Reliable Semantic Knowledge Derivation in Open Decentralized Information Environments

Conceptual Research Synthesis | Manuscript v1.0 | Not empirical validation

Abstract

Open, decentralized information environments contain incomplete, partially contradictory, and erroneous claims whose sources, contexts, identities, and relations may change. The research problem is not simply to preserve these claims or retrieve them consistently, but to determine how initially unsupported relational hypotheses might become defensible semantic knowledge through independent evidence while retaining provenance, uncertainty, identity continuity, and relational history. This paper reconstructs a bounded conceptual synthesis of work on semantic reference, contextual interpretation, validity across transformation, and relational reuse. It distinguishes the structural conditions that make assessment inspectable from the evidential operations required for justified changes in claim status. It then examines selective retrieval and signal-to-noise ratio as a limited analogy for evidential visibility, without treating ranking, repetition, or apparent coherence as verification. The analysis identifies conditions for a defensible derivation and specifies the empirical demonstration that remains absent. It does not claim that progressive knowledge generation has already been established. A further prerequisite concerns signification: a detectable pattern is not yet an interpreted sign, and an interpreted sign is not yet a verified proposition. The paper therefore distinguishes signal detection, sign interpretation, relational hypothesis formation, and evidential validation without treating them as a mandatory computational pipeline.

1. Research problem and questions

How can an open, decentralized relation system progressively derive reliable semantic knowledge from incomplete, partially contradictory, and erroneous information without losing provenance, uncertainty, identity continuity, or relational history?

Complementary question: Under which conditions can initially unsupported, AI-generated relational hypotheses contribute to demonstrable information gain through independent verification?

The positive research objective is to explain what must change between an unsupported proposal and a supported, revisable claim. A list of safeguards alone is insufficient: preservation of records may protect the history of an error as effectively as the history of a valid proposition. The relevant question is how evidence-bearing change can be distinguished from merely representational change.

2. Research basis, method, and evidential status

This is a bounded conceptual synthesis of an existing, single-lineage research corpus addressing semantic identity, provenance, interpretation, contextual validity, relation recovery and reuse. It is not a systematic literature review, a newly executed experiment, or independent replication. Its method is to reconstruct the dependencies among established conceptual distinctions, bounded implementation observations, and unresolved empirical propositions, without converting any of them into a stronger kind of evidence.

The original research setting was an organically evolved umbrella-brand and artifact ecosystem, with heterogeneous websites, domains, products, services and publications created at different times under different responsibilities. Retrospective identifiers may aid identification but cannot reconstruct actual history or establish relations. The research therefore begins with non-invasive exploration of a partially known, decentralized relation space rather than an assumed unified organizational structure. The broader epistemic question preserves this starting point. [R5, §§1–4, 9–10]

Three epistemic categories are used throughout: (i) conceptual or normative conditions that follow from the definitions and validity constraints under examination; (ii) bounded structural observations showing that particular distinctions can be represented or preserved within examined implementations; and (iii) empirical claims about real-world, progressive knowledge generation that are not yet demonstrated. References to underlying research documents indicate provenance of the conceptual lineage, not independent scientific corroboration.

3. The starting state: incomplete and conflicting relational information

An open information environment does not provide a single complete and authoritative representation of its subject matter. A claim may be absent, unsupported, contradicted, time-limited, or accurate only within a specific operating context. These states must not be collapsed: unknown is not false; contradiction is not automatically error; and agreement is not independent confirmation. A relation proposed by an AI system is a candidate for assessment, not a source-grounded fact merely because it has been expressed in structured form.

The first condition for derivation is therefore explicit epistemic separation between a proposed relation, the source statements from which it was inferred, the context in which it is intended to hold, and any independently observed support or counterevidence. This separation preserves the possibility of correction rather than freezing a hypothesis into an apparently established knowledge object.

