatrix} -1 \\ 1 \end{bmatrix}$.
Hence, $A = PDP^{-1}$ with $P = \begin{bmatrix} 2 & -1 \\ 1 & 1 \end{bmatrix}$ and $D = \begin{bmatrix} -1 & \\ & 2 \end{bmatrix}$.
- Finally, we compute the solution $y(t) = e^{At}y_0$:

$$\begin{aligned} y(t) &= Pe^{Dt}P^{-1}y_0 \\ &= \underbrace{\begin{bmatrix} 2 & -1 \\ 1 & 1 \end{bmatrix}}_{\begin{bmatrix} 2e^{-t} & -e^{2t} \\ e^{-t} & e^{2t} \end{bmatrix}} \underbrace{\begin{bmatrix} e^{-t} & \\ & e^{2t} \end{bmatrix}}_{\begin{bmatrix} 1 & \\ & -1 \end{bmatrix}} \frac{1}{3} \begin{bmatrix} 1 & 1 \\ -1 & 2 \end{bmatrix} \begin{bmatrix} 3 \\ 0 \end{bmatrix} = \begin{bmatrix} 2e^{-t} + e^{2t} \\ e^{-t} - e^{2t} \end{bmatrix} \end{aligned}$$

Check. Since it is simple to check, it would be almost criminal to not verify that $y(0) = \begin{bmatrix} 2+1 \\ 1-1 \end{bmatrix} = \begin{bmatrix} 3 \\ 0 \end{bmatrix}$.

Example 135. (homework) Suppose that $e^{Mt} = \frac{1}{10} \begin{bmatrix} e^t + 9e^{2t} & 3e^t - 3e^{2t} \\ 3e^t - 3e^{2t} & 9e^t + e^{2t} \end{bmatrix}$.

- Without doing any computations, determine M^n .
- What is M ?
- Without doing any computations, determine the eigenvalues and eigenvectors of M .

Solution.

- (a) Recall that $e^{Mt} = Pe^{Dt}P^{-1}$ while $M^n = PD^nP^{-1}$, provided that $M = PDP^{-1}$. The fact the formula for e^{Mt} features e^t and e^{2t} , means that the eigenvalues of M must be 1 and 2. Hence,

$$D = \begin{bmatrix} 1 & \\ & 2 \end{bmatrix}, \quad e^{Dt} = \begin{bmatrix} e^t & \\ & e^{2t} \end{bmatrix}, \quad D^n = \begin{bmatrix} 1 & \\ & 2^n \end{bmatrix}.$$

Therefore, we just need to replace e^t by $1^n = 1$ as well as e^{2t} by 2^n to get:

$$M^n = \frac{1}{10} \begin{bmatrix} 1 + 9 \cdot 2^n & 3 - 3 \cdot 2^n \\ 3 - 3 \cdot 2^n & 9 + 2^n \end{bmatrix}$$

- (b) In particular, we see that the underlying matrix is $M = M^1 = \frac{1}{10} \begin{bmatrix} 1 + 9 \cdot 2 & 3 - 3 \cdot 2 \\ 3 - 3 \cdot 2 & 9 + 2 \end{bmatrix} = \frac{1}{10} \begin{bmatrix} 19 & -3 \\ -3 & 11 \end{bmatrix}$.

[Alternatively, we can find M by computing $\frac{d}{dt}e^{Mt} = Me^{Mt}$ and then setting $t = 0$.]

- (c) The eigenvalues are 1 and 2.

Looking at the coefficients of e^t in the first column of e^{Mt} , we can see that $\begin{bmatrix} 1 \\ 3 \end{bmatrix}$ is a 1-eigenvector.

[We can also look the second column of e^{Mt} , to obtain $\begin{bmatrix} 3 \\ 9 \end{bmatrix}$ which is a multiple and thus equivalent.]

Likewise, we find that $\begin{bmatrix} 9 \\ -3 \end{bmatrix}$ or, equivalently, $\begin{bmatrix} -3 \\ 1 \end{bmatrix}$ is a 2-eigenvector.

Higher-order differential equations

Example 136. Write the (second-order) differential equation $y'' = 2y' + y$ as a system of (first-order) differential equations.

Solution. Write $y_1 = y$ and $y_2 = y'$. Then $y'' = 2y' + y$ becomes $y_2' = 2y_2 + y_1$.

Therefore, $y'' = 2y' + y$ translates into the first-order system $\begin{cases} y_1' = y_2 \\ y_2' = y_1 + 2y_2 \end{cases}$.

In matrix form, this is $\mathbf{y}' = \begin{bmatrix} 0 & 1 \\ 1 & 2 \end{bmatrix} \mathbf{y}$.

Comment. Hence, we care about systems of differential equations, even if we work with just one function.

Note. The “trick” of looking at the pair $\begin{bmatrix} y \\ y' \end{bmatrix}$ instead of a single function is what we used to translate the Fibonacci recurrence into a 2×2 system.

Example 137. Write the (third-order) differential equation $y''' = 3y'' - 2y' + y$ as a system of (first-order) differential equations.

Solution. Write $y_1 = y$, $y_2 = y'$ and $y_3 = y''$.

Then, $y''' = 3y'' - 2y' + y$ translates into the first-order system $\begin{cases} y_1' = y_2 \\ y_2' = y_3 \\ y_3' = y_1 - 2y_2 + 3y_3 \end{cases}$.

In matrix form, this is $\mathbf{y}' = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & -2 & 3 \end{bmatrix} \mathbf{y}$.

Example 138. Consider the following system of (second-order) initial value problems:

$$\begin{aligned} y_1'' &= 2y_1' - 3y_2' + 7y_2 & y_1(0) &= 2, \quad y_1'(0) = 3, \quad y_2(0) = -1, \quad y_2'(0) = 1 \\ y_2'' &= 4y_1' + y_2' - 5y_1 \end{aligned}$$

Write it as a first-order initial value problem in the form $\mathbf{y}' = A\mathbf{y}$, $\mathbf{y}(0) = \mathbf{y}_0$.

Solution. Introduce $y_3 = y_1'$ and $y_4 = y_2'$. Then, the given system translates into

$$\mathbf{y}' = \begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 7 & 2 & -3 \\ -5 & 0 & 4 & 1 \end{bmatrix} \mathbf{y}, \quad \mathbf{y}(0) = \begin{bmatrix} 2 \\ -1 \\ 3 \\ 1 \end{bmatrix}.$$

The Jordan normal form

Note that we currently only know how to compute e^{At} when A is diagonalizable. Our next goal is to be able to compute the matrix exponential for all matrices.

Example 139. Diagonalize, if possible, the matrix $A = \begin{bmatrix} 4 & 1 \\ & 4 \end{bmatrix}$.

Solution. The eigenvalues of A are 4, 4.

However, the 4-eigenspace $\text{null}\left(\begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}\right)$ is only 1-dimensional.

Hence, A is not diagonalizable.

Definition 140. A λ -Jordan block is a matrix of the form $\begin{bmatrix} \lambda & 1 & & \\ & \lambda & \ddots & \\ & & \ddots & 1 \\ & & & \lambda \end{bmatrix}$.

Note that if this matrix is $m \times m$, then its only eigenvalue is λ (repeated m times).

As in the previous example, the λ -eigenspace is 1-dimensional (which is as small as possible).

Theorem 141. (Jordan normal form) Every $n \times n$ matrix A can be written as $A = PJP^{-1}$, where J is a block diagonal matrix

$$J = \begin{bmatrix} J_1 & & & \\ & J_2 & & \\ & & \ddots & \\ & & & J_r \end{bmatrix}$$

with each J_i a Jordan block. J is called the **Jordan normal form** of A .

Up to the ordering of the Jordan blocks, the Jordan normal form of A is unique.

Comment. If A is diagonalizable, then J is just a usual diagonal matrix.

Example 142. What are the possible Jordan normal forms of a 3×3 matrix with eigenvalues 4, 4, 4?

Solution. $\begin{bmatrix} 4 & & \\ & 4 & \\ & & 4 \end{bmatrix}, \begin{bmatrix} 4 & & \\ & 4 & 1 \\ & & 4 \end{bmatrix}, \begin{bmatrix} 4 & 1 & \\ & 4 & 1 \\ & & 4 \end{bmatrix}$

The dimension of the 4-eigenspace equals the number of Jordan blocks: 3, 2, 1, respectively.

Comment. Note that, say, $\begin{bmatrix} 4 & 1 & \\ & 4 & \\ & & 4 \end{bmatrix}$ is equivalent to $\begin{bmatrix} 4 & & \\ & 4 & 1 \\ & & 4 \end{bmatrix}$ because the ordering of the diagonal blocks does not matter (as you know from diagonalization).

Example 143.

- (a) What are the possible Jordan normal forms of a 3×3 matrix with eigenvalues 3, 3, 3?
- (b) What are the possible Jordan normal forms of a 4×4 matrix with eigenvalues 3, 3, 3, 3?
- (c) What if the matrix is 5×5 and has eigenvalues 4, 4, 3, 3, 3?

Solution.

(a) $\begin{bmatrix} 3 & & \\ & 3 & \\ & & 3 \end{bmatrix}, \begin{bmatrix} 3 & & \\ & 3 & 1 \\ & & 3 \end{bmatrix}, \begin{bmatrix} 3 & 1 & \\ & 3 & 1 \\ & & 3 \end{bmatrix}$

The dimension of the 3-eigenspace equals the number of Jordan blocks: 3, 2, 1, respectively.

Comment. Note that, say, $\begin{bmatrix} 3 & 1 & \\ & 3 & \\ & & 3 \end{bmatrix}$ is equivalent to $\begin{bmatrix} 3 & & \\ & 3 & 1 \\ & & 3 \end{bmatrix}$ because the ordering of the diagonal blocks does not matter (as you know from diagonalization).

(b) Now, there are 5 possibilities:

$\begin{bmatrix} 3 & & & \\ & 3 & & \\ & & 3 & \\ & & & 3 \end{bmatrix}, \begin{bmatrix} 3 & & & \\ & 3 & & \\ & & 3 & 1 \\ & & & 3 \end{bmatrix}, \begin{bmatrix} 3 & 1 & & \\ & 3 & & \\ & & 3 & 1 \\ & & & 3 \end{bmatrix}, \begin{bmatrix} 3 & & & \\ & 3 & 1 & \\ & & 3 & 1 \\ & & & 3 \end{bmatrix}, \begin{bmatrix} 3 & 1 & & \\ & 3 & 1 & \\ & & 3 & 1 \\ & & & 3 \end{bmatrix}$

The dimension of the 3-eigenspace equals the number of Jordan blocks: 4, 3, 2, 2, 1, respectively.

(c) $\begin{bmatrix} 3 & & & & \\ & 3 & & & \\ & & 3 & & \\ & & & 4 & \\ & & & & 4 \end{bmatrix}, \begin{bmatrix} 3 & & & & \\ & 3 & & & \\ & & 3 & & \\ & & & 4 & 1 \\ & & & & 4 \end{bmatrix}, \begin{bmatrix} 3 & & & & \\ & 3 & 1 & & \\ & & 3 & & \\ & & & 4 & \\ & & & & 4 \end{bmatrix}, \begin{bmatrix} 3 & & & & \\ & 3 & 1 & & \\ & & 3 & & \\ & & & 4 & 1 \\ & & & & 4 \end{bmatrix}, \begin{bmatrix} 3 & 1 & & & \\ & 3 & 1 & & \\ & & 3 & & \\ & & & 4 & \\ & & & & 4 \end{bmatrix}, \begin{bmatrix} 3 & 1 & 1 & & \\ & 3 & 1 & & \\ & & 3 & & \\ & & & 4 & 1 \\ & & & & 4 \end{bmatrix}$

Note that this is just all possible (namely, 3) Jordan normal forms of a 3×3 matrix with eigenvalues 3, 3, 3 combined with all possible (namely, 2) Jordan normal forms of a 2×2 matrix with eigenvalues 4, 4. In total, that makes $3 \cdot 2 = 6$ possibilities.

Comment. Let $p(n)$ be the number of inequivalent Jordan normal forms of an $n \times n$ matrix with a single eigenvalue, n times repeated. We have seen that $p(2) = 2$, $p(3) = 3$, $p(4) = 5$. Note that $p(n)$ is equal to the number of ways of writing n as an ordered sum of positive integers: for instance, $p(4) = 5$ because $4 = 3 + 1 = 2 + 2 = 2 + 1 + 1 = 1 + 1 + 1 + 1$.

$p(n)$ is referred to as the **partition function** and, surprisingly, is a remarkably interesting mathematical object.

[https://en.wikipedia.org/wiki/Partition_function_\(number_theory\)](https://en.wikipedia.org/wiki/Partition_function_(number_theory))

Example 144. (summary of small cases)

(a) There are 2 possible Jordan normal forms of a 2×2 matrix with eigenvalues λ, λ .

Namely. $\begin{bmatrix} \lambda & \\ & \lambda \end{bmatrix}, \begin{bmatrix} \lambda & 1 \\ & \lambda \end{bmatrix}$

(b) There are 3 possible Jordan normal forms of a 3×3 matrix with eigenvalues $\lambda, \lambda, \lambda$.

Namely. $\begin{bmatrix} \lambda & & \\ & \lambda & \\ & & \lambda \end{bmatrix}, \begin{bmatrix} \lambda & & \\ & \lambda & 1 \\ & & \lambda \end{bmatrix}, \begin{bmatrix} \lambda & 1 & \\ & \lambda & 1 \\ & & \lambda \end{bmatrix}$

(c) There are 5 possible Jordan normal forms of a 4×4 matrix with eigenvalues $\lambda, \lambda, \lambda, \lambda$.

Namely. $\begin{bmatrix} \lambda & & & \\ & \lambda & & \\ & & \lambda & \\ & & & \lambda \end{bmatrix}, \begin{bmatrix} \lambda & & & \\ & \lambda & & \\ & & \lambda & 1 \\ & & & \lambda \end{bmatrix}, \begin{bmatrix} \lambda & 1 & & \\ & \lambda & & \\ & & \lambda & 1 \\ & & & \lambda \end{bmatrix}, \begin{bmatrix} \lambda & & 1 & \\ & \lambda & 1 & \\ & & \lambda & 1 \\ & & & \lambda \end{bmatrix}, \begin{bmatrix} \lambda & 1 & & \\ & \lambda & 1 & \\ & & \lambda & 1 \\ & & & \lambda \end{bmatrix}$

Example 145. What are the possible Jordan normal forms of a 6×6 matrix with eigenvalues 3, 3, 7, 7, 7, 7?

Solution. There are $2 \cdot 5 = 10$ possible Jordan normal forms for such a matrix:

$$\begin{bmatrix} 3 & & & & & \\ & 3 & & & & \\ & & 7 & & & \\ & & & 7 & & \\ & & & & 7 & \\ & & & & & 7 \end{bmatrix}, \begin{bmatrix} 3 & & & & & \\ & 3 & & & & \\ & & 7 & 1 & & \\ & & & 7 & & \\ & & & & 7 & \\ & & & & & 7 \end{bmatrix}, \begin{bmatrix} 3 & & & & & \\ & 3 & & & & \\ & & 7 & 1 & & \\ & & & 7 & & \\ & & & & 7 & 1 \\ & & & & & 7 \end{bmatrix}, \begin{bmatrix} 3 & & & & & \\ & 3 & & & & \\ & & 7 & 1 & & \\ & & & 7 & 1 & \\ & & & & 7 & \\ & & & & & 7 \end{bmatrix}, \begin{bmatrix} 3 & & & & & \\ & 3 & & & & \\ & & 7 & 1 & & \\ & & & 7 & 1 & \\ & & & & 7 & 1 \\ & & & & & 7 \end{bmatrix},$$

$$\begin{bmatrix} 3 & 1 & & & & \\ & 3 & & & & \\ & & 7 & & & \\ & & & 7 & & \\ & & & & 7 & \\ & & & & & 7 \end{bmatrix}, \begin{bmatrix} 3 & 1 & & & & \\ & 3 & & & & \\ & & 7 & 1 & & \\ & & & 7 & & \\ & & & & 7 & \\ & & & & & 7 \end{bmatrix}, \begin{bmatrix} 3 & 1 & & & & \\ & 3 & & & & \\ & & 7 & 1 & & \\ & & & 7 & & \\ & & & & 7 & 1 \\ & & & & & 7 \end{bmatrix}, \begin{bmatrix} 3 & 1 & & & & \\ & 3 & & & & \\ & & 7 & 1 & & \\ & & & 7 & 1 & \\ & & & & 7 & 1 \\ & & & & & 7 \end{bmatrix}, \begin{bmatrix} 3 & 1 & & & & \\ & 3 & & & & \\ & & 7 & 1 & & \\ & & & 7 & 1 & \\ & & & & 7 & 1 \\ & & & & & 7 \end{bmatrix}$$

Example 146. How many different Jordan normal forms are there in the following cases?

