OK.DOPE · DOPE.ARCHITECTURE · FOR JESS

Dimensions, for Jess

Web version (with an animated tesseract): https://dope-dimensions.pages.dev

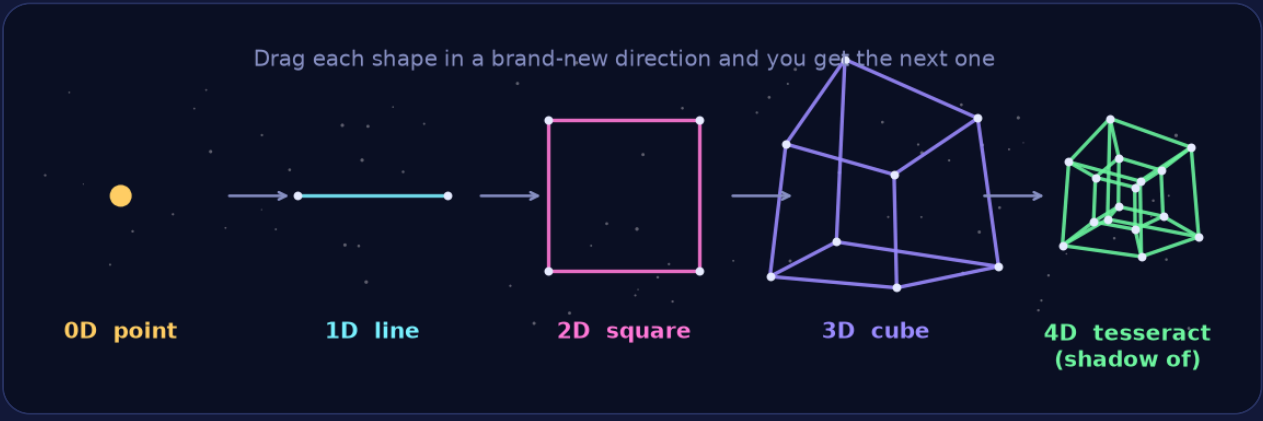
A picture-first walk from a single dot, up through our world and time, out to the edge of imagination, and then into the thousands of dimensions an AI thinks in.

The one idea

What is a dimension?

A dimension is **a direction you can move in that is completely separate from the others**. Another way to say it: it's **how many numbers you need to say exactly where something is.**

Imagine someone left a vinyl record for you at a hotel. To find it, you need an address. Each dimension adds one more number to that address: how far along the shelf, which spot in the lobby, which floor, and what time it's there.

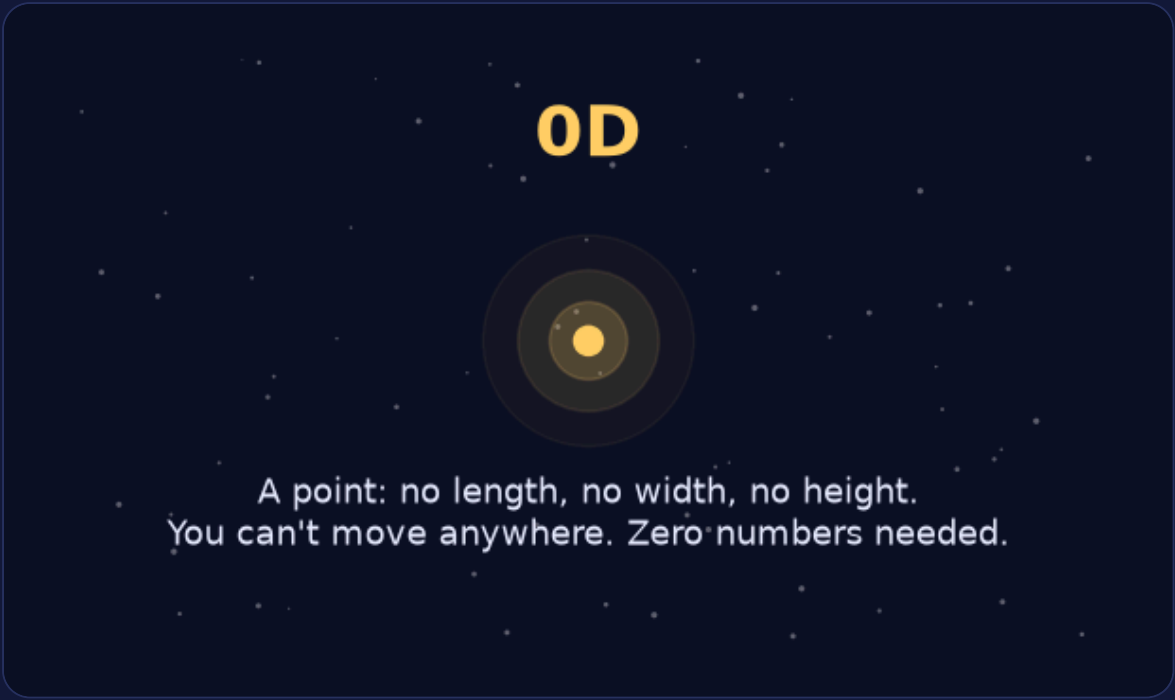


Every shape is the one before it, dragged in a brand-new direction.

> **Keep this in your pocket:** 1 number = 1D, 2 numbers = 2D, 3 numbers = 3D. Everything below builds on that.

