


JOHN RECTOR

Every Direction at *Once*

The geometry underneath
artificial intelligence, and the
future it is building



90°

JOHN RECTOR

Every Direction at *Once*

*The geometry underneath artificial intelligence,
and the future it is building*

Free edition · October 2026 · johnrector.me

© 2026 John Rector. Free to read and share with attribution.

CONTENTS

Before You Begin A Note on What This Book Is

Prologue The Parking Lot

PART ONE The Strangeness

Chapter One Built for Three

Chapter Two All the Orange Is Peel

Chapter Three Everything Is at Right Angles

Chapter Four Everyone Is Equally Far

Chapter Five Every Trap Is Near the Floor

Chapter Six You Only See the Shadow

PART TWO What the Machines Do With It

Chapter Seven A Thousand Dots and One Line

Chapter Eight Meaning Becomes an Address

Chapter Nine The Space Between

PART THREE The Future It Is Building

Chapter Ten Seven Predictions

Chapter Eleven Your Three Dimensions

Claim Ledger What Is Record and What Is Argument

Epilogue The Car Has an Address

Background For Further Reading

BEFORE YOU BEGIN

A Note on What This Book Is

This is a short book about one idea. The idea is that artificial intelligence works for a reason most people have never been told, and that the reason is not size, not data, and not chips. It is geometry.

The machines that now write, translate, search, and reason do all of their work inside spaces with thousands of dimensions. That sounds like a detail. It is the whole story. A space with thousands of dimensions is not a bigger version of the room you are sitting in. It is a different kind of place. It has its own rules, and almost every one of them contradicts something your body taught you before you could talk.

Those rules are old. Mathematicians have known most of them for a century. What is new is that we finally built something that lives there.

I have written about pieces of this on my blog over the past several months. Each post took one strange fact and followed it to a consequence. This book puts the pieces in order. It moves in three parts. The first part is the strangeness: five rules of high-dimensional space, each explained with nothing more than an orange, a box, a hiker, or a parking lot. The second part is why those rules make learning possible. The third part is the future they are already building, stated as predictions you can hold me to.

There are almost no statistics in this book. The few numbers that appear are not evidence. They are the shape of the idea itself, the way the number two is the shape of a pair. You can check every one of them with a pencil, or ask your own AI to check them for you.

You do not need mathematics to read this. You need to be willing to let a few of your instincts be wrong.

PROLOGUE

The Parking Lot

6:40 P.M., A LEVEL YOU DO NOT REMEMBER

You come out of the concert and you cannot find your car. The surface lot was easy, you tell yourself. Rows and columns. You could walk every row in ten minutes. But tonight you parked in the garage, and now there are floors. Five of them. You walk the rows of one floor, then climb the stairs and walk the rows of the next. It takes forty minutes. You find it on the floor you were sure you had already checked.

Everyone has lived some version of that evening. It teaches a lesson so basic we never say it out loud. Every time you add a direction to a search, the search gets harder. Not a little harder. A row is a line. A lot is a grid of lines. A garage is a stack of grids. Each new direction multiplies the places a thing could be.

Now picture a garage that does not stop at floors. Picture one with a thousand independent directions. Not a thousand floors stacked on top of each other, but a thousand different ways to move, each at right angles to all the others. There is no picture for that. You cannot draw it. If you tried to search it the way you searched the garage, one spot at a time, the universe would end before you finished the first corner.

And yet that is exactly where the machines live. Every word an AI reads becomes a point in a space like that. Every image, every question, every customer, every song. The machine does not walk the rows. It does something else, and it finds things anyway. It finds them fast.

The usual explanation is scale. Bigger models, more data, more computers. Scale is real, but it is the answer to how much, not the answer to why. Why does adding dimensions, which should make every search hopeless, make these machines better instead of worse? Why does a space too large to explore turn out to be easy to learn in?


The answer is that high-dimensional space breaks the rules of the garage. It breaks them in specific, learnable ways. Once you know the rules, three things happen. AI stops feeling like magic. It also stops feeling like a trick. And the next ten years become much easier to see coming.

That is what this book is for.

PART ONE

The Strangeness

Five rules your body never learned



EVERY DIRECTION AT ONCE

CHAPTER ONE

Built for Three

MONDAY, ANY KITCHEN

Someone asks you to describe a friend they have never met. You say she is tall, about forty, funny, a little impatient, loves the ocean, hates the phone, and always remembers your birthday. Seven things. You have just placed a person in a seven-dimensional space, and you did not need a single equation to do it.

In 1884 an English schoolmaster named Edwin Abbott published a small novel called *Flatland*. Its narrator is a square. He lives on a plane, a world with length and width and nothing else. One day a sphere passes through his world. The square sees only a circle that appears from nowhere, grows, shrinks, and vanishes. The sphere tries to explain the direction called up. The square cannot picture it. He has no faculty for it. Every intuition he owns was formed on a flat surface, and the new direction is perpendicular to all of them.

We are the square. Our bodies, eyes, and instincts were built for three directions. Everything we believe about space without thinking, about what is near, what is typical, what is in the middle, how hard it is to get lost, was learned in three dimensions and only ever tested there.

That would not matter if nothing important happened anywhere else. But something does now.


A dimension is a question

The word dimension sounds like science fiction. It is not. A dimension is simply an independent way that something can differ. Height is one. Age is another. Whether someone loves the ocean is a third. If you can change

one without forcing a change in another, they are separate dimensions.

