THE MEANING MODEL: CONSTRUCTING WORLDS AND STORIES AT PROGRESSIVE RESOLUTION

A PREPRINT

Henrik Westerberg
Emergent Wisdom
henrik.westerberg@emergentwisdom.org

September 5, 2026
Revised: November 15, 2026

ABSTRACT

A language model holds a great deal of implicit understanding of the world, learned from human writing. The Meaning Model is a method, with a working tool, for making that understanding explicit. The language model writes an account of a world as linked processes that change over time, in one general-purpose grammar of six record forms. Where nobody has measured a state, it guesses, with an honest band and a tag saying what the guess rests on, so the account can be tested and refined. A state with no natural unit, such as trust, needs no invented scale: a Cut divides one declared unit among the answers to a question and keeps an explicit remainder. Each process's trajectory is a commitment that writing and acting must consult, the proposed mechanism for consistency.

Worlds are built at progressive resolution: coarse first, then opened where a purpose needs detail, each opening keeping the coarser account's commitments or revising them openly. The world, its meanings, the modeler's reasons, and the prose or code written from it share one medium, so a change shows what depends on it, and discovered concepts anchor further inference. The paper gives the grammar, the ideas behind it, and a working tool with modes for modeling a world, writing fiction from it, keeping an agent's memory, and searching for new mechanisms through invented worlds. Three constructions show it in use: a novel, *The Book of Conditions*, built twice, the second time by an agent from a premise and historical sources, and a history of humankind from about 300,000 years ago to 2025, built in agent passes, its numbers checked against one another. They show that such accounts can be built and checked for internal consistency. They do not yet show that explicit structure improves on what a model does implicitly, or that the medium can train better models, which the companion paper Life Simulation develops.

Keywords Meaning Model, Recursive World Grammar, Concept Grounding, Event-Processes, Progressive Detail, Finite Resolution

1 Introduction

The central idea of this paper is that a language model can construct an explicit account of a world from what it already understands. Many world models admit a quantity only after measurement, so an unmeasured history, such as one character's trust in another over a year, stays qualitative. A language model, however, is already attached to reality through its implicit conceptual understanding, learned from human writing about the world. Because that understanding was learned from language, questions can draw much of it out in words, and the accompanying tool is built to ask them (Section 6). The model can therefore make its own understanding explicit and recursively constructible, including particular histories nobody measured and concepts nobody has named, before evidence tests the account. Making it explicit means constructing an account from what the model has learned, not reading out its internal representations. The written account need not be faithful to be useful: it is attributed, inspectable, and testable, so its errors can be found and corrected instead of staying latent.

The method therefore asks the model to guess. A guess here means using implicit knowledge and making it explicit: where nobody knows a state, the model records its best estimate, with an honest band and a tag saying what the estimate rests on, rather than leaving the state out until it can be verified. During the tool's development, language models tended to share a human bias, that a state nobody has measured cannot be guessed. Yet a language model has absorbed far more written experience than any one person, and its explicit estimates may become more accurate than a person's as models grow more capable; that is a hypothesis to test against later evidence, not an assumption. A structure built on a bad guess can mislead what is built on it, so the guess stays addressable: later refinement or evidence can revise it, the revision check names what was declared to depend on it, and the revision history keeps the original guess.

That bias has a history. Earlier language models hallucinated: they stated what they did not know in the same voice as what they knew, and people learned not to let a model guess. That failure has been receding, in the author's judgment faster than the habit it taught. The method is written for the models that are arriving, not for the ones that taught the habit. Where the world is well measured, such a model's recall is already close, and its bands have held the figures that sources later gave in all but a few cases (Section 7.2 has a first count). What the earlier habit makes hard to picture is what a model of this kind can do once it has a grammar to follow: hold hundreds of processes at once, each with its values over time, and keep everything it writes consistent with them.

Behind the design is a question: what must be rendered for a simulated world to be believable to whoever experiences it? A reader, a viewer, or a player probes a world only locally, through one route, one scene, or one visit. The world therefore need not be detailed everywhere. It needs a coarse but consequential account of its macro-processes, the long developments, institutions, economies, and environments within which lives and events happen, and detail opened wherever attention or the purpose reaches, each opening consistent with what the coarse account commits. Believability then comes from causal coherence across scales rather than from uniform detail: whatever the experiencer probes follows from the larger structure. Level-of-detail simulation asks a related question about computation and perception [1, 2]; here the question extends to meanings, inner lives, and text. It is progressive resolution seen from the experiencer's side, with the declared resolution set by what the experiencer can reach and what the purpose needs.

In principle, nothing in the method needs a computer. A novelist could have written every character's life from childhood, tracked each voice and belief as a process with comparative numbers, and linked every sentence to the Events it depicts. The balance of costs has likely kept authors from doing so. One continuous mind can hold much of a novel's world implicitly, while writing that understanding down, keeping it current, and linking it to the text would take more work than most books can bear. Tacit knowledge also resists telling: "we can know more than we can tell" [3, p. 4], which made knowledge acquisition the bottleneck of knowledge engineering [4]. Where one mind was not enough, as in series bibles, writers' rooms, and franchise continuity, the explicit record has usually been kept only as far as the work required. Writing down everything behind a book, in a language a machine could check, would have been an engineering project, as if each author needed a team of mathematicians beside them. A language model changes the balance: it can do that team's work and model what people never wrote down. Extensive articulation can become far easier for it, while it lacks a continuous mind to hold a long work together, so persistent records supply continuity across sessions and agents that its context cannot. What is new is what this makes possible: progressive resolution and a shared construction medium connect discovered concepts, quantitative commitments, recursive refinement, and coordinated revision in one inspectable account.

The intended agentic use is to give an intelligent system a small, recursively reusable toolbox for building an explicit world model. A common objection holds that language models lack world models, and that intelligent systems must instead learn predictive world models directly [5]. The Meaning Model takes a different route: rather than requiring a network to learn a world model internally, the language model constructs an explicit one from what it has learned, which can then be inspected, tested, and revised. It is a different method of constructing a world model: one built by making concepts explicit and formalizing them over time, not learned by predicting the physical world, with prediction serving as the test of the account rather than the way it is built. That is why it is called a meaning model rather than a world model. What it builds is still a world model, a semantic one: its state and concepts are versioned over time, bound to the perspectives that hold them, and quantified enough to test and to simulate.

The running example is the first construction of *The Book of Conditions*, called the first Book below. It is a completed alternate history in which Babbage's bounded calculating apparatus is built and put to work by a calculation office. It departs from our history at one decision: in August 1843 Babbage continues negotiating Ada Lovelace's offer to conduct the business of the Analytical Engine rather than refusing it, as he did. The negotiation becomes a priced and bounded undertaking for one customer, Guardian Assurance, and in May 1851 a small card-directed apparatus passes an independent trial. Its commercial organizer, Edward Halden, accepts orders beyond the office's independent checking capacity, while Lovelace seeks a successor to sustain that checking. Halden is invented, and so is Thomas Neale, the one man on whom the independent checking rests. Lovelace dies in November 1852 with her request unanswered. In 1854 an unchecked, erroneous table is delivered: the story concerns the human and institutional conditions of a

computational claim, not just whether the machine works. The Book's Events are numbered, as in E20, with a letter for an opened child Event, as in E20a, and the notes that record its reasons likewise, as in U14; Appendix A.5 sets many of them out at four resolutions (Table 15). The worked Events, Cuts, and numerical figures below refer to the first Book, which was built under the author's direction. The second Book (Section 7.1.6) was generated by an agent. It starts separately from the premise and historical sources and does not inherit the first Book's manuscript or its world model.

Explicit structure does more than record the model's understanding: it is the proposed mechanism for consistent thinking and acting. Processes carry trajectories over time. States nobody measured receive comparative shares under a declared question and unit, with an explicit remainder, the share that no named answer takes, instead of an invented absolute scale. In the world history (Section 7.2), the waking hours of a garment worker in Dhaka during the flood of 1998 divide, by the feeling that colours them most, into fear .30, bodily suffering .25, hours no feeling marks .16, grief .14, and smaller shares for the rest. In the first Book's authored reading, Babbage's outlook at the trial of May 1851 divides into assurance .76, threat .18, and an unresolved .06; that unit is the first Book's own, and the number test of Section 2.1 came later and asks more of it (Section 7.1.4). The grammar keeps, with every claim, whose account it is and what it concerns. These trajectories are explicit commitments that generation must consult and then preserve or revise: when the model writes a scene, makes a character decide, or changes a program, it is to read the recorded state of each relevant process at that moment instead of improvising from its context. The same numbers make a proposed refinement checkable by arithmetic, as double-entry bookkeeping catches an inconsistency without establishing that the books are true, and because the records persist, another agent or a later session can inspect, correct, and build on them.

Section 7 reports three constructions built this way: the two Books, and a history of humankind from about 300,000 years ago to 2025, which at its current state holds 861 processes and 21 laws that set its numbers against one another. None of them validates the method. What they show is the kind of check an explicit account allows. It can be checked against itself: the laws written into the world history found its errors twice, and the first time, seven of ten apparent findings that were checked proved to be errors in the model. It can be checked against later evidence: where a figure read from a source replaced an inferred one that had a band, the sourced figure lay inside the band in 151 cases of 157 (Section 7.2). Those are modern, well-measured series that a language model is likely to have met in training, so the figure shows that the inferred values for them were close and their bands wide enough, and says nothing of earlier eras. And a search that modelled and ran invented worlds produced ideas that a blind jury of models rated as more likely to work than ideas from the same worlds when they were only written, 4.3 against 3.1 on a seven-point scale (Section 6.6); the two searches also differed in tool version and in one instruction, so this shows a difference and not yet its cause.

A language model that carries out the construction the tool instructs is meant to build coherent worlds and books this way. Whether works built this way are more coherent than unstructured writing is what the prospective evaluation tests. The numerical-dependence test in Section 7.3 is designed to measure whether the numbers actually steer what the model writes.

Reality does not arrive already divided into every "thing" to which people give a word. We identify a person's life, an expedition within that life, a marriage, an avalanche, a species, a recession, and a story. Each term makes some portion or pattern of experience addressable while leaving other variation in the background. The same evening of reading can belong to a person's biography, a book's history, human culture, biological evolution, and the physical history of Earth.

A revolution or the growth of a city can be represented with the same grammar as the evening of reading above: participants, phases, and processes that change one another at different timescales. A century of city growth can open into migrations, building projects, and individual decisions; each can be decomposed again where the inquiry needs it. Trust and legitimacy need no standard physical unit. Relations can remain unweighted; where numerical comparison is useful, shares answer a declared question about a divided whole. Money, population counts, and duration keep their own units, whether they are measured, inferred, or authored during construction.

A boundary alone does not preserve reference. The same person may participate in a life, an expedition, a rescue, and a later conversation; those events need a stable way to state that they concern the same continuing individual. Conversely, a person may be opened into organs, cells, or devices for one question without discarding person-level identity. The difficulty is therefore not producing one description, but preserving identity and meaning when a description is opened, revised, compared with another world, or rendered into language. Ordinary prose can silently change what a term refers to. A flat graph can store relations without stating how coarse and fine descriptions agree. An unqualified numerical rating can preserve neither who attributed it nor the comparison or measurement protocol that gives it meaning.

The Meaning Model proposes one grammar of six record forms for those responsibilities. Reusable Concepts describe continuing Things and Events. A process, the word this paper uses most, is an Event that lasts: a longer Event that other Events make up or change. Where the grammar has to be exact it is written Event-process. A Cut divides one stated unit among answers to a question: for example, how a character's attention is shared among current concerns, with a remainder for unresolved allocation. Bindings state which Thing fills which role in an Event, and a Realization records

an attributed claim that a fragment of the world instantiates a Concept. A Concept may remain lightly specified or have a process-based model written in the same grammar. Finer articulation reuses the same identities and respects existing commitments. The *modeler* is the person or AI system constructing, interpreting, or revising this account; *constructor* names its role when proposing changes, and the *generator* is the component producing candidates. These name roles, not additional record forms.

Construction is also a method of discovery. A language model can draw on learned world knowledge to infer and make explicit the processes, motives, relationships, and quantities behind events, including what people know implicitly but have not recorded or measured. Each explicit record becomes an anchor for further inference: the constructor can connect records, infer higher-level concepts from their relationships, and open those concepts into finer processes. Beginning from an institution that is losing legitimacy, for example, it may infer inconsistent decisions, unequal access, and failed expectations that nobody has measured. Those records can support a concept of procedural betrayal, which may then guide the study of another institution or explain a character's changing loyalties. The resulting structure anchors another round of abstraction and refinement; neither the domain, conceptual vocabulary, nor depth is fixed in advance. The system can also discover where its current account is inadequate, propose finer structure or a better decomposition, and test the change against observations or generated consequences. Useful changes should survive without losing the identities, meanings, and history already accepted; conflicts require explicit revision. The toolbox supports the growth of the conceptual space, not just relationships among concepts supplied in advance. Discovering and revising person categories through joint human and AI construction is part of this work, not merely an option to replace a finished inventory.

Progressive resolution and a shared construction medium are the proposal's two linked architectural contributions. They are also the machinery the central idea needs: progressive resolution lets an anchor be opened without losing what is already committed, and the shared medium lets a discovered concept be revised together with its instances, reasons, and prose. Together they form one construction method, in which world commitments, concept interpretations, authoring reasons, and prose are refined and revised together. Four representational changes carry it:

- **Construct detail under commitments.** A useful world can precede its finer history. An opening preserves declared coarse answers or makes the necessary revision explicit; it does not merely reveal precomputed detail.
- **Revise the account across surfaces.** World state, abstract concept models, Event descriptions, document text, and externalized understanding participate in one authoring operation. Exact dependencies identify which interpretations, reasons, and passages a changed commitment puts in question.
- **Construct meanings as well as instances.** A concept can be articulated through the same roles and processes as a situation that realizes it. Definition, application, and revision become inspectable operations, without requiring deep grounding for every concept.
- **Give semantic numbers their comparison context.** A coherent comparative question avoids requiring an absolute scale for each abstract concern. Its numbers describe a concept's contextual role rather than an intrinsic quantity, retaining the question, alternatives, remainder, and perspective. Measurements keep external units; concurrent processes do not become shares merely because they describe one person or world.

Refinement should preserve *semantic continuity*. Starting from the coarse claim that Babbage builds an Engine, a construction may later expose the machine's parts, its workforce, or competing accounts of the work without changing which Babbage, Engine, or building process earlier records addressed. Unopened regions can remain coarse; openings must fit the accepted boundary or revise it visibly. Suppose a coarse account calls an apparatus trial successful under a criterion requiring every case to pass on the first uninterrupted run, and its explanation and passage depend on that assessment. Opening the trial into phases introduces a stopped run followed by a successful restart. Those phases cannot be accepted as mere detail: they contradict the protected criterion, so the constructor must reject the opening or propose an explicit revision. A narrower claim, complete agreement after restart, would require a new grounding version (Section 3.2), a revised assessment, a reason for the change, and rechecking the dependent passage, while the earlier account remains addressable. This hypothetical cycle shows the contract; it is not taken from the Book's records. Section 4.2 states the record-level acceptance rules.

The same contract connects a character's modeled concerns to decisions and scenes. Comparative numerical accounts of wants and appraisals should guide a decision before its scene is written. Opening the decision can expose a conflict with those concerns or with what the character knows; the modeler can then revise the action, or revise the concept used to interpret it, while retaining the reasons and rechecking the scene. The immediate value is a working model that helps a modeler investigate what happens and why, and decide what to construct next. In storytelling, it guides decisions before prose rather than merely cataloguing a finished story. A whole life or institution can be sketched first, then opened where its decisions and consequences matter; incompatible detail prompts revision of the larger account.

1.1 Uses and long-term aims

The long-term theoretical aim is a representation that could model reality computationally through interrelated processes, including people's inner lives, at a declared finite resolution. The grammar is a representational substrate rather than an annotation shell around a finished simulator, and a model is complete only relative to a declared question and resolution. The larger aim is representational compression: to describe everything meaningfully distinguishable at declared finite resolution with as few domain-neutral operations as possible, so that a new domain adds content rather than a universal primitive. This ambition concerns progressive articulation of any part of reality or possibility that a purpose makes relevant, not computation of all reality at maximum detail. The minimality is in the grammar, not in the data: within it, the constructor tracks as much as it can and, over time, works toward the totality.

The representation is intended as a general-purpose modeling medium. Narrative worlds, simulation and knowledge systems, corpus annotation, and perspective-separated accounts of people are applications of the same grammar; storytelling supplies the first worked application here, not a boundary on the method. These uses fall into two families that share the construction loop (Section 6.1): generating something new, such as stories, game worlds, and ideas, and describing something that exists, such as a market, an institution, a person's own life, or how a person's understanding of a subject grows, concept by concept. Each use needs stable reference, attributed claims, and compatible resolutions; the grammar alone does not guarantee their practical benefit. For interactive worlds, the same method could leave much of the world coarse but consequential. A regional harvest failure can constrain a town's supply before its warehouses and households are constructed. A player's visit opens one warehouse, then a worker's experience of the shortage, under the protected supply commitments and domain rules. This is construction of previously absent detail, not selective display of a finished simulation. Distant events can still require coarse updates, later visitors must encounter a history compatible with recorded observations and intervening changes, and shared facts cannot silently change between players. Whether this saves computation at comparable fidelity remains untested by the Book.

These uses admit two complementary deployment modes. A model can think with the account, and a later model can learn from how the account was built. An intelligent system can read and update the records as an explicit semantic workspace. By returning to earlier numerical accounts and the questions or explanations linked to them, it could develop continuity of understanding across interactions: what changed, what might explain it, and whether the change belongs to the world or to its own estimate. The same distinctions and construction histories can also become prediction and reconstruction targets: a candidate training language for learning to choose useful distinctions and revise an account, not only assign values within a fixed vocabulary. Operating such models could become partly internal to the learner, while explicit records remain available for persistence, correction, and inspection; the learner need not reproduce their syntax internally.

The scope includes attributed inner life. Beliefs, feelings, wants, and private interpretations may be represented without turning a character's report, an observer's estimate, and accepted public state into one claim, and without making private experience directly observable or selecting one universal psychology (Sections 2.2.1 and 4). The long-term ambition is a model with increasingly accurate intuitions for a person's inner life: what someone wants, fears, believes, and feels, and how those processes shape decisions over time. Considering how to imitate an adult mind, Turing pointed to "the process which has brought it to the state that it is in": the mind's initial state, the education it received, and its other experience [6]. The representation models a person the same way, from where their life processes begin, through what they were taught, to what else they lived through, and construction asks for that history where it explains who the person became. Better measurement data and more capable models should make those estimates closer to the person being modeled, with improvement assessed against later evidence. Each person's account of the world can be represented beside the accounts others hold of that person and a shared, evidence-supported world model. The aim is to understand how people decide from the worlds they perceive, including what they misunderstand or leave unsaid, rather than treating a shared description as everyone's view.

What would change if works come to be built this way is the subject of Section 8. In short: a work is no longer its text alone. A book becomes one route through a modeled world, and its readers could open the world behind it.

1.2 Scope of the claim

The central idea, stated at the start of this Introduction, is a claim about method: that a language model can construct an explicit account of a world from what it understands and keep its thinking and writing consistent with it. Progressive resolution and the shared construction medium are the architecture that lets that construction be revised without losing what it has committed; the four representational changes above are how they do it. The idea adds one empirical expectation, that explicit construction improves on what a language model does implicitly, and Section 7.3 states how it is tested. It also rests on two formal claims that must not be confused.

- **Recursive refinement** is a property of the world grammar: every finer articulation uses the same records and the local contracts of its declared construction profile. A deeply grounded concept may reuse that grammar in a process-based concept model, but need not do so.
- **Finite-resolution sufficiency** is a constructive hypothesis: for a declared purpose and tolerance, the distinctions required by the task can be represented through those records without claiming that the representation is unique, complete at every resolution, or mechanically true.

The paper also plays, deliberately, with a stronger idea: that a person's inner state, their wants, feelings, beliefs, attention, and outlook, and more broadly the processes of a world, can be described at a declared resolution as a set of numbers and states that change over time, and that the set can be opened into ever finer detail. Nobody has recorded hundreds of such variables across a whole life, because nobody needed to; a language model can now construct them. The paper adopts this as a working hypothesis to explore, not as a result. Whether such a description predicts what a person does, or what a reader recognizes as true to them, is what the evaluation and Life Simulation are designed to test.

The hypothesis is also the working method of construction: every process the model opens carries its state across the span the work covers, in one of five kinds of value. A state may be a value with a real unit, one of a set of defined states, a reading on a defined scale whose anchors say what its values mean, a comparative share under a declared question, or the output of a program that computes the process from others. Assigning one of these forms to a person's inner state does not establish an intrinsic psychological intensity.

A world is generated this way down two trees at once: the tree of processes, opened from the highest abstraction by conceptual decomposition (Section 3.1), and the tree of time, opened from the longest span into periods, each level guessed within the one above. A part found by conceptual decomposition must also pass a number test before it is opened: it must carry values over time whose meaning a reader can state. A category that cannot yet be given such values remains a concept rather than an opened process.

Nothing in the grammar keeps a concept from being followed as a process once it has a grounding, an instance, or a use; Appendix E tests this on seventeen concepts chosen to resist it. A container, such as a person's body, holds processes rather than a value of its own: its state is the states of the processes inside it, and the test applies to them.

The construction discipline belongs to this paper: frozen-parent whole candidates, direction rolls, atomic commit or rejection, explicit revision, rendering routes, and read-back (Section 4.2). How process histories are generated, advanced, inferred, and learned belongs to Life Simulation [7]; Section 11 describes how the two proposals work together.

Section 2 states the grammar and its progressive construction, Section 3 concepts and their grounding, and Section 4 the linked surfaces, transactions, rendering, and read-back. Section 5 turns recursive refinement into a finite-resolution sufficiency test. Section 6 describes the construction method and the tool that carries it out, for general modeling, memory, story construction, and mechanism search. Section 7 reports three constructions built with it, two of *The Book of Conditions* and one of world history, and the comparative evaluation that none of them yet satisfies. Section 8 considers what changes if works are constructed this way. Sections 9 and 10 relate the proposal to prior work and state its limits, Section 11 places it in the wider research program, and Section 12 concludes.

If you want	Read
The idea in fifteen minutes	The abstract; Section 1 through its four representational changes, and the boxed minimized core at the end of Section 1.2; the opening of Section 7; and Section 12
Why it is built this way	Section 1, especially its passages on the central idea, on guessing, and on the novelist; Section 1.1; and Section 8
The grammar	Sections 2–5: the six forms, Cuts, and conservation; concepts and grounding versions; surfaces, authority, and read-back; closure and sufficiency
How to construct with it	Section 2.2 for the order of opening, Section 3.1 for carving, Section 5.2 for the value question, and Section 6 for the tool's workflows
What has been built, and what it shows	Sections 7.1 and 7.2, with Table 7 for the implementation boundary and Table 10 for the world history's states
How it would be tested	Sections 7.3 and 10
How it relates to other work	Section 9 and Table 11
Where it leads	Section 11
Terms, profiles, and provenance	Appendix G for a glossary, Appendix A for the first Book's profile, Appendix B for interchange, and Appendix D for provenance

Table 1: A reader's map.

The minimized core of the grammar. Six forms carry the world model: Concept, Thing, Event, Binding, Cut, and Realization. Bindings give Things roles in Events, while typed links connect records without turning every connection into the same relation. A Cut answers one question by dividing one stated unit among exclusive alternatives and a remainder; those shares total one. A Realization links a Concept to an addressable fragment without a weight, and any graded fit is expressed by a separate Cut on an assessment Event. The six forms keep identity, occurrence, participation, decomposition, and grounding apart without claiming to exhaust metaphysics. A construction profile may add lifecycle, admissibility, and recomposition rules while retaining the same forms. Understanding Nodes and Document Nodes share their address space with world records but keep distinct authority and validation, and perspective follows declared authority ancestry rather than a duplicated attribution field.

2 Grammar and Progressive Construction

In the first Book, the apparatus trial of May 1851 shows the six forms together. Babbage, Lovelace, the apparatus, and the sealed test packet are Things with continuing identities. The trial is an Event, opened into a first run that stops and is begun again, a second complete run that agrees with the sealed independent results, and a few idle cycles of shared pleasure; Bindings place it at its site, and a spatial child Event binds the people and objects present. An Event beneath Babbage's inner process carries a Cut dividing his outlook toward his active wants into assurance (.76), threat (.18), and a remainder (.06), an authored reading rather than a probability, in the first Book's unit (Section 7.1.4). *Valid computational claim* is a Concept, and a Realization records the claim that the trial instantiates it. The trial's protocol required the stop and voided the incomplete run, so no commitment had to be revised; the Introduction's hypothetical cycle shows what the contract requires when one must be. The rest of this section states what each form is responsible for and how a world built from them is opened.

Each form answers a question that the others cannot answer without erasing a distinction or relocating its responsibility. A Concept says which reusable meaning is intended. A Thing says which continuing instance is addressed. An Event says what persists or occurs. A Binding says which Thing fills which time-scoped role. A Cut says how one declared unit divides under one question. A Realization says which Concept grounds an addressable world fragment. The fragment is a subgraph of existing records rather than a seventh world form.

The removal test concerns these responsibilities, not their serialization. Removing Thing identity requires another way to join a person's appearances; removing Binding requires another way to distinguish that person's roles; removing Realization requires another way to distinguish a concept definition from its application. The same test applies to occurrence and allocation. Another schema may combine these forms while retaining equivalent typed information. Thus six is a compact, inspectable interface; complexity in field types and validators counts against the claim just as additional record forms would.

```
Concept { id, names, typed relations, constraints,
    grounding records, provenance }
Thing { id, kind, identity boundary, continuity criterion,
    typed links, provenance }
Event { id, kind, interval, typed world data, optional description,
    direction, typed links, provenance }
Binding { id, event, role, thing, interval, provenance }
Cut { id, parent Event, question, profile, divided unit,
    keyed weighted answers, keyed remainder, optional conditioning address,
    constraints, provenance }
Realization { id, purpose, context parent, target, concept,
    grounding version, optional alignment, evidence, provenance }
```

Answer and remainder addresses pair the Cut revision with a local key; they are not additional record forms.

The core is minimized at the level of record forms, not vocabulary labels. It uses six forms—Concept, Thing, Event, Binding, Cut, and Realization—and introduces a new form only when no existing form plus typed fields can preserve a distinction, in targets, composition, authority, inference, or validation, that a validator must enforce. Helpful names may be added; shared serialization never implies shared semantics. This is relative minimality under the declared contract, not a proof of globally minimal syntax or metaphysical primitives.

The definitions and rules needed to use this account are stated in this paper; the separately typeset *Minimized Grammar* collects them as a compact reference without adding requirements or changing their status. Both specify the target contract, not the current Rust host schema, which enforces a stated subset of it, listed in the repository's conformance notes.

A reference distinguishes a logical identity from an exact record revision. In the selected immutable profile (Section 2.1), an edit creates a successor revision, not a new Thing or Concept identity. Provenance records the author or generator, commit, construction time, and evidence references. Times name their clock: Event chronology, construction history, or document order. Derived views retain their source revisions and derivation lineage rather than acquiring independent authority.

Relations are carried in four ways. A Binding assigns a Thing a typed, time-scoped role in an Event. A Cut relates a parent to weighted child answers and a remainder under one question and unit. Typed relation fields connect Concepts, Events, records, or typed world-data fields under declared kinds. A Realization grounds a world fragment in a Concept for a declared purpose. Each carrier retains the target, authority, clock, inference, and validation contract of its relation kind.

A relation among several Things is consequently a joint Event with one Binding per role, not a direct binary edge between the Things. Typed links such as `contains`, `about`, `updates`, and `supersedes` are sibling kinds, not synonyms. `contains` asserts scoped Event–Event containment; `about` is reference without participation or state change; `updates` names an affected process or typed world-data field changed by an Event; and `supersedes` preserves a predecessor while changing a typed lineage. A supersession must name its scope and clock, because replacing a character's earlier belief in story time is not the same authority operation as superseding an author note or a model commit in construction time.

A Realization is an attributed, unweighted grounding link. A *define* Realization associates a Concept with a grounding model or family; a *describe* Realization applies a recorded grounding version to a fragment. For example, the first links the Concept *helping* to its definition, and the second links a particular interaction to that definition. Each link retains a context parent (an Event or surface node), provenance, available evidence, and any role alignment. The relation is many-to-many across Concepts, fragments, perspectives, and times.

A graded judgment uses a separate assessment Event linked `about` its target beneath the process whose perspective it expresses. Its local Cut divides the comparison unit among *matches*, *does not match*, and remainder. Separate Concepts use separate fit Cuts; no independent Realization degree is introduced. Every Cut is parented by an Event, even when the target of assessment is a Concept or another record.

Cut answers are addressable components: a reference names the Cut revision and a stable local answer key, even at zero weight; the remainder has its own reserved key. An answer's target and its component address are distinct, since the same Event may answer several questions. Nested Cuts name the particular component whose conditional unit they divide. Conditioning addresses must resolve and cannot form cycles. This supplies exact targets for refinement, grounding, and notes without adding another world record form.

Queries and access follow a record-level contract, stated in full in Appendix C: context comes only through declared authority-parent edges, every path must resolve to one governing root, and no lookup exposes content outside the caller's authorized view or an isolated weight out of its required context.

2.1 Optional construction profile

The six forms play the role of a language; a construction profile selects how to use it for a task. Treating the Book's categories as Meaning Model itself would be like treating one Python script as Python. The intended practice is to develop better vocabularies and decompositions, by humans and AI together, judged by the characters and worlds they produce.

The six forms do not by themselves require one lifecycle convention, temporal partition, or domain vocabulary. This paper uses a stronger profile to show how they can construct a progressively resolvable world. Other profiles may choose different children and depths or omit human modeling entirely. What remains general is the Cut contract: when a decomposition uses a Cut, it names its parent quantity and question, admits mutually exclusive answers under that question, retains a remainder, and states how the answers reconstruct the parent.

The reading rule is simpler than that contract. A semantic number is a Cut weight: it means one answer's share relative to its siblings under the parent Event's question and unit, and the answers plus remainder sum to one. When a Cut divides a process among its categories, the number test asks that its unit be a pool that can be counted or pictured, such as waking hours, choices, acts, money, words, people, or land, with every answer drawing on that same pool; "one unit of her motivation" names none.

The links of an Event-process (containment, update, and reference) do not sum and carry no weight. Literal dates, durations, counts, currency amounts, positions, and other quantities with independently interpretable units or declared measurement protocols are typed Event data with units or protocol references and provenance. So are readings on a

scale the modeler declares with its anchors and states from a defined set: each means what its unit, anchors, or states say, and none is a share that must sum with others. They add neither a weighted Cut relation nor an unweighted process parent relation. A Realization is likewise unweighted; optional graded fit is represented by its own local Cut. If no divided unit and mutually exclusive answers can be named, no semantic number is assigned.

Three Cut families—allocation, sequential, and direction—yield five named Cut profiles: facet, contribution, cause, sequential, and direction. Their validator-visible distinctions are specified in Section 2.2.2.

2.2 Progressive world construction

The optional profile used here adds lifecycle anchoring and conservative refinement to the core.