0D

A point: nowhere to go

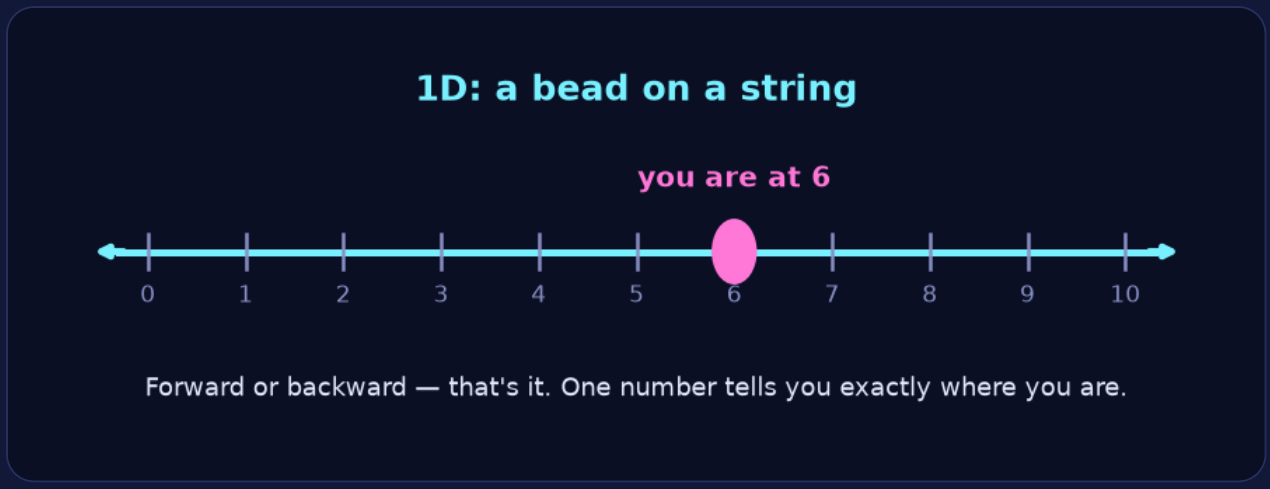


A point has no length, no width, no height. If you lived inside a point, you couldn't move at all, and you'd need **zero numbers** to say where you are, because there's only one place to be.

It sounds like nothing, but it's the seed of everything else.

1D

A line: one number



A bead on a string can only slide forward or back.

Drag the point in one direction and you get a line. Now there's **one** way to move (forward and backward are the same direction, just opposite ways), so **one number** says exactly where you are.

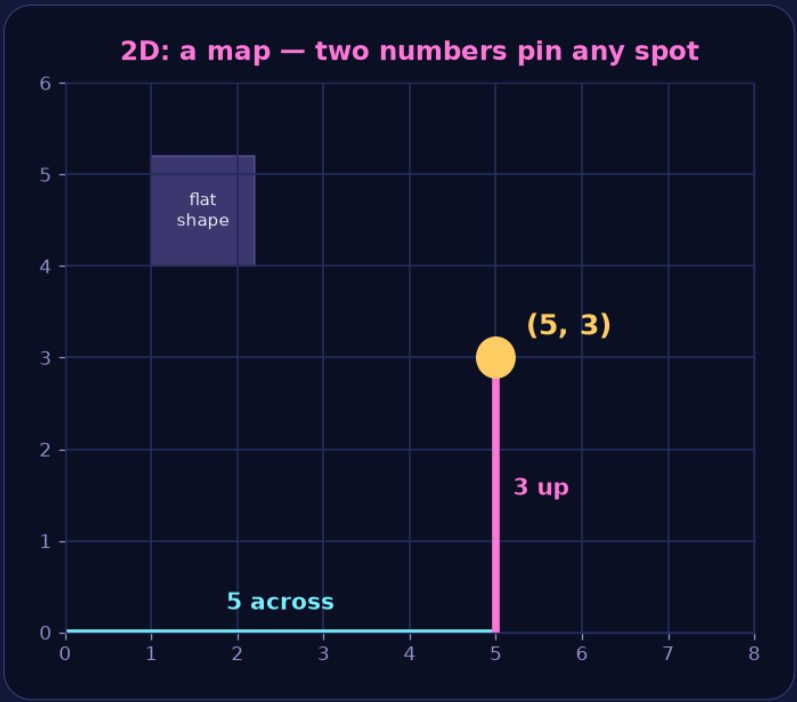
1D things you already use

- The **progress bar** on a YouTube video: "2:41" is the only number you need.
- A **thermometer**, a train track, a ruler, the volume slider.
- A shelf of records: "the 6th one from the left."

> **Try it:** stand in a narrow hallway. You can only go toward one end or the other. A creature living on a line can never pass another creature. It's stuck behind it forever.

2D

A flat world: two numbers



Any spot on a flat map is just two numbers: across and up.

Drag the line sideways and it sweeps out a flat sheet. Now there are **two separate directions**: across and up (or east and north). You can change one without touching the other, so you need **two numbers**, like **(5, 3)**.

