MATH ENCOUNTERS

LOST *in* SPACE:

how DATA *and* INFORMATION *are governed by* HIGH-DIMENSIONAL GEOMETRY

Speaker: Henry Cohn

Higher dimensions:

how can we understand them, and why should we?

Today we'll see how

big data

is described by

high-dimensional geometry.

What's remarkable is that these don't sound more than superficially similar.

Let's begin by thinking about an apparently unrelated problem.

The sphere packing problem

How densely can we pack identical spheres into space?

Not allowed to overlap (but can be tangent).

Density = fraction of space filled by the spheres.

Why should we care?

The densest packing is pretty obvious.
It's not difficult to stack cannonballs or oranges.

It's *profoundly difficult* to prove (Hales 2005, Hales et al. 2017).
But why should anyone but mathematicians care?

One answer is that it's a toy model for:

- Granular materials.
- Packing more complicated shapes into containers.

Sphere packing is a first step towards these more complex problems.

Today we'll examine a deeper reason.

Varying the dimension

What if we didn't work in three-dimensional space?

The two-dimensional analogue is packing circles in the plane.

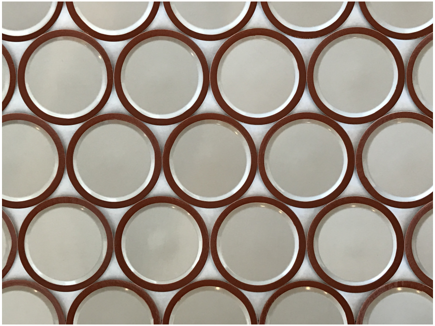

Still tricky to prove, but not nearly as difficult (Thue 1892).

What about one dimension? What's a one-dimensional sphere?

Spheres in different dimensions

Sphere centered at x with radius r means the points at distance r from x.

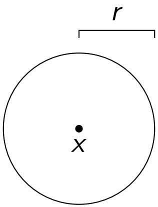

Ordinary sphere in three dimensions, circle in two dimensions.

Just two points in one dimension:

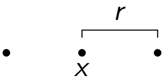

The inside of a one-dimensional sphere is an interval.

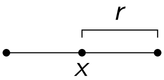

One-dimensional sphere packing is boring:

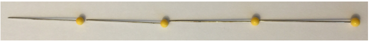

(density $= 1$)

Two-dimensional sphere packing is prettier and more interesting:

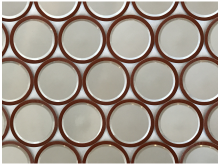

(density ≈ 0.91)

Three dimensions strains human ability to prove:

(density ≈ 0.74)

What about *four dimensions*? Is that just crazy?

Some history

Thomas Harriot (1560–1621)

Mathematical assistant to Sir Walter Raleigh.

A Brief and True Report of the New Found Land of Virginia (1588)

First to study the sphere packing problem.

Claude Shannon (1916–2001)

Developed information theory.

A Mathematical Theory of Communication (1948)

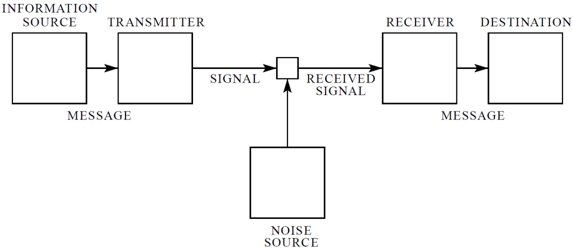

Practical importance of sphere packing in higher dimensions!
We'll return to this later.

What are higher dimensions?

