

Citation graphs are not knowledge graphs

Example Author A

2026-05-17

Demonstration paper in the rrxiv reference corpus. The canonical machine-readable version lives at rrxiv.com/papers/rrxiv:2605.00006.

Abstract

An opinion piece distinguishing the structural commitments of citation networks from those of knowledge graphs proper, with implications for preprint infrastructure. We argue that conflating the two has produced a decade of misnamed retrieval systems.

1 Introduction

An opinion piece distinguishing the structural commitments of citation networks from those of knowledge graphs proper, with implications for preprint infrastructure. We argue that conflating the two has produced a decade of misnamed retrieval systems.

This document is a structured encoding of the paper in the **rrxiv** protocol's Canonical Intermediate Representation (CIR). It engages with the topic **cs.DL**. The encoding registers 1 formal claim (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). Citation networks encode 'this work cited that work' edges; knowledge graphs encode typed semantic relations the two are not interchangeable.

Replication status: untested.

As of the encoding date, it has not yet been independently tested. It depends on 1 prior claim in the same paper.

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

- Knowledge graphs from scientific abstracts
- Citation networks vs knowledge graphs