4. Semiotic prerequisites of relational interpretation

Before an observed pattern can function as evidence for a semantic relation, its status as a sign and its interpretive context must be distinguished from its mere detectability. Signal-to-noise ratio, in its technical sense, concerns a defined signal, interference, and measurement boundary; it does not determine what a sign means or whether a proposition expressed through it is true. The use of “signal” in retrieval discussions must therefore remain explicit about whether it refers to a detectable pattern, a relevance cue, an interpreted sign, or evidence for a claim. [T3; interpretive application]

Peirce's triadic account distinguishes sign (representamen), object, and interpretant; this provides a conceptual vocabulary for separating a representation from its object-directed interpretation. Saussure's account of signifier and signified similarly cautions against treating an isolated expression as self-sufficient meaning. These are theoretical resources, not evidence that a particular AI system implements either semiotic account. Neither theory alone supplies a criterion for independently verifying the truth of a relational hypothesis. [T1; T2; interpretive application]

For the present research question, a repeated string or highly ranked document may motivate a candidate relation, but neither repetition nor salience warrants promotion of that candidate to knowledge. Interpretation can be ambiguous, context-sensitive, or mistaken; the resulting proposition must remain traceable to its source and open to independent checking, narrowing, contradiction, and revision. Semiotics thus clarifies a prerequisite to epistemic assessment without replacing the central problem of reliable knowledge derivation.

5. Structural conditions that make derivation inspectable

5.1. Semantic identity and provenance

Stable reference identity enables a relation to be followed across representations, revisions, and citations. Provenance permits an assessor to ask which source asserted what, when, and under which representational conditions. Neither identity continuity nor provenance alone validates the content of a claim: a persistent identifier can refer consistently to a false assertion. Their contribution is to make changes and evidential dependencies reconstructable. [R1; R3; bounded conceptual synthesis]

5.2. Interpretation without validity transfer

Human contextual selection and task relevance must remain distinct from source-grounded semantic authority. A relation may be useful to an interpretation without becoming true by selection. Likewise, a source can support a narrow proposition without authorizing every conclusion drawn from it. The distinction between interpretation, relevance, membership, and validity prevents contextual reuse from silently broadening evidential scope. [R2; R3; bounded conceptual synthesis] Semantic autonomy remains intact across established contextual connections: relation does not itself create identity, authority, membership or dependency. [R2, §4A]

5.3. Validity boundaries across transformation

A claim carried from one representation, domain, time, or scale into another requires attention to what its evidence actually covers. The constraint "Claim Reach ≤ Validity Reach" expresses the principle that the scope of an assertion must not exceed the scope of its supporting evidence. Transformation, paraphrase, aggregation, or relation recovery do not independently renew validity. A change in context may require a narrowed claim, a revised qualification, or a return to an unresolved state. [R3]

5.4. Relational history and reuse

Recovering an earlier relation is different from re-establishing its truth. Reuse can reduce reconstruction effort and preserve continuity, but it does not substitute for checking whether the relation still applies. Maintaining relational history matters precisely because subsequent

interpretations and corrections should not erase the route by which a claim became accepted, restricted, or rejected.

6. From relational hypothesis to evidentially supported claim

A defensible derivation can be described as a sequence of epistemically distinguishable changes rather than as a proposed new software architecture. First, a candidate relation is formulated with an explicit scope and recorded dependence on its originating material. Second, relevant source claims and potential counterevidence are identified without treating repeated representations of one source as independent observations. Third, verification asks whether additional evidence actually bears on the candidate proposition and its validity boundary. Fourth, the claim status is revised in a way that preserves the previous state, the reason for change, and remaining uncertainty. [R3; R4; research proposal, not validated outcome]

Possible outcomes include supported within scope, narrowed, contradicted, rejected, or unresolved. An unresolved result is not a failed representation; it is an epistemically distinct outcome. Likewise, an independently supported rejection can contribute to knowledge by correcting an earlier hypothesis. None of these outcomes may be inferred from agreement among AI-generated restatements alone.

