

Tree of Reviews Robustness to Noise and Adversarial Perturbations in Multi-Hop QA

Assignee Research

May 30, 2026

Abstract

This report synthesises findings from 10 peer-reviewed papers addressing the following research question: How robust is the Tree of Reviews framework to noisy or adversarially perturbed documents in multi-hop QA compared to flat retrieval methods, measured by F1 score on TriviaQA. Multi-hop question answering is a knowledge-intensive complex problem. Large Language Models (LLMs) use their Chain of Thoughts (CoT) capability to reason complex problems step by step, and retrieval-augmentation can effectively alleviate factual errors caused by outdated and. 10 claims were extracted from source literature; 3 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 5.3/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: Tree of Reviews: A Tree-based Dynamic Iterative Retrieval Framework for Multi-hop Question Answering. Research question: How robust is the Tree of Reviews framework to noisy or adversarially perturbed documents in multi-hop QA compared to flat retrieval methods, measured by F1 score on TriviaQA?.

2 Methodology

Systematic literature search across multiple databases yielded 10 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 5.3/10.

3 Results

10 papers retrieved. 10 claims extracted; 3 independently verified. Quality review score: 5.3/10.

4 Limitations

This report is a machine-generated literature synthesis and does not constitute original research. Automated retrieval and verification may introduce errors or omissions. Review scores reflect automated assessment, not human peer review. Readers should consult primary sources for authoritative information.

5 Extracted Claims

Claim	Verified	Confidence
TOR achieves state-of-the-art performance in both retrieval and response generation on three different multi-hop questions	✓	0.31
Tree of Thought (ToT) enhances the problem-solving capabilities of Large Language Models (LLMs) by introducing a tree-like structure	×	0.08
Asai et al. (2020) trained a retriever that dynamically retrieves information from Wikipedia graphs.	×	0.06
The decomposition of the question and the construction of the tree lack the assistance of external knowledge and information	×	0.06
TOR is the first to propose a retrieval framework that uses a tree-like structure to dynamically initiate requests based on external knowledge	×	0.14
LLMs can decide dynamically whether to initiate further retrieval and what requests to generate based on external knowledge	×	0.08
TOR introduces a tree structure to handle each retrieved paragraph separately, alleviating the misleading effect of irrelevant information	✓	0.31
The diversity of reasoning path extension reduces the impact of a single reasoning error on the whole.	✓	0.24
TOR proposes two tree-based search optimization strategies: pruning and effective expansion.	×	0.06
Pruning and effective expansion strategies demonstrate significant improvements in reducing time overhead and increasing accuracy	×	0.03

References

- <http://arxiv.org/abs/2404.14464v1>
- <http://arxiv.org/abs/2603.05462v1>
- <http://arxiv.org/abs/2510.25621v1>