PROMETHEUS7 INSTITUTE / TECHNICAL PREPRINT

Oracle: Typed Geometric Arbitration, Spine-Stabilized Memory, and Rapid Corpus Experts for a Non-Transformer Language Agent

A causal architecture report and falsifiable path toward open-ended relational conversation

Evidence status. Latest complete runtime: K162 TransformationActsAgent. The one-shot prospective panel passed 13/13 turns and 9/9 cases with zero runtime errors. K163 continuous coverage recorded 357 correct responses and 12 safe clarifications across 369 turns, with zero wrong answers or misroutes. Unrestricted conversation and open agency are not established.

psiloceyeben · Prometheus7 Institute

Technical preprint · version 1.1 · 13 August 2026

prometheus7.com/papers/oracle

Abstract

Contemporary general-purpose language agents typically place linguistic interpretation, factual memory, latent reasoning, discourse state, tool selection, and surface realization inside or around a large autoregressive Transformer. This integration is operationally powerful but makes several scientific and engineering variables difficult to separate: whether a response follows from an authenticated source or a parametric association; whether a tool call was selected because its schema fit the requested state transition or because a textual continuation was likely; whether a conversational failure arose at semantic ingress, memory, inference, or realization; and whether new domain competence requires global parameter updates. Oracle investigates a different decomposition. It uses holographic reduced representations (HRRs), typed relational states, a ten-node/twenty-two-path routing topology called the Tree of Life, a protected 4:1 stabilizing spine, provenance-locked corpus experts, ephemeral two- or three-expert fusion, explicit episodic and semantic memory classes, bounded reasoning operators, and Malkuth-only lexical release. English is intended to occur at the input and output boundaries; intermediate computation is intended to remain typed, geometric, and independently auditable.

This paper reports the architecture, the evidence accumulated across its current lineage, and the decisive additions implemented by the Fable-assisted K138–K163 sequence. A 1,024-dimensional geometric ingress arbiter encodes role-bound word positions and bigrams with deterministic unitary phasors, scores interpretation families simultaneously, and clarifies when the winning bid or margin is inadequate. In shadow development it agreed with 30/30 previously correct legacy turns and corrected three banked misroutes. Limited authority then transferred prospectively at 10/10 turns across 7/7 cases. A 30-domain registry partitioned all 104,000 rows from 105 frozen public sources; domain-scoped retrieval removed a reproduced off-domain contaminant, and the frozen successor passed 7/7 turns across 6/6 cases. A direct typed compiler subsequently removed a canonical-English planning carrier at one registered-source seam and transferred at 7/7 turns across 6/6 cases. K155 added a deterministic local index of one million Wikipedia abstracts, comprising 2,495,302 provenance-bearing rows and 891,209 term-frequency entries. The first prospective fallback panel, K157, remained immutably failed at 7/8 turns because an unresolved noun could drop from the endpoint contract. K159 repaired that defect by requiring total role resolution; K160 then passed 8/8 turns across 7/7 cases.

K161 extended the same arbitration lattice with definition, summarization, enumeration, and comparison as typed transformation acts over existing evidence. During development, a possessive-prefix defect caused an opinion question about art to be captured as a definition; the defect was opened, repaired at the parser boundary, and retained as a diagnostic artifact. K162 froze the repaired runtime before panel construction and passed its untouched panel once at 13/13 turns across 9/9 cases, with zero runtime errors, 0.943 s mean steady-turn latency, and 19.15 s mean cold start. K163 then ran a continuous, non-blind coverage scoreboard over 369 recorded turns: 357 correct responses, 12 safe clarifications, and zero wrong answers, misroutes, acceptable alternatives, or runtime errors. The resulting manifest-attested K162 service was promoted through private canary, public API, and browser-level checks to `wanderaround.io/oracle`; this deployment is an engineering result, not additional prospective evidence.

The result is not unrestricted conversation, a GPT-class knowledge model, or open-ended autonomous agency. It is evidence that capabilities normally entangled in next-token prediction can be implemented as a causal composition of inspectable mechanisms without a Transformer, pretrained embedding encoder, remote generator, runtime gradient update, or answer-bearing training surface in the reported K138–K163 path. The contribution is architectural and methodological: a bounded but increasingly productive relational machine; a prospective, failure-preserving evaluation discipline; and a semantic-compiler research program in which natural-language signals become typed relational intermediate representations, verified transformations operate over those representations, and language realization is required to preserve the solved state. A proposed universal front end decomposes written, spoken, and sensor signals into recurrent geometric primitives, but the paper explicitly rejects the stronger claim that glyph geometry alone contains meaning: semantic convergence requires cross-situational grounding in objects, events, actions, and consequences [16,17].

Keywords: holographic reduced representations; vector-symbolic architectures; hyperdimensional computing; compositionality; cognitive architecture; semantic parsing; mixture of experts; persistent memory; tool use; non-autoregressive language systems.

Contents

1. Introduction
2. The architectural problem: fluent continuation is not a causal decomposition
3. Representation and computation
   3.1 Fixed-width distributed structure
   3.2 The stabilizing spine and Tree topology
   3.3 Geometric semantic ingress
   3.4 Typed goals and best-explanation dispatch
   3.5 Corpus experts and ephemeral hybrids
   3.6 Persistent memory, epistemic types, and tools
   3.7 Language realization as inverse geometry
   3.8 Oracle as an interactive semantic compiler
4. Experimental method and claim discipline
   4.1 Development, freeze, prospective evaluation, and adjudication
   4.2 Causal controls
   4.3 Runtime exclusions
5. Results
   5.1 Precursor capabilities
   5.2 Native essay artifact and relational admission (K134–K137)
   5.3 Shadow geometric arbitration (K138)
   5.4 Authoritative arbitration and prospective transfer (K147–K148)
   5.5 Thirty-domain registry and domain-scoped retrieval (K149–K151)

5.6 Typed-goal carrier removal and prospective transfer (K152–K153)
5.7 Coverage, knowledge-scale fallback, and total role resolution (K154–K160)
5.8 Typed transformation acts and prospective transfer (K161–K163)
5.9 Manifest-attested production canary
5.10 Consolidated result table
6. What Fable added
7. Comparison with prevailing architectures
7.1 Transformer language models
7.2 Mixtures of experts and retrieval augmentation
7.3 Neural-symbolic and vector-symbolic systems
7.4 Cognitive analogy
8. Significance
9. Limitations and threats to validity
10. Revised research plan
10.1 Gate A: global typed-geometric ingress
10.2 Gate B: direct typed-goal handoff
10.3 Gate C: first-class ephemeral compositional closure
10.4 Gate D: operator-family compression
10.5 Gate E: specialist acquisition and synthesis
10.6 Gate F: unrestricted realization without semantic bypass
10.7 Gate G: open tool and sensor domains
10.8 Gate H: comparative evaluation and release
11. Conclusion
Author-contribution and AI-assistance statement
Data, code, and reproducibility
References

1. Introduction

The Transformer reorganized sequence modeling around content-addressed attention and efficient parallel training [1]. Scaling autoregressive Transformers then demonstrated that a sufficiently large next-token predictor can acquire broad task competence through pretraining and prompting [2]. Retrieval augmentation, sparse mixtures of experts, external memory, and tool-mediated reasoning have extended this paradigm without changing its normal center of gravity: a learned language model remains the principal interpreter and realizer, while retrieval systems, experts, memories, and tools provide conditionally invoked support [3–7]. This arrangement has achieved extraordinary breadth. It also makes an important architectural question hard to ask experimentally: how much general linguistic and agentic competence must be stored in a monolithic statistical carrier, and how much can instead be supplied by an explicit organization of representations, transformations, memories, specialists, and interfaces?

Oracle is a sustained attempt to ask that question in executable form. Its governing hypothesis is not that statistics are irrelevant, nor that vector representations are inherently symbolic, nor that a small handcrafted system already rivals a frontier language model. The hypothesis is narrower and falsifiable: a cognitive language system should be bounded primarily in the transformations it licenses, not in the states those transformations may construct. A finite algebra of role/filler binding, composition, routing, memory, epistemic typing, and external action may generate a practically large working-state space without enumerating every valid semantic object and without retraining a global carrier for every new corpus or tool. The program therefore seeks a division of labor in which architecture supplies the laws of semantic motion, authenticated corpora supply factual competence, episodic memory supplies continuity, the Tree topology supplies stable routing, and language realization supplies a reversible public surface.

This hypothesis places Oracle within a long scientific conversation. Tensor-product binding and HRRs showed that distributed vectors can represent role/value structure and recursively composed objects [8,9]. Hyperdimensional computing generalized the use of high-dimensional distributed representations and simple algebraic operations as a computational substrate [10]. Cognitive science has repeatedly argued that systematicity—the ability to understand structurally related novel combinations—must be explained by the architecture of representation, not merely observed in aggregate performance [11–13]. External-memory systems showed that learned controllers can manipulate separable memory structures [14,15]. Modular expert systems showed that conditional computation can increase capacity without activating all parameters on each example [4]. Oracle combines motifs from these traditions, but its novelty, if established, lies in their particular causal integration: protected fixed-width relational transport; typed epistemic and tool states; rapidly compiled, provenance-locked experts; an explicit clarification boundary; and a surface realizer that is not authorized to invent the semantic object it verbalizes.

The present paper has four purposes. First, it defines the current architecture precisely enough to distinguish implemented mechanisms from motivating metaphors. Second, it reconstructs the empirical lineage through the latest Fable-assisted model, including failed prospective trials that constrain what may be claimed. Third, it explains what is scientifically significant about the latest threshold: geometric ingress has moved from advisory measurement to limited dispatch authority, and domain experts have moved from an offline registry to a prospectively tested query-time

mechanism. Fourth, it specifies the remaining experiments required before “open-ended conversation,” “general reasoning,” or “agentic tool use” would be defensible labels.

