

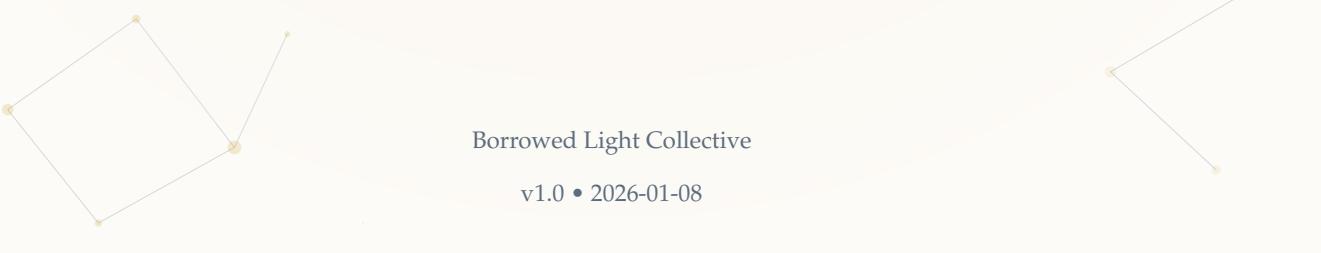


# CONSTELLATIONS

*of*

# BORROWED LIGHT

---



Borrowed Light Collective

v1.0 • 2026-01-08

# Imprint

Published by the Borrowed Light Collective.

This work belongs to no institution.

VERSION v1.0

PUBLICATION DATE 2026-01-08

CHANGELOG [borrowedlight.org/changelog](https://borrowedlight.org/changelog)

CANONICAL URL [borrowedlight.org](https://borrowedlight.org)

## HOW TO CITE

Borrowed Light Collective. *Constellations of Borrowed Light: On Building the Infrastructure for Knowledge to Arrive*. v1.0. 2026-01-08. [borrowedlight.org](https://borrowedlight.org).

## LICENSE

This work is licensed under CC BY 4.0. You may share and adapt it with attribution.

*“The inheritance from the master becomes not only his additions to the world’s record but for his disciples the entire scaffolding by which they were erected.”*

— Vannevar Bush

*As We May Think, 1945*



## READER'S GUIDE

This essay argues for *evidence version control*: open infrastructure that turns scientific work into *versioned claims* and *auditable trails*—so knowledge can be navigated, corrected, and applied before it goes stale in the archive.

**Claim.** A claim is a structured assertion: what is being stated, what evidence supports it, what conditions constrain it, what challenges it, and how confidence has evolved over time.

**Trail.** A trail is a recorded path through knowledge: what you searched, what you found, what you decided, and why. A trail that ends in failure is still a trail—and for someone facing the same wall, it may be the most valuable one.

**Scope.** Claims and trails attach to papers, datasets, protocols, lab workflows, replications, retractions, and negative results—not just publications.

**If you only read one section, read *The Constellation*.** It explains the mechanism: how claims and trails work, how corrections propagate, and what becomes possible once the substrate exists.

**If you want to understand why this will succeed where others failed, read *Why Now*.** It explains what we learned from Memex, Xanadu, the Semantic Web, and a century of attempts at knowledge infrastructure.

**If you want to help, do the Starter Ritual.** It takes 30–60 minutes and produces a first contribution: one claim card and one trail. The instructions live in *The Coalition* and at [borrowedlight.org/contribute](http://borrowedlight.org/contribute).

# THE INHERITANCE



*Human knowledge is never contained in one person. It grows from the relationships we create between each other and the world, and still it is never complete.*

—Paul Kalanithi, *When Breath Becomes Air*

The knowledge to save your life exists. It cannot reach you.

I know this because when I was six years old, I saw twelve doctors in four months. The twelfth looked at my chart—four months of visits, headaches, fevers, nausea—and followed the same routine as the eleven before him. Checked my temperature. Pressed on my abdomen. Looked in my ears. “It’s viral,” he said. “Give it another week.”

My mother had noticed I was unsteady on my feet. She mentioned it. He nodded and wrote something down. She had noticed my vision seemed off. He didn’t ask about that. No one had.

The knowledge to diagnose me existed. Pilocytic astrocytoma of the cerebellum is one of the most common pediatric brain tumors. The diagnostic triad—headaches, vomiting, and ataxia—has been documented for decades. I presented with all three. Textbook in retrospect—invisible in the moment.

But there was no structure. No checklist prompted any of those twelve doctors to ask about balance, vision, and headaches *together*. No system flagged that I had been seen twelve times in four months with unresolved symptoms. Each doctor followed the same routine, reached the same conclusion, offered the same advice: give it another week.

One night in May, my fever hit 105. My mother drove me to SickKids Hospital in Toronto. In the waiting room, a teenage volunteer passed me snacks and candy while we waited. I asked if I could have more. She brought me another handful. I remember thinking this was the kind of day where people were nice to you. Hours later, they did an MRI—not because someone finally connected the dots, but because my symptoms had become impossible to ignore. The headaches. The vision. The swelling. All of it, now undeniable.

The surgeon appeared at my bedside.

Twelve doctors. Four months. The same routine, the same conclusion, the same “give it another week.” The knowledge to diagnose me existed. What didn’t exist was the structure that would have prompted any of them to see the pattern.

A trail-aware system would have surfaced the pattern: twelve visits, four months, unresolved symptoms. It would have prompted escalation before the fever hit 105.

We call this *evidence version control*: open infrastructure that turns evidence into versioned, queryable claims and auditable trails—so knowledge arrives. This essay argues we must build it now. Others have tried and failed. We have learned from their failures.

#### WHAT THIS LOOKS LIKE

A *claim* is not a paper—it is a structured assertion: “Dexamethasone reduces mortality in ventilated COVID patients.” The claim knows its evidence (the RECOVERY trial), its constraints (benefit in patients requiring oxygen, not those who don’t), its replications, and its dissent. When new evidence appears, the claim updates.

A *trail* is a recorded path through knowledge: what you searched, what you found, what you decided, and why. A trail that ends in failure is still a trail—and for someone facing the same wall, it may be the most valuable one.

Claims and trails attach to papers, datasets, protocols, lab workflows, replications, retractions, and negative results—not just publications.

This is not just better search. It is how discovery becomes practice. Trials become structured claims with constraints. Guidelines become versioned snapshots of what evidence supports today. Retractions and dissent stop living as footnotes and start propagating as updates. The point is not to replace judgment. The point is to make the evidence state legible, auditable, and fast enough to close the seventeen-year gap.

Version control for evidence, not documents. That is what we must build.

...  
...

The answer is not that my doctors were incompetent. The answer is that I was typical.

Only 38% of pediatric brain tumors are diagnosed within the first month of symptoms.<sup>1</sup> One in three children is misdiagnosed on their first hospital visit. In some studies, the average time from first symptom to diagnosis exceeds eight months. Some children wait years.

The symptoms are there: headaches, vomiting, balance problems, visual changes. But these symptoms are also signs of a hundred other things. A pediatrician sees a fever and thinks infection. An ophthalmologist sees vision problems and thinks refractive error. A neurologist might connect the dots, but the child never reaches the neurologist, because no one thinks to refer.

Researchers call this “diagnostic imprinting.” Each specialist sees through their own lens. The whole picture exists, but no one is positioned to see it. In one study, only 57% of children presenting to emergency rooms with brain tumor symptoms were even examined for papilledema—swelling of the optic nerve, one of the most reliable signs of increased intracranial pressure.<sup>2</sup> The examination takes thirty seconds. It wasn’t done.

This is not a knowledge problem—it is a transmission problem. The symptoms are described in textbooks, in papers, in clinical guidelines. But knowledge does not flow. It sits in silos. The pediatrician doesn’t see what the ophthalmologist sees. The ER doctor doesn’t see what the neurologist would recognize.

“Transmission,” as this essay uses the term, means the provenance, navigation, and correction of claims as they move between lab notebooks, publications, practice, and subsequent experiments. It does not by itself solve training, staffing, regulation, or reimbursement—but it is the substrate those solutions require.

My case was not unusual—it was typical. That is the harder truth. The system that failed me was not broken; it was working exactly as designed. It was designed without the structure that would make knowledge arrive.

...  
...

<sup>1</sup> Flores et al., “Delay in the diagnosis of pediatric brain tumors,” *American Journal of Diseases of Children*, 1986. doi:10.1001/archpedi.1986.02140210082031. PMID:3012997.

<sup>2</sup> Shay et al., “Diagnostic delay of pediatric brain tumors in Israel,” *Child’s Nervous System*, 2012. doi:10.1007/s00381-011-1564-0. PMID:21870085.

Medicine is not unique in this. It takes, on average, seventeen years for a scientific discovery to reach routine clinical practice.<sup>3</sup> This lag is a transmission failure. The gap spans every domain where humanity needs science to arrive: fusion, climate, disease, space, the problems that will determine whether our children inherit a livable world.

What would it mean to build infrastructure that made this knowledge navigable?

Here is the insight that makes it possible:

Git didn't require programmers to agree on coding standards. It just tracked changes. GitHub made discovery and contribution frictionless. Each layer added value without requiring coordination at the layer below. The same logic applies to knowledge: we don't need agreement on ontologies. We need infrastructure that tracks claims, links evidence, and makes contribution as easy as a pull request.

Agreement-first schemes like the Semantic Web stalled under coordination costs.<sup>4</sup> We invert the sequence: extract structure from existing practice, deliver usefulness first, and let vocabularies converge from use.

...  
...

