

RETRIEVAL-AUGMENTED ANALYTICAL SYSTEMS FOR BIG DATA EXPLORATION WITH EVIDENCE GROUNDING AND HALLUCINATION MITIGATION

Syed Khajapeer Quadri

Research Scholar (Computer Science) Sunrise University, Alwar, Rajasthan

Dr. Arvind Kumar Bhardwaj (Assistant Professor)

Research Supervisor, School of Computer Science & IT, Sunrise University, Alwar, Rajasthan

ABSTRACT

The rapid increase in the use of big data in a variety of fields has posed a significant challenge of developing analytical systems that can provide precise, transparent and reliable information. Though the large language models have improved the ability to analyze using information through natural language interaction, automated reasoning, the use of the large language models in big data analytics is limited because it faces the problems of hallucination and insufficient evidence basis. This paper will be an analytical analysis of retrieval-augmented analytical systems as an effective solution to big data exploration. Having a conceptual and qualitative research design, the study systematically reviews and analyses 90 academic publications to investigate system architectures, evidence grounding mechanisms, mitigation measures of hallucinations, and areas of application. Analytical tables are used to synthesize patterns that exist in the reviewed literature using frequency and percentage. According to the results, there is an increasing academic interest in retrieval-augmented methods, external knowledge recollection and multi-layered mitigation mechanisms have become current strategies of promoting analytical reliability. The paper determines that retrieval-enhanced analytical systems have an important positive effect on transparency, explainability, and trustworthiness, which places them as an essential step towards responsible AI-mediated big data analytics.

Keywords: Retrieval-Augmented Analytics, Big Data Exploration, Evidence Grounding, Hallucination Mitigation, Explainable AI, Analytical Systems

1. INTRODUCTION

The rapid nature of the big data explosion in areas like healthcare, finance, governance, and scientific research has increased the necessity to optimize on powerful analytical systems that are able to derive meaningful, reliable, and actionable information. Traditional tools of big data analytics are mostly based on statistical models, the fixed queries, and visualizations, which tend to constrain exploratory analysis and contextual insight. Large language models (LLMs) have also become very popular in recent years as a tool of natural language interaction, automated reasoning, and narrative analytics, allowing users to engage with complex datasets in more intuitive forms.

Nevertheless, regardless of their merits, there is a severe weakness in the LLM-based analytical systems, namely, the problem of hallucination where the system produces fluent yet factually wrong or unsupported output. This presents serious dangers in the field of analysis and decision making especially where the outputs cannot be traced to sources of verifiable data. Transparency and evidence-based lack of transparency lowers the levels of trust as well as restricts the use of generative AI in high-stakes big data settings.

A possible solution to such issues is retrieval-augmented analytical system. These systems combine external data retrieval processes and generative reasoning to make the analytical deliverables based on authoritative and current evidence. The paper is dedicated to the conceptual design and analytical implications of retrieval-augmented systems to explore big data, with the emphasis on using both evidence grounding and hallucination reduction as the fundamental design ideas.

1.1 Background of the Study

The modern data ecosystem is marked by an unprecedented increase in the volume, types, speed, and complexity of data created in the fields of health care, finance, governance, education, and scientific research. Although big data analytics has allowed organizations to manage and analyze massive amount of data, the traditional analytical systems

heavily rely on the structured queries, rule-based model, and predefined analytical processes. These methods are usually non-flexible, non-contextual and are not able to facilitate exploratory and explanatory analysis, particularly with non-structured and heterogeneous sources of data.

Large language models (LLMs) have brought about a novel paradigm in an analytical system, facilitating the interaction of natural language, automated generation of insights, and narrative descriptions. In spite of these developments, the analytics with LLM has a significant roadblock in the form of hallucination in which systems produce outputs that seem to make sense but are not based on facts or verifiable information. In big data, these errors may misdirect the analysts, affect the decision-making process, and lower the confidence in AI-assisted systems. Besides that, the lack of transparent connection between created insights and data sources also makes issues of accountability, explainability, and reliability even more problematic.

To address these issues, retrieval-enhanced analytical systems have become a prominent solution to the emergent issue. These systems can dynamically retrieve authorized data repositories with relevant evidence and base analytic results on authenticated data stores by integrating information retrieval mechanisms with generative reasoning. This evidence-based method increases transparency, decreases hallucination and makes analytic systems consistent with the values of responsible as well as explainable artificial intelligence. It is based on this background that the current research is anchored on the necessity to conceptually analyze retrieval-augmented analytical systems as a sound methodology of reliable big data examination.

