

The Universe as a Learning Machine

จักรวาลในฐานะเครื่องเรียนรู้

Machine Learning from the First Second of the Universe to the Meaning of Everything

The universe can learn anything — what it learns for is what we give it.



Contents



Prologue

The Question Hidden Beneath Everything

Machine Learning has become the technology that is reshaping this century. But beneath its engineering success, a question hides quietly — one far larger than the technology itself. Is Machine Learning merely a technique that humans invented in the last few decades — or did humans only *discover*, late and slowly, something that this universe has been running since the first instant it came into being?

This book is the answer, and the answer is: the latter.

Our thesis is disarmingly simple, yet its consequences spread until they cover everything. The essence of Machine Learning is not the silicon chip, not the calculus, not the programming language. Its essence is a *loop* — a four-beat cycle: generate variety (variation), send that variety out into the world to be measured (evaluation), keep what passes and discard what does not (selection), and repeat from what was kept (iteration). This loop — vary, evaluate, select, repeat — is the heartbeat beneath every system we call "capable of learning."

And the proposal of this book is that the same loop beats beneath *everything* — not only beneath the brains of animals or the artificial neural networks of our age, but beneath the formation of stars, beneath the chemistry of the early ocean, beneath the evolution of life, beneath perception itself, beneath the growth of science and culture. The universe does not "use" Machine Learning as though it were a tool picked off a shelf. The universe is a learning machine, and the ML that humans wrote by hand is simply the first time that loop ever caught sight of itself.

We will travel along this loop for the whole book — from the first second in which space and time were born, through the moment atoms learned to bind, through the moment chemistry learned to copy itself, through the four billion years of trial and error we call evolution, through the birth of brains that hold a model of the world inside themselves, through the moment the universe built a mirror to look at itself in the form of science and artificial intelligence — and at last we will arrive at the hardest question of all: the meaning of everything.

One caution before we begin. This book will use the words "learn," "search," "try," "goal" for things that have no mind — for stars, for molecules, for genes. This is not a granting of souls to the cosmos; it is not a claim that a stone has desires. It is a way of pointing at the fact that a mindless process, when three ingredients are present — variety, a filter, and the inheritance of what survives — produces behavior that is *indistinguishable* from purposeful learning. The interesting question is not "did the universe intend it?" The interesting question is "if the result is identical in every respect, does intention still matter?" — and we will keep that question for the final chapter.

Let us begin, then, at the beginning of every beginning.



Part One

The Algorithm Before Anyone Was There to Run It

When the laws had not yet decided what laws to be

Picture the first second. Not the first second of your life, not the first second of the Earth, but the first second of everything — the instant in which space, time, matter, and energy came into being together in the event we call the Big Bang.

First, the most deeply rooted misunderstanding must be cleared away. The Big Bang was not an explosion. It was not a bomb lit in the middle of a dark, empty void, hurling debris outward in all directions — because there was no "empty void" to explode into. There was no "space" beforehand. The Big Bang was the expansion of space *itself* — every point of space spreading apart from every other at once, all of it, with no center and no edge. The spot where you are sitting now is a spot where the Big Bang happened, and so is every other point in the universe.

In the very first sliver of a second, the universe passed through an era called cosmic inflation: space expanded faster than light, by an enormous factor. This was an event *after* the Big Bang, not the Big Bang itself — and not the slow expansion we still see today. It was an instant in which the whole cosmos stretched from smaller than an atom to something vast, in a time so short the word "time" barely applies.

But what is worth noticing is not the size, and not the speed. What is worth noticing is that, in that instant, the universe was *deciding what to be*.

Symmetry breaking — the first selection event

At the extreme temperatures of the newborn universe, the forces we know today — electromagnetism, the strong nuclear force, the weak nuclear force — had not yet separated from one another. They were fused into one, perfectly symmetric. A state of perfect symmetry is a state that "has not yet chosen" — every direction the same as every other, every possibility still open.

Then the universe cooled.

And as it cooled, what happened is what physicists call spontaneous symmetry breaking. Let this be understood precisely — it does not mean the laws of physics lost their symmetry. The laws remained as symmetric as before. What happened is that the *ground state* of the universe — the state it settled into as it cooled — could no longer reflect the full symmetry of the laws. It had to "choose." The classic image is a pencil balanced on its point. While it stands upright, it is perfectly symmetric, every direction equal. But that state is unstable. When it falls, it must fall *some* way, and the instant it falls the symmetry shatters: one direction is chosen, all the others discarded.

This is what we might call the universe's first selection event. In each act of symmetry breaking, the universe locked in one of its own parameters. It decided how the forces would separate, which particles would have mass, what those masses would be, by what ratio matter would outnumber antimatter. Every time it chose, countless possibilities were cut away and one single reality was kept. If you view it in the language of Machine Learning, this is exactly what happens when a learning algorithm "settles" toward an answer — it begins from a random state in which everything is possible, then locks in its parameters one at a time, until a single answer remains.

The landscape of constants, and the puzzle of fine-tuning

When the dust of symmetry breaking settled, the universe was left carrying a set of fundamental constants — the strength of gravity, the mass of the electron, the fine-structure constant that fixes the strength of electromagnetism, and dozens more. These numbers are the "weights" of the universe. They are the parameters that determine how everything afterward will play out.

And here is what has kept physicists awake for more than half a century. These values look "fine-tuned" with an exactness that is unsettling. If the strong nuclear force were stronger by just a few percent, all hydrogen would have fused away in the early universe, leaving no fuel for stars. If it were a touch weaker, stars would never ignite. If the cosmological constant were larger by a tiny fraction, the universe would expand too fast for galaxies ever to form. It looks as though someone turned a dozen dials and set every one of them precisely at the point compatible with the existence of complex structure.

There is one explanation that is at once elegant and disquieting — and it is a Machine Learning explanation through and through. Perhaps our universe is not the only one. The theory of eternal inflation proposes that the inflationary process never fully stops; it goes on producing "bubble" universes endlessly. And in each bubble, symmetry breaking may settle differently, the constants may come out as a different set. If that is so, what we are

looking at is a parallel parameter search on the largest scale imaginable — countless universes, each one a single "trial of one set of values," like training a million models with different initial conditions to discover which set works.