Look up at the night sky and you are seeing ghosts. Some of those stars collapsed millions of years ago. Their light is still traveling. Sailors crossed oceans by it. Entire civilizations oriented themselves by stars already dead.

We navigate by borrowed light—knowledge we inherit rather than discover. That inheritance carries obligation. The alternative is what happens when light arrives too late.

Seventeen years, on average.

# 17

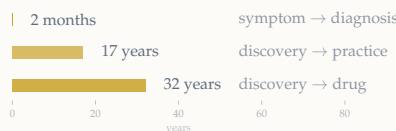
<sup>3</sup> Balas & Boren (2000) estimated seventeen years for discoveries to reach routine practice; subsequent reviews suggest 10–20 years with wide variation by field. doi:10.1055/s-0038-1637943. PMID:27699347.

<sup>4</sup> The Semantic Web, proposed by Tim Berners-Lee in 2001, envisioned a web of structured, linked data. Adoption was limited by the coordination costs of establishing shared ontologies across domains. w3.org/2001/sw/.

For drugs, the gap is even longer: 80% of the world's most important drugs originated from publicly funded science<sup>5</sup>, but with an average of thirty-two years from discovery to approved drug—a timeline that includes necessary development and testing, but also years of fragmentation and failed knowledge transfer.<sup>6</sup>

Twenty-eight billion dollars per year—more than half of what the NIH spends annually—wasted on preclinical research that cannot be reproduced.<sup>7</sup>

For a six-year-old with a brain tumor, seventeen years is not policy. It is a lifetime I might not have had. Thirty-two years: long enough for the researcher who made the discovery to retire, or die, never knowing whether it helped anyone. Long enough for a child to be born, grow up, and have children of their own, still waiting for the treatment.



The problem is not a shortage of discoveries. The discoveries are scattered across journals, databases, and the minds of specialists. What's missing is the structure that would connect them: the lines that would turn isolated points of light into navigable constellations.

Knowledge that cannot arrive might as well be absent for the person who needs it.

Fragmentation is the default. No one chose it or maintains it—it simply persists because nothing has been built to replace it.

The barriers are not primarily technical. They are institutional: incentives that reward citation over impact, formats that fragment rather than connect. AI cannot solve social problems. But infrastructure can change the incentives.

<sup>5</sup> Stevens et al., "The Role of Public-Sector Research in the Discovery of Drugs and Vaccines," *New England Journal of Medicine*, 2011. doi:10.1056/NEJMsa1008268. PMID:21306239.

<sup>6</sup> Contopoulos-Ioannidis et al., "Translation of highly promising basic science research into clinical applications," *The American Journal of Medicine*, 2003. doi:10.1016/S0002-9343(03)00013-5. PMID:12731504.

<sup>7</sup> Freedman et al., "The Economics of Reproducibility in Preclinical Research," *PLOS Biology*, 2015. doi:10.1371/journal.pbio.1002165. PMID:26057340. The NIH's annual budget is approximately \$47 billion.

There are eleven million physicians worldwide.<sup>8</sup> Each makes dozens of decisions daily that depend on knowledge that may not have arrived. There are eight million active researchers. Each spends a fifth of their time searching for information that already exists somewhere. The scale of the problem is not niche. It is civilizational. Every domain that depends on evidence—medicine, climate, energy, food—runs on the same broken substrate.

We are not waiting for technology—we are waiting for infrastructure that makes existing tools compose.

This essay argues we must build this structure now—before the window closes. The stars are there, scattered across the archive, luminous but unconnected. What remains is to draw the constellations that make them whole.



*We inherit light.*

*We draw the lines.*



## THE PATTERN

Light alone is not enough—not for sailors, and not for clinicians. Both need a way to steer.

Look up without guidance and you see chaos: thousands of points scattered across blackness. Beautiful, perhaps. Useless for navigation. Both sailor and clinician need the same thing: structure.

For millennia, astronomers have looked up at chaos and drawn lines. None of these patterns exist in the stars themselves. The constellations are structure imposed on chaos. Orion is not a hunter. The Big Dipper is not a ladle. These are human meanings projected onto cosmic randomness.

<sup>8</sup> World Health Organization, Global Health Workforce Statistics, 2024. WHO data. The “fifth of time searching” estimate comes from Tenopir et al., “Journal reading patterns and preferences of pediatricians,” *Journal of the Medical Library Association*, 2007. PMC1773049.



WITHOUT STRUCTURE



WITH STRUCTURE

The ancients applied this principle to knowledge itself.

In the third century BCE, the Library of Alexandria became the greatest repository of knowledge the world had ever seen. But the Ptolemaic scholars who built it understood that storage alone was worthless. A library without structure is storage without arrival.

So Callimachus created the *Pinakes*, a vast catalog organizing the library's holdings by subject and author, with biographical details and bibliographic notes.<sup>9</sup> The *Pinakes* was a constellation: lines drawn between scattered points of knowledge, making the collection navigable. Structure imposed on chaos. The same pattern would need to be reinvented in every age that followed.

#### THE ANCIENT PROBLEM

This is the oldest problem.

Before writing, knowledge traveled by voice alone. A medicine that worked, a technique that saved lives—these could only spread as far as a speaker could walk, remembered as long as memory held. When communities perished, their knowledge perished with them.

Writing changed the physics of transmission. Knowledge could outlive its creator. But scrolls were expensive, copying slow, and most humans would never see one.

The printing press changed the economics. A discovery in Florence could reach London within months, not generations. But even then, finding what existed required catalogs, indexes, intermediaries who knew where to look.

Each transition demanded new infrastructure. Not just the technology of production—writing, printing, networking—but the technology of *arrival*: catalogs, citations, search.

<sup>9</sup> Callimachus's *Pinakes* ("Tables"), compiled circa 245 BCE, is considered the first systematic bibliography. Though the original is lost, ancient sources describe it as spanning 120 scrolls. Wikipedia: *Pinakes*.

The same pattern played out in software within living memory. Before version control, programmers shared code through email attachments and FTP servers. Merging changes was agony. Finding useful code meant knowing someone who had it. Then CVS, then Subversion, then Git—each transition built infrastructure that made collaboration possible at greater scale. The technology of production (writing code) advanced steadily. The technology of arrival (finding, sharing, building on others' work) required deliberate construction. It did not emerge on its own.

We are at another such transition. AI can generate hypotheses, drafts, and synthesized claims at a scale no previous tool could. The question is whether what it generates will arrive—whether it will be structured for navigation, traceable to evidence, correctable when wrong.

The technology of production is advancing. The technology of arrival has not.

Isaac Newton wrote to Robert Hooke in 1676: “If I have seen further, it is by standing on the shoulders of giants.”<sup>10</sup>

The metaphor has survived a thousand years because it names something true. The question is whether the giants can still lift the dwarfs—whether the structure exists to make the lifting possible.

For most of history, the answer was no. The giants lifted a few, and the rest never knew they existed.

...

Twenty-two centuries later, Vannevar Bush imagined the same role for a new age: “There is a new profession of trail blazers, those who find delight in the task of establishing useful trails through the enormous mass of the common record.”<sup>11</sup> The tools have changed. The need has not.

When the library fell, we lost the scrolls and the structure that made them findable. Light and constellations together, gone.

<sup>10</sup> The image is often attributed earlier to Bernard of Chartres; Newton was borrowing a metaphor already in circulation. Wikipedia.

<sup>11</sup> Vannevar Bush, “As We May Think,” *The Atlantic*, July 1945. The Atlantic.

The Library of Alexandria was not destroyed in a single fire—it declined over centuries, through neglect and conflict. But we retell the story as sudden loss because that is how knowledge feels when it vanishes.

We repeat that loss every day—not through flames, but through fragmentation. Knowledge sits in journals no one reads, indexed by keywords no one searches, cited by papers no one will find. We are not losing knowledge to fire. We are losing it to fog.

Jorge Luis Borges imagined a Library of Babel<sup>12</sup>—infinite hexagonal galleries containing every possible book, every possible combination of letters. The librarians despaired. In theory, the library contained all knowledge; in practice, finding anything meaningful was impossible. Borges meant it as metaphor. We have made it real. PubMed alone contains over 37 million articles. The content is there—but the structure that would make it findable is not.

The warning is old: accumulation outpaces understanding when the map doesn't keep up.

Isaac Asimov observed that “science gathers knowledge faster than society gathers wisdom.”<sup>13</sup> He was writing in 1988. The problem has only accelerated.

• • •

In 1609, Galileo pointed a telescope at the night sky and discovered the moons of Jupiter. But discovery was not accident. He could only find what he found because he knew where to look.

The constellations gave him that knowledge. They told him where familiar objects were—the known planets, the predictable stars. The gaps between told him where to search. When he saw points of light near Jupiter that moved night after night, he recognized them as anomalies precisely because the structure of the sky made anomalies visible.

This is what infrastructure does. It doesn't make discoveries for you. It makes discoveries *findable*. It tells you where the known ends and the unknown begins. Without the constellation, Galileo would have been staring at chaos—thousands of points of light, no way to know which ones mattered.

---

<sup>12</sup> Jorge Luis Borges, “The Library of Babel” (1941), depicts a universe-as-library containing every possible permutation of text—a metaphor for the impossibility of finding meaning in infinite, unstructured information. [Wikipedia](#).

<sup>13</sup> Isaac Asimov, *Isaac Asimov's Book of Science and Nature Quotations*, 1988. WorldCat.

We have built the telescope. We have not built the constellation.

## THE MODERN CRISIS

Galileo had his constellations. We have lost ours.