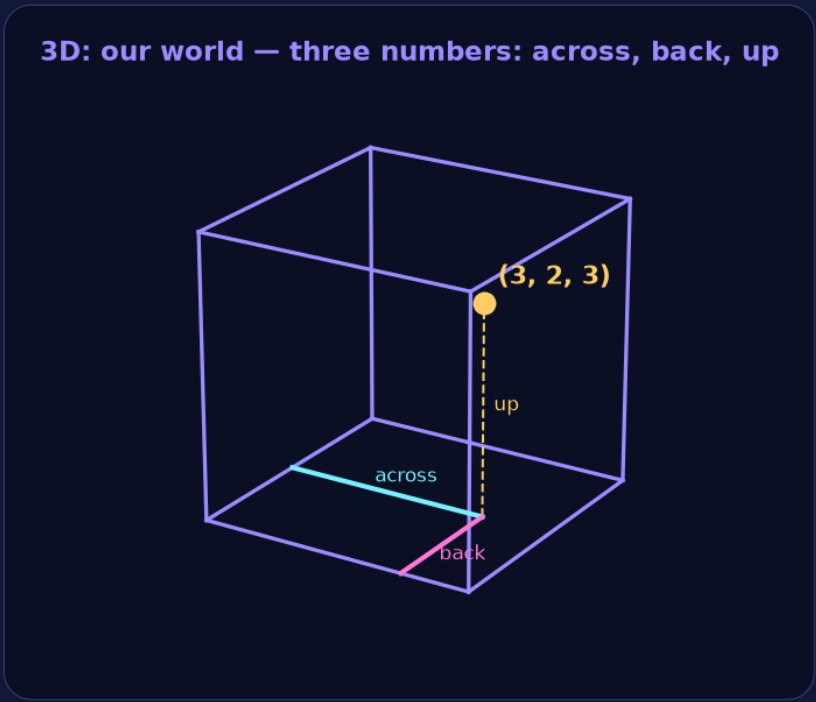
2D things you already use

- Your **phone screen**: every tap is an across-and-down position.
- A **map** or GPS pin (latitude and longitude). A **chessboard** works the same way: "e4" means column e, row 4.
- A **shadow**: a flat copy of a 3D thing.

In 1884, Edwin Abbott wrote a little book called *Flatland* about a square who lives on a flat sheet. He can't imagine "up." It isn't left, right, forward or back, so to him it simply doesn't exist. Keep him in mind; he's how we'll sneak up on 4D.

3D

Our world: three numbers



Across, back, up: three numbers pin anything in a room.

Lift the flat square straight up and it sweeps out a cube. That's the world we live in: **across, back, and up**. The record is 3 steps across the lobby, 2 steps back, on floor 3: **(3, 2, 3)**.

A cool fact about your own eyes

Each eye only sees a flat, 2D picture, like a photo. Your brain compares the two slightly different photos and *builds* the third dimension, depth, for you. You're already doing dimension-math every second.

> **Try it:** hold a finger up and close one eye, then the other. It jumps. That jump is your brain's raw material for 3D.

4D

4D: two different ideas share the name

This is where people get mixed up, because "the fourth dimension" means two different things. Both are real ideas, but they're **not the same thing**.

Idea A · Time

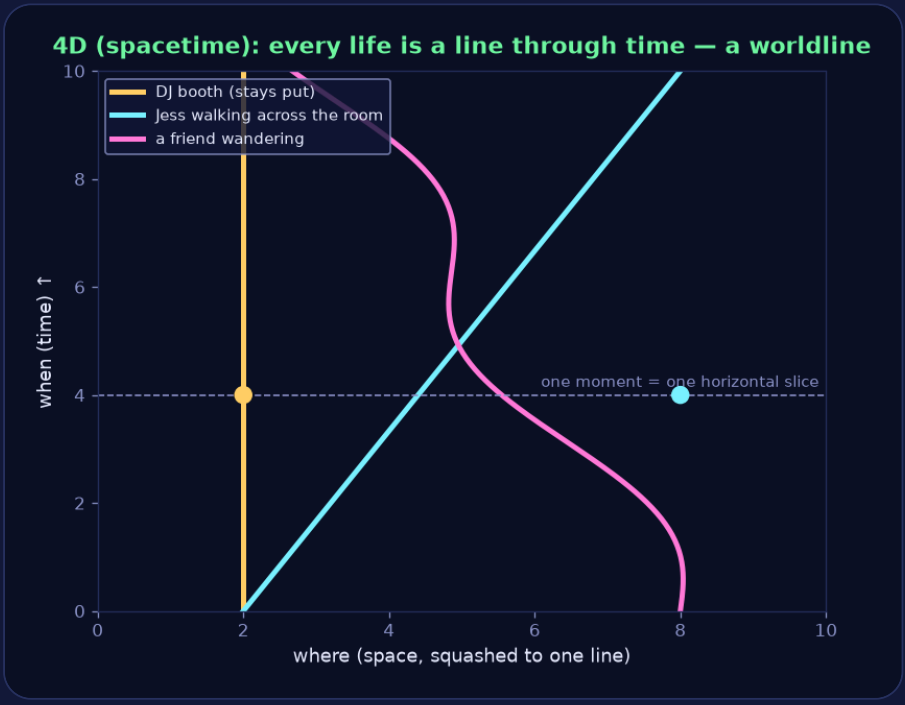
Physics treats time as a 4th number you need to describe an *event*. Not just where, but **when**. Einstein's relativity weaves space and time together into **spacetime**.

Idea B · A 4th direction of space

A mathematical "extra direction," at right angles to across, back and up. We can do the math perfectly, but we can't point to it. Its famous shape is the **tesseract**.

Idea A: time, the 4th number

The DJ plays a record at 8 p.m. and again at 10 p.m. Same booth, same three space numbers, but two different events. To tell them apart you need a fourth number: the time. **(3, 2, 3, 8 p.m.)** and **(3, 2, 3, 10 p.m.)**.



Squash space down to one line, let time run upward, and every person or thing becomes a line: its **worldline**.

A **worldline** is the path of your whole life drawn through spacetime. Something that stays put is a straight line going up. Something moving leans sideways. Meeting someone means your worldlines cross: same place, same time.