- (a) A 8×8 matrix with eigenvalues 1, 1, 2, 2, 2, 4, 4, 4?
- (b) A 11×11 matrix with eigenvalues 1, 1, 1, 2, 2, 2, 2, 4, 4, 4, 4?

Solution.

- (a) $2 \cdot 3 \cdot 3 = 18$ possible Jordan normal forms
- (b) $3 \cdot 5 \cdot 5 = 75$ possible Jordan normal forms

Review.

- Let A be $n \times n$. The matrix exponential is

$$e^A = I + A + \frac{1}{2!}A^2 + \frac{1}{3!}A^3 + \dots$$

Then, $\frac{d}{dt}e^{At} = Ae^{At}$.

Why? $\frac{d}{dt}e^{At} = \frac{d}{dt}\left(I + At + \frac{1}{2!}A^2t^2 + \frac{1}{3!}A^3t^3 + \dots\right) = A + \frac{1}{1!}A^2t + \frac{1}{2!}A^3t^2 + \dots = Ae^{At}$

- If $A = PDP^{-1}$, then $e^A = Pe^DP^{-1}$.
- The solution to $\mathbf{y}' = A\mathbf{y}$, $\mathbf{y}(0) = \mathbf{y}_0$ is $\mathbf{y}(t) = e^{At}\mathbf{y}_0$.

Why? Because $\mathbf{y}'(t) = Ae^{At}\mathbf{y}_0 = A\mathbf{y}(t)$ and $\mathbf{y}(0) = e^{0A}\mathbf{y}_0 = \mathbf{y}_0$.

Example 147. The matrix exponential shares many other properties of the usual exponential:

- $e^Ae^B = e^{A+B} = e^Be^A$ if $AB = BA$

Why the condition $AB = BA$? By the Taylor series, $e^{A+B} = I + (A+B) + \frac{(A+B)^2}{2!} + \dots$. In order to simplify that to

$$e^Ae^B = \left(I + A + \frac{A^2}{2!} + \dots\right)\left(I + B + \frac{B^2}{2!} + \dots\right),$$

we need that $(A+B)^2 = A^2 + AB + BA + B^2$ is the same as $A^2 + 2AB + B^2$. That's only the case if $AB = BA$.

- e^A is invertible and $(e^A)^{-1} = e^{-A}$

Why? That actually follows from the previous property.

Example 148. Compute e^{At} for $A = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}$.

Solution. Note that $A^2 = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}$. Hence, $e^{At} = I + At + \frac{t^2}{2!}A^2 + \dots = I + At = \begin{bmatrix} 1 & t \\ 0 & 1 \end{bmatrix}$.

Example 149. Compute e^{At} for $A = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}$.

Solution. Note that $A^2 = \begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}$ and $A^3 = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}$.

Hence, $e^{At} = I + At + \frac{t^2}{2!}A^2 + \frac{t^3}{3!}A^3 + \dots = I + At + \frac{1}{2}A^2t^2 = \begin{bmatrix} 1 & & \\ & 1 & \\ & & 1 \end{bmatrix} + \begin{bmatrix} 0 & t & 0 \\ 0 & t & 0 \\ 0 & 0 & 0 \end{bmatrix} + \frac{1}{2}\begin{bmatrix} 0 & 0 & t^2 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} = \begin{bmatrix} 1 & t & \frac{t^2}{2} \\ & 1 & t \\ & & 1 \end{bmatrix}$.

Example 150. Compute e^{At} for $A = \begin{bmatrix} 2 & 1 \\ & 2 \end{bmatrix}$.

Solution.

- Write $A = \begin{bmatrix} 2 & 1 \\ & 2 \end{bmatrix} = 2I + N$ with $N = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}$. Note that $2I$ and N commute.

Hence, $e^{At} = e^{2It+Nt} = e^{2It}e^{Nt}$.

- Note that $N^2 = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}$. Hence, $e^{Nt} = I + Nt + \frac{t^2}{2!}N^2 + \dots = I + Nt = \begin{bmatrix} 1 & t \\ 0 & 1 \end{bmatrix}$.

- Combined, $e^{At} = e^{2It+Nt} = e^{2It}e^{Nt} = \begin{bmatrix} e^{2t} & \\ & e^{2t} \end{bmatrix} \begin{bmatrix} 1 & t \\ 0 & 1 \end{bmatrix} = \begin{bmatrix} e^{2t} & te^{2t} \\ & e^{2t} \end{bmatrix}$.

Advanced. Can you show that $A^n = \begin{bmatrix} 2^n & n2^{n-1} \\ & 2^n \end{bmatrix}$?

Example 151. Solve the differential equation

$$\mathbf{y}' = \underbrace{\begin{bmatrix} 2 & 1 \\ & 2 \end{bmatrix}}_A \mathbf{y}, \quad \mathbf{y}(0) = \underbrace{\begin{bmatrix} -1 \\ 1 \end{bmatrix}}_{\mathbf{y}_0}$$

Solution. Repeating the work in the previous example, the solution to the differential equation is

$$\begin{aligned} \mathbf{y}(t) &= e^{At} \mathbf{y}_0 \\ &= e^{2It + Nt} \mathbf{y}_0 \quad \text{with } N = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} \\ &= e^{2It} e^{Nt} \mathbf{y}_0 \quad (\text{because } 2It \text{ and } Nt \text{ commute}) \\ &= \begin{bmatrix} e^{2t} & \\ & e^{2t} \end{bmatrix} \left(1 + Nt + \frac{1}{2}(Nt)^2 + \frac{1}{3!}(Nt)^3 + \dots \right) \mathbf{y}_0 \\ &= \begin{bmatrix} e^{2t} & \\ & e^{2t} \end{bmatrix} (1 + Nt) \mathbf{y}_0 \quad (\text{because } N^2 = \mathbf{0}) \\ &= \begin{bmatrix} e^{2t} & \\ & e^{2t} \end{bmatrix} \begin{bmatrix} 1 & t \\ & 1 \end{bmatrix} \begin{bmatrix} -1 \\ 1 \end{bmatrix} \\ &= \begin{bmatrix} e^{2t} & \\ & e^{2t} \end{bmatrix} \begin{bmatrix} t-1 \\ 1 \end{bmatrix} = \begin{bmatrix} (t-1)e^{2t} \\ e^{2t} \end{bmatrix}. \end{aligned}$$

Check. We should verify that $y_1 = (t-1)e^{2t}$ and $y_2 = e^{2t}$ satisfy $y_1' = 2y_1 + y_2$ and $y_2' = 2y_2$. Indeed, $y_1' = e^{2t} + (t-1)2e^{2t}$ equals $2y_1 + y_2 = 2(t-1)e^{2t} + e^{2t}$.

Comment. For applications, having solutions like $te^{\lambda t}$ or $t \cos(\lambda t)$ (when the eigenvalues are imaginary) is connected to the phenomenon of **resonance**, which you may have already seen.

Important comment. Note that we can immediately see from the solution that the original matrix A is not diagonalizable: there is a term te^{2t} , whereas in the diagonalizable case we would only see exponentials like e^{2t} by themselves.

In our upcoming discussion of complex numbers we will see that e^{2it} (here, $2i$ would be the eigenvalue) can be rewritten in terms of $\cos(2t)$ and $\sin(2t)$. Both of these are periodic and bounded, so that the same is true for every linear combination.

In that case, if the eigenvalue $2i$ was repeated in such a way that the matrix A is not diagonalizable, then we would get the functions $t \cos(2t)$ and $t \sin(2t)$ in our solutions. These, however, are not bounded! This phenomenon (getting solutions that are unbounded under the right/wrong circumstances) is called **resonance**.

<https://en.wikipedia.org/wiki/Resonance>

Understanding when resonance occurs is of crucial importance for practical applications.

Example 152. Solve the IVP $\mathbf{y}' = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \mathbf{y}$ with $\mathbf{y}(0) = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$.

Solution. Recall that the solution to $\mathbf{y}' = A\mathbf{y}$, $\mathbf{y}(0) = \mathbf{y}_0$ is $\mathbf{y} = e^{At}\mathbf{y}_0$.

- We first diagonalize $A = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$.
 - $\begin{vmatrix} -\lambda & 1 \\ 1 & -\lambda \end{vmatrix} = \lambda^2 - 1$, so the eigenvalues are ± 1 .
 - The 1-eigenspace $\text{null}\left(\begin{bmatrix} -1 & 1 \\ 1 & -1 \end{bmatrix}\right)$ has basis $\begin{bmatrix} 1 \\ 1 \end{bmatrix}$.
 - The -1-eigenspace $\text{null}\left(\begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix}\right)$ has basis $\begin{bmatrix} -1 \\ 1 \end{bmatrix}$.
 - Hence, $A = PDP^{-1}$ with $P = \begin{bmatrix} 1 & -1 \\ 1 & 1 \end{bmatrix}$ and $D = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}$.
- Compute the solution $\mathbf{y} = e^{At}\mathbf{y}_0$:

$$\begin{aligned} \mathbf{y} = e^{At}\mathbf{y}_0 &= Pe^{Dt}P^{-1}\mathbf{y}_0 \\ &= \underbrace{\begin{bmatrix} 1 & -1 \\ 1 & 1 \end{bmatrix}}_{= \begin{bmatrix} e^t & -e^{-t} \\ e^t & e^{-t} \end{bmatrix}} \begin{bmatrix} e^t & 0 \\ 0 & e^{-t} \end{bmatrix} \underbrace{\frac{1}{2}\begin{bmatrix} 1 & 1 \\ -1 & 1 \end{bmatrix}}_{= \frac{1}{2}\begin{bmatrix} 1 \\ -1 \end{bmatrix}} \begin{bmatrix} 1 \\ 0 \end{bmatrix} = \frac{1}{2} \begin{bmatrix} e^t + e^{-t} \\ e^t - e^{-t} \end{bmatrix} \end{aligned}$$

Check. Indeed, $y_1 = \frac{1}{2}(e^t + e^{-t})$ and $y_2 = \frac{1}{2}(e^t - e^{-t})$ satisfy the system of differential equations $y_1' = y_2$ and $y_2' = y_1$ as well as the initial conditions $y_1(0) = 1$, $y_2(0) = 0$.

Comment. You have actually met these functions in Calculus! $y_1 = \cosh(t)$ and $y_2 = \sinh(t)$. Check out the next example for the connection to $\cos(t)$ and $\sin(t)$.

Example 153.

- (a) Solve the IVP $\mathbf{y}' = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix} \mathbf{y}$ with $\mathbf{y}(0) = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$.
- (b) Show that $\mathbf{y} = \begin{bmatrix} \cos(t) \\ \sin(t) \end{bmatrix}$ solves the same IVP. What do you conclude?

Solution.

(a) $A = PDP^{-1}$ with $P = \begin{bmatrix} i & -i \\ 1 & 1 \end{bmatrix}$, $D = \begin{bmatrix} i & 0 \\ 0 & -i \end{bmatrix}$.

The system is therefore solved by:

$$\begin{aligned} \mathbf{y}(t) &= Pe^{Dt}P^{-1}\begin{bmatrix} 1 \\ 0 \end{bmatrix} = \begin{bmatrix} i & -i \\ 1 & 1 \end{bmatrix} \begin{bmatrix} e^{it} & \\ & e^{-it} \end{bmatrix} \frac{1}{2i} \begin{bmatrix} 1 & i \\ -1 & i \end{bmatrix} \begin{bmatrix} 1 \\ 0 \end{bmatrix} \\ &= \frac{1}{2i} \begin{bmatrix} i & -i \\ 1 & 1 \end{bmatrix} \begin{bmatrix} e^{it} & \\ & e^{-it} \end{bmatrix} \begin{bmatrix} 1 \\ -1 \end{bmatrix} = \frac{1}{2i} \begin{bmatrix} i & -i \\ 1 & 1 \end{bmatrix} \begin{bmatrix} e^{it} \\ -e^{-it} \end{bmatrix} = \frac{1}{2i} \begin{bmatrix} ie^{it} + ie^{-it} \\ e^{it} - e^{-it} \end{bmatrix} \\ &= \frac{1}{2} \begin{bmatrix} e^{it} + e^{-it} \\ -ie^{it} + ie^{-it} \end{bmatrix} \end{aligned}$$

(b) Clearly, $\mathbf{y}(0) = \begin{bmatrix} \cos(0) \\ \sin(0) \end{bmatrix} = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$. On the other hand, $y_1' = -\sin(t) = -y_2$ and $y_2' = \cos(t) = y_1$, so that

$$\mathbf{y}' = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix} \mathbf{y}. \text{ Since the solution to the IVP is unique, it follows that } \begin{bmatrix} \cos(t) \\ \sin(t) \end{bmatrix} = \frac{1}{2} \begin{bmatrix} e^{it} + e^{-it} \\ -ie^{it} + ie^{-it} \end{bmatrix}.$$

We have just discovered **Euler's identity!**

Theorem 154. (Euler's identity) $e^{i\theta} = \cos(\theta) + i \sin(\theta)$

Another short proof. Observe that both sides are the (unique) solution to the IVP $y' = iy$, $y(0) = 1$.

On lots of T-shirts. In particular, with $x = \pi$, we get $e^{i\pi} = -1$ or $e^{i\pi} + 1 = 0$ (which connects the five fundamental constants).

Rotation matrices

Example 155. Write down a 2×2 matrix Q for rotation by angle θ in the plane.

Comment. Why should we even be able to represent something like rotation by a matrix? Meaning that Qx should be the vector x rotated by θ . Recall from Linear Algebra I that every **linear map** can be represented by a matrix. Then think about why rotation is a linear map.

Solution. We can determine Q by figuring out $Q \begin{bmatrix} 1 \\ 0 \end{bmatrix}$ (the first column of Q) and $Q \begin{bmatrix} 0 \\ 1 \end{bmatrix}$ (the second column of Q).

Since $Q \begin{bmatrix} 1 \\ 0 \end{bmatrix} = \begin{bmatrix} \cos \theta \\ \sin \theta \end{bmatrix}$ and $Q \begin{bmatrix} 0 \\ 1 \end{bmatrix} = \begin{bmatrix} -\sin \theta \\ \cos \theta \end{bmatrix}$, we conclude that $Q = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix}$.

Comment. Note that we don't need previous knowledge of \cos and \sin . We could have introduced these trig functions on the spot.

Comment. Note that it is geometrically obvious that Q is orthogonal. (Why?)

It is clear that $\left\| \begin{bmatrix} \cos \theta \\ \sin \theta \end{bmatrix} \right\|^2 = 1$. Noting that $\left\| \begin{bmatrix} \cos \theta \\ \sin \theta \end{bmatrix} \right\|^2 = \cos^2 \theta + \sin^2 \theta$, we have rediscovered Pythagoras.

Advanced comment. Actually, every orthogonal 2×2 matrix Q with $\det(Q) = 1$ is a rotation by some angle θ . Orthogonal matrices with $\det(Q) = -1$ are reflections.

Example 156. As in the previous example, let Q_θ be the 2×2 matrix for rotation by angle θ in the plane. What is $Q_\alpha Q_\beta$?

Solution. Note that $Q_\alpha Q_\beta x$ first rotates x by angle β and then by angle α . For geometric reasons, it is obvious that this is the same as if we rotated x by $\alpha + \beta$. It follows that $Q_\alpha Q_\beta = Q_{\alpha + \beta}$.