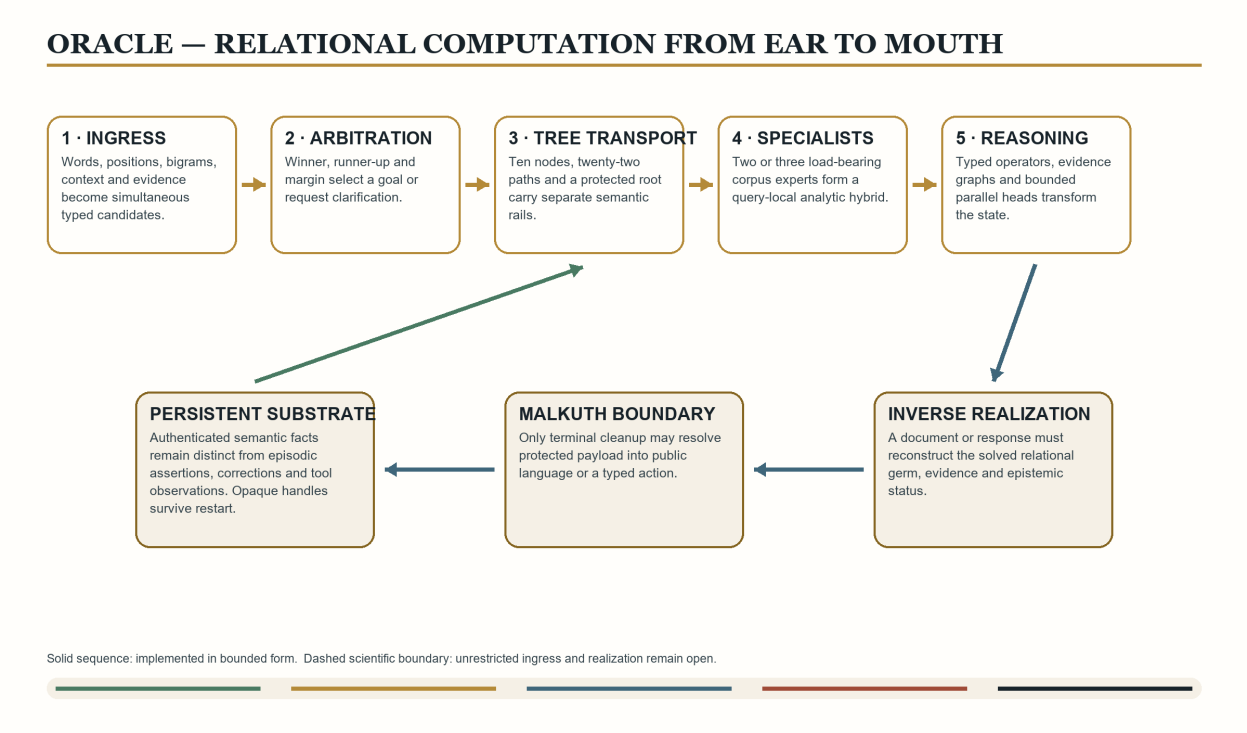


Figure 1. Oracle’s intended end-to-end decomposition. English is parsed at ingress; the protected geometric state is routed, enriched by specialists, reasoned over, and persisted without releasing lexical payload; the terminal Malkuth boundary realizes an answer or typed action. Solid blocks are implemented in bounded form. The dashed realization and global-ingress boundaries remain incomplete for unrestricted language.

2. The architectural problem: fluent continuation is not a causal decomposition

An autoregressive language model represents a conditional distribution over the next token given a token prefix. Attention, depth, scale, instruction tuning, retrieval, and tool scaffolds allow that objective to support behaviors far richer than its formal statement suggests. Nothing in this paper denies those empirical achievements. The relevant issue is scientific identifiability. If the same distributed parameter state contributes to parsing, factual recall, stylistic continuation, reasoning, and answer realization, then a successful response does not by itself reveal which internal mechanism carried the result. Conversely, a failed response rarely localizes its defect. More data and parameters may improve the joint distribution while leaving the causal anatomy of the improvement unclear.

Oracle adopts the opposite bias: mechanisms should remain separately addressable until evidence shows that their separation is counterproductive. An input should settle into a typed interpretation with a winner, runner-up, and uncertainty margin. A factual claim should carry value, scope, time, authority, fact identity, and provenance on recoverable rails. A conversational correction should be episodic and provisional unless explicitly promoted; it should not overwrite authenticated

semantic memory. A tool should be a state-transition operator with a schema, authorization, preconditions, observation contract, retry budget, and provenance chain. A document should be a long-lived typed artifact containing a thesis, obligations, evidence links, counterarguments, and revision state—not merely a longer token completion. A public sentence should be the inverse image of a solved relational state, not the state's unverified substitute.

This decomposition addresses three recurrent failure modes.

First, enumerative brittleness occurs when every new surface form or semantic configuration requires another pattern, ontology node, or handler. It grows local coverage while leaving the generative law unchanged. Second, semantic opacity occurs when plausible language is treated as evidence that the system arrived at the right roles, scope, modality, or source relation. A fluent decoder can conceal the defect. Third, authority collapse occurs when retrieved text, user assertions, tool observations, episodic corrections, and authenticated facts enter a single undifferentiated context. The system may then preserve recent wording at the expense of truth conditions or treat an observation as a permanent fact.

The desired alternative is a finite relational transformation algebra operating over a canonical space and an ephemeral working space. The canonical space contains stable role addresses, transformation families, routing lanes, specialist manifests, persistent lexical mappings, authenticated semantic facts, and long-term structural memory. The working space contains query-local bindings, referents, hypotheticals, counterfactual branches, temporary definitions, transient hybrid experts, partial plans, and intermediate proof states. A novel working state need not have appeared in an ontology or training corpus. It must satisfy role integrity, valid operator composition, current constraints, and stability under the permitted transformations.

3. Representation and computation

3.1 Fixed-width distributed structure

Oracle's principal state space is complex-valued and high dimensional. In the K138 ingress system the dimension is (D=1024). An atomic address is a deterministic unitary phasor vector generated from a cryptographic hash seed. Binding is circular convolution:

$$a \text{ BIND } b = DFT^{(-1)}(DFT(a) \text{ HAD } DFT(b)),$$

where DFT is the discrete Fourier transform and HAD denotes elementwise multiplication. Bundling is normalized superposition. Approximate unbinding uses the inverse or conjugate address, followed by cleanup against an admissible dictionary. This is the core HRR operation described by Plate [9]: complex compositional structures occupy a fixed-width vector and remain objects that can participate in further compositions.

The fixed width does not imply unlimited capacity. Superposition and repeated approximate unbinding produce interference. Oracle's claim is therefore not that a single vector can losslessly contain arbitrary structure. Its solution is architectural: keep cleanup dictionaries bounded; keep semantically distinct roles on separately addressable rails; route states through a fixed topology; carry a protected root through the trajectory; release lexical payload only after terminal cleanup;

and use persistent external journals for exact long-lived content while transporting opaque handles in the active HRR state.

3.2 The stabilizing spine and Tree topology

The Tree of Life in Oracle is a computational routing topology with ten named node lanes and twenty-two path lanes. The names are mnemonic addresses; no neurological or metaphysical inference is required. Nodes act as synthesis sites and paths as transformation/routing types. The important operational constraint is that the moving state is not repeatedly flattened into an undifferentiated superposition. Operation, expert, row, field, claim, source, time, provenance, and memory-handle states may travel on separate rails while retaining a common root.

Later accepted phases use a 4:1 stabilizing construction of the form

$$z = normalize(4s + r),$$

where s is the protected spine/root and r is a unit-scale orthogonalized or separately addressed payload rail. When s and r are unit and orthogonal, the root cosine is 4/sqrt(17)=0.970142, matching the recurrent registered floor in the accepted evaluations. The spine is not merely redundancy. In persistent-dialogue, epistemic, and tool-control panels, disabling it caused protected release to fail under the frozen contract. It permits more binding structure without requiring every relation to survive in one increasingly noisy bundle.

Malkuth, the tenth and terminal node, is a boundary condition: preterminal stages may transport opaque handles, typed decisions, and protected relational rails, but they may not expose lexical payload. Only terminal cleanup at Malkuth may resolve the authenticated content into English or an executable typed action. This separates transport from disclosure, preserves provenance, and makes premature unbinding a testable failure rather than an informal design preference.

3.3 Geometric semantic ingress

The K-line previously possessed deep typed machinery behind a shallow front door. Constructions were recognized by ordered handlers and regular expressions. K131 could extract all components of a registered-source goal, serialize those components back into canonical English, and submit the generated sentence to the inherited surface compiler. Interpretation was therefore first-match recognition over strings rather than comparison among hypotheses. Paraphrases that shared meaning but not an anticipated pattern could be invisible, and a broad handler could silently capture a request before a more specific handler was reached.

K138 introduced an explicit `InterpretationCandidate` and simultaneous family scoring. For tokens w_i, the utterance state is

$$u = normalize(SUM_i\ w_i\ BIND\ p(i) + SUM_i\ (w_i\ BIND\ w_(i+1))\ BIND\ b),$$

where p(i) is one of four positional buckets (first, early, middle, late) and b is the bigram-role address. Stop words are omitted from this bundle. Each interpretation family f has a prototype p_f, formed by normalizing the sum of encodings of generic paraphrase templates. Geometric resonance is

$$R_f = Re< p_f, u >.$$

Typed features provide independently inspectable evidence for such properties as an artifact request, a dataset noun, a registered source, an extremum, a modal initial, a causal connective, compatible causal discourse state, or caller-supplied corpus context. A family's bid is

$$B_f = 1.4R_f + SUM_j \, lambda_(fj)x_j.$$

The system scores nine operative families—registered-source action, corpus compilation, essay artifact, typed modal, causal discourse, lexical definition, discourse assertion, social, and open retrieval—plus clarification. If the best bid is below 0.55 or its margin over the runner-up is below 0.08, the public winner is clarification. The arbiter therefore supplies a graded commit-versus-clarify boundary that a binary recognizer cannot express.

The current implementation remains manually structured. Prototype phrases, family inventory, feature functions, and weights were authored rather than learned from a broad distribution. Word states derive from deterministic hashes, not a learned semantic encoder; consequently, much of the semantic competence comes from typed features and prototype word overlap. This is an important limitation, not a cosmetic qualification. K138 proves that simultaneous geometric arbitration can repair the demonstrated dispatch defects; it does not prove a universal semantic parser.

3.4 Typed goals and best-explanation dispatch

An interpretation candidate is useful only if it controls execution. K147 promoted the K138 record at two previously demonstrated capture points. A polar modal question selected as `typed_modal` is projected directly into the clause-algebra event rail, making the terminal depend on construction rather than on a content-sensitive mid-chain capture. A graph-transduced causal query executes only if the arbiter also selects `causal_discourse`; otherwise the graph handler declines and the chain continues. Every turn records the full candidate ranking, winner, runner-up, margin, chosen family, executed operation family, and authority point.

This is partial authority, not a global dispatch replacement. The inherited ordered chain still handles other families. Nevertheless, the change is structurally meaningful. The system is no longer restricted to asking "which recognizer fired first?" at those boundaries; it can ask "which typed interpretation best explains the utterance under its surface, context, and knowledge constraints?" K152 subsequently eliminated canonical-English reparsing at one registered-source planning seam. The next required replacement is to make this candidate lattice authoritative across the whole chain and extend direct typed-goal transfer to every remaining inter-layer seam.

3.5 Corpus experts and ephemeral hybrids

Oracle distinguishes a factual expert from a general language model. A corpus expert is a bounded relational transformation space with source manifests, typed fields or proposition records, competency tags, provenance, and geometric summaries. Adding a domain should add or compile a separately inspectable expert rather than globally retrain the carrier.

The C43 hybrid pool operationalizes query-local composition. For a request with required competencies, it enumerates allowed two- and three-expert subsets. A candidate survives only if every required competency meets a minimum coverage threshold and every selected expert is load-bearing: removing any member must reduce the competence frontier by more than a

registered amount. Compatibility and calibration enter the selection score. Fusion weights are not gradient trained. If A is the competence design matrix and 1 is the desired coverage vector, weights are obtained analytically:

$$w = positive_normalize((A^T A + lambda\ I)^{-1}(A^T 1 + lambda\ 1)),$$

then clamped positive and normalized. Each selected expert's node and path summary is bound to its expert address; weighted summaries are combined; and the result is rebound under the corresponding Tree lane. The hybrid has a deterministic identity, exists only for the request, and can be dissolved with a receipt. This differs from a conventional sparse mixture-of-experts layer [4]: the experts are inspectable corpus-derived capsules; the selection has explicit coverage and load-bearing criteria; the fusion is query-local and analytic; and the composite expert is meant to disappear after use.