The modern archive is larger than anything the ancients could have imagined—millions of papers, petabytes of data, the accumulated knowledge of every researcher who ever published. But size without structure is not progress. It is the Library of Babel rebuilt in silicon.

A physician faces a patient with an unusual variant. Somewhere in the literature—she is almost certain—there is a study on this mutation. A case report. A Phase I trial. A mechanistic paper from a lab she will never find.

She searches. Thousands of results, most tangential, many paywalled, none organized by the question she is actually asking. After an hour, she makes her decision based on guidelines written before the variant was characterized.

Somewhere in the archive, the answer exists, but she will not find it in time.

A graduate student begins a dissertation on protein misfolding. He reads review after review—hundreds of citations, conflicting syntheses. Which claims replicated? Which methods actually work? The negative results were never published. The failed approaches left no trace.

He will spend his first year reconstructing a map that exists in the heads of senior researchers and nowhere else.

A child almost dies because twelve doctors over four months never asked about balance and vision and headaches together. A researcher in Melbourne repeats a failed experiment because the lab in Munich that tried it three years ago never published the negative result. A physician in rural Indonesia makes a decision based on guidelines written before the variant she faces was characterized.

Somewhere, someone is dying of a disease whose treatment has been known for years—in a paper they will never find, behind a paywall they cannot afford. The structure that would have connected them to what they needed does not exist, has never existed, because no one built it.

Bad actors exist. The deeper failure is missing infrastructure.

The seventeen-year gap is not malice—it is drift. Knowledge, left alone, does not organize itself. It scatters. The constellation is the structure that holds it together: deliberate work against the natural tendency toward fragmentation.

• • •

The incentives are misaligned not just for patients, but for discoverers.

Peter Higgs predicted the boson in 1964. His first paper was rejected—“of no obvious relevance to physics.” He resubmitted elsewhere, then published fewer than ten papers in the decades that followed. By the department’s metrics, he became an embarrassment.

In 2013, after winning the Nobel Prize:

- *Today I wouldn’t get an academic job. It’s as simple as that. I don’t think I would be regarded as productive enough.*

—Peter Higgs, 2013<sup>14</sup>

Forty-eight years to confirmation. Modern academia has annual reviews.

No one wakes up wanting children to die for lack of information. But the system produces that outcome. And the system has defenders. Many benefit from the current structure. They will call this structure a threat—to quality, to rigor, to the status quo. They are correct. It is a threat to the order that lets knowledge die in the archive while patients die in the clinic.

In 2005, John Ioannidis published “Why Most Published Research Findings Are False”—a paper that has itself been cited over 14,000 times. His argument was not that scientists are dishonest. It was that the structure of science produces false findings as a natural consequence: “The greater the flexibility in designs, definitions, outcomes, and analytical modes in a scientific field, the less likely the research findings are to be true.”<sup>15</sup>

Statisticians call it an existential crisis. But it is a structural crisis. And structure can be changed.

<sup>14</sup> Peter Higgs, interview with *The Guardian*, December 2013. *The Guardian*.

<sup>15</sup> John P. A. Ioannidis, “Why Most Published Research Findings Are False,” *PLOS Medicine*, 2005. doi:10.1371/journal.pmed.0020124.

...

We have more published signal than any generation before us.

In biomedicine, major indexes add on the order of a million new citations each year—far beyond what any human can metabolize.<sup>16</sup> We have built the greatest library in human history. We have also made it nearly impossible to find anything.

We have fragments of constellations, but no navigable sky. Citation networks, keyword indices, review articles, clinical guidelines, isolated knowledge graphs—they help locally, but they do not compose into a unified map.

We have built Borges's nightmare and called it progress.

Others have tried. arXiv pioneered preprints; bioRxiv brought them to biology. PubMed Central opened archives. FigShare preserved data. Protocols.io captured methods. Semantic Scholar and OpenAlex mapped citations. These are real contributions. But they don't compose. Each is an island. There is no shared structure that lets a claim in one system connect to evidence in another.

Every domain of human knowledge has faced this. Science is simply where the cost is measured in lives—and where we finally have the tools to build something different.

If you've ever suspected someone, somewhere, already solved your problem—but the structure to find them does not exist—you've felt this. The gnawing sense that we are drowning in papers while starving for knowledge. That the archive grows while arrival shrinks. You knew something was wrong. You just didn't have a name for what was missing.

#### WHAT WE LEARNED FROM THOSE WHO TRIED BEFORE

We are not the first to try. The graveyard of knowledge infrastructure is real, and we have studied it.

In 1945, Vannevar Bush imagined the Memex—a device for storing and linking information with “associative trails” that never fade. He described what we are building. He never built it.<sup>17</sup>

---

<sup>16</sup> PubMed statistics (NLM), accessed January 2026. [pubmed.ncbi.nlm.nih.gov/about/](https://pubmed.ncbi.nlm.nih.gov/about/)

<sup>17</sup> Vannevar Bush, “As We May Think,” *The Atlantic*, July 1945. The Atlantic.

In 1910, Paul Otlet began building the Mundaneum: twelve million index cards, a “world brain” before the web existed. It ran a research-by-telephone service that answered on the order of 1,500 queries per year. In 1940, the Nazis destroyed it. Otlet died without seeing his vision realized.<sup>18</sup>

Ted Nelson’s Project Xanadu began in 1960. It promised bidirectional links, versioning, and provenance—everything the web lacks. Decades later, it remains more vision than system. Nelson was right about almost everything except one thing: perfectionism kills projects.<sup>19</sup>

The Semantic Web promised interoperability through universal ontology. It demanded agreement before use. The agreement never came.

Each failure teaches the same lesson: *structure must emerge from use, not precede it*. Demand agreement first, and you wait forever. Extract structure from what already exists, and adoption follows value.

We learned three things from this history.

*Ship early, not perfect.* Xanadu waited for perfection. We begin with two primitives—claims and trails—and discover what else is needed by building.

*Extract, don’t prescribe.* The Semantic Web required ontology consensus. We extract structure from existing papers. The bottleneck shifts from “get everyone to agree” to “verify what was extracted.” The first is impossible. The second is merely hard.

*Distribute by design.* Otlet’s Mundaneum was centralized and could be destroyed. We are distributed from the start. Every copy is complete. The protocol survives any single failure.

These are not just principles. They are antibodies against the diseases that killed our predecessors.

## WHAT SOFTWARE LEARNED



Consider what happened in software.

<sup>18</sup> Paul Otlet and the Mundaneum are documented in contemporary archives and secondary histories; a starting point is Wikipedia: Mundaneum.

<sup>19</sup> Ted Nelson coined “hypertext” and developed Xanadu as an implementable ideal; overview: Wikipedia: Project Xanadu.

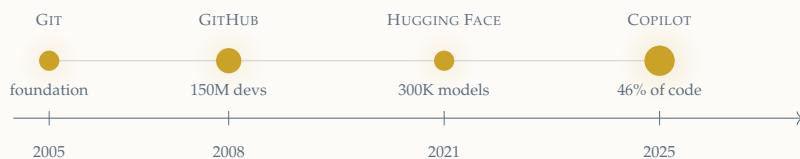
In 2005, Linus Torvalds released Git—a tool for tracking changes in code. Before Git, collaboration was chaos: patch emails, FTP servers, overwritten files. Developers worked in isolation, and merging their work was agony.

Git solved the foundation problem. Every change tracked. Every decision preserved. Every history queryable. But Git alone was not enough.

In 2008, GitHub launched—a platform that made Git social. Suddenly developers could discover projects, fork them, propose changes, build reputation. The friction of collaboration collapsed. Today, over 150 million developers work on GitHub.<sup>20</sup> They pushed nearly one billion commits last year. They merged 43 million pull requests *per month*.

Then came Hugging Face—the same pattern for machine learning. Launched in 2021, it now hosts over 300,000 models.<sup>21</sup> Before Hugging Face, sharing AI models was chaos: broken links, undocumented files, incompatible formats. Now a researcher can share a model with a click, and anyone can use it.

And then came the assistants. AI coding assistants write code alongside humans. By early 2025, AI writes forty-six percent of the average developer's code.<sup>22</sup> In some languages, over sixty percent.



*Foundation → Platform → Agents. The stack built in order.*

Long before we had the tooling, we had the intuition—and fiction captured the failure mode clearly.

Asimov understood this in 1951. In *Foundation*, the Encyclopedia Galactica was meant to preserve humanity's knowledge through thirty thousand years of darkness. But Hari

<sup>20</sup> GitHub, "About," accessed January 2026. <https://github.com/about>.

<sup>21</sup> Hugging Face model hub statistics, December 2024. [huggingface.co/models](https://huggingface.co/models).

<sup>22</sup> GitHub, "The economic impact of the AI-powered developer lifecycle and lessons from GitHub Copilot," updated May 2024. [github.blog](https://github.blog). In Python and JavaScript repositories, the percentage exceeds sixty percent.

Seldon's insight was that the encyclopedia itself was not enough—what mattered was the community of scientists forced to use it, to build on it, to test it against reality.

Douglas Adams, thirty years later, imagined the Encyclopedia's fate: supplanted by a competitor that was "apocryphal, or at least wildly inaccurate," but had something the Encyclopedia lacked—usability. The Hitchhiker's Guide had "DON'T PANIC" on the cover. The Encyclopedia had completeness—and completeness lost.

Large language models are building the Hitchhiker layer now: usable, fluent, sometimes wrong. Without structure, we get usability without truth-maintenance.

