The Distinction–Connection Substrate (Phaneron): A Unified Structural Framework for Meaning, Abstraction, and Resource-Bounded Prediction

Johannes Drescher
Independent Preprint

January 7, 2026

Abstract

We present a unified synthesis of the *Phaneron* research program: a minimal semantics substrate built from two primitives—*Distinction* (nodes exist because they are distinct) and *Connection* (edges exist because nodes co-occur). Concretely, the *Phaneron* is an unlabeled, undirected pseudograph with a single representational layer: nodes carry no attributes; edges carry no labels or weights; and all nuance—roles, n-ary relations, direction/causality, tense/aspect, provenance, even "rules"—is encoded exclusively as *reusable patterns* (finite subgraphs) and versioned rewrite events. We formalize meaning as structure (ideal: automorphism orbits; operational: r-hop neighborhoods and Weisfeiler–Lehman invariants) and optimize structure via a hybrid description-length objective aligned with predictive sufficiency on held-out and future-version tests. We then lift the substrate to agent level via *reflection parity*: an agent's model is "accurate" when it is a task-sufficient, resource-bounded quotient of the world rather than a microstate emulation. Finally, we define *predictability horizons* (singularities) as periods where the minimal compute required for useful forecasts exceeds capacity, and we give an operational detector (Singularity Index). The primary contribution of this paper is *framing*: we separate (i) definitions, (ii) derived properties, (iii) testable hypotheses, and (iv) optional speculative extensions, and we provide a concrete validation plan that can be reproduced with small-scale code and data.

1 Introduction

A recurring failure mode in knowledge representation is *semantic leakage*: we store meaning in a growing zoo of labels, types, edge predicates, weights, and ad-hoc metadata. This makes systems brittle, hard to audit, and hard to unify across modalities (text, vision, action, biology). In practice, we end up with two unsatisfying extremes: (i) richly typed symbolic graphs that do not scale and resist learning; or (ii) opaque embeddings that scale but offer weak inspectability and weak compositionality.

This paper argues for a third option: a *single structural layer* where the only primitives are *Distinction* and *Connection*. Everything else is represented as *structure*—small reusable patterns inside the same substrate. Meaning becomes a property of structural position; learning becomes the discovery and reuse of motifs that compress and predict.

Scope and non-scope. The goal of this synthesis is scientific coherence, not maximal ambition. We focus on the parts of the program that can be stated precisely and tested with reproducible experiments: the substrate, its objective and rewrite loop, and resource-bounded agent-level notions (reflection parity and predictability horizons). Several companion drafts explore applications (language, robotics, multi-agent systems, storage) and speculative extensions (informational physics). Those are summarized as applications or appendices but are not required to accept the core claims.

Category	In this paper
Definitions	The *Phaneron* substrate (Sec. 3), patterns and instances, the ADT (Sec. 4), meaning-as-structure semantics, reflection parity (Sec. 6), predictability horizons and Singularity Index (Sec. 7).
Derived properties	If nuance is encoded only as structure, labeled predicates can be translated into relation patterns; direction/time can be represented by asymmetric motifs plus version-time; horizons follow from bounded compute when novelty outruns consolidation.
Testable hypotheses	MDL-guided motif induction improves compression and predictive performance vs. label-heavy baselines; reflection parity predicts stagewise partition changes; SI detects impending horizons in streaming domains.
Speculative extensions	Informational physics (Universal *Phaneron*), universe-scale constraints, and other metaphysical interpretations are *not* required for the core and are left to separate papers.

Table 1: Claim status contract. The core is structural and testable; speculative extensions are explicitly optional.

Contributions.

- A unified statement of the *Phaneron* substrate, axioms, and meaning-as-structure semantics.
- A formal abstract data type (ADT) view $P = (G, \mathcal{E}, D, R, \mathcal{L})$ that makes the system implementable.
- A rewrite loop (propose $\rightarrow$ score $\rightarrow$ keep/refine/merge) guided by MDL and predictive sufficiency.
- Agent-level adequacy via reflection parity (task-sufficient coarse-graining).
- Predictability horizons and an operational detector (Singularity Index) for resource-bounded singularities.
- A validation plan with three "wedge" demos designed to survive skeptical scrutiny.

2 Status of claims: what is definition, what is hypothesis?

A unifying framework fails fast if it blurs definitions, theorems, and speculation. Table 1 is the contract of this paper.

3 The *Phaneron* substrate

3.1 Axioms

We state the substrate as axioms rather than metaphors.

