Relevance and Contradiction in Candidate-Gene Retrieval

Martin Wheeler, Philip Fleming, Ross Barker

Abstract

The boundary between useful assistance and an unsupported claim becomes visible when the search system filtering away contradictory findings. The article develops a research agenda through critical comparison of published evidence. The primary literature contributes a bounded result: only the supplied bibliographic scope supports discussion here; specific cohorts, variants, effect sizes, datasets, and analysis algorithms remain unverified. Contradictory evidence can be scientifically valuable even when it is less relevant to a search ranker. Synthesis should inspect why findings disagree: populations, measurements, designs, and genuine incompatibility may differ. Retrieve counterevidence deliberately and preserve its eligibility rules. Test whether conclusions change because contradiction is explained or merely because it is removed. The proposed comparison concerns negative studies, retrieved rankings, candidate narratives; it would test whether the specific failure scenario is detected without erasing legitimate uncertainty. An integrated account should preserve the distinction between candidate regions and cellular annotations and any stronger downstream decision.

The interface under examination

An evidence workflow can be fast, reproducible, and still answer the wrong question. In this review, negative studies provides the entry point for testing that possibility against candidate narratives. Consider the possibility of the search system filtering away contradictory findings. This is a proposed stress scenario, not an observed event reported by the cited studies. Its purpose is to expose which evidence contract would be violated even if the final output looked plausible. For candidate regions and cellular annotations, the relevant claim must identify its population or artifact class, its time window, and the action it is supposed to support. A failure at any of these levels may persist even after a local technical improvement. The analysis therefore treats published methods as bounded evidence and asks how their limitations interact with this particular decision. An operational analogy is useful only if it keeps those limitations visible.

Contradictory evidence can be scientifically valuable even when it is less relevant to a search ranker. Synthesis should inspect why findings disagree: populations, measurements, designs, and genuine incompatibility may differ. A coherent narrative is not the same as a complete one. In this article, the relevant contrast is negative studies, retrieved rankings, candidate narratives. The proposed counterexample of the search system filtering away contradictory findings puts that contrast under pressure. Retrieve counterevidence deliberately and preserve its eligibility rules. Test whether conclusions change because contradiction is explained or merely because it is removed. A favorable finding would show that the intended interpretation of candidate narratives remains defensible when negative studies is perturbed under the specified conditions. An unfavorable finding should narrow the claim rather than be discarded as an inconvenient case.

The transition at issue is specific to exploratory hearing and age-related genetic synthesis. It concerns negative studies, retrieved rankings, candidate narratives, and the scenario of the search system filtering away contradictory findings shows why a locally correct result can still become a misleading integrated claim. The input record should retain enough information to decide which part of that scenario violates the intended contract. That includes the original source type, measurement or artifact status, and the reason the proposed interpretation exceeds its support. The evaluator would compare the system's response with a qualified judgment made without seeing that response. Disagreement should be classified rather than erased

by rewriting the prompt. A coherent narrative is not the same as a complete one. A useful negative result would show that the interface cannot preserve this distinction under realistic use; a useful positive result would demonstrate preservation on independent cases rather than merely on the example used to develop the rule. The primary source remains relevant through its particular contribution: Its stated scope separates candidate nomination and cellular plausibility from confirmed causal mechanisms. That contribution should be retained even if the integrated workflow fails. Conversely, successful integration would not remove the source's limitations or turn a proposed extension into an already observed result. This asymmetric interpretation protects both the literature and the user from an unjustified all-or-nothing verdict.

Evidence without scope inflation

(Xiao Yu et al. 2026b) examines how exploratory cross-trait analysis might nominate candidate regions shared by hearing loss and age-related disease and connect them to inner-ear contexts in Exploratory Cross-Trait Genetic Analysis of Hearing Loss and Age-Related Diseases Nominates Cardiometabolic Candidate Regions and Inner-Ear Cellular Contexts. Its methodological basis is the supplied title identifies exploratory cross-trait genetic analysis and cellular-context nomination; an authoritative abstract was not retrieved. The accessible evidence is specific: only the supplied bibliographic scope supports discussion here; specific cohorts, variants, effect sizes, datasets, and analysis algorithms remain unverified. Its stated scope separates candidate nomination and cellular plausibility from confirmed causal mechanisms. Without accessible methods and results, the work cannot substantiate named loci, causal direction, treatment targets, or clinical risk improvements in this review. Reading the work through evidence conflict changes the evaluation question: the issue is whether negative studies supports the intended interpretation of candidate narratives, rather than whether the reported method is impressive in its original task.