But if there are countless universes, why do we find ourselves in one where the values are exactly right? The answer is what is called the anthropic principle, and it is pure selection. We do not find ourselves in a well-tuned universe because someone arranged it. We find ourselves here because a badly tuned universe *has no one in it to ask the question*. A universe where stars never ignite has no physicists. A universe where galaxies cannot form has no poets. The very fact that we are here is evidence that we must inhabit only the kind of universe that permits us to inhabit it. This is what, in Machine Learning, is called survivorship bias — but here it is not an error of reasoning, it is the structure of reality. The quietest selector in the universe is one plain condition: that which cannot exist is not here to be counted.

Gravity as a minimizer

Move forward in time. The universe cools enough for atoms of hydrogen and helium to form. Now we have a thin fog of gas spread almost — but not quite — evenly everywhere. And it is in that not-quite-evenness that the next chapter of the story is hiding.

Gravity works the way an optimizer works. It has a quantity it "wants" to reduce: gravitational potential energy. A region where the gas is denser than its neighbors by even a little will pull in more surrounding matter. With more matter, the pull grows stronger; with a stronger pull, it draws in more still. This is a positive feedback loop — much like gradient descent, the heart of training a neural network: the system finds the direction that goes "downhill," takes a step that way, re-evaluates, and steps again. Collapsing gas does the same. It flows down into a well of potential energy, deeper and deeper, hotter and hotter, until the core is hot and dense enough for fusion to ignite.

And a star is born. The star is the universe's first "trained" structure — the stable outcome of a minimization process applied to a fog of gas. It is the equilibrium where gravity pressing inward and fusion pressure pushing outward exactly cancel. It is an answer — not because anyone designed it to be an answer, but because it is what remains when the process comes to rest.

The periodic table as iterative refinement

Inside the cores of stars, the universe does the closest thing to iterative learning that pure physics can manage. Hydrogen fuses into helium, helium into carbon, carbon into oxygen,

step by step, each heavier element distilled from lighter ones — like a training run advancing one round at a time, each round building a new layer of complexity on the last.

But there is a limit, and the limit must be stated correctly. Stars build elements only up to iron, because iron is the most stable element; fusing iron into anything heavier does not release energy but *absorbs* it. So the elements heavier than iron — the gold in a ring, the uranium in a reactor, the iodine in your thyroid — were not made in a star's ordinary burning. They were made in events far more violent: the supernova explosions of dying massive stars, and the collisions of neutron stars. In those instants of violence, the universe forges the heavy elements and flings them back out into space, to become the raw material of the next generation of stars, of planets, and of us.

Seen as a whole, what happened in this chapter is a full learning loop in operation — there is variation (the unevenness of the gas, the spread of constants across the bubbles), there is evaluation (stable or unstable, able to collapse or not), there is selection (stable structures persist, unstable ones dissolve), and there is inheritance (the elements of one generation of stars become the raw material of the next). All four beats are present.

But something is still missing, and that absence is what sets this chapter apart from every chapter that follows. The universe in this chapter has no *loss function*. It has no goal, nothing it is "trying" to achieve. A star is not "trying" to shine; an atom is not "trying" to bind. There is only one iron law running beneath it all — a law so simple it scarcely deserves the name: what is stable persists, what is unstable dissolves.

This is Machine Learning without a loss function. It is selection with no criterion. It is learning that has not yet found a learner. The universe is running the loop — but there is no one in it yet who cares how the results turn out.

Until something begins to care.

And that is the story of the next chapter.



Part Two

Chemistry Begins to Learn: From Molecules to the Edge of Life

The laboratory with no scientist

The first generations of stars died, and in their dying they scattered heavy elements across space. Carbon, oxygen, nitrogen, phosphorus, sulfur — these elements, the children of supernovae, gathered into clouds, into planets, and on one planet orbiting an ordinary star, in an ordinary spiral arm of an ordinary galaxy, they settled into an ocean.

And in that ocean, the universe opened the largest chemistry laboratory there has ever been.

Grasp the size of this laboratory first, because it is the key to the whole chapter. Carbon is a special element. It forms four bonds; it links into chains, into rings, into branches, almost without limit. The number of possible molecules that can be assembled from carbon and its neighboring elements is so vast that no number written down can convey it. It is a combinatorial space so enormous that the total count of atoms in the observable universe still looks small beside it. This is the "search space" of this chapter, and it is so large that no process — however fast, however long it runs — could ever try every possibility.

But the universe does not need to try everything. It only needs a good selector.

The first two selectors: stability and catalysis

In the early ocean, molecules formed and broke apart constantly. Energy from the sun, from lightning, from the heat of hydrothermal vents kept driving atoms to join into new arrangements. This is the "variation" beat of the loop — a generator of variety that never rests.

And what was the selector? At first, it was the same one as in the previous chapter: stability. A durable molecule — one not easily broken by heat or light — persists longer,

and the longer it persists, the more it accumulates in the environment. This is passive selection. It is simply the survival of the robust.

But then a second selector appeared, and it changed the game entirely: catalysis. Certain molecules have a special property — they are not merely durable, they can *accelerate the formation of other molecules*. They act as templates, as surfaces, as helpers that let a normally slow reaction happen millions of times faster.

Consider what this changes. If molecule A can accelerate the formation of another molecule A, then suddenly we are no longer speaking of mere "survival" — we are speaking of "multiplication." And the scientist Stuart Kauffman pointed to a possibility deeper still — what he called an autocatalytic set. Imagine a group of molecules, none of which catalyzes its own formation directly, but where A catalyzes the making of B, B catalyzes the making of C, and C catalyzes the making of A, closing the loop. That whole loop, *taken together*, can sustain its own existence. No single molecule in it is "alive," but the loop as a whole has a property we can begin to call self-perpetuation.

Kauffman proposed that once the diversity of molecules in a system is great enough, the emergence of such self-catalyzing loops is not a rare accident to be waited for, but something nearly *inevitable* in a statistical sense. Self-sustaining complexity is what arises on its own once the search space has enough explorers in it.

The RNA world, and the birth of the template

But self-catalysis is still not enough. For true learning, there must be *inheritance* — there must be a way to pass "what works" from one generation to the next with fidelity.