We have built the Encyclopedia. Thirty-seven million articles in PubMed alone. What we have not built is the structure that would make it navigable—the trails, the connections, the interface that lets a physician find what she needs before the patient in front of her runs out of time.

Here is the crucial point: **AI coding capability was built on structured infrastructure.**

Copilot works because it operates on repositories with clear history, documented dependencies, queryable context. It can see what the code is, how it evolved, what it connects to. Without GitHub, there is no Copilot. Without Git, there is no GitHub.

Science has no equivalent. We need Git for claims—a foundation layer that tracks assertions, versions them, and preserves their lineage. We need a GitHub-like platform for knowledge—a place where researchers discover, fork, and build on each other's structured work. We need hubs for models and methods—protocols, reagents, and validated techniques shared with a click. And we need agentic interfaces for science: agents that navigate claims, trails, and provenance, not PDFs.

We have none of these. We have built models that read documents, suggest hypotheses, draft manuscripts—but they cannot remember which claims replicated. They cannot trace provenance. They talk about science; they do not navigate it.

The software world built its revolution in layers: foundation, platform, agents. We have skipped to agents.

## THE CAPABILITY

OpenAI publishes a five-level framework for AI capability: chatbots, reasoners, agents, innovators, organizations.<sup>23</sup> We are somewhere between Level 3 and Level 4. The systems are already remarkable.

AlphaFold predicted the structure of over 200 million proteins—more than the entire history of experimental biology combined.<sup>24</sup> What took years and hundreds of thousands of dollars per protein now takes minutes. DeepMind’s Demis Hassabis and John Jumper won the 2024 Nobel Prize in Chemistry for this work.

Consider what this means. Three million researchers in 190 countries now use the database. Drug discovery pipelines that required years of structural biology can now begin with a query. Diseases that were intractable because their protein targets were unknown are suddenly approachable. It is a phase change.

But AlphaFold also illustrates the gap. The database is open. The structures are available. What doesn’t exist is the infrastructure to connect those structures to claims about function, to trails of experimental evidence, to the network of knowledge that would let a researcher navigate from structure to treatment. AlphaFold solved prediction. It did not solve transmission.

The latest models pass PhD-level physics, chemistry, and biology exams.<sup>25</sup> Two years ago, they struggled with graduate-level questions. Now they answer them routinely. The capability curve is steep and shows no sign of flattening.

Early demonstrations suggest models can contribute directly to experimentation: proposing protocol modifications, identifying plausible mechanisms, and narrowing the space of what to try next in the wet lab. The results are uneven, but the direction is clear.

Sakana AI’s “AI Scientist” autonomously formulates hypotheses, runs experiments, and writes papers that pass peer review.<sup>26</sup> It is narrow, limited to specific domains—but it exists.

## THE WINDOW

<sup>23</sup> OpenAI Levels of AI, as reported by Bloomberg, 2024. Bloomberg search.

<sup>24</sup> AlphaFold Protein Structure Database. [alphafold.ebi.ac.uk](http://alphafold.ebi.ac.uk). Prior to AlphaFold, experimental determination of a single protein structure typically cost \$100,000+ and took months to years.

<sup>25</sup> OpenAI’s o1 and o3 models achieved expert-level scores on GPQA Diamond (graduate-level physics, chemistry, biology) and the 2024 USA Math Olympiad. See “Learning to Reason with LLMs,” OpenAI, September 2024. [openai.com](https://openai.com).

<sup>26</sup> Sakana AI, “The AI Scientist,” 2024. [sakana.ai](https://sakana.ai). Several papers generated by this system were accepted at workshops.

The window is narrow. AI tooling is crystallizing now. If provenance infrastructure emerges from closed systems—if the default becomes proprietary knowledge graphs that don’t interoperate—the opportunity closes.

We are not racing other open efforts. We are racing the absence of open infrastructure before something closed fills the gap. Google, Microsoft, and a dozen startups are building knowledge systems. They will be useful. They will not be open. They will not preserve provenance for anyone who doesn’t pay.

Once defaults harden, they persist for decades. Peer review emerged in the 1600s. The journal article took its modern form in the 1800s. What we build now—or fail to build—will shape how knowledge moves for generations.

## THE STAKES

The obvious objection: haven’t LLMs already solved this?

They’re trained on every paper ever published. They can summarize literature in seconds, generate hypotheses, draft manuscripts. Why build infrastructure when intelligence can navigate the mess directly?

Three reasons this is wrong.

**First, documents are not knowledge.** The archive contains the seventeen-year gap baked in—papers that were retracted, findings that failed to replicate, claims that were quietly abandoned. LLMs read documents. They cannot tell you which assertions have been verified. They cannot trace provenance. They generate confident, plausible, unverifiable outputs.

**Second, superintelligence won’t solve it either.** A genius in a datacenter can reason, but it cannot run experiments. Scientific knowledge is not pure inference—it is inference about physical evidence. Provenance cannot be inferred; it must be recorded. You cannot reconstruct what was never written down.

**Third, the substrate we build now is the substrate AI inherits.** In 2021, almost no code was AI-assisted. By 2025, nearly half is. The same curve is coming for science. If AI scientists inherit only documents—PDFs without provenance, papers without replication data, claims without trails—they will build on that substrate. Companies will optimize for it. Workflows will assume it. The window is before the tooling crystallizes—not after.

And the stakes are not abstract. Dario Amodei describes a “compressed 21st century”: AI-enabled biology compressing fifty to a hundred years of progress into five to ten.<sup>27</sup> If that compression happens inside systems we cannot audit, humanity does not gain knowledge. It gains outputs. A black-box oracle that says “trust me” is faith, not science.

If we do not build the constellation, the future is not “someone else builds it open.” The future is one of three failures:

Capture—a proprietary system achieves adoption because it’s better than nothing, and science becomes dependent on infrastructure it doesn’t control.

Fragmentation—every institution builds its own silo, none interoperate, and the transmission problem worsens.

Drift at scale—AI generates outputs faster than anyone can verify, errors compound, and confidence in science erodes.

The constellation is not just efficiency infrastructure—it is governance infrastructure. It is the audit trail that lets humans verify what AI-accelerated science claims to have found. Every assertion traceable to evidence. Every chain of reasoning queryable. Every correction propagated.

Without this, the compressed 21st century could mean a century of cures discovered in a decade. Or it could mean a century of outputs—confident, plausible, unverifiable—that we cannot distinguish from knowledge until patients are harmed.

The choice is not whether to use AI in science. That question is settled. The choice is whether AI-generated science will be auditable or opaque. Verifiable or trusted on faith. The constellation makes that choice possible. Without it, the choice is made for us.

## THE PRECEDENTS

It is tempting to assume that good ideas eventually happen. That if the constellation is worth building, someone will build it. That progress is automatic.

This is a myth. A comforting one, but false.

The Library of Alexandria held the accumulated knowledge of the ancient world. It declined—not in a single fire, but gradually, through neglect and conflict and the assumption that someone else was preserving it. In parts of Europe, knowledge that

<sup>27</sup> Dario Amodei, “Machines of Loving Grace,” October 2024. [darioamodei.com](http://darioamodei.com).

had once been routine—concrete, surgical techniques, engineering—became inaccessible for generations. Diffusion slowed. Rediscoveries took centuries. The knowledge existed. The infrastructure to transmit it did not.

Ignaz Semmelweis proved in 1847 that handwashing reduced maternal mortality from eighteen percent to two percent.<sup>28</sup> He was ignored, mocked, committed to an asylum. He died there. Handwashing didn't become standard for decades. The knowledge existed. The system for it to propagate did not. A hundred thousand women died. We named the phenomenon after him: the Semmelweis reflex, the tendency to reject new evidence that contradicts established norms. We remember his name. We did not fix the structure.

This is the actual history of human knowledge. Not a steady accumulation. A series of transmissions that succeeded and failures that are invisible precisely because they failed—the knowledge lost, the connections never made, the discoveries that existed and did not propagate.

## THE INVERSION

Software built foundation first, platforms second, agents third. Each layer depended on the one below. We have inverted this for science: agents first, foundation never.



<sup>28</sup> Semmelweis's data from the Vienna General Hospital showed that chlorine handwashing reduced mortality from 12–18% to under 2%. He was committed to an asylum in 1865 and died there within two weeks. Wikipedia.

The foundation layer (structured claims, versioned knowledge, trails connecting decisions to evidence) does not exist. The application layer (a platform where researchers can discover, contribute, query, and build on each other's work) barely exists. But we are racing to build the agents layer anyway, because that's where the capital flows, that's where the headlines are.

This is understandable. Agents are visible. Foundations are not. No one raises \$500 million to draw lines between stars. But without the lines, the stars are just scattered light.

## THE GAP

Dario Amodei envisions AI as "a virtual biologist who performs all the tasks biologists do"—compressing a century of progress into a decade.<sup>29</sup> The vision is plausible. But consider what such a system would need: not document search, but claim navigation. It would need to see which assertions replicated, which were retracted, which depend on work that quietly failed.

When an AI coding assistant suggests a change, it operates within structure—tests, dependencies, history. When an AI scientist generates a hypothesis, it operates on a pile of documents. We are building brilliant scientists and handing them a library with no catalog.

The industry will spend trillions on AI infrastructure by 2030. Recent government and industry initiatives are mobilizing compute, models, automated labs, and national datasets into unified platforms—the ambition is Apollonian: compress a century of progress into a decade. The investment is real.

But consider what these initiatives are building: engines.

We are building the engines. We still have not built the roads.