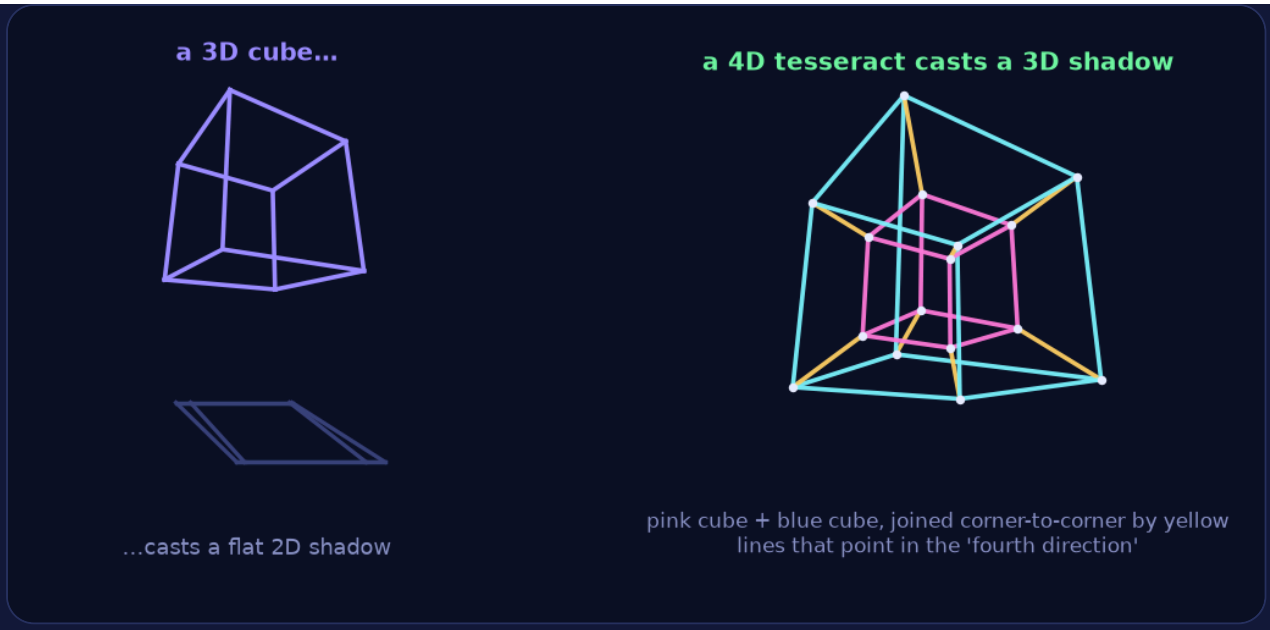
Established: spacetime is well-tested physics. GPS satellites have to correct for relativity's effects on time, or your map pin would drift by kilometres a day.

Idea B: a 4th direction of space, and the tesseract

Follow the pattern we've been using: drag a point to get a line, a line to get a square, a square to get a cube. Now drag the cube in a direction *at right angles to all three*. You get a **tesseract** (also called a hypercube).

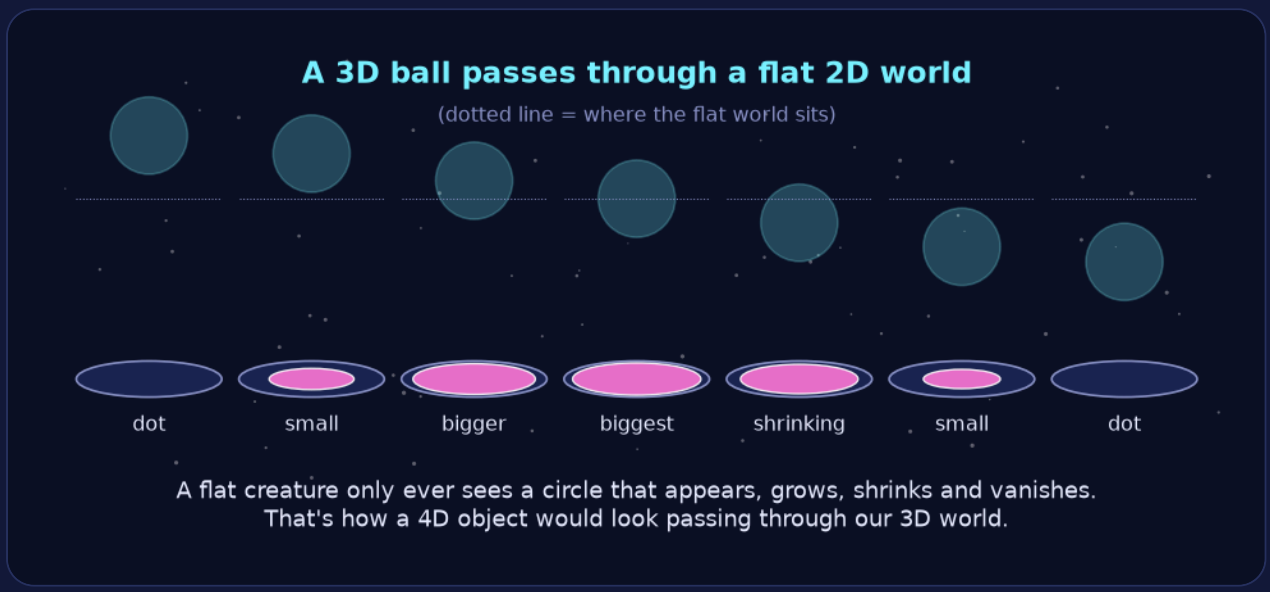
Shape	Dimensions	Corners	Edges
point	0	1	0
line	1	2	1
square	2	4	4
cube	3	8	12
tesseract	**4**	**16**	**32**

Corners double each time. We can't *see* a tesseract, but we can count it perfectly. That's the magic of math: it walks where our eyes can't.



We can only draw a tesseract's *shadow*, the same way a cube leaves a flat shadow on the floor.

The Flatland trick: how a 4D thing would look to us



A ball dropping through a flat world.

If a ball fell through Flatland, the flat square would see a dot appear from nowhere, grow into a circle, shrink, and vanish. He'd never see "a ball," only slices of it.