A useful boundary case is supplied by DRAFT-RL: Multi-Agent Chain-of-Draft Reasoning for Reinforcement Learning-Enhanced LLMs, concerned with how multi-agent reinforcement learning can preserve structural diversity instead of repeatedly refining a single answer (Yuanhao Li et al. 2026b). Its methodological basis is agents generate several chain-of-draft trajectories, peer agents and a learned reward model select promising paths, and actor-critic updates feed those selections into later reasoning. The accessible evidence is specific: the paper evaluates code synthesis, symbolic mathematics, and knowledge-intensive question answering and reports gains in accuracy and convergence. It treats alternative drafts as an explicit exploration space rather than incidental sampling noise. Peer agreement can amplify correlated errors, and additional drafts increase inference and evaluation costs unless diversity is measured rather than assumed. This source enters the present review as a methodological comparison for evidence processing, not as evidence about disease biology, patient outcomes, or clinical efficacy. The transferable question is whether negative studies remains correctly interpreted after a computational choice involving candidate narratives. Its task-specific validation cannot be borrowed to validate the clinical claim.

The relevant question in Towards Explainable RAG: Interpreting the Influence of Retrieved Passages on Generation is how generated answers can be interpreted through the influence of their retrieved passages (Yinghao. “Towards Explainable RAG: Interpreting the Influence of Retrieved Passages on Generation.” *2025 4th International Conference on Robotics Sang 2025). Its methodological basis is the supplied IEEE conference title frames passage-influence analysis for explainable retrieval-augmented generation; implementation particulars were unavailable. The accessible evidence is specific: the bibliographic item is located in the official IEEE record; no unverified attribution algorithm or benchmark result is inferred. It asks for traceability between particular retrieved evidence and generated statements instead of treating a bibliography as an explanation. Passage influence, citation support, causal explanation, and factual correctness are distinct properties; the accessible record cannot establish all of them. This source enters the

present review as a methodological comparison for evidence processing, not as evidence about disease biology, patient outcomes, or clinical efficacy. The transferable question is whether negative studies remains correctly interpreted after a computational choice involving candidate narratives. Its task-specific validation cannot be borrowed to validate the clinical claim.

The sources should be compared through the roles they play, not merged into a single performance claim. Exploratory Cross-Trait Genetic Analysis of Hearing Loss and Age-Related Diseases Nominates Cardiometabolic Candidate Regions and Inner-Ear Cellular Contexts addresses how exploratory cross-trait analysis might nominate candidate regions shared by hearing loss and age-related disease and connect them to inner-ear contexts (Xiao Yu et al. 2026b). By contrast, Towards Explainable RAG: Interpreting the Influence of Retrieved Passages on Generation addresses how generated answers can be interpreted through the influence of their retrieved passages (Yinghao. “Towards Explainable RAG: Interpreting the Influence of Retrieved Passages on Generation.” *2025 4th International Conference on Robotics Sang 2025). That difference matters for evidence conflict: the evidence concerning negative studies is not automatically evidence concerning candidate narratives. A successful software operation, a plausible mechanistic account, and a validated outcome are different achievements. Their integration may improve traceability, but it cannot create an unmeasured link. A coherent review should identify where the sources complement one another and where their scopes remain incompatible. Here the proposed failure scenario is the search system filtering away contradictory findings. The analytical value of that scenario is that it can occur without requiring any cited paper to be wrong: an otherwise valid local finding can be misused at a later interface. Preventing that misuse requires evidence about the interface itself, not simply a larger collection of local results.

Checks for the proposed workflow

Responsible health AI connects prediction with workflow, monitoring, and patient impact (Wiens et al. 2019). The point at which a model reaches practice is not the end of evaluation. New populations, changing measurements, and feedback can alter failure modes. A responsible deployment needs a route for recognizing deterioration and a person authorized to respond. Monitoring should be linked to feasible corrective actions; an alert that cannot change the workflow provides accountability only in a very limited sense. For the comparison involving negative studies, retrieved rankings, candidate narratives, the source supplies an analytical obligation. The review's proposed application remains a hypothesis, with the deliberately difficult scenario serving as a check on scope inflation.

Policy distillation connects compression with sequential decision behavior (Rusu et al. 2015). Compressing a policy concerns sequential behavior rather than one isolated prediction. A small local mismatch can compound across actions, while a locally similar policy can explore differently. Evaluation should retain trajectories and consequential state transitions. The review treats policy transfer as a potential operational tool whose task-level contract still needs checking in the destination environment. An implementation should expose this limitation to reviewers of candidate regions and cellular annotations. The result and its evidence qualification should remain recoverable even when the interface condenses them into a short report.

BM25 remains an important lexical baseline for retrieval systems (Robertson and Zaragoza 2009). Lexical retrieval remains a useful baseline because it is explicit about matching the available terms. Semantic matching can add recall while confusing specialized senses that share broad context. A retrieval evaluation should include both exact-identifier and meaning-based cases. Comparing a new retriever with an appropriately tuned lexical baseline helps determine whether extra computation adds support rather than merely more plausible-looking passages. An implementation should expose this limitation to reviewers of

candidate regions and cellular annotations. The result and its evidence qualification should remain recoverable even when the interface condenses them into a short report.