Seen this way, you have lived in high-dimensional space your whole life. A customer record with forty fields is a point in forty dimensions. A photograph from your phone is a point in millions of dimensions, one for each brightness value of each tiny dot. A song is a point. A sentence is a point. When people say an AI "represents" a word, they mean it gives that word an address: a long list of numbers, often several thousand of them, each one a position along a separate direction.

The number of dimensions is the number of independent ways a thing can differ. Most things that matter differ in thousands of ways.



None of this is mystical. A row in a spreadsheet is a point. The only thing that changes as the row gets longer is that your intuition stops working. And it does not stop working gradually. It breaks.

Why the square cannot be blamed

It is tempting to think that intuition simply runs out of room, that a thousand dimensions is like three dimensions, only more so. That is the mistake this book is about. High-dimensional space is not more of the same. Its basic facts are different. The middle is empty. Every direction is at right angles to every other. Everyone is the same distance from everyone. Traps disappear. And almost everything that moves, moves in a direction you are not watching.

Those five facts are the next five chapters. Each one is simple enough to explain with something you could hold in your hand. Each one, taken seriously, explains part of why the machines work. Together they explain most of it.

The square in Flatland did eventually see the third dimension. The sphere lifted him out of the plane and let him look down. He came home changed, and nobody believed him. We do not need to be lifted anywhere. We only need to stop trusting the instinct that says the room we are in is the only kind of room there is.

CHAPTER TWO

All the Orange Is Peel

SATURDAY MORNING, THE CUTTING BOARD

You slice an orange in half. The peel is thin, a few millimeters, and almost all of what you are holding is fruit. That is the way it should be. The skin is the edge. The inside is the inside.

Here is the first rule, and it is the strangest one.


**In high dimensions, almost everything is near the edge.
The inside is nearly empty.**

Imagine a peel whose thickness is just one percent of the orange's radius. In the orange on your cutting board, that peel holds about three percent of the total volume. The rest is fruit.

Now build the same orange in more dimensions. Keep the peel at one percent of the radius. In ten dimensions, the peel holds almost a tenth of the volume. In a hundred dimensions, it holds almost two thirds. In a thousand dimensions, the peel holds more than ninety-nine point nine nine percent of everything. There is essentially no fruit left. The orange is all peel.

FIGURE 02

How much of the orange is peel



A peel one percent of the radius thick. Share of total volume is $1 - 0.99^n$ for n dimensions. Computed exactly.


Why it happens

The reason is multiplication. To be in the fruit, a point has to be a little way in from the skin in every direction at once. Each direction takes away a small slice of room. In three dimensions, you lose a little three times. In a thousand dimensions, you lose a little a thousand times, and a little multiplied by itself a thousand times is nothing.

If you want the whole idea in one line, the fraction of the orange that is peel is one minus 0.99 multiplied by itself once for every dimension. That is the only formula you need for this chapter. Everything else follows from it.

FIGURE 03

The ball that peaks and vanishes



Volume of a ball of radius one. It grows until five dimensions, then shrinks toward zero while the box around it keeps growing. Computed exactly.

The box that pokes through its own walls

There is a second picture that makes the same point, and it is even stranger. The engineer Richard Hamming liked to use it to show students how badly their intuition fails.

Take a square box, four units on each side. Put four coins in it, one in each corner, each coin of radius one, so they just touch each other and the walls. In the center there is a little gap between the four coins. Fit the largest coin you can into that gap. It is small, much smaller than the others.


Now do the same thing with a cube and eight balls, one in each corner. The ball in the center is a bit bigger than the central coin was, but still well inside the box.

Keep going. In four dimensions, the central ball is exactly as large as the corner balls. In nine dimensions, it is so large that it touches the walls of the box. In ten dimensions, the ball in the middle, the one squeezed between all the others, is bigger than the box that contains them. It pokes out through the walls.

That cannot happen, says the square inside you. It happens. The reason is that the corners of a high-dimensional box are astonishingly far from its center. The diagonal of a box grows with the square root of the number of dimensions, so the corners run away while the walls stay put. The corner balls crowd out toward the corners, and the space between them, the middle, opens into something huge and hollow. A high-dimensional cube is less like a box and more like a sea urchin: a small body and a vast number of long spines.


FIGURE 04

The box that pokes through its own walls



Two dimensions: a small coin in the gap

Hamming's puzzle. Balls of radius one in every corner of a box four units wide; the ball squeezed between them has radius $\sqrt{n} - 1$. At nine dimensions it touches the walls; past that, it pokes out. Computed exactly.




Nobody is average

This is not only a fact about oranges and boxes. It is a fact about anything described by many numbers.

In the early 1950s the United States Air Force had a problem. Its pilots kept losing control of their aircraft, and the cockpits had been built to fit the average pilot. A young researcher named Gilbert Daniels measured thousands of pilots on ten body dimensions that mattered for the cockpit and asked a simple question. How many of them were close to average, within the middle range, on all ten?

The answer was none.

Not a few. None. Each pilot was average on some measurements and not on others, and with ten independent ways to differ, being near the middle on every single one was almost impossible. The middle of a ten-dimensional space is nearly empty. Everyone lives on the peel.