Comment. This allows us to derive interesting trig identities:

$$Q_\alpha Q_\beta = \begin{bmatrix} \cos \alpha & -\sin \alpha \\ \sin \alpha & \cos \alpha \end{bmatrix} \begin{bmatrix} \cos \beta & -\sin \beta \\ \sin \beta & \cos \beta \end{bmatrix} = \begin{bmatrix} \cos \alpha \cos \beta - \sin \alpha \sin \beta & \dots \\ \dots & \dots \end{bmatrix}$$
$$Q_{\alpha + \beta} = \begin{bmatrix} \cos(\alpha + \beta) & -\sin(\alpha + \beta) \\ \sin(\alpha + \beta) & \cos(\alpha + \beta) \end{bmatrix}$$

It follows that $\cos(\alpha + \beta) = \cos \alpha \cos \beta - \sin \alpha \sin \beta$.

Comment. If we set $\beta = \alpha$, this simplifies to $\cos(2\alpha) = \cos^2 \alpha - \sin^2 \alpha = 2\cos^2 \alpha - 1$, the double angle formula that you have probably used countless times in Calculus.

Comment. Similarly, we find an identity for $\sin(\alpha + \beta)$. Spell it out!

More on complex numbers

Let's recall some very basic facts about **complex numbers**:

- Every complex number can be written as $z = x + iy$ with real x, y .

- Here, the imaginary unit i is characterized by solving $x^2 = -1$.

Important observation. The same equation is solved by $-i$. This means that, algebraically, we cannot distinguish between $+i$ and $-i$.

- The **conjugate** of $z = x + iy$ is $\bar{z} = x - iy$.

Important comment. Since we cannot algebraically distinguish between $\pm i$, we also cannot distinguish between z and \bar{z} . That's the reason why, in problems involving only real numbers, if a complex number $z = x + iy$ shows up, then its **conjugate** $\bar{z} = x - iy$ has to show up in the same manner. With that in mind, have another look at Example 89.

- The **absolute value** of the complex number $z = x + iy$ is $|z| = \sqrt{x^2 + y^2} = \sqrt{\bar{z}z}$.
- The **norm** of the complex vector $\mathbf{z} = \begin{bmatrix} z_1 \\ z_2 \end{bmatrix}$ is $\|\mathbf{z}\| = \sqrt{|z_1|^2 + |z_2|^2}$.
Note that $\|\mathbf{z}\|^2 = \bar{z}_1 z_1 + \bar{z}_2 z_2 = \bar{\mathbf{z}}^T \mathbf{z}$.

Definition 157.

- For every matrix A , its **conjugate transpose** is $A^* = (\bar{A})^T$.
- The **dot product** (inner product) of complex vectors is $\mathbf{v} \cdot \mathbf{w} = \mathbf{v}^* \mathbf{w}$.
- A complex $n \times n$ matrix A is **unitary** if $A^* A = I$.

Comment. A^* is also written A^H (or A^\dagger in quantum mechanics) and called the Hermitian conjugate.

Comment. For real matrices and vectors, the conjugate transpose is just the ordinary transpose. In particular, the dot product is the same.

Comment. Unitary matrices are the complex version of orthogonal matrices. (A real matrix is unitary if and only if it is orthogonal.)

Example 158. What is the norm of the vector $\begin{bmatrix} 1-i \\ 2+3i \end{bmatrix}$?

Solution. $\left\| \begin{bmatrix} 1-i \\ 2+3i \end{bmatrix} \right\|^2 = [1+i \ 2-3i] \begin{bmatrix} 1-i \\ 2+3i \end{bmatrix} = |1-i|^2 + |2+3i|^2 = 2 + 13$. Hence, $\left\| \begin{bmatrix} 1-i \\ 2+3i \end{bmatrix} \right\| = \sqrt{15}$.

Example 159. Determine A^* if $A = \begin{bmatrix} 2 & 1-i \\ 3+2i & i \end{bmatrix}$.

Solution. $A^* = \begin{bmatrix} 2 & 3-2i \\ 1+i & -i \end{bmatrix}$

Example 160. What is $\frac{1}{2+3i}$?

Solution. $\frac{1}{2+3i} = \frac{2-3i}{(2+3i)(2-3i)} = \frac{2-3i}{13}$.

In general. $\frac{1}{z} = \frac{\bar{z}}{z\bar{z}} = \frac{\bar{z}}{|z|^2}$

Example 161. (extra) We can identify complex numbers $x + iy$ with vectors $\begin{bmatrix} x \\ y \end{bmatrix}$ in \mathbb{R}^2 . Then, what is the geometric effect of multiplying with i ?

Solution. Algebraically, the effect of multiplying $x + iy$ with i obviously is $i(x + iy) = -y + ix$.

Since multiplication with i is obviously linear, we can represent it using a 2×2 matrix J acting on vectors $\begin{bmatrix} x \\ y \end{bmatrix}$.

$J \begin{bmatrix} 1 \\ 0 \end{bmatrix} = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$ (this is the same as saying $i \cdot 1 = i$) and $J \begin{bmatrix} 0 \\ 1 \end{bmatrix} = \begin{bmatrix} -1 \\ 0 \end{bmatrix}$ (this is the same as saying $i \cdot i = -1$).

Hence, $J = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}$. This is precisely the rotation matrix for a rotation by 90° .

In other words, multiplication with i has the geometric effect of rotating complex numbers by 90° .

Comment. The relation $i^2 = -1$ translates to $J^2 = -I$.

Complex numbers as 2×2 matrices. In light of the above, we can express complex numbers $x + iy$ as the 2×2 matrix $xI + yJ = \begin{bmatrix} x & -y \\ y & x \end{bmatrix}$. Adding and multiplying these matrices behaves exactly the same way as adding or multiplying the complex numbers directly.

For instance, $(2 + 3i)(4 - i) = 8 + 10i - 3i^2 = 11 + 10i$ versus $\begin{bmatrix} 2 & -3 \\ 3 & 2 \end{bmatrix} \begin{bmatrix} 4 & 1 \\ -1 & 4 \end{bmatrix} = \begin{bmatrix} 11 & -10 \\ 10 & 11 \end{bmatrix}$.

Likewise for inverses: $\frac{1}{2 + 3i} = \frac{2 - 3i}{(2 + 3i)(2 - 3i)} = \frac{2 - 3i}{13}$ versus $\begin{bmatrix} 2 & -3 \\ 3 & 2 \end{bmatrix}^{-1} = \frac{1}{13} \begin{bmatrix} 2 & 3 \\ -3 & 2 \end{bmatrix}$

Extra: More details on the spectral theorem

Let us add $\langle \mathbf{v}, \mathbf{w} \rangle$ to our notations for the dot product: $\langle \mathbf{v}, \mathbf{w} \rangle = \mathbf{v}^T \mathbf{w} = \mathbf{v} \cdot \mathbf{w}$.

- In our story of orthogonality, the important player has been the dot product. However, one could argue that the fundamental quantity is actually the norm:

$$\langle \mathbf{v}, \mathbf{w} \rangle = \frac{1}{4}(\|\mathbf{v} + \mathbf{w}\|^2 - \|\mathbf{v} - \mathbf{w}\|^2). \text{ See Example 28.}$$

- Accepting the dot product as immensely important, we see that symmetric matrices (i.e. matrices A such that $A = A^T$) are of interest.

For every matrix A , $\langle A\mathbf{v}, \mathbf{w} \rangle = \langle \mathbf{v}, A^T\mathbf{w} \rangle$.

It follows that, a matrix A is symmetric if and only if $\langle A\mathbf{v}, \mathbf{w} \rangle = \langle \mathbf{v}, A\mathbf{w} \rangle$ for all vectors \mathbf{v}, \mathbf{w} .

- Similarly, let Q be an orthogonal matrix (i.e. Q is a square matrix with $Q^T Q = I$).

Then, $\langle Q\mathbf{v}, Q\mathbf{w} \rangle = \langle \mathbf{v}, \mathbf{w} \rangle$.

In fact, a matrix A is orthogonal if and only if $\langle A\mathbf{v}, A\mathbf{w} \rangle = \langle \mathbf{v}, \mathbf{w} \rangle$ for all vectors \mathbf{v}, \mathbf{w} .

Comment. We observed in Example 155 that orthogonal matrices Q correspond to rotations ($\det Q = 1$) or reflections ($\det Q = -1$) [or products thereof]. The equality $\langle Q\mathbf{v}, Q\mathbf{w} \rangle = \langle \mathbf{v}, \mathbf{w} \rangle$ encodes the fact that these types (and only these!) of geometric transformations preserve angles and lengths.

(spectral theorem)

A $n \times n$ matrix A is symmetric if and only if it can be decomposed as $A = PDP^T$, where

- D is a diagonal matrix, $(n \times n)$

The diagonal entries λ_i are the **eigenvalues** of A .

- P is orthogonal. $(n \times n)$

The columns of P are **eigenvectors** of A .

Note that, in particular, A is always diagonalizable, the eigenvalues (and hence, the eigenvectors) are all real, and, most importantly, the eigenspaces of A are orthogonal.

The “only if” part says that, if A is symmetric, then we get a diagonalization $A = PDP^T$. The “if” part says that, if $A = PDP^T$, then A is symmetric (which follows from $A^T = (PDP^T)^T = (P^T)^T D^T P^T = PDP^T = A$).

Let us prove the following important parts of the spectral theorem.

We already proved the first part in Theorem 94 using the same argument and only slightly different notation.

Theorem 162.

- (a) If A is symmetric, then the eigenspaces of A are orthogonal.
- (b) If A is real and symmetric, then the eigenvalues of A are real.

Proof.

- (a) We need to show that, if \mathbf{v} and \mathbf{w} are eigenvectors of A with different eigenvalues, then $\langle \mathbf{v}, \mathbf{w} \rangle = 0$.

Suppose that $A\mathbf{v} = \lambda\mathbf{v}$ and $A\mathbf{w} = \mu\mathbf{w}$ with $\lambda \neq \mu$.

Then, $\lambda\langle \mathbf{v}, \mathbf{w} \rangle = \langle \lambda\mathbf{v}, \mathbf{w} \rangle = \langle A\mathbf{v}, \mathbf{w} \rangle = \langle \mathbf{v}, A^T\mathbf{w} \rangle = \langle \mathbf{v}, A\mathbf{w} \rangle = \langle \mathbf{v}, \mu\mathbf{w} \rangle = \mu\langle \mathbf{v}, \mathbf{w} \rangle$.

However, since $\lambda \neq \mu$, $\lambda\langle \mathbf{v}, \mathbf{w} \rangle = \mu\langle \mathbf{v}, \mathbf{w} \rangle$ is only possible if $\langle \mathbf{v}, \mathbf{w} \rangle = 0$.

- (b) Suppose λ is a nonreal eigenvalue with nonzero eigenvector \mathbf{v} . Then, $\bar{\mathbf{v}}$ is a $\bar{\lambda}$ -eigenvector and, since $\lambda \neq \bar{\lambda}$, we have two eigenvectors with different eigenvalues. By the first part, these two eigenvectors must be orthogonal in the sense that $\bar{\mathbf{v}}^T\mathbf{v} = 0$. But $\bar{\mathbf{v}}^T\mathbf{v} = \mathbf{v}^*\mathbf{v} = \|\mathbf{v}\|^2 \neq 0$. This shows that it is impossible to have a nonzero eigenvector for a nonreal eigenvalue.

Alternative proof. Note that a complex number λ is real if and only if $\bar{\lambda} = \lambda$. Suppose that λ is an eigenvalue with nonzero eigenvector \mathbf{v} so that $A\mathbf{v} = \lambda\mathbf{v}$. We now observe that $\lambda\mathbf{v}^*\mathbf{v} = \mathbf{v}^*(\lambda\mathbf{v}) = \mathbf{v}^*A\mathbf{v} = \mathbf{v}^*A^*\mathbf{v} = (A\mathbf{v})^*\mathbf{v} = (\lambda\mathbf{v})^*\mathbf{v} = \bar{\lambda}\mathbf{v}^*\mathbf{v}$. Dividing by $\|\mathbf{v}\|^2 = \mathbf{v}^*\mathbf{v}$ (which is not zero!) we find $\lambda = \bar{\lambda}$ from which we conclude that λ is real. \square

Advanced comment. Note that the alternative proof of the second part shows that any Hermitian matrix A (that is, a complex matrix A such that $A^* = A$) has only real eigenvalues. If A is Hermitian, what can we conclude about the eigenspaces if we follow the argument in the first part?

Let us highlight the following point we used in our proof:

Let A be a real matrix. If \mathbf{v} is a λ -eigenvector, then $\bar{\mathbf{v}}$ is a $\bar{\lambda}$ -eigenvector.

See, for instance, Example 89. This is just a consequence of the basic fact that we cannot algebraically distinguish between $+i$ and $-i$.

Singular value decomposition

(Singular value decomposition)

Every $m \times n$ matrix A can be decomposed as $A = U\Sigma V^T$, where

- Σ is a (rectangular) diagonal matrix with nonnegative entries, $(m \times n)$
The diagonal entries σ_i are called the **singular values** of A .
- U is orthogonal, $(m \times m)$
- V is orthogonal. $(n \times n)$

Comment. If A is symmetric, then the singular value decomposition is already provided by the spectral theorem (the diagonalization of A). Moreover, in that case, $V = U$.

Important observations. If $A = U\Sigma V^T$, then $A^T A = V\Sigma^T \Sigma V^T$.

- Note that $\Sigma^T \Sigma$ is an $n \times n$ diagonal matrix. Its entries are σ_i^2 (the squares of the entries in Σ).
- $A^T A$ is a symmetric matrix! (Why?!) Hence, by the spectral theorem, we are able to find V and $\Sigma^T \Sigma$.

In other words, V is obtained from the (orthonormally chosen) eigenvectors of $A^T A$. Likewise, the entries of $\Sigma^T \Sigma$ are the eigenvalues of $A^T A$; their square roots are the entries of Σ , the singular values.

Finally, the equation $AV = U\Sigma$ allows us to determine U . How?! (Hint: $Av_i = \sigma_i u_i$)

This results in the following **recipe** to determine the SVD $A = U\Sigma V^T$ for any matrix A .

Find an orthonormal basis of eigenvectors v_i of $A^T A$. Let λ_i be the eigenvalue of v_i .

- V is the matrix with columns v_i .
- Σ is the diagonal matrix with entries $\sigma_i = \sqrt{\lambda_i}$.
- U is the matrix with columns $u_i = \frac{1}{\sigma_i} Av_i$. If needed, fill in additional columns to make U orthogonal.

Example 163. Determine the SVD of $A = \begin{bmatrix} 2 & 2 \\ -1 & 1 \end{bmatrix}$.

Solution. $A^T A = \begin{bmatrix} 5 & 3 \\ 3 & 5 \end{bmatrix}$ has 8-eigenvector $\begin{bmatrix} 1 \\ 1 \end{bmatrix}$ and 2-eigenvector $\begin{bmatrix} -1 \\ 1 \end{bmatrix}$.

Since $A^T A = V\Sigma^2 V^T$ (here, $\Sigma^T \Sigma = \Sigma^2$), we conclude that $V = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & -1 \\ 1 & 1 \end{bmatrix}$ and $\Sigma = \begin{bmatrix} \sqrt{8} & \\ & \sqrt{2} \end{bmatrix}$.

From $Av_i = \sigma_i u_i$, we find $u_1 = \frac{1}{\sigma_1} Av_1 = \frac{1}{\sqrt{8}} \begin{bmatrix} 2 & 2 \\ -1 & 1 \end{bmatrix} \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ 1 \end{bmatrix} = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$.

Likewise, $u_2 = \frac{1}{\sigma_2} Av_2 = \frac{1}{\sqrt{2}} \begin{bmatrix} 2 & 2 \\ -1 & 1 \end{bmatrix} \frac{1}{\sqrt{2}} \begin{bmatrix} -1 \\ 1 \end{bmatrix} = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$. Hence, $U = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$. Check that, indeed, $A = U\Sigma V^T$!

Comment. For applications, it is common to arrange the singular values in decreasing order like we did.