Definition 1 (Phaneron substrate)**.** *A* Phaneron *graph is an unlabeled, undirected pseudograph $G = (V, E_G)$ (loops and parallel edges allowed). Nodes carry no labels or attributes; edges carry no labels or weights. All domain nuance is encoded only as* patterns*: finite unlabeled graphs that appear as subgraphs of G.*

The following axioms enforce single-layer purity (paraphrased from the substrate preprint):

- **A1 (Distinction).** A node exists iff it is a distinct thing. Nodes carry no labels or attributes.

Figure (placeholder). Minimal relation motif. Two participant nodes connect through an unlabeled intermediate node. Role markers, argument order, tense/aspect, etc. are represented by small additional structure attached to the intermediate node (never by labels or edge direction).

Figure 1: Relation-as-pattern (schematic).

- **A2 (Connection).** An edge exists iff two nodes are connected. Edges carry no labels or weights.
- **A3 (Single-layer).** All nuance is represented as patterns (finite subgraphs); there is no out-of-band metadata.
- **A4 (Pseudograph base).** Loops and multiplicity carry no default semantics; meaning arises only via patterns.
- **A5 (Reification by pattern).** n-ary relations, roles, and "direction" are reified as intermediate subgraphs.
- **A6 (Versioning).** All change is a sequence of delta events (add/remove/merge/split/replace) recorded as structure.
- **A7 (Merge).** Node merges are permitted under structural criteria; merges are replayable (reversible by replay).

3.2 Patterns and instances

Definition 2 (Pattern and instance)**.** *A* pattern *P is a finite unlabeled undirected graph. An* instance *of P in G is an injective homomorphism $f : P \hookrightarrow G$ (equivalently, a subgraph of G isomorphic to P). A pattern dictionary D is a multiset of patterns equipped with a multiset of placements (instances) in G.*

Why undirected? Undirectedness forces all asymmetry (agent/patient, cause/effect, before/after) to be expressed *inside* patterns. This is not a restriction; it is a discipline. It prevents semantic leakage into privileged edge direction, keeping the substrate uniform and audit-friendly.

3.3 Emergence without labels: relations, roles, direction, time

A binary relation between participants a, b can be encoded by inserting an intermediate relation node r connected to both (a minimal K_2-style relation pattern). Higher-arity relations insert a hub or small motif connecting all participants. Roles (agent/patient, subject/object) are encoded by asymmetric micro-motifs attached to the relation hub.

Direction and causality. The base graph has no arrows. A directed relation $A \to B$ is represented as an *asymmetric* pattern that outperforms its symmetric competitor under the objective (Sec. 4.2). In practice, asymmetry arises from (a) version-time precedence in the event log, and (b) world-time precedence encoded as patterns in G.

Proposition 1 (Direction/time as derivable structure)**.** *If (i) edits are versioned in $\mathcal{E}$, and (ii) candidate asymmetric motifs are admitted only when they improve predictive sufficiency while reducing description length, then any stable directed relation can be represented without primitive edge direction: its asymmetry is carried by motifs whose reuse improves the code.*

4 The *Phaneron* as an abstract data type

The substrate becomes a system when we add memory of change, a dictionary, rewrite rules, and an objective.

Definition 3 (Phaneron ADT)**.** *A* Phaneron *system is a tuple*

$$P = (G, \mathcal{E}, D, R, \mathcal{L})$$

where (i) G is an unlabeled undirected pseudograph; (ii) $\mathcal{E}$ is a replayable event log of delta edits and pattern rewrites; (iii) D is a multiset dictionary of patterns with placements in G; (iv) R is a finite set of admissible rewrite schemas (add/remove/merge/split/replace); and (v) $\mathcal{L}$ is an objective that scores states and candidate rewrites.

4.1 Versioning and replay

Versioning is not an afterthought: it is part of the semantics. All edits are logged as first-class event structure, forming a next-event chain. Snapshots are derived patterns (cuts) over the chain. This enables replay, audit, and "future-version" prediction tests (does a rewrite today make tomorrow easier to represent?).

4.2 Objective: MDL aligned with predictive sufficiency

We score a state by a hybrid objective:

$$\mathcal{L}(G, D) = \mathrm{MDL}(G; D) + \lambda\, \mathrm{Risk}_{\mathrm{pred}}(G, D) + \gamma\, \mathrm{Conflict}(G, D). \tag{1}$$

A simple MDL decomposition is

$$\mathrm{MDL}(G; D) = L(D) + L(\text{placements} \mid D) + L(\text{residual edges}).$$

The predictive term can be instantiated as held-out reconstruction (mask edges/pattern placements and predict them), or as *future-version* loss (predict edits or motifs that appear in later $\mathcal{E}$). Conflict is a structural inconsistency measure (e.g., mutually exclusive motifs competing for the same placements, or cycles in precedence motifs that violate a learned acyclicity constraint).