The Air Force stopped designing for the average and started building adjustable seats. That decision is the right lesson in miniature. When things differ in many ways at once, the typical case is not in the middle. The typical case is unusual in its own particular way. There is no average customer, no average patient, no average student, no average reader of this book.

What it means for machines

A machine that learns from millions of examples, each described by thousands of numbers, lives entirely on the peel. It never sees the middle, because nothing is there. Every sentence, every face, every customer is an edge case, and there is room on the edge for all of them to be different without crowding one another.

That is the first hint of why size does not drown these systems. On the surface of a thousand-dimensional orange there is more room than you could fill with every sentence ever written. Things that seemed like they would collide never meet.

We built our institutions, our forms, our sizes and our categories for the middle. The middle was always a fiction. The machines are the first tools we have ever built that do not need it.

CHAPTER THREE

Everything Is at Right Angles

A CROWDED ROOM

You throw a dart at a wall without looking, then turn around and throw another. You would not expect the two darts to have flown in exactly opposite directions, or in exactly the same direction. But they could have. In the room you are standing in, two random throws can point anywhere relative to each other.

Here is the second rule.


In high dimensions, two random directions are almost always at right angles to each other.

Point an arrow in a random direction. Then point a second arrow in another random direction. In ordinary space, the angle between them could be anything from nearly zero to nearly one hundred eighty degrees. It wanders all over the place.

In a hundred dimensions, the angle is almost always within ten degrees of a right angle. In a thousand dimensions, within about three. In ten thousand dimensions, it is ninety degrees to within a degree, almost every time. Random directions do not wander. They are perpendicular.

FIGURE 06

The angle between two random directions



The middle ninety percent of a thousand random pairs at each dimension, from a simple simulation. The band closes in on a right angle.

Why it happens

Two arrows line up when their parts agree. In three dimensions there are only three parts, so a little luck is enough. In a thousand dimensions there are a thousand parts, and each one is a separate coin toss. Some push the arrows toward agreement, some push them apart, and across a thousand independent pushes they cancel almost perfectly. Asking two random directions in high-dimensional space to agree is like asking a thousand strangers to agree on lunch. They will not. On average, they will land exactly on indifference.

Room for more directions than there are dimensions

Here is the consequence that matters most, and it is one of the quiet miracles of this subject.

In three dimensions, you can have exactly three directions that are perfectly at right angles to each other. That is what three dimensions means. If you allow the directions to be almost at right angles, say within a few degrees, you do not get many more.

In high dimensions, if you allow almost, you get an astonishing number. The count of nearly perpendicular directions does not grow in step with the number of dimensions. It grows exponentially with it. Ask only that

every pair stay within about ten degrees of a right angle, and a space of four thousand dimensions, roughly the size used inside many language models, has room for trillions of them.

In 1984 two mathematicians, William Johnson and Joram Lindenstrauss, proved a result that belongs to the same family. They showed that a crowd of points in a huge space can be squeezed into a much smaller one while nearly keeping every distance between them, and that the smaller space only needs to grow with the logarithm of the number of points. A million points do not need a million dimensions to stay apart. A few thousand will do.

Meaning gets its own direction

Now put that next to what an AI does with words.

A language model gives every concept a direction. Not a single slot, not a filing cabinet drawer, but a way of pointing. A concept like royal points one way. A concept like female points another. Plural points another. Past tense another. Because there is room for so many nearly perpendicular directions, the model can give millions of concepts their own direction without them stepping on each other.


Researchers at Anthropic studied this in small networks in 2022 and gave it a name: superposition. When the features it needs are sparse, each one rarely active at the same time as the others, a network with only a handful of dimensions can store many more features than it has dimensions, by laying them along directions that are nearly, but not exactly, at right angles. Each feature interferes with the others only a little, and the interference is the price of the room.

This is why a model with a few thousand dimensions can hold a world. It is not because each dimension stands for one thing. It is because the space has room for far more directions than dimensions, and meaning lives in directions.

The price of almost

Almost perpendicular is not perpendicular. That small gap explains some of the most familiar mistakes AI makes.

When you search for documents that say a product is safe, a system that works by direction may hand you one that says it is not safe. In logic, the word not is a complete reversal. In a space where meaning is direction, not is a short step sideways. The two sentences point almost the same way, because they share almost everything except one small word, and that word moves the arrow only a little.



The machine is not careless. It is living by the rules of its room. In a room where everything is at right angles, a reversal is just another small turn.

CHAPTER FOUR

Everyone Is Equally Far

A NEW CITY, NO MAP

You are looking for coffee. One shop is two blocks away. Another is twenty. The difference is obvious, and it decides where you go. Near and far are the most useful words in the world.

Here is the third rule.


In high dimensions, if points are scattered at random, the nearest one and the farthest one are almost the same distance away.

Scatter a thousand random points on a sheet of paper. Pick a spot and measure the distance to the closest point and to the farthest. The farthest might be fifty times farther than the closest. Near and far mean something.

Now scatter a thousand random points in ten dimensions. The farthest is only about four times as far as the nearest. In a hundred dimensions, about one and a half times. In a thousand dimensions, the farthest point is only about thirteen percent farther than the nearest one. Everyone is roughly the same distance from everyone. Near and far have almost stopped meaning anything.

FIGURE 08

Farthest point ÷ nearest point



One thousand points scattered at random in a cube, measured from one random spot. Logarithmic scale; a ratio of one would mean everyone is exactly equally far. From a simple simulation.

