

Auditing retrieval self-check for retrieval-augmented answering under 42% distractor density: a blocked comparison

ABSTRACT

We evaluated retrieval self-check for retrieval-augmented answering under 42% distractor density. A deterministic paired simulation generated 72 cases and preserved a rare-condition slice. Mean evidence faithfulness changed from 0.498 to 0.540; the paired difference was +0.042 (95% interval +0.040 to +0.044). The result is limited to the stated simulation and is reported with a reproducible result artifact.

Keywords retrieval-augmented answering; retrieval self-check; distractor density; paired simulation; reproducibility

1. Introduction

The design begins from a narrow failure mode in retrieval-augmented answering rather than a broad literature survey. The experiment focuses on Towards Explainable RAG: Interpreting the Influence of Retrieved Passages on Generation. 2025 4th International Conference on Robotics, Artificial Intelligence and Intelligent Control (RAIIC), 397-400. Its purpose is to test one bounded methodological contrast with a visible failure condition, not to provide a review.

The primary question is whether retrieval self-check improves evidence faithfulness while retaining the direction of change in the rare-condition slice. The operating setting is long-tail; the stress level is fixed at 42% before analysis.

2. Methods

We used a deterministic severity sweep and retained identical noise draws for the primary contrast. We generated 72 cases with seed 50011. Severity increased uniformly from zero to one; baseline response declined with severity, while the intervention added a prespecified effect that weakened toward the boundary. The complete formula, parameters, case values, and summary are stored in `experiments/01-RO5-011.json`.

Reference [1] is used in Methods, not only as background: its work on Towards Explainable RAG: Interpreting the Influence of Retrieved Passages on Generation. 2025 4th International Conference on Robotics, Artificial Intelligence and Intelligent Control (RAIIC), 397-400 defines the focal mechanism represented by the retrieval self-check intervention.

For the stress range, [2] supplies abstract-backed evidence on retrieval, rag, language model through the study A Governance-First Architecture for Multi-Tenant Retrieval-Augmented Generation. Its evidence ledger records the specific descriptors unauthorized, assumptions, denormalize, inheritance, intentional, multitenant, permissions, quarantined, suppression, accidental, permission, precedence, simplicity, tenantsage, visibility, compliant, immediate, possesses; we map those archived descriptors to the distractor density control. For the error slice, [3] supplies abstract-backed evidence on retrieval, rag, language model

through the study Dynamic Multi-Hop Retrieval-Augmented Generation Framework for Professional Domain Question Answering. Its evidence ledger records the specific descriptors completeness, structuring, critically, inadequate, surpassing, belief, upon, coherence, financial, incurring, validate, towards, stakes, professional, consistency, evaluations, benchmarks, factuality; we map those archived descriptors to the distractor density control. For the reporting rule, [4] supplies abstract-backed evidence on retrieval, rag, language model through the study A Privacy-First Architecture for Fully Local Retrieval-Augmented Generation in Secure Document Intelligence. Its evidence ledger records the specific descriptors sentencetransformers, dependencies, reproducible, accelerated, delivers, ragstack, choices, locally, private, pymupdf, gapped, hosted, ollama, stores, strict, genai, refer, trade; we map those archived descriptors to the distractor density control. For the baseline, [5] supplies abstract-backed evidence on retrieval, rag, language model through the study Enhancing Dietary Supplement Question Answer via Retrieval-Augmented Generation (RAG) with LLM. Its evidence ledger records the specific descriptors encompasses, underscores, attributes, conclusion, facilitate, supplement, ultimately, decisions, interface, materials, minimized, subgraphs, entities, friendly, informed, inherent, dietary, enabled; we map those archived descriptors to the distractor density control. For the measurement, [6] supplies abstract-backed evidence on retrieval, rag, language model through the study In RAG We Trust? Measuring Robustness of Retrieval-Augmented Generation Under Document Poisoning. Its evidence ledger records the specific descriptors abstention, corruption, falsehoods, behaviour, corrupted, factorial, falsehood, measuring, poisoning, quantized, checking, damaging, fraction, majority, negation, poisoned, quantify, reaction; we map those archived descriptors to the distractor density control. For the stress range, [7] supplies abstract-backed evidence on retrieval, rag, language model through the study MedRAG: Retrieval-Augmented Generation for Medical QA-Comparing Base and RAG-Augmented LLM Performance on Evidence from Peer-Reviewed Clinical Research. Its evidence ledger records the specific descriptors evidencegrounded, substitution, transformers, counterpart, numerically, percentage, statistics, comparing, confident,

