

Auditing retrieval self-check for retrieval-augmented answering under 43% distractor density: a error-stratified estimate

ABSTRACT

We evaluated retrieval self-check for retrieval-augmented answering under 43% distractor density. A deterministic paired simulation generated 72 cases and preserved a median slice. Mean evidence faithfulness changed from 0.542 to 0.582; the paired difference was +0.040 (95% interval +0.038 to +0.042). 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 Robustness of Fine-Tuned LLMs Under Noisy Retrieval Inputs. 2025 6th International Conference on Artificial Intelligence and Electromechanical Automation (AIEA), 417-420. 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 median slice. The operating setting is noisy-input; the stress level is fixed at 43% 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 50091. 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-R05-091.json`.

Reference [1] is used in Methods, not only as background: its work on Robustness of Fine-Tuned LLMs Under Noisy Retrieval Inputs. 2025 6th International Conference on Artificial Intelligence and Electromechanical Automation (AIEA), 417-420 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 Brief Introduction to Retrieval Augmented Generation (RAG). Its evidence ledger records the specific descriptors introduction, omnipotent, dialogues, realize, shuster, aryani, brief, amir, inaccurate, increasing, people, generating, leading, various, application, scenarios, answers, hallucinations; 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 Building a Security and Reliability Evaluation Suite for Retrieval-Augmented

Generation (RAG) Systems. Its evidence ledger records the specific descriptors illustrative, instruments, sacrificing, avoidance, tradeoffs, fairness, monitors, template, toxicity, compute, utility, siloed, hence, suite, miss, conditioning, standardized, calibration; 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 Pool-Gated Retrieval: Beyond Retrieval-Augmented Generation Toward Accountable Evidential Admission. Its evidence ledger records the specific descriptors accountability, epistemically, justificatory, accountable, recognition, commitment, preventing, principled, propagated, registered, structural, widespread, admission, committed, declaring, enforcing, epistemic, exclusive; 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 ACW-RC: Adaptive Confidence-Weighted Retrieval Correction for Robust Augmented Generation. Its evidence ledger records the specific descriptors contradictions, miscalibration, biographical, distribution, eliminating, probability, substitutes, corrective, thresholds, assessing, biography, estimates, evaluator, pubhealth, unlabeled, blending, discrete, expected; 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 FaithGate: A Proposed Architecture for Real-Time Hallucination Detection and Correction in Retrieval-Augmented Generation Systems. Its evidence ledger records the specific descriptors verificationcorrection, pseudoalgorithmic, evaluationmetric, implementability, interdependent, contradictory, independently, specification, disagreement, disconnected, functionally, groundedness, deployments, identifying, regenerates, decomposer, dependence, entailment; 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 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 error slice, [8] supplies abstract-backed evidence on retrieval, rag, language model through the study A2JRAG: Process-Aware Retrieval-Augmented Generation for Public Legal Information Systems. Its evidence ledger records the specific descriptors entityrelationship, representation, twodimensional, demonstration, substantively, procedurally, conditioned, proceedings, transitions, escalation, misaligned, relational, formalize, mandatory, traverses, typically, authored, carolina; we map those archived descriptors to the distractor density control. For the reporting rule, [9] supplies abstract-backed evidence on retrieval, rag, reasoning through the study Operationalizing Evidence Admissibility in Retrieval-Augmented Generation: The DCA-TRAG Architecture. Its evidence ledger records the specific descriptors admissibility, inadmissible, repositories, supersession, eligibility, enforcement, establishes, preliminary, simulations, statistical, downstream, expiration, protection, repository, guarantee, influence, revisions, remained; 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 A Theoretical Analysis of Self-Contained Retrieval-Augmented Generation with Small Language Models. Its evidence ledger records the specific descriptors recommendations, compressing, theoretical, connection, contained, expensive, translate, concerns, concrete, internet, required, together, affects, bounded, develop, finding, helping, amount; 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 median slice. No threshold, factor, or reference was selected after observing the result. The error-stratified estimate 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, noisy-input setting, and parameter table; it does not claim observed field measurements. This keeps the distractor density experiment inexpensive while preserving a rerunnable error-stratified estimate.

4. Results

The baseline mean was 0.542; the intervention mean was 0.582. The paired change was +0.040, with a 95% interval from +0.038 to +0.042. In the prespecified adverse slice, the change was +0.033.

The machine-readable artifact `experiments/01-R05-091.json` contains all 72 paired records for retrieval-augmented answering. Consequently, the +0.040 result can be recomputed directly under the noisy-input 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 noisy-input, the sign of the evidence faithfulness change remained positive in both the full set and the median slice. This supports a focused follow-up on distractor density using measured inputs and the same error-stratified estimate 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 median slice, and the error-stratified

estimate 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 noisy-input generator. A real-data study should retain distractor density and the median 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). Robustness of Fine-Tuned LLMs Under Noisy Retrieval Inputs. 2025 6th International Conference on Artificial Intelligence and Electromechanical Automation (AIEA), 417-420. <https://doi.org/10.1109/aiea66061.2025.11160531>
2. Aryani, A. (2024). A Brief Introduction to Retrieval Augmented Generation (RAG). <https://doi.org/10.59350/8mwjx-hry76>
3. Roy, P., & Roy, D. (2025). Building a Security and Reliability Evaluation Suite for Retrieval-Augmented Generation (RAG) Systems. <https://doi.org/10.20944/preprints202510.0418.v1>
4. Kim, M. H. (2026). Pool-Gated Retrieval: Beyond Retrieval-Augmented Generation Toward Accountable Evidential Admission. <https://doi.org/10.20944/preprints202606.0414.v1>
5. Kumar, C., Deshmukh, S., Khatter, H., Kumar, B., & Pal, Y. (2026). ACW-RC: Adaptive Confidence-Weighted Retrieval Correction for Robust Augmented Generation. <https://doi.org/10.2139/ssrn.6546797>
6. Madhurima, M. (2026). FaithGate: A Proposed Architecture for Real-Time Hallucination Detection and Correction in Retrieval-Augmented Generation Systems. <https://doi.org/10.2139/ssrn.7242858>
7. 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>
8. Martin, T., & Gilchrist, S. (2026). A2JRAG: Process-Aware Retrieval-Augmented Generation for Public Legal Information Systems. <https://doi.org/10.2139/ssrn.7114139>
9. Garg, A., & Esposito, J. (2026). Operationalizing Evidence Admissibility in Retrieval-Augmented Generation: The DCA-TRAG Architecture. <https://doi.org/10.2139/ssrn.7111618>
10. Samal, M. P. (2026). A Theoretical Analysis of Self-Contained Retrieval-Augmented Generation with Small Language Models. <https://doi.org/10.36227/techrxiv.177219989.96070478/v1>