The curse

This is the effect that gave the subject its most famous phrase. In 1957, the mathematician Richard Bellman, working on how to plan decisions over time, ran into the way problems explode as you add variables, and he called it the curse of dimensionality. The phrase stuck because the curse is real. If you try to fill a high-dimensional space with examples, you never can. If you try to find similar things by raw distance, everything looks equally similar.

In 1999 four database researchers, Kevin Beyer, Jonathan Goldstein, Raghu Ramakrishnan and Uri Shaft, wrote a paper with a blunt title: When Is "Nearest Neighbor" Meaningful? Their answer was that, under very broad conditions, as the dimensions pile up, the gap between the nearest and the farthest point shrinks to almost nothing compared with the distances themselves. Asking for the closest match stops making sense.

That should have been the end of the story. A machine that finds things by nearness, in a space where nothing is near, should be useless.

The way out

It is not useless, and the reason is the most important idea in the first half of this book.

Real things are not scattered at random.


The curse describes empty, random space: points thrown anywhere, every direction equally likely. The world is not like that. Faces are not random collections of pixels. Sentences are not random collections of words. Real data, however many dimensions it is described in, tends to lie near a much smaller surface inside the huge space, the way a thread can wander through an enormous room while being, itself, only one dimension thick.

Mathematicians call this the manifold hypothesis. You do not need the word. You need the picture. The room is vast and nearly all of it is empty. The things that actually exist cluster on thin, curved sheets that run through it. On those sheets, near and far mean something again. Two photos of the same face sit close together on the sheet. A photo of a different face sits farther along it.

FIGURE 09


The curse and the way out

The curse: points scattered at random



nothing is near anything

The world: things lie on a thin sheet



along the sheet, near means something again

Left: random points, where nothing is near anything. Right: real data, which lies near a thin sheet running through the room. An illustration.

The curse is a fact about the room. The blessing is a fact about the world. Learning is the work of finding the sheet.

What it means

This is why the old way of searching, by matching exact words, could never scale to meaning, and why the new way works. The machine does not look for the nearest point in the whole room. That would be hopeless. It first learns where the sheets are, and then it measures distance along them. Most of its effort goes into the learning, not the looking.

It also explains why so much of the value in AI comes from a step most people never see: turning raw things into good addresses. A bad address puts everything in the random room, where everyone is equally far. A good address puts things on the sheet, where similar things are neighbors. The quality of an AI system is, to a surprising degree, the quality of its addresses.

CHAPTER FIVE

Every Trap Is Near the Floor

LATE AFTERNOON, A MOUNTAIN IN FOG

You are hiking down and the fog has come in. You cannot see the trail, only the ground at your feet. So you do the only sensible thing. You take a step in whatever direction goes down, then another, then another. After an hour you reach a little hollow where every direction goes up. You are stuck. And through a break in the fog you see the real valley, far below, on the other side of a ridge.

That hiker is the oldest worry in machine learning.

A learning machine improves by doing exactly what the hiker does. It measures how wrong it is, then nudges its settings in whatever direction makes it a little less wrong, again and again, millions of times. Engineers call it gradient descent. The fear, for decades, was the hollow: the machine would settle into a small dip that felt like the bottom and never find the real valley. In a landscape with many settings, people assumed, there would be countless hollows, and the machine would get caught in the first one it met.

Here is the fourth rule. It says the fear was based on three-dimensional thinking.


In high dimensions, almost every place where the ground is flat is a pass, not a trap. Real traps exist almost only near the bottom.

Why traps vanish

To be truly stuck, the ground has to rise in every direction around you. On a mountain you can see, there are only two directions to check, north-south and east-west, so it is easy for both to rise and trap you.

Now imagine a landscape with a million directions. At some flat spot, consider each direction separately. Going one way, the ground might curve up or curve down. Treat it, roughly, as a coin toss. To be in a trap, every single one of a million coins has to come up the same way. The chance of that is one half multiplied by itself a million times. It is so close to zero that the word unlikely is too generous.

What you find instead, nearly everywhere, is a saddle: a place where the ground rises in some directions and falls in others. A saddle feels like a trap if you only test a few directions. It is not one. There is always a way down, and in a space with a million directions there are usually many.



So where are the real traps? Only where nearly every direction genuinely points up, and that happens mostly near the bottom, where there is very little further down to go. In 2014 a group of researchers including Yann Dauphin, Surya Ganguli and Yoshua Bengio laid out this argument for

neural networks. They argued, with evidence from real networks, that the places where training seemed to stall were mostly saddles, not hollows, and that the true dips that remained tended to sit close to the bottom.

The coin toss is a simplification. Real landscapes are not built from independent coins, and the argument is a way of seeing, not a proof. But the shape of it has held up. In high dimensions, being stuck is rare, and being stuck badly is rarer.


The blessing

In 2000, the statistician David Donoho gave a lecture whose title named both sides of the coin: the curses and the blessings of dimensionality. The curses were famous. The blessings, a phrase others had used a few years before but that Donoho made widely known, were the news. Some things get easier as dimensions grow, not harder. This chapter is one of them.

Here is the version you can feel. Give a learning machine more settings, more knobs to turn, and you would expect it to get more confused. There is more room to get lost. Instead, the extra room gives it more ways down. Every new dimension is another possible escape from every saddle. A larger landscape is not a bigger maze. It is a maze with more exits.