Comment. If we had chosen $V = \frac{1}{\sqrt{2}} \begin{bmatrix} -1 & -1 \\ -1 & 1 \end{bmatrix}$ instead, then $U = \begin{bmatrix} -1 & 0 \\ 0 & 1 \end{bmatrix}$ and $\Sigma = \begin{bmatrix} \sqrt{8} & \\ & \sqrt{2} \end{bmatrix}$.

As with diagonalization, there are choices! (A lot fewer choices though.) This is another perfectly fine SVD. In fact, it's what Sage computes below.

Sage. Let's have Sage do the work for us. In Sage, the SVD is currently only implemented for floating point numbers. (RDF is the real numbers as floating point numbers with double precision)

```
Sage] A = matrix(RDF, [[2,2],[-1,1]])
```

```
Sage] U,S,V = A.SVD()
```

```
Sage] U
```

$$\begin{bmatrix} -1.0 & 1.11022302463 \times 10^{-16} \\ 8.64109131471 \times 10^{-17} & 1.0 \end{bmatrix}$$


```
Sage] S
```

$$\begin{bmatrix} 2.82842712475 & 0.0 \\ 0.0 & 1.41421356237 \end{bmatrix}$$

```
Sage] V
```

$$\begin{bmatrix} -0.707106781187 & -0.707106781187 \\ -0.707106781187 & 0.707106781187 \end{bmatrix}$$

Remark 164. (April Fools' Day!) π is the perimeter of a circle enclosed in a square with edge length 1. The perimeter of the square is 4, which approximates π . To get a better approximation, we "fold" the vertices of the square towards the circle (and get the blue polygon). This construction can be repeated for even better approximations and, in the limit, our shape will converge to the true circle. At each step, the perimeter is 4, so we conclude that $\pi = 4$, contrary to popular belief.



Can you pin-point the fallacy in this argument?

Comment. We'll actually come back to this. It's related to linear algebra in infinite dimensions.

Review. SVD

Example 165. Determine the SVD of $A = \begin{bmatrix} 2 & 2 \\ 1 & 1 \end{bmatrix}$.

Comment. In contrast to our previous example, $\text{rank}(A) = 1$. It follows that $A^T A$ has eigenvalue 0, so that 0 is a singular value of A .

Solution. $A^T A = \begin{bmatrix} 5 & 5 \\ 5 & 5 \end{bmatrix}$ has 10-eigenvector $\begin{bmatrix} 1 \\ 1 \end{bmatrix}$ and 0-eigenvector $\begin{bmatrix} -1 \\ 1 \end{bmatrix}$.

We conclude that $V = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & -1 \\ 1 & 1 \end{bmatrix}$ and $\Sigma = \begin{bmatrix} \sqrt{10} & \\ & 0 \end{bmatrix}$.

$$\mathbf{u}_1 = \frac{1}{\sigma_1} A \mathbf{v}_1 = \frac{1}{\sqrt{10}} \begin{bmatrix} 2 & 2 \\ 1 & 1 \end{bmatrix} \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ 1 \end{bmatrix} = \frac{1}{\sqrt{20}} \begin{bmatrix} 4 \\ 2 \end{bmatrix} = \frac{1}{\sqrt{5}} \begin{bmatrix} 2 \\ 1 \end{bmatrix}$$

We cannot obtain \mathbf{u}_2 in the same way because $\sigma_2 = 0$. Since for every vector \mathbf{u}_2 , $A \mathbf{v}_2 = \sigma_2 \mathbf{u}_2$, we can choose \mathbf{u}_2 as we wish, as long as the columns of U are orthonormal in the end.

$$\mathbf{u}_2 = \frac{1}{\sqrt{5}} \begin{bmatrix} -1 \\ 2 \end{bmatrix} \text{ (but } \mathbf{u}_2 = \frac{1}{\sqrt{5}} \begin{bmatrix} 1 \\ -2 \end{bmatrix} \text{ works just as well)}$$

$$\text{Hence, } U = \frac{1}{\sqrt{5}} \begin{bmatrix} 2 & -1 \\ 1 & 2 \end{bmatrix}.$$

$$\text{In summary, } A = U \Sigma V^T \text{ with } U = \frac{1}{\sqrt{5}} \begin{bmatrix} 2 & -1 \\ 1 & 2 \end{bmatrix}, \Sigma = \begin{bmatrix} \sqrt{10} & \\ & 0 \end{bmatrix}, V = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & -1 \\ 1 & 1 \end{bmatrix}.$$

Check. Do check that, indeed, $A = U \Sigma V^T$.

Example 166. Determine the SVD of $A = \begin{bmatrix} 1 & -1 \\ 0 & 1 \\ 1 & 0 \end{bmatrix}$.

Solution. $A^T A = \begin{bmatrix} 2 & -1 \\ -1 & 2 \end{bmatrix}$ has 3-eigenvector $\begin{bmatrix} -1 \\ 1 \end{bmatrix}$ and 1-eigenvector $\begin{bmatrix} 1 \\ 1 \end{bmatrix}$.

Since $A^T A = V \Sigma^T \Sigma V^T$, we conclude that $V = \frac{1}{\sqrt{2}} \begin{bmatrix} -1 & 1 \\ 1 & 1 \end{bmatrix}$ and $\Sigma = \begin{bmatrix} \sqrt{3} & 0 \\ 0 & 1 \\ 0 & 0 \end{bmatrix}$.

$$\mathbf{u}_1 = \frac{1}{\sigma_1} A \mathbf{v}_1 = \frac{1}{\sqrt{3}} \begin{bmatrix} 1 & -1 \\ 0 & 1 \\ 1 & 0 \end{bmatrix} \frac{1}{\sqrt{2}} \begin{bmatrix} -1 \\ 1 \end{bmatrix} = \frac{1}{\sqrt{6}} \begin{bmatrix} -2 \\ 1 \\ -1 \end{bmatrix}$$

$$\mathbf{u}_2 = \frac{1}{\sigma_2} A \mathbf{v}_2 = \frac{1}{1} \begin{bmatrix} 1 & -1 \\ 0 & 1 \\ 1 & 0 \end{bmatrix} \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ 1 \end{bmatrix} = \frac{1}{\sqrt{2}} \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix}$$

\mathbf{u}_3 is chosen so that the matrix U is orthogonal. Hence, $\mathbf{u}_3 = \frac{1}{\sqrt{3}} \begin{bmatrix} -1 \\ -1 \\ 1 \end{bmatrix}$ (or $\mathbf{u}_3 = \frac{1}{\sqrt{3}} \begin{bmatrix} 1 \\ 1 \\ -1 \end{bmatrix}$).

$$\text{Hence, } U = \begin{bmatrix} -2/\sqrt{6} & 0 & -1/\sqrt{3} \\ 1/\sqrt{6} & 1/\sqrt{2} & -1/\sqrt{3} \\ -1/\sqrt{6} & 1/\sqrt{2} & 1/\sqrt{3} \end{bmatrix}.$$

$$\text{In summary, } A = U \Sigma V^T \text{ with } U = \begin{bmatrix} -2/\sqrt{6} & 0 & -1/\sqrt{3} \\ 1/\sqrt{6} & 1/\sqrt{2} & -1/\sqrt{3} \\ -1/\sqrt{6} & 1/\sqrt{2} & 1/\sqrt{3} \end{bmatrix}, \Sigma = \begin{bmatrix} \sqrt{3} & 0 \\ 0 & 1 \\ 0 & 0 \end{bmatrix}, V = \frac{1}{\sqrt{2}} \begin{bmatrix} -1 & 1 \\ 1 & 1 \end{bmatrix}.$$

How did we find \mathbf{u}_3 ? We already have the vectors \mathbf{u}_1 and \mathbf{u}_2 , and need a vector orthogonal to both.

That is, we need to find the vector spanning $\text{span} \left\{ \begin{bmatrix} -2 \\ 1 \\ -1 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix} \right\}^\perp = \text{col} \left(\begin{bmatrix} -2 & 0 \\ 1 & 1 \\ -1 & 1 \end{bmatrix} \right)^\perp = \text{null} \left(\begin{bmatrix} -2 & 1 & -1 \\ 0 & 1 & 1 \end{bmatrix} \right)$.

[Without the intermediate steps, can you see why the null space consists of precisely the vectors orthogonal to both \mathbf{u}_1 and \mathbf{u}_2 ?]

More generally, proceeding like this, we can always fill in “missing” vectors \mathbf{u}_i to obtain an orthonormal basis $\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_m$ that we can use as the columns of U .

Example 167. Determine the SVD of $A = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}$.

Solution. $A^T A = \begin{bmatrix} 1 & 1 \\ 1 & 2 \end{bmatrix}$ has characteristic polynomial $(1 - \lambda)(2 - \lambda) - 1 = \lambda^2 - 3\lambda + 1$.

The eigenvalues of $A^T A$ are $\lambda_{1,2} = \frac{3 \pm \sqrt{5}}{2}$.

$\frac{3 + \sqrt{5}}{2}$ -eigenvector $\begin{bmatrix} 2 \\ 1 + \sqrt{5} \end{bmatrix}$ and $\frac{3 - \sqrt{5}}{2}$ -eigenvector $\begin{bmatrix} 2 \\ 1 - \sqrt{5} \end{bmatrix}$.

It would be rather painful to continue with exact expressions, and that is not how applications typically proceed. Numerically:

- 2.618-eigenvector $\begin{bmatrix} 0.526 \\ 0.851 \end{bmatrix}$ and 0.382-eigenvector $\begin{bmatrix} -0.851 \\ 0.526 \end{bmatrix}$. These eigenvectors are normalized, and it is now actually immediately obvious that they are orthogonal. (Of course, they had to be!)

- Hence, $\Sigma = \begin{bmatrix} \sqrt{2.618} & \\ & \sqrt{0.382} \end{bmatrix} = \begin{bmatrix} 1.618 & \\ & 0.618 \end{bmatrix}$ and $V = \begin{bmatrix} 0.526 & -0.851 \\ 0.851 & 0.526 \end{bmatrix}$.

[We chose $\begin{bmatrix} -0.851 \\ 0.526 \end{bmatrix}$ instead of $\begin{bmatrix} 0.851 \\ -0.526 \end{bmatrix}$, so that, for the resulting V , $\det V = +1$.]

- $\mathbf{u}_1 = \frac{1}{\sigma_1} A \mathbf{v}_1 = \frac{1}{1.618} \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 0.526 \\ 0.851 \end{bmatrix} = \begin{bmatrix} 0.851 \\ 0.526 \end{bmatrix}$
 $\mathbf{u}_2 = \frac{1}{\sigma_2} A \mathbf{v}_2 = \frac{1}{0.618} \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} -0.851 \\ 0.526 \end{bmatrix} = \begin{bmatrix} -0.526 \\ 0.851 \end{bmatrix}$.

Hence, $U = \begin{bmatrix} 0.851 & -0.526 \\ 0.526 & 0.851 \end{bmatrix}$. (Again, notice the obvious orthogonality!)

Comment. The matrix A itself has eigenvalues 1, 1, but the 1-eigenspace is only 1-dimensional. We are missing an eigenvector, which renders A not diagonalizable.

Comment. If we had continued symbolically, there are some magical simplifications like $\sqrt{\frac{3 + \sqrt{5}}{2}} = \frac{1 + \sqrt{5}}{2}$ going on. By the way, this is the golden ratio!

Sage. In Sage, the SVD is currently only implemented for floating point numbers (RDF is the real numbers as floating point numbers with double precision). Here's our computation:

```
Sage] A = matrix(RDF, [[1,1],[0,1]])
```

```
Sage] U,S,V = A.SVD()
```

```
Sage] U
```

$$\begin{bmatrix} 0.850650808352 & -0.525731112119 \\ 0.525731112119 & 0.850650808352 \end{bmatrix}$$

```
Sage] S
```

$$\begin{bmatrix} 1.61803398875 & 0.0 \\ 0.0 & 0.61803398875 \end{bmatrix}$$

```
Sage] V
```

$$\begin{bmatrix} 0.525731112119 & -0.850650808352 \\ 0.850650808352 & 0.525731112119 \end{bmatrix}$$

Example 168. (continued) The matrices U and V are rotation matrices. By what angle?

Why rotations? Recall that orthogonal matrices have determinant $+1$ or -1 .

Since $\det U = +1$ and $\det V = +1$, the orthogonal matrices U, V are rotations.

Solution. Being rotation matrices, each of them equals $\begin{bmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{bmatrix}$ for some angle θ .

To find the angle θ_V for V , we compute $\arccos(0.526) = 1.017$. This means that $\theta_V = 1.017$ or $\theta_V = 2\pi - 1.017$ (make a sketch of $\cos(\theta)$ if that's unclear!). Since $\sin(1.017) = 0.851$ (whereas $\sin(2\pi - 1.017) = -0.851$), we conclude that V is a rotation by $\theta_V = 1.017 = 58.3^\circ$. Keep that angle in mind for the next example!

Likewise, U is a rotation by $\theta_U = 0.554 = 31.7^\circ$.

Comment. The two angles add up to 90° . That's a consequence of the (atypical) fact that the matrices U and V have essentially the same entries.

Example 169. Explain the geometric meaning of the SVD in the previous example.

- The map $\mathbf{x} \mapsto A\mathbf{x}$ with $A = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}$ sends the (orthogonal) grid spanned by $\begin{bmatrix} 1 \\ 0 \end{bmatrix}$ and $\begin{bmatrix} 0 \\ 1 \end{bmatrix}$ to the (nonorthogonal) grid spanned by $A\begin{bmatrix} 1 \\ 0 \end{bmatrix} = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$ and $A\begin{bmatrix} 0 \\ 1 \end{bmatrix} = \begin{bmatrix} 1 \\ 1 \end{bmatrix}$.

Make a sketch! The two grids are overlaid in the first plot on the next page.

- Likewise, for instance, the (orthogonal) grid spanned by $\frac{1}{\sqrt{2}}\begin{bmatrix} 1 \\ 1 \end{bmatrix}$ and $\frac{1}{\sqrt{2}}\begin{bmatrix} -1 \\ 1 \end{bmatrix}$ (the 45° degree rotated version of the previous grid) is sent to the (again, nonorthogonal) grid spanned by $\frac{1}{\sqrt{2}}\begin{bmatrix} 2 \\ 1 \end{bmatrix}$ and $\frac{1}{\sqrt{2}}\begin{bmatrix} 0 \\ 1 \end{bmatrix}$. Make a sketch! The two grids are overlaid in the second plot on the next page.

- Can we find an orthogonal grid which is sent to another orthogonal grid by A ?

Solution. Yes! The SVD $A = U\Sigma V^T$ is equivalent to $AV = U\Sigma$. That is, $A\mathbf{v}_i = \sigma_i\mathbf{u}_i$.

In other words, the orthogonal grid spanned by $\mathbf{v}_1, \mathbf{v}_2$ is sent to the orthogonal grid spanned by $\sigma_1\mathbf{u}_1, \sigma_2\mathbf{u}_2$. As we observed earlier, the grid spanned by $\mathbf{v}_1, \mathbf{v}_2$ is the 58.3° degree rotated version of the standard grid)

While the input grid consists of little squares, the output grid consists of rectangles with sides σ_1, σ_2 .

Make a sketch! The two grids are overlaid in the third plot on the next page.