Admissible rewrites. A rewrite $r \in R$ is admissible when it decreases $\mathcal{L}$ by more than a margin and does not violate hard constraints (e.g., safety constraints in embodied systems). In many cases we use a two-stage gate: (1) cheap structural screening (WL signatures, local statistics), followed by (2) expensive re-scoring on a bounded region.

4.3 Runtime loop

The implementation-level view is a four-stage loop:

$$\text{Distinction} \to \text{Connection} \to \text{Abstraction} \to \text{Consolidation}.$$

Distinctions arrive as new nodes; local co-occurrences become connections; candidate motifs are proposed; and global consolidation accepts/refines/merges motifs based on (1). In companion drafts this is abbreviated as the D→C→A→Consolidation loop.

5 Learning and abstraction by rewrite

5.1 Dictionary induction

Dictionary induction searches for patterns whose reuse reduces description length and improves predictive sufficiency. At minimum, candidates can be sampled from frequent subgraphs; more aggressively, candidates can be proposed by a bounded abductive layer that explains prediction errors with new structure.

Listing 1: Sketch: dictionary induction loop (single-threaded).

```
Input: Phaneron state P=(G,E,D,R,L), budget B
repeat until budget exhausted:
  1) Propose candidate pattern P* from neighborhoods with high loss contribution
  2) Find placements of P* (approx via WL signatures; verify by exact match locally)
  3) Score Delta = L(G, D U {P*}) - L(G, D)
  4) If Delta < -margin and predictive tests pass: accept; else refine or discard
  5) Periodically run merge/split on nodes/patterns to reduce conflict and MDL
```

5.2 Merge and split

Merges reduce graph size by identifying two nodes as the same concept under structural criteria. Splits introduce new distinctions when a node is overloaded (polysemy) or internally inconsistent.

Definition 4 (Merge and split triggers (agent-agnostic))**.** *A merge is suggested when* ΔMDL *is strongly negative and predictive loss does not materially worsen. A split is suggested when local conflict exceeds a threshold and a proposed refinement predicts high expected gain.*

This framing matches the "smaller and better" intuition: cognition develops by becoming *more compressive while staying more predictive.*

6 Agent-level semantics: reflection parity

A substrate is not yet a theory of minds. We need a criterion for when a *Phaneron* is *accurate enough* for an agent.

6.1 Task-indexed coarse-graining

Let the world at time t be a micro-structure W_t (not necessarily a graph, but representable as one). Fix a resource bound B (time, compute, memory) and a goal set $\mathcal{G}$ (tasks/values). Let $Q(\mathcal{G})$ be the set of queries/predictions that matter for $\mathcal{G}$. Define an equivalence relation over world elements:

$$u \sim_{B,\mathcal{G}} v \iff \forall q \in Q(\mathcal{G}) : \; \text{Pred}_W(q \mid u) \approx \text{Pred}_W(q \mid v) \text{ within } \varepsilon(B). \tag{2}$$

Let $\pi_{B,\mathcal{G}} : W_t \to W_t / \sim_{B,\mathcal{G}}$ be the quotient map.

6.2 Reflection sufficiency and parity

Let P_t be the agent's *Phaneron* at time t. We say P_t is *reflection-sufficient* for $(B, \mathcal{G})$ if there exists a homomorphism

$$h : P_t \to \pi_{B,\mathcal{G}}(W_t)$$

that preserves the invariants required by $Q(\mathcal{G})$ within the error budget.

Definition 5 (Reflection parity)**.** *An agent is in reflection parity at time t (for* $(B, \mathcal{G})$*) when* P_t *is reflection-sufficient and* $\mathcal{L}(P_t)$ *is (approximately) minimal among feasible representations under the agent's objective.*

Interpretation. Reflection parity claims that models should be evaluated against the *right target*: the task-sufficient quotient, not the micro-world. This formalizes the slogan: *minds do not copy reality; they mirror it.*

7 Predictability horizons and singularities

A "singularity" is usually discussed as myth. Here we reduce it to a measurable, resource-bounded horizon.