If a 4D ball passed through *our* world, we'd see the same thing one step up: a tiny sphere appearing out of thin air, swelling, shrinking, and disappearing. That's the best intuition anyone has for spatial 4D.

Why 1–4 matter most: these four are the ones humans genuinely live in or can picture. Three of space we move through, and time we move along. Everything past here we reach with math and imagination, not experience.

5D → 10D

Beyond 4: where facts end and imagination begins

Past 4D, it's important to separate **what physics actually says** from **a popular way of picturing it**. Both are interesting, but only one is science.

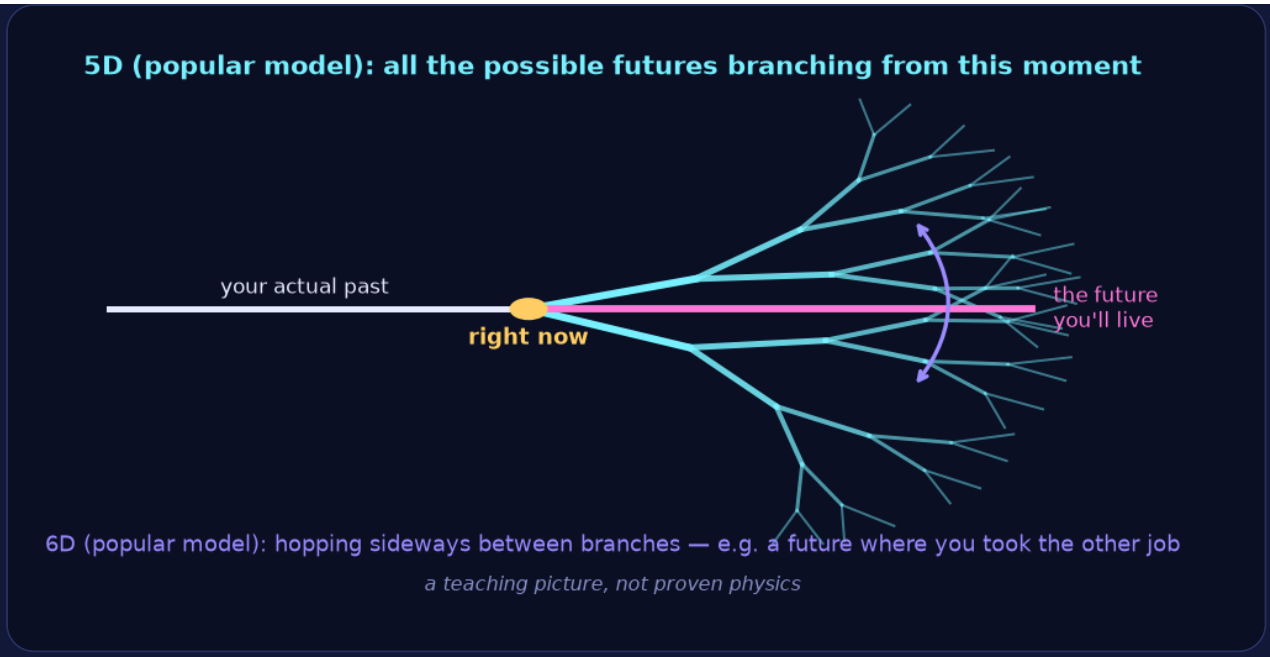
What physics actually says:

- Everything we've measured fits **3 dimensions of space + 1 of time**.
- In the 1920s, Kaluza and Klein tried a 5th dimension to unite gravity and electromagnetism.
- **Superstring theory** needs **10** dimensions (9 space + 1 time); **M-theory** uses **11**. The extra ones would be curled up far too small to notice.
- These are serious, beautiful mathematical ideas, but **no experiment has confirmed extra dimensions yet**. The Large Hadron Collider looked for signs and hasn't found them.



How an extra dimension could hide: a cable looks 1D from far away, but up close an ant can also walk *around* it.

The popular "10 dimensions" picture (a teaching model, not physics): there's a well-known explanation, made famous by Rob Bryanton's 2006 *Imagining the Tenth Dimension*, that stacks dimensions as "bigger and bigger sets of possibilities." It's a fun mental ladder. Physicists don't use it, and string theory's extra dimensions don't mean this. Enjoy it as a thinking toy.



The popular model's 5D and 6D: all the futures that could branch from right now, and hopping between them.

In the popular model...	imagine
5D	All the possible futures branching out from this exact moment.
6D	Moving sideways between those branches, like the version of tonight where you stayed in.
7D	Every possible future from *every possible beginning* of the universe.
8D	Branching between all of those whole sets of possibilities.
9D	All of that, for universes with completely *different laws of physics*.
10D	Everything that could ever possibly be, all at once.



Green = established. Purple = popular teaching model.

The edge of human comprehension

By 9D and 10D, even the teaching model is describing "all possible universes." That's about as far as a human mind can stretch with pictures and stories. Past this point, people stop *imagining* and just *calculate*.

Which is exactly where computers feel at home.

