

How does integrating generative refinement impact answer coherence scores on the CAR benchmark for queries with

Assignee Research

June 10, 2026

Abstract

While Retrieval-Augmented Generation (RAG) mitigates hallucination and knowledge staleness in Large Language Models (LLMs), existing frameworks often falter on complex, multi-hop queries that require synthesizing information from disparate sources. Current advanced RAG methods, employing iterative or adaptive strategies, lack a robust mechanism to systematically identify and fill evidence gaps, often propagating noise or failing to gather a comprehensive context. We introduce FAIR-RAG, a novel agentic framework that transforms the standard RAG pipeline into a dynamic, evidence-driven reasoning

1 Introduction

This paper examines: FAIR-RAG: Faithful Adaptive Iterative Refinement for Retrieval-Augmented Generation. Research question: How does integrating generative refinement impact answer coherence scores on the CAR benchmark for queries with abstract facets compared to standard retrieval-augmented generation?.

2 Methodology

Systematic literature search across multiple databases yielded 14 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 4.5/10.

3 Results

14 papers retrieved. 16 claims extracted; 1 independently verified. Quality review score: 4.5/10.

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
The Reciprocal Rank Fusion (RRF) algorithm is used to re-rank documents without requiring hyperparameter tuning.	×	0.03
The RRF algorithm produces a single, robustly ranked list of the top-5 most relevant documents as candidate evidence.	×	0.03
An LLM agent evaluates each document’s utility with respect to the original user query and discards irrelevant, off-topi	×	0.03
The Structured Evidence Assessment (SEA) module uses a checklist-based methodology to perform granular gap analysis.	×	0.08
The SEA module deconstructs the user’s query into a checklist of discrete, required informational components or ‘finding	×	0.08
The SEA module systematically audits the collected evidence against the checklist, confirming which findings are support	×	0.08
The evidence is deemed sufficient only if all required findings are confirmed.	×	0.08
The ‘unchecked’ items on the checklist constitute a direct, interpretable, and actionable signal for the subsequent Quer	×	0.05
FAIR-RAG is designed to robustly handle complex queries by orchestrating a dynamic, multi-stage workflow.	×	0.08
FAIR-RAG begins with an Adaptive Routing module that analyzes query complexity to determine an optimal execution path.	×	0.07
For non-trivial queries, FAIR-RAG initiates a cyclical process designed to progressively build and validate a context.	×	0.06
The core of the FAIR-RAG cycle is an Iterative Refinement loop where LLM agents intelligently decompose information need	×	0.14
Each cycle in FAIR-RAG culminates in a Structured Evidence Assessment (SEA) module, which acts as an analytical gating m	✓	0.22
The SEA module emulates a cognitive workflow by first deconstructing the user’s query into a checklist of required findi	×	0.08
The SEA module systematically synthesizes the retrieved evidence against the checklist, verifying what is confirmed and	×	0.07
The identified gaps in the SEA module provide a precise, actionable signal for further refinement.	×	0.10

References

- <http://arxiv.org/abs/2109.08925v2>
- <http://arxiv.org/abs/1701.06209v4>
- <http://arxiv.org/abs/2510.22344v1>