

Provenance and Citation Graphs for Engineering-Review Corpora: Sources, Claims, Versions and Dependency Control

Treating an engineering review as a versioned evidence graph so that every claim can be traced, superseded and audited without flattening citation meaning

Juan Seguro

3Dresyns by Resyner Technologies S.L.

© 2026 Juan Seguro. All Rights Reserved.

Correspondence: info@3Dresyns.com · www.3Dresyns.com

ER-394 · Engineering Review · Version 0.1 · 5 October 2026

Abstract

A conventional reference list shows that sources were cited but not which claims depend on which source versions, how claims were transformed, or whether a later review supersedes an earlier interpretation. This review defines a provenance and citation-graph architecture for engineering-review corpora. It distinguishes source identity from source version, citation from evidential support, and derivation from rhetorical relationship. The proposed graph supports differential review, dependency analysis, stale-source detection and controlled supersession. It is designed for engineering corpora in which claims must remain auditable after repeated revisions and cross-paper synthesis.

Keywords: provenance; citation graph; engineering review; claims; version control; evidence graph; dependency

1. References are not dependencies

A bibliography records source membership but usually loses the relationship between a specific claim and the evidence that supports or limits it. CiTO was developed precisely because citations have different factual and rhetorical roles [1]. For engineering qualification, this distinction is operational: a source that defines a method is not equivalent to a source that reports a measured value or contradicts a transfer assumption.

2. The graph needs four independently versioned objects

The minimum objects are source, source version, claim and review version. A DOI may identify a work while a corrected edition or living web standard changes the content actually consulted. Claims also change across review versions. Treating each as an immutable node allows a new review to derive from an older one without rewriting history.

Table 1. Core graph node types.

Node	Immutable identity	Versioned content
source	DOI/standard/record identifier	edition, correction or retrieval state
claim	claim ID	wording, scope, confidence
review	ER identifier	document version
activity	review/edit event	inputs, agent, date
decision	gate or countersignature ID	accepted object hashes

3. Provenance and citation semantics solve different problems

PROV-O represents entities, activities and derivations [2]; citation ontologies describe relationships such as supports, extends or uses method [3]. A review graph benefits from both. Provenance can show that Claim v2 was generated by an editorial activity using Source A v3 and Claim v1, while citation semantics can state whether Source A supports, disputes or only contextualises the claim.

4. Evidence atoms reduce review cost

A claim should be decomposed only to the level needed for independent verification. Over-atomisation creates thousands of trivial nodes; under-atomisation hides mixed support. A practical evidence atom is a sentence or bounded table assertion with one engineering function, a clear scope and explicit source links. The graph can then identify which atoms must be re-reviewed when a source changes.

Table 2. Edge semantics.

Edge	Meaning	Not equivalent to
supports	source directly supports claim	mentions
contradicts	source provides incompatible evidence	different context
usesMethodIn	method adopted from source	supports result
wasDerivedFrom	new object generated from prior object	citation
supersedes	new claim replaces prior claim for defined scope	deletes

5. Dependency control prevents silent cascade errors

When a synthesis paper depends on conclusions from several earlier reviews, a later correction can propagate. The graph should distinguish direct source support from inherited review-to-review dependency. A change-impact query can then traverse only downstream

claims that depend on the modified node. RO-Crate provides a packaging pattern for bounded machine-readable research objects [4], while FAIR principles motivate persistent identifiers and accessible metadata [5].

6. Supersession must be explicit and reversible

A newer review version should not erase the previous evidence path. Supersedes, narrows, broadens and corrects are different relationships. The ledger should preserve both old and new claims, the reason for the change and the authority that accepted it. This supports auditability and prevents a current conclusion from being mistaken for a timeless fact.

Table 3. Automated graph checks.

Check	Flag condition
orphan claim	no support, derivation or explicit author judgement
stale source	claim points to superseded source version
hidden dependency	synthesis imports conclusion without edge
cycle	mutual dependence with no external evidence anchor
impact set	changed node has unreviewed downstream dependants

7. Audit outcome

A corpus passes provenance review when every material engineering claim resolves to a bounded evidence path, every dependency has a type, and version changes can be propagated deterministically. The graph is not a substitute for scientific judgement; it is the infrastructure that makes that judgement inspectable.

8. Scope boundary and neighbouring reviews

ER-394 owns source, claim, version and dependency provenance. ER-393 owns the material-passport container, ER-395 owns living-review update cadence and version governance, and ER-400 owns corpus-level synthesis. ER-394 provides the traceability method used by those reviews but does not perform their lifecycle or corpus decisions.