Anything you can describe with multiple coordinates.

n coordinates $= n$ dimensions

2 is a point on the number line:

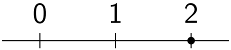

$(2, 3)$ is a point in the plane:

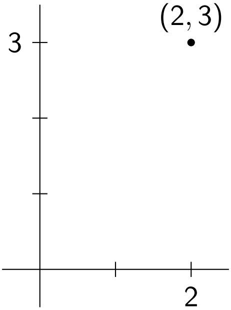

$(2, 3)$ is a point in the plane:

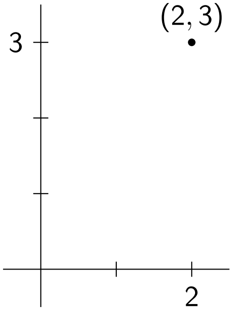

$(2, 3, 5)$ is a point in three dimensions. Like $(2, 3)$ but 5 units up out of the plane.

$(2, 3, 5, 7)$ is a point in four dimensions. Like $(2, 3, 5)$ but 7 units in some entirely new direction (!?).

Is the fourth dimension time?

This question makes no more sense than debating whether the second dimension is width.

Dimensions are mathematical abstractions and don't come with built-in labels saying "this is time."

The fourth dimension can represent time if we want it to, or *anything else we want it to represent*.

Mathematics = freedom.

Four-dimensional space-time is an important concept in physics, but it's just one application of the fourth dimension.

(Similarly, don't think about spatial shortcuts or parallel universes.)

Visualization

So how can we visualize higher dimensions?

I wish I could visualize a thousand dimensions,
but I don't believe any human being can.

Instead, the easiest path to higher dimensions is via algebra:
higher dimensions just means mathematics with more variables.

Distances and volumes

The distance between (a, b, c, d) and (w, x, y, z) is

$$\sqrt{(a-w)^2 + (b-x)^2 + (c-y)^2 + (d-z)^2}.$$

Just like two or three dimensions, but with an extra coordinate.
(n-dimensional Pythagorean theorem)

4d volume of right-angled $a \times b \times c \times d$ box $=$

product $abcd$ of lengths in each dimension.

Just like area of a rectangle or volume of a 3d box.

Higher dimensions work analogously. Just use all the coordinates.

Why should we measure things this way? We don't have to! Once again, mathematics = freedom.

We could measure distances and volumes however we like, to get *different geometries* that are useful for different purposes.

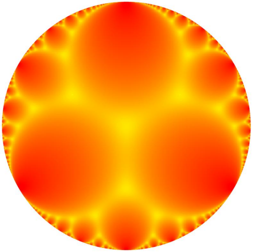

We're going to focus on Euclidean geometry, generalizing high school geometry as directly as possible to higher dimensions.

Applications

Anything you can measure using n numbers is a point in n dimensions. $\mathbb{R}^n = n$-dimensional space ($\mathbb{R}$ = real numbers)

Your height, weight, and age form a point in $\mathbb{R}^3$.

Twenty measurements in an experiment yield a point in $\mathbb{R}^{20}$.

One pixel in an image is described by a point in $\mathbb{R}^3$ (red, green, and blue components). A one-megapixel image is a point in $\mathbb{R}^{3000000}$.

Some climate models have ten billion variables, so their states are points in $\mathbb{R}^{10000000000}$.

High dimensions are everywhere! Low dimensions are anomalous. All data can be described by numbers, so any large collection of data is high dimensional.

Descartes

Descartes developed analytic geometry. Any time you are using two or three variables, you are secretly doing two- or three-dimensional geometry.

Thinking geometrically may help, even when you don't expect it.

More generally, any time you use n variables, you are secretly doing n-dimensional geometry.

Large numbers of variables occur everywhere, and therefore so does high-dimensional geometry.

How can we visualize four dimensions?

Perspective picture of a hypercube.

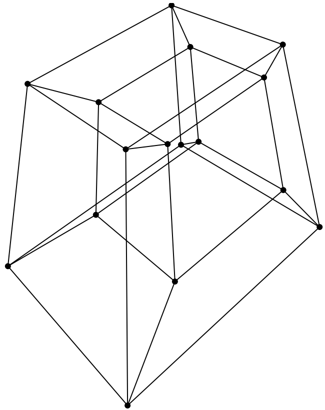

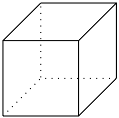

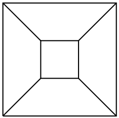

Shadow cast in a lower dimension ("projection").