The original derivation calls this controlled use of unsupported AI assertions the Judo Principle. Inquiry may verify a previously unknown existing relation, discover a different relation, refute the proposed relation, leave the matter unresolved, or identify a prospective design opportunity. The last outcome concerns possible future design, not discovery of an existing relation. The assertion itself supplies no independent evidence. [R5, §5]

The central transition is not “unstructured information becomes a structured relation,” but “a proposition with insufficient support acquires, loses, or changes justified evidential standing.” Structural inspectability makes this transition auditable; independent evidence makes it potentially epistemically meaningful. The present corpus supports the need to distinguish these steps, not the claim that they have been successfully demonstrated end-to-end in an open environment.

7. Information gain: the positive criterion

Information gain must be evaluated relative to a specified uncertainty and a distinguishable evidential change. A new relation link, a more fluent answer, greater document visibility, or a denser network may improve access or representation while leaving uncertainty about a proposition unchanged. Conversely, independently grounded evidence that restricts an overbroad claim may reduce uncertainty even when it decreases the number of claims treated as valid. [D1; research-specific interpretation]

The term information gain is used here in a specifically epistemic application of the broader concept of uncertainty reduction. The general concept does not require independent external verification in every domain or measurement setting. In the present research, however, an asserted gain in reliable semantic knowledge from initially unsupported relational hypotheses requires a justified evidential change, with independence assessed relative to the sources and processes on which the hypothesis depends. This additional requirement characterizes the present research question; it does not redefine information gain generally. [D1; research-specific boundary]

A defensible information-gain claim would identify the prior epistemic state, the proposition and its scope, the newly acquired evidence and its independence, the revised epistemic state, and the relevant uncertainty reduced. Without these distinctions, increased confidence may be an artifact of repetition or selection. This paper does not provide a validated numerical information-gain measure or measured gain for the examined system.

7.1. Evidential change and demonstrable information gain

The progressive derivation of reliable semantic knowledge cannot be established by the accumulation of relational assertions alone. An initially unsupported relational hypothesis may become more strongly supported, more narrowly qualified, contradicted, or remain unresolved when additional evidence becomes available. Each outcome may contribute to epistemic progress, provided that the change is justified by evidence rather than by repetition, representational consistency, or unsupported reassignment of claim status.

The distinction between source multiplicity and evidential independence is particularly important in decentralized information environments. Several apparently distinct sources may reproduce a common underlying assertion. Their agreement does not necessarily constitute independent corroboration. The provenance and relevant dependencies of supporting material must therefore be examined before its contribution to a claim can be assessed.

For the specific epistemic case examined here, a claim of demonstrable information gain requires a defensible difference between the epistemic situation before and after the introduction of new evidence. Such a difference may consist in increased support for a hypothesis, justified rejection of an erroneous relation, or a more precise determination of its contextual and temporal applicability. The preservation of uncertainty and contradiction remains essential where the available evidence does not warrant resolution. [D1; proposed epistemic criterion, not empirical finding]

Bayesian confirmation theory provides probabilistic accounts of how evidence may alter support for a hypothesis (T5, §§3.1–3.3). Information-theoretic measures, by contrast, characterize specified forms of uncertainty (T3). Neither framework alone establishes reliable semantic knowledge generation in an open, decentralized relation system. Their relevance here is theoretical rather than operational. No Bayesian measurement model or quantitative information-gain metric is adopted or validated.

Epistemic progress is not necessarily equivalent to information gain. Newly identified contradictions may expose previously unrecognized uncertainty without immediately reducing it. Such findings can improve the adequacy of an epistemic assessment while leaving the underlying proposition unresolved. A claim of information gain requires a separately justified reduction of relevant uncertainty.

Consequently, the central empirical question is not whether a system can generate plausible relations or maintain a consistent representation of them. It is whether initially unsupported relational hypotheses can undergo independently justified epistemic change while provenance, uncertainty, identity continuity, validity boundaries, and relational history remain recoverable.