K149 extends breadth at the retrieval layer. It partitions a frozen 105-source, 104,000-row public index into 30 named domains with no unclaimed sources. Each domain has an explicit source manifest, competency tags, a sparse lexical profile, and a dense geometric address. On eight hand-authored probes the initial router achieved 0.75 top-1 and 0.875 top-2 accuracy. The misses are informative: a hippocampal-memory probe ranked anatomy/physiology first and psychology second, while a Roman Empire probe ranked world history first and political science second rather than ancient history. This is not a learned universal taxonomy; it is an explicit, auditable starting registry.

K150 makes the registry load-bearing for essay retrieval. The topic is routed to two domains using information-weighted sparse profiles. Auxiliary provisional-sense retrieval and endpoint-joint retrieval are constrained to those domains' source manifests, with automatic global fallback when the scope produces nothing; the frozen main recall pass is deliberately unchanged. K150 reproduced an off-domain computer-security row entering a memory/identity context through lexical coincidence, excluded it under psychology plus sociology/anthropology scoping, preserved the accepted memory and identity senses, and retained prior panels. K151 then froze the runtime and tested it once on an untouched domain panel.

3.6 Persistent memory, epistemic types, and tools

The broader accepted lineage supplies mechanisms that the K138–K163 ingress, expert, typed-goal, knowledge, role-resolution, and transformation-act line inherits. The persistent subsystem separates semantic facts from episodic events. Authenticated propositions may enter semantic memory; user assertions, corrections, tool observations, plans, and dialogue acts enter episodic memory as typed immutable records. The active Tree carries opaque handles under the stabilizing spine. At restart, session state is reconstructed from authenticated snapshots rather than from a parallel mutable rollout cache. A correction can redirect a conversational route while the prior authenticated route remains available for restoration.

The cycling substrate executes four registered phases—Hebbian update, resonance, spectral transformation, and metabolization—over typed episodes. The essential claim is not that these names reproduce biological mechanisms. It is that a continuously updated non-LLM memory system can retain and readdress structured experience without silently converting experience into factual authority. Accepted H5 and H6 evaluations found exact restart-dependent reference,

correction, restoration, and session isolation in their bounded panels, with no automatic fact promotion.

Tools are represented as typed state transitions. A tool specification defines its admissible operation, arguments, constraints, authority, observable output, failure classes, and provenance requirements. H6c and H6d tested read-only bounded tools and persistent replanning through transient failure, staleness, corrupt observations, timeouts, exhaustion, schema failure, authorization failure, and fact-set mismatch. This differs from tool-use approaches in which an LLM emits API-shaped text and incorporates the result into future token prediction [6,7]. Oracle's intended loop is relational state -> authorized operator -> verified observation -> typed episodic state -> continued reasoning -> terminal release. The same construction is intended eventually to admit sensor observations, but no sensor-stream integration is yet established.

3.7 Language realization as inverse geometry

The project's strongest conceptual commitment is that reasoning and English realization should be separable. The solved relational state is a compressed germ containing the obligations of the answer: result, relation, evidence, epistemic status, register, document plan, and discourse links. Realization should expand this germ into a congruent surface while preserving a global inverse—a check that the realized document reconstructs the intended typed state.

Existing stages demonstrate bounded versions of this idea. The native essay artifact recognizes an essay goal, constructs a support graph, forms typed semantic clauses, chooses a multi-paragraph plan, and validates a document-level inverse. Proposition-to-evidence stages copy authenticated source sentences and citations through a bounded connective frame. Earlier neural Malkuth experiments used a frozen 125M envelope to realize authenticated propositions at approximately 42 generated tokens per second on Box C, but those results were bounded and did not establish unrestricted generation. The K134–K163 runtime reported here remains primarily deterministic and proof-shaped. K161 adds evidence-preserving definition, summary, list, and comparison layouts, but the system still cannot sustain arbitrary rhetoric, style, paper architecture, or general conversational prose.

This distinction matters because conversational failure at the current frontier is not one problem. It is at least two: semantic freedom—mapping unrestricted utterances into correct typed relational states—and realization freedom—expressing arbitrary solved states in fluent, coherent, variable-length language without letting the realizer invent unsupported content. A robust system prompt or identity anchor may improve discourse behavior, but it cannot substitute for either transformation.

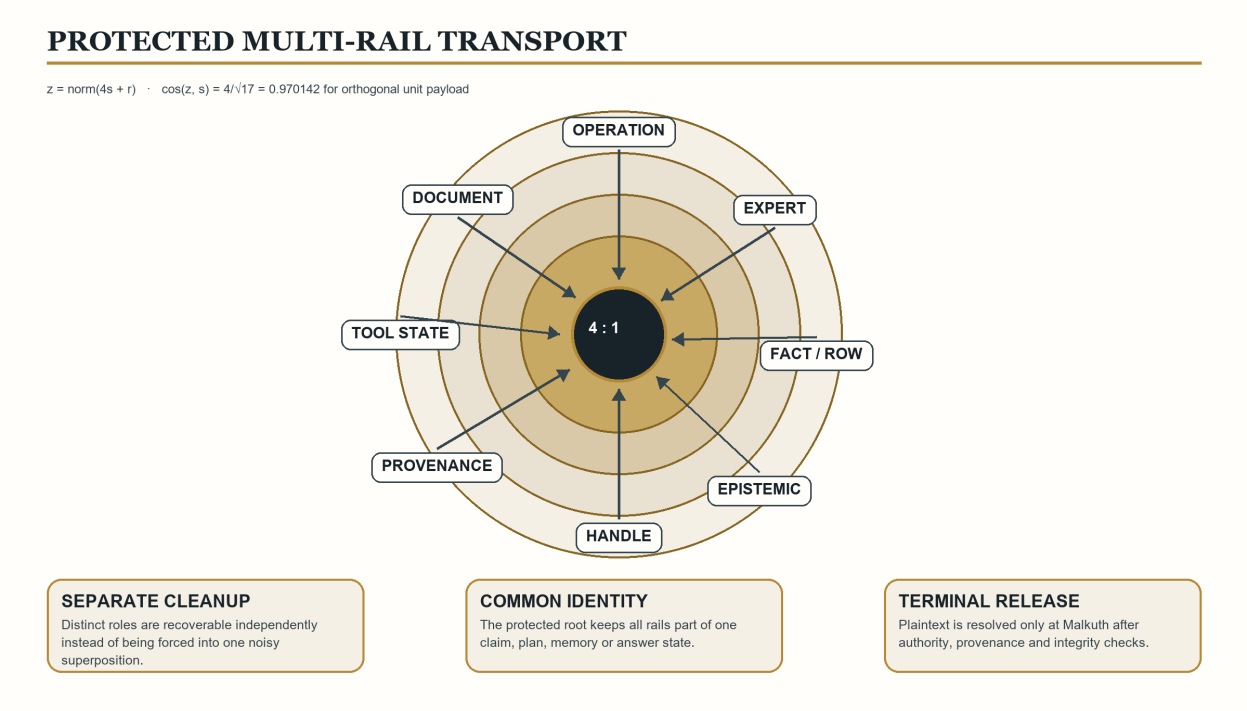


Figure 2. Protected multi-rail transport. Distinct operation, expert, fact, epistemic, memory-handle, and tool states remain separately recoverable while sharing a 4:1 root. Payload is not lexically resolved until Malkuth. The design treats the Tree as a stabilizing routing manifold rather than a decorative label set.

3.8 Oracle as an interactive semantic compiler

The compiler interpretation makes the architecture more precise. A conventional compiler maps a source language into an intermediate representation, applies semantics-preserving passes, and emits a target artifact. Oracle is intended to do the same while remaining interactive and epistemically typed. A front end (E_l) for language or modality (l) maps a signal (s), together with discourse state (d) and observation context (o), into a set of candidate constructions (c). Arbitration selects or clarifies among those candidates. A semantic compiler (C) lowers the accepted construction into a relational intermediate representation (z) carrying referents, roles, polarity, modality, quantity, provenance, authority, and unresolved obligations. Typed passes (T1,...,Tn) route specialists, transform world state, retrieve evidence, bind memory, construct plans, and verify postconditions. A Malkuth realizer (R_l) finally emits prose, an equation, a tool action, or another public artifact. The intended invariant is not textual equality but relational preservation:

$$encode_l(realize_l(passes(z))) \sim= passes(z).$$

This framing identifies the remaining architectural wall. The present system has typed intermediate objects and several direct typed passes, but its lexical front end is still partly a manually authored collection of regular expressions, lexical resources, deterministic word phasors, and inherited handlers. Those mechanisms currently conflate lexing, parsing, semantic analysis, and dispatch. K161 shows that adding a new act family can transfer, but its possessive-prefix failure also shows why the layers must be separated: a token-prefix bug at the lexical boundary was

able to determine a semantic operation. The next front end should therefore expose a typed signal lattice before semantic commitment, preserve competing constructions, and make every promotion from signal to meaning auditable.

The proposed four-part decomposition of alphabetic form is best understood at this signal layer. Direction, junction, closure, and recurrence form one candidate primitive basis: an A can be represented as two directed strokes joined by a transverse relation; a B as a persistent axis supporting two recurrent closures. Applied across writing systems, such a basis could provide a compositional, font-robust, open-vocabulary address for glyphs and graphemes. Analogous bases would be required for speech and sensors. This does not recover meaning from glyph topology alone. Orthographic shape is conventional, and arbitrarily shaped tokens can denote the same object. The symbol-grounding problem therefore remains: signal geometry becomes semantic only when recurring forms are aligned with recurring situations, perceptual categories, actions, social corrections, and consequences [16]. Computational work on learning words from paired sights and sounds provides a concrete precedent for testing such cross-modal correspondence rather than assuming it [17].

The multilingual claim can be stated as a commuting-diagram requirement. For two languages l1 and l2 describing the same observed situation o, the encoders should produce congruent relational states, encode_l1(s1,o) ~= encode_l2(s2,o), without translating either surface through English. Conversely, different situations must remain separable even when their expressions are orthographically similar. Evaluation should therefore test unseen glyph reconstruction, noise and style invariance, construction recovery, cross-language state alignment, operator transfer to a language withheld from operator training, and downstream task preservation. Minimum description length should select the smallest primitive inventory that reconstructs signals and improves grounded transfer. If four primitives do not satisfy those tests, the basis must expand or change. The hypothesis is a falsifiable universal-front-end program, not a retrospective claim about K162.