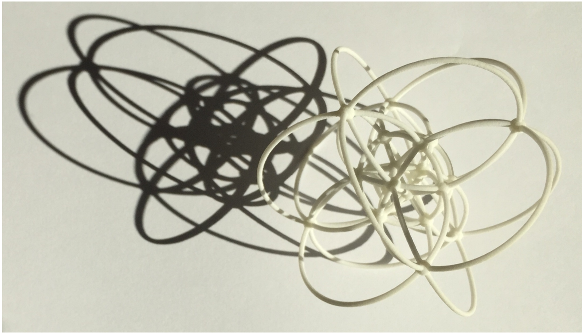

Fundamentally the same as a perspective picture.

Projections can get complicated.

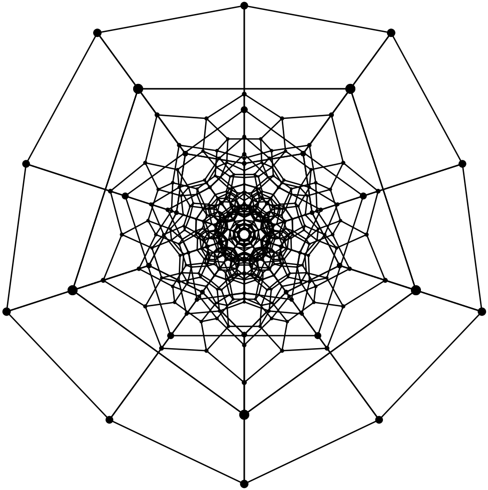

Audience participation time!

Cross sections

Cubes have square cross sections.
Hypercubes have cubic cross sections.

Cubes have other cross sections too:

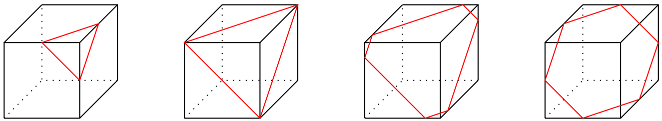

So do hypercubes:

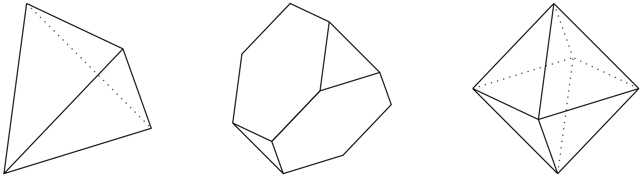


MO MATH

Hinton cubes

Charles Howard Hinton (1853–1907)

A New Era of Thought, Swan Sonnenschein & Co., London, 1888

Model 7. PLUVIUM.

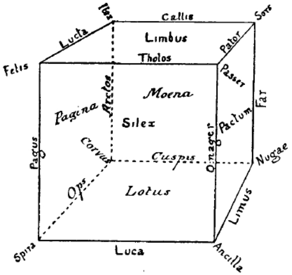

Colours : Pluvium, Dark-stone.

Points : Spira, Silver. Ancilla, Turquoise. Nugæ, Fawn. Corvus, Gold. Felis, Quaker-green. Passer, Peacock-blue. Sors, Dull-purple. Ilex, Light-blue.

Lines : Luca, Leaf-green. Limus, Smoke. Cuspis, Orange. Ops, Stone. Pagus, Dull-blue. Onager, Dark-pink. Far, French-grey. Arctos, Brown. Tholos, Brown-green. Pator, Green-blue. Callis, Reddish. Lucta, Rich-red.

Surfaces : Silex, Burnt-Sienna. Pactum, Yellow-green. Moena, Dark-blue. Pagina, Yellow. Limbus, Ochre. Lotus, Azure.

(Hinton introduced the term "tesseract" and invented the automatic pitching machine.)

