

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 187. 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 188. 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:

- $\underbrace{D(f(x) + g(x))}_{(f(x)+g(x))'} = \underbrace{D(f(x))}_{f'(x)} + \underbrace{D(g(x))}_{g'(x)}$
- $\underbrace{D(cf(x))}_{(cf(x))'} = \underbrace{cD(f(x))}_{cf'(x)}$

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 189. 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 190. 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 191. 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 192. (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:
 \Rightarrow on the order of N^2 operations
- if there are N pivots total:
 \Rightarrow 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.

Just for fun and curiosity!

Recall that we introduced the **dimension** of a vector space as the number of vectors in a/any basis. In Calculus, on the other hand, you learn about curves (1-dimensional), surfaces (2-dimensional) and solids (3-dimensional).

The reason that Linear Algebra is relevant for curved objects like surfaces is that locally these (typically) do look flat (like a plane), so that our tools apply at least locally.

What should a 1.5 dimensional thing look like?


Something between a curve and a surface...

(Note that our linear algebra approach to dimension is not helpful.)







Here is a candidate.



Continuing this process, results in the **Koch snowflake**, a **fractal**:



- Its perimeter is infinite!
Why? At each iteration, the perimeter gets multiplied by $4/3$.
- The table below indicates that its boundary has dimension $\log_3(4) \approx 1.262!!$

the effect of zooming in by a factor of 3			
		$\times 3$	$d = 1 = \log_3(3)$
		$\times 9$	$d = 2 = \log_3(9)$
		$\times 4$	$d = \log_3(4) \approx 1.262$

Does this have any practical relevance? Surprisingly, yes!

Have you ever wondered why perimeters of countries are missing from wikipedia? Or, why the coastline of the UK is listed as 11,000 miles by the UK mapping authority but 7,700 miles by the CIA Factbook?

Some of the fun can be found at: https://en.wikipedia.org/wiki/Coastline_paradox