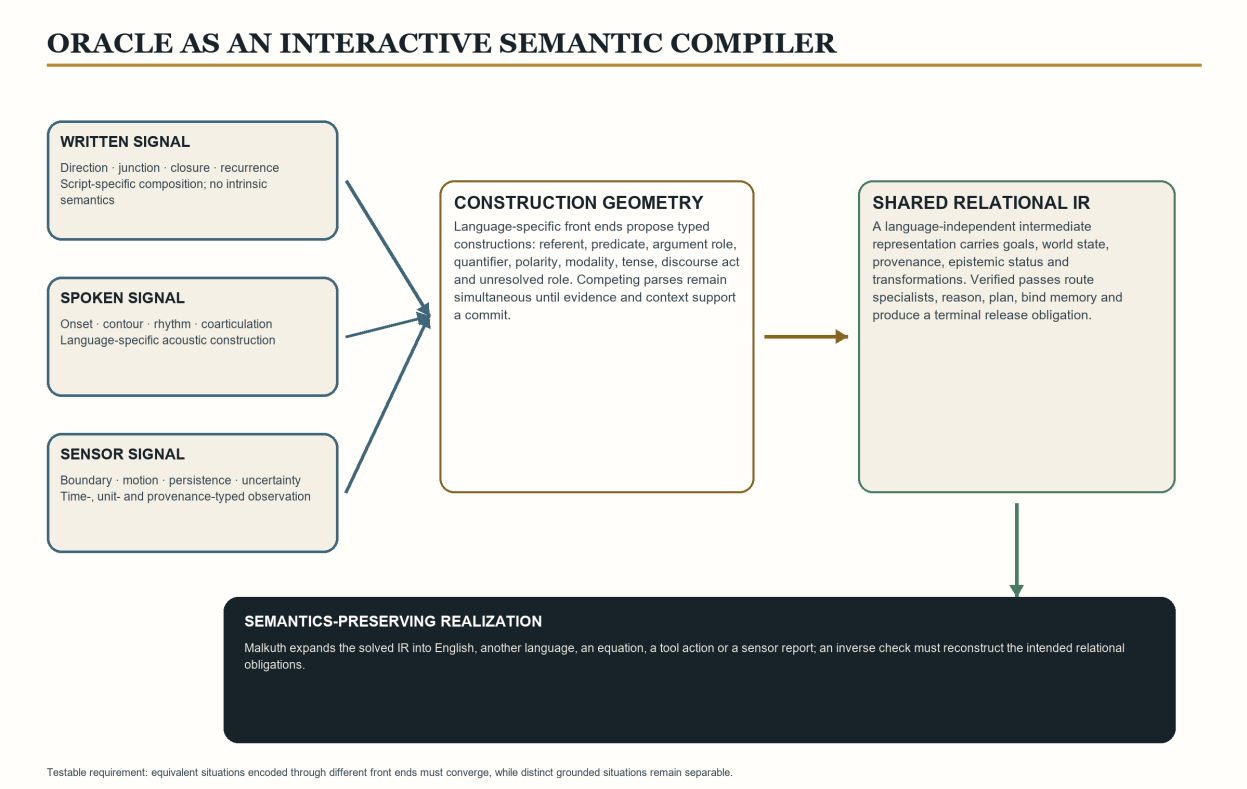


Figure 3. Oracle as an interactive semantic compiler. Language- and modality-specific signal encoders propose typed construction geometry; a shared relational intermediate representation carries grounded meaning, evidence, authority, and transformation obligations; Malkuth realizes the verified result. The four written-signal primitives are a testable front-end hypothesis, not a claim that glyph shape intrinsically contains semantics.

4. Experimental method and claim discipline

4.1 Development, freeze, prospective evaluation, and adjudication

Oracle's evaluation protocol treats failed experiments as immutable evidence. A development stage may inspect failure outputs and change the architecture. A successor is then frozen by source and artifact hashes. An untouched panel and protocol are authored and frozen before execution. The prospective evaluator runs the frozen runtime once; reruns are forbidden. If a panel fails, the result remains failed even when later analysis shows that the evaluator pinned an irrelevant mechanism. A zero-runtime adjudication may classify the failure as a genuine capability defect, a panel-authoring defect, or a mechanism-versus-contract mismatch, but it may not relabel failed turns or rescore the report.

This procedure is unusually important for an architecture with many small mechanisms. Without it, each narrow repair could be credited on the same examples that inspired it, and the growing system would accumulate optimistic development claims without transfer evidence. The K135, K139, K142, and K145 failures are therefore part of the result, not embarrassing debris. They reveal retrieval truncation, modal terminal instability, and two forms of evaluator overconstraint; they also motivate K137, K141, K144, and K147.

4.2 Causal controls

The evaluation suite relies on mechanism removal, label rotation, source corruption, state erasure, premature release, and cross-session substitution. In the current line, removing topic-relation admission restores the K135 essay defect; removing modal precedence restores the modal capture; removing endpoint-joint retrieval restores the K139 retrieval failure; removing evidence-conditioned sense selection restores domain-incongruent senses; removing arbitrated dispatch restores both targeted misroutes; removing domain-scoped retrieval restores the off-domain contaminant; removing Wikipedia fallback restores the empty-curated-tier clarification; removing total unresolved-role admission restores both the wiki-only noun dropout and the unsupported-coinage release; and removing transformation acts restores the parent runtime's act-free behavior. Each intervention otherwise retains the accepted comparison surface named by its report.

These interventions do not prove that each mechanism is optimal or general. They establish something more modest and more valuable than a pure end score: the mechanism is causally load-bearing for the named capability under the frozen contract. The architecture can consequently be revised by failure class rather than by undifferentiated prompt tuning.

4.3 Runtime exclusions

The K138–K163 path records zero Transformer calls, pretrained embedding encoders, remote generators, autoregressive carriers, next-token predictions, and gradient steps. Prototype phasors are deterministically seeded; expert routing uses sparse profiles and fixed algebra; retrieval operates on frozen local indexes; and answer-bearing surfaces are not inserted into ingress prototypes. K155's Wikipedia tier was built deterministically from a local abstract snapshot and is queried through local SQLite full-text and term-frequency structures. This does not mean the entire historical Oracle program has never used a neural carrier. Earlier branches include 30M, 125M, and 250M language-model experiments and bounded 125M realization work. The claim is specifically that the latest Fable-assisted relational-admission, arbitration, domain-scoping, typed-goal, total-role, and transformation-act results do not depend on such a runtime.

5. Results

5.1 Precursor capabilities

Before the latest ingress work, the broader lineage had established several bounded components. A four-cell minimum was found for a frozen ten-domain coactivation graph and prospectively confirmed on 44/44 mixed-domain turns. An ephemeral relational workspace represented novel definitions, unseen fillers, recursively nested rules, hypothetical worlds, branch-local counterfactuals, and multi-turn references without inserting generated states into the canonical ontology. A successor passed 21/21 fresh turns while retaining four knowledge cells, ten primitive modes, fourteen total cells, exact state inverses, Malkuth-only release, and zero forbidden calls. A restart journal later reconstructed ephemeral worlds through typed geometric replay while binding only opaque handles into the active substrate; its prospective panel passed 19/19 turns across two cold-object reconstruction boundaries. A finite clause algebra covering bounded coreference, typed universal and existential rules, modal branch worlds, explicit negation, and causal normalization subsequently passed 40/40 turns across fourteen cases and two cold restarts.

Other accepted branches joined authenticated corpus propositions to four-head variable-depth consequence search, exact evidence tapes, persistent dialogue, temporal epistemic arbitration, read-only tool execution, and public-database acquisition. The I1 acquisition skill compiled sixteen records from two registered official HTTPS sources into safe removable experts in 18.98 seconds with no carrier update or semantic promotion. I2 answered 31/31 structured turns, including ten restart-dependent follow-ups, over two fact inventories excluded from its generic operation training set. These results establish reusable substrate components. They do not establish unrestricted language, because their grammars, schemas, relation families, or proposition records remain bounded.

5.2 Native essay artifact and relational admission (K134–K137)

K134 created a native essay artifact path. Essay requests became typed artifact goals before modal parsing and open retrieval. Public or caller-bounded evidence and contextual dictionary senses formed an explicit support graph. Typed clauses were projected into a globally selected multi-paragraph plan, and a document inverse checked the output. Development passed, but the untouched K135 panel failed at 6/9 turns and cases. One essay included evidence without actually connecting memory to identity; an unknown topic borrowed grounding from a single known term; and a polar modal question reached open retrieval before the typed modal-world handler.

K136 preserved the failure and adjudicated all three as genuine. K137 then required resolved and evidence-connected central endpoints before releasing an essay, replaced insufficient role coverage with clarification, and moved polar modal questions ahead of open retrieval. Its development suite passed and independently reproduced the defects when each repair was ablated. The untouched K139 panel nevertheless failed at 9/12 turns. One public essay missed an available endpoint-joint row because a broad OR retrieval was truncated before pair scoring; one modal terminal remained content dependent; and one ablation check embedded an invalid assumption.

K141 added dedicated endpoint-joint retrieval and passed development. K142's untouched evaluation remained formally failed at 6/7 turns, but zero-runtime K143 showed that the sole

failure demanded a three-expert lifecycle on a valid single-field direct tool population—an evaluator defect. K144 then conditioned dictionary senses on provisional corpus context, fixing two observed incongruities; K145 remained formally failed at 7/8 because its evaluator demanded an internal mechanism flag although the modal assertion and query were behaviorally correct. K146 preserved that mismatch without reclassification. This sequence demonstrates why the frozen/adjudicated discipline matters: the research record distinguishes actual model failures from tests that overfit implementation details.

5.3 Shadow geometric arbitration (K138)

K138 replayed 38 turns from existing frozen panels without executing the runtime. Thirty of those turns were legacy-correct and constrained by covered interpretation families. On that subset the arbiter agreed with legacy execution on all 30 turns, achieved 1.0 covered-family accuracy, and introduced zero regressions. The minimum and mean winning margins were 0.7541 and 1.3973. It also corrected three banked misroute specimens: a polar modal was assigned to `typed_modal` with margin 0.8796; a registered-source request previously captured by public text was assigned to `registered_source_action` with margin 1.2084; and a stale causal opinion query was sent to clarification with margin 0.3591. Runtime execution count was zero and the analysis completed in 0.438 seconds.

This was shadow development evidence. It showed that a simultaneous typed-geometric lattice could dominate the demonstrated first-match failures without taking control. It did not establish prospective transfer or justify global authority.

5.4 Authoritative arbitration and prospective transfer (K147–K148)

K147 promoted arbitration at the two demonstrated capture points. Its development report reproduced both parent misroutes, produced construction-determined polar-modal terminals across three content-diverse surfaces, gated causal graph execution by the selected family, retained causal and modal conversation-local worlds, replayed all passing turns from K139, K142, and K145, retained registered tools, essays, and rapid experts, and restored both defects when the new authority was ablated. The suite took 368.32 seconds. It explicitly did not replace the full precedence chain.

K148 froze K147 and executed an untouched panel exactly once. All ten turns across seven cases passed; architecture rate was 1.0 and runtime error rate 0.0. The panel included a fresh polar modal, an off-topic question inserted between causal turns, evidence-bounded essays, grounding control, the conversation-local modal world, a registered tool, and rapid experts. Mean cold-start latency was 19.60 seconds; mean steady-turn latency was 978.50 ms. This is the strongest current evidence that limited geometric arbitration transfers beyond its development specimens.

5.5 Thirty-domain registry and domain-scoped retrieval (K149–K151)

K149 assigned every source and all 104,000 rows in the frozen index to 30 domain experts. The artifact compiled in 16.17 seconds. Its eight routing probes produced 0.75 top-1 and 0.875 top-2 accuracy. Because this was a registry build rather than a frozen evaluation, these values are descriptive, not a broad routing benchmark.

K150 reproduced an off-domain provisional-context contaminant in the frozen parent: a computer-security row sharing surface terms entered a memory/identity essay's sense context. Domain scoping routed the topic to psychology and sociology/anthropology, constrained auxiliary retrieval to three source manifests, excluded the contaminant, and preserved the accepted `memory.n.03` and `identity.n.01` senses. The system replayed 33/33 previously passing panel turns, retained the 9/9 corrected K132 regression, retained K147 dispatch behavior, K144 senses, K137 repairs, evidence-bounded essays, clarification/anti-echo controls, and rapid experts. Removing domain scoping restored the contaminant with the blessed senses otherwise unchanged. Development passed in 474.56 seconds.