This is part of why the field kept discovering, to its own surprise, that bigger models were easier to train well, not harder. In 2019 Mikhail Belkin and his colleagues described a curve that had been showing up in practice and that the textbooks said should not exist. As models grew past the size where they could memorize their training examples perfectly, their performance on new examples, after getting worse, began to improve again. They called it double descent. The textbook had been drawn in a small room. The machines were working in a large one.

FIGURE 11

Double descent

The shape only, not data. Error on new examples gets worse near the size where a model can memorize its examples, then improves again as models keep growing. Named by Belkin and colleagues in 2019.

What it means

The hiker in the fog is still a good picture of how learning works. One step at a time, downhill, never seeing the whole mountain. What changed is the mountain. In a space with enough directions, the fog does not trap you, because there is almost always a way down that you have not tried yet.

That is a strange thing to be true, and it has a consequence that goes well beyond machines. The more independent ways you have to improve, the less likely you are to be truly stuck. A person, a company, or an institution that can only change along two or three directions gets trapped in hollows all the time. One that can change along many rarely does.

CHAPTER SIX

You Only See the Shadow

A QUARTERLY REVIEW

The dashboard has twelve numbers on it. Eleven of them are flat. The twelfth moved a little. Everyone agrees it was a quiet quarter. Six weeks later the business is in a different place entirely, and nobody can say when it happened.

Here is the fifth rule. It is the one I find most beautiful, and the one with the longest reach.


In high dimensions, almost all of any change happens in directions you are not watching.

Every measurement is a shadow. When you watch one number, you are looking at the whole of something, projected onto a single direction. When you watch twelve numbers, you are looking at its shadow on twelve directions. That is fine in a world of three dimensions, where twelve gauges would cover everything with room to spare.

Now suppose the thing you care about lives in a thousand dimensions, and it changes in some direction you did not choose in advance. How much of that change shows up in any one gauge? On average, only one part in a thousand of its total strength, measured the way physicists measure energy. Twelve gauges catch about twelve parts in a thousand. The other ninety-nine percent of the change, roughly, happened completely, in directions nobody was looking at.

FIGURE 12

You only see the shadow



Why it happens

It is the second rule wearing different clothes. A random direction is almost perpendicular to any fixed direction you pick. Your gauges are fixed directions. The change is a random direction. So the change is almost perpendicular to every gauge you own. The gauges barely flicker. The thing itself has moved a long way.

That is why the quarter looked quiet. Nothing was hidden. Nothing was faked. Every number on the dashboard was correct. The change was simply at right angles to the dashboard.

Something I have been writing about for a long time

Readers who know my work on the Reality Equation will recognize this shape. In that framework, Expectation is not a single number. It has a real part, Prediction, the thing habit forecasts, and a part at right angles to it, the Ideal. And the lesson I keep returning to is that you only feel the real part. A change in the perpendicular direction does not make anything bigger or smaller in the part you can feel. It turns you. And you cannot feel a turn.

I mean that as a statement about experience, not as physics, and I want to be careful to keep the two separate. The mathematics in this chapter is plain geometry. The Reality Equation is a proposal about how surprise is lived. But the two rhyme. Two dimensions are enough to show that a perpendicular change can be invisible to the one you feel. A thousand dimensions make it the general case. In a world of many directions, almost every real change is a turn, and almost every turn is unfelt by any instrument pointed the old way.

What it means for machines

A learning machine is, in a sense, the first instrument that watches all the directions at once. That is what it means for it to represent something with thousands of numbers instead of twelve. It does not need you to decide in advance which gauges matter. It sees the shadow on every direction it has, and most of what it notices are changes no human dashboard was built to catch.

This is why machines now find patterns people cannot see, in a heartbeat, a photograph of an eye, a protein, a smell. The patterns were never hidden in some secret place. They were sitting at right angles to everything we had chosen to measure.

It also tells you where the next surprises will come from. Not from the numbers you watch. From the directions you never thought to.

The five rules together

Before we go on, here they are in one place.

FIGURE 13

The five rules

- 01

The middle is empty

Almost everything is near the edge. Nobody is average.
- 02

Everything is at right angles

Room for far more directions than dimensions.
- 03

Everyone is equally far

So real things must live on thin sheets.
- 04

Traps live near the floor

More directions means more ways down.
- 05

You only see the shadow

Almost every change is perpendicular to your gauges.

Everything in Part One on one page.

None of these is a matter of opinion. You can check each one with a computer in an afternoon. What follows from them, though, is where it gets interesting.

PART TWO

What the Machines Do With It

Finding the sheet in the empty room



CHAPTER SEVEN

A Thousand Dots and One Line

A SCHOOL SCIENCE FAIR


A student has dropped a ball from a balcony a thousand times and written down how long each fall took from each height. She could hand the judges the notebook. Instead she draws one curve through the dots and says: give me any height, even one I never tried, and I will tell you the time.

That moment, the notebook becoming the curve, is the difference between memory and intelligence. It is also the cleanest way to see what the first half of this book has been building toward.

A thousand dots in a notebook can answer a thousand questions perfectly and no others. Ask about the height she never tried and the notebook is silent. The curve is smaller than the notebook, and it knows more. It has thrown the dots away and kept the shape they were lying on. That is compression, and the right kind of compression does not lose information. It gains the ability to go somewhere new.