Retrieval-augmented generation makes external evidence part of the generation process (Lewis et al. 2020). Making documents available to a generator changes its information access but does not guarantee that a statement is supported. A useful analysis separates document selection from generation conditioned on those documents. Missing evidence, irrelevant context, and selective quotation create different failures. The proposed workflow should preserve retrieved passages and evaluate claim support independently of fluency or the mere presence of citations. This is methodological support for examining evidence conflict, not a claim that the source studied exploratory hearing and age-related genetic synthesis. A separate application-specific test is required at the interface described in this review.

Dense retrieval provides semantic matching but must be compared with strong lexical baselines (Karpukhin et al. 2020). Semantic retrieval can connect questions and passages whose wording differs. That advantage should be tested against cases where wording differs because the underlying population or mechanism differs. Retrieval accuracy is not identical to answer accuracy, and neither automatically establishes evidence eligibility. A good evaluation retains the ranked candidates so later errors can be localized to retrieval rather than attributed to the generator by default. Within exploratory hearing and age-related genetic synthesis, this distinction is relevant to evidence conflict. The proposed failure scenario is the search system filtering away contradictory findings; its evaluation must preserve the original source's methodological scope.

Replication and transparent workflows are prerequisites for reliable biomedical claims (National Academies 2019). Reproducing a calculation and independently replicating a finding are different achievements. Recoverable data and code allow one to inspect the first; independent observations and design choices inform the second. A review should avoid calling several reanalyses of one dataset separate replications. Clear reporting of analysis changes also makes a failed replication informative instead of allowing it to disappear as an implementation dispute. This is methodological support for examining evidence conflict, not a claim that the source studied exploratory hearing and age-related genetic synthesis. A separate application-specific test is required at the interface described in this review.

FAIR principles support reusable cohorts, signals, and analysis metadata (Wilkinson et al. 2016). Reusable evidence needs more than a downloadable file. Identifiers, metadata, provenance, and meaningful access conditions help another researcher determine what an artifact contains and whether it can answer a new question. Reuse does not mean universal permission or universal scientific comparability. The information needed to judge those boundaries should travel with the artifact rather than remain in private workflow knowledge. The connection to relevance and contradiction in candidate-gene retrieval is a constraint on interpretation. It would be inappropriate to borrow the source's methodological authority as direct confirmation of the integrated domain claim.

Causal inference requires a clearly specified intervention, time order, and exchangeability assumptions (Hernán and Robins 2020). A prediction can be accurate without identifying what an intervention would change. A useful causal analysis begins by defining the hypothetical comparison, eligibility, and time zero. Its assumptions concern the process generating exposure and observation, not merely the flexibility of a fitted model. In this review, causal vocabulary is appropriate only when that comparison can be specified and defended. The connection to relevance and contradiction in candidate-gene retrieval is a constraint on interpretation. It would be inappropriate to borrow the source's methodological authority as direct confirmation of the integrated domain claim.

Epidemiologic design makes confounding, selection, and measurement error part of interpretation (Rothman, Greenland, and Lash 2008). Confounding, selection, and measurement error are distinct problems and need not move an estimate in the same direction. An analysis should describe them before deciding which adjustment is adequate. A larger dataset can reduce random error while leaving systematic error almost untouched. Precision should therefore be interpreted alongside the quality of the design and the eligibility of its population. This is methodological support for examining evidence conflict, not a claim that the source studied exploratory hearing and age-related genetic synthesis. A separate application-specific test is required at the interface described in this review.

False-discovery control is central when many molecular features are screened (Benjamini and Hochberg 1995). Multiplicity arises from the family of questions actually searched. Controlling false discoveries within a named family does not automatically cover later tissue filters, alternative endpoints, or narrative choices. The proposed audit should retain exploratory branches so the family is visible. Independent confirmation then addresses a different question from whether a discovery threshold was correctly applied during screening. For the comparison involving negative studies, retrieved rankings, candidate narratives, the source supplies an analytical obligation. The review's proposed application remains a hypothesis, with the deliberately difficult scenario serving as a check on scope inflation.

Research priorities and negative results

To make relevance and contradiction in candidate-gene retrieval empirically criticizable, the next study could proceed as follows. Retrieve counterevidence deliberately and preserve its eligibility rules. Test whether conclusions change because contradiction is explained or merely because it is removed. These procedures are proposed, not completed. The case set would include the scenario of the search system filtering away contradictory findings, with its validity judged independently of the system's response. The relevant comparison axes are negative studies, retrieved rankings, candidate narratives.