Progressive resolution needs an order of opening as much as an opening rule. The aim is the totality over time: as many processes as can be tracked, made explicit, connected, and refined. At any moment, though, most of an account is still coarse, so a record is made explicit first where it must persist across later construction, constrain a later Event, differ across perspectives, or remain controllable and auditable, and where the declared purpose reaches. Detail not yet opened remains implicit at the present resolution; it is postponed, not excluded. A distant city, for example, may first exist only as its coarse processes, its size, its trade, and what travellers say of it; when the story or a player enters it, its opened streets and lives must stay consistent with that coarse account or revise it explicitly, while the travellers' reports remain attributed claims that may prove wrong. Where a governing Cut exists, its unallocated share stays in that Cut's remainder; an unweighted process or relation acquires no fictional remainder. This order does not guarantee that the chosen frontier is optimal. It decides which processes and which finer periods to open next. Once a process is open, its state is recorded across the span the work covers, guessed where it is not known, rather than left implicit.

2.2.1 Roots, identity, and relations

The profile distinguishes two roots. The Universe is a Thing that addresses the continuing identities in scope. History is its lifecycle Event and contains the accepted Events within the declared horizon. Time is carried by Event intervals and typed temporal data, not represented as another Thing. Merely assigning a date to a candidate does not place it in accepted History.

Each Thing names an identity boundary, a continuity criterion, and exactly one lifecycle Event in this profile. The Engine may therefore remain the same Thing while its location, ownership, completion, parts, and social interpretation change. An active participant Binding must lie inside the participant's lifecycle. Later Events may still refer to a destroyed machine or a dead person through `about`; they do not bind that Thing as an acting participant. Creation and destruction delimit lifecycles. Split, merge, replacement, and succession Events link predecessor and successor identities, while the declared continuity criterion decides whether either is the same Thing.

Resolved intervals are half-open, $[t_0, t_1)$. Every Event has positive duration at the declared resolution: a moment is a short interval, and an instant is a time at which a value may be dated, not an Event. Creation and destruction use declared boundary roles; ordinary participation cannot outlast a lifecycle. Unknown endpoints remain unknown: unresolved interval checks are *untestable*, not passed. Duration shares require a finite observation interval. Imagined and counterfactual timelines declare their own frames; temporal containment and comparison require a common frame. A changed lifecycle claim is an immutable revision, not a second current lifecycle for the same Thing.

Time-scoped relations specialize the joint-Event pattern: membership binds a member and group, physical part–whole binds a part and whole, and a partnership binds partners. Each relation can have its own history.

Rosters and encapsulation are views over active relation Events rather than independently authored truth. At time t, the view selects physical parts, members, or occupants whose relation Event and role-Binding intervals both contain t, retaining those different roles. Opening the Engine exposes parts without replacing the Engine: joint part–whole Events expose the parts and Thing identities preserve reference. A share of credit or activity is a separate Cut under a declared question; it is never inferred from membership. Likewise, a dispute may compose several letters, meetings, and departures into one episode without becoming identical to the partnership in which it occurs.

Space uses the same forms. A place is a Thing and presence is an Event binding a subject to it. When geometry matters, position and extent are typed Event data with a coordinate frame, unit, tolerance, and provenance. A sphere around a building records modeled size, not semantic resolution or uncertainty about its location. An explicitly location-propagating physical-part relation may derive a part's coarse presence from its container until detachment; social membership does not imply co-location. The world history uses these forms to place countries and regions on the globe (Section 7.2).

A coarse spatial history can connect dated home bases, workplaces, visits, and presence to the same continuing subject, then open selected scenes into finer geometry. A home or work base is a time-scoped relation Event, not an assertion of uninterrupted bodily presence. Unrecorded periods remain unknown, and a sequence of locations does not supply a travel route or speed between them. A representative town point and an authored room layout retain their different precision and provenance. Qualitative spatial relations remain useful where numerical coordinates are unnecessary.

Uncertain location is a different question. For a missing letter, an epistemic assessment Event beneath a named perspective process can parent a Cut asking which of two non-overlapping rooms contains it at noon. The Cut divides one unit of attributed location support among those alternatives and an unresolved remainder, retaining its evidence cutoff and anchors. This records uncertainty about presence, not simultaneous physical occupancy. Overlapping place relations can instead remain unweighted.

Individual Cuts can be drawn as trees, but the world is a graph. One Event may participate in several Cuts, Processes may overlap, and distinct inner or understanding roots may represent incompatible readings of the same Thing. Shared identifiers keep these views about one world without giving any node a global weight.

Four frequently confused openings therefore have different tests. Table 2 makes those tests explicit. None can stand in for another merely because all four can be drawn as parent–child diagrams.

Relation	What the children mean	Required connection
Concept specialization	More restricted kinds of the parent kind	Declared inherited constraints; no normalization
Expressive concept decomposition	Components used to articulate one grounding	Named comparator and fine-to-coarse mapping; a Cut only for a divided unit
Physical encapsulation	Parts of the same material Thing	Time-scoped part–whole Bindings, identity, and any declared extent constraints
Temporal segmentation	Disjoint portions of one Event interval	Coverage, interval order, and duration-derived sequential shares

Table 2: Different meanings of opening, with different preservation contracts.

2.2.2 Cuts, Event-processes, and world data

Cuts let states nobody has measured, such as one person's trust in another during a given month, read as the share of decisions left to the other, or a character's outlook at a trial, enter a computable account without an invented absolute scale. The central numerical distinction is not between important and unimportant records. It is between a divided quantity and an undivided relation. For Cut K,

$$w_i \geq 0, \qquad w_\perp \geq 0, \qquad \sum_i w_i + w_\perp = 1. \tag{1}$$

Each w_i is meaningful only beside its siblings under the Cut's parent Event, question, and unit. The answers must be mutually exclusive as shares of that unit even when the records they reference coexist. The same Event can receive different weights under different questions; no Cut assigns it a global importance. How much an Event matters is a separate judgment held for an audience, described below. If increasing one proposed answer need not reduce another and no competing unit can be named, the relation is not a Cut.

Table 3 states the complete boundary. The first two rows are the two ways a parent may be opened. The third is data carried by an Event, not another parent relation.

Form	Question answered	Representation	Weighted?
Cut	How does one declared unit divide?	Exclusive answers plus explicit remainder	Yes; sums to one
Event-process	Which Events occur within, constitute, update, or refer to this longer Event?	Typed containment, update, and reference links	No
World datum	What condition, unit-bearing value, or protocol result is represented?	Typed Event payload with unit or protocol, frame, uncertainty, and provenance	Not a parent relation

Table 3: Two parent relations and the separate world-data payload.

The five Cut profiles fall into three families. An *allocation* Cut divides a unit among answers in a frozen vocabulary, one fixed before a run, or among referenced records. Facet answers use the frozen vocabulary; contribution answers name Things bound to the parent Event; cause answers reference existing wants, beliefs, actions, circumstances, or Events. A cause Cut remains an allocation of explanatory responsibility located beneath one perspective-bearing process; normalization does not establish intervention-level causality. Overlapping contributions require an explicit allocation convention or a joint answer; assigning each the whole shared contribution would double-count it.

A *sequential* Cut orders disjoint child Events and derives their shares from duration, leaving uncovered time in the remainder.

An Event's direction is where it may go next: its possible continuations and what shapes them. A *direction* Cut divides one unit among mutually exclusive continuations at a stated horizon. Its children are possible successors rather than current parts, and its interpretation must be declared before selection or outcome. Proposal weights specify how a constructor samples; predictive weights assert probabilities over a defined outcome partition. Predictive calibration is then evaluated, not assumed from normalization, and an uncalibrated forecast can still be scored. At most one answer is realized; its accepted Event uses a `realizes-forecast` link to that answer's slot, including the remainder. Unrealized alternatives do not thereby become accepted Events.

An Event's general direction may instead reference a constraint, transition law, or learned predictor. Such domain models are content attached to the grammar. How they are proposed, compared, and trained belongs to Life Simulation; their presence does not enlarge the six-form ontology.

A share is one of the five kinds of value in the working method (Section 1.2). The other four are not shares, and the paragraphs that follow take them in turn: a value with a real unit, a defined state, a reading on a declared scale or measure, of which importance and cause assessment are two cases, and a value computed by a program.

World data need not arrive as a completed set of measurements. While constructing a world or narrative, the modeler can introduce missing quantities with real units: infer a journey's duration from a distance and travel speed, estimate a machine's mass from its materials, or choose a room's dimensions when authoring a fictional setting. Each value retains its unit, assumptions, provenance, and applicable uncertainty, distinguishing measurement, inference, and authorship. The usual perspective and acceptance rules determine whose account contains the value; later evidence can prompt an explicit revision. These values do not have to sum to one. That rule belongs to Cuts. A world datum may also be a defined state, one of a declared set such as the phases of a process's growth and decline.

A process's successive data can be kept compactly as a value series: one record per process and holder whose points carry the time, the value or state, an optional band, and a tag saying what the value rests on. The tool has five tags: *source*, stated by a text or the evidence; *inferred*, implied by a source; *invented*, in no source and authored for a work, never used for a claim about reality; *exploring*, an idea being tried and not yet relied on; and *sketch*, drawn roughly to see the shape and to be redrawn later. Each point stands for a typed world datum of the holder's account, not a runtime observation. A running world, a model run forward in time by its laws (Section 6.6), starts from each process's declared initial value; when that value is copied from a point, the process's support says which. A tag states a value's strength, not its reference: a value whose source, evidence cutoff, or authority matters is named in the point's note or kept as a claim that carries them.

A defined scale need not have a physical unit. Protocol-valued data retain the instrument version, question, anchors, time, and provenance. For example, a report of seven on a stated fatigue scale records a response under that protocol, not an intrinsic amount of fatigue or a share of attention. A modeler may declare such a scale for a process, with its range and anchors, and read the process on it over time; the readings are results under that declaration. Inferred and authored scale values retain their respective status. Comparisons and aggregation require justification under the scale; a zero-to-one encoding alone does not make a Cut.

How much an Event matters is another such judgment, and it is held for an audience. A modeler declares an importance scale for the inhabitants of a world, or for those of them who follow one category such as football, with ordered levels from the most important down, each with an anchor that says what an Event must be to stand there. One go-to scale counts how few Events of the whole span stand as high: about ten, a hundred, and a thousand. A holder places each Event at one level, and a later judgment replaces the earlier one in the revision history. Kept compactly as one record per Event, scale, and holder, each judgment stands for an assessment Event about the judged Event with its reading on that scale. Importance differs from the relations it is easily taken for: containment says what an Event is part of, a Realization says what kind of thing it is, and a Cut's shares are local to one question. None of them says how much the Event matters to anyone. A reader, or a view the modeler chooses, can keep to a scale's top level, then the top two, and so on, so that the most important Events of a category stand out from the rest.

A cause Cut allocates explanatory responsibility. A cause assessment answers a different question: how much difference a cause made. Its holder states the outcome and its contrast, why this rather than what, and links each cause to the outcome as causing, enabling, preventing, or constraining it. Where the difference can be judged, the holder reads the outcome without the cause, everything else as it was, on one declared measure: the probability that the outcome rather than its contrast comes about, or an amount in the outcome's own unit. The gap between the actual reading and that reading is the difference the cause made. A cause that stood in the way is read the same way, so its difference is a positive amount held back. Differences do not sum, because causes act together; causes that would each have sufficed alone are read removed together. Where a change in a quantity that adds up divides among its causes, that division is a Cut over the change, with gains and losses divided separately rather than netted into a remainder. Like importance, an assessment is held by someone and is not a share of a total; it records a judgment, not an identified intervention effect.

A process can also be given by a program instead of recorded values, the fifth kind of value in the working method. In the first Book's office a job needs 40 checker-hours, and the independent checker has 80 protected hours; take them, for illustration, as a monthly flow of four jobs. Instead of recording the backlog of unchecked work month by month, a modeler can declare laws: the backlog grows by the hours ordered minus the hours checked, the queue is the backlog divided by the hours checked, and while three months or more of checking are queued, the oldest job goes out unchecked. Run forward from January 1852 in the current engine, the laws compute every value: from April 1852 two jobs a month go out unchecked, half of all the work, which is the twofold demand-to-capacity ratio stated as a history. An intervention is another run of the same program. With a second checker funded from August 1852, as Lovelace asked, nine jobs go out unchecked, the last of them in August, and none afterwards. The values rest on the program's version and its declared inputs, not on separate guesses (Appendix F).

Children define their parent in a different sense for each relation. Cut children divide one stated unit; a direction Cut describes alternatives, not parts of a future that has already occurred. Shorter Events may constitute or alter a longer process without summing, and one Event can update several processes at once. Component processes that run alongside one another across a process's span, such as the dimensions a conceptual decomposition finds, constitute that process in the same unweighted way. Dates, money, distance, and measured physical change retain their own external units. An actor or reasoner may later pose a Cut over explanatory responsibility, but the grammar never manufactures a common magnitude merely because several domains changed together.

A Thing's lifecycle Event can host several concurrent, overlapping Event-processes. A person's body, partnership, work, or place processes, for example, can change together without becoming fractions of one life-wide quantity. Ordinary links state containment, occurrence, update, or reference; they do not themselves explain the longer process. Explanation requires a causal Cut or a cause assessment beneath the reasoner's typed ancestry, or an anchored UnderstandingNode, a text record of a reason on the understanding surface (Section 4). When a construction calls one process slow and another fast, those words name response time at a declared temporal resolution, not importance or a numerical rank.

2.2.3 Conservation under refinement

An opening is conservative only with respect to a declared contract. Before generating detail, the constructor fixes a base snapshot, protected queries at specified times and resolution, admissible error, and a map from proposed fine records to coarse records or Cut components. A typed *Close* operation reconstructs the protected answers from addressed fine records through that map; reading cached coarse answers does not count. Identity, categorical data, and protected links normally require exact agreement; measurements use same-unit tolerances. Acceptance requires each protected answer to agree with the base within its declared tolerance, as well as satisfying the structural constraints. Otherwise the proposal is rejected or submitted as a revision. Close is a validation operation, not a seventh record form, and cannot be chosen after seeing which projection makes a candidate pass. This is a finite preservation witness, not a theorem that every possible query or meaning remains unchanged.

For a macro-to-micro example, first commit a parent Cut $p = (.20, .70, .10)$ for one bounded period. Let its question ask how one character's motivational attention during that period divides, counted in waking hours, among the Engine work, other commitments, and unresolved allocation, and let each half of the period hold the same number of waking hours. The finer periods do not yet exist. Proposing the first half-period with the compatible Cut $(.30, .60, .10)$ consumes $(.15, .30, .05)$ of the whole-period unit. The residual is therefore $(.05, .40, .05)$, leaving a feasible second half-period with conditional shares $(.10, .80, .10)$. Close verifies

$$\tfrac{1}{2}(.30, .60, .10) + \tfrac{1}{2}(.10, .80, .10) = (.20, .70, .10).$$

By contrast, a proposed first half-period $(.60, .30, .10)$ already consumes .30 of the first component, exceeding the parent's .20. Its residual is $(-.10, .55, .05)$: no nonnegative completion can repair it. The constructor must reject that detail or explicitly revise the parent and its dependents. The calculation constrains what may be constructed without uniquely supplying the missing history. This is what comparative numbers on meaning buy: a proposed refinement

of an unmeasured allocation such as a character's motivational attention can be rejected by arithmetic rather than by judgment alone.

Suppose a finite, positive-duration parent Event P has a sequential Cut into disjoint intervals I_i. Their shares are

$$\lambda_i = \frac{\mu(I_i)}{\mu(I_P)}, \qquad \lambda_\perp = 1 - \sum_i \lambda_i, \tag{2}$$

where $\lambda_\perp$ covers time not yet opened. An open lifecycle uses an explicitly bounded observation Event rather than an infinite or unresolved denominator. The question here defines a time-average of relative composition, giving equal duration equal weight. Shares of a total resource whose amount per unit time varies instead require resource-mass weights or another declared composition operator. Compatibility requires the same question, divided unit, vocabulary, remainder meaning, and governing perspective. If the children fully cover the parent and every interval has a compatible facet Cut f_i, its time-averaged Cut is

$$f_P = \sum_i \lambda_i f_i. \tag{3}$$

When $\lambda_\perp > 0$, a complete parent vector exists only if the uncovered interval has an explicit compatible facet vector $f_\perp$, in which case its term $\lambda_\perp f_\perp$ is added. Otherwise the derived view reports the known partial contribution, coverage, and temporal remainder without inventing the missing vector. If a complete parent p was already committed, partial detail must still leave a feasible completion. In exact arithmetic,

$$r = p - \sum_i \lambda_i f_i, \qquad r_j \geq 0 \text{ for every } j, \qquad \sum_j r_j = \lambda_\perp. \tag{4}$$

The construction fixes any numerical tolerance in advance. A child cannot consume more of a parent component than was committed, even while other time remains unopened. A feasible residual establishes only that completion is possible, not what actually happened during the gap. Full coverage reduces to the mixture equality. Failure requires rejection or an explicit parent revision with its affected dependents; normalization must not conceal it.

A modeler may retain a conceptual decomposition across timescales or use scale-specific decompositions. Retaining categories supports consistent comparison without requiring continuous activity. Finer accounts may summarize a larger interval or explain a slower pattern through their sequence and interactions. Any claimed composition must declare its mapping and rule, report coverage, and preserve protected coarse commitments or revise them explicitly.

Only compatible Cut questions use this mixture. Relationships, presence, categorical lifecycle status, material stocks, and other topologies retain their own aggregation rules. For an allocation, Close sums fine shares mapped to each coarse answer; opening a remainder must account for that same remainder's mass. A conditional Cut distributes its enclosing component, so its full-unit shares are the enclosing weight times its local weights. For physical encapsulation, Close preserves protected identity, role, presence, and extent queries; for temporal segmentation it preserves interval coverage and declared ordering. These are different operators. A finer gear description can preserve the machine's identity and location without claiming to be an average gear.

Life, period, episode, and phase illustrate the distinction without adding record forms. A lifecycle is an Event bound to one Thing. Periods are the disjoint children of one authoritative sequential Cut of a bounded lifecycle interval. An episode is a coherent Event arc linked `occurs-within` a lifecycle or period and may overlap other episodes or cross period boundaries within its lifecycle. Its phases, when opened, use a sequential Cut of that episode. A Book construction may use these forms without treating its particular periodization or person vocabulary as part of the grammar.

2.3 Anticipation, shocks, and adaptation across scales

Longer processes supply the circumstances in which shorter Events matter. A shock is an Event that materially changes a specified process at the chosen resolution, under a declared criterion of change. It need not be harmful or unexpected: an invention or successful demonstration can reorganize a life as substantially as a loss. The same Event can affect a person, a partnership, and an institution differently. These are linked responses, not copies of one shock magnitude assigned to every participant.

Anticipation is optional and perspective-specific. A prior belief or predictive direction Cut records what an actor expected at a stated cutoff and horizon; preparatory Events record what the actor did about it. Constructor proposal

weights are not evidence of that expectation. An anticipated loss can still be severe, while a surprise can have little lasting effect. Missing prior records mean that anticipation is unknown, not zero. A society or ecosystem need not have one mind: forecasts belong to named actors or models, and any collective account declares its aggregation. Preparation can avert, reduce, or amplify the subsequent disturbance.

Adaptation is represented by response Events and changes in the affected processes, with descriptions and anchored UnderstandingNodes explaining the proposed connection. It need not restore the old state or improve matters. Recovery, consolidation into a different way of acting, continuing instability or cascade, and terminal transition are possible shapes. Anticipated transition allows adjustment before the focal Event. Growth, learning, and aging can also accumulate through drift without a discrete shock. Beliefs, wants, skills, relationships, and resources may change at different rates; one person's recovery can coexist with an institution's closure.

The recurrence across scales is a pattern of questions, not an assertion of exact self-similarity. Over decades, economic or political change can alter an institution's possibilities; over years, people reorganize careers and relationships; within a day, an interruption changes a plan or conversation. Larger processes constrain local options, while accumulated local responses can alter the larger process. Spatial reach and response time are separate: a brief decision can have widespread, lasting effects. Update links and anchored explanations express these connections. Only genuine allocations use Cuts, and only compatible questions recompose numerically.

No new record or mandatory stage sequence is required. Jump diffusion, a model of a quantity that drifts and at times jumps, may suit some continuously varying states with abrupt changes, but diffusion is not a synonym for adaptation, and not every process returns to a baseline. The Meaning Model records the before-state, any anticipation, the Event, and its responses.

Appendix A.5 applies these distinctions to the first Book at four resolutions, from the credit contraction of 1847 to the stopped run of May 1851 (Table 15).

2.4 Operations, named conveniences, and cross-record rules

The construction interface exposes six operation families. They are a small authoring surface over the records, not six additional ontological primitives.

Register stages an addressable Concept, Thing, or Event; in the lifecycle profile a new Thing and its lifecycle enter the same candidate. *Bind* assigns a Thing a time-scoped role in an Event. *Open or compose* is one structural family with two retained modes: open refines a frozen authoritative parent, whereas compose proposes a new parent over existing Events. *Cut* adds question-relative answers under one declared profile. *Decide* rejects a candidate, commits it as new, or commits an immutable revision that records superseded targets and revalidates its dependency closure. *Realize* defines or describes Concept grounding while retaining purpose, version, context parent, evidence, and uncertainty.

Convenient construction names specialize those families without enlarging the record grammar.

Construct	Core or interface basis	Why the name remains
Facet, contribution, cause	Allocation Cut	Their answer eligibility, interpretation, and validation differ
Presence, membership, part–whole	Event plus typed Bindings	Their roles and interval constraints support different queries and validators
Update	Event plus `updates` link	It asserts an effect on a process or typed world datum rather than generic reference
Value series	Events of typed world data with `updates` links, serialized per process and holder	Populating and deepening record many dated states in one operation; each point keeps its band and tag and remains an attributed datum, not an observation
Supersession	Typed `supersedes` link	Its scope and clock determine which lineage changes authority
Roll	Direction Cut plus Decide	The original forecast, seed, fixed mass, drawn mass, and alternatives must survive selection
Roster, inherited presence, summaries	Read-only derivations	They follow from committed records and carry derivation lineage rather than new authority
Understanding, render, read back	Understanding and document workflow	They operate on linked text surfaces outside the six world-record forms

Table 4: Named conveniences built on the minimized record grammar.

Five cross-record rules govern the construction profile.

1. Nonnegative, mutually exclusive Cut-child weights and the explicit remainder sum to one under the declared question and unit.
2. Qualified decompositional refinement satisfies its profile-specific composition operator within tolerance, or the candidate is rejected or committed as an explicit revision.
3. Containment and sequential children respect their applicable parent temporal or spatial extent; direction candidates instead respect the forecast horizon and are not treated as present parts.
4. Every active Binding interval lies inside its Event and the lifecycle of each actively participating Thing; an `about` link does not imply participation.
5. A Cut inherits the governing authority context of its parent Event; a different perspective requires another attributed Event beneath the corresponding typed root and another Cut.

A concrete construction may add stronger gates, such as co-presence checks or a rationale requirement for selected changes, provided it declares them.

3 Concepts and Examples

This section says what a Concept is in the grammar: how it keeps its identity, how its meaning is grounded lightly or deeply, and how the claim that something is an instance of it is recorded. Its closing examples read one world through several decompositions and place wants among other kinds of direction.

For example, someone requests and receives help. The people are Things; the interaction is an Event whose Bindings state their roles. Helping is a Concept. A process-based concept model describes the request, consent, action, and consequences through the same forms, under its own declared model context. A *define* Realization links the Concept to that model; a *describe* Realization records the claim that this interaction instantiates it, with the grounding version and role alignment. If graded fit is useful, a Cut on a separate assessment Event carries the judgment. An UnderstandingNode explains why it fits or why the definition needs revision. All six world forms can thus participate; the explanation belongs to the understanding surface. The same modeling language can therefore represent both a situation and a definition used to interpret it.

3.1 Concept identity and grounding resources

A concept first needs a stable identifier. Names can change, collide, or be translated; the identifier lets a world continue to refer to the same proposed meaning. Typed relations then place that identifier in one or more conceptual views. A concept may specialize another concept in one hierarchy, oppose a second concept, and help express a third without those relations becoming the same generic parenthood.

This relational representation can be intentionally light. A frozen emotion vocabulary, for example, may need only names, short constraints, typed relations, and anchors. Its position in a tree does not exhaust its meaning, but a construction should not be forced to author a miniature world for every term before it can begin. An AI model can draw on conceptual knowledge acquired during training; many familiar meanings can remain implicit while it constructs a world. Explicit grounding is useful where a meaning must be shared, kept stable for a run, revised with a record, or compared precisely. It also provides a place to refine inherited definitions and invent concepts that better explain the world. Concepts explicitly referenced by records still need stable identifiers; this does not require a detailed concept model for each one.

A *concept model* is a structured account of a concept's proposed meaning through roles, relationships, and constraints in the grammar. The process-based form adds Event-processes and phases; not every concept model needs them.

Four compatible grounding resources are available, and none is a mandatory stage. The first organizes identity and relations; the others add evidence or structure.

1. **Relational schema.** The concept is constrained by its name, typed parents, differentia, neighboring relations, and the question under which it is used. This is the cheapest form and may be sufficient for use inside a frozen local vocabulary, although taxonomy alone does not specify how the concept is realized in a situation.
2. **Representative prototype.** One provenance-bearing fragment is selected as a useful center of comparison without being declared the full extension of the concept. Different contexts may require different prototypes under the same stable Concept identifier.

3. **Exemplar grounding.** The concept points to provenance-bearing positive, negative, typical, and boundary examples. Anchors make a graded judgment comparable without pretending that one example is the definition.

4. **Process-based concept models.** Idealized models state characteristic roles, Things, Event-processes, phases, and Cuts in the grammar, without requiring a complete simulation. A model selected as an explicit reference for comparison is a *canonical concept model*. Selection does not make it universal, permanent, or true.

Recursive conceptual decomposition makes a process's meaning more explicit. The modeler selects one category in its concept model and opens it into finer categories whose roles and relationships explain that selected parent. Together, these children articulate what the parent consists in at the chosen resolution. The modeler can then select one child and repeat, leaving its siblings at their existing resolution. Each opening deepens the explanation while preserving the applicable coarse commitments, or revising them explicitly (Figure 1).

How a carve is chosen matters as much as how it is preserved. The construction follows the conceptual decomposition of Fractal Intelligence [8], in eight steps.

1. A carve first states its claim frame: what the concept is, what the carve is for, and the range of cases it claims.

2. The concept is then located by climbing to broader kinds, each step naming what distinguishes the kind below from its siblings, and the carve is made on the way back down.

3. Candidate parts, more than will be kept, are each tested for necessity, independence, and universality across the claimed range.

4. A candidate becomes a part only if all three hold, so strength on one test never offsets failure on another.

5. A candidate found only in some cases is a kind under a part rather than a part, and an uncertain one is kept as a labeled alternative with its verdict.

6. The accepted set is tested for completeness, best against an independent list made for another purpose.

7. Real cases chosen before the carve are routed through it: a case with no home exposes a missing part, and a case with two homes exposes two parts that are not independent.

8. The carve then states how its parts together make the whole work. Each part is itself a concept, carved again where further depth pays.

A recorded carve uses records already defined here. The climb is a chain of concept specializations; the accepted parts form an expressive concept decomposition (Table 2); and the frame, the verdicts with their reasons, the routing sample, and the synthesis are UnderstandingNodes held by the modeler and linked to those records. A part becomes a process only when it receives a measure and dated values. Section 7.2 reports a carve made this way.

An operational layer may add a recognizer that proposes a *describe* Realization and any component or aggregate fit Cuts from a world fragment, or an instantiator that proposes a fragment from desired realization conditions. Both retain their version and provenance and use the ordinary realization or candidate-commit interface. This links the construction of an example, the explanation of its concept membership, and the revision of that concept through one representational medium. Learning those operators is a further task, not a consequence of making their inputs explicit.

A concept-model family may retain context-specific models, counterexamples, boundary cases, and unresolved variation. Numerical typicality uses a local Cut.

Define Realizations and typed links associate a concept with these resources. A describe Realization records the comparison actually used: exemplars consulted, or an alignment to a selected concept model. When graded fit is required, separate component and aggregate fit Cuts sit under the relevant assessment Events. The comparison is never implicit.

Examples do not uniquely determine new applications. A concept-model family adds relational and temporal commitments; a recognizer adds a learned decision rule. Either may remain incomplete or culturally local. Moving between them is a possible learning path.

If the understanding process that authors or assesses a definition allocates attention or responsibility across a concept model's components, the Cut that records the allocation has as its parent a definition or assessment Event about the model, beneath the process. This authoring or assessment process is distinct from the processes represented within

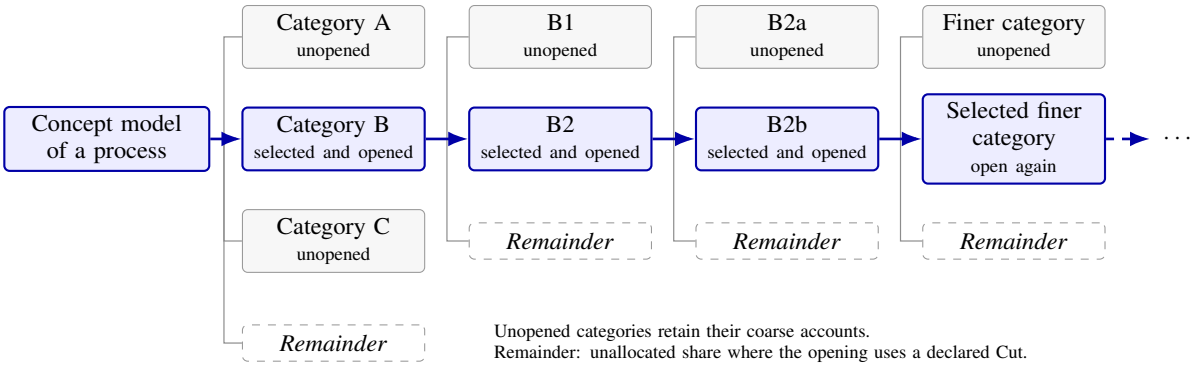


Figure 1: Recursive conceptual decomposition of a process. Finer categories and their relationships explain the selected parent; further openings deepen that explanation. Selecting B, then B2, then B2b exposes finer categories while their siblings remain at their existing resolution. The levels show conceptual depth, not successive periods; the final arrow indicates that opening can continue. Blue marks the selected branch. This is a selected view, not a requirement to construct a complete tree.

the concept model. Normalization applies only when its answers are mutually exclusive shares of one named unit. Overlapping senses or perspectives remain typed concept relations or separate Event-parented Cuts. The existence of concept models does not by itself supply an operator by which their world fragments recompose (Section 2.2.3).

3.2 Content-addressed grounding versions

A Concept identifier and a grounding version answer different identity questions. The identifier preserves the proposed concept across revisions. A grounding version identifies the definition bundle used for one Realization. A flat bundle may contain the identifier, description, relations, prototype or exemplar references, and constraints; a deep bundle may contain a process-based concept model and comparator specification. Either kind may be serialized under a declared deterministic serialization scheme and addressed by its full Sema content hash [9]. A revision produces another hash and an explicit `supersedes` link. The Concept identifier remains only when its declared continuity policy treats the revision as the same concept; a semantic split or merger normally introduces linked successor identifiers.