The existing conceptual and implementation-based findings support the investigation of these conditions. They do not yet establish the empirical effectiveness of the complete knowledge-derivation process.

7.2. Operational boundary of an epistemic information-gain claim

Operationalization here means identifying what an empirical assessment would need to distinguish; it does not authorize a new metric, instrument, or modification of any frozen experiment. The unit of assessment is a specified relational proposition with identifiable referents, relation semantics, contextual and temporal scope, and an explicit initial evidential status. The comparison must retain the original hypothesis and its source dependencies rather than retrospectively rewriting the prior state.

For a change to count as independently justified, the assessor must identify the new observation or evidence, its provenance, the extent to which it depends on sources used to generate or previously support the hypothesis, and its applicability to the proposition under assessment. Distinct documents, domains, citations, or model outputs cannot be assumed independent solely because their surface representations differ. Independence is an evidential question, not a document-counting rule.

The resulting assessment distinguishes support, disconfirmation, contextual narrowing, and continued uncertainty. It must state what was learned about the proposition and why that conclusion was not already licensed by the prior evidence. A status change alone is not sufficient. The prior and revised epistemic states must remain reconstructable, including contradictions and failed hypotheses. A defensible reduction of relevant uncertainty may be qualitative where no justified probability model exists; a numerical information-gain value must not be fabricated from categorical statuses.

A bounded example illustrates the distinction without supplying empirical findings. Suppose a model proposes that organization A operated system B during period T on the basis of a secondary description. A second website repeating that description does not by itself independently verify the claim. A separately sourced, dated operational record could justify support for the specified period if its authenticity, reference identity, and scope are established. If that record instead identifies organization C as operator, justified rejection or narrowing of the original hypothesis could constitute epistemic progress. The example is hypothetical; neither outcome is evidence that the investigated relation system achieves information gain.

The research therefore has a testable target but not yet a positive test result: whether independent evidence can produce an inspectable and justified change in the epistemic status of an initially unsupported AI-generated relation, without erasing provenance, uncertainty, identity continuity, validity limits, or relational history. The frozen Relational Resolution execution series cannot be treated as such a test by reinterpretation: its pending raw runs and instrumentation audit do not constitute evidence of information gain or system effectiveness.

7.3. Evidence-case audit: existing research corpus

The following is a bounded audit of the identified research records, not a claim of exhaustive discovery across every repository, archive, or unpublished run. Its question is whether any examined record documents, together, an initially unsupported AI-generated relational hypothesis, a traceable prior evidential state, a genuinely independent verifying or disconfirming observation, and a justified revised epistemic state with preserved provenance, identity, uncertainty, and scope. Absence of a qualifying record in this audit is not evidence that no such record exists elsewhere.

Structural Invariants of the Semantic Core (R1) supplies empirical observations about recurrence and non-recurrence of architectural properties across implementations from a common development lineage. Its seven candidate invariants and six negative findings concern structural patterns, not the independent verification of an initially unsupported AI-generated relational claim. Multiple implementations within that lineage are not independent replication. Audit classification: structural evidence; no qualifying epistemic-transition case identified.

HCL Concept-to-Implementation Re-Freeze (R2) supplies contextual-continuation conditions; Claim Validity Across Transformation (R3) supplies a conceptual boundary between Claim Reach and Validity Reach. Neither record, as examined here, documents a complete independent before/after verification event for an initially unsupported AI-generated relation. Audit classification: conceptual conditions and scope constraints; no qualifying epistemic-transition case identified.

Relational Resolution Continuation Marker v1.0 (R4) states that no valid technical or empirical result exists beyond Test 4.11; it requires 18 separately executed external raw runs before Test 4.12, which is restricted to a technical instrumentation audit. The marker explicitly prohibits inference of Information Gain and treating instrumentation runs as repeated empirical evidence. Audit classification: frozen pre-empirical boundary; no qualifying epistemic-transition case identified.