An evidence-linking protocol would use a set of individually judged claims, eligible supporting passages, plausible but ineligible passages, and genuine counterevidence. The evaluation would retain rankings and intermediate interpretations. Source deletion and substitution would test whether support changes coherently. Fluency, citation presence, source influence, and actual support would be recorded separately, so an answer cannot receive full credit simply for sounding well documented.

For the exploratory genetic source, obtaining and inspecting the primary methods is a prerequisite to stronger attribution. The proposed review can presently test evidence-role preservation, but it cannot validate unverified genetic regions, sample construction, or tissue mappings. Any later biological protocol would need to be designed around the actual reported candidates rather than invented examples assigned to the publication.

Interpretive boundaries

Access limits are material for Exploratory Cross-Trait Genetic Analysis of Hearing Loss and Age-Related Diseases Nominates Cardiometabolic Candidate Regions and Inner-Ear Cellular Contexts, Towards Explainable RAG: Interpreting the Influence of Retrieved Passages on Generation. Where only bibliographic scope or a limited abstract is available, the discussion does not assign unverified datasets, algorithms, effect sizes, or clinical recommendations. This review is a bounded conceptual synthesis rather than a systematic review with an exhaustive search protocol. It cannot establish the effectiveness of the proposed integrated workflow. Its specific contribution is to articulate how evidence conflict changes the interpretation of negative studies, retrieved rankings, candidate narratives. Technical sources support computational comparisons only; they do not validate biological mechanisms or patient benefit. Reviews and correspondence retain their publication roles and are not counted as new experiments. Proposed stress cases are illustrative, not observed failures attributed to the papers. Future work should therefore test the weakest inference described here, preserve contrary results, and disclose the cost of checking rather than treating additional automation as intrinsically beneficial.

Conclusion

Negative studies deserves to remain distinguishable from candidate narratives. That distinction is the central safeguard identified by this review. Retrieve counterevidence deliberately and preserve its eligibility rules. Test whether conclusions change because contradiction is explained or merely because it is removed. The remaining problem can be stated as a testable contrast involving negative studies, retrieved rankings, candidate narratives, with the search system filtering away contradictory findings as a deliberately challenging alternative. Further work should test the comparison on independent cases, measure correction and review costs, and retain failed transitions as evidence. Until then, the integrated account is a research agenda rather than a new empirical guarantee.

References

1. Yu, Xiao, et al. “Exploratory Cross-Trait Genetic Analysis of Hearing Loss and Age-Related Diseases Nominates Cardiometabolic Candidate Regions and Inner-Ear Cellular Contexts.” *Intelligent Translational Medicine*, 2026, https://doi.org/10.1016/j.itm.2026.100002.

2. Li, Yuanhao, et al. “DRAFT-RL: Multi-Agent Chain-of-Draft Reasoning for Reinforcement Learning-Enhanced LLMs.” *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 40, no. 35, 2026, pp. 29530-37, https://doi.org/10.1609/aaai.v40i35.40195.

3. Sang, Yinghao. “Towards Explainable RAG: Interpreting the Influence of Retrieved Passages on Generation.” *2025 4th International Conference on Robotics, Artificial Intelligence and Intelligent Control (RAIIC)*, IEEE, 2025, pp. 397-400, https://doi.org/10.1109/raiic65850.2025.11170170.

4. Wiens, Jenna, et al. “Do No Harm: A Roadmap for Responsible Machine Learning for Health Care.” *Nature Medicine*, vol. 25, 2019, pp. 1337-40.

5. Rusu, Andrei A., et al. “Policy Distillation.” *arXiv*, 2015, arXiv:1511.06295.

6. Robertson, Stephen, and Hugo Zaragoza. “The Probabilistic Relevance Framework: BM25 and Beyond.” *Foundations and Trends in Information Retrieval*, vol. 3, no. 4, 2009, pp. 333-89.

7. Lewis, Patrick, et al. “Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks.” *Advances in Neural Information Processing Systems*, vol. 33, 2020, pp. 9459-74.

8. Karpukhin, Vladimir, et al. “Dense Passage Retrieval for Open-Domain Question Answering.” *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*, 2020, pp. 6769-81.

9. National Academies of Sciences, Engineering, and Medicine. *Reproducibility and Replicability in Science*. National Academies Press, 2019.

10. Wilkinson, Mark D., et al. “The FAIR Guiding Principles for Scientific Data Management and Stewardship.” *Scientific Data*, vol. 3, 2016, article 160018.

11. Hernán, Miguel A., and James M. Robins. *Causal Inference: What If*. Chapman & Hall/CRC, 2020.

12. Rothman, Kenneth J., Sander Greenland, and Timothy L. Lash. *Modern Epidemiology*. 3rd ed., Lippincott Williams & Wilkins, 2008.

13. Benjamini, Yoav, and Yosef Hochberg. “Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing.” *Journal of the Royal Statistical Society: Series B*, vol. 57, no. 1, 1995, pp. 289-300.