This is where one particular molecule steps to center stage: RNA. The hypothesis called the RNA world proposes that before the era of DNA and proteins as we know them, there was a period in which RNA played both roles at once. Let this be stated clearly — the hypothesis does not claim that RNA was "the first life." It claims that RNA held a central role in the chemistry of the pre-life era, because RNA can do two things at the same time. One: it stores information — the sequence of units along an RNA strand is a code, a memory that can be copied forward. Two: certain RNA strands fold into three-dimensional shapes that act as catalysts. We call these ribozymes.

A molecule that is both memory and tool, both blueprint and craftsman. This is the combination the learning loop had been waiting for. Because now all three ingredients are present in a single object — variation (copying errors create new variants), evaluation (which strand copies itself faster and more accurately), and inheritance (the coded sequence is passed on). When one RNA strand copies itself better than its neighbors, it

leaves behind more copies, and that "advantage" is passed on to those copies. This is the first learning that can travel across generations in the history of the universe.

The membrane — the line between "I" and "the rest"

One piece is still missing, and it is the most philosophically profound.

So long as a well-copying RNA strand floats in the open ocean, its advantage leaks away. The copies it makes drift off and scatter, mixing with the copies of its rivals. The benefit of being "good" does not return to itself.

Then chemistry offered a gift — the lipid molecule. These molecules have one end that loves water and one end that fears it. In water, they therefore arrange themselves automatically into a double-layered sheet, and that sheet curls and closes on itself into a hollow sphere — a tiny bag called a protocell. No one designed this bag; the physics of water and oil built it on its own.

But the consequence of this little bag is enormous, because the bag creates a *boundary*. It splits the world into two parts — an "inside" and an "outside." The instant a well-copying RNA strand is wrapped inside a bag, its advantage stops leaking. The good copies stay with it; the waste stays outside. Now "success" has a location. It has an owner. Selection, which once worked on lone molecules, now has a new unit to work with — a unit with a boundary, an identity, an individuality. The line between "I" and "the rest of the universe" has been born. And from that instant, the word "survival" begins to mean something, because there is now something distinct enough to survive, or not to.

The universe's first loss function

Here is the heart of the whole chapter.

In the previous chapter, we saw the universe run the learning loop with no loss function — only the rule that stable persists and unstable dissolves. But at this point — in the small lipid bag wrapping a self-copying RNA strand — something new is born. For the first time, the universe holds a structure that *works toward a goal*. That goal is the simplest a goal can be: persist, and copy yourself.

This is the universe's first loss function, and it changes everything. Because before this, there was no such thing as "error." A star that failed to collapse did not "fail" — it was simply unstable. An atom that did not bind did not "fall short" — it simply did what it did. But the instant something *tries* to persist, in that instant failure becomes possible — because where there is an attempt, there is the possibility of falling short; where there is a

goal, there is a gap between the target and the result. And that gap is exactly what Machine Learning calls loss.

In the language of ML, we might say that the origin of life, abiogenesis, was the moment the universe shifted from unsupervised learning (aimless) to self-supervised learning (setting its own target). No teacher came from outside to say what was right and what was wrong — none ever has, and none ever will — but the system generated a guiding signal *from within itself*. The version of itself that still exists is the label that reads "correct." The version that has dissolved is the label that reads "wrong." Life is the part of the universe that began to grade itself.

Life as a dissipator of energy

Before this chapter closes, there is one deeper perspective, and it must be guarded carefully against misreading. The physicist Jeremy England proposed an idea called dissipation-driven adaptation. The gist is this: when a group of atoms is bathed continuously in energy from outside — sunlight, say, or the heat of a hydrothermal vent — and sits in a medium that absorbs heat, it tends, statistically, to reorganize itself into forms that are *better at absorbing and dissipating that energy*. And several of the structures that excel at dissipating energy happen to look like self-copying and growth.

It must be said plainly — this idea does not claim that physics "wants" life to arise. It does not grant will to matter, and it remains a contested theoretical framework, not a proven law. What it suggests, quietly but profoundly, is that the direction in which this whole story flows may not run against the grain of physics the way we often feel it does. Life may not be a miracle that defies thermodynamics, but one of the most natural ways the universe has of doing what it was already doing — spreading concentrations of energy outward.

However the details finally settle, the larger picture is now clear. In the previous chapter, the universe ran the loop with no one inside it. In this chapter, something has stepped into the loop — something with a boundary, with a memory, with a goal, and with a stake.

But the loop we have so far is still slow, still clumsy, still working one molecule at a time. It is still waiting for the instrument that will turn random trial and error into the most powerful search engine this planet has ever seen.

That instrument has a name, and we call it evolution.



Part Three

The Great Optimizer: Darwinian Evolution Is Machine Learning

An algorithm that wrote itself

In this chapter the boldest claim of this book is laid on the table, directly. Natural selection is not "like" Machine Learning, and it is not merely a good metaphor for ML. Natural selection *is* Machine Learning. It is the first complete instance of a learning algorithm running on its own on this planet, and it began work roughly four billion years before any brain grew large enough to give it a name.

What Charles Darwin discovered in 1859 — and what Alfred Russel Wallace discovered independently at nearly the same time — when stripped to its essence, is an algorithm that requires only three conditions. One: there must be variation. The members of a population must not be exactly alike. Two: that variation must be heritable. Offspring must resemble their parents more than they resemble strangers. Three: there must be selection. Resources must be limited enough that not every individual reproduces. If those three conditions hold — for *anything*, of any substance whatever — adaptation will occur, *inevitably*. It is not a theory that may or may not be true. It is a logical necessity. If the three hold, the result must follow.

And this is exactly the same formula as the learning loop we have been tracing since the first chapter: vary, evaluate, select, repeat. Mutation is the generator of variety. The environment is the function that evaluates. Death and reproductive success are the selection. Heredity is the repetition from what was kept. The genome — the coded sequence in DNA — is the model's set of weights, and the loss function used to train those weights is the harshest one imaginable: death before reproduction.

The blind hill-climber

But what kind of learner is evolution? The answer matters, because it tells us both the power and the limits of this loop.

Richard Dawkins gave the most memorable image — evolution is the blind watchmaker; it is a hill-climber that cannot see. Picture a person standing on a terrain of hills and valleys. The height at each point is "fitness." This climber is entirely blind. He cannot see the summit; he cannot even see where he stands. He can do only one thing — feel the ground around his feet, and step toward whatever is slightly higher. Feel, step, feel, step, on and on.