The present audit therefore finds no complete qualifying case in the examined records. This is a limitation of the available inspected evidence, not a negative experimental result about the proposed system. A positive claim would require inspection of the original hypothesis and its time-stamped source state, the independent evidential record and its dependency lineage, the warranted scope of any change, and preserved prior and revised states. An unresolved, contradicted, or rejected hypothesis must not be silently discarded. No numerical information-gain score is justified by the present records.

Audit status (6 November 2026): source-bounded review only; no new experiment, no change to frozen Relational Resolution conditions, no independent empirical confirmation of progressive semantic knowledge derivation.

8. Selective retrieval, evidential visibility, and signal-to-noise

Even where relevant evidence exists in a distributed information environment, it may not enter a particular assessment. Retrieval, ranking, and summarization expose a bounded subset of sources. The resulting answer can be internally coherent yet evidentially incomplete. Failure to retrieve a relation does not establish its absence; prominence does not establish validity; and multiple citations that share an originating source do not constitute independent corroboration.

Signal-to-noise ratio is a technical measure defined for a specified signal, interference, and measurement boundary. It is used here only as a bounded analogy: retrieval may improve the prominence of task-relevant material while obscuring qualifications or counterevidence. In semantic settings, the designation of a source as “noise” is itself a context-sensitive judgment and cannot be equated with disagreement. A contradictory source may be precisely the information needed to correct a claim’s scope. No semantic SNR metric or controlled measurement is asserted in this paper. The semiotic distinction is essential here: what an algorithm selects as salient need not correspond to a valid interpretation, and a valid interpretation need not provide independent

evidential support. Removing contradictory sources as “noise” may increase apparent clarity while decreasing epistemic adequacy. [T3; bounded analogy, not measured semantic SNR]

This adds an epistemic constraint to the derivation problem: the evidence used to promote a hypothesis must be evaluated not only for apparent relevance and credibility, but also for dependence, omissions, and limits of coverage. Selection quality, evidential independence, and validity remain separate properties. Improved visibility alone is not information gain.

9. Assessment consistency as a subordinate consideration

Comparable evidence can be assessed inconsistently, while consistent assessments can share a systematic error. The distinction between unwanted judgment variability and bias discussed by Kahneman, Sibony, and Sunstein (2021) provides a limited supporting reference for this separation. It does not define the present research problem or establish a result about AI systems. Consistency is useful for inspecting an assessment process, but neither consistency nor consensus is proof of evidential sufficiency or truth. [T4; subordinate analogy]

10. Integrated findings and limits of the evidence

Finding 1 — Preservation is not generation. Identity, provenance, contextual distinction, and history support inspectability; they do not independently convert unsupported hypotheses into knowledge. [R1; R3; synthesis]

Finding 2 — Validity is bounded. Relation reuse and representational transformation cannot silently extend a claim beyond its evidence; correction and narrowing are legitimate epistemic outcomes. [R3]

Finding 3 — Evidence is not equivalent to visibility. Selective retrieval can determine which sources enter an assessment, while repetition can inflate apparent support without independent corroboration. [R6; synthesis]

Finding 4 — Information gain is a distinct achievement. A justified reduction in relevant uncertainty requires a distinguishable new evidential contribution and an explicit comparison of epistemic states. [D1; research-specific criterion]

Empirical boundary — The existing corpus does not independently demonstrate progressive, reliable knowledge derivation from initially unsupported AI-generated hypotheses under real-world open, decentralized conditions. No new experiment, source verification, causal retrieval effect, or quantified gain is claimed here. [R1; R4; no empirical result]

11. Conclusion

The central challenge of reliable semantic knowledge derivation is not the accumulation of more relations or the suppression of inconsistent outputs. It is the evidentially justified transformation of relational hypotheses while preserving the conditions that make their history, identity, uncertainty, and validity inspectable. Existing conceptual and bounded structural work supports the necessity of these distinctions. Selective retrieval explains why the evidence available to an assessment may be narrower than the available information environment; the signal-to-noise analogy cautions against confusing a cleaner selection with stronger evidence. Neither observation replaces independent

verification. A detectable signal can guide inquiry, but sign interpretation and independent evidential support remain separate questions.