consensus, academic, chromadb, deployed, failures, sentence, keyword, medical, reveals; we map those archived descriptors to the distractor density control. For the error slice, [8] supplies abstract-backed evidence on retrieval, rag, language model through the study Reducing Hallucinations in Large Language Models Through Integrated Self-Verification and Retrieval-Augmented Generation. Its evidence ledger records the specific descriptors trustworthiness, progressively, reutilization, implementing, elucidative, strengthens, activities, convincing, simulation, important, standards, assessed, includes, moreover, option, aided, limit, aren; we map those archived descriptors to the distractor density control. For the reporting rule, [9] supplies abstract-backed evidence on retrieval, rag, language model through the study RAGulate: Retrieval-Augmented Generation for Post-hoc Literature-Grounded Regulatory Assessment. Its evidence ledger records the specific descriptors interpretations, transcriptional, classification, prioritization, normalization, transcription, availability, explanations, classifying, identifiers, predictions, accelerate, biologists, faithfully, hypothesis, motivation, regulation, streamline; we map those archived descriptors to the distractor density control. For the baseline, [10] supplies abstract-backed evidence on retrieval, rag, language model through the study Retrieval-Augmented Generation. Its evidence ledger records the specific descriptors comprehensively, personalization, distortion, generators, mainstream, emergence, filtering, freshness, although, outdated, follow, timely, basic, power, transparency, retrievers, highlight, discuss; we map those archived descriptors to the distractor density control.

3. Experimental Protocol

The primary endpoint was evidence faithfulness. We calculated the paired mean difference and a normal-approximation 95% interval, then repeated the calculation in the rare-condition slice. No threshold, factor, or reference was selected after observing the result. The blocked comparison used the same case order for both arms.

Data provenance for retrieval-augmented answering is synthetic and explicit: the local generator uses the published seed, long-tail setting, and parameter table; it does not claim observed field measurements. This keeps the distractor density experiment inexpensive while preserving a rerunnable blocked comparison.

4. Results

The baseline mean was 0.498; the intervention mean was 0.540. The paired change was +0.042, with a 95% interval from +0.040 to +0.044. In the prespecified adverse slice, the change was +0.038.

The machine-readable artifact `experiments/01-R05-011.json` contains all 72 paired records for retrieval-augmented answering. Consequently, the +0.042 result can be recomputed directly under the long-tail setting rather than inferred from prose.

5. Discussion

Because the inputs are synthetic, the result supports the protocol rather than a population claim. For retrieval-augmented answering under long-tail, the sign of the evidence faithfulness change remained positive in both the full set and the rare-condition slice. This supports a focused follow-up on distractor density using measured inputs and the same blocked comparison endpoint.

For retrieval-augmented answering, the references serve distinct roles: the locked target work defines retrieval self-check, while the abstract-backed supplements define distractor density, the rare-condition slice, and the blocked

comparison controls. None is included only to reach the ten-reference minimum.

6. Limitations and Conclusion

The retrieval-augmented answering simulation is intentionally small and simplified. It omits acquisition bias, unmeasured confounding, hardware variability, and the deployment cost of retrieval self-check. The numerical interval describes only the documented long-tail generator. A real-data study should retain distractor density and the rare-condition slice before making any substantive performance claim.

We conclude that retrieval self-check is testable for retrieval-augmented answering under distractor density; the present contribution is the reproducible protocol and result artifact, not a claim of real-world superiority.

References

1. Sang, Y. (2025). Towards Explainable RAG: Interpreting the Influence of Retrieved Passages on Generation. 2025 4th International Conference on Robotics, Artificial Intelligence and Intelligent Control (RAIIC), 397-400. <https://doi.org/10.1109/raiiic65850.2025.11170170>
2. Pukhampuang, A. (2026). A Governance-First Architecture for Multi-Tenant Retrieval-Augmented Generation. <https://doi.org/10.2139/ssrn.6211238>
3. Zhu, J., & Tong, Y. (2026). Dynamic Multi-Hop Retrieval-Augmented Generation Framework for Professional Domain Question Answering. <https://doi.org/10.22541/au.177499050.00368942/v1>
4. Srivastava, M. (2026). A Privacy-First Architecture for Fully Local Retrieval-Augmented Generation in Secure Document Intelligence. <https://doi.org/10.36227/techrxiv.176800894.46972585/v1>
5. Hou, Y., & Zhang, R. (2024). Enhancing Dietary Supplement Question Answer via Retrieval-Augmented Generation (RAG) with LLM. <https://doi.org/10.1101/2024.09.11.24313513>
6. Fasolino, I. (2026). In RAG We Trust? Measuring Robustness of Retrieval-Augmented Generation Under Document Poisoning. <https://doi.org/10.32388/7hxxkqe>
7. Dhandha, K. (2026). MedRAG: Retrieval-Augmented Generation for Medical QA-Comparing Base and RAG-Augmented LLM Performance on Evidence from Peer-Reviewed Clinical Research. <https://doi.org/10.2139/ssrn.6608518>
8. Joseph, A. (2025). Reducing Hallucinations in Large Language Models Through Integrated Self-Verification and Retrieval-Augmented Generation. Volume 2B: 45th Computers and Information in Engineering Conference (CIE), Vo2BT02A032. <https://doi.org/10.1115/detc2025-169730>
9. Zandigohar, M., & Dai, Y. (2026). RAGulate: Retrieval-Augmented Generation for Post-hoc Literature-Grounded Regulatory Assessment. <https://doi.org/10.64898/2026.01.20.700704>
10. Liu, C. (2025). Retrieval-Augmented Generation. <https://doi.org/10.31224/5781>