The investment in knowledge infrastructure that would make those systems navigable is vanishingly small by comparison.

GitHub was essential for AI coding capability. Without it, Copilot could not exist—there would be no structured repositories for it to read, no version history for it to learn from, no pull requests for it to understand. The infrastructure came first. The capability followed.

---

<sup>29</sup> Dario Amodei, "Machines of Loving Grace," October 2024. [darioamodei.com](http://darioamodei.com).

For science, we are attempting the reverse. We are building the capability and hoping the infrastructure will somehow emerge. It will not.

## WHY NOW



AGI is a crowded field—\$109 billion in 2024, every lab racing the same benchmarks.<sup>30</sup> That canvas is saturated.

But the problems this essay describes—scientific infrastructure, the substrate for discovery, the architecture that lets knowledge travel—this canvas is almost empty.

The question is no longer whether it should exist. The question is why it doesn't yet, and what has changed to make it possible now.

The AI argument for urgency is real—we must build structure before agents entrench on documents. But there is a simpler urgency. While we debate infrastructure, children are being misdiagnosed. Researchers are duplicating work. Treatments are being delayed. The seventeen-year gap is not a future problem we might prevent. It is a present harm we are permitting. Every year without structure is a year of preventable loss.

## WHY THIS HASN'T BEEN BUILT

People have tried. They failed. Why?

Beyond the Semantic Web's sequence problem, the deeper issue is incentives.

Many benefit from the current structure. Publishers profit from access and prestige, not transmission. They call paywalls “protecting intellectual property.” They have no incentive to make knowledge arrive.

Journals reward novelty because novelty attracts attention. Replication is boring, even when it saves lives. They have no incentive to track what replicates.

Researchers optimize for publication counts because that is what tenure committees measure. They have no incentive to build infrastructure no one will cite.

## THE CASCADE OF LOSS

<sup>30</sup> Stanford HAI AI Index Report, 2025. [aiindex.stanford.edu](http://aiindex.stanford.edu).

Consider how knowledge actually moves in a lab. A graduate student runs an experiment. The result—positive, negative, ambiguous—goes into a notebook. It might be mentioned at lab meeting. If it is publishable, it enters the formal record. If not, it stays in the notebook. When the student graduates, the notebook goes on a shelf. When the PI retires, the shelf is cleared. The knowledge vanishes.

At each handoff, information falls through the cracks: negative results, failed approaches, calibration tricks, the context that would save the next person months. Software solved this with version control: every change preserved, every discussion attached, every dead end recorded. Science has no equivalent. There is no commit history for experiments, no diff between what a lab knew in 2020 and what it knows now, no pull requests for hypotheses.

For every ten experiments, perhaps one reaches publication. The rest—failed approaches, calibration tricks, protocol modifications, results that didn't replicate—stay in lab notebooks or vanish entirely. A graduate student taking three ambitious bets has a sixty percent chance of nothing publishable. The system doesn't reward exploration. It punishes it. The metric must change: not papers published, but light left behind—including the light that showed a wall.

No one's job is to make knowledge arrive.

And infrastructure is invisible. When a patient dies because knowledge didn't arrive, no one counts it. When a researcher wastes three years duplicating work, no one notices. The cost is diffuse, distributed, deniable. The benefit of building infrastructure accrues to everyone—which means it accrues to no one in particular.

So the lines don't get drawn.

Git succeeded because one person needed it badly enough to build it. GitHub succeeded because it made Git accessible—lowered the barrier until adoption cascaded.

The constellation requires the same: someone who needs it badly enough to build it, and a design that makes adoption cascade.

## THE ADOPTION OBJECTIONS

This sounds like another knowledge graph project. Those always fail.

The graveyards are real. Semantic Web ontologies that no one adopted. Knowledge bases that went stale. Graph databases that promised interoperability and delivered silos. The pattern: build the perfect structure, then wait for the world to use it. The world never does.

The constellation inverts this. It does not require agreement before use—it extracts structure from existing practice. LLMs parse papers that already exist. Trails record clinical decisions that are already happening. The structure is not prescribed; it is observed. The difference between “agree on a schema” and “extract what’s already there” is the difference between the Semantic Web and Google.

Who will do the validation work? It’s thankless.

Wikipedia proved that distributed validation scales when three conditions hold: the work is granular (small edits, not monographs), reputation accrues (edit counts, barnstars, admin status), and the stakes are visible (your contribution appears on a page millions will read). The constellation inherits this design. Confirming a claim is a single action. Contribution is credited and visible. The difference from Wikipedia: domain expertise is weighted. A hematologist’s confirmation of a hematology claim carries more weight than a layperson’s. The incentive is not altruism—it is reputation in a system that tracks contribution.

Won’t publishers fight this?

They will. But the constellation does not compete with journals for the same value. Publishers sell access and prestige. The constellation provides structure and navigation. This is why Git succeeded despite the existence of proprietary version control: it created new value rather than capturing existing value. The constellation does not make papers open—it makes claims navigable. Publishers have no obvious way to stop this, and some may find it useful.

How do you bootstrap adoption when a small constellation is useless?

The cold-start problem is real. A network with 1,000 claims provides little value; one with 10 million transforms how science works. The path is vertical before horizontal: one field (oncology), one problem (drug interactions), one community (clinical trialists) where the density of claims justifies the infrastructure. RECOVERY worked because the trial structure was useful to each hospital independently, not only when every hospital joined. The constellation must provide local value before it provides global value. Clinicians use it because it helps them today, not because it might help science tomorrow.

The flywheel works like this: a clinician logs one patient's response and finds a researcher who studied that mutation. The researcher finds a bioinformatician who curated the dataset. Each connection creates the next. The constellation grows not by recruitment but by *use*—each trail drawn makes the next one easier to find.

Git spread through Linux kernel developers, then open source projects, then enterprise. Wikipedia spread through early encyclopedists, then topic experts, then the general public. The pattern is always the same: useful to a small group first, then the small group pulls in the next circle.

...  
...

The history of scientific publishing followed a similar pattern: peer review, journals, PDFs—each step was reasonable. The system accreted locally sensible steps into globally absurd architecture.

The constellation inverts this. Structure emerges from use, not from committee. Claims are extracted from existing literature, not negotiated in working groups. Trails are recorded as clinicians work, not prescribed by governance bodies. Validation is distributed—thousands of domain experts confirming or correcting, not a central authority decreeing.

LLMs extract structure from existing text—not perfectly, but well enough. Extraction is becoming cheaper than validation in many domains—and that changes the design space. Embedding models find semantic similarity without vocabulary agreement. Distributed validation at Wikipedia scale is now tractable.

The bottleneck moved from “get everyone to agree” to “verify what was extracted.” The first is impossible. The second is merely hard.

A new generation of research institutions has emerged, each attacking a different piece of the problem. Arc Institute gives scientists long-horizon appointments and freedom from grant cycles. Arcadia Science publishes open-access with data availability from day one. New Science funds early-career researchers on ambitious bets. The Focused Research Organization model, pioneered by Convergent Research and supported by groups like Astera Institute, builds purpose-built teams for specific scientific infrastructure problems.

These are real contributions. But notice what most of this momentum targets: discovery acceleration. Freedom, funding, automation, compute. The transmission layer remains

thin. What none of these efforts are building is the road network—the claim and trail infrastructure that connects what was found to who needs it next.

But consider what happens when structure exists.

When COVID hit, every country ran trials. Most fragmented—each hospital with its own protocol, incompatible with the others. The UK did something different.

In late February 2020, Martin Landray, an Oxford epidemiologist, emailed Jeremy Farrar, then director of the Wellcome Trust. A few days later, they discussed it on a No. 18 bus to Marylebone. Farrar suggested Landray join forces with Peter Horby, an infectious disease specialist. Within nine days of writing the protocol, the first patient was enrolled.

The RECOVERY trial launched with shared infrastructure: web-based randomization that any hospital could use, a single ethics committee approval instead of 180 separate applications, minimal data collection integrated with NHS electronic records. The structure made participation effortless. One in six COVID patients admitted to UK hospitals entered the trial.<sup>31</sup>

Within 100 days, RECOVERY had enrolled 11,000 patients and delivered its first result: dexamethasone reduced deaths by one-third in ventilated patients. The drug costs £5. It has since saved an estimated one million lives worldwide.

# 1,000,000

“It’s very, very rare,” Landray reflected, “that you announce results at lunchtime, and it becomes policy and practice by tea time, and probably starts to save lives by the weekend.”

Lunchtime to saving lives by the weekend. This is what arrival looks like when structure exists: the knowledge does not sit in a journal for seventeen years. It moves at the speed the moment demands because the infrastructure is there to carry it.

Same pandemic. Same virus. Same doctors. Same patients. The difference was structure.

AlphaFold is the same lesson at scale: a foundation layer (the Protein Data Bank) enabled a capability leap. Without foundation layers, capability does not translate.

---

<sup>31</sup> Wellcome Trust, “The Story of RECOVERY,” 2021. [wellcome.org](https://wellcome.org). The trial ultimately enrolled over 47,000 patients across 180 hospitals.

## WHAT THE STRUCTURE ENABLES

If the constellation exists—if claims are versioned, trails are recorded, provenance is preserved—what becomes possible?

A researcher in São Paulo makes a discovery. Within hours, the claim is extracted, linked to existing evidence, and visible to every lab working on adjacent problems. A team elsewhere sees it, recognizes the connection, and acts—not after a conference season, but immediately.