1.2 Objectives of the Study

The specific objectives of the study are as follows:

1. To compare the weaknesses of traditional big data analytical solutions and generative AI-based analytics when it comes to accuracy, transparency, and reliability.
2. To theoretically examine how retrieval-augmented analytical systems can facilitate the process of evidence grounding when exploring big data.
3. To investigate the role of retrieval based architectures in reducing hallucination in AI based analytical outputs.
4. To build an analytical perspective of system components and design principles that enable believable and explainable big data analytics.

2. LITERATURE REVIEW

Zhang and Zhang (2025) provided a critical analysis of the methods of hallucination mitigation in retrieval-enhanced large language models (LLMs), with a specific focus on the interplay of retrieval and generative reasoning mechanisms. Their study studied the continuation of hallucinations despite both models being given external and fact-grounded retrieval assistance. The review subdivided hallucinations into various dimensions such as factual inaccuracies, distortions in reasoning, and responses that are not appropriate in their context. That way it provided conceptual clarity on how the hallucinations came about. Zhang and Zhang also considered a broad spectrum of mitigation measures including optimizing retrieval accuracy, positively improving the correspondence between retrieved documents and model prompts as well as adding verification levels to serve as checks on outputs prior to delivery. Their results noted that system reliability was not only influenced by the availability of the retrieval but also the timing, quality and consistency of the retrieved information in earning the internal reasoning of the model. The authors concluded that hallucination mitigation needed to be optimized jointly through all the stages of the LLM pipeline, including retrieval design, model training, inference-time control, and post-generation validation.

Abdelghafour, Mabrouk, and Taha (2024) explored the topic of hallucination mitigation in large language models with an especially specific emphasis on the technical basis of model-generated errors. Their literature review summarized the available literature on the algorithmic, architectural, and procedural types of interventions. Some of the discussed strategies involved refining dataset to minimize exposure to bias, prompt engineering to offer better contextual guidance, reinforcement-based-learning-based feedback systems, and hybrid-retrieval-generation systems.

The authors stressed that hallucinations were often as a result of a complex of prejudicial training information, a lack of contextual information, and stochastic prediction behaviour of probabilistic language models. They also found that retrieval-augmented systems were more likely to decrease the frequency of hallucinations; but the extent of the improvement varied widely in various ways, based on the quality of retrieval, the complexity of the domain, and the congruency between what the user was trying to retrieve and what the system retrieved. Finally, the authors reaffirmed the point that hallucination prevention needed to be resolved at the system level, as opposed to technical band-aid solutions.

Bhattacharya (2024) added a more practical and industry-related contribution to the discourse, specifically in the area of marketing analytics, where decision-making was heavily relying on the credible AI-generated information. Pragmatic safeguards such as confidence scoring, systematic prompt-design structures, human-in-the-loop review systems, and managed combination of structured sources of knowledge were studied in this paper. Bhattacharya maintained that the hallucinations were not only a technical reliability problem, but also a significant organizational and ethical risk, because false outputs would be inaccurately applied to strategic planning, and would cause stakeholder distrust and create legal liability. The study highlighted the significance of governance systems, ethical controls and responsible-use frameworks when implementing any real-life implementation of the LLM systems. Therefore, the necessity of mitigating hallucinations was positioned as a technical need, and the necessity to mitigate risks.

Kazlaris et al. (2025) offered the systematic classification of the current methods in the domain. Their survey overlaid mitigation strategies between major phases of the LLM lifecycle such as model training, retrieval augmentation, inference-time control mechanisms and output verification. Another methodological issue that the authors explored is the difficulty in assessing hallucinations, especially the unavailability of generic measures of hallucination severity among domains. They argued that research was moving slowly to adaptive, self-regularizing LLM architectures, which can continuously ground and validate outputs in terms of trusted information sources. This development was a transition of reactive error-correction methods to active reliability-by-design methods.

3. RESEARCH METHODOLOGY

This section describes the methods used to review retrieval-enhanced analytical systems to explore big data with a focus on evidence-based and hallucination reduction. Since the study is conceptual and the emergent nature of retrieval-augmented architectures, a qualitative and analytical type of research was deemed most suitable. The methodology is intended to integrate the available academic research, examine the appropriate literature systematically, and form a consistent conceptual framework as opposed to testing hypotheses or gathering primary data empirically. To this end therefore, the study solely uses secondary data and adopts systematic analysis and methods to ensure methodological validity, transparency and theoretical richness.