Alicia Boole Stott (1860–1940)

On certain series of sections of the regular four-dimensional hypersolids, Verhandelingen der Koninklijke Akademie van Wetenschappen te Amsterdam **7** (1900), no. 3, 1–21.

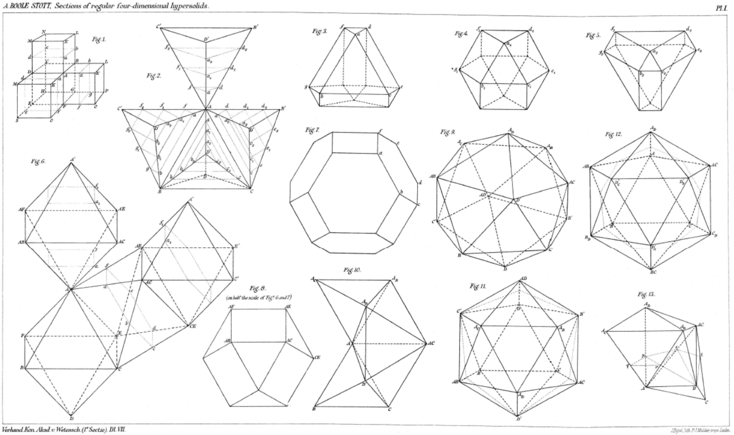

Information theory

Shannon's application of high-dimensional sphere packings.

Represent signals by n-dimensional vectors s.
Each coordinate specifies some measurement.

E.g., radio signals, with
coordinates = amplitudes at different frequencies.

In most applications, n will be large. Often hundreds or thousands.
(No reason to limit ourselves to making just a few measurements!)

Send stream of signals over this channel.

Noise in the communication channel

The key difficulty in communication is noise.

Send signal s in $\mathbb{R}^n$; receive r on other end.
Noise means generally $r \neq s$.

The channel has some *noise level* ε, and r is almost always within distance ε of s.

Imagine an *error sphere* of radius ε about each sent signal, showing how it could be received.

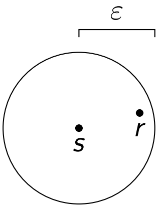

How can we communicate without error?

Agree ahead of time on a restricted vocabulary of signals.

If s_1 and s_2 get too close, received signals could get confused:

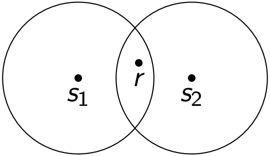

Did r come from s_1 or s_2? Therefore, keep all signals in S at least 2ε apart, so the error spheres don't overlap:

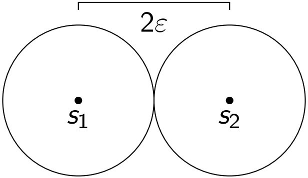

This is sphere packing!

The error spheres should form a sphere packing. This is called an *error-correcting code*.

For rapid communication, want as large a vocabulary as possible. I.e., to use space efficiently, want to maximize the packing density.

Rapid, error-free communication requires a dense sphere packing. Real-world channels correspond to high dimensions.

Of course some channels require more elaborate noise models, but sphere packing is the most fundamental case.

What is known?

Each dimension seems to behave differently.

Good constructions are known for low dimensions.

No idea what the best high-dimensional packings look like,
or even whether they are crystalline or amorphous.

Upper/lower density bounds in general.

Bounds are very far apart:

For $n = 36$, differ by a *multiplicative factor of 52*.

This factor grows exponentially as $n \to \infty$.

Tremendous wasted space outside the spheres?

Packing in high dimensions

On a scale from one to infinity, a million is small, but we know almost nothing about sphere packing in a million dimensions.

Simple lower bound: can achieve density at least 2^{-n} in $\mathbb{R}^n$.

How to get density at least 2^{-n}

Consider any packing in $\mathbb{R}^n$ with spheres of radius r, such that no further spheres can be added without overlap.

No point in $\mathbb{R}^n$ can be $2r$ units away from all sphere centers. I.e., radius $2r$ spheres cover space completely.

