

Critique Agents in High-Stakes Counseling Summaries

Abstract

A central question for maternal-infant evidence communication is not simply whether an output is accurate, but whether it survives scrutiny of human accountability. This literature-based conceptual analysis does not report a new experiment. The focal evidence concerns how evidence concerning HIV, pregnancy, and infant feeding should be interpreted in ongoing clinical communication. Accountability requires knowing who can authorize an action, challenge a summary, and respond to error. A critic model does not become an accountable professional by producing a second opinion. Map responsibility and escalation routes. Test whether unsupported claims can be stopped before they enter a consequential decision and whether correction reaches affected users. The proposed comparison concerns reasoning review, sensitive communication, professional oversight; it would test whether the specific failure scenario is detected without erasing legitimate uncertainty. Implications are restricted to the evidence roles relevant to maternal-infant evidence communication, leaving broader efficacy and deployment claims unresolved.

The interface under examination

A narrow reading of critique agents in high-stakes counseling summaries reveals a broader problem of evidence design. The question is what an informed user can conclude from reasoning review when professional oversight remains uncertain. Consider the possibility of a self-correction system treated as a clinical authority. 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 clinical correspondence and changing patient records, 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.

Accountability requires knowing who can authorize an action, challenge a summary, and respond to error. A critic model does not become an accountable professional by producing a second opinion. Automation should expose uncertainty to someone with actual authority. In this article, the relevant contrast is reasoning review, sensitive communication, professional oversight. The proposed counterexample of a self-correction system treated as a clinical authority puts that contrast under pressure. Map responsibility and escalation routes. Test whether unsupported claims can be stopped before they enter a consequential decision and whether correction reaches affected users. A favorable finding would show that the intended interpretation of professional oversight remains defensible when reasoning review is perturbed under the specified conditions. An unfavorable finding should narrow the claim rather than be discarded as an inconvenient case.

The strongest comparison would isolate the disputed transition in reasoning review, sensitive communication, professional oversight. The proposed counterexample, a self-correction system treated as a clinical authority, identifies what should not be accepted as equivalent evidence. 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. Automation

should expose uncertainty to someone with actual authority. 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 relevance is the role of clinical discussion in identifying interpretation and communication obligations around sensitive longitudinal decisions. 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.

Checks for the proposed workflow

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. For the comparison involving reasoning review, sensitive communication, professional oversight, 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.

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. An implementation should expose this limitation to reviewers of clinical correspondence and changing patient records. The result and its evidence qualification should remain recoverable even when the interface condenses them into a short report.

Medical machine learning must fit data provenance, care pathways, and prospective evaluation (Rajkomar, Dean, and Kohane 2019). The clinical relevance of a model depends on how its output changes a care process. Development data can contain information that becomes available only after the intended decision. A review must identify that availability boundary before comparing algorithms. Prospective use also changes who is measured and who is treated, so evaluation of a frozen historical dataset cannot establish the entire future workflow. This is methodological support for examining human accountability, not a claim that the source studied maternal-infant evidence communication. A separate application-specific test is required at the interface described in this review.

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 reasoning review, sensitive communication, professional oversight, 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.

Calibration separates discrimination from confidence used in clinical escalation (Guo et al. 2017). Discrimination and calibration deserve separate evaluation. A score can order cases well yet systematically

overstate its certainty. Aggregate calibration can also hide local miscalibration near the threshold where action changes. An audit should retain predicted probabilities and observed outcomes rather than report only correct classifications, and it should evaluate calibration on data not reused to choose the correction. Within maternal-infant evidence communication, this distinction is relevant to human accountability. The proposed failure scenario is a self-correction system treated as a clinical authority; its evaluation must preserve the original source's methodological scope.

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. For the comparison involving reasoning review, sensitive communication, professional oversight, 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.

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. For the comparison involving reasoning review, sensitive communication, professional oversight, 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.

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. An implementation should expose this limitation to reviewers of clinical correspondence and changing patient records. The result and its evidence qualification should remain recoverable even when the interface condenses them into a short report.

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 human accountability, not a claim that the source studied maternal-infant evidence communication. 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. An implementation should expose this limitation to reviewers of clinical correspondence and changing patient records. The result and its evidence qualification should remain recoverable even when the interface condenses them into a short report.