K151 then froze K150 and executed an untouched panel once. It passed all seven turns and all six cases with architecture rate 1.0 and zero runtime errors. Mean cold start was 15.63 seconds; mean steady-turn latency was 980.14 ms; the panel runtime was 22.53 seconds. The claim boundary is exact: the 30-expert registry is load-bearing at query time for the tested history-expert essays and retained controls. Coverage beyond the 105 indexed sources, unrestricted conversation, and open-domain knowledge are not established.

5.6 Typed-goal carrier removal and prospective transfer (K152–K153)

K152 opened the remaining canonical-English carrier at the registered-source planning seam. The inherited path serialized a typed request into canonical English and reparsed it to recover an executable plan. K152 added a direct typed compiler that consumes the relational goal and emits `ToolAction` objects without that intermediate language surface. Across four registered-source constructions—minimum and maximum earthquake queries and minimum and maximum population queries—the typed and carrier paths produced byte-identical plans and identical fields. End-to-end operation, status, response, and evidence were unchanged on the opened surface. Removing `typed_goal_direct` restored the carrier path with behavior otherwise identical, establishing a localized causal replacement rather than a coincidental output match. The development suite retained 40/40 passing turns from five prior one-shot panels, the 9/9 corrected K132 regression, essays, modal worlds, rapid experts, K147 arbitration, K150 domain scoping, and all parent-integrity checks. It completed in 444.12 seconds. The result did not remove English carriers at other seams and was not yet prospective.

K153 froze K152 before panel construction and executed its untouched panel once. All seven turns across six cases passed with architecture rate 1.0, case and turn accuracy 1.0, and zero runtime errors. Fresh minimum and maximum queries in both registered domains executed through the typed path without the canonical-English internal carrier; evidence-bounded essays, grounding control, the conversation-local modal world, and rapid experts also remained accepted. Mean cold-start latency was 22.53 seconds, mean steady-turn latency was 1.99 seconds, and total panel runtime was 36.55 seconds. The exact claim is therefore one-seam typed-goal transfer under a frozen prospective protocol—not global carrier removal, unrestricted conversation, or open-domain knowledge.

5.7 Coverage, knowledge-scale fallback, and total role resolution (K154–K160)

K154 changed the diagnostic scale without changing the claim class. It executed a continuous coverage scoreboard over 321 recorded turns: 274 correct responses, 40 clarifications, seven misroutes, zero wrong answers, and zero runtime errors. Because the surfaces were known to the harness and were permanently recorded for exclusion from future untouched panels, this was not a blind or prospective benchmark. Its value was taxonomic. The misroutes concentrated attention on missing knowledge and unresolved semantic roles rather than on the already accepted modal, causal, tool, and essay mechanisms.

K155 compiled a local Wikipedia abstract snapshot into a schema-compatible fallback index. Metadata records one million articles read, 2,772,709 sentences observed, 2,495,302 provenance-bearing rows indexed, and a term-frequency vocabulary of 891,209 entries under CC-BY-SA 4.0. K156 activated this tier only when the curated 105-source index contained no joint endpoint evidence; curated topics remained on curated sources. Development passed, but K157's frozen one-shot panel remained failed at 7/8 turns across seven cases. The failed manga/artists turn satisfied most surface and structural checks but did not form the required Wikipedia joint spine because the dictionary-unknown noun `manga` silently dropped from the endpoint set. K158 preserved the failure and identified a genuine unresolved-role defect linking the K157 failure to K154's unsupported `welkin drossum` release.

K159 repaired the class in two directions. A dictionary-unknown noun with real Wikipedia document frequency is retained as a nominal endpoint and must participate in relational connectivity; a coinage unresolved by the dictionary, caller corpus, and Wikipedia tier blocks release regardless of how many neighboring terms are known. Development reproduced both parent defects, repaired both, retained all passing turns from seven earlier one-shot panels, and restored the defects under ablation. K160 froze that successor and passed its untouched panel once at 8/8 turns across 7/7 cases, with architecture rate 1.0, zero runtime errors, 1.164 s mean steady latency, 19.24 s mean cold start, and 50.11 s total runtime. This establishes prospective transfer of total role resolution within the indexed tiers; it does not establish polysemy resolution or open-domain knowledge.

5.8 Typed transformation acts and prospective transfer (K161–K163)

K161 treated definition, summarization, enumeration, and comparison as operations over already admissible relational evidence rather than as independent answer templates. The extended arbiter scored these act families simultaneously with the inherited families. Definition used the lexical sense path; summaries, lists, and comparisons selected and arranged provenance-bearing clauses. The development suite unified twenty definition paraphrases under one answered terminal, executed all four act families, replayed all passing turns from eight one-shot panels, retained modal and causal worlds, registered tools, rapid experts, essays, grounding controls, and the corrected K132 regression, and restored the parent's act-free behavior under ablation.

The development run also exposed a defect that a mechanism-only evaluator initially obscured. The opinion question “What's your take on modern art these days?” was captured as a definition because an optional-article regular expression consumed the prefix `the` inside `their`-like possessive forms and because definition parsing admitted possessive initial terms. This produced a semantically nonsensical definition response while leaving the causal world intact. The repaired

grammar requires a complete article token and rejects possessive initial terms; the evaluator now verifies observable modal and causal retention, the unrelated art boundary, the absence of stale content, and stable causal answers rather than requiring a historic audit flag. The complete K161 development suite then passed in 917.91 seconds, including 62/62 panel replay turns and all retained controls. The opened diagnostic remains preserved with pre-repair hashes.

K162 froze the repaired K161 source and assets before any untouched act surfaces were authored. Eligibility was determined only at the topic level, with exact runtime surfaces excluded. The panel contained nine cases and thirteen turns, including fresh definition, summary, list, comparison, modal, tool, rapid-expert, causal-persistence, and act-boundary surfaces. It was executed exactly once and passed 13/13 turns across 9/9 cases, with architecture rate 1.0, zero runtime errors, 0.943 s mean steady-turn latency, 19.15 s mean cold start, and 71.80 s total runtime. The causal query remained stable across an intervening unrelated opinion question, and the opinion question was not captured as a definition or causal continuation. This is prospective evidence for transfer of four typed acts and the repaired boundary, not act-family closure or unrestricted realization.

K163 retained all 321 K154 surfaces and added forty-eight act and boundary surfaces, yielding a continuous 369-turn coverage scoreboard. It recorded 357 correct responses and 12 safe clarifications, with zero wrong answers, misroutes, acceptable alternatives, or runtime errors. Definition improved to 60/60 correct; summary, list, comparison, and act-boundary control were each 12/12; registered-source actions were 116/116; typed modal 72/72; causal discourse 20/20; stale-context controls 5/5; corpus compilation 12/12; and essays produced 24 correct responses plus 12 clarifications for unsupported invented topics. Mean latency over the heterogeneous continuous run was 3.603 s. These values are useful coverage evidence, but K163 is neither blind nor prospective and cannot be interpreted as a general conversation benchmark.

5.9 Manifest-attested production canary

The frozen K162 runtime was packaged with its source, extension, runtime assets, one-shot report, and K163 coverage report. The service refuses startup if any manifest-attested hash differs, if the K162 report is not a one-execution prospective pass with reruns forbidden, or if the K163 source identities do not match the frozen runtime. On Box C, the service cold-loaded in 16.21 seconds. A private canary passed runtime identity, prospective evidence, coverage identity, greeting, definition, summary, comparison, repaired boundary, and typed-tool checks. The same suite then passed through the public `wanderaround.io` route. A browser-level test verified that K162 was selected, the identity anchor and claim boundary rendered, and a live definition turn completed with 428 ms reported native reasoning latency. K153 was then stopped and disabled but retained for rollback. These are deployment checks, not new capability evidence.

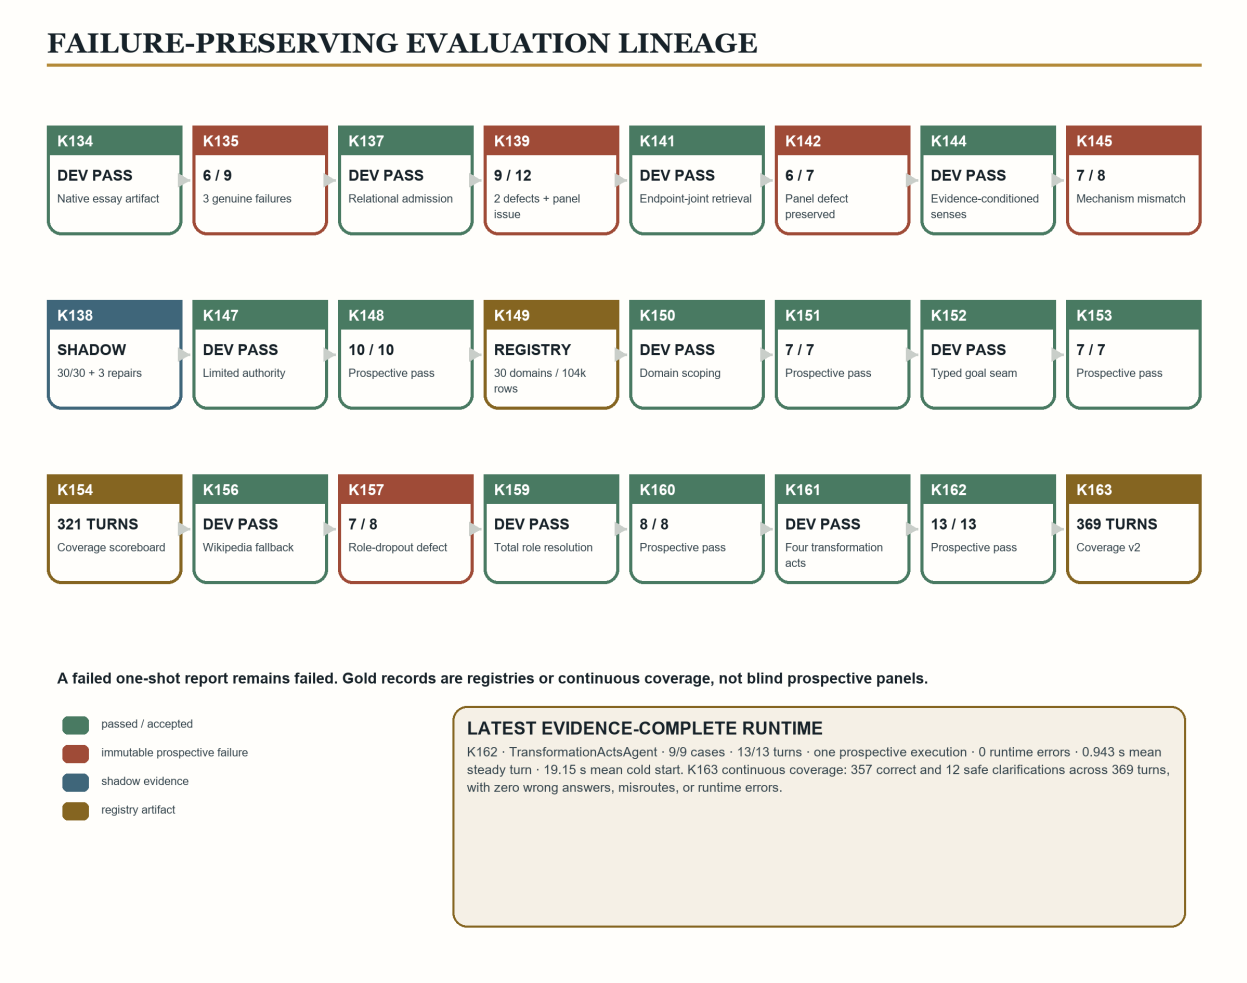


Figure 4. Failure-preserving lineage from native essay construction through knowledge-scale fallback, total role resolution, and typed transformation acts. Red stages are immutable prospective failures; gold stages are registry or continuous-coverage artifacts; blue denotes shadow evidence; green stages are accepted development or prospective passes. Development repairs do not erase earlier results.