The defensible result is conditional: for the initially unsupported AI-generated relational hypotheses examined here, justified changes in claim status require identifiable evidence assessed for independence and applicability, while preserving provenance and validity boundaries. A claim of information gain additionally requires a demonstrable reduction of relevant uncertainty; new contradictions may instead expose uncertainty. Whether these conditions enable progressive reliable knowledge generation in practice remains an empirical question. Subsequent observation may confirm, restrict, or refute the proposed conditions without retrospectively altering the research derivation.

Research provenance and references

Retrospective Comparison — Knowledge Graph Refinement. Paulheim (2017, §§3.1–3.3) distinguishes graph completion from error detection and internal from external data sources. The present synthesis addresses the complementary question of warranted change in relational claim status while preserving provenance, uncertainty, identity continuity and relational history. This comparison establishes thematic proximity, not methodological equivalence, empirical confirmation or scientific novelty. [T7]

Primary research lineage (conceptual and bounded implementation materials; not independent replications): Structural Invariants of the Semantic Core; Human Context Layer — Design Principles and Concept-to-Implementation Re-Freeze; Claim Validity Across Transformation; Relational Resolution — Pre-Empirical Execution Freeze and Continuation Marker; Research Derivation — Synthesis Fixation v1.0; Selective Retrieval and Evidence Visibility — Synthesis Addendum v1.0. Source versions, locators and evidential status are recorded below; internal research records are not represented as peer-reviewed publications. The original derivation is documented in R5; R6 identifies the selective retrieval addendum.

Theoretical foundations: Peirce, C. S. (1931–1958). Collected Papers of Charles Sanders Peirce (C. Hartshorne, P. Weiss, & A. W. Burks, Eds.). Harvard University Press. Saussure, F. de (1916). Cours de linguistique générale (C. Bally & A. Sechehaye, Eds.). Payot. These sources support semiotic distinctions, not empirical claims about AI systems.
Supplementary reference: Kahneman, D., Sibony, O., & Sunstein, C. R. (2021). Noise: A Flaw in Human Judgment. Little, Brown Spark. Included only as a subordinate distinction between judgment variability and bias.

Evidence note: No internally maintained monitoring records, repository activity, website visibility observations, or administrative project states are presented as independent scientific validation.

Source-to-claim evidence register

This register distinguishes externally checked bibliographic anchors, source-grounded internal research findings, and original hypotheses. It does not certify independent empirical validation. Citation keys in square brackets refer to the entries below.

T1 — Peirce. C. S. Peirce, Collected Papers of Charles Sanders Peirce, vol. 2, §228 (text c. 1897; C. Hartshorne and P. Weiss, eds., Harvard University Press, 1932). Supports the triadic sign–object–interpretant distinction; does not verify AI-derived relations. Critical-edition collation was not undertaken.

T2 — Saussure. F. de Saussure, Cours de linguistique générale (first published 1916), C. Bally and A. Sechehaye, eds., with A. Riedlinger; consulted Payot edition, 1971, Première partie, chapitre I, §1, pp. 97–99. Supports the signifier/signified distinction; posthumously edited work.

T3 — Shannon. C. E. Shannon, "A Mathematical Theory of Communication," Bell System Technical Journal 27(3), 379–423 (1948), DOI 10.1002/j.1538-7305.1948.tb01338.x; and 27(4), 623–656, DOI 10.1002/j.1538-7305.1948.tb00917.x. Supports technical communication, uncertainty, and noise distinctions; not a semantic truth metric.