Learning is not storing the examples. It is finding the surface they lie on, so you can stand on it anywhere.

FIGURE 14

A thousand dots and one line

Fall time grows with the square root of the height. The dots are illustrative; the curve is the law. The curve answers for a height the notebook never covered.

The sheet, again

Chapter Four ended with a picture: a vast room, nearly empty, with the things that actually exist lying on thin sheets running through it. The student's curve is a one-dimensional sheet in a two-dimensional room. The ball's fall time depends on its height in a single smooth way, so all the true dots lie on a single line.

A language model is doing the same thing in a room with thousands of dimensions. Every sentence it has read is a dot. The sentences people actually write are not scattered across the room. They lie on an enormous, folded, many-layered sheet, the sheet of things that make sense. Training is the slow work of finding that sheet. Once it is found, the model can stand at a spot on it where no dot ever was, and what it says there will still make sense, because it is still on the sheet.

That is why a model can answer a question nobody has ever asked. It is not looking the answer up. It is standing at a new place on a surface it has learned the shape of.

Two kinds of compression

Newton's law of gravity fits on a postage stamp. It is a compression of every falling apple and every orbiting moon into one short line. It is a closed form: a formula you can write down and read.

A large AI model is also a compression, but it does not fit on anything. Its shape is spread across billions of numbers, none of which means anything on its own. It is a distributed form. You cannot read it the way you read Newton. You can only ask it questions and watch where it stands.

For three hundred years, science meant finding closed forms. The pattern was small enough to write, or we did not count it as found. The machines have broken that rule. They find patterns whose shape is too large and too folded to ever be written as a formula, and they use those patterns anyway. Some of the most consequential knowledge in the coming decades will be knowledge no human can read, held in sheets no human can draw, and it will still work.

Why dimensions are the point

It would be easy to think of all those dimensions as a cost, a necessary evil of complicated data. The rules in Part One say the opposite.

The edge rule says there is room on the sheet for every case to be different. The right-angle rule says there are enough directions for every concept to have its own. The trap rule says that finding the sheet is possible, because the search rarely gets stuck. The shadow rule says the machine can notice what no fixed set of gauges would. And the distance rule, the curse, is exactly what makes the sheet worth finding: in the empty room, nothing is near anything, so the only useful nearness is nearness along the sheet.

The dimensions are not the problem the machine has to overcome. They are the room it needs to do the work.

CHAPTER EIGHT

Meaning Becomes an Address

A BASEMENT, A BANKER'S BOX

You are looking for a letter your grandmother wrote. You know what it said. You remember how it made you feel. But the box is filed by year, and you do not remember the year. So you open every folder.

For all of human history, finding something meant knowing what someone else had called it. A book had a call number. A document had a filename. A website had keywords. A product had a category. To find the thing, you had to guess the name the filer used. If you guessed wrong, the thing might as well not exist.

That is changing at the root. In the machines, the address of a thing is computed from the thing itself. Turn a sentence into its several thousand numbers and those numbers are where it lives. No one chose them. No one filed it. And because of the rules in Part One, things that mean alike end up near each other on the sheet, whatever words they happened to use.

For the first time, where a thing is stored is decided by what it means, not by what someone called it.


The famous example

In 2013 Tomas Mikolov and his colleagues, first at Microsoft Research and then at Google, published methods for giving words addresses, and their word vectors did something that startled people. Take the address of king. Subtract the address of man. Add the address of woman. The closest

word to where you land is queen.

FIGURE 15

Meaning has directions



The famous word-vector result from 2013, drawn as a parallelogram. In real vectors the arithmetic lands close to queen, not exactly on it.

That is not a trick of the particular words. It means the direction from man to woman was roughly the same direction as the one from king to queen. The machine had found, without being told, that a concept like gender is a direction in the space, and that you can move along it. The relationship was not stored anywhere. It was simply the shape of the room.

The example is imperfect. The arithmetic only lands close, not exactly, and plenty of other attempted analogies fail. But it was the first time most people saw that meaning could be geometry: that words have locations, that relationships are directions, and that thought might look, from the inside, like movement.

The small words problem

There is a cost, and it is worth knowing because you will run into it.

A sentence's address is a kind of blend of all its words. The words that carry the most meaning in logic, not, only, except, never, are often the smallest. In the blend, each word is one voice in a room. The longer the passage, the more voices, and the quieter the one word that changes

everything.

So an AI search for documents that say a contract allows something may return a document that says it allows everything except that thing. Both sit close together on the sheet, because they share nearly all their words and nearly all their direction. A careful system adds a second step that reads closely once the neighborhood is found. The address gets you to the right street. Reading gets you to the right door.

What it means

The consequence is bigger than better search.

Every system we built around names starts to loosen. Filing systems. Keyword advertising. Job titles matched against job descriptions. Library categories. Product taxonomies. Legal research by exact phrase. All of them were workarounds for the fact that machines could not tell what things meant, so people had to label everything in advance and hope the labels matched.

When the address comes from meaning, the label becomes optional. You will be able to find your grandmother's letter by asking for the letter where she talked about the farm and sounded tired. The box will not need to be filed by year. It will not need to be filed at all.