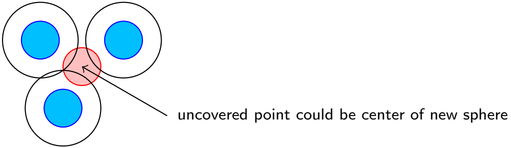

Doubling the radius multiplies the volume by 2^n.

Thus, the radius r packing has density at least 2^{-n} (since the radius $2r$ packing covers all of space). Q.E.D.

This is very nearly all we know!

Nonconstructive proof

This proof shows that the best packing in $\mathbb{R}^n$ must have density at least 2^{-n}.

But it didn't describe where to put the spheres in an actual packing. It's a *nonconstructive proof*.

In fact, nobody has any idea how to build such a good packing. Our constructions are all much, much worse.

Philosophical conundrum:
Are high-dimensional sphere packings intrinsically undescribable?
Or do we just lack the tools to describe them?

Packing in high dimensions (large n)

We just showed:
density at least 2^{-n}.

Hermann Minkowski (1905):
at least $2 \cdot 2^{-n}$.

...[many further research papers]...

Keith Ball (1992):
at least $2n \cdot 2^{-n}$.

Stephanie Vance (2011):
at least $\frac{6}{e} n \cdot 2^{-n}$ when n is divisible by 4.

Akshay Venkatesh (2013):
at least $\frac{e^{-\gamma}}{2} n \log \log n \cdot 2^{-n}$ for certain dimensions.

For comparison, the best upper bound known is $2^{-0.599n}$, due to Grigorii Kabatiansky and Vladimir Levenshtein (1978).

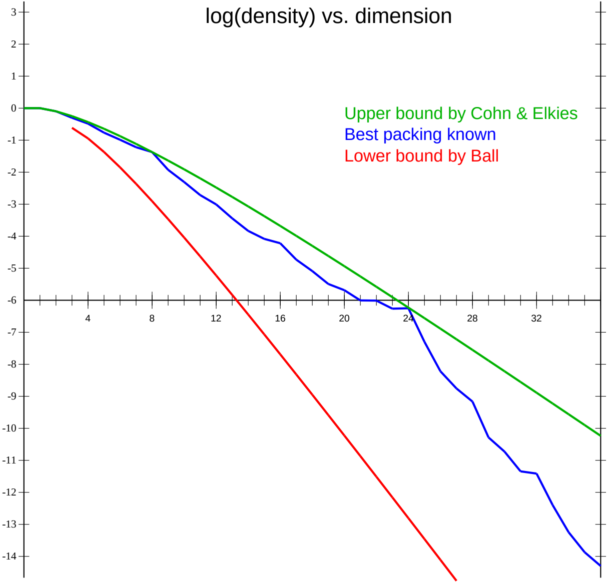

log(density) vs. dimension
Upper bound by Cohn & Elkies
Best packing known
Lower bound by Ball
3
2
1
0
-1
-2
-3
-4
-5
-6
-7
-8
-9
-10
-11
-12
-13
-14
4
8
12
16
20
24
28
32

The most remarkable packings

Certain dimensions have amazing packings.

$\mathbb{R}^8$: E_8 root lattice
$\mathbb{R}^{24}$: Leech lattice [named after John Leech (1926–1992)]

Extremely symmetrical and dense packings of spheres.

Connected with many areas in mathematics, physics
(such as string theory, modular forms, finite simple groups).

What makes these cases work out so well?

Sphere packing problem is solved for only two cases with dimension greater than 3. Breakthrough in March, 2016:

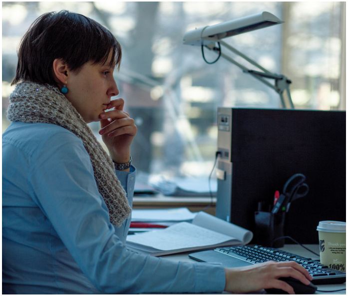

Theorem (Viazovska).
The E_8 root lattice achieves the greatest possible sphere packing density in $\mathbb{R}^8$, namely $\pi^4/384 \approx 25.36\%$.