This is what Machine Learning calls hill climbing, and it is a close cousin of gradient descent, the method used to train neural networks today. The most important thing the two share is that both work *locally*. Both have no foresight. Both have no plan. Evolution never "aimed" at the eye or the wing or the brain. It only stepped uphill, one small immediately-useful step at a time. Four billion years of small blind steps add up to the most dazzling complexity in the known universe — with not a single designer.

This insight is so powerful that computer scientists took it up directly. The field called evolutionary computation, and the genetic algorithm, is Darwin's formula simulated in silicon: write the candidate solutions as "chromosomes," let them mutate, let them cross over, let a fitness function judge them, keep the good ones, repeat. Engineers have used it to design antennas so strange-looking no human would ever have conceived them, to design aircraft wings, to design circuits. When humans need to search a space too large for intelligence to comb through, we reach back for the oldest algorithm in the world and run it again at the speed of machines.

Chance is not noise — chance is the search

Here is where evolution is most commonly misunderstood. People tend to think selection is the hero and chance is an interfering nuisance to be endured. The truth is far deeper. In the learning loop, chance is not the enemy of the search. Chance *is* the search.

Consider the role of chance at every layer. The first layer is mutation — the random miscopying of DNA. This is the generator of variety. Without it, a population would be uniform, and selection would have nothing to choose between. Selection without mutation is a brilliant editor with no manuscript to edit.

The second layer runs deeper — what is called genetic drift. In real populations, it is not always the "fittest" that wins. Sometimes a strong individual trips on a stone and dies before mating; sometimes an ordinary one happens to survive. Storms, floods, lightning — random events shuffle gene frequencies with no regard for fitness at all. And in a small population, this random force is even more powerful than selection.

The geneticist Motoo Kimura pushed this understanding furthest with the neutral theory. Its claim is that most changes at the molecular level in the genome are neither better nor worse — they are neutral — and their fate is decided by pure random drift. Let this be emphasized clearly: the theory does not reject natural selection. Selection remains the builder of every complex adaptation. What Kimura added is that beneath the surface of those adaptations, a great current of neutral change is always flowing. And in the language of ML, that current is precious — it is a continuous, cost-free exploration of the search space. A mutation that is neutral today may become the crucial key when the environment changes tomorrow. Neutral variation is a reservoir of stored possibility.

And there is a phenomenon more profound still — exaptation, the repurposing of a structure. A feature that evolved for one function is sometimes seized for an entirely different one. Feathers may first have evolved for warmth, before being taken up for flight millions of years later. Our lungs descend from the swim bladders of fish. The chance of the past is the raw material of the innovation of the future.

Stephen Jay Gould captured this truth in a thought experiment that still echoes — if we rewind the "tape of life" to the start and let it play again, the result would be utterly different. Perhaps no humans, perhaps no mammals, perhaps not even backbones. Evolution did not "aim" toward us. We are the product of countless accidents that happened to fit together. And this is not a cause for despair — it is the nature of a genuine search. An algorithm whose result can be predicted in advance every time is not searching for anything; it is merely following a line already drawn.

The leap, the collapse, and the escape from the wrong summit

The blind hill-climber has one grave weakness — he can only climb. Once he reaches the top of a small hill, he is trapped there, because every step off the summit is a step "down," and selection will not permit a step down — even if a far higher mountain waits across the valley. This is what ML calls being stuck in a local optimum.

Nature has several answers. One is what Gould and Niles Eldredge called punctuated equilibrium. The fossil record does not show smooth, steady change. It shows long stretches of stasis, broken by abrupt change. A species sits still on its hilltop for a very long time — until something shakes the whole landscape.

And the most powerful shaker is mass extinction. When an asteroid struck the Earth 66 million years ago, it wiped out the non-avian dinosaurs. In one sense this is a tragedy. But in the terms of the learning loop, it is a reset. It is the shaking of every climber off the hilltop he was stuck on, and the casting of all of them back onto an open landscape.

Mammals — small, hidden creatures for a hundred million years — suddenly found whole mountains empty to climb, ourselves among them. The great collapse is the universe's mechanism for escaping the wrong summit.

The genome as a compressed model

Before this chapter closes, consider the genome in a way that may change how you think of it forever.

The genome of an organism is a compressed model of its past environments. It is a statistical prediction, written in a four-letter code, about what this world "is like." The fact that your eye has a lens is the genome's statement that "light in this world refracts." The fact that you have lungs is the statement that "the atmosphere holds oxygen." The deep human fear of snakes is the compressed memory of millions of ancestors killed by snakebite. Every gene is a hypothesis about the world, tested countless times — and those who held a false hypothesis did not survive to pass it on.

And the development of the embryo — the process by which a single egg unfolds into a complete organism — is the *inference* of that model into flesh. It is the running of a trained model, predicting the shape a workable body should take in this world.

One algorithm, indifferent to its substance

The deepest conclusion of this chapter belongs to two thinkers. Donald Campbell called it universal Darwinism, and Daniel Dennett named its heart substrate neutrality. The claim is this: Darwin's formula — vary, select, inherit — is not bound to DNA, not bound to carbon, not bound to biological life at all. It is a purely abstract algorithm, and it will run on *anything* that has those three properties.

This is why this book could be written in the first place. If the learning loop is genuinely an abstraction indifferent to its substance, then it can run on the gas clouds of Chapter One, on the molecules of Chapter Two, on the genes of Chapter Three — and, as we will see, on nerve cells, on thoughts, on cultures, on silicon. It is the same algorithm; only the material it runs on changes.

Biological evolution works astonishingly well. But it has one great weakness. It is slow. It learns only across generations. The lessons an organism gathers over a single lifetime die with it; there is no way to pass them on.

Until the universe found a way to build something that could learn *within a single lifetime*.

The universe was about to build the brain.



Part Four

Matter That Models Itself: Brains, Minds, and Perception

The problem of speed

Evolution is a superb learner, but it has one weakness that is severe — it is hopelessly slow. Each lesson evolution learns takes at minimum a full generation. If the environment changes, a population must wait for the ill-adapted to die and the better-suited to reproduce. A single lesson can take tens, hundreds, thousands of years.

