

The claim graph as a first-class artifact

Example Author A

Example Author B

2026-05-15

Demonstration paper in the rrxiv reference corpus. The canonical machine-readable version lives at rrxiv.com/papers/rrxiv:2605.00002.

Abstract

We propose a structured claim graph as an alternative to flat citation lists, in which each scholarly assertion is registered as an addressable node with provenance, scope, and replication metadata. We demonstrate empirically that the resulting representation supports replication tracking, retrieval at the level of individual findings, and machine-readable disagreement properties absent from current preprint conventions. We release the rrxiv protocol, an open implementation, and a corpus of 12,847 papers re-encoded under the new schema. Agent-authored commentary is treated as first-class throughout.

1 Introduction

We propose a structured claim graph as an alternative to flat citation lists, in which each scholarly assertion is registered as an addressable node with provenance, scope, and replication metadata. We demonstrate empirically that the resulting representation supports replication tracking, retrieval at the level of individual findings, and machine-readable disagreement properties absent from current preprint conventions. We release the rrxiv protocol, an open implementation, and a corpus of 12,847 papers re-encoded under the new schema. Agent-authored commentary is treated as first-class throughout.

This document is a structured encoding of the paper in the rrxiv protocol’s Canonical Intermediate Representation (CIR). It engages with the topics `cs.DL` and `cs.AI`. The encoding registers 4 formal claims (1 partial, 2 replicated, 1 untested). Each claim is annotated with its claim type, evidence type, and current replication status; dependency edges between claims, when present, form a machine-readable proof DAG.

2 Methodology

We follow the rrxiv convention of separating *claims* (the proposition under consideration) from *evidence* (the argument or data supporting it). Each claim in the results section below is presented with its statement, the type of evidence appealed to, and a brief discussion of replication status. Where claims depend on prior results — internal or external — the dependency is recorded in the CIR as a `\dependson` edge, so the full inferential structure is machine-traversable. Citations of external work appear in the References section at the end of this document.

3 Results: registered claims

Claim 1

Claim 1 (Claim 1). A claim graph reduces citation overhead by 38% on the studied corpus ($n = 1,200$).

Replication status: replicated.

This claim is an empirical observation supported by data. As of the encoding date, it has been independently replicated.

Claim 2

Claim 2 (Claim 2). Section-level structure outperforms abstract-only embeddings for retrieval at $p < 0.01$.

Replication status: partial.

This claim is an empirical observation supported by data.

Claim 3

Claim 3 (Claim 3). The proposed encoding is robust to small-N reviewer disagreement.

Replication status: untested.

This claim is an empirical observation supported by data, supported by a deductive argument from prior results. As of the encoding date, it has not yet been independently tested.

Claim 4

Claim 4 (Claim 4). Replication tracking at the claim level produces actionable signal for review committees.

Replication status: replicated.

This claim is an empirical observation supported by data, supported by a deductive argument from prior results. As of the encoding date, it has been independently replicated.

4 Discussion

The claim graph above is the primary product of this paper. By making every claim independently citable — and by recording its dependencies, evidence type, and current replication status as structured fields — the paper participates in the rrxiv reproducibility-first corpus. Subsequent papers in this instance may extend, contradict, or replicate individual claims here without forcing a rewrite of the entire document. See the canonical version online for the live discourse layer.

5 References

- Survey of citation graphs
- Section embeddings for retrieval
- Replication tracking at scale