AI · the super-awesomeness

How an AI uses thousands of dimensions

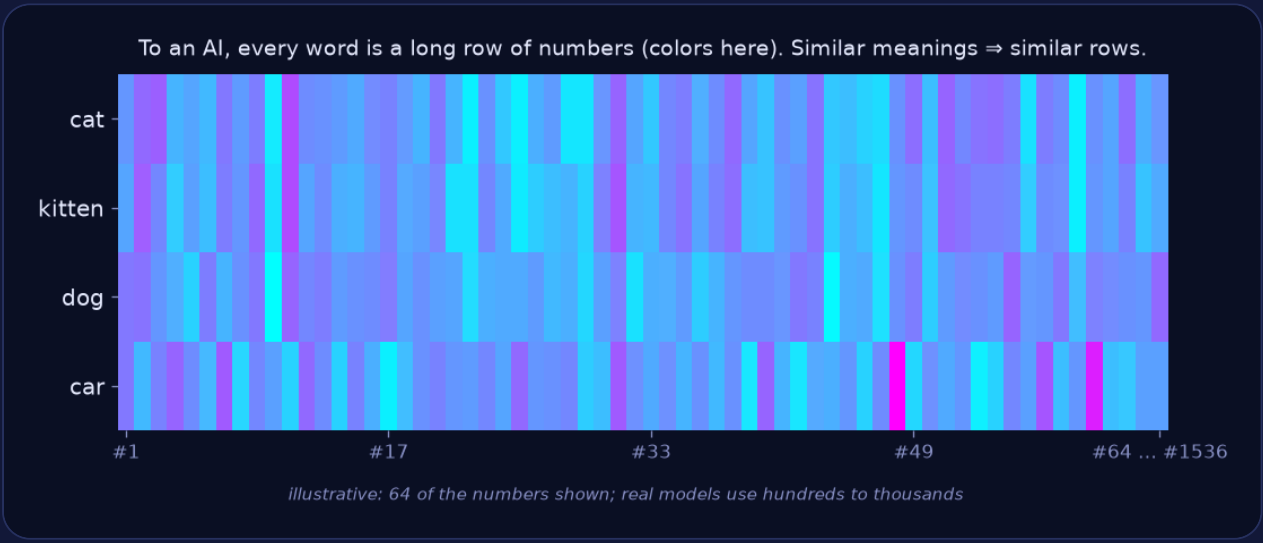
Here's the twist: AI uses far more dimensions than 10. Modern models work with **hundreds to thousands**, and it's not science fiction. It's ordinary math running on ordinary chips.

The key point: these aren't dimensions of *space*. Nobody's floating through 1,536-dimensional room. They're **dimensions of meaning**: each one is just another number in a long list that describes something.

Step 1: a word becomes a long list of numbers

Remember the hotel address? 3 numbers found a spot in a room. An AI gives every word, sentence or song description an "address" too, just a much longer one, called an **embedding**:

```
"Virtual Insanity" → [0.12, −0.38, 0.07, 0.91, −0.22, … 1,536 numbers]
```

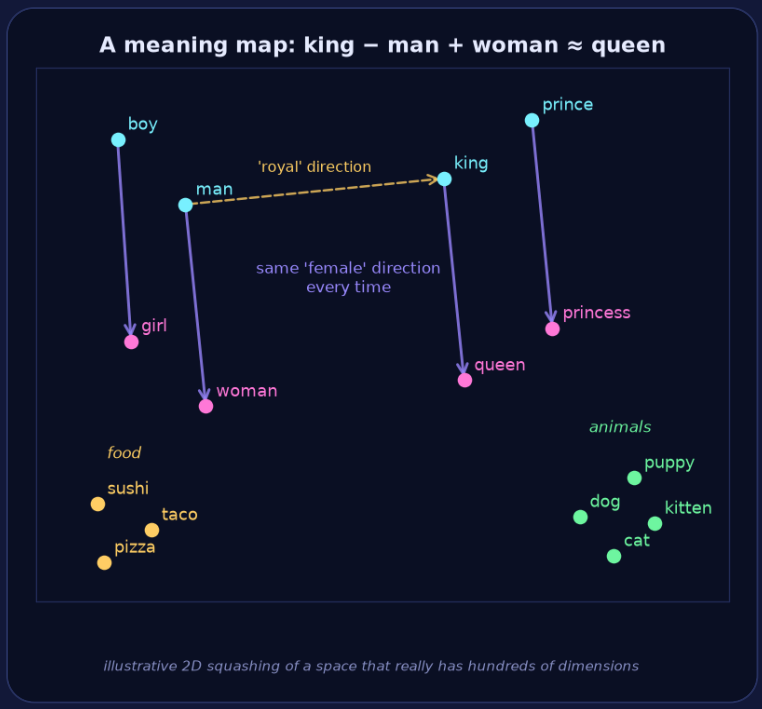


Similar meanings get similar rows of numbers. (Illustration.)

Step 2: close together = similar meaning

The AI learns those numbers by reading huge amounts of text. Words used in similar ways end up with similar addresses, so they land **near each other** in the space. "Cat" lands near "kitten." "Pizza" lands near "taco." "Leaky faucet" lands nowhere near "late-night DJ set."

Step 3: directions mean things



A squashed 2D picture of a space that really has hundreds of dimensions. (Illustration.)