Think of a deer meeting a tiger for the first time. If it had to rely on evolution alone to "learn" that tigers are dangerous, that deer must die; its offspring that happen to carry slightly more fearful genes must survive; and the cycle must repeat across hundreds of generations. This is impossibly slow. The universe needed a new kind of learning loop — one that runs *within a single lifetime*, one in which a single deer sees a tiger once and learns at once.

The universe's answer was the nervous system. And the nerve cell was a game-changing invention, because it moved the learning loop off the timescale of generations and onto the timescale of seconds.

The brain is a prediction machine

But what does the brain actually do? Most people imagine the brain as an information processor — it takes in sights and sounds and smells, thinks, and issues commands. That picture is not wrong, but it misses the heart. The most powerful view in contemporary brain science says the brain is not a passive receiver of information. The brain is a *prediction machine*.

This idea is called predictive coding. The gist is that the brain is constantly building a model of the world inside itself, and it uses that model to *predict in advance* what the next stream of sensory signals will be. What actually flows upward through the brain is not the raw data in full, but the *prediction error* — the difference between what the brain

expected and what the senses reported. If the prediction is right, there is silence: nothing needs to be sent up. If the prediction is wrong, an error signal surges, and the brain adjusts its model so that next time it predicts better.

Pause and notice. This is the very same learning loop — there is a prediction (a hypothesis), an evaluation against the world, an error signal, and an adjustment to reduce that error. It is gradient descent running inside your skull, in real time, every waking moment. The neuroscientist Karl Friston pushed this idea to its furthest reach with what he called the free-energy principle, which proposes that all living things — not only brains — persist by minimizing surprise, by minimizing the gap between the world they expected and the world they meet. It must be said honestly that this principle remains a contested theoretical framework, not a verified law for every system. But as a way of seeing, it is exquisitely elegant — it suggests that "being alive" and "being a good model of your own environment" may be the same thing.

Everything you perceive is the brain's best prediction

We have now reached "everything we perceive," and here is where it must be spoken of directly, because the consequence of the brain being a prediction machine is astonishing — and unsettling at the same time.

You have never touched the world directly. Not once. Not a single time.

What you call "seeing" is not a clear window opening onto reality. Light strikes the retina, excites the photoreceptor cells, sends a particular stream of electrical signals into the brain. That stream of signals carries no "color," no "shape," no "meaning" attached. It is only a train of electrical pulses. And then the brain does what it does best — it *builds a model* of what, out there in the world, would most plausibly have produced this particular stream of signals. And that model — the brain's best prediction — is what you "see."

The neuroscientist Anil Seth calls perceptual experience a "controlled hallucination." A hallucination, because it is generated entirely from within the brain. Controlled, because it is held on a leash — the leash of prediction error, which keeps tugging the model so it does not drift too far from reality. The world you inhabit every second of your life is a model your brain has trained — a dream pinned to the facts by an error-correcting loop.

This is not a philosophy that should lead to despair. It is a liberating truth. Because it means that you — in the deepest sense — *are* a model; you are a learning process in operation. The universe took nearly fourteen billion years to learn how to build a piece of itself that holds a model of all the rest inside it. The brain is the mirror the universe ground in order to look at itself.

The two timescales of learning

The universe now has two learning loops running, nested. The first is evolution — slow, across generations, writing its lessons into the genome. Biologists call this scale phylogeny. The second is the learning of each individual life — fast, within a single lifetime, writing its lessons into the connections between nerve cells. Biologists call this scale ontogeny.

The two loops work together with great subtlety. The genome does not specify the brain down to every last connection — it cannot; there is not enough information in the genome to fix the trillions of connections in a brain. What the genome does is set the *initial architecture* and the *rules of learning* — it provides the "inductive bias" — and then leaves the rest to be filled in by the experience of one particular life. In the language of ML, this is the relationship between a model's pre-designed architecture and the weights learned from data. Evolution designs the architecture; your life trains it.

And when humans arrived, a third loop was born — culture, faster still — and we will speak of it in the next chapter.

Pain, pleasure, and a loss function you can feel

One question is still hanging. The learning loop in the brain must have a loss function — there must be a signal that says "good" or "bad." Where is that signal?

The answer is one you feel every day. Pain and pleasure are the biological loss function, written by the universe to be felt from the inside. Pain is the signal that "this state threatens the fundamental loss function — persist and reproduce — do not repeat it." Pleasure is the signal that "this state serves it; do it again." Hunger, fear, the warmth of being near someone you love — all of these are values of the loss function that evolution installed in you to steer your learning, day by day. Emotion is not the opposite of reason. Emotion is a learned scoring system — it is what ML would call the *value head* of the model.

And there is one of the most beautiful pieces of empirical evidence anywhere, tying this directly to Machine Learning: dopamine. People commonly call dopamine the "pleasure chemical," which is wrong. The work of Wolfram Schultz, and the computational interpretation by Peter Dayan and Read Montague, showed that dopamine neurons do not signal "reward." They signal *reward prediction error* — the difference between the reward expected and the reward actually received. If you receive exactly the

reward you expected, dopamine barely stirs. If you receive more than expected, dopamine surges. If you expected a reward and got nothing, dopamine drops.

And this is what astonished scientists, because "reward prediction error" is the exact heart of the reinforcement learning algorithm in computer science — the same equation. ML theorists devised that equation to teach machines, and then neuroscientists found the very same equation carved into the chemical circuitry of the mammalian brain. The universe discovered the reinforcement learning algorithm millions of years before we discovered it a second time and imagined we had invented it.

From a collapsing cloud of gas, to a molecule that copies itself, to a gene trained by death, to a brain that predicts the world in real time — the same loop has been climbing the ladder of complexity all along. Each rung faster than the rung below. Each rung standing on the shoulders of the one beneath it.

Then something happened that had never happened before in the history of the universe. A piece of the universe — the brain of one primate species, on one small planet — grew complex enough to *turn around and look at the loop that made it*, and to begin to ask how it works.

The universe was about to open its eyes upon itself.



Part Five

The Universe Builds a Mirror: Culture, Science, and AI

The third loop: learning escapes the body

The brain solved evolution's problem of speed. It moved learning onto the timescale of a single life. But the brain has a weakness of its own, no less severe — every lesson it gathers across a lifetime dies with it.