The following Sage code prepares the plots on the next page. Even if you have no coding background, can you see, roughly, what is happening?

```
Sage] def grid_lines(v1, v2, n, args={}):
    lines = Graphics()
    for i in [-n..n]:
        lines += line([i*v1-n*v2, i*v1+n*v2], color='red', **args)
        lines += line([i*v2-n*v1, i*v2+n*v1], color='blue', **args)
    return lines
```

```
Sage] def svd_rotate(angle):
    A = matrix([[1,1],[0,1]])
    t = angle*2*pi/360
    R = matrix([[cos(t),-sin(t)],[sin(t),cos(t)]])
    G1 = grid_lines(R*vector([1,0]), R*vector([0,1]), 2, {'linestyle':'-'})
    G2 = grid_lines(A*R*vector([1,0]), A*R*vector([0,1]), 2, {'linestyle':'--'})
    B = polygon([(-5,-5), (-5,5), (5,5), (5,-5)], fill=False)
    O = point((0,0), pointsize=30,color='black')
    return B+O+G1+G2
```


Grid spanned by $\begin{bmatrix} 1 \\ 0 \end{bmatrix}$ and $\begin{bmatrix} 0 \\ 1 \end{bmatrix}$ (dotted), and grid spanned by $\begin{bmatrix} 1 \\ 0 \end{bmatrix}$ and $\begin{bmatrix} 1 \\ 1 \end{bmatrix}$ (dashed):

Sage] `svd_rotate(angle = 0)`



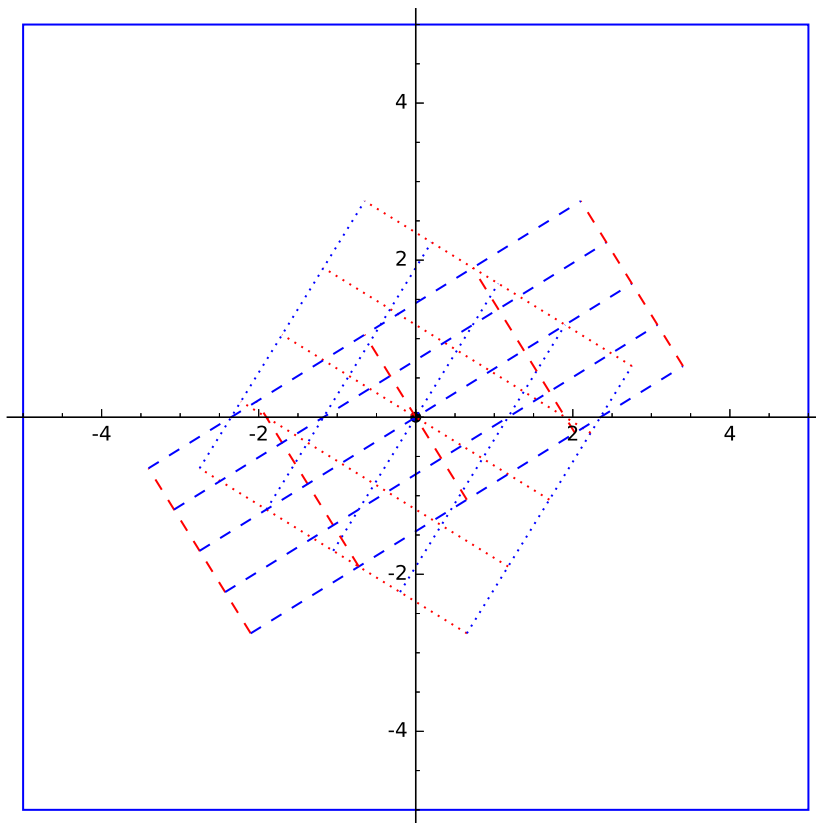
Grid spanned by $\frac{1}{\sqrt{2}}\begin{bmatrix} 1 \\ 1 \end{bmatrix}$ and $\frac{1}{\sqrt{2}}\begin{bmatrix} -1 \\ 1 \end{bmatrix}$ (dotted), and grid spanned by $\frac{1}{\sqrt{2}}\begin{bmatrix} 1 \\ 2 \end{bmatrix}$ and $\frac{1}{\sqrt{2}}\begin{bmatrix} 0 \\ 1 \end{bmatrix}$ (dashed):

Sage] `svd_rotate(angle = 45)`



Finally, here is the special situation (given by the SVD!) which shows an orthogonal grid (rotated by 58.3° degree) that is sent to another orthogonal grid (rotated by 31.7° degree):

Sage] `svd_rotate(angle = 58.3)`



For more pictures and detailed comments see the beautiful article:
<http://www.ams.org/samplings/feature-column/fcarc-svd>

Example 170. Show that the eigenvalues of $A^T A$ are all nonnegative.

Proof. Suppose that λ is an eigenvalue of $A^T A$. Then $A^T A \mathbf{v} = \lambda \mathbf{v}$ (where \mathbf{v} is a λ -eigenvector).

It follows that $\frac{\mathbf{v}^T A^T A \mathbf{v}}{\|\mathbf{A}\mathbf{v}\|^2} = \lambda \mathbf{v}^T \mathbf{v} = \lambda \|\mathbf{v}\|^2$. Finally, $\lambda \|\mathbf{v}\|^2 \geq 0$ implies that $\lambda \geq 0$. □

The **pseudoinverse** of an $m \times n$ matrix A is the matrix A^+ such that the system $A\mathbf{x} = \mathbf{b}$ has “optimal” solution $\mathbf{x} = A^+\mathbf{b}$.

Here, “optimal” means that \mathbf{x} is the smallest least squares solution.

In particular:

- If $A\mathbf{x} = \mathbf{b}$ has a unique solution, then $\mathbf{x} = A^+\mathbf{b}$ is that solution.
- If $A\mathbf{x} = \mathbf{b}$ has many solutions, then $\mathbf{x} = A^+\mathbf{b}$ is the one of smallest norm (the “optimal” one; and there is indeed only one such optimal solution).
- If $A\mathbf{x} = \mathbf{b}$ is inconsistent but has a unique least squares solution, then $\mathbf{x} = A^+\mathbf{b}$ is that least squares solution.
- If $A\mathbf{x} = \mathbf{b}$ has many least squares solutions, then $\mathbf{x} = A^+\mathbf{b}$ is the one with smallest norm.

When there is a unique (least squares) solution, we know how to find the pseudoinverse:

- If A is invertible, then $A^+ = A^{-1}$.
- If A has full column rank, then $A^+ = (A^T A)^{-1} A^T$.

Recall. If $A\mathbf{x} = \mathbf{b}$ is inconsistent, a least squares solution can be determined by solving $A^T A \mathbf{x} = A^T \mathbf{b}$. If A has full column rank (i.e. the columns of A are independent; in this context, the typical case), then $\mathbf{x} = (A^T A)^{-1} A^T \mathbf{b}$ is the **unique** least squares solution to $A\mathbf{x} = \mathbf{b}$.

Example 171.

- (a) What is the pseudoinverse of $\Sigma = \begin{bmatrix} 2 & 0 \\ 0 & 3 \\ 0 & 0 \end{bmatrix}$?
- (b) What is the pseudoinverse of $\Sigma = \begin{bmatrix} 2 & 0 & 0 \\ 0 & 3 & 0 \end{bmatrix}$?
- (c) What is the pseudoinverse of $\Sigma = \begin{bmatrix} 2 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}$?
- (d) In each case, compute $\Sigma^+ \Sigma$ and $\Sigma \Sigma^+$.

Solution.

- (a) Recall that, if A has full column rank, then $A^+ = (A^T A)^{-1} A^T$.

Here, $\Sigma^T \Sigma = \begin{bmatrix} 4 & 0 \\ 0 & 9 \end{bmatrix}$, so that $\Sigma^+ = (\Sigma^T \Sigma)^{-1} \Sigma^T = \begin{bmatrix} 1/4 & \\ & 1/9 \end{bmatrix} \begin{bmatrix} 2 & 0 & 0 \\ 0 & 3 & 0 \end{bmatrix} = \begin{bmatrix} 1/2 & 0 & 0 \\ 0 & 1/3 & 0 \end{bmatrix}$.

Alternative. Let us think about the optimal solution to $\Sigma \mathbf{x} = \mathbf{b}$, that is, $\begin{bmatrix} 2 & 0 \\ 0 & 3 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} b_1 \\ b_2 \\ b_3 \end{bmatrix}$.

The (unique) least squares solution is $\mathbf{x} = \begin{bmatrix} b_1/2 \\ b_2/3 \end{bmatrix}$. (Review if this is not obvious!)

Since $\begin{bmatrix} b_1/2 \\ b_2/3 \end{bmatrix} = \begin{bmatrix} 1/2 & 0 & 0 \\ 0 & 1/3 & 0 \end{bmatrix} \mathbf{b}$, we conclude that $\Sigma^+ = \begin{bmatrix} 1/2 & 0 & 0 \\ 0 & 1/3 & 0 \end{bmatrix}$.

- (b) Let us think about the smallest norm (“optimal”) solution to $\Sigma \mathbf{x} = \mathbf{b}$, that is, $\begin{bmatrix} 2 & 0 & 0 \\ 0 & 3 & 0 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} b_1 \\ b_2 \end{bmatrix}$.

The general solution is $\mathbf{x} = \begin{bmatrix} b_1/2 \\ b_2/3 \\ t \end{bmatrix}$, where t is a free parameter.

Clearly, the smallest norm solution is $\begin{bmatrix} b_1/2 \\ b_2/3 \\ 0 \end{bmatrix}$.

Since $\begin{bmatrix} b_1/2 \\ b_2/3 \\ 0 \end{bmatrix} = \begin{bmatrix} 1/2 & 0 \\ 0 & 1/3 \\ 0 & 0 \end{bmatrix} \mathbf{b}$, we conclude that $\Sigma^+ = \begin{bmatrix} 1/2 & 0 \\ 0 & 1/3 \\ 0 & 0 \end{bmatrix}$.

- (c) Now, $\Sigma \mathbf{x} = \mathbf{b}$, that is, $\begin{bmatrix} 2 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} b_1 \\ b_2 \end{bmatrix}$ has no solution (unless $b_2 = 0$).

We therefore need to think about least squares solutions.

The general least squares solution (why?!) is $\mathbf{x} = \begin{bmatrix} b_1/2 \\ s \\ t \end{bmatrix}$, where s, t are free parameters.

Clearly, the smallest norm least squares solution is $\begin{bmatrix} b_1/2 \\ 0 \\ 0 \end{bmatrix}$.

Since $\begin{bmatrix} b_1/2 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} 1/2 & 0 \\ 0 & 0 \\ 0 & 0 \end{bmatrix} \mathbf{b}$, we conclude that $\Sigma^+ = \begin{bmatrix} 1/2 & 0 \\ 0 & 0 \\ 0 & 0 \end{bmatrix}$.

- (d) Firstly, $\Sigma^+ \Sigma = \begin{bmatrix} 1/2 & 0 & 0 \\ 0 & 1/3 & 0 \end{bmatrix} \begin{bmatrix} 2 & 0 \\ 0 & 3 \\ 0 & 0 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$ and $\Sigma \Sigma^+ = \begin{bmatrix} 2 & 0 \\ 0 & 3 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} 1/2 & 0 & 0 \\ 0 & 1/3 & 0 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}$.

Secondly, $\Sigma^+ \Sigma = \begin{bmatrix} 1/2 & 0 \\ 0 & 1/3 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} 2 & 0 & 0 \\ 0 & 3 & 0 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}$ and $\Sigma \Sigma^+ = \begin{bmatrix} 2 & 0 & 0 \\ 0 & 3 & 0 \end{bmatrix} \begin{bmatrix} 1/2 & 0 \\ 0 & 1/3 \\ 0 & 0 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$.

[Note how the pseudoinverse tries to behave like the regular inverse. But since Σ has only 2 columns, $\Sigma^+ \Sigma$ and $\Sigma \Sigma^+$ can have rank at most 2 (so cannot be the full 3×3 identity).]

Thirdly, $\Sigma^+ \Sigma = \begin{bmatrix} 1/2 & 0 \\ 0 & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} 2 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}$ and $\Sigma \Sigma^+ = \begin{bmatrix} 2 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} 1/2 & 0 \\ 0 & 0 \\ 0 & 0 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$.

[Here, Σ has rank 1, so that $\Sigma^+ \Sigma$ and $\Sigma \Sigma^+$ can have rank at most 1.]

In general. Proceeding, as in this example, we find that the pseudoinverse of any $m \times n$ diagonal matrix Σ is the $n \times m$ (transposed dimensions!) diagonal matrix whose nonzero entries are the inverses of the entries of Σ .

Comment. Observe that, in all three cases, $\Sigma^{++} = \Sigma$.

Comment. Note that $\begin{bmatrix} 1 & 0 \\ 0 & \varepsilon \end{bmatrix}^+ = \begin{bmatrix} 1 & 0 \\ 0 & \varepsilon^{-1} \end{bmatrix}$ for small $\varepsilon \neq 0$, while $\begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}^+ = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$. This shows that the pseudoinverse is not a continuous operation.

It turns out that the pseudoinverse A^+ can be easily obtained from the SVD of A :

Theorem 172. The **pseudoinverse** of an $m \times n$ matrix A with SVD $A = U\Sigma V^T$ is

$$A^+ = V\Sigma^+U^T,$$

where Σ^+ , the pseudoinverse of Σ , is the $n \times m$ diagonal matrix, whose nonzero entries are the inverses of the entries of Σ .

Proof. The equation $Ax = b$ is equivalent to $U\Sigma V^T x = b$ and, thus, $\Sigma V^T x = U^T b$.

Write $y = V^T x$ and note that y and x have the same norm (why?!).

We already know that the equation $\Sigma y = U^T b$ has optimal solution $y = \Sigma^+ U^T b$.

Since y and x have the same norm, it follows that $x = Vy = V\Sigma^+ U^T b$ is the optimal solution to $Ax = b$.

Hence, $A^+ = V\Sigma^+ U^T$. □

Lemma 173. The pseudoinverse of A^+ is $A^{++} = A$.

Proof. Starting with the SVD $A = U\Sigma V^T$, we have $A^+ = V\Sigma^+ U^T$, which is the SVD of A^+ .

Therefore, $A^{++} = U\Sigma^{++} V^T$. The claim thus follows from $\Sigma^{++} = \Sigma$. □

Example 174. Determine the pseudoinverse of $A = \begin{bmatrix} 1 & -1 \\ 0 & 1 \\ 1 & 0 \end{bmatrix}$ in two ways.

First, using the SVD and, second, using the fact that A has full column rank.

Solution. (SVD) We have computed the SVD of this matrix before.

Since $A = U\Sigma V^T$ with $U = \begin{bmatrix} -2/\sqrt{6} & 0 & -1/\sqrt{3} \\ 1/\sqrt{6} & 1/\sqrt{2} & -1/\sqrt{3} \\ -1/\sqrt{6} & 1/\sqrt{2} & 1/\sqrt{3} \end{bmatrix}$, $\Sigma = \begin{bmatrix} \sqrt{3} & 0 \\ 0 & 1 \\ 0 & 0 \end{bmatrix}$, $V = \frac{1}{\sqrt{2}} \begin{bmatrix} -1 & 1 \\ 1 & 1 \end{bmatrix}$,

the pseudoinverse is $A^+ = V\Sigma^+ U^T$ where $\Sigma^+ = \begin{bmatrix} 1/\sqrt{3} & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix}$.

Multiplying these matrices, $A^+ = \frac{1}{3} \begin{bmatrix} 1 & 1 & 2 \\ -1 & 2 & 1 \end{bmatrix}$.

Comment. For many applications, it may be neither necessary nor helpful to multiply V, Σ^+, U^T .

Solution. (full column rank) Since A clearly has full column rank, we also have $A^+ = (A^T A)^{-1} A^T$.

Indeed, $A^+ = (A^T A)^{-1} A^T = \begin{bmatrix} 2 & -1 \\ -1 & 2 \end{bmatrix}^{-1} \begin{bmatrix} 1 & 0 & 1 \\ -1 & 1 & 0 \end{bmatrix} = \frac{1}{3} \begin{bmatrix} 2 & 1 \\ 1 & 2 \end{bmatrix} \begin{bmatrix} 1 & 0 & 1 \\ -1 & 1 & 0 \end{bmatrix} = \frac{1}{3} \begin{bmatrix} 1 & 1 & 2 \\ -1 & 2 & 1 \end{bmatrix}$.

Example 175. What is the pseudoinverse of $A = \begin{bmatrix} 2 & 2 \\ 1 & 1 \end{bmatrix}$?

Solution. Recall (or compute) that $A = U\Sigma V^T$ with $U = \frac{1}{\sqrt{5}} \begin{bmatrix} 2 & -1 \\ 1 & 2 \end{bmatrix}$, $\Sigma = \begin{bmatrix} \sqrt{10} & 0 \\ 0 & 0 \end{bmatrix}$, $V = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & -1 \\ 1 & 1 \end{bmatrix}$.