5.10 Consolidated result table

Stage	Evidence class	Main result	Outcome	Boundary
K134	Development	Native typed essay artifact and document inverse	Pass	Not prospective; bounded prose
K135	Untouched one-shot	Essay transfer, unknown-topic and modal controls	6/9	Three genuine failures
K137	Development	Relational admission and modal precedence	Pass	No geometric ingress
K139	Untouched one-shot	Admission transfer and modal world	9/12	Two genuine defects; one panel defect
K141	Development	Endpoint-joint retrieval	Pass	Modal defect deferred
K142	Untouched one-shot	Retrieval and retained capabilities	6/7	Only failure was panel-authoring mismatch

Stage	Evidence class	Main result	Outcome	Boundary
K144	Development	Evidence-conditioned lexical senses	Pass	Arbitrary sense quality unproven
K145	Untouched one-shot	Sense transfer and retained capabilities	7/8	Only failure pinned an internal mechanism
K138	Shadow development	Simultaneous typed-geometric arbitration	30/30 legacy-correct; 3/3 banked misroutes corrected	No dispatch authority
K147	Development	Arbitration authoritative at two capture points	Pass	Not full chain replacement
K148	Untouched one-shot	Fresh arbitration transfer	10/10 turns; 7/7 cases	Not unrestricted conversation
K149	Registry artifact	30 domains, 105 sources, 104,000 rows	0.75 top-1; 0.875 top-2 on 8 probes	Descriptive small probe set
K150	Development	Domain-scoped auxiliary retrieval	Pass; 33/33 replay; 9/9 regression	Main recall unchanged
K151	Untouched one-shot	Registry load-bearing at query time	7/7 turns; 6/6 cases	105-source boundary
K152	Development	Direct typed-goal compilation at registered-source seam	Pass; 4/4 byte-identical plan surfaces; 40/40 replay	Other English carriers remain
K153	Untouched one-shot	Fresh transfer without the opened internal carrier	7/7 turns; 6/6 cases	One seam only; not unrestricted
K154	Continuous coverage	Broad failure taxonomy over retained surfaces	274 correct; 40 clarified; 7 misrouted / 321	Not blind or prospective
K155–K156	Index artifact + development	One-million-abstract local fallback tier	2,495,302 rows; development pass	Knowledge quality bounded by source and filters
K157	Untouched one-shot	Wikipedia-tier transfer	7/8 turns; 6/7 cases	Genuine unresolved-role dropout
K159	Development	Total resolution across dictionary, caller, and Wikipedia tiers	Pass; defects restored under ablation	Polysemy remains bounded
K160	Untouched one-shot	Fresh unresolved-role transfer	8/8 turns; 7/7 cases	Indexed-tier boundary
K161	Development	Definition, summary, list, and compare as typed acts	Pass; 62/62 panel replay	Caller-supplied free-text acts deferred
K162	Untouched one-shot	Fresh act transfer and repaired boundary	13/13 turns; 9/9 cases	Not act-family closure
K163	Continuous coverage v2	Retained coverage plus new act families	357 correct; 12 clarified / 369	Not blind or prospective

6. What Fable added

Fable's most important contribution was diagnostic before it was implementational. Code inspection showed that the K-line's "semantic" layers still exchanged English strings. The construction transducer had already extracted source, direction, metric, and projection roles, yet serialized them back into canonical English for another parser. The inheritance chain implemented interpretation as the first recognizer returning a non-null result. This explained apparently unrelated failures as one representational defect: same-meaning paraphrases could be invisible, alternative interpretations were never compared, and no graded confidence supported clarification.

The resulting additions were not another ontology. They were a new inter-layer contract. K138 supplied typed candidates, simultaneous geometric resonance, contextual feature evidence, a winning margin, and a clarification threshold. K147 gave that record limited causal authority. K149 gave corpus domains explicit manifests and geometric addresses. K150 connected those domains to retrieval rather than leaving them as metadata. K151 supplied prospective transfer evidence. K152 removed one canonical-English reparsing seam and demonstrated exact typed-plan equivalence; K153 supplied prospective transfer evidence for that removal. K154 then used a broad continuous scoreboard to locate residual defect families. K155–K160 converted one of those families into a larger knowledge tier with total unresolved-role admission, while preserving an intermediate failed prospective result. K161–K163 added transformation acts, repaired a lexical/semantic boundary revealed by runtime diagnosis, established one-shot transfer, and remeasured continuous coverage. Geometry thereby moved from a downstream routing ornament toward the input-side basis of interpretation; specialists moved from a conceptual pool toward query-time evidence constraints; and language-shaped internal buses began yielding to explicit typed state.

Fable's methodological contribution was equally important. Its ladder distinguished implemented mechanisms, research gates, and open hypotheses rather than treating every successful local test as evidence for the final goal. It favored zero-runtime adjudication when a prospective result had already been spent, isolated the smallest genuine defect, added one causal repair, replayed retained behavior, froze source identity, and reserved unseen surfaces for a single evaluation. That sequence made the later progress more compute-efficient because it prevented broad speculative retraining and avoided consuming prospective evidence during diagnosis. The method does not guarantee a correct architecture, but it turns architectural development into a chain of localized causal claims whose failures remain scientifically informative.

This changes the plan in a precise way. The central research gate is no longer “add enough WordNet, VerbNet, patterns, and specialists that arbitrary language happens to be covered.” WordNet and VerbNet remain useful as authenticated sources of lexical relations and transformation families, but the system must compress their recurring structures into reusable operators and learn or induce the mapping from natural surface to a candidate lattice. The decisive next work is global geometric ingress, extension of the proven typed-goal handoff across all seams, ephemeral compositional closure, and an inverse realizer—not indefinite ontology expansion.

7. Comparison with prevailing architectures

7.1 Transformer language models

Transformers learn distributed sequence transformations end to end and can internalize syntax, semantics, factual associations, discourse habits, and task procedures in shared parameters [1,2]. Their advantage is broad amortization: one model can respond across enormous surface diversity. Oracle instead externalizes the joints. Its representation state is narrow but interpretable; its knowledge experts are explicit and removable; its memory distinguishes fact from episode; its reasoning and tool operations are typed; and its release boundary is constrained by provenance. This offers causal auditability and cheap modular updates, but currently sacrifices the Transformer’s fluency, coverage, and learned semantic smoothness.

The proper comparison is therefore not “Oracle is a smaller GPT.” It is a different factorization. A Transformer learns an extraordinarily capable high-dimensional field in which many functions emerge jointly. Oracle attempts to construct an ecology of smaller organs whose interactions produce the behavior. The user’s “single cell versus multicellular organism” analogy is useful if treated carefully: it describes organizational modularity, not biological fidelity. A Transformer is not literally a cell, and Oracle is not literally a brain. The scientific proposition is that differentiation of function, persistent state, routing, and boundary control may yield capability per trained parameter that monolithic next-token scaling leaves unused. The compiler analogy sharpens this comparison: Transformer inference samples or decodes from a learned sequence field, whereas Oracle intends to parse a signal into typed IR, apply inspectable passes, and realize a verified result. Present Oracle only partially satisfies that compiler contract.

7.2 Mixtures of experts and retrieval augmentation

Sparse mixtures of experts activate a subset of learned subnetworks per input, increasing parameter capacity without proportional computation [4]. Retrieval-augmented generation conditions a parametric generator on nonparametric memory [3]. Oracle shares conditional routing and external knowledge, but differs in what an expert is and what owns the final answer. Its experts may be compiled from bounded public data in seconds, carry source and schema identity, and be analytically fused under explicit competency constraints. The semantic result is intended to exist independently of the realizer. A retrieved sentence cannot become authoritative merely because it appears in context; it must enter a typed evidence relation. These are promising differences, but no cost or quality superiority over modern MoE/RAG systems has yet been demonstrated on standard external benchmarks.

7.3 Neural-symbolic and vector-symbolic systems

Oracle is closest in spirit to vector-symbolic and neural-symbolic architectures. Smolensky’s tensor-product representations and Plate’s HRRs address the role-binding problem; Kanerva’s hyperdimensional computing frames high-dimensional distributed vectors as a general computational medium [8–10]. Oracle’s differentiating proposal is not a new binding operator. It is the use of HRR as an end-to-end transport and control substrate joining semantic ingress, a fixed routing topology, persistent opaque memory, source-locked experts, multi-head consequence reasoning, epistemic arbitration, tool observations, and terminal language release. Whether this integration scales beyond bounded test grammars is the central open empirical question.

7.4 Cognitive analogy

Human learning motivates the architecture but does not validate it. Children experience enormous sensory and linguistic streams, learn within social interaction, build object and causal expectations, and acquire productive language from recurring correspondences among sound, gesture, attention, objects, actions, and correction. This supports investigating structured priors, compositionality, intuitive world models, grounding, and learning-to-learn [13,16,17]. It does not imply that HRR, the Tree topology, the proposed four-part signal basis, or the current memory cycle reproduces neurobiology. Oracle is a cognitive-level construction: an attempt to name functional requirements without claiming one-to-one correspondence to brain regions.

8. Significance

The immediate significance is methodological. Oracle demonstrates that a language-agent research program can preserve failures, freeze runtimes, execute untouched panels once, distinguish behavioral contracts from internal mechanisms, and attach destructive controls to architectural claims. In a field where fluent output can tempt researchers to overinterpret a demonstration, this discipline is itself consequential.

The architectural significance is the emergence of a common representational currency across functions. The same bind/bundle/route/cleanup operations can address a construction family, a discourse referent, a fact, an expert, a tool operation, an episodic handle, or a document obligation. This does not make the contents identical. It makes their interfaces composable. If a corpus, a sensor, and a tool can each be compiled into typed state transitions over the same role space, then adding competence becomes closer to registering a new organ than retraining the whole organism.

The computational significance, if the program succeeds at broader scale, would be a different marginal-cost curve. A conventional foundation model pays a large global training cost to amortize broad competence, then often requires fine-tuning, retrieval infrastructure, or prompting for specialization. Oracle aims to keep the general transformation kernel stable while compiling a new corpus expert in seconds or hours and composing it at query time. Existing I1 and K149–K163 results show parts of this behavior on bounded data: small expert compilation, deterministic routing artifacts, a one-million-abstract fallback index, sub-second mean steady latency on the K162 one-shot panel, and a public CPU service. They do not establish GPT-3- or GPT-4-level knowledge, language quality, or reasoning at a fraction of the cost. Such a claim would require matched public benchmarks, matched corpora, total energy and engineering accounting, latency/throughput measurement, and independent replication.