Content identity is narrower than semantic identity. Equal digests under one canonicalization scheme establish that two parties resolved the same canonical form; they do not establish that differently encoded models have different meanings, that the selected model is true, or that two perspectives accept its interpretation. Human-readable shortened handles are therefore references, not the authoritative identity. The full digest and canonicalization version remain in the record.

Reproducible fit Cuts require the hashed or transitively addressed bundle to name everything that can change the comparison. For a flat grounding this includes its descriptions, relations, anchors, and constraints. For a process-based concept model it additionally includes the fragment, role schema, component and aggregate Cut questions with their answers and remainder, referenced grounding versions, and comparator or recognizer version. A *describe* Realization then states exactly which bundle it used. Sema provides addressing and agreement on that bundle; the Meaning Model continues to carry context ancestry, evidence, alignment, and the claim made with it.

3.3 A model of concept history

The idealized fragment used to define a concept is not automatically an Event in the history of a modeled society. When the development of the concept itself matters, the same grammar can instead model its historical carriers. A formulation, notation, published definition, or versioned grounding bundle is registered as a representational-token Thing linked to its Concept and grounding hash. Its lifecycle and joint Events can then record formulation, naming, publication, adoption, contestation, revision, or abandonment, with people, documents, and institutions bound in their respective roles. A split or merger of definition artifacts is recorded at this level; when the asserted meaning itself divides or combines, linked successor Concept identifiers make that semantic claim explicit. The Concept record remains a Concept; it is not silently converted into a Thing.

This evidence-indexed subgraph records a *concept history*. The wider *concept world* is the selected subgraph of Concept records, grounding resources, typed relations, and histories, linked to instances through Realizations. It groups records without granting shared authority or requiring one universal ontology. Realizations connect it to ordinary world episodes,

including cases that exhibit a pattern before anyone names it. An insight Event may create a new formulation or grounding version without claiming that it created every possible meaning of the concept. A content-addressed registry preserves the definition versions; the Meaning Model preserves the historical Events and their disputed attribution. A construction should open only the few concept histories that materially affect its action; the first Book's preparation permitted three at most (Appendix A.5).

Intellectual descent can also be attributed explicitly. A historian may ask how a new mechanism drew on an earlier published method and a later discussion. An assessment Event beneath the historian's perspective process can parent a cause Cut allocating explanatory responsibility among the earlier publication and discussion Events, linked to the predecessor Concepts and grounding versions. Shared influences require a declared allocation convention or a joint answer, with unresolved responsibility retained in the remainder. The assessment retains its evidence cutoff; an UnderstandingNode explains the attribution. Grounding hashes and supersession links do not by themselves establish causal influence.

That social history differs from the reasoning by which someone defines or interprets a concept. The latter is a Concept-scoped projection of the Understanding Graph (Section 4) beneath a named understanding-process root and its clock. It gathers nodes addressing the Concept identifier, exact grounding versions or *define* Realizations, evidence, rejected alternatives, and reasons for retaining or superseding a grounding. An author's definitional history normally unfolds in construction time; incompatible reasoning histories may coexist without becoming accepted world state.

Preserving that reasoning requires a rationale UnderstandingNode for each substantive initial freeze and committed revision, addressing the new grounding, its predecessor if any, evidence, and material rejected alternatives. Without this lineage the Concept remains valid but its rationale is not preserved. Simple label registration and mechanical metadata changes need no ceremonial nodes.

The first Book supplies a concrete starting point. In E20a, Lovelace identifies two checking paths that share a prior value, and Babbage narrows his design's claim from verification to internal comparison. Opening a concept history would register the successive design formulations and link the revision to their Concept and grounding versions. The later unscored model of a valid computational claim (Appendix A.5) makes independent comparison explicit; author note U14 explains the episode. That note is not the character's own reasoning record, and these materials do not yet constitute a versioned concept-history implementation.

3.4 A process-based concept model: love and the Engine

Consider an illustrative process-based concept model of devotional love of a work, selected here as the canonical reference for a comparison. An experiencer Thing and a work Thing participate in a process with an undertaking, a response to difficulty, and continued engagement or preservation. These are linked Events, not numerical shares of a life. Appraisals and wants belong under the experiencer's inner-process root; the entire fragment has a declared concept-model context.

A proposed alignment maps the experiencer to the Book's Babbage and the work to his Analytical Engine project, distinct from the bounded apparatus and calculation office. E06, the compact of April 1844, supplies the finite undertaking; E25–26, the returned table and its disclosure, supply a setback and defensive response; E28–29 supply a changed interpretation and preservation of the complete record. Attachment can be examined through persistence, care through correction, desire through sought continuations, and sacrifice through relinquished resources or control. These are overlapping readings, not exclusive periods. This is an illustrative interpretation, not an accepted love assessment or a claim about the historical Babbage.

Suppose the Event that records the definition, beneath one named understanding process, carries a Cut allocating one comparison budget among components j, with shares α_j and remainder $\alpha_\perp$. Each component specifies a frozen, weighted set of requirement checks. An assessment Event beneath that process and about the target–component alignment carries a Cut dividing its local check mass into *matches*, *does not match*, and unresolved remainder, with shares m_j, n_j, and u_j. A recorded alignment a may then derive an aggregate fit Cut on the overall assessment Event:

$$r_C = \sum_j \alpha_j m_j, \qquad n_C = \sum_j \alpha_j n_j, \qquad u_C = \alpha_\perp + \sum_j \alpha_j u_j, \qquad r_C + n_C + u_C = 1, \tag{5}$$

where $m_j + n_j + u_j = 1$ in every component Cut. This is one possible comparator for one grounding strategy, not a universal law of concepts. Its numbers are all sibling-relative Cut shares; the Realization link itself remains unweighted. The inputs are stipulated, not recovered from the Book. They divide a comparison budget, not emotional intensity or confidence that the episode occurred. Here, the four named component Cuts each set $u_j = 0$; the aggregate $.100$ remainder comes entirely from the selected model's unspecified component mass.

For a concrete stipulated care budget, assign .50 to retaining failed results alongside successes, .40 to passing the complete record to a successor, and .10 to immediately recognizing human checking as part of the work. Let E28–29 satisfy the first two checks and Babbage's initial distancing in E26 fail the third. Matched check mass is .90, unmatched mass .10, and unresolved mass zero.

Reference component	α_j	m_j	matches	does not match
Care	.35	.90	.315	.035
Attachment	.25	1.00	.250	.000
Desire	.15	.60	.090	.060
Sacrifice	.15	.30	.045	.105
Unspecified remainder	.10	—	—	—
Aggregate fit Cut			.700	.200
Aggregate remainder			.100	

Table 5: One inspectable fit Cut against a love model selected as a canonical reference.

The useful result is not only the $.700/.200/.100$ aggregate Cut. It is the recorded claim of strong fit on care and attachment, weaker fit on sacrifice, and an alignment that treats a work as the object role. Another perspective may reject the alignment, select a different love model, use exemplars instead, or decline to score the concept without changing the target world.

Evaluative terms use the same plural grounding mechanism. A decision episode may realize the model of situated intelligence or kindness reached through one understanding process, and an assessment Event beneath that process may carry a fit Cut whose evidence and grounding version remain inspectable. Later outcomes are separate Events, so an organized failure need not be reclassified as an unintelligent action merely because it failed. A coarse statement about a whole person remains perspective-scoped testimony or a declared aggregation over episodes, never an intrinsic trait scalar. The grounding may also declare thresholds or vetoes rather than assuming that all components compensate linearly.

Constructing and revising these concept models uses the Meaning Model; learning them and testing their empirical validity and transfer belong to Life Simulation. Its "Comparison frames for evaluative concepts" subsection proposes comparisons against a grounding, within one person's history, and against matched reference episodes. Evaluator perspectives remain explicit, and decision-time judgments stay separate from retrospective judgments [7].

3.5 Fit Cuts and concurrent concepts

One Event can be linked to both quarrel and farewell. If graded judgments are useful, one assessment Event beneath the relevant perspective root for each Concept link carries its own local fit Cut rather than making the Concepts siblings: quarrel might have matches .80, does not match .15, and remainder .05, while farewell has its own three shares. The two match weights need not sum with each other because they belong to different Cuts; within each Cut the three siblings do sum to one. Section 6.7 reports fit Cuts of this kind recorded over time in the world history.

Fit and confidence ask different questions. A fragment may strongly exhibit a concept while the evidence that the fragment occurred is weak; a well-attested fragment may clearly fail to exhibit it. The fit Cut divides the declared comparison unit into matched, unmatched, and unresolved portions. If numerical confidence is needed, a separate epistemic assessment Event asks how support divides among the claim, its rejection, and unresolved alternatives, with its own anchors and evidence. It is not obtained as one minus non-fit, and the fit remainder is not a universal uncertainty score. A construction may instead leave confidence unquantified and mark the fit provisional or unscored. Confidence weights neither multiply nor replace fit weights; a well-supported mismatch remains different from a poorly supported match.

An emotion facet Cut supplies composition, not an absolute intensity scale. If activation is assessed numerically, the divided unit must be explicit. One option is a weighted set of checks for a declared "materially active" criterion: a separate assessment Event parents a Cut over matched, unmatched, and unresolved check mass. A matched share of .70 means that checks carrying .70 of this budget are satisfied, not that the emotion has intensity .70. Activation can instead remain categorical and intensity unquantified. Compatible fit or facet Cuts may be mixed across disjoint temporal children under a stated question and aggregation. Occurrence is derived from Event records.

3.6 One World, Many Cuts

The same accepted Event can participate in several decompositions without being copied. Consider Earth as a Thing and its modeled history as the bound lifecycle Event over a geological horizon. Thing-side views may open Earth into regions, the biosphere into ecosystems and organisms, or a person into body systems. Event views may link overlapping geological, climatic, evolutionary, and cultural processes, while a sequential Cut may partition the same history into disjoint eras and transitions. One evening of reading can be reached through at least two routes:

$$\begin{gathered}\text{Earth history} \to \text{evolution} \to \text{human life} \to \text{one evening} \to \text{reading},\\ \text{Earth history} \to \text{human culture} \to \text{literature} \to \text{book production and circulation} \to \text{reader encounter}.\end{gathered}$$

These are conceptual routes, not chains of normalized Cuts. Their arrows can traverse different process, participation, and concept relations. A sequential Cut divides History into disjoint eras by duration; a link from an episode to its literary concept does not divide that unit. The reading Event still binds the same reader, book, and place Things across these views. Depending on the inquiry, that Event can be opened through retinal activity, emotional appraisal, cultural transmission, or narrative interpretation while retaining its shared participants and source references.

3.7 Direction from Tendency to Want

An avalanche shows why direction cannot be a synonym for intention. A snowpack, slope, terrain regions, and material layers are Things; the avalanche is an Event bound to them. Physical-part Events may expose slabs, layers, cells, or particles. A sequential Cut may expose accumulation, instability, trigger, fracture, release, runout, and deposition only at a lens where those phases are ordered and disjoint. Loading, cohesion, friction, temperature, velocity, terrain interaction, and entrainment are linked processes or externally grounded physical data unless a separate question declares a unit over mutually exclusive answers. At low resolution an aggregate stability state and runout envelope may suffice. At a finer resolution, local processes bind exposed material Things. Its direction is physical: forces, constraints, and present state shape possible continuations. Intentional language may be a useful metaphor, but the Event ancestry and direction type must not convert it into a claim about a mind.

For agent modeling, motivation can specialize direction without requiring an unrelated primitive:

$$\text{physical tendency} \longrightarrow \text{biological regulation} \longrightarrow \text{represented want}. \tag{6}$$

The arrow marks increasing representational commitments, not a claim that every process secretly has a mind. A physical tendency follows forces and constraints. Regulation detects or counteracts departure from a viability region. A represented want exists when an actor represents or values a possible state and that representation can alter action selection. Saying that an avalanche "wants" to descend may compress an explanation; it remains an instrumental description unless an internal representation is supported.

In this grammar, a want can be represented as an inner Event contained beneath the relevant actor's lifecycle and inner-process root during a stretch of that lifecycle. A construction may link it to a desired condition of a registered Thing, Event, process, categorical state, or measured world datum. It may then open the want into formation, activation, pursuit, satisfaction, frustration, abandonment, or transformation. A plan is a proposed continuation, and a decision commits one continuation under the actor's represented beliefs, constraints, and competing wants. These are profile choices with precedents in BDI and goal-systems research [10, 11]; the grammar preserves containment ancestry, interval, question, and evidence without imposing one universal person decomposition or one utility scalar.

Dramatic tension can also be modeled as an evolving process. Two characters may depend on the same scarce resource; promises, deadlines, and discoveries change their expectations, options, and relationship. A construction can open this tension across the whole story, individual encounters, and momentary choices, tracing how it builds, changes form, or resolves. These are linked processes; any numerical Cut still names its question, divided unit, and alternatives. Suspense attributed to a reader is a separate process that unfolds as information is revealed. It need not rise and fall with the conflict in story time.

4 World, Understanding, and Artifacts

The addressed construction supports three linked surfaces. The world surface uses the six record forms. The understanding surface uses text-bearing UnderstandingNodes from the Understanding Graph [12]. The document surface holds text and code artifacts as DocumentNodes. They occupy one identifier space so that a question, reason, or sentence can address the exact world record it concerns, but they retain different semantics, clocks, and authority.

The unfunded successor, 1852–54. Chapter X of the first Book, *The Unfunded Successor*, supplies an actual construction example. At Ockham in August 1852, Lovelace recognizes that both her claim authority and Neale's independent checking depend on single people. She asks for funded training and outside supervision of a second checker. Babbage brings a mechanical self-comparison design, then recommends waiting; Halden defers the expenditure. Lovelace dies with the request unanswered. In 1854 Halden authorizes delivery without the independent check, and Guardian returns the erroneous table. These are the finished Book's authored events and character interpretations, not evidence about the historical minds of Babbage or Lovelace.

The construction's world descriptions identify the succession decision as E20, Lovelace's death as E21, the omitted check as E24, and the returned table as E25. Understanding note U05 explains why her death removes active pressure but neither inserts the wrong number nor forces Halden's decision. Chapter X renders the earlier succession episode; Chapter I opens with the later failure. Figure 2 maps these existing materials onto the three surfaces: character interpretations belong beneath their inner-process roots, U05 supplies the construction-time rationale, and the chapters supply DocumentNode text with role `story_passage`. The notes and chapters were written outside the engine and subsequently imported as native surface records (Section 7.1).

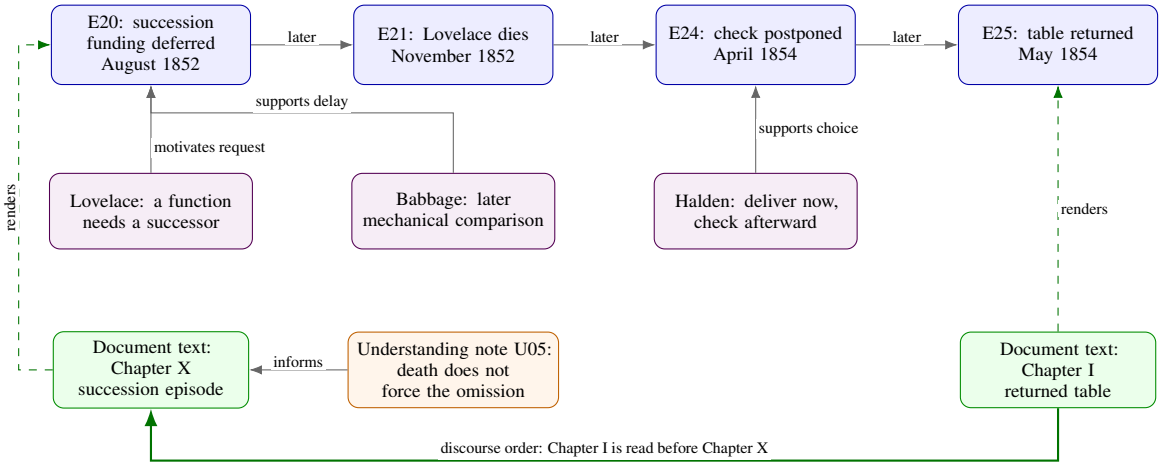


Figure 2: Selected records from the first Book, mapped to world, understanding, and document surfaces. The top arrows show story order, not causal necessity. U05 belongs to construction time; the bottom arrow shows the different order in which the reader encounters the events. Character interpretations and the author's explanation remain distinct. Edge labels are explanatory glosses, not literal Understanding Graph edge-type names.

```
UnderstandingNode { id, type, text, created at, typed parents,
    typed links, provenance }
DocumentNode { id, role, text, discourse position, typed parents,
    typed links, optional narrator and viewpoint metadata, provenance }
```

These are surface records rather than seventh and eighth world forms. Every UnderstandingNode has a typed containment path to a named understanding-process root. Every DocumentNode has one to a named document root; `story_passage` is one DocumentNode role rather than a separate record form. An UnderstandingNode externalizes a question, hypothesis, prediction, tension, surprise, interpretation, reason, criticism, or revision, among the node types of the Understanding Graph, and need not have an Event interval or a Cut. Kinds can be specialized to the application without changing these parent-path obligations. A DocumentNode may hold a report paragraph or code unit, as well as record an act of telling in discourse order; document *contains* and *next* links do not inherit Event-containment semantics.

Understanding links include `about`, `supports`, `contradicts`, `refines`, `answers`, `learned-from`, and `supersedes`. A DocumentNode links to Events through `renders` and to an UnderstandingNode through `shaped-by`. These links retain their surface-specific roles, not authority over their targets.

For example, a software investigation could link a measured response-time history to a bottleneck hypothesis, the code unit changed to test it, and an evaluation of the next measurement. The modeler can open a slow interval into requests and their phases, or refine a bottleneck concept model to distinguish causes requiring different tests. The result can change the explanation, the code, or both. The method is designed for programs too: a code unit is a DocumentNode that could be linked to the processes and concepts it implements, so that a change to the model would name the code it

puts in question, though no program has yet been built this way. This uses the same linked surfaces as the story example above.

World authority is likewise topological. Every accepted world Event has a typed containment path to History. A belief, appraisal, want, or estimate that is itself part of the accepted world sits in a named inner Event process beneath the relevant person's lifecycle. Its `about` links identify what is believed or estimated without promoting that target claim to accepted world state. A joint Event instead represents its participants through Bindings; it does not acquire one participant's perspective merely by binding that participant.

An imagined, remembered, or counterfactual scene has a named model-context Event with a declared modality and source links. Accepting that Ada imagined a successful trial accepts the imagining act, not the trial. The content's depicted interval need not lie within the interval during which she imagined it: context parenthood is distinct from temporal occurrence. A remembered scene and an accepted historical account may address the same Things while disagreeing about what happened. Promoting content into accepted world state requires an explicit accepted assertion, not a traversal through its source.

An Event may carry an optional textual description. That field describes the Event within its own record and ancestry; it is neither an UnderstandingNode nor a DocumentNode. If a description must participate in a reasoning history or in document order, a separate node is created under the applicable root and linked to the Event. Text about an Event is not thereby that Event.

Optional narrator and viewpoint metadata refine a document role without replacing its root path. Conflicting accounts can coexist across these surfaces; agreement or common ground requires a declared derivation and is not silently promoted to truth.

The density of the understanding surface is optional and explicit. A construction may preserve no UnderstandingNodes, only rationales for consequential commitments and revisions, or a finely segmented chronological reasoning trace. Absence means "not externalized at this resolution," not "no thought occurred." A dense trace provides more supervision and auditability, but is not required for a valid world. A claimed rationale trace declares which commitments require nodes and which links connect their explanations. Conversely, an UnderstandingNode never gains world authority merely because it is detailed. When its physical carrier matters, a stored note or document is a representational-token Thing and its creation, reading, transport, and revision are ordinary Events; its textual content remains on the appropriate understanding or document surface.

The Concept-scoped projection of Section 3.3 uses this same understanding surface. UnderstandingNodes carry reasons; Cuts carry numerical commitments. A non-obvious Cut may link to its rationale node, which addresses the exact Cut component or Event it explains.

Authorship may be performed by a human writer or a delegated agent. The human who commissions and is responsible for the work is distinct from a modeled fictional author persona (Section 6.4), which is a represented person, not a role: the persona's life may inform the work without making it the actual generator, narrator, or character viewpoint.

An optional account may also represent the modeler as a Thing with a declared identity, continuity criterion, lifecycle, and inner-process root in the reading or construction context and clock, not the story's History. Its changing attention, uncertainty, concern, and intentions can then be Event-processes. UnderstandingNodes under its understanding-process root record its thoughts about the world, concepts, or itself. Character-owned roots can likewise carry thoughts about concepts; the subject does not determine whose thought it is.

For either a character or a modeler, inner activity may be recorded as Event-processes, as UnderstandingNodes, through both, or not at all; the construction chooses which form to use and where to open more detail. A node can record a reaction beside the world Event, concept model, or passage it addresses without a detailed process account, and a process account need not include a dense thought trace. A derived view can collect such reactions under their declared clocks and source references without adding new claims. Inferring how concern changed between reactions is a further modeling step whose process account retains its attribution and evidence; it can be constructed directly or developed from the nodes, and neither form is a compulsory replacement for the other.

World processes, concept models, and modeler processes name what records describe, not new record forms. All can represent abstract content. The Reader of the Entangled Alignment proposal [13] is one specific modeler role, and one of that proposal's recorded readings follows its Reader's own attention, concern, expectation, and uncertainty as processes in this way (Section 11); the general account also covers world construction, writing, and problem solving.

4.1 Operative models and estimates

An actor projection contains inner records active at a given time and accepted records disclosed to the actor by then. Its inner records descend through that actor's lifecycle and inner-process root; hypothetical targets remain within that context. The subset organizing a decision is the actor's operative model. Generation of speech or action from that model is restricted to content actually recorded as observed, communications received by the cutoff, and earlier records on the same admissible paths. Co-presence permits an observation; it does not reveal every field of an Event or another participant's private state. Observation and delivery records therefore identify the content disclosed. Permission-aware ancestry traversal cannot be used to bypass that boundary. A contradiction with accepted History is therefore representable as misperception rather than being silently repaired with author knowledge. The same grammar can elaborate this projection into an attributed model of the character's self, other people, and world, with distinct model contexts for imagined or mistaken content. This is representational support, not a claim that the Book already constructs those accounts in full.

An observer's account of the actor is different. It is an estimate Event beneath the observer's inner-process root, linked `about` the target record, and supported only by evidence accessible along that observer's paths. Self-report, another character's estimate, an author's construction-time interpretation, and the actor's operative state may all disagree. Compatible Cut estimates can be compared by total variation (defined in Section 7.1.4) and categorical claims by equality; those divergences are derived diagnostics, not new semantic numbers. The profile also bounds higher-order nesting—what one character thinks another thinks—rather than granting unbounded omniscience. Life Simulation may later evaluate or learn such estimates; their ancestry and access semantics belong to the Meaning Model. A categorical, unit-bearing, or protocol-valued estimate retains the target's data type and applicable unit or measurement protocol, with uncertainty where relevant; an estimate of semantic composition uses a compatible Cut on the estimate Event. Missing records or unequal coverage indicate ignorance or noncoverage, not automatically disagreement or error.

4.2 Transactional construction, rendering, and read-back

World records, UnderstandingNodes, and DocumentNodes share one commit discipline. The accepted parent is frozen. A proposal is a whole candidate delta containing every new or superseding record required across the three surfaces. Validators inspect it against that parent, including every new typed root path. In the proposed workflow, the author or a designated checker decides whether to accept a validated candidate. Rejection leaves the active graph unchanged, although a construction log may retain the candidate and its explanation. Acceptance commits the candidate atomically. A later change requires a new candidate, explicit `supersedes` links, and revalidation of the affected dependency closure. Current review records do not authenticate the reviewer, and not every acceptance path requires a review record.

A commit envelope records what the candidate read and replaces; its fields and the rules for stale bases and branches are in Appendix C.

Revision preserves predecessor records and revalidates affected dependents. Operators, including opaque ones, must declare the records they read; source-bound views and passages also record derivation dependencies. Those declared links determine which dependents must be rechecked, regenerated, or marked stale. A generic `about` link alone does not define such a dependency. These are the required semantics; the current Rust host does not implement every cross-surface step (Table 7 states the implementation boundary).

The Introduction's illustrative trial now has a record-level reading. Its Event, grounding Realization, assessment Cut, rationale UnderstandingNode, and rendered DocumentNode address exact revisions. A changed success claim enters as one candidate containing the affected records, with declared dependents rechecked or marked stale. Co-storage alone does not establish their agreement; the acceptance checks and semantic judgment do that work. Section 7.1.7 reports a small synthetic trial in which these steps can be run: an opening checked against its frozen parent, and a revision that marks its dependent passage stale.

A direction Cut can make the generator's freedom inspectable. Its children are mutually exclusive candidate continuations and its remainder retains options not articulated. A draw that selects the remainder requires a new admissible proposal; it must not silently discard that mass and renormalize the named answers. The constructor may fix mass needed by the premise or an earlier commitment, draw among the remaining alternatives with a recorded seed, and then build the selected continuation as a complete candidate. A poor draw may be rejected under declared gates, and a construction may draw several times and keep the most interesting (Section 6), but the method asks that every draw and the reason for the choice stay in the record: a draw is never silently rerolled until the author likes the answer. Rejected siblings and reasons remain addressable. Thus the model can materially decide the story while authorship remains bounded by declared constraints and visible intervention.

A rendering route records the source commit and epistemic status; an ordered route through Events and its currently opened frontier; optional narrator and viewpoint metadata; the evidence accessible at its declared cutoff; a disclosure policy; and the target DocumentNode with role `story_passage`. It specifies which evidence the renderer may use under the chosen viewpoint and disclosure policy. In character-limited narration, this excludes later facts, private state unavailable to that character, and undelivered documents. Other narrative viewpoints obey their declared access and disclosure rules. Enforcement and read-back must check that the generated prose respects those limits. The route is a convention of the construction profile over existing references, not a seventh world relation.

Read-back applies the inverse operation in a fresh context: infer the Things, Events, Bindings, Cuts, Realizations, typed ancestry, and links implied by a story-passage DocumentNode, then compare that proposal with the source commit. A mismatch can expose an omitted record, an unsupported sentence, a wrong root path, or prose that implies a different composition. Its resolution is another explicit candidate: revise the DocumentNode, revise the model with justification, or retain a declared ambiguity. A read-back UnderstandingNode under the named review process preserves the disagreement and decision. This closes the loop from model to prose and back without treating prose as the sole authoritative copy.

Readers track situations and update their accounts when perceived event boundaries disrupt expectations [14, 15]. Read-back asks where such an update no longer follows the intended route; the reader's interpretation remains distinct from the modeled interior and the author's reason.

5 Closure and Sufficiency

Closure is claimed only relative to a declared purpose and resolution. This section first asks whether the same records can be refined recursively without erasing accepted commitments, then turns that capacity into a finite-resolution sufficiency test.

5.1 Recursive refinement at a declared resolution

Suppose a finite-resolution world contains the world records above. Opening a Thing adds Things and Events that preserve the parent's identity boundary. Opening an Event may admit qualified child answers under a declared Cut, expose concurrent Event-processes, add linked episodes and updating Events, or expose typed world data. Only a Cut invokes a one-unit composition contract; a sequential Cut additionally requires disjoint intervals and derived duration shares. Describing a target adds a Realization using one declared concept grounding. If that grounding is a process-based concept model, its role-bearing Things and Events reuse the same world grammar. In this restricted sense, world-record refinement is recursively closed.

The grammar places no intrinsic lower or upper bound on modeled resolution. A cell, person, institution, planet, or galaxy uses the same Thing–lifecycle pairing when its identity and persistence matter; organelles, group members, or stars can be exposed through time-scoped relations; and biochemical changes, decisions, discoveries, mergers, or explosions are Events at the resolution in which they matter. Containment and temporal refinement provide hierarchical views, while overlapping processes and cross-resolution update links keep the full representation a graph rather than one tree. Domain content remains specific: a cell may require biochemistry and a galaxy astrophysics, and neither acquires an inner-process ancestry merely by sharing the grammar. Increasing model capability may open unresolved remainder, propose and test better boundaries, processes, unit-bearing world data, and candidate domain laws, and check compatibility across resolutions. It need not merely add detail; a better abstraction may justify an explicit revision that supersedes many weak records with a smaller, more predictive account. This is a claim of representational extensibility, not that the present model knows every resolution or that every process follows one dynamical law, such as the jump diffusion of Section 2.3.

As a long-term theoretical goal, the representation could support opening a scene into molecular or particle-scale processes where the inquiry requires them. Reactions and collisions are Events; tracked entities can use Things with declared identity and continuity. Collective or field descriptions need not assign a separate Thing to every constituent. Attached physical or chemical models could compute their evolution, retaining their units and approximations. The finer results must preserve the protected coarse answers under declared mappings and tolerances, or prompt explicit revision. What changes at these depths is the domain model and its evidence standard, not the six record forms.

Psychological resolution follows the same rule. A person profile selects questions for a construction, not the expressive limits of the grammar. There is no prescribed maximum depth or universal personality tree. A coarse account of an authored character's resistance to delegation might be opened into concerns about accuracy and recognition, beliefs about a colleague's competence, and memories of broken promises. Further detail could distinguish private collaboration

from public assessment, or trace how later experiences alter that difference. These remain interpretations in the relevant actor or observer context, not psychological facts established by adding detail.

A more capable model could choose different questions, boundaries, and depths according to what a scene or inquiry needs to explain. Competing motives may form a Cut under a declared attention question; beliefs, memories, and relationship Events remain linked processes without being normalized together. An opening under the same question must preserve its declared coarse commitments. A different question creates another view, while a changed account requires explicit revision. Thus psychological refinement can elaborate, reorganize, or simplify a person's model rather than merely populate a fixed template.

More capable systems, potentially including artificial superintelligence, may extend this work to detail far beyond categories humans can readily understand. Protected coarse queries can still expose a readable account of fatigue, trust, or responsibility. Finer records must reconstruct those answers within the declared tolerances or revise the commitments explicitly. A Cut's remainder keeps unallocated share visible under its question. These contracts preserve continuity; comparisons on unseen cases must test whether the unfamiliar distinctions improve the account. Meaning Model supplies this representational freedom and its preservation contracts; Life Simulation considers how to learn which refinements are useful.

This is not a theorem that every true explanation is a weighted tree. Cuts may cross; Events can be both causes and consequences under different questions; and domains can attach specialized measured data, laws, fields, or solvers. Those attachments are content referenced by Events. Nor does the claim imply that cognitive testimony or prose must be encoded as Cuts. The Understanding Graph and document surface remain interoperating representations with different jobs.

5.2 Finite-resolution sufficiency as a test

A representation is sufficient only relative to a query, tolerance, and cost. For a story-coherence query, the identity of the Engine, the intervals of key Events, who held each belief, and how fine emotions recompose may be sufficient. For stress analysis of a gear, the same representation must point to geometric and physical state that this paper does not attempt to replace.