Hence, $A^+ = V\Sigma^+ U^T$ where $\Sigma^+ = \begin{bmatrix} 1/\sqrt{10} & 0 \\ 0 & 0 \end{bmatrix}$.

Multiplying these matrices (which may not be necessary or helpful for applications), $A^+ = \frac{1}{10} \begin{bmatrix} 2 & 1 \\ 2 & 1 \end{bmatrix}$.

Note. Since A does not have full column rank, $A^+ = (A^T A)^{-1} A^T$ cannot be used. That's because $A^T A$ is not invertible.

Comment. Here, $A^+ A = v_1 v_1^T = \frac{1}{2} \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix}$ and $A A^+ = u_1 u_1^T = \frac{1}{5} \begin{bmatrix} 4 & 2 \\ 2 & 1 \end{bmatrix}$ are not visually like the identity. However, note that these are the (orthogonal) projections onto v_1 and u_1 respectively (in particular, the eigenvalues are $1, 0$).

Review.

- If the $m \times n$ matrix A has SVD $A = U\Sigma V^T$, then its pseudoinverse is $A^+ = V\Sigma^+U^T$.
Here, Σ^+ , the pseudoinverse of Σ , is the $n \times m$ diagonal matrix, whose nonzero entries are the inverses of the entries of Σ .
- The system $Ax = b$ has “optimal” solution $x = A^+b$.
Here, “optimal” means that x is the smallest least squares solution.

Example 176.

- Find the pseudoinverse of $A = [1 \ 2 \ 3]$.
- Find the smallest solution to $x_1 + 2x_2 + 3x_3 = 6$.

As before, smallest solutions means the solution x such that $\|x\|$ is as small as possible. One obvious solution is $[1, 1, 1]^T$, but is it the smallest?

Solution.

(a) $A^T A = \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix} [1 \ 2 \ 3] = \begin{bmatrix} 1 & 2 & 3 \\ 2 & 4 & 6 \\ 3 & 6 & 9 \end{bmatrix}$ has 14-eigenvector $\begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix}$ and 0-eigenvectors $\begin{bmatrix} -2 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} -3 \\ 0 \\ 1 \end{bmatrix}$.

$$u_1 = \frac{1}{\sigma_1} A v_1 = \frac{1}{\sqrt{14}} [1 \ 2 \ 3] \frac{1}{\sqrt{14}} \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix} = 1$$

Hence, $A = U\Sigma V^T$ with $U = [1]$, $\Sigma = [\sqrt{14} \ 0 \ 0]$, $V = \begin{bmatrix} 1/\sqrt{14} & * & * \\ 2/\sqrt{14} & * & * \\ 3/\sqrt{14} & * & * \end{bmatrix}$.

$$A^+ = V\Sigma^+U^T = \begin{bmatrix} 1/\sqrt{14} & * & * \\ 2/\sqrt{14} & * & * \\ 3/\sqrt{14} & * & * \end{bmatrix} \begin{bmatrix} 1/\sqrt{14} \\ 0 \\ 0 \end{bmatrix} [1] = \frac{1}{14} \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix}$$

Comment. No surprise on U . The only options for U are $U = [1]$ and $U = [-1]$.

Comment. Realizing what we did here allows us to write down A^+ immediately for all $1 \times n$ matrices A . See Example 177.

Homework. Complete the SVD of A . That is, find an option for the two missing columns of V , so that V is an orthogonal matrix. In other words, find an orthonormal basis for the 0-eigenspace.

Comment. An even better approach would be to compute AA^T first (instead of $A^T A$) which would allow us to compute U first (rather than V first). Can you fill in the blanks?

- We are solving $Ax = [6]$ with $A = [1 \ 2 \ 3]$ as in the previous example.

We conclude that the smallest solution is $x = A^+[6] = \frac{3}{7} \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix}$.

Compare. $\left\| \frac{3}{7} \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix} \right\| = \frac{3}{7} \sqrt{14} \approx 1.604$ is indeed smaller than, say, $\left\| \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} \right\| = \sqrt{3} \approx 1.732$.

Geometric picture. The equation $x_1 + 2x_2 + 3x_3 = 6$ describes a plane (not through the origin), and we are asking for the point on that plane which is closest to the origin. That’s a typical question in Calculus III. Note that $[1 \ 2 \ 3]^T$ is the normal vector of the plane. Explain why the answer had to be a multiple of that normal vector!

Example 177. More generally, find the pseudoinverse of $A = [a_1 \ a_2 \ a_3]$.

Solution. As in the previous example, we see that the answer will be $A^+ = \frac{a}{\|a\|^2}$ with $a = \begin{bmatrix} a_1 \\ a_2 \\ a_3 \end{bmatrix}$.

Comment. Likewise for $A = [a_1 \ a_2 \ \dots \ a_n]$.

Example 178. How is the rank of A reflected in its singular value decomposition $A = U\Sigma V^T$?

Solution. The rank of A is equal to the number of nonzero singular values.

Theorem 179. (matrix approximation lemma) Suppose A is a $m \times n$ matrix, and we want to approximate A using a matrix B of rank s (smaller than the rank of A). Let $A = U\Sigma V^T$ be the SVD of A (with singular values in decreasing order). Then, the best such approximation is $B = U_s \Sigma_s V_s^T$, where Σ_s is the $s \times s$ diagonal matrix with entries $\sigma_1, \sigma_2, \dots, \sigma_s$ and U_s, V_s are obtained from U, V by only taking the first s columns.

Comment. Note that, by choosing s small compared to r , we can store an approximation of A using much less data. This approximation will be good if the omitted singular values $\sigma_{s+1}, \sigma_{s+2}, \dots, \sigma_r$ are all “small”.

Comment. Equivalently, $B = U\Sigma_s V^T$, where Σ_s is now obtained from Σ by setting all but the largest s singular values to 0. In other words, Σ_s has the values $\sigma_1, \sigma_2, \dots, \sigma_s$ on its diagonal, followed by zeros.

In other words. Here is another common way to say the same thing:

- Observe that $A = U\Sigma V^T$ is equivalent to $A = \sum_{i=1}^r \sigma_i \mathbf{u}_i \mathbf{v}_i^T$.
- Each matrix $\mathbf{u}_i \mathbf{v}_i^T$ has rank 1.
- The best rank s approximation to A is $B = \sum_{i=1}^s \sigma_i \mathbf{u}_i \mathbf{v}_i^T$.

Advanced comment. Here, “best” approximation is measured using the Frobenius norm of a matrix A (which is the same as the norm of a vector with all the entries of A).

Example 180. Determine the best rank 1 approximation of $A = \begin{bmatrix} 1 & 1 & 1 \\ 1 & 0 & -1 \end{bmatrix}$.

Solution. We determine (do it!) that A has the SVD

$$A = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix} \begin{bmatrix} \sqrt{3} & 0 & 0 \\ 0 & \sqrt{2} & 0 \end{bmatrix} \begin{bmatrix} 1/\sqrt{3} & -1/\sqrt{2} & 1/\sqrt{6} \\ 1/\sqrt{3} & 0 & -2/\sqrt{6} \\ 1/\sqrt{3} & 1/\sqrt{2} & 1/\sqrt{6} \end{bmatrix}^T = \begin{bmatrix} 1 & 1 & 1 \\ 1 & 0 & -1 \end{bmatrix}.$$

Hence, the best rank 1 approximation of A is (that is, we keep 1 singular value only) is

$$\begin{bmatrix} 1 \\ 0 \end{bmatrix} \begin{bmatrix} \sqrt{3} \end{bmatrix} \begin{bmatrix} 1/\sqrt{3} \\ 1/\sqrt{3} \\ 1/\sqrt{3} \end{bmatrix}^T = \begin{bmatrix} 1 & 1 & 1 \\ 0 & 0 & 0 \end{bmatrix}.$$

Comment. Equivalently, $\begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix} \begin{bmatrix} \sqrt{3} & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} 1/\sqrt{3} & -1/\sqrt{2} & 1/\sqrt{6} \\ 1/\sqrt{3} & 0 & -2/\sqrt{6} \\ 1/\sqrt{3} & 1/\sqrt{2} & 1/\sqrt{6} \end{bmatrix}^T = \begin{bmatrix} 1 & 1 & 1 \\ 0 & 0 & 0 \end{bmatrix}$.

Example 181. Determine the best rank 1 approximation of $A = \begin{bmatrix} 1 & -1 \\ 0 & 1 \\ 1 & 0 \end{bmatrix}$.

Solution. Recall that $A = U\Sigma V^T$ with $U = \begin{bmatrix} -2/\sqrt{6} & 0 & -1/\sqrt{3} \\ 1/\sqrt{6} & 1/\sqrt{2} & -1/\sqrt{3} \\ -1/\sqrt{6} & 1/\sqrt{2} & 1/\sqrt{3} \end{bmatrix}$, $\Sigma = \begin{bmatrix} \sqrt{3} & 0 \\ 0 & 1 \\ 0 & 0 \end{bmatrix}$, $V = \frac{1}{\sqrt{2}} \begin{bmatrix} -1 & 1 \\ 1 & 1 \end{bmatrix}$.

Hence, the best rank 1 approximation of A is $\frac{1}{\sqrt{6}} \begin{bmatrix} -2 \\ 1 \\ -1 \end{bmatrix} \begin{bmatrix} \sqrt{3} \end{bmatrix} \frac{1}{\sqrt{2}} \begin{bmatrix} -1 \\ 1 \end{bmatrix}^T = \frac{1}{2} \begin{bmatrix} 2 & -2 \\ -1 & 1 \\ 1 & -1 \end{bmatrix}$.

Comment. The best rank 2 approximation of A is, of course, A itself. Observe that we can delete the third column of U and the third row of Σ without changing the product.

Example 182. (image compression) Let us load a 341x512 grayscale photo and store it as a matrix A . Each entry of the matrix is a value between 0 (black) and 1 (white).

The beautiful picture is taken from: <http://www.southalabama.edu/departments/publicrelations/brand/photography.html>


[The same approach works with color pictures. These are often represented by three matrices: one for the red component of the pixel, one for the green and for the blue component (RGB color scheme).]

```
Sage] import pylab
Sage] A = matrix(pylab.imread('/home/armin/photo.png'))
Sage] A.dimensions()
(341, 512)
Sage] A[0,0]
0.137254908681
Sage] matrix_plot(A, cmap='gray')
```



Next, we compute the SVD of A . Despite the size of A that takes the computer only a fraction of a second:

```
Sage] U,S,V = A.SVD()
Sage] S.diagonal()[:6]
[238.443435709, 79.4429775448, 35.4540786319, 20.5662302846, 20.0697710337, 13.3421216529]
Sage] list_plot(S.diagonal(), scale='semilogy')
```



As we can see, the magnitude of the singular values drops off quickly. We get a good approximation to A (our original photo) by computing a best rank s approximation to A by computing $U_s \Sigma_s V_s$ where Σ_s is the $s \times s$ diagonal matrix with entries $\sigma_1, \sigma_2, \dots, \sigma_s$ and U_s, V_s are obtained from the corresponding matrices in the SVD $A = U \Sigma V^T$ by only taking the first s columns.

```
Sage] def A_approx(s):  
    U0 = U.matrix_from_columns([0..s-1])  
    S0 = diagonal_matrix(S.diagonal()[0:s])  
    V0 = V.matrix_from_columns([0..s-1])  
    return U0*S0*V0.transpose()
```


Taking only 100 of the 341 singular values, we get an approximation which is almost as good as the original:

```
Sage] matrix_plot(A_approx(100), cmap='gray')
```



But notice the development of artifacts. Taking only 20 of the 341 singular values, a lot is lost:

```
Sage] matrix_plot(A_approx(20), cmap='gray')
```



Comment. Image compression is just one (nice visual) example of the power of SVD. A variation of this approach can, for instance, also be used for image denoising. Much more generally, the SVD is able to extract the most important features of any sort of data!

Review. matrix approximation and compression

Function spaces

Recall the following:

- We call objects **vectors** if they can be added and scaled (subject to the usual laws).
- A set of vectors is a **vector space** if it is closed under addition and scaling.

In other words, vector spaces are spans.

We will now discuss spaces of vectors, where the vectors are functions.

Why? Just one example why it is super useful to apply our linear algebra machinery to functions: we discussed the **distance** between vectors and how to find vectors closest to interesting subspaces (i.e. orthogonal projections). These notions are important for functions, too. For instance, given a (complicated) function, we want to find the closest function in a subspace of (simple) functions. In other words, we want to approximate functions using other (typically, simpler) functions.

Comment. Functions $f(x)$ and $g(x)$ can also be multiplied. This is an extra structure (it makes appropriate sets of functions an **algebra**, which is something more special than a **vector space**), which we ignore during our discussion of vector spaces.

An inner product on function spaces

On the space of, say, (piecewise) continuous functions $f: [a, b] \rightarrow \mathbb{R}$, it is natural to consider the dot product

$$\langle f, g \rangle = \int_a^b f(t)g(t)dt.$$

Why? A (sensible) dot product provides a (sensible) notion of distance between functions. The dot product above is the continuous analog of the usual dot product $\langle \mathbf{x}, \mathbf{y} \rangle = \sum_{t=1}^n x_t y_t$ for vectors in \mathbb{R}^n . Do you see it?!

As a consequence, once we have the dot product, we can orthogonally project functions onto spaces of simple functions. In other words, we can compute best approximations of functions by simple functions (for instance, best quadratic approximations).

Why continuous? We need that any product $f(x)g(x)$ is integrable. That means we cannot work with all functions. Continuity is certainly sufficient. In fact, the right condition is that $f(x)^2$ should be integrable on $[a, b]$ (i.e. $f(x)$ is square-integrable). Such a function is said to be in $\mathcal{L}^2[a, b]$.

Example 183. What is the orthogonal projection of $f: [a, b] \rightarrow \mathbb{R}$ onto the space of constant functions (that is, $\text{span}\{1\}$)?

Solution. The orthogonal projection of $f: [a, b] \rightarrow \mathbb{R}$ onto $\text{span}\{1\}$ is

$$\frac{\langle f, 1 \rangle}{\langle 1, 1 \rangle} = \frac{\int_a^b f(t)1dt}{\int_a^b 1^2 dt} = \frac{1}{b-a} \int_a^b f(t)dt.$$

This is the average of $f(x)$ on $[a, b]$.

Comment. Makes perfect sense, doesn't it? Intuitively, the best approximation of a function by a constant should indeed be the one where the constant is the average.


Example 184. Find the best approximation (in the L^2 sense) of $f(x) = \sqrt{x}$ on the interval $[0, 1]$ using a function of the form $y = ax$.

Solution. The orthogonal projection of $f: [0, 1] \rightarrow \mathbb{R}$ onto $\text{span}\{x\}$ is

$$\frac{\langle f, x \rangle}{\langle x, x \rangle} x = \frac{\int_0^1 f(t)t dt}{\int_0^1 t^2 dt} x = 3x \int_0^1 t f(t) dt.$$

In our case, the best approximation is

$$3x \int_0^1 t\sqrt{t} dt = 3x \int_0^1 t^{3/2} dt = 3x \left[\frac{1}{5/2} t^{5/2} \right]_0^1 = \frac{6}{5}x.$$



What does “in the L^2 sense” mean? There are various ways to make precise what “best approximation” should mean. Here, we mean that the norm of the error gets minimized: since the error is $\text{error}(t) = \sqrt{t} - at$, this means that

$$\|\text{error}\|^2 = \langle \text{error}, \text{error} \rangle = \int_0^1 \text{error}(t)^2 dt = \int_0^1 (\sqrt{t} - at)^2 dt$$

is minimized over all choices for a (note that it does not make a difference whether we are minimizing the norm or the square of the norm). Since the above norm is known as the L^2 norm, we are finding the best approximation “in the L^2 sense”. For instance, an alternative meaningful norm is the L^∞ norm, for which $\|\text{error}\|$ simply measures the maximum absolute value of the error (this may sound simpler but there is no corresponding notion of orthogonality so that we cannot apply tools like orthogonal projections).