7.1 Minimal compute and capacity

Let $C(t, H, \varepsilon)$ denote the minimal compute required to achieve forecast error $\leq \varepsilon$ over horizon $[t, t+H]$, minimized over feasible algorithms/models. Let K denote the agent's effective capacity (time/energy/compute and coordination).

Definition 6 (Prediction horizon)**.** *A horizon occurs at time t if there exists $H > 0$ such that $C(t, H, \varepsilon) > K$ while $C(t - \delta, H, \varepsilon) \leq K$ for small $\delta > 0$.*

Definition 7 (Singularity Index)**.** *The Singularity Index is*

$$\mathrm{SI}(t, H, \varepsilon) = \frac{C(t, H, \varepsilon)}{K}.$$

Sustained SI ≥ 1 *indicates a predictability horizon for (H, ε).*

7.2 Phaneron interpretation: novelty outruns consolidation

In a *Phaneron* system, horizons appear when novelty influx $\phi(t)$ outruns consolidation rate $\mu(t)$: the dictionary cannot keep up, MDL and conflict rise, and semantic throughput drops. Increasing capacity, improving proposal mechanisms, or inducing better abstractions pushes the horizon outward.

7.3 Operational detector (sketch)

In practice, C is not known; we estimate $\hat{C}$ by fitting rolling cost–error curves and tracking proxy signals: dictionary churn, residual-edge fraction, conflict, WL-color dispersion, and "semantic FPS" (effective update steps per decision). A horizon alarm requires (i) sustained SI crossing, (ii) change-point detection, and (iii) corroborating proxy stress signals.

8 Optional: awareness and intrinsic objective (brief)

Several companion drafts propose a self-anchored definition of awareness and an intrinsic equilibrium objective. We include only the minimal form required for implementation and testing.

Definition 8 (Structural awareness (sketch))**.** *Fix a self pattern s and an active motif set $\mathcal{A}_t$ admitted into the current consolidation pass. A node x is* in awareness *at time t if it is reachable from s via bounded paths whose internal relation nodes lie in $\mathcal{A}_t$, and if x is in the current edit set for consolidation.*

A natural intrinsic objective is to seek equilibrium in the self neighborhood by improving prediction and compression while reducing conflict, subject to floors that prevent degenerate equilibria (e.g., "kill the sensor") and subject to hard value constraints. This can be treated as an engineering hypothesis for inner alignment, not as metaphysics.

9 Applications (short map)

The core substrate is domain-agnostic. The research program includes several applications that instantiate the same ADT:

- **Language (LMMs).** A bidirectional bridge between surface strings and a versioned *Phaneron* meaning graph; evaluate paraphrase invariance, cross-lingual alignment, and explainable QA.
- **Embodied robotics.** Sensors stream into distinctions/connections; affordances and skills are motifs; actions are rewrite schemas; versioning enables replay and safety audit.
- **Multi-agent systems.** Messages are patterns (often Minimal Explanation Subgraphs); common ground grows as overlapping motif sets with compatible predictions.
- **Storage (Phaneron-Store).** ID-less, versioned graph storage at extreme scale via content-addressed state DAGs, pattern-first adjacency, and cross-time compression.
- **Bioinformatics (genomics).** Treat regulatory grammar as reusable motifs over sequence and multi-omics context, targeting interpretable enhancers-to-genes patterns.

10 Validation plan: three wedge demos

A theory of representation that cannot win a small benchmark is not ready to claim unification. We propose three minimal, reproducible validations.

10.1 Demo A: Labeled KG → *Phaneron* → queries still work

Translate a small labeled knowledge graph (RDF-style triples) into unlabeled relation patterns. Measure (i) compression ratio, (ii) query-answer accuracy vs. baseline, and (iii) robustness under predicate renaming and synonym injection. *Success criterion:* structure-only representation preserves query semantics while improving compression/predictive tests.

10.2 Demo B: Streaming novelty and consolidation spikes

Stream observations into a *Phaneron* and induce motifs online. Introduce a controlled distribution shift (new sensor or new domain). Track MDL, residual edges, dictionary size, conflict, and semantic FPS. *Success criterion:* a transient spike followed by consolidation and recovery, with interpretable motifs capturing the shift.

10.3 Demo C: Multi-agent common ground by pattern exchange

Two agents observe partial worlds and exchange Minimal Explanation Subgraphs under a message budget. Measure coordination score, convergence speed, and message efficiency. *Success criterion:* structure exchange yields faster alignment than token exchange under comparable budgets.