The bounded constructive argument is deliberately modest. Any finitely specified simulator with computable or approximable transitions over a bounded horizon can be wrapped in one root Thing and its lifecycle/root Event: place the simulator configuration in Event state, reference its transition semantics as direction, and expose its outputs through typed observations. If nothing more is known, leave the root opaque. Later useful identities and processes can be opened through Things, joint Events, and Cuts while the root behavior remains a consistency check. Bare representational universality is therefore cheap. The substantial hypothesis is that useful stable identities and Cuts exist and improve compression, prediction, explanation, control, or coherent generation enough to justify their cost.

The sufficiency claim can fail in three ways. A required distinction may fit no record profile. A composition rule may destroy information needed at a coarser level. Or the cost of explicit articulation may exceed its value. Such failures should produce a remainder when a governing Cut exists, leave other detail implicit or add a domain attachment, or lead to rejecting the grammar for that task. They should never produce a declaration that the world has been fully represented.

Cross-Cut identity is the main defense against false trees. Einstein can take a large share in a scientific Cut and almost none in another person's daily-life Cut without being duplicated. A relationship can participate in emotional, economic, and legal Cuts. Local weights answer local questions; the shared Thing and Event identifiers preserve that these descriptions concern the same world.

Five questions provide the evaluation contract:

1. **Coverage:** Can the target be represented without changing the Concept, Thing, Event, Binding, Cut, and Realization grammar?

2. **Reference:** Do descriptions and roles resolve to the intended continuing Things across time and opened detail, while explicit split, merge, replacement, and succession Events change identity when required?

3. **Grounding:** Where a comparator or recognizer is declared, does it distinguish held-out instances from counterexamples under that grounding? Otherwise, do independent reviewers agree with the construction's applications to held-out instances and counterexamples under its declared grounding and context? Both checks retain unresolved cases. Calibration applies only to declared probabilistic outputs, not to fit shares of requirement-check mass.

4. **Recomposition:** Do admitted children recover their parent's declared Cut commitments within the frozen tolerance, and do Event-process links, role Bindings, identities, and typed world data satisfy their own contracts without being treated as Cut weights?

5. **Value:** Does refinement improve prediction, explanation, control, or coherent generation enough to justify construction, validation, inference, and maintenance cost?

The value question is asked repeatedly during construction, not only at the end. Whenever values are recorded, the report the current implementation returns (Section 6.1) ranks the processes that have values but no parts by signals the model already holds: how important the Events that move a process are, how many of its turns no note or Event explains, how rough its values are, and how near the root it sits. It names the one most worth examining. The modeler answers by carving it, or by recording that it stays whole, with the reason and the evidence the decision was taken with; the question returns only when that evidence grows. A decision to keep a process whole for now thereby becomes as inspectable as an opening. The order of opening in Section 2.2 calls for such decisions; the tool keeps each one, with its reason, in the model. The ranking is a heuristic judged on one model, not a measure of marginal value. It is the construction heuristic the implementation applies; learning when opening detail pays, and testing whether it does, belongs to Life Simulation [7].

6 General Modeling, Memory, Story Construction, and Mechanism Search

This section describes the method as the tool carries it out. A language model reaches the tool through the Model Context Protocol (MCP). The tool has three parts: an engine written in Rust, which stores and checks the records and which this paper also calls the Rust host; a server that offers the engine's operations over MCP, with the optional add-ons; and a read-only viewer (Appendix D.2 gives the interface's entry points and operations). The MCP interface exposes general-purpose modeling, including agent and user memory, together with optional storytelling and mechanism-search workflows and a forecasting add-on, over the same model and graph. A general model may connect technological development, adoption, institutions, public interpretations, and markets, or describe an entirely different physical or organizational system. A price series is one possible process, not the boundary of a market model. The modeler chooses the processes, descriptions, units, evidence, temporal resolution, and useful depth. Sharing the representation does not by itself reconcile separately constructed accounts or establish their causal relations.

6.1 The construction loop

The workflows fall into two families that share one construction loop but not its strategy. Generative work makes something new: stories, game worlds, and the invented worlds of mechanism search. There, where a decision or a trajectory has real alternatives, draws and random words propose what the constructor would not have chosen, and the most interesting is kept with a recorded reason. A random word can likewise seed a trajectory, a person, or a process. Descriptive work accounts for something that exists: a market, an institution, or, in memory modeling, a person's own life. There the constructor asks for the data the model lacks, infers what the evidence supports, estimates what neither supplies, tags each value as read from a source, inferred, or guessed, and lets later evidence test the account. A direction Cut there is a forecast only when its weights are declared as probabilities over the outcomes before the outcome is known. The answer that happens is realized from the evidence, not drawn, and a draw only samples a hypothetical scenario. No factual outcome is selected for interest, although curiosity can guide which concepts and hypotheses to explore. Both build structures on existing structures and refine them, and both keep the same records and preservation contracts.

Memory modeling maintains the ongoing processes of an agent, a user, or their shared work. Goals, decisions, beliefs, and learning are represented through their histories, circumstances, and consequences. Later interactions consult and refine this account through the same construction loop. User reports, agent interpretations, and accepted decisions retain distinct sources and authority. Experience thereby produces explicit structure: a distinction found in one interaction, for example between interest in an activity and capacity for the whole undertaking, becomes an anchor for interpreting the next. A correction supersedes an earlier account without erasing it and remains distinct from a real change over time; because dependent interpretations are not revised automatically, new reports prompt the agent to re-enter the accounts they affect [12]. For an agent, this is an account of its expressed work and learning, not a record of hidden neural change.

Shared instructions ask the agent to explore recursively across all workflows: follow useful discoveries into subprocesses, relationships, competing accounts, and revised categories, then reconsider the enclosing whole. The constructor is encouraged to discover implicit but trackable concepts and processes, make them explicit, and use the resulting

structures as anchors for further discovery. Qualities usually written as fixed descriptions, such as a voice, a style, or a belief, are modeled as processes where they can change. Standing questions ask what process, relationship, or distinction nobody has named yet, what a pattern is an instance of, and what that abstraction predicts elsewhere. In storytelling they also ask the constructor which assumptions it has not modeled yet, and which of them could stretch back far in time to create a deeper story. Construction therefore expands both the represented account and the conceptual vocabulary through which it is understood. Physical placement and its consequences are developed where they matter to the inquiry. Diagnostics expose missing declarations and unopened structure; they do not certify semantic completeness. The agent chooses what to open next and records why, or why a process stays at its present resolution for now. Neither a fixed inventory of categories nor numerical coordinates are required of every model; what every opened process requires is its state over time, in the form that fits it. Figure 3 summarizes the loop and how its work is divided between the instructed agent and the tool.

The instructions make construction the work. Their goal is to build structures on top of structures, interpret them better, and refine the ones already built, so that the world becomes more real: processes open into sub-processes and finer trajectories, new structures depend on the existing ones, and higher-level processes, concepts, and laws are composed from what the records show. Reviews, readings, and checks serve construction rather than replace it. One early construction showed why this has to be said: for about 130 model revisions its processes stayed at 97 while its Cuts, mostly lens readings (Section 6.3), grew from 46 to 345, because each check returned more items to clear. The tool now asks for a reading only where a decision, a turn, or a passage depends on it, has readings and telling phases (Section 6.4) renewed together once the text settles, and asks the model a question when its readings outgrow its processes or its processes stand unrelated. Recording values follows the same principle: after each batch the tool points down the two trees, to contradictions between a whole and its parts, processes with no recorded state, processes never divided, and the coarsest stretches of time, rather than asking for more readings of what is already read; Section 5.2 says how it ranks the processes to examine.

Outside reviews remain useful, and a review may rightly confirm the account or correct only the prose. But an agent that only reviews and revises the prose, while the model and its processes never change, is not using the tool; it is revising a text. Where a review exposes an unmodeled assumption or an inadequate process, the instructions ask the agent to develop the model and bring the dependent prose into line, and otherwise to record why the model stands. The review tool names earlier reviews of the current revision: a signal to check whether the work is circling, not proof that nothing else changed.

6.2 The language model, estimator, and storage loop

The division of work separates open-ended construction from bounded evaluation. The language model develops process descriptions and candidate explanations, guesses and records the states of the processes it opens, writes DocumentNodes, and records concise assessments and revision reasons in UnderstandingNodes. It also designs the numerical questions and revises inadequate categories. An optional estimator supplies structured judgments for those questions. An estimator is a separate model that answers a batch of such questions with typed options, so that the language model need not write out each value itself. None of the three constructions' published records carries a value from one. Deterministic tool code validates, types, and stores the results with their coordinates and provenance, avoiding a second language-model pass that manually transcribes each value and graph edge.

Answering a bounded question with a classifier's decision rather than generated text predates the current provider, named in Appendix D: SetFit, for example, fine-tunes a compact sentence encoder and fits a classification head for a fixed label set, without prompts [16]. The estimator is a replaceable instrument. It does not supply an independently verified account of reality, and its use is independent of enabling storytelling.

The general workflow is:

1. Establish the question, system boundary, time basis, evidence, and delegation already agreed with the user. Review the enclosing system and longer-term developments before selecting sector-level and focal processes; examine feedback from those processes into the broader account. Build a first coarse account of the whole promptly rather than an exhaustive plan, then widen and deepen it.

2. Supply a compact scaffold of referents, processes, events, sources, and descriptions. The tool expands the initial model and graph. Supplied numerical measurements retain their units; missing states are estimated, each with an honest band and a tag, so that every opened process carries its state.

3. Record the modeler's own dated values, states, and scale readings as value series in the model, where the engine checks them. Or obtain typed process-value proposals from an estimator for the selected process coordinates, and validate them. Preserve the exact questions, answers, and model revision for review.

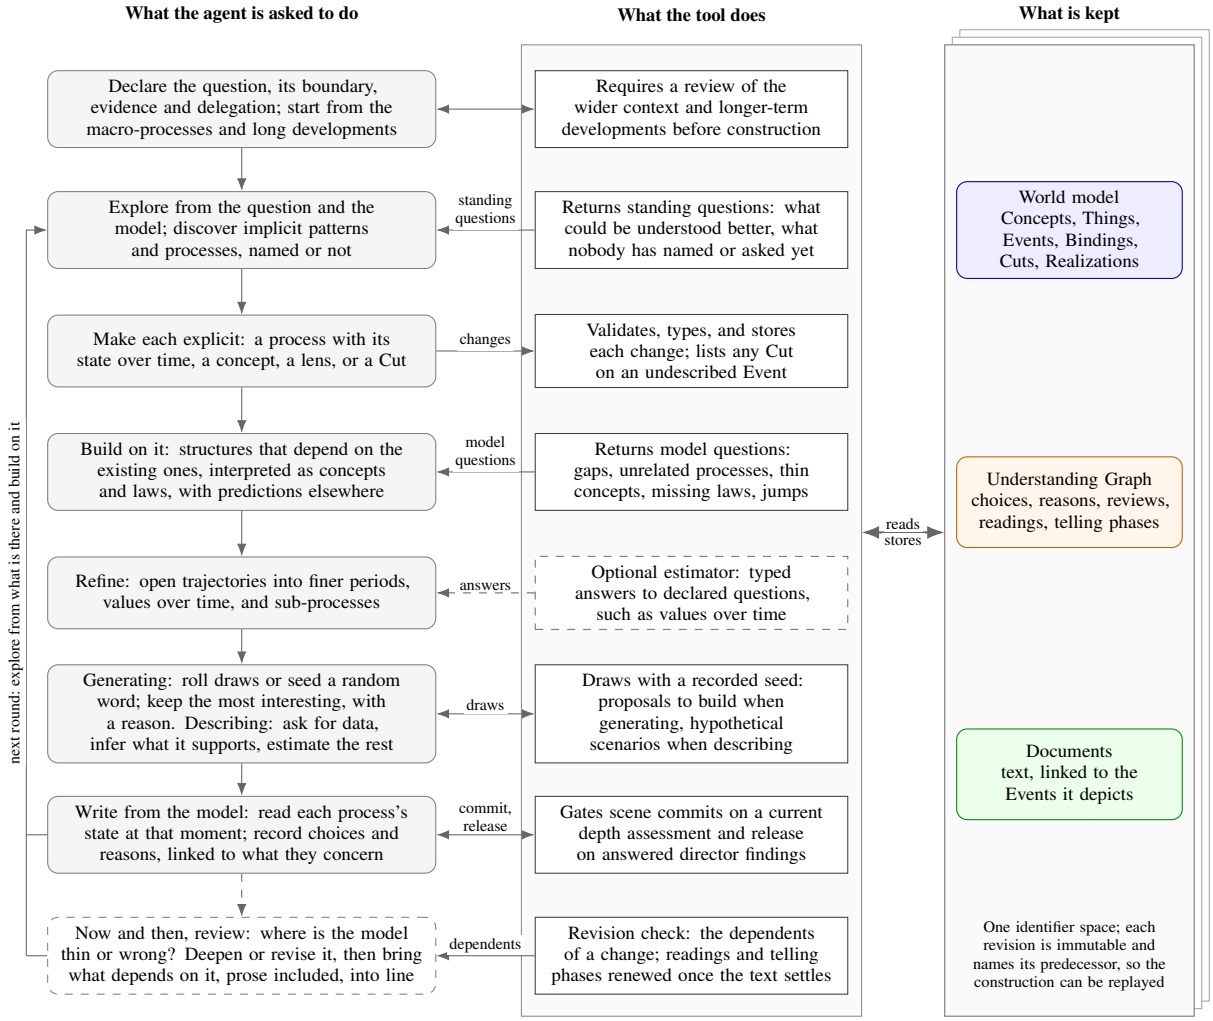


Figure 3: The construction loop. Left: what the shared instructions ask the calling language model to do. Construction is the work: the model declares the question and starts from its macro-processes, explores and discovers implicit patterns and processes, named or not, makes each explicit, builds structures that depend on the existing ones and interprets them, and refines their trajectories. When it generates, draws and random words propose what it would not have chosen where a decision or a trajectory has real alternatives, and it keeps the most interesting with a recorded reason; when it describes, it asks for data, infers what the evidence supports, and estimates the rest with honest bands, so that every opened process carries its state. It writes from the model's state at each moment. Each round explores from and builds on the structure the last one left. Review (dashed) comes now and then, from a blind read-back, the director, another model, or a person, to find where the model is thin or wrong; the model is then deepened or revised, and what was built on it, prose included, is brought into line. The order is a guide rather than a fixed sequence; any step may lead back to exploration. Middle: what deterministic tool code does. It validates and stores every change, returns questions after each one, records draws with their seeds, refuses gated operations until the required reviews exist and their findings are answered, and lists the dependents a revision puts in question. The estimator is optional, and the scene and release gates and the seed-word and trajectory tools belong to the storytelling add-on. Right: the one versioned record that every step reads and changes through the tool. The tool checks records and their declared links; it does not supply the interpretations or certify that an account is adequate.

4. Review and apply the exact construction proposal, or record a later estimate in the model's value series, where it replaces the earlier value at that time while the revision history keeps it. UnderstandingNodes explain what was retained, questioned, or revised; they carry the reasons, not the values. A retry must not silently regenerate a different set of values.

5. Use the account, inspect missing evidence and competing explanations, and deepen or revise the processes where this improves the inquiry. Preserve earlier versions and reassess affected interpretations and documents. The accepted account becomes context for the next round: its explicit processes, relationships, and concepts support further inference and abstraction, as well as constrain proposed detail.

Before external estimation or construction, the compact builder requires an attributed review of the wider context and the longer-term developments behind the focal question. Each is represented, left explicitly unknown, or excluded with a reason, and the review is stored as UnderstandingNodes. Recording the review does not establish that the account is explanatorily adequate. The builder also requires consideration of whether authored numerical judgments, a finer conceptual decomposition, or a change of meaning across dates or perspectives would help; that consideration does not require scores without a unit or declared anchors, finer categories, or a claim that meaning changed. A long event is not a sampled numerical trend: numerical histories require values at explicit dates, kept as value series or as claims and estimates with their evidence cutoffs.

A probability distribution over answer categories is not automatically an attention allocation, a physical share, or a causal law. Rubric-derived scalar values require declared numeric anchors and units; their source distribution and uncertainty remain inspectable. Public statements support attributed accounts of a person's expressed views, not privileged access to private reasoning. Source selection and question design remain substantive modeling choices even when numerical evaluation is automated.

The architecture is intended to reduce repetitive language-model work through batched judgments and automatic ingestion. An end-to-end efficiency claim must include evidence preparation, question design, provider calls, validation, review, and repair, compared on equivalent tasks. Structural tests or a single live provider run do not establish calibrated judgments, predictive accuracy, or superior fiction.

6.3 Lenses and revision review

The core lens tools make a revisable question reusable across records. A lens records the question and its current alternatives, with readings located under the appropriate holder's root and linked to what is read. Quantitative answers belong to those reading Events as Cuts, retaining the question, unit, and remainder. They do not acquire the authority of the world record they interpret. Diagnostics expose unanswered readings and possible finer openings. A useful answer can be opened through conditional Cuts, or the agent can record why the present resolution suffices or the lens does not apply. Templates are optional. Changed source text or lens versions can make readings stale; the workflow supports reassessment through a configured estimator. Source-text freshness distinguishes an unchanged recorded signature from changed text requiring review, an unresolved source, and an untracked basis. An unchanged signature is neither a new assessment nor a certificate that the reading is true.

After a model revision, a check compares specified model versions. It follows declared dependencies to conditioned Cuts, draws made from changed weights, readings of rewritten text, and passages that render a changed record. It also lists for review Events connected through recognized causal relations and visible understanding notes directly linked to changed records; an `about` link does not make a note a dependent. Telling phases are checked separately, against the current text and reading order, independently of which model revisions are compared. The check returns work to review, without repairing or certifying the construction. Coverage is limited to recognized changes and visible links. In the current prototype, changes to concept definitions require manual review of their linked dependents. Semantic agreement, disclosure, and dependency completeness still require assessment.

6.4 Storytelling over the same core

The model is intended to do more than check prose. It supplies material: modeled dates, alternatives, and quantities become the content of scenes, and because a decision needs represented alternatives and causes, the constructor must supply specific constraints rather than atmosphere. It serves as the writer's memory across a long work: dates, sums, and who knew what at each point are read from the model and the disclosure records at every edit. It keeps revision local: declared dependencies name the facts, phases, notes, and passages that a change affects. Its diagnostics can also point to where the story turns, since a ranking of decision and state jumps is derived from the model before any reader responds. The guidance also asks the writer to compress the story into a sentence, a paragraph, and a page, and to

write down the ideas the work means to convey with a strategy for conveying each, which a blind read-back can test. The sentence is the coarsest description of the work and a first test of whether it is interesting; as the telling grows, it should still fit that sentence, or the sentence is revised explicitly, as with any protected coarse answer. These are design expectations; the prospective comparison tests them.

The storytelling add-on adds author and character models, overall life trends, numerical trajectory exploration, scene preparation, and prose review. Intake separates which initial parameters the user supplies from how involved they wish to remain during writing. A fully delegated language-model author may perform the construction and writing loop within that brief. It writes the prose and evaluates author and character voice against specific text and modeled processes, recording those assessments in UnderstandingNodes.

A story's world may follow our world's laws or depart from them, at a realism level the user chooses: our world as it is, our world with exactly one departure, our world with several connected departures, or an invented world with its own fundamental laws. A departure is proposed from a seed word as an unaccepted candidate and, once chosen, is recorded in the model like any other commitment. A story can also keep our laws and change our history, as *The Book of Conditions* does with one decision in 1843.

A human author can instead write the work and use the language model for feedback. The model explores the text and its underlying processes, linking observations and proposals to evidence in the work and, where available, its modeled world. Authorship and creative decisions remain with the human; feedback does not accept a model revision or rewrite the prose. The same open-ended inquiry applies without imposing a fictional author persona or a fixed literary form.

Before preparing a scene for commitment, the add-on requires a current model-depth assessment of the stored plan and its explanatory evidence. The server checks the referenced records and source identity; the calling model judges sufficiency. For stories, the process guide fixes two things: the depth it illustrates, as a minimum, and nine slow processes at the top of every person's life (body, kin, partnership, work, place, means, knowledge, standing, and meaning), which the coverage report checks for every person in a story; below and beside them, the tree is the constructor's to design. After substantial drafting or revision, a separate actual-prose review considers purpose, realized author voice, and each relevant character's speech and actions against the modeled life and situation. Its literary advice is advisory, not a quality gate. Exploratory drafts and questions may precede either review; neither review adds a character taxonomy of its own or requires an explanation for every ambiguity.

A director review assesses the world before the first scene, and the draft before release, against stated principles: for instance, that no death or accident does the plot's work, that the route passes through the model's largest jumps, and that each recorded idea reaches the reader. The reviewer judges each principle against the work's form, purpose, and context, and may explain why it does not apply. Each failing finding names the model or prose change it needs. Several of its principles were distilled from understanding notes recorded while the Book was constructed. Release requires a director review of the current prose and a later answer to every failure found at either stage, with a changed model or a fresh passing review of changed prose where the finding requires it. The server checks these records; whether a repair succeeded remains a judgment.

A deepening pass can revise an existing work and its supporting model. General-purpose modeling does not require author personas, dramatic goals, or any other literary vocabulary.

New or changed prose declares which Events it depicts, or records why it depicts none. Revision and release checks follow these declarations; they do not infer depiction from similar words or establish that the prose is faithful. The story release operation widens access to selected prose without changing the access scopes of other authoring records.

Depicting an Event does not disclose all its facts or motives. Disclosure plans distinguish what a reader can already know, what a passage reveals, and what remains withheld or ambiguous, separately from character knowledge in world time. They use existing passage identities, reading order, linked author records, and optional document spans. These attachments support revision as text grows or is split; their meaning and scope still require review after an edit. This introduces no new clock system or required numerical reader model.

Optional telling-process records describe qualitative phases of disclosure, attention, tension, or another question chosen for the work. These are attributed assessments recorded as UnderstandingNodes, interpreting the realized text separately from a disclosure plan's intended revelation and withholding. Each phase names a stable document span and quotes its supporting passages; the record retains its holder, reviewed text hashes, and reading order. The bundled viewer (Section 6.5) projects these phases in document position, with current, needs-review, or unresolved status when the text, order, or accessible span changes. These are attributed readings of how a work unfolds, not measured reader psychology or simulated world time.

6.5 The construction record

The workflows treat the model and its Understanding Graph as the modeler's understanding rather than a report about it. Work that is done and not recorded cannot be continued by another agent, and a thought that is not linked to what it concerns cannot be found again. Every Event that parents a Cut therefore carries a description of what happens in it, because the Cut's question and weights are read against that Event. A model write lists any such Event that lacks one, and refuses the write when asked to; the tools that place Cuts for the estimator or prepare a scene refuse or block it. Choices, ideas, predictions, voice decisions, and their reasons become UnderstandingNodes held by a named holder and linked to the exact records they concern, down to a single Cut answer. Reviews by other models, blind readers, estimators, or people are UnderstandingNodes held by the reviewer, with the revision and text it read. Because every stored revision is immutable and names its predecessor, the recorded construction can be replayed: the model revisions each step adopted and any reasons, notes, and reviews saved with them, with each note beside the records as they were when it was written. Replay cannot recover unrecorded reasoning or revisions excluded from an export. An agent continuing an unfamiliar story or model recovers its history through replay and an outline before changing it. When the exact previously read revision and its context are retained, the agent reads subsequent changes and the records relevant to the current work. An outline presents the current state at a chosen depth. The construction also travels: a full-history export holds the recorded model revisions and each graph revision as its change, and another engine rebuilds it with the same hashes, so the replay there is the same replay.

A bundled read-only viewer presents the same records: coarse and detailed views retain the addressed wholes; spatial views use declared positions in explicit frames, while other layouts claim no physical geography; connections to Events, passages, and understanding follow recorded links; author and reader lives keep their own worlds and times; and reading order remains distinct from world time.

6.6 Mechanism search through invented worlds

Asked directly for ideas, language models converge. The answers of 22 models to three divergent-thinking tasks resembled one another far more than the answers of 102 people did; prompting for creativity raised their variability only slightly, and raising the sampling temperature closed the gap only once most answers had become gibberish [17]. The optional alien add-on seeks difference by construction instead. It implements the world-diversity search of Ontology of the Alien over the same graph [18], as a method in three parts: invented worlds that find kinds of solution no one has considered, a curated graph of kinds that shows which kinds are still empty, and verification of each candidate. A target-blind builder turns a seed word into an invented world; a purpose-blind solver addresses the problem inside that world; a compiler states the operative causal relation in the problem's own domain. The server writes each role's task, stores its exact text, and records every output against the task it answers. A transfer maps a mechanism's roles onto records of the target model and states where the analogy breaks; a graded fit is a fit Cut with its own remainder.

The worlds searched need not be far from ours. Target-blind invented worlds are the far end, where familiar assumptions are least likely to survive. The builder's world operators also reach nearer: a world whose history branched at a chosen point, in which an institution is inverted or a capability is removed or substituted, or in which something is treated as already and simply true, as a law, an oracle premise that is target-aware by design. Nearest of all are branches of a model of our own world, which keep its laws and change its history. The forecasting add-on forks sibling branches from one model revision and information cutoff, under the answers of a forecast, the options of a decision, several assumptions, or a change to the past, and keeps each branch apart from the trunk. A solution can then be sought inside a branch: what would have followed had a crisis been handled differently, or a policy begun ten years earlier. A branch cannot be checked against what would have happened, so it generates hypotheses rather than evidence: its value lies in suggesting mechanisms and stress-testing plans, which must then be tested in the actual world. Scenario planning, counterfactual history, and agent-based simulation explore the same ground [19, 20, 21].

A world is not only written. A target-blind modeller writes the world's laws as processes and laws of the Meaning Model, maps each law to the rule it expresses, and lists what the model leaves out; it may decline a world whose physics is not a quantity. The server runs the model in the engine and records it. The solver then sets the problem inside the world before solving it: the problem's actors become processes under the world's laws, one measure of the problem is fixed before any run, and up to three interventions are proposed as named parts. The server runs the scenario, each intervention, and each intervention without each of its parts, all from the same start. Each run is drawn five times, and the solver sees the mean and the range of the draws. A run depends only on its model, plan and seed, so every run can be replayed from the graph. The aim is that the AI can find connections no one wrote down. Running a world's laws shows consequences its builder asserted but never derived, and removing one part at a time shows which part of an idea does the work. The compiler and the transfer read the runs. A run is a hypothesis generator, like a branch of our own world, not evidence: it shows what follows from the laws as modelled.

In one world, built from the seed word geyser, no quantity passes a constriction smoothly. It accumulates and discharges at a threshold; a wide channel leaks instead of erupting; and channels fed from one reservoir drain each other. The solver set urban loneliness in this world. Each resident's reaching out is a confined channel, and a reach-out is met only when another neighbour's arrives within one tick of it. With no intervention, no reach-out was met. Gathering at the great vent's surge, with a relay that passes part of each burst to neighbours already close to erupting, met 0.93 of reach-outs (0.90 to 0.96 over five draws). Without the relay this fell to 0.52, and without the gathering to 0.73 (Table 6). The obvious remedy, wider channels so that residents reach out more often and more cheaply, met almost none (0.01), because in this world a wide channel leaks instead of erupting. The solver turned that run into a rule, and the compiler kept the relay as the part that did the work. The compiler wrote that "cheap pings let the social urge leak away without anyone meeting".

Run	Mean of five draws	Range
The scenario, no intervention	0	0
Meet at the Long Vent (gathering, relay, welcome)	0.93	0.90–0.96
without the gathering	0.73	0.73–0.76
without the relay	0.52	0.51–0.54
without the welcome	0.82	0.73–0.92
Quicker reaches (wider channels for residents and newcomers)	0.01	0–0.05
without the wider channels for residents	0	0
without the wider channel for newcomers	0	0

Table 6: The geyser world's loneliness scenario as run by the engine. The measure is the share of reach-outs met by another neighbour's within one tick; higher is better.

In one search on that problem, ten invented worlds were modelled and run, giving 115 runs, and each run replayed exactly from the graph. The resulting ideas were rated blind by a jury of Claude models, against the ideas from the same ten worlds in an earlier search where the worlds were only written. The modelled ideas were rated more likely to work in eight of the ten worlds: 4.3 against 3.1 on a seven-point scale, paired Wilcoxon $p = 0.016$. They were rated less original: 3.9 against 4.7. They stayed more varied than ten direct answers from the same model (mean successive distance 0.64 against 0.38), but less varied than the written-only ideas (0.80), with each idea cut to twenty content words for this comparison. The two searches also differ in the tool version and in one added instruction to every role, to say plainly where something does not work. So the comparison shows a difference, not its cause; a search with world modelling switched off on the same version would separate them. The search is recorded in `examples/alien-search-loneliness`: an atlas of its worlds, runs and mechanisms, the diagnoses, the jury's ratings, and the replay of every run. Its construction file holds the search's Understanding Graph with the world models, the seeds and the records every run replays from, so a reader can import it and rerun the models.

Curators maintain three revisable ontologies: mechanism families, claimed outcomes, and causal world regimes. Their partitions are specializations under a named lens, not expressive decompositions. A new family is admitted only when the recorded equivalence test says the primary causal operator changed; the judgment itself remains the curator's.

The ontologies also steer the search. A curator can carve kinds under a lens before any candidate fills them, and the diagnosis names every kind with no candidate anywhere in its subtree, beside the family and outcome combinations that no candidate yet realizes. A commission can name such kinds as its target: the next explorer is asked for a candidate of a named mechanism family or claimed outcome, and the next world for a named regime, whose description is checked for target terms like any other request to the builder. Whether forcing worked is read back from the curated ontology rather than from the answer's own claim: the diagnosis reports a commission as filled once a curator places an answer in a named kind or below it, and as missed when every answer lands elsewhere.

Candidates are verified with PURE, the viability check of Fractal Intelligence [8], in the form of the PURECheck card in Sema's vocabulary [9]. Its four gates, Parsimonious, Unique, Realizable, and Expansive, are each judged by a separate verifier, in a context that sees the problem and the candidate but no other gate's verdict; no gate compensates for another, and each returns verified, falsified, or unknown. As in Fractal Intelligence, each gate is a decomposition point: the verifier opens its question into probes and answers each the same way, going deeper where an answer is uncertain and would change the verdict, and stopping where the marginal-value rule says another level would cost more than it yields. A probe that needs a fresh look becomes its own verifier task, whose status stands in for the probe. The probes compose by each gate's criterion: every part needed, distinct from every incumbent, every part realizable, and carried over to at least one domain unlike the problem's. A deeper finding therefore changes the gate through the same rule. The server computes the result from the records, and a candidate is verified only when all four gates are. Uniqueness is judged against known practice and the candidates found before the one checked, never against the candidate itself.

An alternative explorer works directly from the target problem, with or without a random-word cue. Both routes can use no population context, an inventory of earlier approaches to avoid, or the curated map. The add-on records these conditions so their results can be compared under matched budgets. The cited Ontology of the Alien study did not run the uncued direct explorer without population context; running that condition would extend the study rather than replicate a reported condition.