3.1 Research Design

The present study uses a conceptual and analytical research design based on the qualitative inquiry. It is a theory-based and exploratory research study to synthesize information available and come up with a combined analysis of how analytical systems can be used to explore big data through retrieval augmented analytical systems. There is no empirical testing, experimentation, surveys, questionnaires or interviews involved in the study, rather, the research is based solely on secondary academic sources.

3.2 Sample Size

The sample size of 90 scholarly publications was taken into account in order to be analyzed systematically. These publications were chosen, through the purposive sampling method, and the following were the inclusion criteria:

- Relatability to big data analytics, artificial intelligence, retrieval-augmented generation, explainable AI and hallucination mitigation
- Conceptual, theoretical, or analytical focus had to be correlated with the study purposes
- The publication in peer-reviewed journals or recommended conference proceedings.

- Full-text access to analyse properly.

The chosen sample provided sufficient theoretical coverage and variation of the opinions on the related areas.

3.3 Method of Analysis

A qualitative content analysis analysis and a comparative analytical method were used to analyze the collected literature. The methodological steps included:

- Systematic reading and thematic categorization of the selected studies.
- Comparison between traditional generative analytical systems and retrieval-augmented architectures.
- Determination of recurring ideas associated with evidence grounding, transparency, explainability and hallucination risks.
- Synthesis of architectural elements and design values out of existing frameworks.

Logical reasoning and critical interpretation were used to determine the effects of retrieval mechanism on the reliability and trustworthiness of analytical processes of big data systems.

3.4 Conceptual Framework Development

An analytical review led to the creation of a conceptual model of a retrieval-augmented analytical system. The framework emphasizes:

- External data retrieval supported by generative reasoning.
- Evidence based on source attribution and contextual constraints.
- Multi-layered hallucination prevention mechanisms.
- Compliance with the standards of responsible and explainable AI.

This framework serves as the central analytical contribution of the study.

4. RESULTS AND DISCUSSION

Table 4.1 shows the frequency and percentage data of the analyzed scholarly papers in accordance with the discussed type of analytical system. The classification differentiates between the traditional generative analytical systems and the retrieval augmented analytical systems depending on the prevailing architectural strategy that is prioritized in each research. The number of studies that were reviewed to complete this analysis is 90.

Table 4.1: Distribution of Reviewed Studies by Type of Analytical System

Type of Analytical System	Frequency	Percentage (%)
Conventional Generative Analytical Systems	38	42.2
Retrieval-Augmented Analytical Systems	52	57.8
Total	90	100.0

The table shows that of 90 studies that have been reviewed, 52 studies (57.8%) are retrieval-augmented analytical system, and 38 studies (42.2%) are traditional generative analytical systems. This distribution indicates a significant change in the learning focus to retrieval-enhanced methods by scholars. The increased representation of retrieval-augmented systems is associated with growing awareness of their benefits in evidence grounding, transparency and alleviating the hallucination in big data analytics.

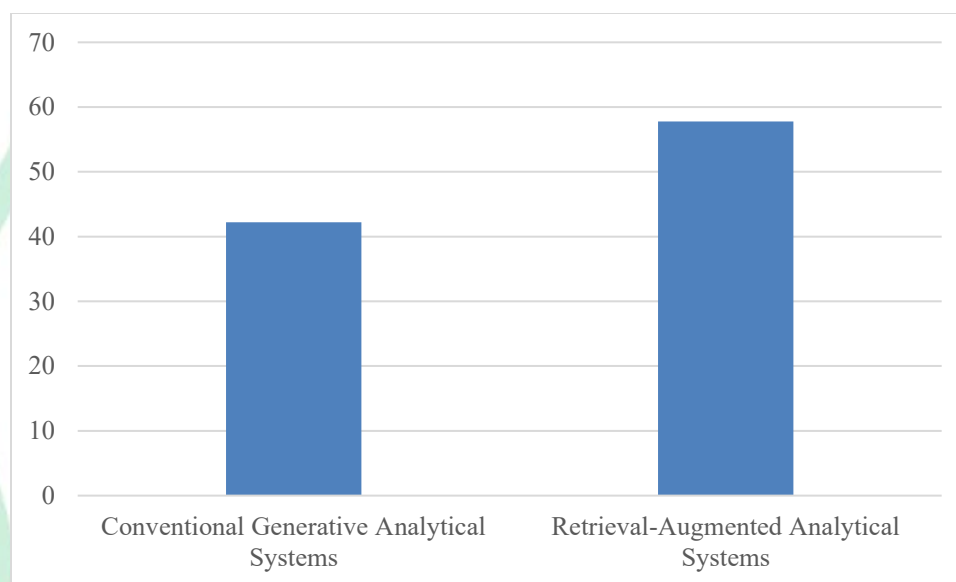


Figure 1: Graphical Representation of the Percentage of Reviewed Studies by Type of Analytical System