CHAPTER NINE

The Space Between

A STUDIO, AN EMPTY CANVAS

A designer has a client who wants something warm but modern, playful but serious, like a hotel lobby but also like a kitchen. Ten years ago that brief was a joke. Today she types it, and a picture appears that is all of those things at once, and is not quite like anything that existed before.

People argue about whether AI can be creative. The geometry gives a simpler way to say what it is actually doing.

On the sheet of things that make sense, every known thing is a dot. Between the dots there is surface: places where nothing has been made yet but where anything made would still make sense. Most of the sheet is like that. The dots are sparse. The surface between them is enormous.

Most new things are not far away. They are between.

When the designer asks for warm but modern, the machine goes to the neighborhood of warm things, then moves along the direction of modern, and stops at a place on the surface nobody has stood on. The result is new in the plain sense that it did not exist. It is also familiar, because every direction it moved along was learned from things that did exist.


Two kinds of new

That points to a distinction that will matter a great deal.

The first kind of new is between. It is a place on the sheet nobody has visited, surrounded by places people have. Machines are extraordinarily good at this kind. They can visit more of the surface in an afternoon than humanity visited in a century, and almost everything they find there will make sense.

The second kind of new is off the sheet. It is the thing that does not lie on the surface of what made sense before, until someone puts it there. Perspective in painting. The number zero. Atonal music. The first time someone said that the earth moves. These did not sit between existing ideas. They changed the shape of the sheet itself.

A machine that has learned the sheet can wander it endlessly. It is less clear that it can step off it on purpose, and it is very clear that when it steps off by accident we have another name for it. When a model states a false fact with total confidence, it has usually moved to a place in the room that is not on the sheet of true things, while still being on the sheet of things that sound true. The two sheets lie very close together. In many places they are a hair apart.



What it means


Between-creativity is about to become nearly free. Ideas that are combinations of existing ideas, designs that blend known styles, solutions that join known parts in a new arrangement, will be produced in such quantity that they stop being scarce.

That does not make human creativity less valuable. It moves it. When the space between the dots fills in, what becomes rare is the decision about which direction to move in, and the occasional act of putting a new dot where no sheet was. The value goes to taste, to judgment, and to the few people willing to be wrong in an interesting direction.

PART THREE

The Future It Is Building

Seven predictions you can hold me to




CHAPTER TEN

Seven Predictions

A TUESDAY IN 2031

You need the proposal you sent a client two springs ago, the one with the pricing you regret. You do not remember its name, its folder, or the client's exact company name. You say: the proposal where I undercharged the marina people. It is on your screen before you finish the sentence. You realize you have not named a file in years.

A book about geometry should end with something you can check. Here is what I think the five rules do to ordinary life once machines that live by them are everywhere. Each prediction follows from a rule in Part One. Each one is stated plainly enough that you can tell, in a few years, whether I was wrong.



1. Filing ends.

Within five years, most people will stop naming and organizing the things they will need to find again. Folders, filenames, tags and labels will survive the way card catalogs survived: as something older people remember using. You will find things by describing what they meant to you, because that is where they are stored. This is the address rule from Chapter Eight, and it is the first prediction because it is already starting.

2. The keyword retires.

Anything that today is bought, sold or matched by exact words, search advertising, job listings, legal research, product categories, will move to matching by neighborhood on the sheet of meaning. The businesses that sell words will have to learn to sell directions. The ones that do not will discover that their customers stopped typing the words.

3. The average stops being designed for.

Forms, sizes, plans, lessons, treatment schedules and contracts were built for the middle because the middle was the only thing a human designer could reason about. The edge rule says the middle is empty. As machines make it cheap to fit a thing to one person's own position among thousands of dimensions, the default will flip from one size with exceptions to fitted with a standard version for the rare person who wants one. The adjustable cockpit seat becomes the model for everything.

4. The dashboard is replaced by the turn report.

Organizations will keep their gauges, but they will stop trusting them alone. The shadow rule says almost every real change happens at right angles to whatever you chose to measure. The most valuable report in a company will be the one a machine writes about what moved in the directions nobody was watching. Leaders who read it will be surprised less often. Leaders who do not will keep having quiet quarters that turn out not to have been quiet.

5. Knowledge that works but cannot be read becomes normal.

More and more of what we rely on, in medicine, materials, weather, biology and finance, will come from distributed patterns held in models rather than formulas written on paper. We will use them the way we use a trusted doctor's judgment: because of their record, not because we can follow every step. The argument over whether to trust a result we cannot explain will be one of the defining arguments of the next decade, and the side that wins it will be the side with the better track record, not the better explanation.

6. Between-creativity becomes abundant, and direction becomes the scarce skill.

Combinations of existing ideas, styles and solutions will be produced in such volume that they stop being worth much on their own. What gains value is the choice of which direction to move in, the taste to know which of a thousand blends is the right one, and the rare willingness to put a new point where no sheet was. The person who can say "not that, further this way" will be worth more than the person who can make the thing.

7. The errors that matter will be the small ones.

As machines get better, the mistakes people actually notice will narrow to two kinds, both geometric. The first is the small-word mistake: not, except, only and never treated as a short step instead of a reversal. The second is the confident step off the sheet of true things onto the sheet of things that merely sound true. Neither is fixed by size alone. Both are fixed by pairing the machine's sense of location with a slower, closer reading. The systems people trust will be the ones that do both.