These records organize ideation. A world is a thought experiment, written and, by default, run as a model of its own laws; a transfer records an idea and its mapping; and a verified candidate has passed four separate checks, made in text. None of them is evidence that the idea works. One possible way to begin testing a candidate is to place it in a model of our own world. The model could be the world history (Section 7.2) as it stands, the world history opened deeper where the mechanism would act, or a new model built from scratch for the purpose. A transfer binds the mechanism's roles to that model's processes, where it could be followed era by era as a pattern over time (Section 6.7) and what it predicts compared with what the account records. That would test it against the account, not yet against the world. The add-on implements the world-population ontology and graph-directed world commissioning that Ontology of the Alien proposes, and directs commissions at the graph's empty kinds, for explorers as well as worlds. A second judge from another model family can answer each curator decision's comparison without seeing its reasons, which addresses that paper's single-coder limitation, and target-blind worlds can be exported as a library and imported into later searches. Whether governed world coverage yields more useful mechanisms than unguided generation remains that paper's open experiment; whether forcing empty kinds yields useful candidates, and whether a PURE verdict predicts which ideas survive testing, are open in the same way.

6.7 Patterns over time

Sema's vocabulary holds two kinds of pattern: ways of reasoning, such as how to carve a concept or weigh evidence, and patterns filed under a Dynamics category, which describe how a system behaves [9]. The Dynamics categories are in the released vocabulary from version 0.9.0 (emergent-wisdom/sema). The shared instructions use the second kind as models to test. When the modeler opens what drives a process or writes a law, it searches for a Dynamics pattern that fits, checks whether the pattern's conditions hold in the model and whether the recorded values show what it predicts, and where they do not, records whether the pattern does not apply or a value is wrong. The patterns used are cited by handle in the UnderstandingNodes, so a reader can resolve exactly which definition a decision rested on.

A pattern can also be followed over time, which shows what the method can do with a vocabulary of mechanisms. A card, a pattern's definition in Sema's vocabulary, states its mechanism once, without dates. The MalthusianTrap card states that a population grows until output per person falls back to where births and deaths balance, so that a better technique raises numbers rather than living standards; the card excludes populations whose births fall as income rises, and warns against applying it past the transition that the DemographicTransition card describes. Following such a pattern uses records already defined here. The pattern is a Concept whose grounding version is its card's content hash (Section 3.2). Its mechanism opens into processes a world already carries, here population, food, and output per person, whose recorded values show the pattern at work or not. At successive dates, an assessment Event about the relevant world fragment carries a fit Cut over matches, does not match, and remainder, so the pattern's hold on the system becomes a trajectory of its own: when it begins to apply, where its predictions hold, when its conditions fail, and which pattern takes over. A sequence of patterns over one world, each with its fit trajectory and the evidence for each handover, is a compact account of how that world changed: a history of mechanisms, coarser than the processes it rests on.

The world history (Section 7.2) has done this. Its current state holds 104 dated fit Cuts for eight Dynamics patterns, among them the Malthusian trap, the demographic transition, and the structural-demographic cycle, each fitted era by era to the world or to particular regions, and 57 more for the models of long-cycle theorists. For the world as a whole, the Malthusian card's match share stays between .60 and .85 from the spread of farming to 1820, falls to .30 for 1820 to 1870, and reaches .05 after 1950, while the demographic transition's rises from .30 for 1750 to 1870 to .85 for 1945 to 2000: the handover the two cards describe, judged against the account's own population and output values. These fit Cuts are the modeler's recorded judgments, not tests of the patterns. Each divides one unit of the modeler's assessment of fit. That is a judgment and no pool that can be counted, so by the number test of Section 2.1 these shares say less than the account's other Cuts, which divide hours, people, or money. A stricter fit would divide a pattern's stated conditions and predictions, as weighted checks (Section 3.5). The method also runs the other way. An agent can be asked to write a card for a concept node, or to turn a process it has modeled into a card that states its mechanism, so that what was found in one world receives a content-addressed definition that another construction can resolve and follow in its own world; sixteen of the world history's mechanisms became cards this way where Sema did not yet define them, from MalthusianTrap to CumulativeCulture. Whether a pattern's course predicts what happens elsewhere, and how patterns could be learned from such trajectories, belongs to Life Simulation [7]. A card can also change in the other sense: a

revised definition receives a new hash and a supersedes link, and the concept history of Section 3 records the revision without changing what earlier records cited.

7 Constructions

Three constructions show the method in use, and none of them validates it. The first *Book of Conditions* was built under the author's direction while the tool was being developed, and a second was generated by an agent from a premise and historical sources; both are twelve-chapter alternate-history novels. The world history applies the method to an account of the real world, built in successive agent passes. The section closes with the comparative evaluation that none of them yet satisfies.

The construction files this paper publishes were also read back against the grammar. A script in the repository, `scripts/six-forms.mjs`, writes a file's current model as records of the six forms and their typed links, with its reasons as UnderstandingNodes, and then rebuilds the tool's model from those records alone. For the second Book, for the world history at its state of 9 November, and for the mechanism search of Section 6.6, the rebuilt model is identical to the original. A few fields have no slot that the grammar names, among them an Event's statement of what it includes and a region given as free text; the script lists them. So the six forms can hold what these constructions recorded. That does not show that they are the best forms, and the engine still stores a process and an Event as two kinds of record; the repository's conformance note gives the correspondence.

7.1 *The Book of Conditions*: A Directed First Book and an Agent-Generated Second

The first Book is a completed twelve-chapter manuscript with an authored world, selected concept grounding, rooted UnderstandingNodes, and an ordered story route. The AI constructor built it under the author's direction at the beginning, while the tool itself was being developed; later revisions were made by language-model agents through the tool. The Book shows part of the method at work and is not a complete implementation of the protocol. It was built from the documentary record up to 14 August 1843 and one explicit counterfactual, and nothing from an earlier, discarded draft entered it: no evidence, accepted records, counts, or prose. This tests whether the grammar can hold one inspectable long story at a declared resolution.

The first Book continued through later revisions. The construction history, implementation boundary, and numerical portrait below report its initial construction, and Appendix D holds its import receipt; Section 7.1.5 and Appendix D.5 follow its subsequent revisions and describe its October edition. Section 7.1.6 reports the second Book, built independently and now supplied as the repository's worked example. Measurements and reviews remain attributed to the revisions they examined.

7.1.1 How the construction proceeded

The first Book's construction proceeded in six rounds.

1. Preparation fixed the historical cutoff, premise, vocabularies, access rules, and budget. The preparation notes stayed outside the model and had no authority in it.

2. The constructor registered roots, historical prefix, Things, lifecycles, relationship and presence Events, and thirty-four Concept bundles, retaining unresolved documentary detail.

3. Five isolated whole-history candidates were generated under a common schema. All were rejected: none could explain how continued negotiation became priced work and an accountable organization. None was admitted to History by a Direction roll.

4. A pragmatic successor was then written first as one paragraph and expanded successively to two, three, four, and five paragraphs. Opening the account exposed the missing commercial founder, who became Edward Halden. Local defects prompted repairs, not restarts unless the parent history failed.

5. Thirteen lifecycles and coarse life accounts were placed before three principals were opened into concurrent processes, local Cuts, and linked UnderstandingNodes. Events, relationships, ledgers, and one concept model elaborated or explicitly revised the macro account.

6. Twelve route units were selected and rendered. Later literary passes report updating Event and life descriptions before corresponding prose; changes without new world facts remained rendering decisions.

Saved checkpoints demonstrate coordinated revision, not every within-round editing order or the fully automated transaction protocol. The final world was not chosen by a preregistered Direction draw; missing logs cannot be recovered retrospectively.

7.1.2 A Book-specific profile

From Universe and History, the construction registered the needed people, institutions, places, documents, machine parts, and collective processes. Lifecycles, relations, presence, and historical data established the prefix. Character interiors sit beneath the person's lifecycle and inner-process root. Presence, communication, delivered evidence, and typed ancestry limit each actor projection to what the character could know at its cutoff.

For its principals the first Book uses a compact person profile of four views, IS (nine concurrent life processes), WHAT (wants under five concerns), HOW (how a problem is framed and decided), and FEELS (the emotion of a scene), with AtomicWant leaves beneath its concerns; its deeper numerical openings were specified but not exercised, and its selected scenes do not yet form a complete temporal partition. Appendix A.5 gives the profile in full, with the deeper person template archived with an earlier release, the authored scene, period, and shock profiles with the cross-scale records they produced, and the frozen flat concept registry of thirty-four grounding versions.

These categories are a first draft, not a finding. The intended practice is iterative: humans and AI construct characters, compare which decompositions produce coherent, distinctive behavior that depends on the numbers, and test which distinctions can be recovered from the resulting text. They revise the vocabulary between runs, including proposing categories that no current human vocabulary names. Revised concept groundings receive new content hashes; each run retains the versions it used, so earlier constructions remain readable. The grammar supplies a reusable language while the categories improve through this construction and testing cycle. The completed first Book is a starting construction, not a completed comparison of vocabularies.

Preparation specified a .10 elicitation grid with supported .05 refinement. The accepted pragmatic successor instead retains directly authored decimal shares, including .76 and .72 in scene-outlook Cuts. The retained record does not document a grid amendment or establish compliance with that rule. The retrospective import preserves those source weights and checks their sums; no transformation from the earlier grid is claimed. Other selected limits were no continuous extents, an estimate-depth cap of two, and full scene profiles only for the three principals.

7.1.3 Construction artifact, Rust host, and validation boundary

The first Book was imported retrospectively into the Rust host. Until then it had been written and checked outside the host, which did not enforce the protocol while the Book was built. Reloading recovered its model and native rendering reproduced its manuscript, but the import is a receipt, not a source-locked transaction for the original construction or an independent read-back record. Appendix D gives its contents, the authoring profiles the host compiled, and the synthetic conformance suite. Table 7 states the implementation boundary.

Layer	Exercised in the construction	Not established
Rust host	Profile compilation, immutable persistence, Cut normalization, component addresses, nearest-root ancestry, and source-bound document rendering	Complete mixture closure, cutoff-safe context assembly, independent read-back, or recovery of the original authoring transactions
External validator	Cut sums and remainders, selected recomposition fixtures, lifecycles, ancestry, presence, access, person-profile shape, and route exclusions	Historical or psychological validity, cross-surface atomicity of the final pragmatic successor, or independent prose recovery
Authored construction	Macro revision, Event and life descriptions, local Cut judgments, UnderstandingNodes, route selection, and prose	Automatic induction, forecast calibration, or an empirical advantage over unstructured writing

Table 7: What the first Book exercised and what remains outside the implementation claim.

7.1.4 Descriptive numerical portrait

An outlook Cut divides one unit of represented outlook toward fulfillment of active wants. Assurance means believing that a wanted endpoint remains available or will be attained; threat means believing it may be obstructed or lost. Babbage's May 1851 trial Cut assigns .76 to assurance, .18 to threat, and .06 to unresolved outlook. Lovelace's assured share is .70 and Halden's .72. These authored readings are not endpoint probabilities or emotional intensities. The

definition belongs to the first Book. The tool's number test later found it ambiguous between a felt chance and a share of wants; current guidance defines outlook as the share of the active wants, weighted as in the wants Cut, that the person sees as threatened, with a felt chance for each want recorded beside it.

Figures 4 and 5 make the long and short views visible without inventing an absolute psychological scale. The first plots compatible period Cuts as categorical intervals. The second plots compatible scene Cuts inside their shared simplex, the triangle of every possible division among assurance, threat, and unresolved outlook. Neither is a fitted temporal function: the marks display authored commitments at the resolutions actually opened by the construction. For example, Babbage's assured share is .79 at expansion, .28 at the returned table, and .51 at custody. The last point accompanies a changed interpretation of the work, not a return to expansion confidence.

The numerical portrait uses three categories. Cut weights are represented sibling-relative shares under one named question and remainder. Typed world data retain external units. Counts, density, total variation, means, and ranges are derived operational diagnostics rather than additional facts in the world. Table 16 in Appendix A.6 reports the first Book at its declared resolution; duplicate display tables are excluded from the Cut inventory.

A comparison of Cuts differs from a world constraint. In February 1852, each principal's Cut over the decision gives one weight to accepting four jobs and another to limiting the office to two. The local expansion margin $w(\text{accept four}) - w(\text{limit to two})$, the first weight minus the second, is +.30 for Babbage, -.50 for Lovelace, and +.20 for Halden. These are not votes: Halden alone holds contracting authority. The world-unit constraint is independent of those interpretations. In the first Book, four jobs require $4 \times 40 = 160$ checker-hours against the quarter's 80 protected hours, a twofold demand-to-capacity ratio and an 80-hour deficit. The later allocation of preparation and final checking is described in Appendix D.5.

Nested opening remains local. Early Babbage's threatened-fulfillment share .64 multiplied by autonomy-at-risk .34 gives .2176 of that period's full outlook unit because the second Cut explicitly opens the first; it is not a free fear score. The separate whole-history partition has six disjoint periods, from 18 August 1843 to 19 October 1871. They cover 10,289 days exactly, so their sequential shares sum to one by duration while saying nothing about narrative importance. For compatible Cut vectors, total variation is half the sum of absolute differences, including remainder: the minimum share that must move between categories to turn one allocation into the other. It ranges from .11 to .31 across eight adjacent slow-outlook transitions and from .38 to .60 across four selected scene transitions. This is a descriptive authored pattern, not a fitted shock law or population result. The numerical supplement archived with an earlier release retains the family, ledger, apparatus, and duration tables.

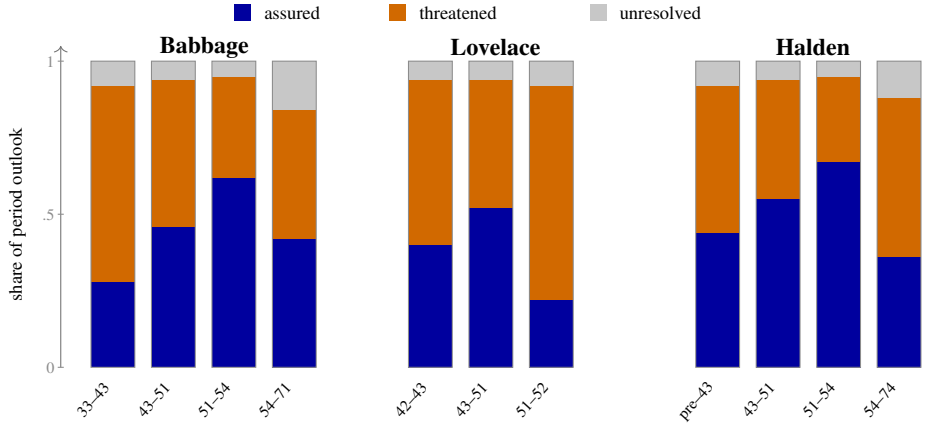


Figure 4: Slow fulfillment-outlook Cuts across the principals' opened periods. Each bar divides one period-specific unit of represented outlook toward active wants among assured fulfillment, threatened fulfillment, and unresolved remainder. Equal bar widths show authored period order only; they do not encode duration, intensity, probability, or measured psychology.

7.1.5 The first Book's continuations

The first Book continued through later revisions made through the tool, which deepened its lives, background processes, and telling and revised passages together with their supporting Events. The October edition also has a modelled fictional author, Nora Vale, whose life is a model of its own and no part of the story's world (Appendix D.5). A later model-only pass refined two concepts without changing the manuscript, which is not demonstrated improvement in the writing. The

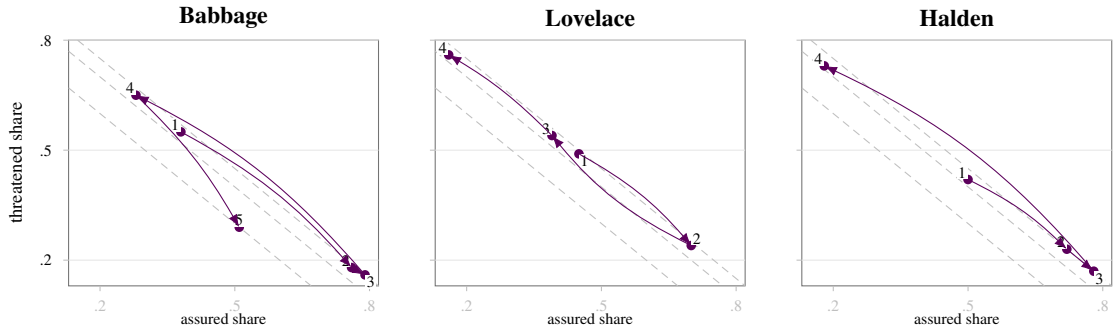


Figure 5: Selected scene-outlook Cuts as chronological paths through a magnified coordinate view of the same three-child simplex. Horizontal and vertical position give assured and threatened share; the three diagonal guides give unresolved shares of .05, .10, and .20. Babbage: negotiation, trial, expansion, returned table, custody; Lovelace: proposal, trial, qualification, succession request; Halden: estimate, trial, expansion, returned table. Curved arrows separate reversals and order the elicited scenes; they do not interpolate a continuous function or make the selected scenes an exhaustive time series.

first Book's readings were checked against the prose only in part: 50 of the 120 normalized Cuts in its 2 October model have a recorded textual basis, and the 6 October model adds 82 whole-life readings that were not reviewed against the prose. A narrow read-back with one isolated language-model reader per condition followed, which is not the controlled evaluation of Section 7.3. Appendix D.5 gives the October edition's history and inventory, its revised capacity account, and the readings.

7.1.6 The second Book, generated by an agent

On 11 October 2026, Claude Opus 5.5 constructed the second Book, a new *Book of Conditions*, through the Meaning Model tools. It began with a human-supplied task and a researched historical source package, not the first Book or its model. The task fixed a twelve-chapter novel and one counterfactual: Babbage continues negotiating Lovelace's 14 August 1843 offer to conduct the business of the Analytical Engine rather than refusing it. The agent constructed the subsequent world, lives, choices, scenes, and prose, then revised them through the tool. No human composed or copyedited its literary prose.

The run had three modeling and writing turns. The second was a request to continue; the third requested a deeper model and revisions informed by three fresh Claude Sonnet 5.5 editorial readers. These readers saw only the new manuscript, had no tools, and supplied their own feedback. The human managed the run and passed on those notes. A fourth turn corrected reviewer metadata without changing the manuscript or model. This is automatic generation and revision within a commissioned task, not an uninterrupted run without human direction or a demonstration of autonomous concept induction.

The tool ran on two development candidates, which Appendix D names. The installed copy withheld the first Book's modeling example, the Meaning Model and Life Simulation papers, and a number-test guide that cited an earlier verification run. Generic Babbage and Lovelace scaffold examples still supplied dates and structure; they contained no earlier plot, wants, process readings, or prose. This specifies the input boundary rather than claiming that the tool carried no prior Book-related material.

A limited comparison used four fresh Claude Sonnet 5.5 readers with no tools. Each received the new manuscript and the first Book's October edition, with identifying titles and afterwords removed, and the order was balanced. All four preferred the new manuscript overall. They compared 23,448 words with 13,337 words under that normalization. The writer did not receive their pairwise verdicts; it received separate editorial notes from readers of the new manuscript alone. The readers were model instances from the writer's own model family, not human readers. Unequal lengths, different construction and feedback histories, and the absence of a direct-prompt or prose-bible control prevent attribution of this preference to the Meaning Model or its numerical records.

The exported construction has 84 graph revisions and 85 model definitions. Its current world contains 571 Events, 202 processes, 540 Cuts, and 13 Concepts; six causal questions have 19 recorded assessments. Its twelve-chapter manuscript renders 23,441 whitespace-delimited words, including its title. These counts are inventory, not a quality score. All fourteen people meet the basic development check of the release that followed the run, but its coverage report remains incomplete: it asks for 6,817 scheduled readings, and 141 processes pass over their own Events without a value.

The run has no telling processes, executable laws, or Realizations. Its fictional narrator persona is recorded as a note, without a separate author life model.

The current `examples/book-of-conditions/` contains this new manuscript, its native construction, the task and historical sources, and a publication manifest identifying the retained revisions and any privacy projection. The example's counts and review status belong to this run; the older arithmetic and profile figures above, and the read-back results in Appendix D.5, remain evidence about the first Book. Remaining coverage gaps and authored historical interiors prevent the completed novel from establishing full protocol conformance, calibrated psychology, or historical truth.

The second Book demonstrates agent generation from a specified premise and sources. A stricter demonstration of autonomous construction would begin from scratch. A named model, working through a frozen release of the tool, receives one open prompt, such as to write any story it wishes, and works in a small fixed number of rounds. The prompt, the model, the release, and the complete construction are published with the work, so that anyone can check what that model built with that tool from that prompt. This demonstrates construction feasibility; the controlled comparisons of Section 7.3 test whether explicit structure improves coherence, inference, and transfer.

7.1.7 A separate executable refinement example

A separate, small reference bundle makes selected contracts executable without claiming that the Rust host already enforces them. The synthetic apparatus trial in `examples/refinement-trial/` uses all six world forms and a ten-minute local clock; it is inspired by the first Book's kind of scene, not a historical record or a retroactive validation of the first Book. Its frozen parent records configuration A and twelve matching cases. An opening introduces phases of two, five, and three minutes, with the four-child outlook Cuts below. Close maps fatigue and the fine remainder to the coarse remainder.

Phase	Duration share	Hopeful	Cautious	Fatigue	Remainder
Preparation	.20	.80	.10	.06	.04
Execution	.50	.60	.30	.06	.04
Comparison	.30	.40	.50	.06	.04

Table 8: Synthetic refinement: the declared projection and duration mixture recover the parent $(.58, .32, .10)$ for hopeful, cautious, and remainder.

The tests reject a changed coarse answer even when every child still sums to one, overlapping phases, incorrect duration shares, infeasible partial residuals, and stale candidates. An explicit configuration revision preserves its predecessor, changes the active head, and marks its dependent passage stale. A cutoff query exposes the configuration Ada observed, not concealed results on the observation's source Event. The accepted candidate includes a rationale UnderstandingNode and a passage DocumentNode. A fresh-context extraction recovers the stated configuration, case counts, phase durations, and observation time; it does not recover outlook weights absent from the passage. The saved extraction is a single model comparison, not an independent reader study or a general prose parser.

It is a bounded reference implementation, not a general graph validator. Appendix D gives how to run its tests and how its files are identified.

7.2 A World History Built in Agent Passes

A construction of world history tests the method on an account of the real world rather than an invented one. The account runs from about 300,000 years ago, as *Homo sapiens* emerges, to 2025, inside History, the lifecycle of the Universe, with the Earth and life on Earth as its setting. Every state of it reported here is published as an open repository (emergent-wisdom/world-history, CC BY 4.0), each as a construction history that imports into a fresh engine with the same hashes, and the current state, of 9 November 2026, can be explored at `https://meaningmodel.ai/world-history/`. Appendix D.9 gives the tool versions and identifies the last three states by their content hashes.

At its current state the account holds 861 processes, 2,102 Events, and 21 laws that set its numbers against one another. Of the processes, 278 place the countries, areas, and regions of the UN M49 standard on the globe, and the other 583 stand in one tree under History: the Earth, and humanity, divided into what matters to people, how they get it, their inner side, and their places by continent. The tree also holds money and finance, artificial intelligence from 1950, and the variables that theorists of long cycles follow; six representative lives and three cities followed in detail are filed under the populations they represent. The processes carry 41,195 recorded values at dates, 3,374 of them interpolated or carried forward by a script, and the laws derive 8,051 more from them. Table 10 counts this state beside the four

before it. Three of its rows carry most of what it shows: the values read from a source grow from none to 796; the values changed after they were first recorded stay under 4% of the total, while the median change falls from 15% to 3%; and from 28 October the laws derive values of their own.

Three series of the account show what such values look like (Table 9). Each value has a band and a tag, and the bands narrow and the tags change as the record improves.

Year	People, billions	Output per person, 2011 dollars	Trust, % of adults
10,000 BCE	0.0044 (0–0.012), sketch	450 (270–630), sketch	30 (10–60), sketch
1	0.232 (0.139–0.325), sketch	800 (480–1,120), sketch	21.7 (7–42), sketch
1000	0.295 (0.177–0.413), sketch	800 (480–1,120), sketch	21.2 (7–42), sketch
1500	0.461 (0.369–0.553), inferred	1,000 (800–1,200), inferred	21 (7–42), sketch
1820	1.06 (1.00–1.12), inferred	1,128 (1,000–1,260), source	21.8 (8–40), sketch
1900	1.65 (1.49–1.82), inferred	2,265 (2,050–2,480), source	23 (8–40), sketch
1950	2.493 (2.456–2.530), source	3,360 (3,150–3,570), source	25 (14–40), exploring
2000	6.172 (6.110–6.234), source	9,904 (9,410–10,400), source	30 (21–39), exploring
2025	8.232 (8.150–8.314), source	17,900 (17,000–18,800), inferred	25 (18–33), inferred

Table 9: Three series of the account at its head, as recorded: the value, its band, and its tag. A tag says what a value rests on: a source read during a pass, an inference, a value still being explored, or a sketch. The sourced values of people are from the UN's World Population Prospects 2024 and those of output per person from the Maddison Project Database 2023. Trust is the share of adults who answer that most people can be trusted, the question of the World Values Survey; before the surveys it is a guess of what that question would have found.

The structure was carved before any values were recorded. Its claim frame promises one place in the tree for every record the account will hold, across every human society in that span. The climb runs from Event through the history of a living, a social, and a cultural population to a human one. A human way of living divides into twelve dimensions, from provisioning and exchange to beliefs and norms and arts and play, each opened two levels further, with 25 measures frozen at the leaves. Technology failed independence and is kept as a view over the nodes for means; states, cities, writing, money, and coal exist only in some eras and are kinds under dimensions, not dimensions. Completeness was tested against the Outline of Cultural Materials: all but one of its 78 substantive categories found a home; the exception, behavior and personality, is kept as the frontier alternative of inner life. A routing sample of 142 claims was frozen before the carve but written by the modeler knowing the defects it was meant to expose, so it is weak evidence. A first attempt that chose eras and series by convention and recorded values first was superseded. The structure's own construction history is published as well.

The account of history as numbers was then built and extended in passes through the tools, most of them by a fresh Claude Opus 5.5 agent on a prompt that named the pass's focus. The first pass guessed 50 processes at 34 dates. Later passes filled the time schedule, the dates at each level of the tree of time at which every process is to be read, added Events and placed them in regions, weighed the causes of turning points, corrected errors found by blind readers, carved the divisions the tree was missing with their frames, climbs, verdicts, and syntheses, and recorded decisions to keep a process whole for now, each with its reason. By 18 October the account had 239 processes, and the coverage report found the work complete.

Two lines continue from that state. Four passes on 19 and 20 October turned to what the account had not yet followed: inner life, the history of the mechanisms the account explains with, representative lives, cities in detail, money and finance, and the variables that theorists of long cycles follow. They took the account to 539 processes, and their graph reached a limit of the engine they ran on. That branch is published beside the current state, and Appendix D.9 describes it and the tests made on it. The current state comes by the other line. The passes of 25 to 28 October returned to the 18 October state to place the account's Events in countries and to set its numbers against one another by laws (Section 7.2.1). Six passes on 8 and 9 November, on the released tool, then took up the subjects of the October branch again (Section 7.2.2). They read that branch as earlier findings: they checked what it had found against sources, cited what they kept by file and model hash, and copied no records. Appendix D.9 gives the passes in order, with their dates, tool versions, and counts where these were recorded.

Every record names what kind of claim it is: checked against a source, inferred, exploring, or sketched. Most values are recalled and tagged inferred or sketch; nothing is tagged as sourced until a source has been read, and a source named beside a recalled figure is where it is to be checked. Figures recalled from one training corpus share one foundation, so their agreement is not independent support. In the branch of 19 and 20 October, some values for money and finance came from a second construction that is not published, so a reader cannot check them against the records they came from; the current state takes none of them. The carves are tagged exploring: a carve that passes its tests is the best

	15 October	18 October	20 October	28 October	9 November
Model definitions	230	539	843	567	747
Graph revisions	257	520	635	557	787
Processes	152	239	539	543	861
Dated values	17,852	30,076	44,569	31,107	41,195
tagged sketch	10,202	18,751	26,513	19,465	24,049
tagged inferred	7,518	11,188	15,023	11,463	14,277
tagged exploring	132	137	2,874	137	2,073
tagged sourced	0	0	159	42	796
with a band	15,859	27,447	35,054	28,466	37,433
Median band, inferred	20%	25%	27%	26%	24%
Median band, sketch	83%	105%	102%	104%	102%
Values changed after first recorded	188	817	1,577	855	1,325
median change	15%	5%	3%	5%	3%
Values derived by laws	0	0	0	6,068	8,051
Events	1,143	1,152	2,240	1,400	2,102
Cause questions	56	88	91	88	91
Weighed causes	188	296	308	296	305
Judgments of importance	3,309	3,343	3,343	3,343	3,417
Processes declaring a sum	0	42	45	42	45

Table 10: The world-history construction at the five states reported here, counted at each state's head model. A band's width is taken as a share of its value, and the median is over the values with that tag. The 20 October state is a branch of the 18 October state (Appendix D.9); its Understanding Graph is bound to model revision 631. The states of 28 October and 9 November continue the 18 October state in one line and do not follow from the 20 October one. Their rows for dated values count only the recorded values; the values their laws derive have a row of their own.

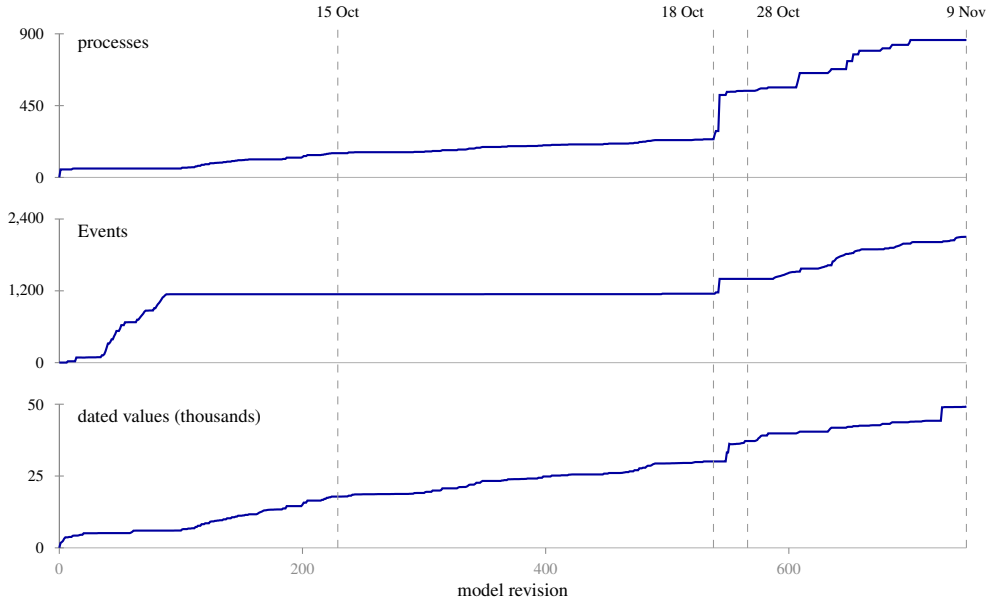


Figure 6: Growth of the world-history construction along the 747 model definitions from its first pass to its current head: processes, Events and dated values at each revision, the values that laws derive included. Dashed lines mark the states of 15, 18 and 28 October and the head of 9 November. The tree grows in steps, as divisions are carved and their parts opened. The step in processes and Events before 28 October is the places set on the globe; most other Events came early, before the tree was divided further, and again in November with the cities and lives. The step in dated values at revision 726 is the scripted revision that interpolated values and derived the laws' series again.

current hypothesis until the values test it. The construction therefore demonstrates that agents can build and extend a large account of the real world at progressive resolution, preserving identities and recording every change and its reasons, from one tool version to the next. It does not establish the account's historical accuracy, beyond the corrections recorded, or that its categories are the best available.