11 Related work (brief)

The *Phaneron* intersects multiple lines of work: Minimum Description Length and algorithmic compression; the Information Bottleneck; neurosymbolic and knowledge-graph representations; predictive processing and active inference; and graph canonicalization/WL invariants. The distinguishing choice is not "graphs" per se but a *single unlabeled layer* in which all semantics are patterns rather than labels, enabling uniform rewrite/compress/predict operations.

12 Limitations and open questions

- **Complexity.** Subgraph matching is costly; WL invariants are incomplete; practical systems require heavy approximation and caching.
- **Noise and spurious motifs.** MDL helps, but real data can induce brittle patterns; robust priors and cross-version tests are critical.
- **Time and dynamics.** Encoding temporal structure purely as motifs is principled but needs clearer best practices and benchmarks.
- **Safety and values.** "Hard constraints" need a concrete, non-brittle instantiation; value learning remains open.
- **Scope control.** Physics-scale extensions are conceptually interesting but must not be conflated with the core testable framework.

13 Conclusion

The Distinction–Connection substrate (*Phaneron*) proposes a disciplined minimalism: represent all semantics as structure in an unlabeled undirected graph; learn by inducing reusable motifs that compress and predict; evaluate minds by reflection parity against task-sufficient quotients; and detect singularities as resource-bounded predictability horizons. The immediate next step is not more grand claims but reproducible wedge demos. If the substrate cannot outperform or match label-heavy baselines on small, clear tasks, it should be revised. If it can, the path opens toward larger applications in language, robotics, and collective intelligence.

A Appendix A: Notation cheat sheet

$G = (V, E_G)$: base pseudograph.
$P = (G, \mathcal{E}, D, R, \mathcal{L})$: Phaneron ADT with event log $\mathcal{E}$, dictionary D, rewrites R, objective $\mathcal{L}$.
$N_r(v)$: r-hop neighborhood.
$\chi_t(v)$: WL color at refinement step t.
$\pi_{B,\mathcal{G}}$: task-indexed quotient map of the world.
$C(t, H, \varepsilon)$, K: minimal compute vs. capacity; $\mathrm{SI} = C/K$.

B References

References

[1] J. Rissanen. "Modeling by shortest data description." *Automatica*, 1978.

[2] P. Grünwald. *The Minimum Description Length Principle.* MIT Press, 2007.

[3] N. Tishby, F. Pereira, and W. Bialek. "The information bottleneck method." 1999.

[4] K. Friston. "The free-energy principle: a unified brain theory?" *Nat. Rev. Neurosci.*, 2010.

[5] B. Baars. *A Cognitive Theory of Consciousness.* Cambridge University Press, 1988.

[6] S. Dehaene. *Consciousness and the Brain.* Viking, 2014.

[7] B. Weisfeiler and A. Leman. "The reduction of a graph to canonical form and the algebra which appears therein." 1968.

[8] H. Ehrig et al. *Fundamentals of Algebraic Graph Transformation*. Springer, 2006.

[9] R. Besold et al. "Neural-Symbolic Integration: A Roadmap." *arXiv*, 2017.

[10] P. Hitzler et al. *Foundations of Semantic Web Technologies*. Chapman & Hall/CRC, 2009.

[11] T. Brown et al. "Language Models are Few-Shot Learners." *NeurIPS*, 2020.

[12] J. Drescher. "The Distinction–Connection Substrate (Phaneron): Undirected Graphs for Meaning, Noise, and Emergence." Preprint, 2025.

[13] J. Drescher. "The Phaneron as an Abstract Data Type: Formal Definition, Expressivity, and Algorithms." Preprint, 2025.

[14] J. Drescher. "Reflection Parity: Task-Sufficient Coarse-Graining in Phaneron Minds." Preprint, 2025.

[15] J. Drescher. "Predictability Horizons: A Resource-Bounded Theory of Singularities." Preprint, 2025.

[16] J. Drescher. "Large Meaning Models (LMMs): Phaneron as a Language-Agnostic Structural Semantics Substrate." Preprint, 2025.

[17] J. Drescher. "Phaneron-Store: ID-less Graph Storage for Single-Layer Semantics at Trillion–Quadrillion Scale." Preprint, 2025.

[18] J. Drescher. "Common Ground by Construction: Pattern Exchange and Collective Intelligence in Phaneron Agents." Preprint, 2025.

[19] J. Drescher. "Embodied Phaneron: A Single-Substrate Architecture for Perception, Affordances, and Safe Skill Composition." Preprint, 2025.

[20] J. Drescher. "The Distinction–Connection–Abstraction Hypothesis: A Structural Missing Link for Human Cognition." Preprint, 2025.