CHAPTER ELEVEN

Your Three Dimensions

A LAB IN CHARLESTON, A WEEKDAY

Someone sits across the table and asks the question everyone eventually asks. But how does it actually know? And I have a choice. I can talk about parameters and training data, and their eyes will glaze. Or I can pick up an orange.

I have a strange enthusiasm, and I know it. Very few people get excited about the geometry of high-dimensional space. I do, because I think it is the most honest answer to the question everyone is asking about AI. It is not a bigger brain. It is not a trick. It is a different kind of room, with rules we have known about for a century and never needed until now, and those rules happen to be the ones that make learning possible.

That answer changes how I think about my own place in it.

We live in three dimensions. The machines live in thousands. It would be easy to conclude that we are simply outclassed, that a mind that sees every direction at once has no use for one that sees three. I do not think that follows. Look again at the rules. The machine can visit any place on the sheet. It can watch every direction. It can find the surface the examples lie on. What none of the rules gives it is a reason to go one way rather than another.

Direction is what a person brings. Not because people are special in some mystical sense, but because the question of what matters is not a question about the shape of the room. It is a question about why you are in it.

I have felt this change in my own work. There was a point where I noticed that I had become the slowest part of everything I was doing, the step

every task was waiting on, and I started telling my AI, more and more often, you decide. It was the right call, and it is the same lesson as this book. The machine is better than I am at moving through the room. My job is to point.

So here is the practical advice, if a book like this is allowed any. Stop competing with the machine on its own ground. Do not try to visit more of the sheet than it can. Do not try to watch more directions. Spend your attention, which is the one thing you have that it does not need, on choosing where to look and saying which way is forward. That is your three dimensions. It is enough.

CLAIM LEDGER

What Is Record and What Is Argument

Already true

The five rules in Part One are mathematics, not opinion. Most of the volume of a high-dimensional ball lies near its surface. Random directions in high dimensions are nearly perpendicular. Random points in high dimensions are nearly equidistant. Critical points in high-dimensional landscapes are overwhelmingly saddles under broad assumptions. Any fixed set of measurements captures only a small share of a random change. Language models do represent meaning as positions and directions in spaces with thousands of dimensions, and search by meaning is already replacing search by exact words in many products.

What has to happen

For the predictions to come true, the address has to become trustworthy enough that people stop keeping their own labels as a backup. Fitted products have to become cheaper to deliver than standard ones, not just possible. Organizations have to be willing to act on a report about directions they never chose to measure. And trust in unreadable knowledge has to be earned case by case, through records that people can check.

Where I am probably wrong

The trap rule rests on a simplification. Real learning landscapes are not built from independent coin tosses, and there are well-studied cases where training does get stuck. If the next generation of models stops improving with size, it will be partly because the room ran out of exits, and Chapter

Five will need rewriting.

I may also be wrong about timing. Filing and keywords are habits, and habits outlive their reasons. Five years may be ten.

And I may be wrong about direction being the human's job. If machines learn to choose what matters as well as they learn to move, the last chapter of this book is the first thing that will age. I do not think that is close. I would rather say it plainly than hedge every sentence against it.

EPILOGUE

The Car Has an Address

6:40 P.M., THE SAME GARAGE, A FEW YEARS FROM NOW

You come out of the concert and you do not look for your car. You never knew where you left it. You did not need to. The place it is, among all the places it could be, was never something you had to remember. It was an address, and the address was kept for you.

The garage still has its floors. The room still has its thousands of directions. Nothing about the geometry changed. What changed is that we finally built something that lives comfortably where we cannot, and it turned out the strange rules were not a wall. They were the reason it works.

The middle was always empty. The directions were always at right angles. The traps were always rarer than we feared. Most change was always happening where we were not looking.

We just could not see it from three.

BACKGROUND

For Further Reading

- Edwin A. Abbott, *Flatland: A Romance of Many Dimensions* (1884).
- Richard Bellman, *Dynamic Programming* (Princeton University Press, 1957), where the curse of dimensionality is named.
- Richard W. Hamming, *The Art of Doing Science and Engineering* (1997), chapter 9, "n-Dimensional Space."
- Gilbert S. Daniels, *The "Average Man"?* (Wright Air Development Center, 1952).
- William B. Johnson and Joram Lindenstrauss, "Extensions of Lipschitz mappings into a Hilbert space" (1984).
- Kevin Beyer, Jonathan Goldstein, Raghu Ramakrishnan and Uri Shaft, "When Is 'Nearest Neighbor' Meaningful?" (ICDT 1999).
- David L. Donoho, "High-Dimensional Data Analysis: The Curses and Blessings of Dimensionality" (AMS lecture, 2000).
- Tomas Mikolov, Wen-tau Yih and Geoffrey Zweig, "Linguistic Regularities in Continuous Space Word Representations" (NAACL 2013).
- Yann Dauphin and colleagues, "Identifying and attacking the saddle point problem in high-dimensional non-convex optimization" (NIPS 2014).
- Charles Fefferman, Sanjoy Mitter and Hariharan Narayanan, "Testing the manifold hypothesis" (Journal of the American Mathematical Society, 2016).
- Mikhail Belkin, Daniel Hsu, Siyuan Ma and Soumik Mandal, "Reconciling modern machine-learning practice and the classical bias-variance trade-off" (PNAS, 2019).

- Nelson Elhage and colleagues, "Toy Models of Superposition" (Anthropic, 2022).