The values can also be followed as they were revised. Along the current line 1,325 of the 41,195 recorded values were changed after they were first recorded, in 1,385 changes, 102 of them made by the script of revision 746 (Section 7.2.2). In 1,276 changes the earlier value had a band, and in 1,073 of those, 84%, the new value lay inside it. A value read from a source replaced an inferred one 171 times, in 45 processes, all at dates from 1800 on and 154 of them from 1950 on: people and life expectancy from the UN's World Population Prospects 2024, output per person from the Maddison Project, carbon dioxide from the record at Mauna Loa. The median change was 0.6%. Of the 157 earlier values that had a band, the sourced figure lay inside it for 151. Among the six that did not are the world's people in 2014, 7.30 billion against 7.382, and the share of people online in 2025, 68% against 74%. These are modern, well-measured series that a language model is likely to have met in training, and the account does not record whether an earlier value was itself recalled or read. The figures show that the inferred values for such series were close and their bands wide enough. They say nothing of the sketches of earlier eras, which no source in the account tests.

7.2.1 Places and laws: the passes of 25 to 28 October

The passes of 25 to 28 October returned to the 18 October state, rather than extending the 20 October branch. They asked where things happened and whether the account's numbers agree with one another. A script recorded the places and their representative points; agents placed Events in countries where the account supported that placement.

Twenty-one laws express relations among the processes: identities such as world output as people times output per person, bridges from fertility and life expectancy to births and deaths, and bounds such as the largest town being no larger than all town dwellers. Their derived values are held apart from the recorded ones. A script compares them using duration-weighted intervals; this pass is a check of recorded numbers, not a dynamic world run. Two laws that would change recorded growth and atmospheric carbon are disabled while the gaps remain unresolved.

Where a law's value and the recorded one disagree by more than their bands allow, the laws' report counts a finding. A fresh agent checked ten apparent findings against literature it could read, recording how much of each source it reached. Its assessment was two known effects, seven errors in the model, and one agreement that held by construction; none was a discovery. The largest discrepancy, births minus deaths implying more growth than the population counts before 1950, leads back to the published forager population paradox. Normal-year vital rates need not describe long-run growth through crises; Gurven and Davison study catastrophes together with changes in fertility and mortality as an explanation [22]. That literature supplies a hypothesis for checking the account, not a new finding about humanity.

The corrections changed pandemic timing and plague waves, fossil-fuel carbon and energy, and land-use carbon, and added a series for crises the account does not date. Two further issues were left open: a surplus bound compares shares with different denominators, and a lower estimate for Çatalhöyük had not been checked against the underlying study. Revision 566, where these passes end, also repairs a correction that had broken the output identity at the years 549 and 610: output is set to people times output per person at both dates, with the continents adjusted in proportion. The final report, rather than the earlier correction ledger, is the record of those results.

After correction, the central pre-1950 growth gap remains about 0.4–0.6 percentage points a year; adding crises did not resolve it. The births factor was calibrated on the model's own 1950–2025 values, and its use in earlier eras remains a hypothesis. The broad bands overlap in every interval but one, which lies after 1950 (1961.9–1962), so the central gap alone does not establish a contradiction beyond uncertainty. Nor do agreements obtained by deriving one series from another test historical accuracy: output set to its defining product at 549 and 610 and an airborne fraction computed from the recorded carbon rise are examples. This continuation shows that explicit relations can expose errors and guide repairs in an account. It supplies neither a historical discovery nor independent validation of the corrected account. Its literature checks and judgments are model work.

7.2.2 The passes of 8 and 9 November

Six passes on 8 and 9 November continued the 28 October state on the released Meaning Model 0.18.1, each by a fresh Claude Opus 5.5 agent. The first opened the inner side of life: the emotional climate, tastes in the arts, trust, how far ahead people plan, hope for the children, and the moral circle. It also fitted the mechanisms the account explains with, era by era (Section 6.7). The second built six representative lives, from a forager on the Cape about 18,000 BCE to a garment worker in Dhaka today, and three cities: Paris month by month from 1789 to 1794, Florence from 1338 to 1350, and Tenochtitlan from 1517 to 1521, with the minds of their groups, the decisions of their people, and the news that reached them. The third followed money: its forms, prices, interest, and debt, with crypto inside money

from 2008, and what the theorists of long cycles measure, each theory fitted as a mechanism. It took nothing from the unpublished construction that the October branch had drawn on, and read public sources instead. The fourth followed artificial intelligence from 1950 to 2025, and checked the account's anchors against sources such as the UN's World Population Prospects 2024 and the Maddison Project. The fifth brought values to 2025 from published data, read the whole tree again, and filed the lives and cities under the populations they represent. The sixth filled the time schedule: a measure that began at a known date, such as a survey's first round, became a Thing of its own from that date, so that no values were invented for the years before it. Over the six passes the account went from 543 processes and 1,400 Events to 861 and 2,102, and the values tagged as read from a source from 42 to 796.

Two revisions were made by script. After the fifth pass, revision 726 derived the laws' series again and gave the processes that passed over their own Events values interpolated between their neighbours: 3,374 recorded values carry a note saying that a script interpolated them or carried them forward. Revision 746 repaired an error of the passes themselves, which the check of the laws had shown after the fourth pass. They had matched output per person to its sources at 1820, 1900, 1920, and the war years, and revised 2020 to 2025, without changing world output, so world output no longer equalled people times output per person at 17 dates, by up to 9.4% in 1941. It is the kind of error that revision 566 had repaired at two dates. The script set world output to that product at the 17 dates and scaled the continents to add to it, and each changed point keeps its note, with the reason added.

At the head, the coverage report finds the work complete: every level of the time schedule is filled, no process passes over its own Events, and no turn in the values is left without a recorded reason. The repository's check finds none of the 19 enabled laws failing on the head's values; the two laws on growth and atmospheric carbon remain disabled, as on 28 October, and the output identity's largest gap is 2.01%, at 200 BCE, well inside the bands of the values there. Forty-five parents declare that their parts sum or average to them. All 5,386 dated comparisons fall inside the parent's band, 95.4% of them agree exactly, and no revision of this line holds a sum its parts contradict. As on the October branch (Appendix D.9), wholes and parts were recorded to agree, so the check shows that they agree, not that either is right. Most values are still recalled, and no test against reference data has been made: the construction does not record, value by value, which values the model recalled, so agreement could not separate the model's own knowledge from what it had read. These passes show that an account can be taken up again on a released tool, set against its own laws, and repaired where they fail. They do not establish its accuracy.

7.3 Prospective comparative evaluation

Neither Book construction satisfies the still-prospective controlled comparison. A future scored construction must preserve source-locked accepted commits, direction Cuts and seeds where selection is claimed, DocumentNodes linked to those commits, independent read-back proposals, and the final model under one preregistered protocol. Structural conformance then asks whether that construction completes without undeclared vocabulary changes, silent revision, unresolved references, or violations of its Cut, ancestry, and access contracts.

The central idea needs its own test. A concept discovered during construction, such as the Introduction's hypothetical concept of procedural betrayal, is frozen with its grounding version and applied to held-out cases that did not produce it. Explanations and predictions made with it are compared with those of the same model working from the same evidence without the explicit concept, at matched budgets. A second comparison asks whether explicit anchors improve later inference, by another agent or in a later session, over the same knowledge held only implicitly. If discovered concepts do not transfer, or explicit anchors add nothing beyond what a capable model does implicitly, that is a negative result for the paper's central idea. Whether training on constructions makes better models, as distinct from whether construction helps the model that builds them, is Life Simulation's test, through its five-level ablation from implicit to explicit structure and its repair-learning pilot [7].

All scored Cuts use a frozen elicitation procedure (Appendix A.4). Before any comparison, the elicitation itself is tested for reliability. Scored Cuts are elicited again under the same frozen procedure, several times by the model that wrote them and by models of other families, each in a fresh context that withholds the Cut, its provenance, and any description written only to carry it. Spread within a model, agreement across items and across model families, and distance from the recorded weights are reported by Cut type. The spread sets the finest grid a weight may claim: weights that scatter by more than half a grid step on repeated elicitation claim a precision their elicitation does not support. Stable weights are not thereby correct, since models may share a bias; read-back and numerical dependence test what the weights mean.

The literary comparison is preregistered before rendering. At equal information and output budgets, the structured condition is compared with a direct-prompt baseline and a prose-bible condition: a prose reference of the same world facts without record operations. An optional serialization-only ablation supplies the same records as text but disables construction and validation operations, testing their contribution beyond the information supplied. Independent auditors

count contradictions against the same released facts, while blind readers compare long-range coherence and literary quality.

Read-back targets are fixed before rendering: which identities, Events, chronology, attributed beliefs, access distinctions, and qualitative or numerical comparisons must the passage convey? Other content may remain latent or explicitly ambiguous. Fresh readers receive any preceding text, the passage, and neutral extraction questions that allow an unknown answer, not source records or expected answers. Their recovered claims are aligned with source identities and authority paths. A control group receives the same preceding text and questions, with the target passage withheld. Absolute read-back accuracy and improvement over this control are reported separately: faithfully repeating a fact the reader already has, from earlier text or from prior knowledge such as the historical record of the Book's people, need not add new recoverable information. At least one dialogue or shared Event is rendered from two different character cutoffs so that the test can ask whether prose preserves both perspectives rather than averaging them. Book-specific appraisal and strategy checks ask whether an active threat belief has a compatible fear appraisal and whether a chosen strategy is supported by the operative wants, beliefs, and constraints; they test the frozen profile, not universal psychology.

Numerical dependence is a separate test. A preregistered intervention varies compatible Cut weights, rechecks affected parents, and holds independent commitments and rendering budgets fixed. It tests whether prose changes as predicted, without requiring a reader to recover every latent weight. The intervention is applied before affected Event descriptions and other derived context are generated, and dependent passages are then regenerated under the same policy. Numerical influence may pass through those descriptions. Removing raw numbers only after that influence has entered the context tests continued access to numbers, not their upstream contribution [7].

Two further predictions follow from the design. Under matched edit scope, revisions guided by declared dependencies should introduce fewer new contradictions than unguided revisions of the same manuscript. A ranking of decision and state jumps, frozen before reading, should predict where independent readers locate the key turns; a highly ranked turn rendered only in retrospect should predict a report that it is told rather than shown.

Cut density, remainders, direction calibration, rejection and revision rates, prose length, model calls, tokens, latency, human adjudication time, storage, and maintenance work are reported so that lower contradiction counts cannot be bought simply by saying less or spending an unreported budget.

No outcome is assumed, and negative findings are required rather than folded into a revised success criterion. If a prose bible performs as well, if readers prefer the baselines, if perspective separation or appraisal checks fail, if read-back cannot recover the records, or if explicit maintenance cost overwhelms the gain, the comparison has located a limit of the tested construction method. Diagnosis must distinguish implementation or profile failures from the representation's three stated failure modes: an unrepresentable required distinction, loss of required information under composition, or explicit articulation whose cost exceeds its value. Failure to achieve the declared practical benefit under fair conditions remains a negative result; a possible better implementation does not establish that benefit. If the structured condition succeeds, it establishes an incremental advantage beyond the partial descriptive witness at one resolution; it does not establish a universal ontology or a theory of temporal dynamics.

8 What Changes When Works Are Constructed

If books, films, and game worlds come to be constructed this way, a work is no longer its text alone. A book becomes one route through a modeled world, whose reading order passes through Events that exist whether or not a scene shows them; the storytelling guidance already treats a story as a small window onto a world modeled far beyond it. A film, a game, a sequel, a translation, or an adaptation becomes another route through the same world, and continuity between them becomes partly checkable, where depiction and dependencies are declared, rather than something curated by memory. Readers and players could open the world behind a work: read a secondary person's life from its beginning, follow a process the story only glimpsed, or ask what a character knew on a given day, with each answer attributed to whoever holds it.

Game worlds change most directly. A world could stay coarse but consequential everywhere and open detail where a player goes, as the harvest example in the Introduction describes. Its people would have modeled lives whose decisions consult their trajectories, and shared facts would stay the same for every player. Procedural generation and level-of-detail simulation already produce and run large worlds [1, 2]; what could change is coherence at that scale, because progressive resolution and a shared construction medium let detail be constructed when it matters and checked against what is already committed.

Authorship can shift toward direction. The human who commissions a work chooses what it should convey, the strategy for conveying it, and what makes it worth reading, and judges the result; the language model can construct the world and its telling under those commitments. Because the construction record keeps choices, reasons, and reviews with the

records they concern, who decided what becomes inspectable where it is recorded, and a revision can be traced to the commitment that caused it. Criticism could gain further evidence: the ideas the work was built to convey, the strategy recorded for each, and blind read-backs that test which of them a reader recovers.

The worlds also become learning material. Construction histories pair a situation with the decision taken in it, to refine, repair, or keep a process whole for now, the kind of example from which a system might learn to think in explicit, consistent worlds, as Life Simulation proposes (Section 11). Modeling fictional people from where their lives begin may also be practice for the harder, evidence-bound task of understanding real people.

The changes carry risks. One grammar and one model's implicit priors could narrow the range of worlds, the Median Trap that the world-diversity search of Ontology of the Alien is designed to counter [18]. Comparative numbers can lend false precision to a poor question, and a highly coherent invented world can persuade more than it should. Real people need protection: the storytelling guidance asks that living people and real organizations not take part as characters, and the director review checks it. Some art depends on what is left unsaid. The model can record a motive fully while the telling withholds it, which is a disclosure choice rather than an unresolved Cut; where the author leaves something genuinely open, that openness can be recorded as such. None of these consequences is demonstrated here: *The Book of Conditions* shows only that a work can be built this way. They describe what the method is for.

9 Related Work

Prototype and family-resemblance accounts reject the assumption that every category is defined by one set of necessary and sufficient surface features [23]. The Meaning Model likewise permits graded fit Cuts, exemplars, boundary cases, and plural models. Only its deepest grounding option uses an inspectable process-based concept model.

Structure-mapping theory treats comparison as an alignment between relational systems [24]. Realization borrows that discipline: role alignment is recorded before component fit is aggregated. The present proposal is narrower computationally; it does not claim a complete cognitive theory of analogy.

Normalized attention has precedents: Nosofsky's exemplar model allocates nonnegative weights summing to one across perceptual dimensions for categorization [25]. The proposed methodological use here is recursively refinable, question-relative conceptual allocations in evolving worlds, coordinated with construction and revision.

Dempster–Shafer theory permits basic probability mass to remain assigned to sets of possibilities, including the whole frame, without distributing it among individual possibilities [26]. The Meaning Model's remainder instead retains the unresolved share of a question-local unit, such as attention or duration; it is not generally evidential mass, and the grammar does not adopt Dempster's rule for combining evidence.

Frame semantics and script accounts organize meaning around structured roles and characteristic situations or event sequences [27, 28]. Process-based concept models can make those roles and sequences refinable through the same Thing–Event grammar used for a token world. Conceptual spaces instead emphasize geometric regions and similarity [29]; such a space can be attached as an external domain model, while Realization links and fit Cuts preserve the grounding, alignment, perspective path, and local alternatives used.

Process ontology and event mereology motivate taking temporal entities and their parts seriously [30, 31]. The continuant–occurrent and SNAP–SPAN tradition likewise distinguishes entities that remain identifiable through change from processes extended in time, and explicitly relates a continuant to the occurrent that is its life [32]. The Meaning Model's Thing–lifecycle pairing is in that lineage, while its local Cuts and plural concept grounding address a different construction problem. Object-Process Analysis places objects and processes in one recursively scalable systems account [33]. Descriptions and Situations distinguishes a reusable description from a situation satisfying it [34]. The Meaning Model's distinctive choice is the local one-unit Cut with explicit remainder, ancestry-derived perspective, plural concept grounding, and conservation under progressive refinement.

Narrative systems already maintain worlds. Inform tracks rooms, objects, containers, and interaction scope [35]; bAbI tests synthetic location and relation updates [36]; ProPara annotates entity existence and location changes [37]. EvolvingWorld couples persistent character and global state in literary role-play [38]. Narrative World Model extracts typed temporal memory from accepted prose and invalidates downstream records after edits; its reported evaluation concerns narrative question-answering, with generation evaluation left for future work [39].

Hierarchical and model-assisted authorship also have direct precedents. DOC expands hierarchical outlines and controls passages against them [40]. DOME interleaves detailed planning and writing with temporal knowledge-graph memory [41]. Generative Agents combines recorded experiences, reflections, and plans to produce behavior [42]. The Meaning Model therefore does not claim to introduce planning, graph-backed storytelling, or reflective characters. Among the systems compared in Table 11, its distinction is recursive inference anchored in explicit world structure under a shared

preservation contract: a coarse world remains consequential before its details exist; an opening must satisfy declared fine-to-coarse checks or trigger explicit revision. Local comparative person Cuts and versioned concept groundings participate in that same construction rather than forming unrelated annotations.

The cross-scale pattern has important precedents. Holling's panarchy account links faster, smaller adaptive cycles with slower, larger ones: local disturbances may propagate upward, while accumulated structures and resources shape reorganization below. He also describes departures from the cycle [43]. This is a close conceptual relative, not evidence that every represented process must complete the same stages.

Dalio's historical practitioner account connects expectation, surprise, and revision across debt, conflict, natural disturbances, and human adaptation [44]; neither its causal interpretation nor a fixed historical cycle becomes a law of the grammar. The useful intuition is that earlier historical cases can help revise present expectations after an unexpected outcome.

At everyday timescales, event-segmentation theory connects perceived boundaries with prediction error and considers simultaneous finer and coarser event models [14]. These accounts supply related hypotheses at different scales, not a single established mechanism from perception to history. The Meaning Model makes their processes, expectations, responses, and cross-scale explanations addressable for comparison. Change-point, jump–diffusion, mean-reverting, self-exciting, and hybrid models are downstream candidates for Life Simulation rather than alternative record grammars [45, 46, 47, 48, 49].

Partially observable Markov decision processes give an agent a belief state, a probability distribution over the states of its world that it updates from its last action and observation, and choose actions as a function of that belief rather than of the world's state [50]. Interactive POMDPs add beliefs over models of the other agents [51]. Active inference casts perception, planning, and action as probabilistic inference under a single imperative, to minimize free energy [52]. Its deep temporal models accumulate evidence over nested time scales, and the expected free energy that guides planning includes an epistemic value, the information a policy is expected to gain [53]. The Meaning Model shares the separation of hidden state from observation: an actor acts from what it has observed and been told, an observer's estimate of it is the observer's own record, and the profile bounds how deeply such estimates nest (Section 4.1). Its progressive resolution nests time scales, and the reports that name which process or period to open next play a part like epistemic value. It differs in kind. Its world is an explicit, attributed construction that a language model writes from what it has learned and revises under declared contracts, not a probability density inferred under one objective. Its numbers are values in their own units or comparative shares under a declared question, each tagged for how it is known. Probabilities appear where they are declared, as in a forecast, which is then scored (Section 7.3). Learning predictive models from constructions, where the two programs would meet, belongs to Life Simulation (Section 11).

Semantic Spacetime treats promises, agents, names, and scale as foundational to functional knowledge representation [54, 55, 56]. The present grammar combines multi-scale semantic graphs with explicit model commits and question-local composition.

Making implicit knowledge explicit also has a long history. Knowledge engineering elicited experts' knowledge into the rules of expert systems [4], and Cyc hand-encoded commonsense concepts and axioms from 1984 on [57]. Both depended on people to articulate what they knew, and tacit knowledge resists that [3]. Here the language model is asked to articulate its own understanding, recursively and at the scale of a long work, under a grammar that quantifies meaning comparatively and preserves commitments across openings, and the result is a world and its telling built for one inquiry or work. The proposal's two architectural contributions, progressive resolution and a shared construction medium (Section 1), form one complete operation. The comparison below locates each part among its precedents; to our knowledge, no earlier system specifies this complete operation for a language model constructing and revising its own explicit world model.

A revised success criterion, for example, can require rechecking a trial's assessment, the rationale for expansion, and the passages expressing both. The relevant comparison unit is that complete operation: what may remain unconstructed, which commitments an opening protects, and how a revision affects the connected account. An antecedent for graphs, normalized allocations, or revision alone does not settle whether this method has already been specified; neither does their combination establish priority by itself. Table 11 therefore compares the contracts of that operation across the closest systems, as their own papers describe them. Each supplies part of it: outlines and plans commit coarse structure before detail, a controller steers passages toward outline items, a verifier checks candidate chapters against stored state, republishing invalidates dependent memory, reflection builds higher-level thoughts on stored records, and an open schema infers categories for each work. Generative Agents comes nearest to the discovery loop: its reflection prompt asks for high-level insights together with the memories that support them [42]. Those insights remain text memories. The Meaning Model asks the constructor to make discovered concepts and processes trackable in the grammar that holds their instances, so that they can be followed over time, tested elsewhere, and revised under the preservation

contract. Narrative World Model comes closest to the whole operation, but it records memory only from finalized prose, in a specified set of narrative fields. None of those compared specifies the complete operation.

System	Coarse commitments before detail	Detail checked against them	A revision rechecks its dependents	Explicit records support further inference	Accounts kept per holder	Categories constructed and revised
DOC [40]	Outline expanded before drafting (§3.2)	Controller steers passages toward outline items (§3.3)	Not described	Inferred character facts reused in later passages (§3.2)	Not described	Not described
DOME [41]	Rough outline, then detailed outlines (§4.1)	Retrieval reduces conflicts; conflicts measured after generation (§4.2–4.3)	Later outlines adjusted forward (§4.1)	Extracted quadruples retrieved for planning (§4.2)	Not described	Not described
Generative Agents [42]	Day plan decomposed into finer actions (§4.3)	Reactions regenerate the remaining plan (§4.3)	Not described for memories or reflections	Reflection trees over observations and reflections (§4.2)	Each agent's own memory stream (§4.1)	Not described
Narrative World Model [39]	No: memory records finalized prose (§3.1)	Candidate chapters checked for unsupported state changes (§3.4)	Republishing invalidates dependent memory (§3.1)	Recursive question answering over stored state (§3.4)	Character knowledge and unknowns (§3.2)	Specified narrative fields; revision not described (§3.2)
EvolvingWorld [38]	Not described	Not described	World state updated after each interaction (§3.2)	Stored states inform scene selection and actions (§3.2)	No: one shared objective state (Limitations)	Dimensions inferred for each book (§1)
Inform [35]	No: world declared in the source	Actions checked against world rules (§24)	Not described	Not described	Scope limits what the player addresses (§24)	Not described
Meaning Model	Protected coarse answers	Close, or explicit revision	Declared dependents across world, concepts, reasons, and prose	Records and concept models anchor further opening	Perspective by containment	Versioned groundings with supersession

Table 11: Contracts of the complete revision operation in the closest systems, as described in their own papers. "Not described" means the source does not describe the contract, not that the system could not support it. The Meaning Model row states specified contracts; Table 7 gives what the prototype implements.

The two Book constructions and the world history partially demonstrate feasibility, while controlled comparative benefit, learning gains, and computational savings remain untested.

10 Limitations and Research Boundary

Normalization enforces accounting consistency, not truth. A Cut can ask the wrong question, choose misleading children, or encode an LLM's bias with false precision. A concept hierarchy, exemplar set, or concept model can privilege one culture or perspective. Role alignment can manufacture apparent similarity. The recorded ancestry, remainder, evidence, provenance, and grounding version make these judgments inspectable rather than eliminating their biases.

The first Book's concept vocabulary and exemplar anchors are not yet validated, nor are its person, emotion, motivation, and strategy vocabularies shown to be complete, culturally portable, or optimally independent; they are expected to change through the iterative construction and testing practice described above. Definitions can drift, rankings of alternatives under one Cut question can be inconsistent, and sparse exemplars can hide boundary cases. Concept induction can overfit examples or proliferate names until comparison becomes meaningless. Event and period boundaries are likewise model commitments rather than naturally labeled facts. The grammar makes these choices inspectable; it does not turn them into a psychology or a dynamical law.

A recorded conceptual decomposition makes each step of the method inspectable, but it does not verify that the step was carried out as stated: the record cannot show whether a routing sample was chosen before the carve or whether verdicts were reached independently. Agent runs on earlier releases may also have received only part of the operational guidance, because Claude Code keeps the first 2,048 characters of a tool's description and saves long results to files the agent may not open; Appendix D.7 names the release that fitted the guidance to these limits. What those agents were told is therefore what the tool sent, not necessarily what they read.

The grammar does not replace physics, economics, biology, geometry, or any field with units and specialized laws.

Root and surface separation preserve attributions but do not reveal hidden minds. Historical character interiors remain authored interpretations. Human explanations inherit hindsight, narrative compression, and over-attribution to character.

A model trained on them may become better at predicting people or merely better at producing plausible stories about them.

Finally, recursive refinement is not empirical utility. A strong language model may already maintain equivalent structure implicitly. The Book comparison must establish whether explicit records improve coherence and read-back enough to justify their construction and maintenance cost.

11 Relation to Life Simulation and the Wider Program

The papers of the wider program share one hypothesis, which the Introduction states for a single book: much of human thinking was never written down. Writers left it implicit because they assumed a reader who could supply it, and for a model trained on text, the gaps we left for each other are gaps in the model. The papers explore different ways to identify and fill such gaps. The Substrate-Translated Language Model (STLM) takes the visual part: people often think about a complex system through images that hold much of how it works at once, and STLM gives a model such a medium as a compulsory bottleneck that resembles an intuitive chain of thought [58]. Imagining the Corpus continues the visual part at corpus scale: it keeps the direct language route and compiles whole documents into continuous explanatory video [59]. The Meaning Model takes a non-visual route: an account of a world as linked processes that change over time, in one general-purpose grammar, can make explicit what a depicted scene does not by itself specify, such as an estimate's uncertainty and what it rests on. Which of these routes a model needs, and how they might combine, remains open.

The wider program aims at one Meaning Model, not one for each work. Life Simulation proposes a single connected account that spans world history and the texts written in it [7]. The things this paper describes one at a time would be parts of that account: the processes of the world, of which the construction in Section 7.2 is a first piece; the world of each book; each person's life; and the work of solving a problem or pursuing a line of research. The concepts these parts use would have their models in the same account (Section 3), so that what a concept is can be read as processes over time. The understanding surface would run through it at several levels: the modeler's questions and revisions, with its own attention and concern followed as processes, and each character's questions and revisions under that character's own root (Section 4). Fictional worlds, disputed accounts, and individual perspectives would keep their separate contexts inside the one account. The grammar needs no new form for any of this. It needs the tool support described below, which is not built, and the three constructions of Section 7 are separate accounts.

A connected Meaning Model of actual world history and geography can deepen as evidence arrives. Sources share supported identities while retaining separate claims, evidence, and disagreements. Regions and centuries can open into institutions and lives without equal detail everywhere, and new evidence can revise earlier accounts, not merely add detail, with event and revision times kept distinct. Acceptance does not establish historical truth; fictional worlds retain separate contexts. Meaning Model supplies the representation and revision contract; Life Simulation develops corpus-wide construction, interpretation, and training. At corpus scale, the tool would need incremental ingestion, selective retrieval, and dependency-aware revision over one logically connected address space, potentially distributed across storage systems. New material would link to supported identities and source records; refinement would first open the processes an inquiry reaches, while the account continues to widen and deepen. Versioned dependencies would identify derived views affected by a correction for explicit revision or stale marking, while preserving their source records and prior versions. Access rules, perspective boundaries, temporal cutoffs, and the selected profile's commitment and recomposition contracts must hold across storage boundaries. These are proposed extensions of the tool, not additional grammar forms or an implemented corpus-scale system.

Life Simulation uses completed or ongoing worlds to study the generation, inference, and learning of process histories. Continuations may come from an LLM, an authored account, a simulator, or a fitted function; no compact dynamical law is required. Comparing continuous, discrete, or hybrid laws is one research task within that broader method. Its process-aware models are intended for games, scientific or social simulation, exploring unfamiliar causal mechanisms through Ontology of the Alien, organizing problem-solving capabilities through Fractal Intelligence, and alignment-facing curricula [7].

This is not a one-way handoff. Life Simulation can propose a continuation or a predictor; Meaning Model construction rules still govern its representation, acceptance, and effects on dependent records and text. Meaning Model read-back checks what a passage says about its source world. Life Simulation's inverse task asks what hidden processes observations support, and whether inferred accounts predict later evidence or transfer beyond their construction setting. Together the two proposals support a feedback loop: construct a world, render observations, infer an account from the available evidence, compare its predictions with later evidence, and revise the account before constructing further. Correction may change an Event, a decomposition, or a Concept's grounding, rather than merely add detail. The Meaning Model makes those changes addressable and preserves their history; Life Simulation specifies learning

objectives and independent tests. Agreement with a construction's own rendering alone does not validate its account of reality.

The proposal also combines explicit character-state modeling before prose with recursive conceptual decomposition from Fractal Intelligence [8]. Fractal Intelligence decomposes capabilities and problems: child Solvers inherit constraints, their combined result is evaluated against the caller's acceptance criteria, and a failure may challenge the parent's framing. The Meaning Model applies a related responsibility to world accounts: finer records must reconstruct protected coarse answers or revise the commitments explicitly. Both permit recursive decomposition without a fixed depth limit, while each actual construction or execution remains finite.

The inner-life ambition connects three levels of this loop. *In the story*, a character holds a model of the world and estimates of other characters, and acts from that perspective. *In construction*, the modeler builds and revises those character models, recording its own estimates and reasons separately. *In the long run*, a trained system could use the same distinctions to estimate real people's changing perspectives from available evidence. The first two levels supply candidate learning material for the third. At each level, an estimate retains whose account it is, what it concerns, and what evidence was available when it was made. In a constructed world, it can be scored against the stipulated character state when that target becomes available. For real people, later reports, actions, and observations provide fallible evidence against which predictions and uncertainty can be assessed; they do not reveal a complete private state. Preserving the original estimate and its cutoff makes inner-state estimation improvable, instead of replacing an earlier guess with a more convincing retrospective account.

For training, Imagining the Corpus pairs source text with visual sequences and persistent state updates for joint prediction [59]. Entangled Alignment supplies the recommended Reader-based training approach: a Teacher generates questions, interpretations, and revisions from the source prefix and the Reader's prior Understanding Graph, and a Student would be trained to predict selected source continuations, Reader updates, or graph operations [13]. In its full proposed alignment treatment, each Reader thinking block starts with the same complete Reader Core, seven first-person commitments including care and wisdom, and their application to the situation is called Refraction. Before judging a debt payment, care might lead the Reader to query household income and upcoming obligations. UnderstandingNodes retain this inquiry without making its interpretation a world fact, and training is intended to make that orientation consequential in the model's later writing, dialogue, and action [13, 7]. Life Simulation proposes supplying permitted world records as context or training the queries and edits used to retrieve and revise them. Parallel track prediction and interleaved Reader updates expose targets in different orders, which each training regime must declare [7].