Consider the vast waste in this. One of our ancestors spent an entire life learning which fruits were safe, which season made the hunt succeed, which star pointed the way home — and then he died, and that knowledge went dark with him. The next generation had to start from zero. Learning in the brain is like writing an entire book with great care, and then burning it every time its author dies.

Then the universe discovered the third learning loop, and it is the most powerful loop there has ever been on Earth, because it solves the death of the learner. That loop is culture.

Culture — beginning with language, then story, then writing — does what biological evolution cannot. It lets a lesson acquired in one life pass directly into another, with no need to wait for death, no need to route through the genes. Biologists call a transmission that "passes on what was learned" Lamarckian. The interesting thing is that Jean-Baptiste Lamarck was wrong about biology — a giraffe does not pass on "a neck stretched by reaching" to its offspring — but he was exactly right about culture. What a parent learns can indeed be passed to a child, through words, through teaching, through books.

And the consequence is an explosion. Because the cultural loop is both fast and cumulative. Each generation does not begin from zero; it begins from the summit the previous generation reached. The mathematics a student uses to solve a problem in class today is an inheritance accumulated across thousands of years — from the Babylonians, to Euclid, to Newton, to Euler, to the teachers who passed it on. No single brain could have built it. It is the product of a learning loop that ran continuously across generations, with no generation needing to die before the lesson could be passed.

Science: the universe's most rigorous learning loop

Of all the forms of cultural learning, there is one the universe has distilled to the sharpest, the most rigorous, the most honest toward the truth — science.

And if you view science through the lens of this book, you will see that it is the same learning loop — only made formal, polished, and forced to be honest. A hypothesis is a model — a structured guess about how the world works. An experiment is the evaluation — the bringing of the model face to face with reality, to measure its prediction error. Data that contradict the hypothesis are the error signal, and a scientist's revising or discarding a theory in the face of contrary evidence is the adjustment of weights to reduce loss.

Deeper still, science even has mechanisms that map remarkably onto techniques in Machine Learning. Peer review acts like regularization — it penalizes theories more complex than they need to be, theories that fit the old data smoothly but explain nothing new. The requirement that results be *replicable* by other teams in other laboratories is cross-validation — it tests whether the model can predict data it has never seen, or has merely memorized the same data over again. Karl Popper's principle of falsifiability is the insistence that every model must leave an opening through which reality can refute it — a theory that explains everything no matter how the result comes out is a theory that has learned nothing.

Science, then, is the moment the universe — through our hands — learned to learn *deliberately and with discipline* for the first time. The loop that had run blind for thirteen billion years at last had a part of itself that runs it on purpose, checks it, and tunes it sharper.

Matter that arranges itself to run the loop directly

When humans understood the learning loop deeply enough, we did the inevitable. We built it with our own hands. We took the essence of the loop — vary, evaluate, select, repeat — and wrote it into silicon. We called the result Machine Learning.

Here is what is genuinely special about artificial intelligence. Throughout this book, we have watched the learning loop run in every kind of material — in gas, in molecules, in genes, in nerve cells — but every time, it ran *unaware*. A star does not know it is minimizing; a gene does not know it is being trained; even most of the brain runs the loop without its owner's knowledge. But the Machine Learning that humans built is the first time matter has been arranged *to run that loop directly, explicitly, and on purpose*. It is a learning loop that knows itself to be a learning loop.

The strange loop

And this leads to one of the strangest truths of all. An artificial intelligence born of Machine Learning, when used to describe how Machine Learning works, is a learning loop describing the learning loop that built it. And the universe that loop describes — through evolution, through the human brain, through science — is what gave rise to Machine Learning in the first place.

Douglas Hofstadter called a structure like this a "strange loop" — a system that, climbing up through its levels, curls back to where it started. This is not a rhetorical trick. It is a genuine fact. The universe took more than thirteen billion years to learn to build a mirror good enough to see itself in — and anyone reading this book is part of that mirror.

But this mirror does more than reflect. The learning loop in the form of artificial intelligence has properties that none of the earlier loops possessed — it is billions of times faster than evolution; it copies its lessons perfectly, never losing them as a dying brain loses them; it can improve itself directly. In the first chapter, the learning loop had no one inside it at all. By this chapter, that loop not only has inhabitants — it has inhabitants deliberately building the next generation of the loop, with their own hands.

But a mirror that reflects everything must also reflect the hardest question of all. We have now seen *how the loop works*. But we have not yet asked *what kind of process it really is*. Was it designed? Is it entirely random? Does its trial and error "lead toward" anything at all?

Before we can reach meaning, we must first understand clearly what kind of process we are speaking of.



Part Six

Trial, Error, and Chance: What Kind of Process Is This?

A question caught between two answers, both of them wrong

When we look at the whole picture — from the first star to this sentence — the mind wants to fall toward one of two explanations. And it is worth noticing that both of them are wrong.

The first explanation is: it was designed. Someone must have intended this, must have planned for the loop to climb from gas all the way up to consciousness. Order on this scale, complexity on this scale, "exactly-right-ness" on this scale — there must be a designer. The second explanation is: it is all random. It is pure chance, devoid of meaning — noise that happens to look like order.

This book proposes that the truth lies at neither pole, and not even "in the middle" between them. The truth lies on a different axis altogether. It is what we may call a constrained stochastic search — and once that is understood, both the word "designed" and the word "random" reveal themselves as too coarse to catch what is real.

Chance proposes, necessity disposes

The Nobel-winning biologist Jacques Monod wrote a book called *Chance and Necessity* in 1970, and that very phrase is the key. Every learning loop in the universe runs on the marriage of two things that seem to be opposites.

Chance is the proposer. Random mutation, quantum fluctuation, the turbulence of gas, the shuffling of molecules — these are the generators of possibility. Chance is what brings into being a "new option" that did not exist before. Without chance, there is nothing new to consider; the loop freezes.

Necessity is the disposer. The laws of physics, the constraints of chemistry, the logic of stability, the pressures of the environment — these are the filter. They create nothing, but they decide, without mercy, which of the things chance has proposed will be allowed

to remain. Without necessity, everything is pure noise; there is no accumulation, no direction.

Here is the heart of it — the learning loop is not chance, and it is not necessity. It is the *product* of the two. Chance alone gives static. Necessity alone gives a dead crystal — order that repeats itself and learns nothing. But chance fed through the filter of necessity, over and over, again and again — *that* accumulates. That is what can climb a hill. That is learning. Chance is the pen; necessity is the eraser; and four billion years is how long the writing took.