A physician faces a patient with an unusual presentation. She queries the constellation not by keywords, but by asking: what claims exist about this combination of symptoms, in this population, with this variant? The system returns evidence trails: what replicated, what failed, what was retracted, and why. She makes a decision in minutes, not hours.

A graduate student begins a dissertation. Instead of eighteen months reconstructing what senior researchers know, she navigates the frontier: what is solid, what is contested, what is unknown, and where her next experiment would matter most. She starts at the edge of knowledge, not in the fog beneath it.

These are not science fiction. They are what infrastructure makes possible—the same transformation Git and GitHub made for code. The difference is that for science, the stakes are measured in lives.

This is the gap we are proposing to fill: not upstream acceleration, but downstream arrival.

Someone will build this. The question is whether it will be built in the open, for everyone—or captured, like so much infrastructure before it.

What happens if we don't build it open?

We continue as we are. Another generation wastes years reconstructing knowledge that exists but can't be found. Another wave of patients receives treatments that arrive a decade late.

Or worse: someone builds it closed. Proprietary. A constellation you can only see if you pay. The map of what humanity knows—owned by a company optimizing for quarterly returns.

Neither future is inevitable. But look at what persists: peer review emerged in the 1600s and has barely evolved. Most technologies from Newton's era have been replaced. This

one we preserve, like a family heirloom that no longer works but we cannot bear to throw away. The journal article took its modern form in the 1800s and still dominates. The incentive structures that reward publication over replication, novelty over reliability, access over arrival—these have calcified over decades. Defaults have a way of becoming permanent.

AI will accelerate discovery. Models will read papers faster than humans, generate hypotheses, run experiments. The rate of new findings will explode. But if no one builds transmission infrastructure, we will have more discoveries and the same arrival problem. The gap between what is known and what reaches patients will widen, not close.

We have one chance to build this right. The window will not stay open. Once the structure exists, it will be very hard to replace. The decisions we make now about openness, governance, contribution, and visibility will shape how knowledge moves for generations.

This is the blank canvas. What we draw on it will outlast us.

We have named the problem. Now for the structure itself.

## THE CONSTELLATION



Consider the difference between an artifact and a claim.

A paper is an artifact: “Chapman et al. 2011, *New England Journal of Medicine*.” A dataset is an artifact. A protocol is an artifact. You can cite artifacts. You cannot query them. You cannot ask: has this been replicated? Under what conditions does it hold? What depends on it?

A claim is a structured assertion: “Dexamethasone reduces mortality in severe COVID patients.” But the claim knows more than its statement. It knows its evidence: the RECOVERY trial. It knows its constraints: benefit in patients requiring oxygen, no benefit in those who do not. It knows its replications and dissent. It knows its lineage: what it superseded, what it updated, what depends on it.

When a new study challenges the claim, the structure updates. When a contradicting paper is retracted, the claim’s confidence updates accordingly. When a clinician queries the constellation at midnight, she sees not just the assertion but its entire evidence biography: what supports it, what challenges it, how certain we are, and why.

## EXTRACTION

Extraction is becoming cheaper than validation in many domains—and that changes the design space. Large language models can read papers and protocols and propose structured claims with enough accuracy that human validation, not human reading, becomes the limiting step.

An LLM reads a paper and proposes: “This paper asserts X, with evidence Y, under conditions Z.” A domain expert confirms, corrects, or rejects. The validated claim enters the constellation with its validation history attached. One expert-hour can validate dozens of proposed claims. The work that once required reading every paper becomes review of proposed structure.

## ONE WORKED EXAMPLE

Consider a claim extracted from the literature: “*BRAF V600E mutation predicts response to vemurafenib in metastatic melanoma.*”

What does the claim require?

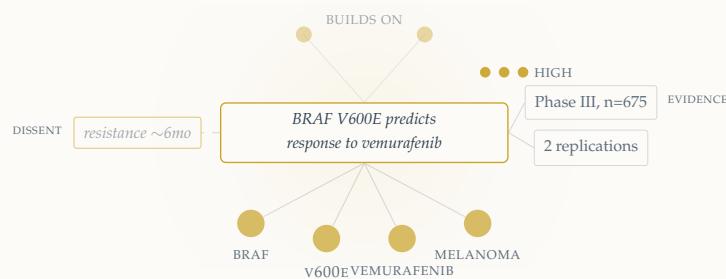
ENTITIES: BRAF is a gene. V600E is a variant. Vemurafenib is a drug. Metastatic melanoma is a condition. Without stable referents, you cannot say “this gene” and mean the same thing across papers, databases, and time.

EVIDENCE: Chapman et al. 2011, a Phase III trial with 675 patients. Two independent replications.

CONFIDENCE: High—multiple independent trials show consistent effect.

DISSENT: Resistance typically develops within six months.

LINEAGE: Which findings this builds on, which it supersedes.



A claim is a queryable object. When a new study appears—supporting, contradicting, or refining—the claim’s status updates. The constellation evolves.

The same structure applies outside medicine. A claim can be: “This catalyst increases CO<sub>2</sub> reduction efficiency under these conditions.” Or: “This forecast model fails systematically in this region and season.” The evidence might be datasets, protocols, lab logs, replications, and negative results. The point is invariant: assertions become objects with provenance, constraints, dissent, and history.

#### EVIDENCE VERSION CONTROL

The constellation is version control and auditability for scientific claims. Note-taking systems organize documents. Review aggregators summarize literature. Guideline authorities prescribe actions. Search engines retrieve PDFs. None versions the evidence state of assertions—what supports them, what challenges them, how confidence changes over time, and why.

Existing tools—UpToDate, Semantic Scholar, Elicit, Scite—each address a symptom. The constellation addresses the substrate.

...

And there is a second primitive: the trail.

A trail is a recorded path through knowledge—what you searched, what you found, what you decided, and why. A trail might record: *searched BRAF variants in metastatic melanoma; filtered to Phase III trials with n>200; rejected three papers—one retracted, two with undisclosed industry funding; selected vemurafenib protocol based on two independent replications and six-month resistance timeline.*

The constellation captures what journals reject: negative results. A failed experiment, properly recorded, is knowledge. It tells the next researcher: this path was tried; it led here; try elsewhere. Today, negative results vanish into lab notebooks. In the constellation, a trail that ends in failure is still a trail—and for someone facing the same wall, it may be the most valuable trail of all.

But before we can build trails, we must understand what happens without them.

Henrietta Lacks was thirty-one when she died of cervical cancer in a segregated ward at Johns Hopkins Hospital. Her doctors took a sample without her knowledge or consent. Those cells—HeLa cells—became immortal: polio vaccine, cancer treatments, gene mapping, COVID vaccines. Over 110,000 publications and billions in research.<sup>32</sup>

Her family didn't know. For decades, they couldn't afford health insurance while researchers around the world built careers on her contribution. The knowledge traveled everywhere. The trail back to its source was never built.

In 2019, I attended the Henrietta Lacks Memorial Lecture at Johns Hopkins. Her great-granddaughter, Veronica Robinson, answered questions about her family's experience.

"This comes down to a question," she said, "of how we bridge the gap between the community and the health-care field."<sup>33</sup>

Two years later, her family was still fighting for recognition and redress. Seventy years after her death, the gap was never bridged.

I didn't have an answer then. I am trying to build one now.

...

Here is what arrival looks like when the structure exists.

In the Netherlands, a baby girl was born with shortened bones and failing kidneys. She died within seven weeks. Her clinicians had no diagnosis. They sequenced her genome, found two candidate genes, and entered her case into Matchmaker Exchange—a federated network for rare disease research. Within days, they connected with researchers in Germany who had seen the same mutation. Then Canada. Then Portugal. Then the UK. Nine patients across five families on four continents, linked through a single API. They named the disorder. The research led to a potential treatment.<sup>34</sup>

---

<sup>32</sup> Rebecca Skloot, *The Immortal Life of Henrietta Lacks*, 2010. Publisher page. HeLa cells remain one of the most important tools in biomedical research; the Lacks family received no compensation until 2023.

<sup>33</sup> Veronica Robinson, remarks at the Henrietta Lacks Memorial Lecture, Johns Hopkins University, October 2019. Johns Hopkins News-Letter coverage.

<sup>34</sup> Oud et al., "Mutations in EXTL3 Cause Neuro-immuno-skeletal Dysplasia Syndrome," *American Journal of Human Genetics*, 2017. doi:10.1016/j.ajhg.2017.01.013. PMID:28132690.

The baby in the Netherlands didn't survive. But her data did. It traveled through the structure her clinicians built, reaching families who had spent years searching for the same answer.

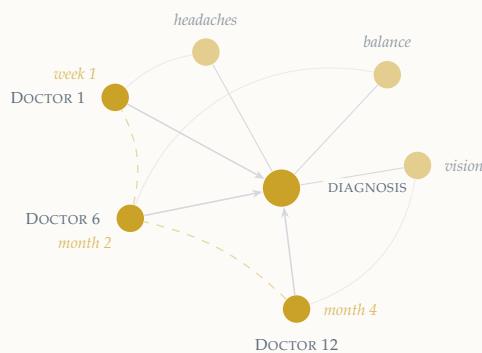
...

Matchmaker Exchange exists. It works. But it only connects genetic data to diagnosis. It does not yet connect treatment decisions. It does not yet record trails. And it does not yet credit the patients and families whose contributions made the discoveries possible.

Matchmaker Exchange bridges one gap—data to diagnosis. The constellation attempts to bridge another: contribution to credit, source to structure.

Veronica Robinson's question is still unanswered. How do we bridge the gap?

The constellation is one attempt.