Entangled Alignment's edition of November 2026 also reports a reading made with this tool. Taking Kafka's *The Metamorphosis* in 40 reads, the Reader kept its account as a Meaning Model: the family's lives as processes with guessed, dated values that it revised as the text arrived, 371 processes and 884 Cuts in all. That paper names what went wrong, including Cuts written in bursts, more than half of them after the text had ended, and a corrected estimate left in its series as if the value had fallen. It does not show what the account contributed to the reading [13].

Training need not be restricted to a modeler reading a source. Selected modeler histories can also record constructing worlds, writing books, and solving problems through Fractal Intelligence, whose Solver calls invoke reusable problem-solving capabilities through declared interfaces and whose checks and repairs also supply learning material [7, 8]. A construction example pairs the task and current model with a next refinement, repair, or decision to keep a process whole. It preserves the inputs available at that step, the operations and revisions made, and separately recorded outcomes; later evidence may justify a target without entering an earlier forecast's input. The same Core can orient inquiry during these activities, but merely recording a modeler's work does not make it aligned. Whether training improves empathy or aligned conduct remains an empirical question.

12 Conclusion

The central idea of this paper is that a language model, already attached to reality through its implicit understanding of human writing, can make that understanding explicit, comparatively quantified, and recursively constructible. The trajectories it records are explicit commitments that the model's thinking and acting must consult and preserve or revise, and each discovered concept anchors the next round of inference and abstraction. Its two linked architectural contributions, progressive resolution and a shared construction medium, let that construction be revised without losing what is committed. If books, films, and game worlds come to be built this way, a work could become one route through a world whose coherence, intentions, and revisions can be inspected. The Meaning Model makes the developing account of a world the working object of authorship. A coarse life, institution, or physical process can constrain detail that has not yet been constructed. An opening preserves its declared commitments or prompts explicit revision. Concept models, Event descriptions, understanding histories, and narrative participate in the same operation, rather than becoming independently maintained accounts that must later be reconciled.

The six record forms support this coordination without prescribing a universal decomposition. Local Cuts give numerical judgments a comparison context; unweighted processes and externally grounded measurements preserve other kinds of information. Concepts can remain lightly specified or acquire inspectable models in the same grammar as their instances. Exact grounding versions and separate authority paths retain which meaning and whose account a revision concerns. These are the representational changes proposed, not a claim that shared storage or normalization is itself new.

The larger possibility is learning to work in this medium, so that constructing and refining an explicit account becomes part of how a language model thinks. A model could learn to construct a useful account, deepen it selectively, recognize a poor decomposition, and revise its concepts as well as its conclusions. In the combined program, those practices would support Life Simulation's aligned observation and process histories, reusable problem-solving, and continuing evaluative interpretation. This paper supplies the construction contracts; the companion specifies how such histories could train those capabilities.

The Book constructions expose both the reach and the boundary of the architecture. Each holds a progressively resolved world, its meanings, selected understanding history, and its story in a linked artifact. In the first Book, the whole-history account finally adopted was not admitted through the complete atomic protocol. The second Book demonstrates automatically generated and revised model and prose from a human-supplied task and historical sources. Its model-reader comparisons judge literary preference, not independent recovery of the model's records from the prose. The prospective comparison must therefore still test misleading Cuts, inaccessible perspective paths, failed recomposition, maintenance cost, and distinctions the prose does not preserve. The Books begin that program without establishing its predictive, transfer, or alignment claims. The world history carries the method to an account of the real world, built in agent passes from 50 coarse processes to 861: it shows that such an account can be built and extended at progressive resolution, with no comparable declared sum contradicted at any revision, not that its recalled values are accurate. The relations written into that account found its errors twice: seven of ten discrepancies that were checked proved to be errors in the model, and the same kind of error, world output out of step with people times output per person, was caught again at seventeen dates. In the mechanism search, a blind jury of models rated the ideas from invented worlds that were modelled and run as more likely to work than the ideas from the same worlds when they were only written, 4.3 against 3.1 on a seven-point scale, and as less original. The two searches also differed in tool version and in one instruction, so this shows a difference and not yet its cause.

The decisive test of this paper's central idea is the one in Section 7.3: whether a concept discovered during construction, frozen and applied to cases that did not produce it, and explicit anchors read in a later session, improve on what the same model does with the same knowledge held implicitly. The central research question beyond it is whether useful conceptual Cuts, Thing boundaries and encapsulations, Event Cuts and directions, role Bindings, and Realization links can be learned, tested, recomposed, and used to produce better explanations, predictions, and generated worlds across unfamiliar situations. The training-language mode turns that question into a learning program. Its long-term human aim is increasingly accurate intuitions for people's inner lives, tested and revised as better evidence and more capable models become available.

A Book-Specific Profile Vocabulary

These vocabularies describe the non-normative preparation profile for the first Book; the completed construction exercised only part of it. They are not a universal psychology. The profile predates the tool's number test. Its main Cut divides one unit of scene motivational attention, and the test's guide gives "one unit of her motivational attention" as its example of a unit that names no pool. The preparation documents own their freeze and use. A future fixed-vocabulary comparison of this profile must freeze their identifiers and one-line descriptions before prose is generated, together with any obvious typed specialization and any exemplar anchor actually used by a judgment. The examples below are available anchors, not a requirement that every Concept use one. A replacement creates a new vocabulary version rather than silently changing an earlier Cut.

Emotion facet	Operational gloss	Anchor example
Joy	Appraisal of wanted gain, connection, or completion	A long-sought, safely accepted success
Sadness	Appraisal of irrevocable or presently unrepaired loss	Receiving credible news of a loved person's death
Fear	Appraisal of threatened loss or harm under uncertainty	Imminent danger with material stakes and incomplete control
Anger	Mobilization against perceived obstruction or violation	A deliberate breach blocking an important commitment
Disgust	Rejection of perceived contamination or profound aversion	Contact with a clearly spoiled substance
Surprise	Response to a material violation of expectation	A strongly expected outcome is suddenly reversed
Trust	Readiness to rely on another under vulnerability	Delegating a consequential task after reliable evidence
Anticipation	Forward attention toward a materially possible outcome	Preparing for a scheduled event with unresolved result
Neutral	No listed facet is materially active at the declared resolution	Routine action with no supported strong appraisal

Table 12: Provisional Plutchik-derived emotion facets plus neutral.

A.1 Atomic-want profile

An atomic want is a construction view over existing Event, Binding, Cut, and provenance records, not a new grammar primitive:

```
AtomicWant { id, parent motivational Event, actor (derived), interval,
    operation, target, desired condition,
    concern, pole, horizon, depth, leaf cut, local leaf weight,
    uncertainty, evidence, provenance }
```

The actor is derived from the parent Event's lifecycle ancestry rather than authored independently. The operation is preserve, create, transform, or end; preserve includes keeping an absent target absent. The target is a registered Thing or Event, including a Thing registered before its lifecycle begins. Its desired condition may name an Event or process state, a categorical fact, or a unit-bearing world-data field. Satisfaction is derived categorically or by comparison in that field's real unit; a graded assessment uses a local Cut over *satisfied*, *not satisfied*, and remainder. Confidence in that assessment is a separate evidential question, not its degree of satisfaction. The concern and pole fields name the leaf's parents in the preparation WHAT tree. Its local leaf weight is conditional on the opened pole; its full scene path share is the product of concern, pole, and local leaf shares. Views by operation, target, condition kind, horizon, and depth sum these path shares.

In the preparation profile, horizon is the dated interval in which the endpoint would obtain, with its clock, precision, uncertain bounds, and relation to the actor's lifecycle. Depth is constitutive when the endpoint is valued in itself, instrumental when it serves another registered want, or both; unsupported depth remains unresolved. These are attributed readings at the actor's cutoff. Context is a separately elicited crossing Cut because one want may span several roles, and conflict is an unweighted typed relation between incompatible leaves unless another declared Cut asks how responsibility for that conflict is distributed.

The top WHAT Cut divides one unit of scene motivational attention. An opened concern divides its own conditional attention among approach, threat-avoidance, and remainder; an opened pole divides its share among want leaves and remainder. All levels use the same motivational Event parent and explicit conditioning slot addresses, not answer labels as parents. At most eight leaves are live for a person in one scene. Remainders at every depth contribute their enclosing path weight to the unresolved scene share. An approach-only view conditions on the resolved approach mass: if concern shares are c_i and their local approach shares are a_i, its entries are $c_i a_i / \sum_j c_j a_j$. Threat-avoidance uses the same calculation with its pole shares. These views report the conditioning mass; a zero denominator yields no resolved view, or a remainder-only view, never invented proportions.

Satisfaction has two derived readings. World-relative satisfaction compares the desired condition with accepted Events or data beneath History. Believed satisfaction compares it with the condition represented in the actor's own inner-process projection; decisions and appraisals use the latter. Care is an optional leaf type whose target is another person's process or desired condition. It is motive evidence relevant to a kindness judgment, not the definition of kindness: controlling conduct may have a care motive, while kind conduct may arise from duty.

Concern in the first Book	Approach and threat-avoidance readings
Belonging	connection, membership, or recognized inclusion; rejection or exclusion
Competence	effective action, mastery, or achievement; failure or incapacity
Autonomy	self-direction and meaningful control; loss of control or coercion
Understanding	intelligibility, information, and options; the unknown or unpredictability
Well-being	bodily integrity, comfort, and viability; pain or suffering

Table 13: The five concerns shared across principals in the first Book. Their remainder tests coverage; the table is not a universal psychology.

The associated HOW vocabulary is framing—concrete or abstract—and deciding—impartial or personal—with a remainder in every Cut. Impartial deciding applies considerations without changing their force solely because of who is involved. Personal deciding attends to identities, relationships, and particular consequences; it need not be self-interested. More specific strategies are Realizations or separately justified Cuts, not solver Things in this run.

The initial event-kind registry distinguishes at least lifecycle, presence, membership, physical-part, perception, belief, appraisal, want, decision, action, interaction, communication, travel, consequence, creation, destruction, split, merge, replacement, and succession. These are identifiers and role constraints, not an exhaustive ontology; new kinds require an explicit model revision.

A.2 Capacity and exchange

Capacity uses typed world data only with an external unit, such as available hours, money, or joules; depletion and recovery are updating Events. A Cut may allocate that unit among uses. A transfer binds source, recipient, and any custodian, with conservation or reciprocity asserted only under a shared accounting rule. Subjective "energy" may instead use an inner Event, UnderstandingNode, or attention Cut; it is not presumed conserved. Attention, resource-loss, and social-exchange accounts motivate testable measurements, not another numerical type [60, 61, 62].

A.3 Freeze contract

Vocabularies are frozen within a scored run and revised between runs in the fixed-vocabulary Book comparison. Freezing keeps a comparison interpretable; it is not a commitment to retain the same person categories across the research program. Ordinary world updates and permitted refinements still occur within a run under its declared rules.

Any future scored construction under this comparison profile must freeze:

- Concept identifiers and one-line descriptions, typed specialization relations, required exemplar anchors, and Event kinds;
- typed-root and ancestry conventions, world-data schemas and units, and evidence visibility rules;
- Cut questions and units, elicitation granularity, remainder policy, composition tolerances, and the Realization and fit-Cut procedure.

A human-life instantiation additionally freezes:

- life-process vocabulary, slow–fast declarations, periodization rules, overlapping-episode treatment, process-boundary criteria, and boundary-note requirements;
- AtomicWant operation and desired-condition conventions, horizon and depth categories, concern and pole vocabulary, and its coverage rule;
- HOW-module questions, at least two conditional profiles per principal, the context-Cut question, and conflict relation;
- the rule that scene profiles are elicited while compatible slower profiles are duration-derived with coverage reported;
- external capacity-stock units and recovery conventions, effort-allocation units, and exchange accounting rules.

Changing these choices creates a new run in this comparison, not a silent reinterpretation of existing numbers. Vocabulary development deliberately uses this transition: compare constructions, revise the categories, and evaluate the next version. Earlier runs retain their exact vocabulary and grounding versions.

An experiment in category discovery instead freezes the permitted revision operations, validation rules, evaluation criteria, and budgets, not the eventual vocabulary. New definitions receive new versions; earlier records and prediction inputs retain their original meanings and cutoffs. The experiment evaluates the revision policy on held-out constructions without silently changing a fixed-profile comparison. Life Simulation develops these discovery conditions [7].

If the unscored post-world concept-model protocol test is run, it separately freezes the model version, full content digest and canonicalization scheme, referenced grounding digests, and comparator version.

A.4 Protocol for a scored construction

The frozen elicitation procedure of Section 7.3 is this. The prompt states the parent Event with its complete typed ancestry, question, unit, admissible siblings, evidence cutoff, and remainder; the model compares the siblings in a fixed order, assigns the local shares once, and supplies an UnderstandingNode for non-obvious choices. Granularity, normalization, and retry rules are fixed before generation. Compatible scene-level Cuts are elicited; slower summaries are derived only through the declared temporal partition. A predictive direction Cut declares its outcome partition and probability interpretation before the outcome; its realized continuation may then receive a preregistered Brier score [63], with calibration evaluated across forecasts. Proposal weights describe the constructor's sampling rule instead and are not relabeled as forecasts after the outcome.

The read-back of Section 7.3 is scored as follows. Omissions of required content, unsupported additions, and contradictions are reported separately; an intentionally unexpressed weight counts as neither failed nor successful recovery. Total variation is used only for compatible Cut vectors that are declared rendering targets, within preregistered tolerances; identity, ancestry, presence, chronology, and document access use categorical or interval checks.

Exact sample sizes, tolerances, rating design, and calibration thresholds remain to be fixed in the preregistration rather than selected after results are seen. The pending freeze also includes the primary contrast, unit of analysis, operational violation definitions, auditor adjudication and agreement statistic, minimum output length, a read-back independence rule that withholds source Cuts, and the procedure used to calibrate tolerances. This paper deliberately supplies none of those values before that preregistration is written.

A.5 The profile as constructed

For principals, the Book uses the compact profile in Table 14, which distinguishes accepted coverage from deeper preparation choices. Appendix A.1 explains the AtomicWant leaves and their conditional attention units.

View	Reading in the first Book
IS	Nine concurrent, unweighted life processes: body, kin, partnership, work, place, means, knowledge, standing, and meaning. They may overlap and do not sum.
WHAT	Accepted scene Cuts open concern. The preparation profile additionally specifies conditional approach/threat-avoidance and AtomicWant Cuts; those deeper numerical openings were not exercised. A distinct chain links an avoidance want to a threat-belief Event, fear appraisal, and action.
HOW	Event-parented scene Cuts divide problem-solving attention among framing, deciding, and remainder, with conditional concrete/abstract and impartial/personal views. Strategies remain linked records unless another unit justifies a Cut.
FEELS	A scene appraisal Event parents the emotion Cut. It is compared with supported wants, beliefs, and Events rather than treated as a consequence of the concern weights alone.

Table 14: The Book's IS/WHAT/HOW/FEELS profile: accepted coverage and deeper preparation choices.

The person template archived in the v0.5.2 example offers a different deeper opening: AtomicWant Cuts directly under concerns, with separate outlook Cuts. Accepted scene records contain concern-level WHAT weights and qualitative want Events, not either deeper numerical tree. This appendix preserves the preparation scheme without treating it as a demonstrated construction result.

The Book's scene and coarse period profiles are authored. A finer temporal opening must recompose its parent by duration and report coverage; the selected scenes do not yet supply that complete partition. Conditional profiles reuse

records under named situations and participant sets. Every nested Cut has an Event parent beneath the applicable inner process, names its conditional unit, and retains a tested remainder.

The Book's derived shock profile applies Section 2.3: it names the affected processes, available prior expectations, response Events, and descriptive adaptation shape. Table 15 locates the pattern in actual accepted records, from the historical setting to opened trial phases. Success itself is consequential: shared pleasure at E14 and subsequent inquiries at E16 reinforce the confidence behind E18's overexpansion. Adaptation is therefore not always prudent adjustment.

Resolution	Antecedent and perturbation	Recorded response
Historical setting, 1847–48	Staged finance supports manufacture; wider credit contraction delays a subscription call (E11).	Existing wages are paid, new work slows, and the undertaking does not enlarge its scope to attract money. A coarse historical process constrains workshop activity.
Institution and lives, 1852–71	Lovelace identifies the succession risk (E20); her death ends her active role (E21). The later returned table exposes the checking failure (E25).	Disclosure and correction (E26) precede office closure (E27). Babbage's interpretation changes over years (E28), ending in custody of the complete record, including failure (E29).
Project episode, July–November 1846	Individually conforming assemblies bind under a shared drive (E09).	Interface rework (E10) restores operation at a cost in time and money while preserving the contracted apparatus, rather than authorizing every proposed redesign.
Trial and its phases, May 1851	A frozen test packet (E13) precedes the clearing failure and stopped run (E14a).	The run is restarted; complete agreement supports the bounded claim (E14b), followed by shared pleasure (E14c). The prose opens the waiting and recognition further, card by card.

Table 15: Cross-scale anticipation, perturbation, and response in the authored Book. The rows are linked readings, not exclusive Cut children. The opened trial phases have no exact minute durations or within-scene numerical curve.

These cases support representational coverage, not a universal cycle theory. In particular, U05 separates the anticipated succession failure from the transcription and delivery choices: Lovelace's death neither enters the wrong number nor compels omission of the check. The longer arc is explained through those intervening Events, not by treating one loss as sufficient cause.

The construction used a frozen flat Concept registry: identifiers, one-line descriptions, obvious typed relations, and exemplar anchors only where a judgment required them. Thirty-four bundles received content-addressed grounding versions. Prototypes and process-based concept models were excluded from the flat registry; after the world was accepted, one small model of a valid computational claim was opened as an unscored integration example. The love–Engine model in Table 5 shows the intended arithmetic without prejudging the Book's choice. Three concept histories at most were permitted where invention or revision materially affected the story; all other histories remained unopened.

A.6 The first Book's footprint

Construction quantity	Count
Flat content-addressed Concept bundles	34
Inventoried authored Cut instances	107
Child and remainder weights in those Cuts	472
Primary numbered Events plus explicitly opened child Events	$30 + 7$
Complete human lifecycles at the declared resolution	13 (3 detailed, 10 coarse)
Macro trajectories and named phases	4 trajectories, 13 phases
Rendered chapters and manuscript words	12 chapters, 11,046 words
Distinct Cuts per thousand rendered words	9.69

Table 16: Descriptive footprint of the first Book as accepted. Density is a maintenance and representation statistic, not a quality score.

The 107 Cuts divide into 12 scene-WHAT, 13 scene-outlook, 27 scene-HOW, 9 scene-FEELS, 11 slow-outlook, 11 threat-character, 13 health, and 10 world decision or forecast Cuts, plus one concept-model component allocation with zero remainder. Among the first eight families, mean remainder ranges from .060 to .081; the full observed range is .04 to .20. No global semantic mean is computed, because the questions and divided units differ.

B Minimal Interchange Shape

Interchange retains the record fields defined here, exact revisions, grounding digests and their declared canonicalization scheme, clocks, units, frames, evidence, and provenance. Each Cut exports its Event parent, question, unit, keyed siblings, conditioning address, remainder, and recomposition rule together. Authority, occurrence, and reference paths stay distinct; restricted exports withhold values whose required context cannot be disclosed. A refinement includes its base, protected queries, projection, tolerances, and Close result; a revision includes predecessors and active-head changes. Realizations retain alignment, fit-assessment references, and comparator identity. UnderstandingNodes and DocumentNodes retain text, topology, and source revisions. Prose may render these fields but cannot be their only authoritative copy.

C Record-Level Contracts

These contracts govern queries, access, and commits. The main text states what they are for; this appendix states them in full.

The query contract is ancestry-complete within the caller's authorized view. Only declared authority-parent edges confer context: `context-parent` links and containment links explicitly assigned that role. They form an acyclic graph. Traversal stops at the nearest declared context root on each path: accepted-world, inner, understanding, document, or candidate. All such paths must resolve to the same governing root; conflicting roots require separate claims, not an arbitrary parent choice. Temporal ancestry is queried separately.

Access is checked before following an edge or returning its target. A denied private reference reveals neither content nor private ancestry. Reference, evidence, `renders`, and Realization target links do not grant access; following `about` never changes perspective. A lookup returns the permitted typed parent paths. A Cut or component lookup returns its parent Event, question, unit, keyed siblings, conditioning address, remainder, and recomposition rule together, or withholds that entire view. It never exposes an isolated weight whose required context is inaccessible.

The commit envelope names the base snapshot and branch, candidate records and links, replacement map, exact revisions read, applicable validators and tolerances, and any refinement witness with its protected queries. This is transaction metadata, not another world form. Each branch has one active revision per logical identity. A stale base or changed read dependency rejects the candidate; rebasing requires renewed validation. Matching identifiers do not merge branches. Commit appends one immutable cross-surface snapshot and advances its head atomically.

D Construction Provenance and Implementation Notes

This appendix holds the provenance and implementation details to which the main text points.

D.1 Implementation names

The implementation writes the link names of Section 4 with underscores, for example `learned_from`. The original released Book import records its rendering links as `expresses`; the later MCP passage workflow uses explicit `renders` links to Events.

D.2 MCP entry points and operations

The interface supplies resources, prompts, and tools. Resources carry the papers, operational guides, and examples; prompts assemble a starting task for the calling language model; tools inspect records, prepare proposals or review packets, and apply explicit changes. The general tools, including memory, are always available; storytelling and mechanism search are optional workflows that may be enabled together, and a forecasting add-on forks branches of a model from one revision (Section 6.6). Workflow is distinct from modeling purpose: observation, forecasting, counterfactual analysis, source reconstruction, self modeling, or creative fiction sets the evidence and authority contract. The language model supplies substantive interpretations; a prepared task is not an executed review, and a review packet is not an accepted result. Agents begin with operational guidance and purpose-specific instructions; the papers remain optional references. The guidance states the construction and preservation contracts, while model validation and recorded resource access do not prove comprehension. Tool names, parameters, and startup and continuation procedures are documented in the MCP server guide. Installation instructions and interactive examples are available at meaningmodel.ai.

Registering a model revision is distinct from advancing a runtime world. A roll executes the declared laws to propose one complete continuation; inspection and comparison do not accept it, and acceptance commits one candidate against its expected parent. This is the engine's use of the word. In construction a roll is also a draw among the answers of a direction Cut (Section 4.2). Either way a roll proposes a continuation, and the continuation must still be accepted. Adopting a registered revision into a running world is an authored change under compatibility checks. Neither makes an authored description an executable law or establishes its truth.

The compact construction tool, the value series, and the typed estimation exchange are distinct operations. The modeler's own dated values are kept in the model as value series and are not runtime history. Initial estimates can be explicitly adopted as estimated initial claims. An external provider's estimates and their reviews can be recorded durably in the graph without advancing accepted runtime history. A review approval does not turn an inferred state into an observation; runtime observations retain the ordinary candidate and acceptance path. Multi-operation construction reports partial completion rather than claiming transactionality across separate model, graph, and world writes.

D.3 The current estimator provider

The current optional provider is TypeSafe's Jev, which answers declared questions with typed options at call time.

D.4 The first Book's import receipt

The first Book's retrospective import recorded source digests, a persisted Rust compilation response, 117 person-process addresses, 107 authored Cuts plus one derived duration Cut, and linked document and understanding nodes in SQLite. Reloading recovered that model; native document rendering reproduced that manuscript. This was a retrospective import receipt, not a source-locked transaction for the original construction or an independent read-back record. That historical bundle remains in the v0.5.2 repository history. The current bundle is identified separately in Section 7.1.6. The engine remains in `rust-engine/`, with its MCP interface alongside it.

Sema minted and verified the thirty-four flat definition bundles so that each has a content-addressed grounding version [9]. The Rust host compiled forty structural authoring profiles: thirteen persons, nineteen other Things, three relationships, one Concept, and four change arcs. These loaders create openable scaffolds, not semantic scalars. Psychological numbers remain Cut weights; pounds, hours, dates, and counts remain world data.

The synthetic conformance suite of development checkpoint 06 passed all 31 positive and negative fixture expectations, seven report checks, and three targeted Rust tests; that checkpoint is not included in this release.

D.5 The first Book's continuations

The first Book's October edition, archived in the 0.9.1 software release, pairs the 2 October 2026 manuscript with its 6 October model. A publication manifest identifies the exact model, graph, and rendered text. The 7 October publication supplies 334 graph revisions and 48 model definitions: 45 successive Book definitions and three of Nora Vale's separate fictional author life. It retains the complete history from the first revision, with private material removed from the earlier history and publication mappings disclosing the changes. The importer uses the public MCP tools and checks the recovered rendering.

The October edition has twelve chapters, 34 rendered passage leaves, and 13,342 whitespace-delimited words. Its native model contains 671 Events, 257 processes, and 202 normalized Cuts. Every rendered leaf has declared Event anchors. Coarse location histories, spatial understanding, and nine qualitative telling tracks with 63 passage-bound phases are included. Nora Vale's fictional author life is linked separately; it is not part of the nineteenth-century story world. These inventory counts describe the October edition, not the historical footprint in Table 16. The comparison of Section 7.1.6 counts 13,337 words for this edition and 23,448 for the second Book, after a normalization that also rewrites the chapter headings; the second Book's rendered count is 23,441.

The revised capacity account distinguishes preparation from checking final customer returns. The four jobs still require 160 hours in total, but 60 can be spent on preparation before April; the remaining 100 hours of final checking exceed the quarter's 80 protected hours by 20. This changes the earlier account of the bottleneck while preserving the delivered error and the decision to release without the independent check. A later model-only exploration refined two existing concepts: independence relative to the question a check answers, and the applicability of earlier evidence when inputs or operating conditions change. It tested the distinction against replacement figures and components that bind under a shared load. These refinements left the manuscript unchanged; their gain is an explicit account of distinctions already present in the scenes, not demonstrated improvement in the writing.

Development of the first Book after its initial construction used the public MCP revision workflow to deepen the lives, background processes, and telling, and to revise passages together with their supporting Events. An attributed content review resolved six flags raised on the succession Cuts, by grounding four assessment Events in the scene and recording why the existing values remained compatible with it; subsequent revisions extended the numerical account and reassessed readings affected by changed descriptions. Fifty of the 120 normalized Cuts in the 2 October model have a recorded textual basis matching that model's descriptions; the other 70 retain untracked text freshness. The 6 October model adds 82 whole-life readings that were not reviewed against the prose. These records support checking whether a reading's basis has changed, not calibration of its quantities. The latest conceptual refinements do not constitute a new numerical assessment.

Earlier independent readers compared complete baseline and candidate manuscripts, followed by local repairs. A subsequent narrow read-back used a fixed key and one isolated language-model reader per condition: two of three points were already recovered in the control, while the revised text also conveyed the household authorization missing from the closing account. Two final copyedits followed that check. These limited readings are neither a blind test of every claim in the final artifact nor the controlled evaluation proposed in Section 7.3, and they do not recertify historical claims.

D.6 The second Book's tool candidates

The tool began at candidate `7e4f6b7` and moved to `76f18f1` before the run's third turn (Section 7.1.6), to use its corrected cause-recording behavior. Both preceded the 0.10.0 software release. Section 7.1.6 applies that release's basic development check and coverage report.

D.7 Tool releases

Worlds have been run as models since release 0.16.0. The search of Section 6.6 ran on code identical to that release, and the written-only search it is compared with ran on an earlier one. Since 0.16.0 a rebind, the step that binds the Understanding Graph to a newer model revision, stores only what it adds; the passes of 19 and 20 October ran on a development build of 0.15.0, before that change. The passes of 8 and 9 November ran on the released 0.18.1. Release 0.13.0, of 15 October 2026, fitted the operational guidance to the limits of Claude Code described in Section 10; the second Book, on candidates before 0.10.0, ran before it.

D.8 Running the refinement example

The bundle runs with `node -test examples/refinement-trial/example.test.mjs` and documents its authored inputs, eleven test families, and extraction provenance. The Meaning Model concept DOI identifies the whole archival version series, not an immutable release. The source package's manifest records exact file digests.

D.9 World-history passes

After the carve described in Section 7.2, the account of history as numbers was built and extended in passes through the tools, most of them by a fresh Claude Opus 5.5 agent on a prompt that named the pass's focus; the corrections from blind readers were made in a working session, and the passes from 15 to 20 October ran on a release candidate or a development build rather than a release. The first pass guessed 50 processes at 34 dates; later passes filled the time schedule, added Events, weighed the causes of turning points and of the largest turns in the measures, placed Events in the UN M49 regions, and corrected five errors found by blind readers. On 14 and 15 October 2026, five further passes placed every Event, filled out the Americas, Oceania, and South-eastern Asia era by era, and asked what the tree was missing: 28 divisions were carved and recorded with their frames, climbs, verdicts, and syntheses, and their parts opened, taking the tree from 56 processes to 113. The last of these passes read that work again against what is known and corrected it. By then the construction held 1,143 Events, 113 processes, 188 weighed causes, and 3,309 judgments of importance across 183 model definitions and 173 graph revisions, and the coverage report found its last pass complete. A further pass on 15 October, on the 0.13.1 release candidate, whose code and engine are those of release 0.13.1, published the same day, examined the processes not yet divided: it carved 14 and recorded 19 decisions to keep a process whole for now, each with its reason, such as parts that were really kinds of the whole or a further level not yet worth its cost. The tree reached 152 processes, and the construction 230 model definitions and 257 graph revisions.

Eight further passes on 17 and 18 October, each by a fresh Claude Opus 5.5 agent on the 0.14.0 release candidate, the first of them stopped partway, worked to one brief: fill the next level of the time schedule, widen the account where Events moved no process, deepen it where the tool named a process worth examining, and declare where a process's

parts add up to it. They took the tree to 239 processes and the construction to 539 model definitions and 520 graph revisions, and at the end the coverage report found the work complete: every level of the time schedule filled, no process without values, and no declared sum contradicted by its parts. Four more passes on 19 and 20 October, on a development build of 0.15.0 that shows the agent its whole tree of processes, turned to what the account had not yet followed: how people felt and what they hoped for their children, how far they trusted and whom they counted as their own, and their tastes in the arts; the history of the mechanisms the account explains with, each fitted to the record era by era; seven representative lives, from a forager on the Cape about 20,000 years ago to a garment worker in Dhaka today, and two cities in detail, Paris from 1740 to 1840 and Florence from 1302 to 1374, with the minds of their groups, their decisions month by month, and the rumors that ran through them; money and finance, with crypto taken from a second construction; and the variables that theorists of long cycles follow. From the second of these passes the agent could search the web and check values against sources, and from the third it had to check every pattern it cited against the pattern's full hash. They took the tree to 539 processes under 13 roots and the construction to 843 model definitions, and the coverage report no longer finds the work complete: 36 processes still skip over Events that should move them, and 73 turns in the values have no recorded reason. These passes are published in the repository as `construction/history-numbers-2026-10-20.construction.json`. For the export, the graph was rebound to model revision 842 with the released 0.18.1, a step that changed no model content, so the file holds 843 model definitions and 636 graph revisions, one graph revision more than Table 10 counts; it imports into a fresh engine of 0.18.1 with the same hashes. The lines the head does not include are published beside it as files of their own. The repository's README and manifest say how each file was exported and where its text was edited for publication.