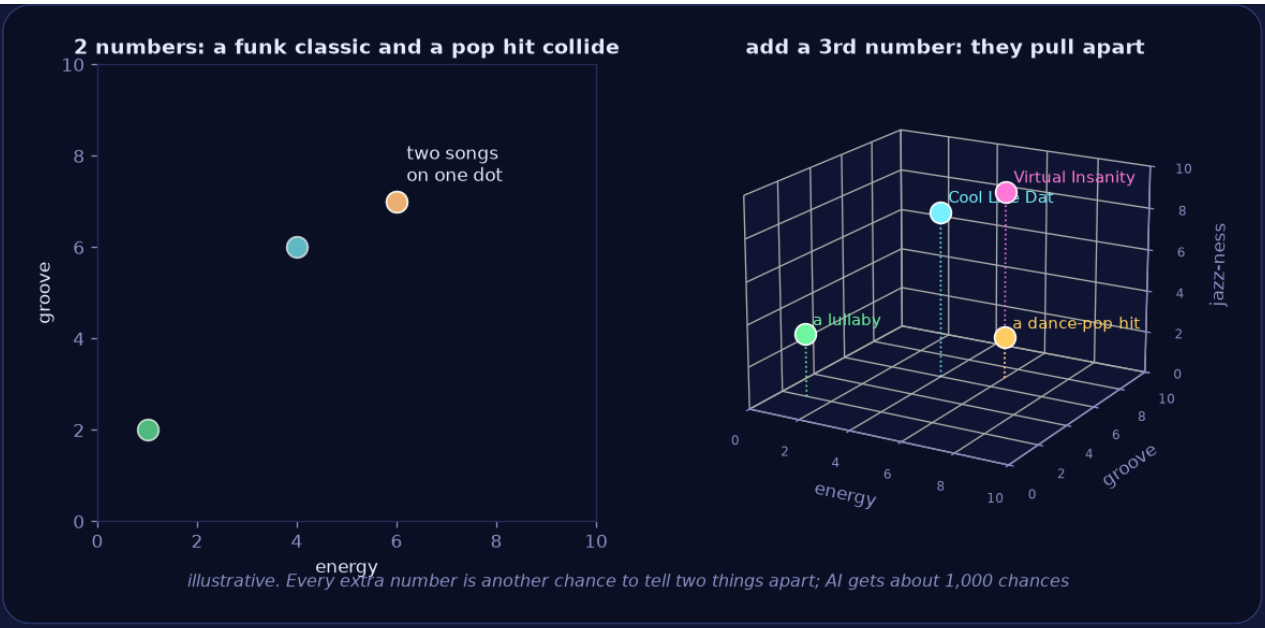
The famous discovery (word2vec, 2013): the step from **man** → **woman** points in nearly the same direction as **king** → **queen**. So you can do arithmetic with meaning:

king − man + woman ≈ queen

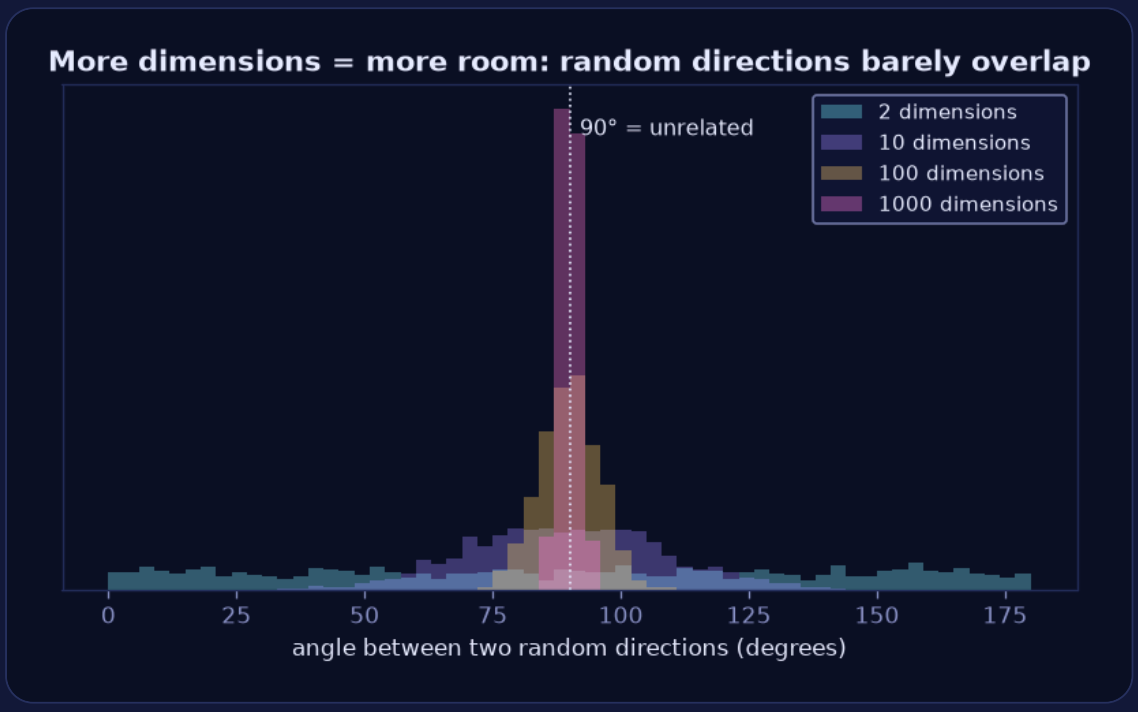
It's approximate. It works for some relationships and fails for others, but it shows that **directions in the space carry meaning**, like "more royal," "more female," "more upbeat," "more 90s."

Step 4: why so many dimensions?

Think about everything that makes a song a song: tempo, mood, era, groove, vocals, language, instruments, how it makes you feel at 2 a.m. Three numbers can't hold all that. A thousand can.