9. Claim-level provenance is finer than document-level citation

A document-level reference list answers which sources were consulted; it does not answer which source supports a particular clause or which transformation produced a derived statement. Claim-level provenance therefore represents the assertion as an addressable object and links it to one or more evidence objects with typed relations such as supports, contradicts, extends, usesMethodFrom or suppliesBoundaryFor. CiTO demonstrates that citation relations can carry semantics beyond a generic link [1,3], while PROV-O records derivation and activity history [2]. The graph should preserve both because a citation relation and a production-history relation answer different questions.

10. Multi-source assertions need explicit aggregation lineage

Engineering reviews often contain a sentence whose conclusion depends on several sources. Treating that sentence as if it came from one paper hides the aggregation step. Nanopublication approaches illustrate how individual assertions can be packaged with provenance and publication information; recent work extends that idea to multi-source assertions and tracks the evidence contributing to support or refutation [6]. For an ER corpus, the corresponding rule is simple: when a claim is synthesized, the record should identify the component evidence and the editorial operation that created the synthesis. This allows a later source correction to trigger review of only the affected derived claims.

11. Dependency propagation turns provenance into change control

A graph becomes operational when it can answer impact questions. If a source is retracted, corrected or superseded, the system should enumerate every evidence atom derived from it, every claim that uses those atoms, every table or conclusion that depends on those claims, and every later review that imports the conclusion. The propagation should not automatically invalidate all descendants; instead it should classify dependencies as critical, supporting or contextual. A critical edge requires re-evaluation, a supporting edge may reduce confidence without changing the conclusion, and a contextual edge may require no scientific change. This prevents both silent cascade failure and unnecessary corpus-wide reopening.

12. Versioned claims and reversible supersession

Scientific correction is safer when supersession is additive rather than destructive. The old claim remains addressable, the new claim receives a new version identity, and an explicit relation states why the old one was replaced. The same principle applies to datasets, tables and review versions. RO-Crate can package related research objects and their context [4], while FAIR principles reinforce persistent identification and reusable metadata [5]. In an engineering corpus, reversibility matters because a later audit may need to reconstruct what evidence was available when an earlier decision was made.

13. Per-article bidirectional audit protocol

This first protocol operates at the individual-review level. It samples one article graph in both directions: forward tracing begins with a source and verifies every derived evidence atom and claim, while reverse tracing begins with a high-level conclusion and verifies that each dependency resolves to immutable or versioned evidence. The per-article audit also checks orphan claims, citations with no typed role, stale supersession links and version identifiers that resolve to different bytes. It deliberately does not test corpus-wide dependency cycles or release snapshots; those belong to the corpus-level protocol in section 19. Passing this audit does not prove the science true; it proves that the article reasoning is traceable enough to be challenged, corrected and reproduced.

14. Choosing the granularity of an evidence atom

Too coarse a provenance unit hides dependencies; too fine a unit creates an unmanageable graph. A practical evidence atom is the smallest statement whose support can change independently. A numerical result and its test conditions often belong together because changing the method changes interpretation. A broad conclusion containing several separable propositions should be split. The granularity rule should be stable across the corpus so that one article is not represented at paragraph level and another at individual-word level. The test is operational: if one supporting source were corrected, could the affected assertion be revised without rewriting unrelated content? If yes, the atom is probably sufficiently granular.

15. Graph quality metrics

Provenance graphs require QA just as manuscripts do. Useful metrics include proportion of claims with at least one typed evidence edge, proportion of evidence objects with immutable identifiers or hashes, number of orphan nodes, unresolved version links, cycles in derivation relationships, superseded objects still used by current claims and claims whose dependency path crosses a withdrawn source. These metrics do not judge truth, but they reveal where the corpus cannot explain itself. A release gate can therefore fail on traceability even when the prose appears complete.

16. Human review remains authoritative

Automation can enumerate dependencies, detect broken links and propagate change candidates, but it should not decide whether a scientific conclusion survives a source correction. That decision requires interpretation of redundancy, relevance and mechanism. The graph is therefore a decision-support instrument: it identifies which claims require review and presents the evidence paths. Human reviewers decide whether the finding is unchanged, narrowed, downgraded or withdrawn. Recording that disposition as a new provenance activity completes the loop and preserves accountability.

17. Dependency propagation and impact analysis

A citation graph becomes an engineering control when a change can be propagated to affected claims. If a source is corrected, retracted or superseded, the graph should identify every claim that depends directly on it and every review conclusion that depends on those claims. Propagation should stop where an independent evidence path still supports the conclusion, and it should escalate where the affected source is the sole support. The impact report should distinguish 'source changed', 'claim support changed' and 'conclusion changed'. This prevents a source-level event from automatically invalidating an entire corpus while still ensuring that genuine single-point dependencies are visible. [1-4]

18. Provenance completeness versus evidentiary adequacy

Complete provenance does not guarantee a good claim. A perfectly traceable statement may still rely on a weak or inapplicable source, while a strong study can be poorly linked in the graph. ER-394 therefore separates two gates. The provenance gate asks whether every assertion can be traced to the exact source version and transformation. The evidence gate asks whether that source actually supports the proposition within the required scope. Both must pass for a qualification claim. Keeping the gates separate is valuable operationally because a provenance defect can often be repaired without changing science, whereas an evidentiary defect may require claim narrowing or new literature.