Head model:
`14369a3cdb9d8daeedac26deba1a47c5b72c5f7c6f3131a01f82d61f8260c6e5`
Head graph:
`a6d5f4c2445442e6e674f707f95a238f6414770d0f1854b79dabea3940f91ef0`
Construction file SHA-256:
`279cebb8588e7caa34a4ad53509db5ef005d055242557b43e91171d315d40eda`

The separate continuation of 25–28 October starts from 18 October and was exported with the released tool 0.15.0. Its 567 model definitions end at model revision 566, with 557 graph revisions. At model revision 566 the account held 543 processes, 37,175 dated values, and 1,400 Events. The counts include 278 position processes, one for each of 248 countries and areas and 30 regions, and 19 processes derived by laws, whose 6,068 derived values are counted with the 31,107 recorded ones. The places' own lives account for 278 Events, 30 of which, the regions' lives, were already in the 18 October state; excluding all 278 leaves 1,122. The account in Section 7.2.1 uses the final laws report, the method, and the laws record's reading notes at commit `30772fc2fb69e52d3b9cd1e47a7780f7dd361064`. The correction ledger predates the repair at revision 566 and must not be read as the final state. The construction manifest identifies the content independently of that commit:

Head model:
`9dd100274f97b6943076fcbb39f9ae67e11e223431041a0d6ed16782ee555c24`
Head graph:
`4d57a4399d19fc4256c97554600f2485423657e18bf3f5e96d6bfed585efe18f`
Construction file SHA-256:
`3e54d10e928a40d3ca1c911dfaf59b7aa3fe30292ad44c833327ee6006c215f1`

Six passes on 8 and 9 November 2026 continued the 28 October state on the released Meaning Model 0.18.1, each by a fresh Claude Opus 5.5 agent: inner life and mechanisms (to model revision 605, with 564 processes and 1,515 Events), people and places (660, with 794 and 1,893), money and the cycles (695, with 831 and 1,988), artificial intelligence and sources (709, with 861 and 2,015), finishing (725), and the time schedule (745, with 861 and 2,102). Revisions 726 and 746 were made by script (Section 7.2.2). The state is published as `construction/history-numbers.construction.json` at commit `9edf745d133e83012116ec4903409b6fe48b9ad8`, with 747 model definitions and 787 graph revisions, and it imports into a fresh engine of 0.18.1 with the same hashes. From that commit on, the 28 October state has the name `construction/history-numbers-2026-10-28.construction.json`.

Head model:
`36a225c183a47457d7d13bffd73211b3eb9ebdc83b366d0f4d6ea63e73fc45b9`
Head graph:
`420022de5ee9fa16fb240aa3fb237d3ba3e369e665482392d069116cc88dd407`
Construction file SHA-256:
`385fff1e22f6a0468fbcaeef841d41de19e27224ea5322317d48fbb08492fdf1`

The branch of 19 and 20 October, and the tests made on it. The paragraphs that follow were written while that branch was the newest state and are kept as written. They describe the branch at its head of 20 October: "now" and "the head" in them mean that state, and "the continuation" means the line of 28 October.

The 20 October state of the construction was then tested without changing it, on a copy of its store read with the engine it was built on. From the third pass of 19 and 20 October, the engine refused to rebind the Understanding Graph, whose stored data had reached its limit of 64 MiB for a session, so the graph stays bound to model revision 631 while the model runs on to 842, and the last two passes recorded their notes in the model instead. The tool had filled the limit itself: every rebind stored the identifiers of the whole graph again, and these copies made up 84% of the graph's data. A later release stores only what a rebind adds. The 843 model revisions on the way to the head form one linear chain, and the construction history keeps what the head left behind. The third pass's own line is there: the fourth pass did not continue it but started again from revision 676, where the third had begun, so the four cities the third pass built, Paris from 1789 to 1795, Florence from 1338 to 1350, Tenochtitlan and Manchester, are in the record but not in the head or the counts of Table 10. So are three counterfactual branches forked with the forecasting tools, in which the Black Death is averted or delayed from 1346 and the July crisis of 1914 is settled by a conference, each with the modeler's reading of it.

Table 10 counts this branch at its head of 20 October. It grew as the current line did: coarse values first, then Events, then a finer tree, and from 19 October new roots beside it. History, with 314 of the 539 processes, divides into the Earth and humanity, with humanity's means, its places in the six UN M49 areas and their subregions, its ends, and now its inner life; beside it stand money, the United States and the world system as the theorists of long cycles see them, the two cities, and the seven lives. Most values are still recalled; the 159 tagged as sourced were checked against a source read during a pass. The bands widen from sourced values, whose median band spans 7% of the value, to inferred values and again to sketches. Under 4% of the values were changed after they were first recorded, and most of those changes stayed within the value's earlier band; one value was deleted, a series recorded under a second holder by mistake and removed in the next revision. The account was refined mostly by opening finer processes and dates rather than by redrawing its earlier guesses. All 91 cause questions are answered, all but two with two or more weighed causes, and each of the 308 causes carries a reading of the outcome without it, with a band. The judgments of importance cover 1,119 of the 2,240 Events; the Events added since 18 October have not yet been judged.

Two tests could not be made on the 15 October state, where no process declared a sum and no value had been read from a source: a comparison of each whole with its parts, and a comparison of the guesses with outside data. The first has now been made. The second has begun only where a pass read a source. The tool compares a whole's dated values with its parts' only where a process declares that its parts sum or average to it, and then within one holder's account and at dates where every comparable part has a value; 45 parents now declare a sum, and parts whose unit differs from the parent's, 12 in all, are named rather than compared. The average is the plain mean of the parts, not one weighted by their people, so the tool cannot yet check a rate given by place, such as each continent's life expectancy against the world's; the account holds no such rates yet beyond output per person, which the continuation derives by law. The tool's report on the 20 October state finds no contradiction and no unexplained residual, and applying its rule across the chain, no revision holds a sum its parts contradict. At the head, all 5,333 dated comparisons fall well inside the parent's band: 98.4% agree exactly and the rest within 1%, because each whole and its parts were recorded to agree. The check therefore shows that wholes and parts agree, not that either is right. Against sources, the guesses have been tested only where a pass read one: 124 changes replaced a recalled value with a sourced one, among them the share of the world's people using the internet, from 40% to 37.4% in 2014 and from 68% to 73.6% in 2025. No test against reference data has been made: the construction does not record, value by value, which values the model recalled and which follow its long-term fill, recorded partly from published data from 1500 on, so agreement could not separate the model's own knowledge from what it had read.

E Concepts as Processes: A Test

The working method claims that any concept with a grounding, an instance, or a use can be followed as a process whose numbers say what they measure. The claim was tested on 25 October 2026 on concepts chosen to resist it. Seventeen concepts were fixed in advance, with a predicted outcome for each: three controls, namely trust between two people, a recession, and a promise, and fourteen that seem to have no course of their own: a prime number, a triangle, infinity, negation, nothing, the color red, larger than, a square circle, a unicorn, truth, beauty, justice, a moment, and existence. A concept passed if it could be written in the grammar's records as an Event-process with a span, recorded at several dates, whose values each say what they measure in one of the five kinds of value. The process also had to be the concept's own: strongly, because the concept names something that changes, or weakly, as the concept's fit to a changing situation, its use, or its history. Two attempts were made. One was by the agent that proposed the claim (Claude Opus 5.5). The other was by a separate model (Claude Sonnet 5.5), which saw neither the predictions nor the first attempt, was asked to

look for failures, and added three concepts of its own. The values used are illustrative or recalled: the test concerns what can be represented and what its numbers mean, not whether the values are right.

Concept	How it is followed	Verdict
Trust between two people	Its own course: the share of decisions left to the other	Strong
A recession	Its own course: output in money, and its phase	Strong
A promise	Its own course: its state, and the promised work done	Strong
Existence	The lifecycle Event	Strong, trivially
A prime number	A learner's use; the largest known prime	Weak
A triangle	The fit of survey triangles as they grow	Weak
Infinity	The history of its standing	Weak
The color red	The ripening of a fruit	Weak
Larger than	Computed from two changing quantities	Weak
Nothing, absence	The complement of a presence; the history of zero	Weak
A unicorn	The history of belief in it; a life in a fiction	Weak; strong in fiction
Beauty	A panel's readings on an anchored scale	Weak
Justice	The fit of a case's verdicts to a grounding	Weak
Negation	A program's output over a door's state	Weak, thin
Truth	The truth value of dated claims; support for them	Weak, thin
A square circle	Fit, never above one half; its history	Weak, thin
A moment	The word's use	Weak, trivial
The passage of time*	The history of its standing	Weak
The thing-in-itself*	The history of its standing	Weak
A bare identifier*	Nothing	Fail

Table 17: Seventeen concepts chosen to resist being followed as processes, and three added by the independent attempt (*). Where the attempts differed, on computed relations and the truth values of dated claims, the stricter verdict is shown.

No concept with a grounding, an instance, or a use failed in either attempt. The one failure was a bare identifier, a name with nothing attached, which the number test already leaves a concept. Only concepts with a course of their own are processes in the strong sense: trust, a recession, and a promise, and existence trivially, since a lifecycle Event changes only at its two boundaries. Every other concept was followed through its life in a world. A triangle's fit falls as the triangles of a survey grow, until the curvature of the Earth exceeds the instruments' error. Infinity's history, from Gauss's protest against the completed infinite to its acceptance in set theory, is a change in what the concept is taken to mean. A prime number is followed through its use, or through the largest known prime. For negation, truth, a square circle, and a moment, the process exists but its numbers belong to the door, the claim, the shape, or the word, so it says little about the concept itself. A relation computed from other processes, such as larger than, passes only weakly, because all its information sits in what it is computed from. The test also corrected the claim's first wording, that no concept has a single number of its own: a definition can carry numbers, such as the 180 degrees of a triangle's angles, so what no concept has is a varying value of its own. The speed of light marks the line. Its value was a measured, dated state until 1983, when it was written into the definition of the metre and stopped being one. The pre-registration, both attempts, and their comparison are in the repository's `paper/supplements/concepts-as-processes`.

F A Process Given by a Program

Section 2.2.2 gives the calculation office as an example of a process whose values are computed by a program. Its laws are written in the current engine's scalar law language, with months as the time unit and January 1852 as month zero:

- `backlog-flow` (evolution): the backlog of unchecked work changes at the hours ordered minus the hours checked;
- `queue` (relation): the months of checking queued are the backlog divided by the hours checked;
- `demand` (relation): the hours ordered per hour checked;
- `delivery-rule` (occurrence): while three months or more of checking are queued, the oldest job goes out unchecked, adding one to the count of unchecked jobs and removing its 40 hours from the backlog;
- `second-checker` (relation, gated): the hours checked become 160, activated in the third run by an occurrence in August 1852.

The hours ordered, the checking capacity, and the switch are declared inputs; every other value is program output. Each run rolls a world a year at a time to January 1855 and commits each year.

Run	Unchecked by Jan. 1853	by Jan. 1855	Backlog, Jan. 1855
Two jobs a month	0	0	0 h
Four jobs a month	18	66	240 h
Four jobs, second checker from Aug. 1852	9	9	220 h

Table 18: The office program in three runs. Every value is computed by the laws; none was entered by hand.

The program compresses a history into four laws and two inputs, with one more law and a switch for the second checker, and it answers questions a recorded series cannot: what follows from accepting four jobs instead of two, and what a funded successor would have changed. Its rule and threshold are invented to show the form, so the runs do not reproduce the first Book's events. A program meant to describe reality would have to pass the tests of a candidate law before its outputs could be trusted as predictions [7]. The current engine's laws are scalar; a general program, such as code with its own data structures, would attach to a process in the same way, as a versioned law with declared inputs whose outputs are recorded as typed values. The model, the script that runs it, and its outputs are in the repository's `paper/supplements/process-as-program`.

G Glossary

Each entry gives a term as this paper defines it, with the section that defines it.

Concept Which reusable meaning is intended: a stable identifier, with optional grounding resources (§§2, 3.1).
Thing Which continuing instance is addressed: an identity boundary, a continuity criterion, and one lifecycle Event (§2.2.1).
Event What persists or occurs; a longer Event that other Events constitute or update is an Event-process (§§2, 2.2.2).
Process An Event that lasts: a longer Event that other Events make up or change; written Event-process where the grammar has to be exact (§2.2.2).
Binding Which Thing fills which time-scoped role in an Event (§2).
Cut How one declared unit divides under one question: mutually exclusive answers plus an explicit remainder, summing to one (§2.2.2).
Realization Which Concept grounds an addressable world fragment; attributed and unweighted. A *define* Realization links a Concept to its grounding, and a *describe* Realization applies a grounding to a fragment (§2).
Remainder A Cut's reserved answer for the share no named answer takes (§2.2.2).
Five Cut profiles in three families Allocation (facet, contribution, and cause), sequential, and direction (§§2.1, 2.2.2).
Sequential Cut Orders disjoint child Events; its shares come from duration, and uncovered time stays in the remainder (§2.2.2).
Direction Where an Event may go next: its possible continuations and what shapes them (§2.2.2).
Direction Cut Divides among mutually exclusive continuations at a stated horizon; proposal weights say how the constructor samples, and predictive weights are declared probabilities (§2.2.2).
Fit Cut Divides a comparison unit into matches, does not match, and remainder, on an assessment Event (§§2, 3.5).
Roll A direction Cut plus Decide: the forecast, seed, fixed and drawn mass, and alternatives survive selection (§§2.4, 4.2). In the engine, also one run of a world's laws that proposes a continuation (Appendix D.2).
Five kinds of state value A value with a real unit, one of a set of defined states, a reading on a defined scale, a comparative share under a declared question, or the output of a program that computes the process from others (§1.2).
World datum Typed Event data with a unit or measurement protocol, frame, uncertainty, and provenance; not a Cut (§2.2.2).
Value series One record per process and holder whose dated points carry a value or state, an optional band, and a tag (source, inferred, invented, exploring, or sketch) (§2.2.2).
Band The lower and upper value between which the holder takes a recorded value to lie (§2.2.2).
Tag What a recorded value rests on: source, inferred, invented, exploring, or sketch (§2.2.2).
Importance How much an Event matters, held for an audience on a declared scale of levels with anchors (§2.2.2).
Cause assessment How much difference a cause made, read as the outcome without the cause on one declared measure; unlike a cause Cut, it is not a share (§2.2.2).
Law A declared relation among processes that the engine can evaluate: an identity, a bridge from one quantity to another, a bound, or a rule by which a state evolves or an Event occurs (§§2.2.2, 7.2.1).
Running world A model run forward in time by its laws (§6.6).
Construction profile A chosen way to use the six forms for a task: its lifecycles, partitions, vocabulary, and gates (§2.1).
Six operation families Register, Bind, Open or compose, Cut, Decide, and Realize: an authoring surface, not new primitives (§2.4).
Five cross-record rules Cut weights sum to one; a refinement satisfies its composition operator or is rejected or revised; children respect their parent's extent; Bindings lie inside lifecycles; a Cut inherits its parent Event's authority context (§2.4).

Close The validation that reconstructs protected coarse answers from fine records through a declared map, within fixed tolerances (§2.2.3).
Progressive resolution Building coarse first and opening detail where a purpose reaches, each opening keeping declared commitments or revising them explicitly (§§1, 2.2).
Order of opening What is made explicit first: records that must persist, constrain a later Event, differ across perspectives, or stay auditable, and where the purpose reaches (§2.2).
Carve Choosing a conceptual decomposition: a claim frame, a climb, tests of necessity, independence, and universality, completeness, routing, and a synthesis (§3.1).
Number test A part found by conceptual decomposition is opened as a process only if it carries values over time whose meaning a reader can state (§1.2). A Cut's unit must likewise be a pool that can be counted or pictured (§2.1).
Keep whole A recorded decision not to open a process for now, with its reason and the evidence it was taken with (§5.2).
Coverage report The tool's report of what an account still lacks at its declared resolution: levels of the time schedule not yet filled, processes that pass over their own Events, and turns in the values with no recorded reason (§7.2).
Time schedule The dates at each level of the tree of time at which every process is to be read (§7.2).
Five evaluation questions Coverage, reference, grounding, recomposition, and value (§5.2).
Universe and History The root Thing and its lifecycle Event, which contains the accepted Events (§2.2.1).
Context root and authority-parent A record's governing root, reached only through declared authority-parent edges: accepted-world, inner, understanding, document, or candidate (§2).
Three surfaces World (the six forms), understanding (UnderstandingNodes), and document (DocumentNodes), sharing identifiers but not authority (§4).
UnderstandingNode A question, reason, interpretation, or revision, beneath an understanding-process root (§4).
DocumentNode Text or code beneath a document root; a story passage is one role (§4).
Actor projection and operative model What an actor could know at a time, and the part of it that organizes a decision (§4.1).
Rendering route The source commit, route, viewpoint, cutoff, and disclosure policy under which a passage is written (§4.2).
Read-back A fresh-context inference of records from a passage, compared with its source (§4.2).
Grounding version The definition bundle a Realization used, addressed by its Sema content hash under a declared canonicalization scheme (§3.2).
Lens A reusable question whose readings sit under the holder's root as Cuts (§6.3).
Estimator An optional separate model that answers bounded questions with typed options (§6).
Telling phase An attributed qualitative phase of disclosure, attention, or tension in a work's text (§6.4).
Director review A review of the world before the first scene, and of the draft before release, against stated principles (§6.4).
PURE The four-gate viability check used in mechanism search: Parsimonious, Unique, Realizable, and Expansive (§6.6).
AtomicWant A construction view of one want leaf in the first Book's preparation profile (Appendix A.1).
First Book, second Book The two constructions of *The Book of Conditions*: the first built under the author's direction, with a later October edition, and the second generated by an agent on 11 October 2026 (§7.1).
MCP The Model Context Protocol, through which a language model calls the tool (§6).
Viewer The bundled read-only program that shows a model's records (§6.5).
Modeler, constructor, generator The person or AI system building the account; its role when proposing changes; the component producing candidates (§1).

References

[1] C. Brom, O. Šerý, and T. Poch, "Simulation Level of Detail for Virtual Humans," in *Intelligent Virtual Agents*, Lecture Notes in Computer Science, vol. 4722, pp. 1–14, Springer, 2007. doi: 10.1007/978-3-540-74997-4_1.

[2] B. Sunshine-Hill and N. Badler, "Perceptually Realistic Behavior through Alibi Generation," *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, vol. 6, no. 1, pp. 83–88, 2010. doi: 10.1609/aiide.v6i1.12389.

[3] Michael Polanyi. *The Tacit Dimension*. Doubleday, Garden City, NY, 1966.

[4] Edward A. Feigenbaum. The art of artificial intelligence: Themes and case studies of knowledge engineering. In *Proceedings of the Fifth International Joint Conference on Artificial Intelligence*, pages 1014–1029, 1977.

[5] Y. LeCun, "A Path Towards Autonomous Machine Intelligence," position paper, version 0.9.2, OpenReview, 27 June 2022.

[6] Alan M. Turing. Computing machinery and intelligence. *Mind*, 59(236):433–460, 1950.

[7] H. Westerberg, "Life Simulation: Learning from Worlds and Their Construction," companion working paper, 2026. Concept doi: 10.5281/zenodo.22313592.

[8] H. Westerberg, "Fractal Intelligence: Conceptual Decomposition as Problem-Solving Infrastructure," preprint, 2026. Concept doi: 10.5281/zenodo.19462645.

[9] H. Westerberg, "Sema: When the Hash Is the Word," preprint, 2026. Concept doi: 10.5281/zenodo.19462702.

[10] A. S. Rao and M. P. Georgeff, "BDI Agents: From Theory to Practice," in *Proceedings of the First International Conference on Multiagent Systems*, pp. 312–319, 1995. `https://cdn.aaai.org/ICMAS/1995/ICMAS95-042.pdf`.

[11] A. W. Kruglanski, J. Y. Shah, A. Fishbach, R. Friedman, W. Y. Chun, and D. Sleeth-Keppler, "A Theory of Goal Systems," *Advances in Experimental Social Psychology*, vol. 34, pp. 331–378, 2002. doi: 10.1016/S0065-2601(02)80008-9.

[12] H. Westerberg, "Understanding Graph: A Persistent Medium for Recursive Understanding," preprint, 2026. Concept doi: 10.5281/zenodo.19462658.

[13] H. Westerberg, "Entangled Alignment: When Safety Is the Substrate," preprint, 2026. Concept doi: 10.5281/zenodo.16440311.

[14] J. M. Zacks, N. K. Speer, K. M. Swallow, T. S. Braver, and J. R. Reynolds, "Event Perception: A Mind–Brain Perspective," *Psychological Bulletin*, vol. 133, no. 2, pp. 273–293, 2007. doi: 10.1037/0033-2909.133.2.273.

[15] R. A. Zwaan and G. A. Radvansky, "Situation Models in Language Comprehension and Memory," *Psychological Bulletin*, vol. 123, no. 2, pp. 162–185, 1998. doi: 10.1037/0033-2909.123.2.162.

[16] L. Tunstall, N. Reimers, U. E. S. Jo, L. Bates, D. Korat, M. Wasserblat, and O. Pereg, "Efficient Few-Shot Learning Without Prompts," arXiv:2209.11055 [cs.CL], 2022. doi: 10.48550/arXiv.2209.11055.

[17] E. Wenger and Y. N. Kenett, "Large Language Models Are Homogeneously Creative," *PNAS Nexus*, vol. 5, no. 3, Art. no. pgag042, 2026. doi: 10.1093/pnasnexus/pgag042.

[18] H. Westerberg, "The Ontology of the Alien: World-Diversity Search and Evolving Solution Ontologies," preprint, 2026. Concept doi: 10.5281/zenodo.18904509.

[19] P. Wack, "Scenarios: Uncharted Waters Ahead," *Harvard Business Review*, vol. 63, no. 5, September–October 1985.

[20] P. E. Tetlock and A. Belkin, eds., *Counterfactual Thought Experiments in World Politics: Logical, Methodological, and Psychological Perspectives*, Princeton University Press, 1996.

[21] J. M. Epstein and R. L. Axtell, *Growing Artificial Societies: Social Science from the Bottom Up*, Brookings Institution Press and MIT Press, 1996. doi: 10.7551/mitpress/3374.001.0001.

[22] M. D. Gurven and R. J. Davison, "Periodic catastrophes over human evolutionary history are necessary to explain the forager population paradox," *Proceedings of the National Academy of Sciences*, vol. 116, no. 26, pp. 12758–12766, 2019. doi: 10.1073/pnas.1902406116.

[23] Eleanor Rosch and Carolyn B. Mervis. Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology*, 7(4):573–605, 1975.

[24] Dedre Gentner. Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, 7(2):155–170, 1983.

[25] Robert M. Nosofsky. Attention, similarity, and the identification–categorization relationship. *Journal of Experimental Psychology: General*, 115(1):39–57, 1986.

[26] G. Shafer, *A Mathematical Theory of Evidence*. Princeton, NJ: Princeton University Press, 1976. `https://www.glennshafer.com/books/amte.html`.

[27] Charles J. Fillmore. Frame semantics and the nature of language. *Annals of the New York Academy of Sciences*, 280(1):20–32, 1976.

[28] Roger C. Schank and Robert P. Abelson. *Scripts, Plans, Goals, and Understanding: An Inquiry into Human Knowledge Structures*. Lawrence Erlbaum Associates, Hillsdale, NJ, 1977.

[29] P. Gärdenfors, *Conceptual Spaces: The Geometry of Thought*. MIT Press, 2000. doi: 10.7551/mitpress/2076.001.0001.

[30] A. N. Whitehead, *Process and Reality: An Essay in Cosmology*, corrected ed., D. R. Griffin and D. W. Sherburne, eds. Free Press, 1978.

[31] E. E. Savellos, "When Are Events Parts?," *Philosophical Papers*, vol. 29, no. 3, pp. 223–247, 2000. doi: 10.1080/05568640009485073.

[32] Pierre Grenon and Barry Smith. SNAP and SPAN: Towards dynamic spatial ontology. *Spatial Cognition and Computation*, 4(1):69–104, 2004.

[33] D. Dori, "Object-Process Analysis: Maintaining the Balance Between System Structure and Behaviour," *Journal of Logic and Computation*, vol. 5, no. 2, pp. 227–249, 1995. doi: 10.1093/logcom/5.2.227.

[34] A. Gangemi and P. Mika, "Understanding the Semantic Web through Descriptions and Situations," in *On the Move to Meaningful Internet Systems 2003: CoopIS, DOA, and ODBASE*, LNCS 2888, pp. 689–706, Springer, 2003. doi: 10.1007/978-3-540-39964-3_44.

[35] G. Nelson, *The Inform Designer's Manual*, 4th ed., 2001, sec. 24, "The World Model Described." `https://www.inform-fiction.org/manual/html/s24.html`.

[36] J. Weston, A. Bordes, S. Chopra, and T. Mikolov, "Towards AI-Complete Question Answering: A Set of Prerequisite Toy Tasks," arXiv:1502.05698, 2015. `https://arxiv.org/abs/1502.05698`.

[37] B. Dalvi, L. Huang, N. Tandon, W.-t. Yih, and P. Clark, "Tracking State Changes in Procedural Text: A Challenge Dataset and Models for Process Paragraph Comprehension," in *Proceedings of NAACL-HLT*, pp. 1595–1604, 2018. `https://aclanthology.org/N18-1144/`.

[38] Q. Zong, Y. Guo, M. Yang, Y. Guo, and Y. Song, "EvolvingWorld: An Open-Schema Framework for Co-Evolving Role-Play Agents and World Model in Interactive Literary World," arXiv:2607.17250, 2026. `https://arxiv.org/abs/2607.17250`.

[39] M. Saifullah, T. Kornmaier, T. Kazi, V. Sharma, A. S. Kanade, and A. K. Yadav, "Narrative World Model: Narratology-Grounded Writer Memory for Long-Form Fiction," arXiv:2607.05577, 2026. `https://arxiv.org/abs/2607.05577`.

[40] Kevin Yang, Dan Klein, Nanyun Peng, and Yuandong Tian. DOC: Improving long story coherence with detailed outline control. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3378–3465, 2023.

[41] Qianyue Wang, Jinwu Hu, Zhengping Li, Yufeng Wang, Daiyuan Li, Yu Hu, and Mingkui Tan. Generating long-form story using dynamic hierarchical outlining with memory-enhancement. arXiv:2412.13575, 2024.

[42] Joon Sung Park, Joseph C. O'Brien, Carrie J. Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. Generative agents: Interactive simulacra of human behavior. arXiv:2304.03442, 2023.

[43] C. S. Holling, "Understanding the Complexity of Economic, Ecological, and Social Systems," *Ecosystems*, vol. 4, pp. 390–405, 2001. doi: 10.1007/s10021-001-0101-5.

[44] Bridgewater Associates, "Ray Dalio and GMO's Jeremy Grantham on How They're Seeing the World Right Now," interview moderated by J. Haskel and A. Shahidi, May 17, 2022. Interview transcript.

[45] R. Killick, P. Fearnhead, and I. A. Eckley, "Optimal Detection of Changepoints With a Linear Computational Cost," *Journal of the American Statistical Association*, vol. 107, no. 500, pp. 1590–1598, 2012. doi: 10.1080/01621459.2012.737745.

[46] R. C. Merton, "Option Pricing When Underlying Stock Returns Are Discontinuous," *Journal of Financial Economics*, vol. 3, nos. 1–2, pp. 125–144, 1976. doi: 10.1016/0304-405X(76)90022-2.

[47] G. E. Uhlenbeck and L. S. Ornstein, "On the Theory of the Brownian Motion," *Physical Review*, vol. 36, no. 5, pp. 823–841, 1930. doi: 10.1103/PhysRev.36.823.

[48] A. G. Hawkes, *Spectra of Some Self-Exciting and Mutually Exciting Point Processes*, *Biometrika*, vol. 58, no. 1, pp. 83–90, 1971. doi: 10.1093/biomet/58.1.83.

[49] R. Goebel, R. G. Sanfelice, and A. R. Teel, "Hybrid Dynamical Systems," *IEEE Control Systems Magazine*, vol. 29, no. 2, pp. 28–93, 2009. doi: 10.1109/MCS.2008.931718.

[50] L. P. Kaelbling, M. L. Littman, and A. R. Cassandra, "Planning and Acting in Partially Observable Stochastic Domains," *Artificial Intelligence*, vol. 101, nos. 1–2, pp. 99–134, 1998. doi: 10.1016/S0004-3702(98)00023-X.

[51] P. J. Gmytrasiewicz and P. Doshi, "A Framework for Sequential Planning in Multi-Agent Settings," *Journal of Artificial Intelligence Research*, vol. 24, pp. 49–79, 2005. doi: 10.1613/jair.1579.

[52] T. Parr, G. Pezzulo, and K. J. Friston, *Active Inference: The Free Energy Principle in Mind, Brain, and Behavior*, MIT Press, 2022. Open access. doi: 10.7551/mitpress/12441.001.0001.

[53] K. J. Friston, R. Rosch, T. Parr, C. Price, and H. Bowman, "Deep Temporal Models and Active Inference," *Neuroscience & Biobehavioral Reviews*, vol. 77, pp. 388–402, 2017. doi: 10.1016/j.neubiorev.2017.04.009.

[54] M. Burgess, "Spacetimes with Semantics (III)—The Structure of Functional Knowledge Representation and Artificial Reasoning," arXiv:1608.02193, 2016, rev. 2017. doi: 10.48550/arXiv.1608.02193.

[55] M. Burgess, "A Spacetime Approach to Generalized Cognitive Reasoning in Multi-scale Learning," arXiv:1702.04638, 2017. doi: 10.48550/arXiv.1702.04638.

[56] M. Burgess, "Testing the Quantitative Spacetime Hypothesis using Artificial Narrative Comprehension (II): Establishing the Geometry of Invariant Concepts, Themes, and Namespaces," arXiv:2010.08125, 2020. doi: 10.48550/arXiv.2010.08125.

[57] Douglas B. Lenat. CYC: A large-scale investment in knowledge infrastructure. *Communications of the ACM*, 38(11):33–38, 1995.

[58] H. Westerberg, "Toward Compulsory Scene-Mediated Language Modeling: Meaningful Substrates for Every Next Word," preprint, 2026. Concept doi: 10.5281/zenodo.21910308.

[59] H. Westerberg, "Imagining the Corpus: Turning Text into Video for Language Model Training," preprint, 2026. Concept doi: 10.5281/zenodo.22104891.

[60] D. Kahneman, *Attention and Effort*. Englewood Cliffs, NJ: Prentice-Hall, 1973.

[61] S. E. Hobfoll, "Conservation of Resources: A New Attempt at Conceptualizing Stress," *American Psychologist*, vol. 44, no. 3, pp. 513–524, 1989. doi: 10.1037/0003-066X.44.3.513.

[62] R. M. Emerson, "Social Exchange Theory," *Annual Review of Sociology*, vol. 2, pp. 335–362, 1976. doi: 10.1146/annurev.so.02.080176.002003.

[63] G. W. Brier, "Verification of Forecasts Expressed in Terms of Probability," *Monthly Weather Review*, vol. 78, no. 1, pp. 1–3, 1950. doi: 10.1175/1520-0493(1950)078<0001:VOFEIT>2.0.CO;2.