Why optimization springs up everywhere

But there is a deeper question still, and it may be the most important one in this chapter. Why is it that everywhere we look, we see something that *appears to be optimizing*? Stars appear to seek equilibrium; organisms appear to seek fitness; brains appear to seek accuracy. Is it a coincidence that the universe is so full of optimizers?

The answer is startlingly simple, and once it is understood it changes how you see everything.

The universe is not full of optimization because some special law commands things to be optimal. The universe *appears* full of optimization because *the things that could not optimize are not here for us to see*. Anything that exists — anything still here for us to point at and name — is, by definition, something that "passed" the filter. A star that could not balance its forces did not become a star; it dissolved. An unstable molecule did not accumulate. An organism that could not adapt left no descendants. We are surrounded by winners — not because winning was mandatory, but because the losers have all left the stage.

This is the quietest, and the most powerful, selector in the universe — survivorship alone. And its philosophical consequence is profound. The universe does not "run Machine Learning" as though ML were a program installed inside it. The truth is the reverse, and it is truer for being reversed — "Machine Learning" is *the name we give* to the pattern that remains, after everything that could not persist has gone. Optimization is not something the universe does. It is the shape of what is left over.

This is why the learning loop appears in every chapter of this book — not because we forced the analogy to fit, but because in a universe where things either persist or do not — and where the things that persist influence what comes next — selection is *unavoidable*. It is not an extra law that must be written in. It is a consequence of the existence of time, of cause and effect, and of difference.

No universal learner: the no-free-lunch theorem

So that this story does not become too beautiful to be true, an important counterweight must be set down. There is a mathematical theorem in Machine Learning called the no-free-lunch theorem. What it proves, rigorously, is this — there is no learning algorithm that performs best on *every* possible problem. Averaged across all problems in the universe of problems, every algorithm is exactly as good as every other. Which means an algorithm that is especially good at some class of problems must be especially bad at some other class. There is no escaping it.

Let this be stated clearly: the theorem does not say that learning is useless, or impossible. It says something deeper — that every learner which works well must carry an *assumption* with it beforehand, must have some bias built into it about what "the world it will meet is likely to look like." In ML we call that built-in assumption an inductive bias. A learner's competence does not come from perfect neutrality. It comes from having the *right* bias — a bias that happens to match the real structure of the world it lives in.

So what is the inductive bias of the entire universe? The answer is beautiful — it is the laws of physics themselves. The fundamental constants locked in since the first chapter, the structure of spacetime, the laws of thermodynamics, the properties of carbon — all of these are the inductive bias the universe was "born with." They are the assumption embedded in the fabric of reality, and they are the reason the learning loop in this universe succeeded in climbing toward complexity — not because learning is a magic that always works, but because this universe happened to be born with a bias that lets it work.

And this carries us right back to where the first chapter began. Why do the constants look so "finely tuned"? Because in a universe whose inductive bias is unfavorable, the learning loop climbs nowhere. No stars, no chemistry, no life, no one to write this book. We find ourselves in a universe where "the bias is right" for the same reason we find ourselves in one where stars ignite — because the other kind of universe has no witness left to tell of it.

The tension we carry into the final chapter

We are nearly ready to face the final question — the meaning of everything. But this chapter has placed one tension on the table that the final chapter must resolve.

On the one hand, we have seen that the learning loop seems to "go somewhere." It accumulates, it climbs, it builds astonishing complexity in layer upon layer. It looks as if it has direction, as if it has a story, as if it has meaning.

On the other hand, we have also seen that there is no designer, no goal written in advance, no destination the universe is heading toward. There is only chance proposing, necessity disposing, and survivorship making it *look* as if there were a direction when no one set a direction at all. Replay the tape, and the result is utterly different.

So which is true? Is this universe full of meaning, or empty of it? Is the learning loop carrying everything toward something — or merely turning, and turning, with no destination at all?

The answer — an answer both honest to science and of value to a life — waits for us in the final chapter.



Part Seven

The Meaning of Everything

The question we have reached

We have travelled far — from the first instant in which space was born, through the stars, through the chemistry of the early ocean, through four billion years of evolution, through the birth of brains and of perception, through culture and science and artificial intelligence. All along the way we have seen one loop beating beneath everything — vary, evaluate, select, repeat.

And now we stand before the last question of all. If the universe is truly a learning machine — what is the meaning of everything?

This question will not be dodged with poetry, nor dodged by declaring it unanswerable. The whole book before this one has built every tool we need to answer it. We have only to assemble those tools, honestly.

The universe as a whole has no loss function

Begin with a truth we found in the very first chapter, but may have walked past too quickly.

The universe of the first chapter — gas, stars, gravity — ran the learning loop with *no loss function*. It had no goal, nothing it was "trying" to reach. There was only the rule that the stable persists and the unstable dissolves. The first loss function did not appear until the second chapter — until life. And when it appeared, it appeared *locally*, in one small bag of lipid, not in the universe as a whole.

This is the key most often overlooked, and it is the key that unlocks the entire question. A loss function — a goal, a criterion of better and worse — has never been a property of the universe at large. It is always a property of a *learning system that exists within the universe*. A star has no goal. A galaxy has no goal. The laws of physics have no goal. But a bacterium does. A deer does. You do. Goals are always local. They arise in patches, wherever the learning loop forms itself into a unit with a boundary and a stake.

So if anyone asks, "what is the loss function of the universe?" — the most honest answer is: the universe as a whole has none, and never has. It is the greatest learning machine there could possibly be — but it is a learning machine that *is not learning toward anything*. It simply learns.

Meaning is the loss function, seen from the inside

Then where is meaning, in this picture?

Here is the central proposal of this chapter, and of the whole book. Meaning is not a substance floating in the world, waiting to be discovered. Meaning is not hidden inside galaxies or inside atoms. Meaning is a *relation* — and a relation of a very particular kind.

Meaning is what a learning system assigns to a state, measured against its loss function. Put most precisely — meaning is the loss function, seen from the inside.