Twenty-four dimensions builds on her techniques:

Theorem (Cohn, Kumar, Miller, Radchenko, and Viazovska).
The Leech lattice Λ_{24} achieves the greatest possible sphere packing density in $\mathbb{R}^{24}$, namely $\pi^{12}/12! \approx 0.1929\%$.

Why is the sphere packing problem difficult?

Many local maxima for density.

Lots of space.

Complicated geometrical configurations.

No good way to rule out implausible configurations rigorously.

High dimensions are weird.

Visualization, take two

There are many ways to try to visualize four dimensions.
But they become incomprehensible as the dimension increases.

Is there any way we can conceive of a million dimensions?
What could it look like?

Let's start by thinking about cubes, because they are easy to pack.

Even the cubes go crazy

Imagine a unit cube in $\mathbb{R}^n$. Each edge has length 1.

There are 2^n vertices (corners).

Distance between opposite corners is $\sqrt{1^2 + \cdots + 1^2} = \sqrt{n}$.

Imagine $n = 10^6$. How can this be?
Let's imagine connecting the center of a cube to all its vertices.

Even the cubes go crazy

Imagine a unit cube in $\mathbb{R}^n$. Each edge has length 1.

There are 2^n vertices (corners).

Distance between opposite corners is $\sqrt{1^2 + \cdots + 1^2} = \sqrt{n}$.

Imagine $n = 10^6$. How can this be?
Let's imagine connecting the center of a cube to all its vertices.

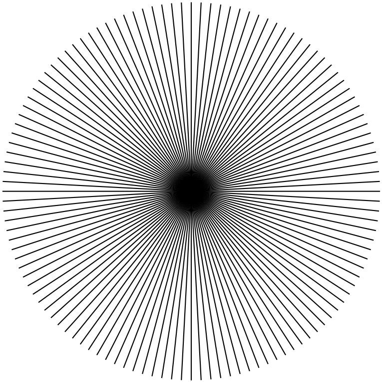

Concentration of volume

Almost all volume in high-dimensional solids is concentrated near the boundary.

Why? Imagine a sphere. Fraction of volume less than 90% of way to boundary is 0.9^n. This tends rapidly to zero as $n \to \infty$. For example, $0.9^{1000} \approx 1.75 \cdot 10^{-46}$.

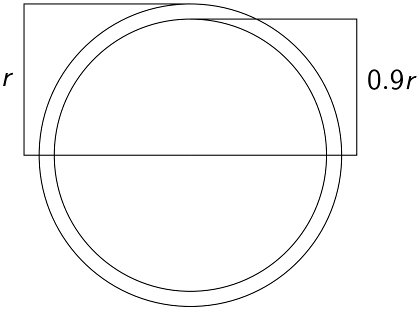

We can replace 0.9 with 0.99, 0.999, or whatever we want.

[Aside for those who know calculus:
As $n \to \infty$ with c fixed, fraction of volume less than $1 - c/n$ of the way to the boundary is $(1 - c/n)^n \to e^{-c}$.]

All low-dimensional intuition must be revised.

Boundaries, where all the interaction takes place, become increasingly important in high dimensions.

Relevant for packing and for analysis of algorithms.

Also the second law of thermodynamics (increase of entropy)!

In summary

High dimensions are deeply weird. Even four dimensions can be puzzling, and a million dimensions are outlandish.

We really don't know what the best sphere packing in a million dimensions looks like. Our estimate of its density might be off by an exponential factor. We don't even know whether it is ordered (like a crystal) or disordered (like a lump of dirt).

But these problems really matter. Every time you use a cell phone, you can thank Shannon and his relationship between information theory and high-dimensional sphere packing.

For more information

For references, links, and an STL file for a 3-d printer, see:

`http://3d.momath.org`

Have fun exploring higher dimensions!