Evidence without scope inflation

(Xiao Yu et al. 2026c) examines how evidence concerning HIV, pregnancy, and infant feeding should be interpreted in ongoing clinical communication in HIV, Pregnancy, and Infant Feeding. Its methodological basis is a JAMA correspondence contribution rather than a newly reported cohort or randomized intervention. The accessible evidence is specific: the bibliographic record establishes a short correspondence item; inaccessible full-text details are not reconstructed or treated as clinical recommendations. Its relevance is the role of clinical discussion in identifying interpretation and communication obligations around sensitive longitudinal decisions. The correspondence does not provide independent new trial evidence here; this review offers no individualized HIV or feeding advice. Reading the work through human accountability changes the evaluation question: the issue is whether reasoning review supports the intended interpretation of professional oversight, rather than whether the reported method is impressive in its original task.

Towards Explainable RAG: Interpreting the Influence of Retrieved Passages on Generation asks 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 reasoning review remains correctly interpreted after a computational choice involving professional oversight. 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. HIV, Pregnancy, and Infant Feeding addresses how evidence concerning HIV, pregnancy, and infant feeding should be interpreted in ongoing clinical communication (Xiao Yu et al. 2026c). 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 human accountability: the evidence concerning reasoning review is not automatically evidence concerning professional oversight. 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 a self-correction system treated as a clinical authority. 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.

Research priorities and negative results

A research agenda should address the specific negative case of a self-correction system treated as a clinical authority. Map responsibility and escalation routes. Test whether unsupported claims can be stopped before they enter a consequential decision and whether correction reaches affected users. These procedures are proposed, not completed. The case set would include the scenario of a self-correction system treated as a clinical authority, with its validity judged independently of the system's response. The relevant comparison axes are reasoning review, sensitive communication, professional oversight.

A discriminating protocol would specify a preferred interpretation and at least one credible alternative before choosing measurements. It would include negative controls and observations that could favor the alternative. Reviewers would classify whether disagreement concerns data quality, model structure, or the scientific question itself. The comparison should reduce uncertainty about that distinction rather than add more observations equally compatible with every account.

The correspondence contribution should remain a source of interpretive questions, not a dataset of treatment outcomes. A source-review study can assess whether summaries preserve that role. A separate clinical or measurement study would be required to test patient outcomes, with professional oversight and appropriate ethical design. No automated-review score would substitute for that independent evidence.

Interpretive boundaries

Access limits are material for HIV, Pregnancy, and Infant Feeding, 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 human accountability changes the interpretation of reasoning review, sensitive communication, professional oversight. 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

An analysis of critique agents in high-stakes counseling summaries should end with a bounded claim: useful evidence can be preserved without pretending that every downstream inference has been tested. Map responsibility and escalation routes. Test whether unsupported claims can be stopped before they enter a consequential decision and whether correction reaches affected users. The remaining problem can be stated as a testable contrast involving reasoning review, sensitive communication, professional oversight, with a self-correction system treated as a clinical authority 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. "HIV, Pregnancy, and Infant Feeding." *JAMA*, vol. 336, no. 2, 2026, p. 169, <https://doi.org/10.1001/jama.2026.6033>.
2. 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/raiiic65850.2025.11170170>.
3. Rothman, Kenneth J., Sander Greenland, and Timothy L. Lash. *Modern Epidemiology*. 3rd ed., Lippincott Williams & Wilkins, 2008.
4. 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.
5. Rajkomar, Alvin, Jeffrey Dean, and Isaac Kohane. "Machine Learning in Medicine." *New England Journal of Medicine*, vol. 380, 2019, pp. 1347-58.
6. Wiens, Jenna, et al. "Do No Harm: A Roadmap for Responsible Machine Learning for Health Care." *Nature Medicine*, vol. 25, 2019, pp. 1337-40.
7. Guo, Chuan, et al. "On Calibration of Modern Neural Networks." *Proceedings of the 34th International Conference on Machine Learning*, 2017, pp. 1321-30.
8. 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.
9. Lewis, Patrick, et al. "Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks." *Advances in Neural Information Processing Systems*, vol. 33, 2020, pp. 9459-74.
10. 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.
11. National Academies of Sciences, Engineering, and Medicine. *Reproducibility and Replicability in Science*. National Academies Press, 2019.
12. Wilkinson, Mark D., et al. "The FAIR Guiding Principles for Scientific Data Management and Stewardship." *Scientific Data*, vol. 3, 2016, article 160018.