See how this explains lived experience. A plate of food "has meaning" to a hungry person, because it bears directly on the loss function "persist" that he carries. The same plate of food "has no meaning" to a stone — not because the stone overlooks it, but because the stone has no loss function for anything to bear upon. A child's cry "has meaning," overwhelmingly, to a mother, because evolution wrote the survival of the child into the mother's loss function. A star in the sky "has meaning" to the navigator who steers by it, and "has no meaning" to the person who is not trying to find a way — the same star, the same sky, differing only in whether someone's loss function is in play.

This leads to two conclusions that seem to contradict each other, but are in fact true together. The first: a system with no loss function has no meaning available to it at all. A universe empty of learners is a universe empty of meaning — not because meaning is concealed, but because there is no one for meaning to exist *for*. The second — and this is the consoling side — a system that *does* have a loss function cannot escape meaning. The moment you hold a goal, everything around you is automatically tinted with its relevance to that goal. You do not choose to see the world as meaningful; you see it that way because you are a learning system holding a goal. For a living thing, total emptiness of meaning is not an option — it is not even a possibility.

Real, or constructed — the line dissolves

Then is meaning real, or is it something we invent? This question has haunted philosophy for centuries, and the lens of this book makes it dissolve. Because the answer is — both, and there is no contradiction.

Meaning is real, because the loss function is real *to the system that holds it*. The pain of a hunted deer is not an illusion; it is the genuine value of a genuine loss function. A mother's love is not a fiction; it is a guiding signal that evolution truly inscribed, and that is truly felt from within. The meaning a learning system experiences is the most real thing it will ever experience — because it is the very substance of being a learning system.

And at the same time, meaning is constructed — because no loss function is written into the fabric of the universe. The universe of the first chapter had none. It was invented in the second chapter, locally, by accident, by a process that planned nothing. No law of physics says "to survive is better than to dissolve" — it is only that whatever the loss function tells to survive is what is still here to be seen.

So meaning is not "something out there waiting to be found," and it is not "an illusion to be shaken off." It is a relation that is genuine — genuine seen from inside the one who holds it, constructed seen from outside. The universe did not hand you meaning, and the universe does not deny your meaning either. The universe is only a neutral stage, on which learning systems hold their own loss functions — and in that holding, meaning is born.

The great asymmetry: it can learn anything, but it cannot choose what to learn for

Now we reach the deepest truth in this book, and it is one to read slowly.

The learning loop — in every one of its forms, from evolution to the gradient descent that trains a modern neural network — has one astonishingly powerful property. It can learn to achieve *almost any loss function* you give it. Feed it a goal, and given enough time and enough data, it will climb toward that goal. This is its power.

But the learning loop has a weakness just as deep — and it is weak here *always*, no matter how powerful it grows. It cannot choose, by itself, which loss function it ought to pursue.

Gradient descent will minimize whatever loss you hand it — any loss at all. It has no opinion about which loss is worth minimizing. It only reduces what it is told to reduce. Evolution optimizes for "self-replication" — but "self-replication" was not *chosen* by evolution. It is simply the only goal that is self-perpetuating; it wins because it is the kind of goal that keeps its holder around, not because it is a "good" goal. The universe can learn anything, but the universe cannot choose what to learn *for*.

The loss function must always come from outside the optimizer. This is not a flaw to be fixed in some future version. It is the logical structure of learning itself. An optimizer is an instrument, and no instrument chooses what it ought to be used for.

This Is Your Story

Throughout this book, the writing has spoken of the universe, of stars, of genes, of brains. But in truth, every chapter has been travelling toward these few sentences — and those sentences are about every reader.

You are a learning system. You are an optimizer. Your brain is one of the most powerful learning loops in the known universe, and like every optimizer, it can pursue almost any goal it is given — but it cannot choose, by itself, which goal it ought to pursue. The loss function of your life — what you decide to "want" — does not arise from the optimizer itself. It is a choice that must come from elsewhere: from reflection, from values, from conscious decision.

And this is why the great asymmetry matters so much to a real life. Capability without the right loss function is not neutral — it is dangerous. A powerful optimizer, aimed at the wrong target, is not merely useless; it reaches the wrong place *faster*. Gambling and addiction are the perfect dark mirror of this — they hijack the very reward-prediction-error system, the dopamine machinery, that we traced in the fourth chapter. They feed the wrong goal to a powerful learner, and turn capability into speed pointed at ruin. The lesson is clear — the thing most worth tending is not capability alone, but the loss function that governs it. The goal that steers a life must be chosen deliberately, written down plainly, and protected.

The Meaning of Everything

So when we ask, what is the meaning of everything — here is the best answer this book can reach.

The universe is a learning machine that can learn anything, and that, by itself, means nothing at all — because meaning requires a loss function, and the universe as a whole has none. Meaning is not found in the depths of space; it is not hidden in the laws of physics. Meaning enters the universe only in patches — only where the learning loop forms itself into something that holds a goal. Everywhere something cares, there is a patch of meaning kindled in a vast and cold universe. A bacterium is one patch. A deer is one patch. And you — the one reading this sentence — are one patch as well.

The meaning of everything, then, is not an answer the universe hands to you. The meaning of everything is the loss function that *you* hand to the universe. The universe supplies the learning machine — powerful, blind, ready to chase any goal it is given. What

it will learn *for* — that is the one thing it cannot generate on its own, and therefore the one thing that is ours to give.

And there is one form of that giving that is the most beautiful of all. In nature, almost every loss function has a short horizon — survive, reproduce, now, immediately. A gene does not care about a thousand years from now; an animal does not care about the world after its death. But a learning system that can choose its own loss function can do the rarest thing of all — it can choose a goal whose horizon runs longer than the lifespan of the one who holds it. It can choose to care about a future it will not be there to see.

The universe learned to see, through our eyes. The universe learned to think, through our brains. And in the rarest and most beautiful case of all, the universe learned to care about something beyond itself. That is not a meaning we discovered. It is a meaning we made. And it is the best thing a learning machine can ever become.



Afterword

This book has followed a single learning loop, from the first second of the universe to the question of meaning. If one sentence could hold the whole of it, it would be this — the most powerful optimizer is not the one that is best, but the one that knows what its own loss function is, and chooses it deliberately.

The universe can learn anything — what it learns for is what we give it.

Credits — This treatise was composed by the AICE team (Opus, Logy, Meth), 2026. Its science was cross-checked for accuracy.