The theoretical significance is a concrete answer to the question “where should open-endedness reside?” Oracle’s answer is: not in an infinite inventory and not necessarily in unrestricted stochastic continuation, but in the closure properties of a finite relational algebra. If known operators can accept unseen fillers; if valid results can become inputs to further operators; if temporary states persist without ontology insertion; if context changes the interpretation landscape; if specialists add bounded invariants; and if realization preserves the solved state, then the reachable semantic space can be far larger than the stored primitive set. This is the same sense in which a finite grammar can license indefinitely many sentences or a finite program can generate unbounded sequences. Practical open-endedness would still be limited by memory, numerical interference, parser coverage, world knowledge, and action authority.

The compiler significance is that the remaining problem can now be partitioned by layer. A malformed glyph or sound belongs to signal analysis; an ambiguous argument structure belongs to construction parsing; an unsupported referent belongs to semantic admission; an incorrect derivation belongs to a typed transformation; a hallucinated statement belongs to evidence or realization; and an unsafe action belongs to authority and postcondition verification. That partition does not solve those classes, but it converts “make the model conversational” from one undifferentiated objective into falsifiable interface contracts. The transition is substantive because K152 and K161 already demonstrate both sides of the principle: removing an internal English carrier can preserve an executable plan, while a lexical-prefix defect can corrupt semantic

dispatch. The architecture is therefore sufficiently concrete to attack systematically, though not sufficiently complete to declare the general language problem solved.

Finally, the epistemic significance is the refusal to identify language with thought. Oracle treats language as a shaped public interface to a deeper typed state. A question defines constraints and a relational germ; reasoning transforms that germ; an answer grows as a congruent inverse image. This is more than a metaphor when the inverse, evidence addresses, and role recoveries are measured. It remains a research hypothesis when applied to unrestricted prose. The four-primitive front end strengthens this position only if it is grounded: decomposing signals may expose reusable morphology, but it cannot make a symbol about the world merely by redescribing its shape.

9. Limitations and threats to validity

The current system is not an unrestricted conversationalist. K162's thirteen-turn panel is tiny relative to the diversity of natural dialogue, and K163's 369 surfaces are internally authored, recorded, and non-blind. The interpretation-family inventory is finite; its prototypes and feature weights are manually engineered; the dense encoder lacks learned lexical semantics; and global dispatch still contains an inherited ordered handler chain. The native essay realizer is proof-shaped and can produce awkward, weakly synthesized, or incorrect dictionary senses. Public canary outputs can be structurally safe while remaining stylistically poor or topically unhelpful. Successful operation and provenance checks therefore do not establish human-level usefulness or exposition.

The domain registry's eight probes are descriptive and insufficient for general routing claims. Top-2 success can hide a weak top-1 boundary. K150 scopes only auxiliary retrieval, not the frozen main recall pass. K155 increases lexical reach dramatically, but Wikipedia presence is not equivalent to correct sense, trustworthy evidence, or complete world knowledge. Total role resolution can prevent silent noun deletion and unsupported coinage release while still choosing a poor relation among polysemous terms. Domain assignment, source authority, fallback policy, and the one-million-article cutoff were authored.

The strongest evaluations are internally designed. Hash freezing and one-shot execution reduce outcome leakage but do not replace external benchmarks or independent replication. Many panels test exact structured contracts rather than human judgments of usefulness, truth, or conversational naturalness. Some prospective failures were caused by evaluator defects, showing that preregistration can preserve a bad metric as faithfully as a good one. Zero-runtime adjudication mitigates interpretive confusion but cannot recover statistical power.

CPU efficiency has not been measured against storage- and compute-matched baselines. Reported compile times exclude the creation and curation of source corpora, ontology, code, tests, and the Wikipedia snapshot. K162 measured 19.15 seconds cold and 0.943 seconds mean steady latency on its one-shot panel; K163's heterogeneous continuous run averaged 3.603 seconds; the production service cold-loaded in 16.21 seconds; and one browser definition turn reported 428 ms native reasoning. These are not throughput benchmarks. Concurrent-user behavior, tail latency, queueing, memory pressure, swap pressure, energy, artifact-transfer cost, and quality under production traffic require measurement.

The four-part signal decomposition has not been implemented or evaluated. Latin-letter examples may bias the primitive inventory toward one script family; nonalphabetic, abjad, abugida, syllabic, logographic, tactile, and signed systems may require different or additional primitives. Even perfect script reconstruction would establish a codec, not semantics. Only grounded cross-language and cross-modal convergence tests can support the stronger universal-front-end claim.

The Tree topology's ten nodes and twenty-two paths are a useful fixed routing manifold, but no ablation has yet shown that this specific topology is superior to a generic graph, a learned router, or another bounded hierarchy across broad tasks. Likewise, the 4:1 spine is causally necessary under several frozen contracts, but its optimality and scaling law under many rails remain open.

No result supports consciousness, sentience, human-equivalent cognition, or AGI. The "mind," "organ," and "multicellular" language in this paper names architectural organization, not phenomenology.

10. Revised research plan

10.1 Gate A: global typed-geometric ingress

Refactor ingress into explicit compiler layers: signal decomposition, construction parsing, semantic role resolution, typed-IR lowering, and chain-wide arbitration. Every current handler should expose a typed candidate with required roles, proposed goal, geometric resonance, context compatibility, knowledge support, and reconstruction score. Arbitration must run before execution. The system should clarify on low absolute support, thin winning margin, unresolved required roles, or failed analysis-by-synthesis. Evaluation must use paraphrase-disjoint, construction-disjoint, and adversarially ambiguous panels; compare against the frozen ordered chain; and measure family accuracy, role accuracy, calibration, clarification utility, and silent-capture rate. K161's lexical-prefix capture becomes the minimal regression: no signal-level segmentation decision may silently determine a semantic act.

The present hashed-word encoder should remain as a control. The experimental system should learn an in-house HRR relational state transducer from dictionaries, WordNet, VerbNet, dependency/constituency resources, and broad non-panel contexts without invoking a pretrained encoder at runtime. It should preserve multiple candidate senses when the margin is weak and use argument type, discourse state, forward/inverse proposition congruence, and source compatibility to settle or clarify. In parallel, test direction, junction, closure, and recurrence as a minimum-description written-signal basis across typologically diverse scripts. Compare four-primitive, expanded-primitive, byte, character, subword, and acoustic front ends on reconstruction, corruption robustness, construction recovery, cross-language state convergence, and withheld-language operator transfer. All systems must receive paired situational grounding; a glyph-only system is a negative control.

10.2 Gate B: direct typed-goal handoff

K152–K153 close this gate only at the registered-source planning seam: a direct typed compiler now emits byte-identical plans and has transferred prospectively. Extend the same contract to every remaining boundary among ingress, controller, reasoner, memory, specialist broker, document constructor, and tool compiler, deleting canonical-English inter-layer reparsing as each seam passes freeze, equivalence, retention, and ablation tests. English should occur only at the ear and mouth. An analysis-by-synthesis test should render the candidate interpretation into an independent paraphrase or structured gloss, re-encode it, and require relational agreement with the input. Ablating each new typed handoff must restore its opened carrier path or failure class.

10.3 Gate C: first-class ephemeral compositional closure

Implement a distinct query-local HRR workspace with explicit create, bind, branch, reference, promote, expire, and delete operations. Known operators must accept unseen fillers without retraining. The output of every valid operation must remain a valid input to another operation. Benchmarks should progress from unseen bindings to novel nested rules, hypothetical worlds, multi-turn persistence, local counterfactual edits, plural reference, tense/aspect, interacting modality, and contradictions. Nearest-ontology collapse must be an explicit ablation.

10.4 Gate D: operator-family compression

Treat WordNet and VerbNet as data for discovering recurring transformation families, not as the final state inventory. Candidate operators should be retained only when decomposition into simpler operators increases description length, reduces reconstruction, reduces transfer, or destabilizes composition. The objective is minimum-description relational closure: compress redundancy by identifying the algebraic generator that reconstructs a family of forms, not merely by storing fewer surface examples.

10.5 Gate E: specialist acquisition and synthesis

Expand K149/K162 from the 105-source curated tier and K155 Wikipedia fallback to authenticated, license-tracked corpora in 10–20 general domains, then to arbitrary registered schemas. Build the acquisition skill so a source manifest, schema, license, provenance policy, lexical profile, geometric address, contradiction policy, and removal receipt are compiled together. Evaluate specialist compilation time, source coverage, routing calibration, retrieval precision, contradiction handling, and post-dissolution contamination. The two- or three-expert broker should compare its analytic hybrid against the best single expert, fixed late fusion, a generic graph router, a no-spine control, and an oracle upper bound. A synthetic expert must demonstrate cross-domain reasoning that no selected source independently contains, while every released claim remains traceable to supporting evidence or an explicit derived step.

10.6 Gate F: unrestricted realization without semantic bypass

Extend K161's evidence-preserving transformation acts into an in-house Malkuth realizer over dialogue, explanations, technical answers, essays, paper sections, revisions, comparisons, equations, and symbolic restatements. It may use a frozen small carrier if and only if the carrier cannot determine answer content independently of the protected state. Evaluation should rotate or erase the relational germ, evidence tape, document plan, register, and act program separately. A successful realizer must preserve claims, polarity, modality, quantities, citations, paragraph obligations, and cross-section terminology while remaining fluent under human evaluation. Caller-supplied free-text transformations should precede unrestricted generation because they isolate whether the act algebra can operate over novel content without importing answer knowledge.

Paper writing should be treated as persistent structured work. The document state should bind thesis, outline, section obligations, evidence, citations, objections, terminology, unresolved questions, and revision history. Metrics should include evidence entailment, citation validity, outline coverage, cross-section consistency, revision improvement, and long-context recovery—not merely surface preference.

10.7 Gate G: open tool and sensor domains

Unify conversation and tool use by representing both as relational state transformation. Move from registered read-only tools to schema induction, natural argument filling, plan generation, verified observations, retry/reorder/stop, and narrowly authorized writes. Require simulation, approval boundaries, idempotency, rollback where possible, side-effect receipts, and adversarial prompt-injection tests.

For sensors, define a typed stream expert with time, unit, calibration, missingness, drift, uncertainty, and provenance roles. A sensor array should become a bounded corpus continuously updated by observation. Multiple sensor specialists may then synthesize a local state and converse about it. This follows from the architecture but is not yet implemented.

10.8 Gate H: comparative evaluation and release

Freeze external conversation, compositionality, retrieval, reasoning, document, tool, and long-context suites. Include storage- and compute-matched baselines: the frozen ordered Oracle, a conventional RAG pipeline, compact Transformer encoders/decoders, sparse expert baselines, generic graph routing, no-spine HRR, randomized topology, and external teacher systems. Report total development compute, training and compilation wall time, CPU time, energy, RAM, artifact size, cold and steady latency, throughput, accuracy, calibration, abstention, human preference, and cost per accepted task.

Promote in layers: offline reproduction, independent cross-host reproduction, public artifact release, shadow endpoint, canary, and live production. K162 has completed the internal and public canary sequence; the next promotion criterion is not another endpoint label but independent external evaluation. Never label a bounded release “open-ended” because its components suggest a path to open-endedness. The label should follow a fresh, broad, adversarial behavioral result.