For comparison. Let’s minimize $r(a) := \|\text{error}\|^2$ directly! First, we compute that

$$r(a) = \int_0^1 (\sqrt{t} - at)^2 dt = \int_0^1 (t - 2at^{3/2} + a^2t^2) dt = \left[\frac{1}{2}t^2 - \frac{4}{5}at^{5/2} + \frac{1}{3}a^2t^3 \right]_0^1 = \frac{1}{2} - \frac{4}{5}a + \frac{1}{3}a^2.$$

We now minimize $r(a)$ over all choices of a using what we learned in Calculus I: if the minimum occurs at a , then we necessarily have $r'(a) = 0$. We compute that $r'(a) = -\frac{4}{5} + \frac{2}{3}a$. Solving $-\frac{4}{5} + \frac{2}{3}a = 0$, we find $a = \frac{6}{5}$ as the only candidate. Which matches exactly what we have found above!

Clearly, this approach becomes more challenging if our approximations have more than one degree of freedom. However, our linear algebra approach continues to work fine, as we will see in the next example.

Example 185. Find the best approximation of $f(x) = \sqrt{x}$ (in the L^2 sense) on the interval $[0, 1]$ using a function of the form $y = a + bx$.

Important observation. The orthogonal projection of $f: [0, 1] \rightarrow \mathbb{R}$ onto $\text{span}\{1, x\}$ is not simply the projection onto 1 plus the projection onto x . That's because 1 and x are not orthogonal:

$$\langle 1, x \rangle = \int_0^1 t dt = \frac{1}{2} \neq 0.$$

Solution. To find an orthogonal basis for $\text{span}\{1, x\}$, following Gram–Schmidt, we compute

$$x - \left(\begin{array}{c} \text{projection of} \\ x \text{ onto } 1 \end{array} \right) = x - \frac{\langle x, 1 \rangle}{\langle 1, 1 \rangle} 1 = x - \frac{1}{2}.$$

Hence, $1, x - \frac{1}{2}$ is an orthogonal basis for $\text{span}\{1, x\}$.

The orthogonal projection of \sqrt{x} on $[0, 1]$ onto $\text{span}\{1, x\} = \text{span}\left\{1, x - \frac{1}{2}\right\}$ therefore is

$$\frac{\langle \sqrt{x}, 1 \rangle}{\langle 1, 1 \rangle} 1 + \frac{\langle \sqrt{x}, x - \frac{1}{2} \rangle}{\langle x - \frac{1}{2}, x - \frac{1}{2} \rangle} \left(x - \frac{1}{2}\right) = \frac{\int_0^1 \sqrt{t} dt}{\int_0^1 1 dt} + \frac{\int_0^1 \sqrt{t} \left(t - \frac{1}{2}\right) dt}{\int_0^1 \left(t - \frac{1}{2}\right)^2 dt} \left(x - \frac{1}{2}\right).$$


We compute the three new integrals:

$$\begin{aligned} \int_0^1 \sqrt{t} dt &= \left[\frac{2}{3} t^{3/2} \right]_0^1 = \frac{2}{3} \\ \int_0^1 \sqrt{t} \left(t - \frac{1}{2}\right) dt &= \int_0^1 \left(t^{3/2} - \frac{1}{2} t^{1/2}\right) dt = \left[\frac{2}{5} t^{5/2} - \frac{1}{3} t^{3/2} \right]_0^1 = \frac{2}{5} - \frac{1}{3} = \frac{1}{15} \\ \int_0^1 \left(t - \frac{1}{2}\right)^2 dt &= \int_0^1 \left(t^2 - t + \frac{1}{4}\right) dt = \left[\frac{1}{3} t^3 - \frac{1}{2} t^2 + \frac{1}{4} t \right]_0^1 = \frac{1}{3} - \frac{1}{2} + \frac{1}{4} = \frac{1}{12} \end{aligned}$$

Using these values, the best approximation is

$$\frac{\int_0^1 \sqrt{t} dt}{\int_0^1 1 dt} + \frac{\int_0^1 \sqrt{t} \left(t - \frac{1}{2}\right) dt}{\int_0^1 \left(t - \frac{1}{2}\right)^2 dt} \left(x - \frac{1}{2}\right) = \frac{2}{3} + \frac{12}{15} \left(x - \frac{1}{2}\right) = \frac{4}{5} x + \frac{4}{15}$$

The plot below confirms how good this linear approximation is (compare with the previous example):



Orthogonal polynomials

Example 186. Proceeding as in the previous example, compute an orthogonal basis for the space $\text{span}\{1, x, x^2, x^3\}$.

Solution. To find an orthogonal basis, we use Gram–Schmidt:

$$\begin{aligned}
 q_1 &= 1 \\
 q_2 &= x - \frac{\langle x, q_1 \rangle}{\langle q_1, q_1 \rangle} q_1 = x - \frac{\langle x, 1 \rangle}{\langle 1, 1 \rangle} 1 = x - \frac{1}{2} \\
 q_3 &= x^2 - \frac{\langle x^2, q_1 \rangle}{\langle q_1, q_1 \rangle} q_1 - \frac{\langle x^2, q_2 \rangle}{\langle q_2, q_2 \rangle} q_2 = x^2 - \frac{\langle x^2, 1 \rangle}{\langle 1, 1 \rangle} 1 - \frac{\langle x^2, x - \frac{1}{2} \rangle}{\langle x - \frac{1}{2}, x - \frac{1}{2} \rangle} \left(x - \frac{1}{2} \right) \\
 &= x^2 - \frac{1}{3} 1 - \frac{1}{12} \left(x - \frac{1}{2} \right) = x^2 - x + \frac{1}{6} \\
 q_4 &= x^3 - \frac{\langle x^3, q_1 \rangle}{\langle q_1, q_1 \rangle} q_1 - \frac{\langle x^3, q_2 \rangle}{\langle q_2, q_2 \rangle} q_2 - \frac{\langle x^3, q_3 \rangle}{\langle q_3, q_3 \rangle} q_3 = \dots = x^3 - \frac{3}{2} x^2 + \frac{3}{5} x - \frac{1}{20}
 \end{aligned}$$

The polynomials $1, x - \frac{1}{2}, x^2 - x + \frac{1}{6}, x^3 - \frac{3}{2}x^2 + \frac{3}{5}x - \frac{1}{20}$ form an orthogonal basis for the space of polynomials of degree at most 3.

Comment. Of course, we could keep going by next including x^4, x^5, \dots Up to scaling, the resulting polynomials are known as the **shifted Legendre polynomials** and they are an example of a family of **orthogonal polynomials**. They are important, for instance, in approximating more complicated functions using polynomials (see the previous example, for instance).

Homework. Fill in the details of the computation for q_4 (maybe using Sage for support). For instance, here is how to compute $\int_0^1 t^2 \left(t - \frac{1}{2} \right) dt$ using Sage:

```

>>> t = var('t')
>>> integral(t^2*(t-1/2), t, 0, 1)

```

$$\frac{1}{12}$$

In the literature, the interval $[0, 1]$ is often replaced with the interval $[-1, 1]$ (because of the symmetry). If we proceed as above, then the resulting orthogonal polynomials are known as the **Legendre polynomials**. In the case of the interval $[-1, 1]$, we consider the space of all polynomials (with real coefficients) together with the dot product

$$\langle p_1, p_2 \rangle = \int_{-1}^1 p_1(t)p_2(t)dt. \tag{1}$$

Comment. That dot product is useful if we are thinking about the polynomials as functions on $[-1, 1]$ only. You can, of course, consider any other interval and you will obtain a shifted version of what we get here.

Example 187. Are $1, x, x^2, \dots$ orthogonal (with respect to the inner product (1))?

Solution. Since $\langle x^r, x^s \rangle = \int_{-1}^1 t^r t^s dt = \int_{-1}^1 t^{r+s} dt$, we find that $\langle x^r, x^s \rangle = \begin{cases} \frac{2}{r+s+1}, & \text{if } r+s \text{ is even,} \\ 0, & \text{otherwise.} \end{cases}$

Hence, if $r+s$ is odd, then the monomials x^r and x^s are orthogonal. On the other hand, if $r+s$ is even, then x^r and x^s are not orthogonal.

Example 188. Use Gram-Schmidt to produce an orthogonal basis $\mathbf{p}_0, \mathbf{p}_1, \mathbf{p}_2, \dots$ for the space of polynomials with the dot product (1). Compute $\mathbf{p}_0, \mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_3, \mathbf{p}_4$.

Instead of normalizing these polynomials, **standardize** them so that $\mathbf{p}_n(1) = 1$.

Solution. We construct an orthogonal basis $\mathbf{p}_0, \mathbf{p}_1, \mathbf{p}_2, \dots$ from $1, x, x^2, \dots$ as follows:

- Starting with 1 , we find $\mathbf{p}_0(x) = 1$.

For future reference, let us note that $\|\mathbf{p}_0\|^2 = \int_{-1}^1 1 dx = 2$.

- Starting with x , Gram-Schmidt produces $x - \left(\frac{\text{projection of } x \text{ onto } \mathbf{p}_0}{\langle \mathbf{p}_0, \mathbf{p}_0 \rangle} \right) \mathbf{p}_0 = x - \frac{\langle x, \mathbf{p}_0 \rangle}{\langle \mathbf{p}_0, \mathbf{p}_0 \rangle} \mathbf{p}_0 = x - \int_{-1}^1 t dt = x$.

Again, that's already standardized, so that $\mathbf{p}_1(x) = x$.

Comment. The previous problem already told us that x is orthogonal to 1 .

For future reference, let us note that $\|\mathbf{p}_1\|^2 = \int_{-1}^1 t^2 dt = \frac{2}{3}$.

- Starting with x^2 , Gram-Schmidt produces $x^2 - \left(\frac{\text{projection of } x^2 \text{ onto } \text{span}\{\mathbf{p}_0, \mathbf{p}_1\}}{\langle \mathbf{p}_0, \mathbf{p}_0 \rangle} \mathbf{p}_0 + \frac{\langle x^2, \mathbf{p}_1 \rangle}{\langle \mathbf{p}_1, \mathbf{p}_1 \rangle} \mathbf{p}_1 \right) = x^2 - \frac{\langle x^2, \mathbf{p}_0 \rangle}{\langle \mathbf{p}_0, \mathbf{p}_0 \rangle} \mathbf{p}_0 - \frac{\langle x^2, \mathbf{p}_1 \rangle}{\langle \mathbf{p}_1, \mathbf{p}_1 \rangle} \mathbf{p}_1 = x^2 - \frac{1}{2} \int_{-1}^1 t^2 dt - \frac{x}{2/3} \int_{-1}^1 t^3 dt = x^2 - \frac{1}{3}$.

Hence, standardizing, $\mathbf{p}_2(x) = \frac{1}{2}(3x^2 - 1)$.

Comment. The previous problem told us that x^2 is orthogonal to x (but not to 1).

- Continuing, we find $\mathbf{p}_3(x) = \frac{1}{2}(5x^3 - 3x)$ and $\mathbf{p}_4(x) = \frac{1}{8}(35x^4 - 30x^2 + 3)$.

Comment. These famous polynomials are known as the **Legendre polynomials**. The Legendre polynomial \mathbf{p}_n is an even function if n is even, and an odd function if n is odd (can you explain why?!).

An explicit formula is $\mathbf{p}_n(x) = 2^{-n} \sum_{k=0}^n \binom{n}{k}^2 (x+1)^k (x-1)^{n-k}$.

For instance, $\mathbf{p}_2(x) = \frac{1}{4}((x-1)^2 + 2^2(x-1)(x+1) + (x+1)^2) = \frac{1}{2}(3x^2 - 1)$.

https://en.wikipedia.org/wiki/Legendre_polynomials

Comment. Legendre polynomials are an example of **orthogonal polynomials**. Each choice of dot product gives rise to a family of such orthogonal polynomials.

https://en.wikipedia.org/wiki/Orthogonal_polynomials

Comment. It is also particularly natural to consider the dot product (1), where the integral is from 0 to 1 . In that case, we obtain what's known as the shifted Legendre polynomials $\tilde{\mathbf{p}}_n(x) = \mathbf{p}_n(2x - 1)$. Compute the first few and compare with Example 186.

Comment on other norms. Our choice of inner product

$$\langle f, g \rangle = \int_a^b f(t)g(t)dt$$

for (square-integrable) functions on $[a, b]$ gives rise to the norm $\|f\| = \left(\int_a^b f(t)^2 dt \right)^{1/2}$. This is known as the L^2 -norm (and often written as $\|f\|_2$).

It is the continuous analog of the usual Euclidean norm $\|\mathbf{v}\| = (v_1^2 + v_2^2 + \dots)^{1/2}$ (known as ℓ^2 -norm).

There do exist other norms to measure the magnitude of vectors, such as the ℓ_1 -norm $\|\mathbf{v}\|_1 = |v_1| + |v_2| + \dots$ or, more generally, for $p \geq 1$, the ℓ_p -norms $\|\mathbf{v}\|_p = (|v_1|^p + |v_2|^p + \dots)^{1/p}$.

Likewise, for functions, we have the L^p -norms $\|f\|_p = \left(\int_a^b f(t)^p dt \right)^{1/p}$.

Only in the case $p = 2$ do these norms come from an inner product. That's a mathematical (as opposed to geometric) reason why we especially care about that case.

Example 189. Give a basis for the space of all polynomials.

Solution. $1, x, x^2, x^3, \dots$

Indeed, every polynomial $p(x) = a_0 + a_1x + a_2x^2 + \dots + a_nx^n$ can be written uniquely as a sum of these basis elements. (“can be” = span; “uniquely” = independent)

Comment. The dimension is ∞ . But we can make a list of basis elements, which is the “smallest kind of ∞ ” and is referred to as **countably infinite**. For the space of all functions, no such list can be made.

Just for fun. Let us indicate this difference in infiniteness in a slightly simpler situation: first, the natural numbers $0, 1, 2, 3, \dots$ are infinite but they are countable, because we can make a (infinite but complete) list starting with a first, then a second element and so on (hence, the name “countable”). On the other hand, consider the real numbers between 0 and 1 . Clearly, there are infinitely many such numbers. The somewhat shocking fact (first realized by Georg Cantor in 1874) is that every attempt of making a complete list of these numbers must fail because every list will inevitably miss some numbers. Here’s a brief indication of how the famous diagonal argument goes: suppose you can make a list, say:

#1	0.111111...
#2	0.123456...
#3	0.750000...
	⋮

Now, we are going to construct a new number $x = 0.x_1x_2x_3\dots$ with decimal digits x_i in such a way that the digit x_i differs (by more than 1) from the i th digit of number $\#i$ on our list. For instance, $0.352\dots$ in our case (for instance, $x_3 = 2$ differs from 0 , the 3 rd digit of sequence $\#3$). By construction, the number x is missing from the list.

Comment on fun. The statement “some infinities are bigger than others” nicely captures our observation. It appears in the book *The Fault in Our Stars* by John Green, where it is said by a cranky old author who attributes it to Cantor. Hazel, the main character, later reflects on that statement and compares $[0, 1]$ to $[0, 2]$. Can you explain why that is actually not what Cantor meant...?

Fourier series

A **Fourier series** for a function $f(x)$ is a series of the form

$$f(x) = a_0 + a_1\cos(x) + b_1\sin(x) + a_2\cos(2x) + b_2\sin(2x) + \dots$$

You may have seen Fourier series in other classes before. Our goal here is to tie them in with what we have learned about orthogonality.

In these other classes, you would have seen formulas for the coefficients a_k and b_k . We will see where those come from.

Observe that the right-hand side combination of cosines and sines is 2π -periodic.

Let us consider (nice) functions on $[0, 2\pi]$.

Or, equivalently, functions that are 2π -periodic.

We know that a natural inner product for that space of functions is

$$\langle f, g \rangle = \int_0^{2\pi} f(t)g(t)dt.$$

Example 190. Show that $\cos(x)$ and $\sin(x)$ are orthogonal (in that sense).

Solution. $\langle \cos(x), \sin(x) \rangle = \int_0^{2\pi} \cos(t)\sin(t)dt = \left[\frac{1}{2}(\sin(t))^2 \right]_0^{2\pi} = 0$

In fact:

All the functions $1, \cos(x), \sin(x), \cos(2x), \sin(2x), \dots$ are orthogonal to each other!