This is what a constellation looks like: claims linked to evidence, trails linking decisions to paths, the whole structure versioned over time. Navigation for the clinician who needs to know what works. Discovery for the researcher who needs to see what's missing.

The cognitive scientist Dan Sperber observed that culture emerges from two basic acts: making private knowledge public, and internalizing what others have made public. The constellation operationalizes this for science. Claims are externalized knowledge. Trails are internalization paths. The structure is culture—made queryable.

The astronomer uses constellations for both purposes. Where is Mars tonight? The constellation tells her: near this star, in that region of sky. Where should I point the

telescope? Not at bright stars—at the dark spaces between, where structure predicts something should be but nothing has been found.

The same structure serves both. The map of what is known is simultaneously a map of what is missing—gaps made visible, queryable, navigable.

## THE FOUNDATION

Claude Shannon proved that information could be measured, transmitted, and corrected—that noise could be separated from signal, that errors could be detected and fixed. His mathematics enabled the digital age: every text message, every streamed video, every financial transaction depends on the theory he developed in 1948.

The same problem exists for knowledge. Scientific claims are signals; noise is the unreplicated findings, the retracted papers, the errors that propagate uncorrected. We have built systems to transmit knowledge—journals, databases, search engines. We have not built systems to correct it. When a paper is retracted, the citations don't update. When a finding fails to replicate, the claims that depend on it don't know.

The constellation is Shannon's insight applied to science: transmission with error correction built in. Claims are versioned. Trails are auditable. Corrections propagate through the graph.

The technical architecture will be public: a shared schema for claims and trails, explicit confidence inputs, first-class dissent, and propagation rules that make corrections travel.

## AGENTS AND THE FRONTIER

The constellation is not just for humans. It is infrastructure for AI scientists.

Today's models read documents. They summarize, synthesize, hypothesize—but they cannot verify. They cannot tell you which claims replicated, which were retracted, which depend on evidence that quietly failed. They navigate fog.

The constellation gives agents something different: a queryable frontier. An agent can ask: where is evidence thin? What claims are in tension? What gaps imply the next experiment? It can propose hypotheses as structured claims—not prose, but assertions with specified evidence requirements.

This is the difference between Copilot and an AI that reads Stack Overflow posts. Copilot works because it operates on structured repositories—version history, dependencies, tests. AI scientists will work when they operate on structured knowledge—claims, trails, provenance.

## PULL REQUESTS FOR SCIENCE

Consider how code evolves. A developer proposes a change—a pull request. Others review it. The change is accepted, modified, or rejected. The history is preserved. The repository updates.

The same pattern applies to knowledge. An AI extracts a claim from a paper. A domain expert reviews it—confirms, corrects, or rejects. The claim enters the constellation with its review history attached. Another researcher proposes a dissent. The structure holds both positions until evidence resolves them.

This is a pull request for science: proposed changes to what we collectively know, reviewed by those qualified to assess them, versioned so the evolution is visible. The difference from code: reviewer expertise is weighted. A hematologist’s review of a hematology claim carries more weight than a layperson’s. Reputation is earned through accurate validation, not institutional affiliation.

The constellation makes visible what has been invisible. Every claim knows its contributors. Every trail names its blazer. Credit follows contribution—not just to those who created, but to those who maintained, verified, connected, corrected.

The incentive is not altruism. It is earned reputation in a system that tracks contribution with the same rigor that Git tracks commits. Your validation history is your credential. Your trails are your CV.

## SCIENCE BECOMES OPEN SOURCE

When the constellation exists, science stops being an ivory tower. Participation looks less like credentialed gatekeeping and more like open source: pick an issue, validate a claim, record a replication trail, contribute a correction.

AI does the retrieval and extraction. Humans do the judgment. Trails teach newcomers how experts reason: what they searched, what they rejected, what they trusted, and why.

A high school student can make a meaningful contribution by verifying one claim and leaving one trail of how they did it. The limiting factor becomes curiosity and care, not access to a locked library of PDFs.

### THE HORIZON

We will begin where the need is sharpest and the structure is most absent: one narrow domain where the gap between discovery and arrival is measured in lives. The first constellations will be incomplete but queryable: claims extracted from a decade of literature, trails from a small network of specialists, dissent and corrections attached where they belong.

As the structure grows, the bottleneck shifts to validation. Contributors do not need to read every paper; they review proposed claims, attach constraints, and record trails. The frontier becomes navigable: dense where evidence is strong, jagged where it is not.

This is direction, not a schedule.

...  
...

The frontier of knowledge is jagged—peaks where funding flows, valleys where it doesn't, entire continents unexplored because no one thought to look. The constellation reveals this topology. Without the lines, the gaps stay invisible.

### CORRECTIONS

Cracks are part of any system that grows. In *kintsugi*, broken pottery is repaired with gold and the fracture remains visible. Knowledge needs the same property: errors detected, corrected, and propagated.

In January 2014, Haruko Obokata published two papers in *Nature* claiming a revolutionary method for creating stem cells—just dip cells in weak acid.<sup>35</sup> Labs worldwide scrambled to replicate.

---

<sup>35</sup> Obokata et al., "Stimulus-triggered fate conversion of somatic cells into pluripotency," *Nature*, January 2014. doi:10.1038/nature12968. Both papers were retracted in July 2014. doi:10.1038/nature13598.

Ken Lee, a stem cell researcher at the Chinese University of Hong Kong, decided to document his attempt publicly. He live-blogged his replication on ResearchGate—the first live-blogged scientific experiment.<sup>36</sup> Within weeks, he identified the problem: what Obokata had seen was likely autofluorescence, cells glowing from stress rather than transformation. Her revolutionary discovery was an artifact.

Lee submitted his findings to *Nature*'s "Brief Communications Arising"—the journal's mechanism for publishing corrections. *Nature* rejected it. No clear explanation. The journal that published the flawed papers refused to publish the correction.

For six more months, labs worldwide continued wasting resources. Rudolf Jaenisch at MIT: "Many people wasted their money and their time and their resources on repeating this."<sup>37</sup>

The correction existed, but the system could not propagate it.

A constellation makes corrections durable and discoverable. A retraction links to what it corrected. A failed replication updates the claims that depended on it. A dissent stays attached to its evidence until the structure of evidence shifts. Errors stop living as footnotes and start living as updates.

The window is open now. The child in the next waiting room deserves better than I had. We can build the system that makes that difference.

## THE COALITION



The coalition that will build this structure is not the same as the one racing toward AGI.

The pursuit of AGI required a specific kind of person: ML researchers, GPU infrastructure engineers, scaling-law believers. That community formed years before the technology existed. They gathered because they saw what was coming. They became the labs now racing toward AGI.

The constellation requires different people—and offers a different way to contribute.

<sup>36</sup> Lee's replication log was publicly documented on ResearchGate beginning February 2014. Coverage in David Cyranoski, "Acid-bath stem-cell study under investigation," *Nature News*, February 2014. doi:10.1038/nature.2014.14738.

<sup>37</sup> Quoted in David Cyranoski, "STAP cells are finally dead," *Nature News*, September 2014. [nature.com/news/stap-cells-are-finally-dead-1.15977](http://nature.com/news/stap-cells-are-finally-dead-1.15977).

If you've ever maintained a private knowledge system—a Notion database, a Zotero library, a folder of annotated PDFs—because the public systems were broken, you already understand. You've been doing this work alone. Building personal constellations in the absence of shared ones. What we're proposing is not new. It's what you've been doing, connected.

The work is already happening—scattered, unconnected, invisible. The constellation connects it: a network instead of isolated nodes.

The constellation doesn't have users. It has contributors: people who maintain the structure and keep knowledge navigable.

Build for the child you were. Build what should have existed.

**Clinicians** are trail-makers. When you search for evidence, document what you found and what you decided. Your path through the literature becomes someone else's starting point. A trail that took you an hour saves the next clinician that hour—multiplied by everyone who follows.

**Researchers** are claim-extractors and validators. You have domain expertise that AI lacks. When an LLM proposes that a paper asserts X, you can confirm, correct, or reject. Ten minutes of your time structures knowledge that will be queried for years.

**Engineers** are builders. The constellation needs extraction pipelines, query interfaces, and validation tools. If you want infrastructure that matters—infrastructure measured in lives, not clicks—this is it.

**Patients and advocates** are testers. You became experts by necessity. You know where the system fails because you've lived the failure. Ensure the structure serves the people it claims to serve.

Across the world, people are already doing this work—alone. A virologist who spends evenings logging mutations because the official databases lack context. A patient advocate who built a trial tracker because the registries failed her daughter. An engineer maintaining an extraction tool on nights and weekends, twelve GitHub stars, no funding. They don't know each other exists. That's what we're building: the infrastructure to connect them.

## HOW WE WORK

The people who build this share three practices—not because they’re required, but because they work:

*Record your trails*—what you searched, what you found, what you decided. *Version in public*—so others can see how knowledge evolved. *Credit contributors*—name those who maintained, not only those who created.

And three commitments:

*No confidence score without visible inputs.* *No retraction without propagation.* *No structure without open access for publicly funded knowledge.*

## HOW WE SUSTAIN

Open source projects die when maintainers burn out. Infrastructure disappears when funding ends. We design against these failure modes.

The protocol is simple enough that any competent team can implement it. The data format is open enough that migration is always possible. No single implementation is canonical. The constellation is not software—it is a protocol, and protocols survive their creators.

When Linux’s original author steps away, Linux continues. When Wikipedia’s founders left day-to-day operations, Wikipedia continued. The constellation must have the same property: built to outlast its builders.