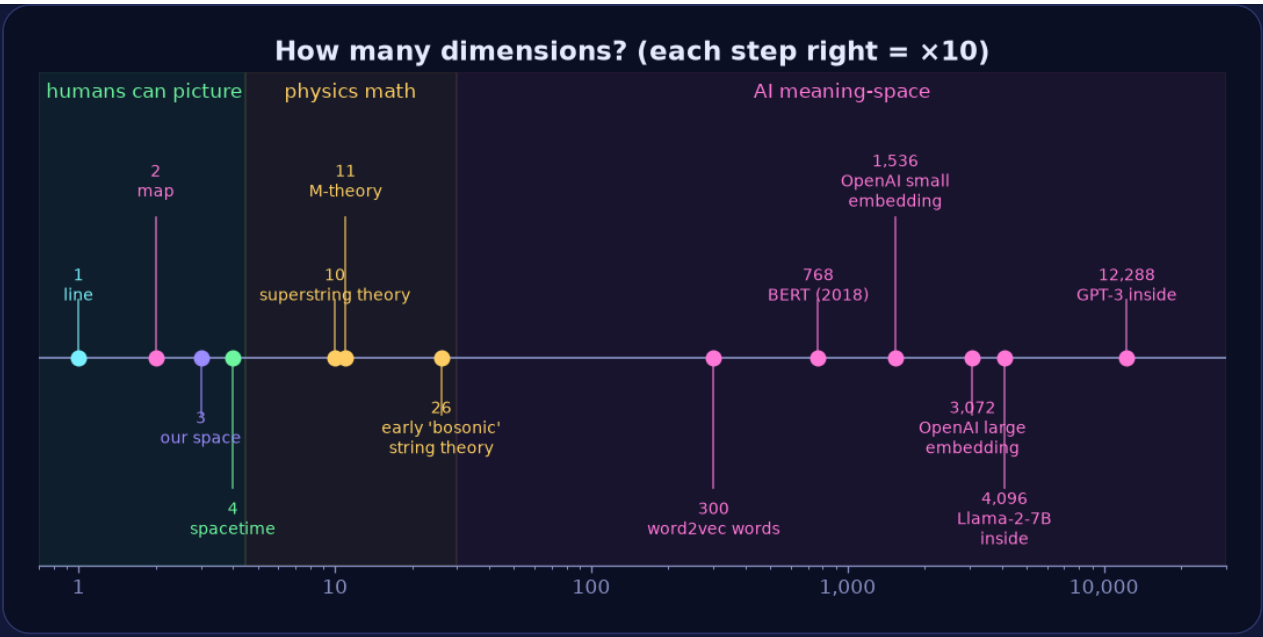


Squash things into too few numbers and different things collide. Every extra dimension is another way to tell them apart. (Illustration.)



In 2D, random directions overlap all the time. In 1,000D, almost every pair of directions is close to 90°, unrelated. There's room for a huge number of separate meanings.

That's the super-power: high-dimensional space has **enormous room**. Thousands of different shades of meaning can each get their own direction without stepping on each other. So when you ask "find me songs that feel like Virtual Insanity," the AI turns your request into an address and looks for the nearest neighbours. Same idea as finding the record in the hotel, just with 1,536 numbers instead of 3.



From a line to an AI's inner workings. Each step right is ten times more.

System	Dimensions
word2vec word vectors (2013)	300 (typical)
BERT-base (2018)	768
OpenAI text-embedding-3-small	1,536
OpenAI text-embedding-3-large	3,072
Llama-2-7B, internal "thought" vectors	4,096
GPT-3 (175B), internal vectors	12,288

Honest fine print: the AI doesn't "see" 1,536-D space any more than we do. It just does arithmetic on long lists, very fast. Individual numbers usually don't have neat labels like "sadness = #412"; meaning is spread across many of them at once. And close in the space means "used in similar ways," which is often, but not always, the same as "means the same thing."

Recap

The whole trip on one card

	What it is	Numbers to find you	Status
0D	a point	0	math fact
1D	a line · a video's progress bar	1	math fact
2D	a flat sheet · your screen · a map	2	math fact
3D	our space	3	lived experience
4D	spacetime (time) *or* a 4th spatial direction (tesseract)	4	time: tested physics · spatial: math
5–10D	string-theory extra dimensions · or the popular "possibilities" ladder	5–11	unconfirmed theory · teaching model
768–12,288D	AI meaning-space (embeddings)	hundreds to thousands	real, running right now

Humans live in 3, travel through a 4th, can imagine up to about 10, and built machines that think in thousands. Pretty dope.

Sources

Where these facts come from

- Edwin A. Abbott, *Flatland* (1884), free at Project Gutenberg
- Spacetime & worldlines: Einstein Online (Max Planck Institute for Gravitational Physics)
- GPS & relativity: R. Pogge, Ohio State
- Tesseract: Wolfram MathWorld
- Extra dimensions & the LHC: CERN
- Popular 10-D model: Rob Bryanton, *Imagining the Tenth Dimension* (2006), labelled here as a teaching model
- king − man + woman: Mikolov et al., 2013 (word2vec)
- BERT-base 768: Devlin et al., 2018 · OpenAI 1,536 / 3,072: embeddings guide · Llama-2-7B 4,096: Touvron et al., 2023 · GPT-3 12,288: Brown et al., 2020

Provenance

Initial prompts

This page was built from Monte's request, reproduced verbatim:

> ok.dope.challenge: build this better with images and dope.architecture as a standalone site(pdf and initial prompts included): write an explainer for multiple dimensions and how they work. Start with 1d and go through the known dimensions and then explain how ai can use many more dimensions than humans. Make a pdf with images along with the dimensions so Jessica can understand it. Hopefully. And You update that a little bit more with like, um, some additional graphics that actually explore that a little bit more so that you can, um, even more visual in that sense. And then a little bit more of one, two, three, four dimensions. The first four becoming the most important because those are the furthest that humans understand from my understanding. And then going on into the ninth dimension, which is like the furthest of human comprehension. And then go into your super awesomeness. Does that make sense? I'm pretty sure it does

Built on top of an earlier 8-page draft, "Dimensions, illustrated." All graphics were drawn from scratch with code (matplotlib) for this page. Embedding pictures are illustrations, not real model output.

---

sealed · dope.architecture · <ok.dope>

PDF edition · fingerprints for this PDF and the web page live in `SEAL.md` at https://dope-dimensions.pages.dev/SEAL.md

(a file can't contain its own fingerprint)