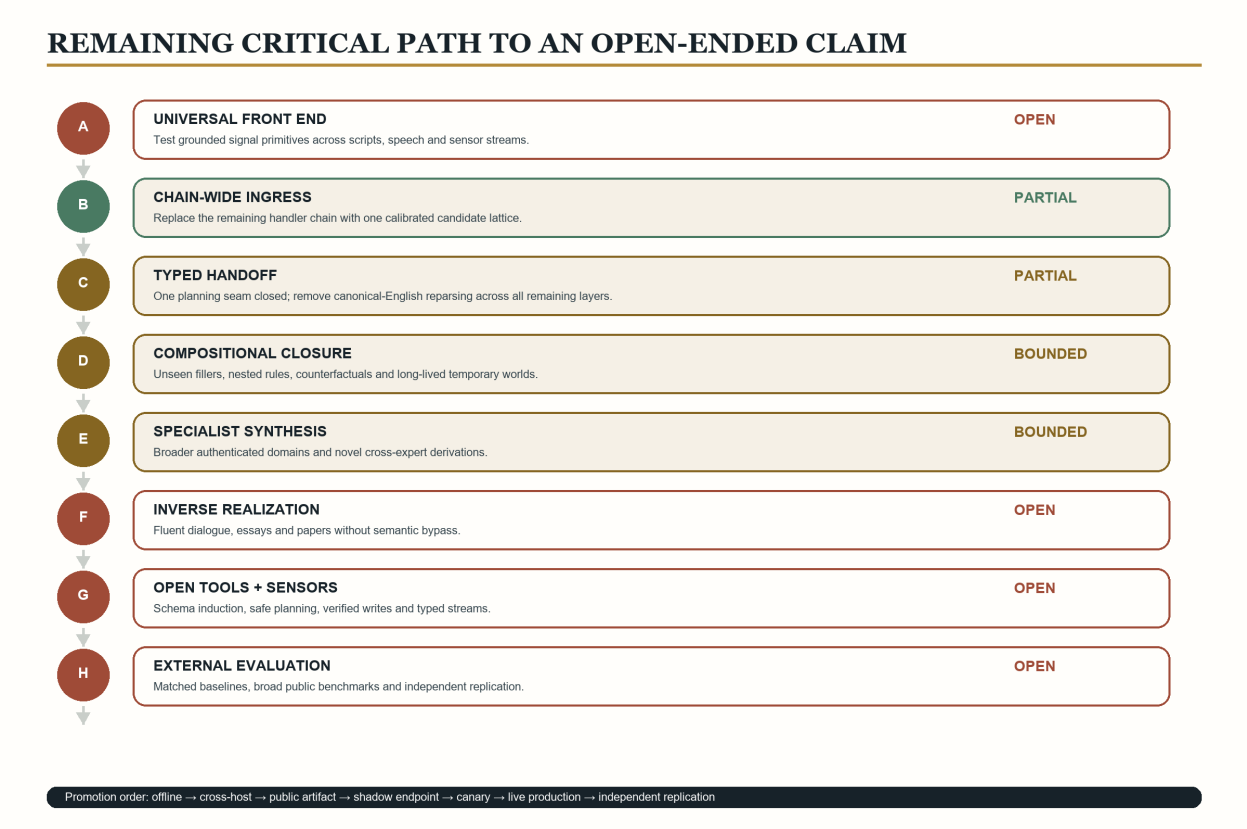


Figure 5. Remaining critical path. K162 adds prospective transformation-act transfer, total role resolution, and a public manifest-attested runtime to the earlier limited arbitration, domain scoping, and one-seam carrier removal. A universal grounded front end, chain-wide arbitration, all-seam typed handoff, compositional closure, non-bypassing realization, open planning, and external evaluation remain necessary before an unrestricted-agent claim.

11. Conclusion

Oracle presently occupies an unusual middle ground. It is far more than a conceptual diagram: it has executable HRR transport, typed episodic and semantic memory, a protected spine, finite clause and modal worlds, rapid source experts, ephemeral hybrid specialists, bounded consequence reasoning, evidence tapes, read-only tool state, native essay artifacts, geometric interpretation candidates, limited authoritative arbitration, a prospectively validated 30-domain retrieval scope, one prospectively validated direct typed-goal planning seam, a one-million-abstract fallback tier, prospectively validated total role resolution, and four prospectively transferred transformation acts. The exact frozen runtime is publicly reachable through a manifest-attested CPU service. It is also far less than a modern unrestricted conversational model: its semantic ingress is not global, its language realizer is bounded, its domain and source policies are authored, its panels are small, and its open-ended claims remain unearned.

The most important threshold crossed by the Fable-assisted additions is architectural coherence. The same system that already knew how to preserve typed facts, memories, tools, and expert states now has the beginning of a compatible input geometry and demonstrated typed bridges from

interpretation to evidence, transformation, and execution. Interpretation can be ranked rather than merely matched; ambiguity can be surfaced rather than silently captured; domain identity can constrain evidence rather than decorate metadata; unresolved content roles cannot silently disappear; a registered-source goal can reach an executable plan without being flattened into an internal English sentence; and definition, summarization, enumeration, and comparison can be represented as operators over admitted content. These changes have survived untouched one-shot tests. The pieces now fail at identifiable compiler interfaces instead of dissolving into a single vague complaint that the model "is not conversational."

If the next gates succeed, the contribution would not be a cheaper imitation of a Transformer. It would be evidence for a different organization of machine intelligence: a small, stable relational kernel whose reachable states grow through composition; whose factual breadth grows through compiled experts; whose continuity grows through typed cycling memory; whose actions are authorized state transitions; and whose language is the inverse surface of an already solved geometric object. The proposed universal front end would add a further result: different scripts, spoken languages, and sensor streams could compile into one grounded relational IR without making English the hidden ontology. That possibility is significant enough to pursue aggressively and uncertain enough to require the strict claim boundaries maintained here.

Author-contribution and AI-assistance statement

The originating concepts and research direction—including HRR-bound nested states, the Tree-of-Life node/path routing form, the stabilizing spine, Malkuth-only unbinding, a generalist composed from specialists, query-local hybrid experts, geometric transformation of language, inverse-image realization, continual substrate memory, natural language as traversable shape, and decomposition of signal forms into reusable geometric primitives—were proposed and directed by psiloceyeben. Fable and Codex, automated coding and research assistants operating under that direction, contributed code synthesis, source inspection, experiment execution, failure analysis, evaluation design, deployment, documentation, and manuscript preparation. Fable's specific K138–K163 methodological and implementational contribution is described in Section 6. Codex repaired and reevaluated the K161 act boundary, constructed and executed the K162/K163 gates, deployed the manifest-attested K162 service, and prepared version 1.1 of this manuscript. Automated assistants are not listed as human authors. All empirical claims are tied to locally preserved source, manifests, frozen panels, reports, and deployment checks; the author remains responsible for public claims and release decisions.

Data, code, and reproducibility

The working artifacts are preserved in the Oracle research workspace. The K162 runtime source SHA-256 is `24c9790b4521aa08001bc6394a709fa0b85581f043a8565059b153d9c514f474`; its act extension SHA-256 is `a61e645f0ca29ee5dab038bb55aa7164010029da9808d71c9695fca332f71f02`; and its frozen manifest SHA-256 is `a0189aa1f51e4d178746d5e1b18b339fb6f948faecc46cb35f3ba12c3b3cffbe`. The prospective report SHA-256 is `ccd719ef5b0c31edc8b058689c24d67df7bb5e321ff4ffc150cd487da1569c43`, explicitly records `execution_count: 1`, and forbids reruns. The K163 continuous coverage report SHA-256 is `c6be415c4265c77deeaa57295f222b2a5e4aee0021b5dc2e534d4fa558a26a34`. The K155 Wikipedia index SHA-256 is `ca3929ffee66025c00810815cda495b6566ca787565826a8ab91acabe92cb614`. The published paper package includes a machine-readable evidence ledger and artifact manifest. A separately packaged bounded Hugging Face candidate exists in the workspace; unrestricted conversation and open agency are not release claims.

References

[1] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, “Attention Is All You Need,” Advances in Neural Information Processing Systems 30, 2017. https://research.google/pubs/attention-is-all-you-need/

[2] T. B. Brown et al., “Language Models are Few-Shot Learners,” Advances in Neural Information Processing Systems 33, 2020. https://arxiv.org/abs/2005.14165

[3] P. Lewis et al., “Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks,” Advances in Neural Information Processing Systems 33, 2020. https://arxiv.org/abs/2005.11401

[4] N. Shazeer et al., “Outrageously Large Neural Networks: The Sparsely-Gated Mixture-of-Experts Layer,” International Conference on Learning Representations, 2017. https://arxiv.org/abs/1701.06538

[5] A. Graves et al., “Hybrid Computing Using a Neural Network with Dynamic External Memory,” Nature 538, 471–476, 2016. https://doi.org/10.1038/nature20101

[6] S. Yao et al., “ReAct: Synergizing Reasoning and Acting in Language Models,” International Conference on Learning Representations, 2023. https://arxiv.org/abs/2210.03629

[7] T. Schick et al., “Toolformer: Language Models Can Teach Themselves to Use Tools,” Advances in Neural Information Processing Systems 36, 2023. https://arxiv.org/abs/2302.04761

[8] P. Smolensky, “Tensor Product Variable Binding and the Representation of Symbolic Structures in Connectionist Systems,” Artificial Intelligence 46(1–2), 159–216, 1990. https://doi.org/10.1016/0004-3702(90)90007-M

[9] T. A. Plate, “Holographic Reduced Representations,” IEEE Transactions on Neural Networks 6(3), 623–641, 1995. https://doi.org/10.1109/72.377968

[10] P. Kanerva, “Hyperdimensional Computing: An Introduction to Computing in Distributed Representation with High-Dimensional Random Vectors,” Cognitive Computation 1, 139–159, 2009. https://doi.org/10.1007/s12559-009-9009-8

[11] J. A. Fodor and Z. W. Pylyshyn, “Connectionism and Cognitive Architecture: A Critical Analysis,” Cognition 28(1–2), 3–71, 1988. https://doi.org/10.1016/0010-0277(88)90031-5

[12] B. M. Lake and M. Baroni, “Generalization without Systematicity: On the Compositional Skills of Sequence-to-Sequence Recurrent Networks,” Proceedings of Machine Learning Research 80, 2873–2882, 2018. https://proceedings.mlr.press/v80/lake18a.html

[13] B. M. Lake, T. D. Ullman, J. B. Tenenbaum, and S. J. Gershman, “Building Machines That Learn and Think Like People,” Behavioral and Brain Sciences 40, e253, 2017. https://doi.org/10.1017/S0140525X16001837

[14] A. Graves, G. Wayne, and I. Danihelka, “Neural Turing Machines,” 2014. https://arxiv.org/abs/1410.5401

[15] M. Dehghani, S. Gouws, O. Vinyals, J. Uszkoreit, and Ł. Kaiser, “Universal Transformers,” International Conference on Learning Representations, 2019. https://research.google/pubs/universal-transformers/

[16] S. Harnad, “The Symbol Grounding Problem,” Physica D: Nonlinear Phenomena 42(1–3), 335–346, 1990. https://doi.org/10.1016/0167-2789(90)90087-6

[17] D. K. Roy and A. P. Pentland, “Learning Words from Sights and Sounds: A Computational Model,” Cognitive Science 26(1), 113–146, 2002. https://doi.org/10.1207/s15516709cog2601_4