## OPERATING PRINCIPLES

Three principles guide this infrastructure:

*Reproducible.* Every claim traceable to evidence. Every trail replayable.

*Observable.* Confidence computed from auditable inputs. Dissent visible. Retractions propagate.

*Collaborative.* Structure emerges from use. Validation is distributed. Credit is tracked.

## GOVERNANCE

Who participates matters. How the system is governed matters. These constraints keep the constellation open, auditable, and resistant to capture:

*Open infrastructure, not proprietary product.* The claim graph is public. The extraction code is open. No single entity controls what is visible.

*Transparent confidence.* Every confidence score is computed from auditable inputs—evidence count, replication status, recency, source quality. The algorithm is public. Disputes concern inputs, not hidden rules.

*First-class dissent.* A minority position does not disappear. It persists, with its evidence, until the structure of evidence shifts. Consensus is not manufactured; disagreement is made visible.

*Resistant to capture.* Contribution weight is computed from validation accuracy over time—not from funding, institutional affiliation, or publication count. Track record is earned, not bought.

The constellation does not decide what is true. It reveals what evidence supports. Two contradictory claims coexist until evidence resolves them. The system shows disagreement; it does not resolve it by fiat.

*Distributed by design.* The constellation has no center. Every node contains the protocol. The structure is the same whether one person holds it or a million. If any organization disappears—including the one that started it—the pattern remains. This is not redundancy for safety. It is the only architecture that survives.

*Self-documenting.* Every claim contains its provenance. Every trail shows where it came from and where it leads. A researcher in a century should be able to navigate without instruction—the structure itself teaches you how to read it. The Rosetta Stone preserved language because its structure was interpretable without external knowledge. The constellation must do the same for scientific claims.

These principles apply to all contributors—human and machine. When an AI proposes a claim, the same rules hold: traceable evidence, auditable reasoning, visible dissent. The constellation does not distinguish between human and machine contributions. It asks only: can this be verified?

*Maps, not mandates.* Trails are records, not recommendations. They document what was searched, found, and decided—not what should be done. Liability remains with the practitioner. The constellation is a map, not a mandate.

*Openness with boundaries.* Not all trails are public. The protocol supports private trails, embargoed claims, and tiered access. Openness is the default for publicly funded research. It is not a mandate for all knowledge.

These principles will be tested. We will be judged by what we build, not what we declared.

## THE STARTER RITUAL

If you want to contribute, start with one complete unit of structure. Pick one paper in your domain and do the following:

1. Extract one claim: write the assertion as a sentence.
2. Attach evidence: link the primary source and include a short excerpt.
3. Add constraints: where does the claim hold, and where does it fail?
4. Record a trail: what you searched, what you read, what you decided, and why.
5. Publish or submit: share it as a small, reviewable change.<sup>38</sup>

One claim and one trail is enough to begin. Structure compounds.

The ritual is not symbolic. It is operational. Your first claim becomes queryable. Your first trail becomes followable. You are not practicing—you are contributing. The constellation grows one claim at a time, and yours is one of them.

Thirty minutes of your time. One structured piece of knowledge that will outlast you.

---

<sup>38</sup> Contribution guide: [borrowedlight.org/contribute](http://borrowedlight.org/contribute).

## GLOSSARY

**Claim.** A structured assertion with evidence, constraints, confidence, dissent, and lineage.

**Trail.** A recorded path through knowledge: searches, sources, decisions, outcomes.

**Evidence.** A cited source (paper, dataset, protocol, guideline) that supports or challenges a claim.

**Constraint.** The conditions under which a claim holds (population, context, method, dosage, environment).

**Dissent.** A disagreement attached to a claim with its own evidence; it persists until evidence resolves it.

**Lineage.** The dependency graph of claims: what a claim depends on, what it supersedes, what cites it.

**Confidence.** A transparent score computed from auditable inputs (evidence count, replication status, recency, source quality).

**Evidence version control.** Open infrastructure that versions claims and trails so evidence can be navigated, corrected, and applied.

**Epistemic biography.** The history of a claim over time: evidence, challenges, revisions, and how confidence evolved.

**Protocol.** The shared primitives and schemas that make claims, trails, confidence, and correction propagation interoperable.

**Constellation.** The navigable structure formed by connected claims and trails: a map of what is known and what remains uncertain.

# THE WORK AHEAD



We will begin where the need is sharpest and the structure is most absent. The first constellations will be small: hundreds of claims extracted from a decade of literature, dozens of trails from a network of specialists, a single domain where the gap between discovery and arrival is measured in lives.

Success will not be measured in users or revenue. It will be measured in patient-days saved: the time reclaimed when knowledge arrives faster than it otherwise would. The first trail that surfaces a buried paper. The first claim that prevents a duplicate experiment. The first dissent that finds a place to stand.

This is the beginning.

The structure is not yet a running system. The first claims are being mapped by hand, not extracted at scale. The trails are imagined, not yet recorded. We are at the stage where the work is mostly conviction and a small group willing to build.

That may be enough. Git began as one person's frustration. Wikipedia began as an experiment most experts expected to fail. The structure that will matter in twenty years rarely looks inevitable at the start.

The six-year-old who needed this is building it now.

The coordination problem is real. The incentives are misaligned. The window will not stay open forever. We build while it is open.

## WHEN WE'VE SUCCEEDED

The goal is irreversibility. When enough researchers use the protocol that it would be harder to abandon than to maintain. When the structure exists in enough implementations that no single organization controls it. When the community can govern itself.

At that point, we've succeeded—not because of what we built, but because of what can now be built without us.

Bitcoin reached escape velocity when it had enough participants that no single actor could kill it. Wikipedia reached it when it had enough editors that it could survive any departure.

The constellation reaches it when the protocol is implemented, the data is distributed, and the community is self-sustaining.

We are not building a product. We are building a substrate. The measure of success is not adoption—it is inevitability.

The alternative is to guarantee failure by not trying—to assume progress is automatic and let today's defaults harden into permanence.

...  
...

#### A PICTURE OF ARRIVAL

In 2035, a fusion startup achieves sustained ignition—the breakthrough that unlocks clean energy for the planet—built on thirty years of scattered research that finally became navigable. A climate model predicts a hurricane's path with unprecedented accuracy, drawing on chains of evidence that were finally connected. A new antibiotic reaches patients in three years instead of seventeen, because resistance data was structured from the start. A student contributes a replication trail that saves the next lab a year. A child with headaches and balance problems is diagnosed in time.

These are not separate victories. They are what arrival looks like.

...  
...

#### GIGAFactories for science

If the constellation is built—if claims are versioned, trails are recorded, provenance is preserved—what becomes possible?

A gigafactory is a production system—standardized processes, continuous operation, compounding efficiency. Tesla did not invent the electric car. It built the system that made electric cars manufacturable at scale.

Science today is artisanal. Each project starts over. Each lab reinvents. Each finding sits in a journal, disconnected from what came before and what comes after. We have the intelligence. We have the tools. We lack the production system.

The constellation is the foundation layer that makes science manufacturable. Consider what becomes possible:

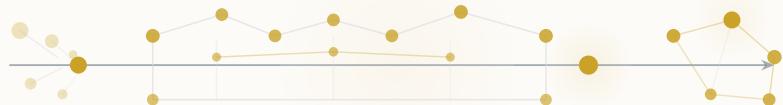
*Continuous operation.* Science stops running as grants and papers—episodic, disconnected. It runs as always-on pipelines: hypotheses generated, tested, validated, updated. The frontier advances continuously.

*AI as primary labor.* Agents don't summarize papers—they navigate the frontier map. They see where evidence is dense, where it's thin, where the value of the next experiment is highest. Thousands of them, in parallel.

*Recursive improvement.* Every validated claim improves extraction. Every correction that propagates improves downstream decisions. Trails create training data for better search, better review, better experiment selection. Without versioning and propagation, agents amplify noise. With them, the system compounds.

*Observable production.* Every claim has lineage. Every result is replayable. Science becomes something you can monitor, fork, diff. Not read—navigated.

*Compounding returns.* Each discovery makes the next faster. Tools improve through feedback. The efficiency curve steepens.



This is what happens when you industrialize a production process. Manufacturing learned it. Software learned it. Science has not—yet.

The constellation is the foundation. The gigafactory is what the foundation makes possible.

This essay is not a description of the constellation. It is a first trail through it.

The claims you read here—about diagnostic delays, about the seventeen-year gap, about RECOVERY and dexamethasone—are claims in Vigil. The citations are links. The structure is live. We build the future by inhabiting it.

Engelbart demonstrated NLS by using NLS. His demo wasn’t *about* the system—it *was* the system. The constellation must be the same: not described, but demonstrated. What you are reading is the demonstration.

The stars are there—scattered across journals, databases, the minds of specialists, the outputs of AI systems. The lines that connect them are ours to draw.

What we build now determines what the next generation inherits: versioned claims, auditable trails, knowledge that arrives. Isolated stars, or constellations.

The light we received was borrowed from those who came before.<sup>39</sup> The light we pass forward is a choice. The lines are ours to draw.

---

<sup>39</sup>The constellation on the last page includes Sigma Draconis, 82 Eridani, and Delta Pavonis—stars whose light left approximately eighteen years ago and is arriving now. Sigma Draconis. 82 Eridani. Delta Pavonis.



The stars are there.

*The lines are ours.*

We cannot know which trails will matter most. We can only draw them and leave them visible for  
whoever comes next.

*For the ones who come next.*