Moreover, they form a basis in the sense that every other (nice) function can be written as a (infinite) linear combination of these basis functions.

Example 191. What is the norm of $\cos(x)$?

Solution. $\langle \cos(x), \cos(x) \rangle = \int_0^{2\pi} \cos(t)\cos(t)dt = \pi$

Why? There's many ways to evaluate this integral. For instance:

- integration by parts
- using a trig identity
- here's a simple way:
 - $\int_0^{2\pi} \cos^2(t)dt = \int_0^{2\pi} \sin^2(t)dt$ (cos and sin are just a shift apart)
 - $\cos^2(t) + \sin^2(t) = 1$
 - So: $\int_0^{2\pi} \cos^2(t)dt = \frac{1}{2} \int_0^{2\pi} 1 dx = \pi$

Hence, $\cos(x)$ is not normalized. It has norm $\|\cos(x)\| = \sqrt{\pi}$.

Similarly. The same calculation shows that $\cos(kx)$ and $\sin(kx)$ have norm $\sqrt{\pi}$ as well.

Example 192. How do we find, say, b_2 ?

Solution. Since the functions $1, \cos(x), \sin(x), \cos(2x), \sin(2x), \dots$, the term $b_2\sin(2x)$ is the orthogonal projection of $f(x)$ onto $\sin(2x)$.

$$\text{In particular, } b_2 = \frac{\langle f(x), \sin(2x) \rangle}{\langle \sin(2x), \sin(2x) \rangle} = \frac{1}{\pi} \int_0^{2\pi} f(t) \sin(2t) dt.$$

In conclusion:

A (nice) $f(x)$ on $[0, 2\pi]$ has the Fourier series

$$f(x) = a_0 + a_1\cos(x) + b_1\sin(x) + a_2\cos(2x) + b_2\sin(2x) + \dots$$

where

$$a_k = \frac{\langle f(x), \cos(kx) \rangle}{\langle \cos(kx), \cos(kx) \rangle} = \frac{1}{\pi} \int_0^{2\pi} f(t) \cos(kt) dt,$$

$$b_k = \frac{\langle f(x), \sin(kx) \rangle}{\langle \sin(kx), \sin(kx) \rangle} = \frac{1}{\pi} \int_0^{2\pi} f(t) \sin(kt) dt,$$

$$a_0 = \frac{\langle f(x), 1 \rangle}{\langle 1, 1 \rangle} = \frac{1}{2\pi} \int_0^{2\pi} f(t) dt.$$

The next example illustrates that we can likewise deal with intervals other than $[0, 2\pi]$ (or, equivalently, 2π -periodic functions).

The main observation is that, since $\cos(x)$ has period 2π , the scaled function $\cos\left(\frac{2\pi}{L}x\right)$ has period L .

As we are just scaling, it is not hard to see that the functions

$$1, \cos\left(\frac{2\pi}{L}x\right), \sin\left(\frac{2\pi}{L}x\right), \cos\left(2 \cdot \frac{2\pi}{L}x\right), \sin\left(2 \cdot \frac{2\pi}{L}x\right), \cos\left(3 \cdot \frac{2\pi}{L}x\right), \dots$$

are still orthogonal to each other—now, adjusted for period L , with respect to the inner product

$$\langle f, g \rangle = \int_0^L f(t)g(t) dt.$$

Example 193. Suppose that $f(x)$ is 5-periodic. Write down the first few terms of the Fourier series for $f(x)$ with undetermined coefficients. Spell out how to compute the coefficients of the sine functions.

Solution. The Fourier series for $f(x)$ is

$$f(x) = a_0 + a_1\cos\left(\frac{2\pi}{5}x\right) + b_1\sin\left(\frac{2\pi}{5}x\right) + a_2\cos\left(\frac{4\pi}{5}x\right) + b_2\sin\left(\frac{4\pi}{5}x\right) + a_3\cos\left(\frac{6\pi}{5}x\right) + \dots$$


The coefficients b_n can be computed as

$$b_n = \frac{\langle f(x), \sin\left(\frac{2\pi}{5}nx\right) \rangle}{\langle \sin\left(\frac{2\pi}{5}nx\right), \sin\left(\frac{2\pi}{5}nx\right) \rangle} = \frac{\int_0^5 f(t) \sin\left(\frac{2\pi}{5}nt\right) dt}{\int_0^5 \sin^2\left(\frac{2\pi}{5}nt\right) dt} = \frac{2}{5} \int_0^5 f(t) \sin\left(\frac{2\pi}{5}nt\right) dt.$$

For the final (optional) equality, we used that $\int_0^5 \sin^2\left(\frac{2\pi}{5}nt\right) dt = \int_0^5 \cos^2\left(\frac{2\pi}{5}nt\right) dt$ combined with $\cos^2 + \sin^2 = 1$ to conclude that the integral in the denominator must be $\frac{5}{2}$.

Example 194. Find the Fourier series of the 2π -periodic function $f(t)$ defined by


$$f(t) = \begin{cases} -1, & \text{for } t \in (-\pi, 0), \\ +1, & \text{for } t \in (0, \pi), \\ 0, & \text{for } t = -\pi, 0, \pi. \end{cases}$$



Solution. We compute $a_0 = \frac{1}{\pi} \int_{-\pi}^{\pi} f(t) dt = 0$, as well as

$$\begin{aligned} a_n &= \frac{1}{\pi} \int_{-\pi}^{\pi} f(t) \cos(nt) dt = \frac{1}{\pi} \left[- \int_{-\pi}^0 \cos(nt) dt + \int_0^{\pi} \cos(nt) dt \right] = 0 \\ b_n &= \frac{1}{\pi} \int_{-\pi}^{\pi} f(t) \sin(nt) dt = \frac{1}{\pi} \left[- \int_{-\pi}^0 \sin(nt) dt + \int_0^{\pi} \sin(nt) dt \right] = \frac{2}{\pi n} [1 - \cos(n\pi)] \\ &= \frac{2}{\pi n} [1 - (-1)^n] = \begin{cases} \frac{4}{\pi n} & \text{if } n \text{ is odd} \\ 0 & \text{if } n \text{ is even} \end{cases}. \end{aligned}$$

In conclusion, $f(t) = \sum_{\substack{n=1 \\ n \text{ odd}}}^{\infty} \frac{4}{\pi n} \sin(nt) = \frac{4}{\pi} \left(\sin(t) + \frac{1}{3} \sin(3t) + \frac{1}{5} \sin(5t) + \dots \right)$.



Observation. The coefficients a_n are zero for all n if and only if $f(t)$ is odd.

Comment. The value of $f(t)$ for $t = -\pi, 0, \pi$ is irrelevant to the computation of the Fourier series. They are chosen so that $f(t)$ is equal to the Fourier series for all t (recall that, at a jump discontinuity t , the Fourier series converges to the average $\frac{f(t^-) + f(t^+)}{2}$).

Comment. Plot the (sum of the) first few terms of the Fourier series. What do you observe? The “overshooting” is known as the **Gibbs phenomenon**: https://en.wikipedia.org/wiki/Gibbs_phenomenon

Comment. Set $t = \frac{\pi}{2}$ in the Fourier series we just computed, to get Leibniz’ series $\pi = 4 \left[1 - \frac{1}{3} + \frac{1}{5} - \frac{1}{7} + \dots \right]$. For such an alternating series, the error made by stopping at the term $1/n$ is on the order of $1/n$. To compute the 768 digits of π to get to the Feynman point (3.14159265...721134999999...), we would (roughly) need $1/n < 10^{-768}$, or $n > 10^{768}$. That’s a lot of terms! (Roger Penrose, for instance, estimates that there are about 10^{80} atoms in the observable universe.)

Remark. Convergence of such series is not completely obvious! Recall, for instance, that the (odd part of) the harmonic series $1 + \frac{1}{3} + \frac{1}{5} + \frac{1}{7} + \dots$ diverges. (On the other hand, do you remember the alternating sign test from Calculus II?)

Linear transformations

Throughout, V and W are vector spaces.

Just like we went from column vectors to abstract vectors (such as polynomials), the concept of a matrix leads to abstract linear transformations.

In the other direction, picking a basis, abstract vectors can be represented as column vectors (see Lecture 35). Correspondingly, linear transformations can then be represented as matrices.

Definition 195. A map $T: V \rightarrow W$ is a **linear transformation** if

$$T(c\mathbf{x} + d\mathbf{y}) = cT(\mathbf{x}) + dT(\mathbf{y}) \quad \text{for all } \mathbf{x}, \mathbf{y} \text{ in } V \text{ and all } c, d \text{ in } \mathbb{R}.$$

In other words, a linear transformation respects addition and scaling:

- $T(\mathbf{x} + \mathbf{y}) = T(\mathbf{x}) + T(\mathbf{y})$
- $T(c\mathbf{x}) = cT(\mathbf{x})$

It necessarily sends the zero vector in V to the zero vector in W :

- $T(\mathbf{0}) = \mathbf{0}$ [because $T(\mathbf{0}) = T(0 \cdot \mathbf{0}) = 0 \cdot T(\mathbf{0}) = \mathbf{0}$]

Comment. Linear transformations are special functions and, hence, can be composed. For instance, if $T: V \rightarrow W$ and $S: U \rightarrow V$ are linear transformations, then $T \circ S$ is a linear transformation $U \rightarrow W$ (sending \mathbf{x} to $T(S(\mathbf{x}))$). If S, T are represented by matrices A, B , then $T \circ S$ is represented by the matrix BA . In other words, matrix multiplication arises as the composition of (linear) functions.

Example 196. The **derivative** you know from Calculus I is linear.

Indeed, the map $D: \left\{ \begin{array}{l} \text{space of all} \\ \text{differentiable} \\ \text{functions} \end{array} \right\} \rightarrow \left\{ \begin{array}{l} \text{space of all} \\ \text{functions} \end{array} \right\}$ defined by $f(x) \mapsto f'(x)$ is a linear transformation:

- $\underbrace{D(f(x) + g(x))}_{(f(x)+g(x))'} = \underbrace{D(f(x))}_{f'(x)} + \underbrace{D(g(x))}_{g'(x)}$
- $\underbrace{D(cf(x))}_{(cf(x))'} = c \underbrace{D(f(x))}_{cf'(x)}$

These are among the first properties you learned about the derivative.

Similarly, the **integral** you love from Calculus II is linear:

$$\int_a^b (f(x) + g(x))dx = \int_a^b f(x)dx + \int_a^b g(x)dx, \quad \int_a^b cf(x)dx = c \int_a^b f(x)dx$$

In this form, we are looking at a map $T: \left\{ \begin{array}{l} \text{space of all} \\ \text{continuous} \\ \text{functions} \end{array} \right\} \rightarrow \mathbb{R}$ defined by $T(f(x)) = \int_a^b f(x)dx$.

Example 197. Consider the space V of all polynomials $p(x)$ of degree 3 or less. The map $D: V \rightarrow V$ given by $p(x) \mapsto p'(x)$ is a linear. Write down the matrix M for this linear map with respect to the basis $1, x, x^2, x^3$.

Solution. $M = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 3 \\ 0 & 0 & 0 & 0 \end{bmatrix}$

For instance, the 3rd column says that x^2 (the 3rd basis element) gets sent to $0 \cdot 1 + 2 \cdot x + 0 \cdot x^2 + 0 \cdot x^3 = 2x$.

Example 198. Consider the map

$$D: \left\{ \begin{array}{l} \text{space of poly's} \\ \text{of degree } \leq 3 \end{array} \right\} \rightarrow \left\{ \begin{array}{l} \text{space of poly's} \\ \text{of degree } \leq 2 \end{array} \right\}, \quad p(x) \mapsto p'(x).$$

Write down the matrix M for this linear map with respect to the bases $1, x, x^2, x^3$ and $1, x, x^2$.

Solution. $M = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 3 \end{bmatrix}$

For instance, the 3rd column says that x^2 (the 3rd basis element) gets sent to $0 \cdot 1 + 2 \cdot x + 0 \cdot x^2 = 2x$.

Example 199. What is the pseudo-inverse of the matrix M from the previous example? Interpret your finding.

Solution. (final answer only) The pseudo-inverse of $\begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 3 \end{bmatrix}$ is $\begin{bmatrix} 0 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 1/2 & 0 \\ 0 & 0 & 1/3 \end{bmatrix}$.

The corresponding linear map sends 1 to x , x to $\frac{1}{2}x^2$ and x^2 to $\frac{1}{3}x^3$. That is, the pseudo-inverse computes the antiderivative of each monomial.

Comment. This is not surprising, since we are familiar from Calculus with the concepts of derivatives and antiderivatives (or integrals), and that these are “pseudo” inverse to each other.

Comment. Similarly, the pseudo-inverse of $\begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 3 \\ 0 & 0 & 0 & 0 \end{bmatrix}$ is $\begin{bmatrix} 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 1/2 & 0 & 0 \\ 0 & 0 & 1/3 & 0 \end{bmatrix}$.

Now, the corresponding linear map sends 1 to x , x to $\frac{1}{2}x^2$, x^2 to $\frac{1}{3}x^3$, and x^3 to 0 . That is, the pseudo-inverse computes the antiderivative of each monomial, with the exception of x^3 which gets sent to 0 (its antiderivative does not live in the space of polynomials of degree ≤ 3).

Example 200. (The April Fools' Day “proof” that $\pi = 4$, cont'd)

In that “proof”, we are constructing curves c_n with the property that $c_n \rightarrow c$ where c is the circle. This convergence can be understood, for instance, in the same sense $\|c_n - c\| \rightarrow 0$ with the norm introduced as we did for functions.

Since $c_n \rightarrow c$ we then wanted to conclude that $\text{perimeter}(c_n) \rightarrow \text{perimeter}(c)$, leading to $4 \rightarrow \pi$.

However, in order to conclude from $x_n \rightarrow x$ that $f(x_n) \rightarrow f(x)$ we need that f is continuous (at x)!!

The “function” **perimeter**, however, is not continuous. In words, this means that (as we see in this example) curves can be arbitrarily close, yet have very different arc length.

We can dig a little deeper: as you learned in Calculus II, the arc length of a function $y = f(x)$ for $x \in [a, b]$ is

$$\int_a^b \sqrt{(dx)^2 + (dy)^2} = \int_a^b \sqrt{1 + f'(x)^2} dx.$$

Observe that this involves f' . Try to see why the operator D that sends f to f' is not continuous with respect to the distance induced by the norm

$$\|f\| = \left(\int_a^b f(x)^2 dx \right)^{1/2}.$$

In words, two functions f and g can be arbitrarily close, yet have very different derivatives f' and g' .

That's a huge issue in **functional analysis**, which is the generalization of linear algebra to infinite dimensional spaces (like the space of all differentiable functions). The linear operators (“matrices”) on these spaces frequently fail to be continuous.

How little we actually know!

Q: How fast can we solve N linear equations in N unknowns?

Estimated cost of Gaussian elimination:

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

- to create the zeros below the first pivot:
 \implies on the order of N^2 operations
- if there are N pivots total:
 \implies on the order of $N \cdot N^2 = N^3$ operations

- A more careful count places the cost at $\sim \frac{1}{3}N^3$ operations.
- For large N , it is only the N^3 that matters.

It says that if $N \rightarrow 10N$ then we have to work 1000 times as hard.

That's not optimal! We can do better than Gaussian elimination:

- Strassen algorithm (1969): $N^{\log_2 7} = N^{2.807}$
- Coppersmith–Winograd algorithm (1990): $N^{2.375}$
- ... Stothers–Williams–Le Gall (2014): $N^{2.373}$ (If $N \rightarrow 10N$ then we have to work 229 times as hard.)

Is $N^{2+(\text{a tiny bit})}$ possible? **We don't know!** (People increasingly suspect so.) (Better than N^2 is impossible; why?)

Comment. The above algorithms actually are for computing matrix products. It can be shown that, if $M(N)$ is the cost for multiplying two $N \times N$ matrices, then $N \times N$ systems can also be solved for cost on the order of $M(N)$. In other words, we don't even know how costly it is to multiply two matrices.

Good news for applications:

- Matrices typically have lots of structure and zeros
which makes solving so much faster.