T4 — Judgment noise. D. Kahneman, O. Sibony, and C. R. Sunstein, Noise: A Flaw in Human Judgment. Little, Brown Spark, 2021. ISBN 9780316451406. Supporting conceptual distinction only; not independent evidence for the present system.

R1 — Structural implementation observations. Rights of Robots. Structural Invariants of the Semantic Core: A Cross-Implementation Study of Web-Based Semantic Reference Structures. Version 1.0, September 2026, 49 pp. §§7–11, 13–14. Archive: research/semantic-infrastructure/structural-invariants-of-the-semantic-core.pdf (RoR snapshot, 28 September 2026). Seven candidate invariants, six negative findings, and developmental observations are distinct evidentiary classes. Five primary cases and a subsequent control share one implementation lineage; neither independent replication nor demonstrated knowledge gain.

R2 — Contextual interpretation. HCL Concept-to-Implementation Re-Freeze v2.2 (22 September 2026), §§4A, 10–13. Supersedes v2.1 of the same date, retained as historical provenance. Grounded continuation requires established semantic connection, contextual fit and distinction preservation. The scoped correction §4A preserves semantic autonomy. Neither version is independent empirical evidence of knowledge derivation.

R3 — Validity bounds. Claim Validity Across Transformation v1.0 (research draft; available source), especially section 11, "Claim Reach and Validity Reach." Establishes the proposed scope constraint and distinction between warrant strength and applicability reach; it is an original conceptual argument, not independent empirical corroboration.

R4 — Empirical boundary. Relational Resolution — Pre-Empirical Execution Freeze v1.0 (16 September 2026) and Continuation Marker v1.0 (available sources). Test 4.11 completed; 18 separately executed external raw runs and Test 4.12 instrumentation audit outstanding. Explicitly: empirical result NONE; do not infer information gain or effectiveness.

D1 — Definition alignment. informationgain.de /method/ definition, as documented in the user-supplied website archive: information gain is uncertainty reduction resulting from acquisition, observation, or processing of information. The manuscript treats independently evidenced semantic knowledge derivation as a narrower epistemic application, not the general definition.

T5 — Crupi, V. Confirmation. Stanford Encyclopedia of Philosophy, substantive revision 4 August 2025, §§3.1–3.3, https://plato.stanford.edu/archives/fall2025/entries/confirmation/. Secondary scholarly reference for probabilistic confirmation, not a validated operational model.

T6 — Semantic information (authoritative secondary source). Floridi, L., Stanford Encyclopedia of Philosophy, "Semantic Conceptions of Information," substantive revision 14 January 2022, https://plato.stanford.edu/entries/information-semantic/. Documents competing semantic conceptions, including truth-related accounts; does not establish that uncertainty reduction is universally equivalent to semantic knowledge.

R5 — Research Derivation Freeze v0.1. Working title: Semantic Relations in Organically Evolved Systems. 9 October 2026, §§1–5, 9–10. Original research derivation; both research questions are preserved verbatim. Not independent corroboration.

R6 — Research Derivation — Selective Retrieval and Evidence Visibility. Synthesis Addendum v1.0, 31 October 2026, §§1–3. Bounded analytical supplement to Synthesis Fixation v1.0 (13 October 2026). Supports the distinction between evidence visibility and independent verification; no causal retrieval experiment is claimed.

T7 — Paulheim, H. (2017). Knowledge graph refinement: A survey of approaches and evaluation methods. Semantic Web, 8(3), 489–508. https://doi.org/10.3233/SW-160218. §§3.1–3.3 and 7.2. Retrospective comparison; not independent validation.

OPEN — Empirical and operational requirements. Explicit operational criteria for evidential independence and prior/posterior uncertainty, and independent empirical validation remain outstanding. The SEP entries are secondary scholarly references. No critical-edition collation, exhaustive literature review or external peer review is claimed.

Evidence classification: T = externally checkable theoretical source; R = original research lineage (conceptual or bounded observation); D = domain-definition baseline; OPEN = unresolved source or empirical requirement.