19. Corpus-scale dependency and release audit protocol

This second protocol operates above the individual article and complements, rather than duplicates, the bidirectional audit in section 13. A periodic corpus audit samples both high-degree and apparently isolated nodes. It repeats core integrity checks where necessary - identifier resolution, source-version availability, claim-to-source linkage and supersession chains - then adds controls that only exist at corpus scale: dependency cycles, cross-article impact propagation, reproducibility of graph queries and preservation of a machine-readable release snapshot. The union of sections 13 and 19 is therefore intentional: section 13 tests whether one review can be reconstructed claim by claim; section 19 tests whether the released corpus remains a coherent dependency system. A corpus release should preserve the snapshot needed to reproduce the dependency state that existed when articles were countersigned.

20. Source correction and retraction workflow

When a source is corrected, retracted or materially superseded, the corpus should not rely on manual memory. The affected source node is marked with the event, the exact version and the effective date; direct claim dependencies are then enumerated. Each claim is reviewed for whether the change is irrelevant, narrows support, removes support or reverses the proposition. Downstream reviews are updated only where the affected claim contributes materially to their conclusions. The audit should retain both the old and new dependency states so that earlier publication decisions remain reconstructable. This workflow is especially important for large review corpora because a single widely reused source can otherwise create an invisible cascade of stale claims. [1-4]

21. Claim confidence should not be encoded as provenance

Provenance records where an assertion came from and how it was transformed; confidence records how strongly the evidence supports the assertion. Combining them into one field obscures both. ER-394 therefore recommends separate attributes for source lineage, evidentiary status and reviewer disposition. A claim may have impeccable provenance but weak external validity, or strong evidence but incomplete version tracking. The graph should allow these conditions to coexist. This separation also permits independent reviewers to change an evidentiary judgement without rewriting the historical source lineage, preserving auditability.

22. Release package for a provenance-controlled corpus

A corpus release should preserve the article set, claim-source graph, version manifest, supersession map and a machine-readable list of open findings. Hashes bind the graph to the exact article and source objects that were reviewed. A human-readable release note identifies major dependency changes since the previous release. The package does not need to expose restricted source bytes, but it should preserve identifiers and access conditions sufficient to reproduce the review route. This architecture lets later auditors distinguish an unavailable source from a source that was never identified and a changed article from one that merely inherited a new corpus-level label.

23. Limitations and research priorities

Claim-level provenance increases audit power but also creates annotation cost and the possibility of false precision. Research priorities are methods for choosing evidence-atom granularity, automated assistance that preserves human verification, and empirical measures of how provenance completeness affects review error and update time. Corpus-scale stress tests should simulate source correction, retraction, inaccessible links, superseded articles and conflicting source versions, then measure whether the graph identifies the correct downstream claims without excessive false alarms. Another priority is separation of public provenance from restricted evidence where source bytes cannot be redistributed. ER-394 therefore treats provenance infrastructure as a change-control system whose performance should be tested, not as an assumption that more links automatically produce better science. [1-6]

References

1. Shotton D. CiTO, the Citation Typing Ontology. *Journal of Biomedical Semantics*. 2010;1(Suppl 1):S6. doi:10.1186/2041-1480-1-S1-S6.
2. Lebo T, Sahoo S, McGuinness D, et al. PROV-O: The PROV Ontology. W3C Recommendation; 30 April 2013. <https://www.w3.org/TR/prov-o/>.
3. Peroni S, Shotton D. FaBiO and CiTO: ontologies for describing bibliographic resources and citations. *Journal of Web Semantics*. 2012;17:33-43. doi:10.1016/j.websem.2012.08.001.
4. Sefton P, O Carragain E, Soiland-Reyes S, editors. RO-Crate Metadata Specification 1.1. Research Object Crate Community; 2023. doi:10.5281/zenodo.7867028.
5. Wilkinson MD, Dumontier M, Aalbersberg IJJ, et al. The FAIR Guiding Principles for scientific data management and stewardship. *Scientific Data*. 2016;3:160018. doi:10.1038/sdata.2016.18.
6. Menotti L, Marchesin S, Giachelle F, Silvello G. Provenance-driven nanopublications: representing source lineage and trust networks for multi-source assertions. *International Journal on Digital Libraries*. 2025;26:24. doi:10.1007/s00799-025-00431-x.

Competing interests

The author is Founder, CEO and CTO of 3Dresyns by Resyner Technologies S.L.; ER-394 is a methodological engineering review and does not provide a comparative commercial recommendation.

Data availability

ER-394 is a literature-based engineering review; no new experimental dataset was generated.