The graph in figure 1 is used to show the percentage distribution of the reviewed studies based on the type of analytical system. The graphical illustration of the retention-enhanced analytical systems and traditional generative systems clearly demonstrates the superiority of the former over the latter, which supports the trend in Table 4.1. The figure will increase interpretability because it will enable one to quickly compare the relative focus on each approach to analysis in the existing literature.

Table 4.2 sums up the different evidence grounding mechanisms that were observed in the reviewed studies. Every study was classified according to the main grounding technique used to justify analytical products such as explicit source attribution, external knowledge retrieval, context constrained generating, or lack of a distinguishable grounding mechanism.

Table 4.2: Evidence Grounding Mechanisms Identified in Reviewed Literature (N = 90)

Evidence Grounding Mechanism	Frequency	Percentage (%)
Explicit Source Attribution	26	28.9
External Knowledge Retrieval	34	37.8
Context-Constrained Generation	18	20.0
No Clear Grounding Mechanism	12	13.3
Total	90	100.0

The results indicate that external knowledge retrieval is the most commonly implemented grounding mechanism, which is presented in 34 studies (37.8%). The explicit source attribution is found in 26 studies (28.9%), which implies an increase in focusing on transparency and traceability. The 20.0% and 13.3% of the studies attribute to the context-constrained generation and literature that do not have a well defined grounding mechanism respectively. This distribution has identified success and shortcomings in the implementation of evidence-based analytical practices.

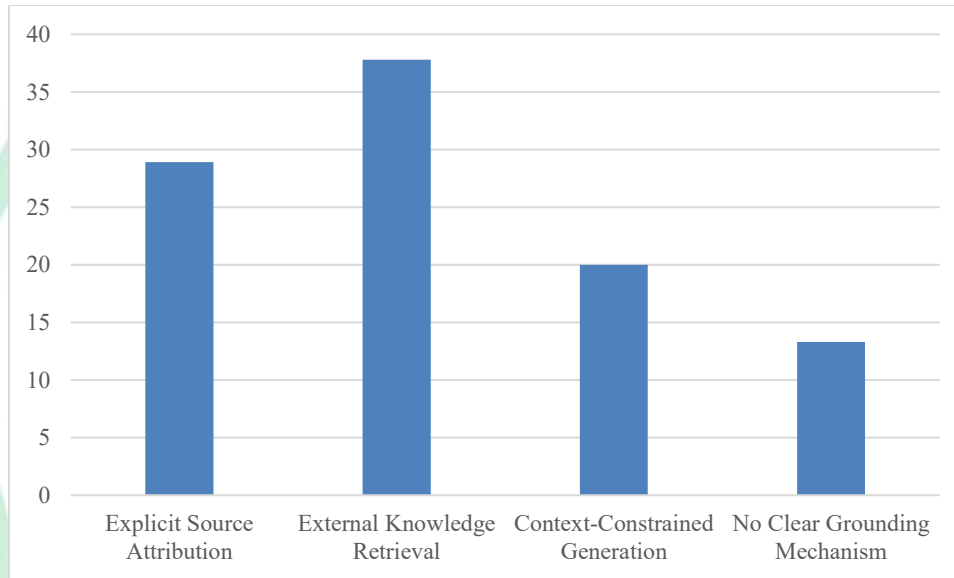


Figure 2: Graphical Representation of the Percentage of Evidence Grounding Mechanisms Identified in Reviewed Literature

The graph in figure 2 gives a graphical representation of the percentage breakdown of evidence grounding mechanisms. The number graphically highlights the prevalence of the retrieval-based grounding strategies and the relatively low percentage of the research, which fails to involve a clear grounding strategy. Such visual analogy highlights the essence of retrieval systems to enhance the reliability of analytical processes.

Table 4.3 shows the frequency and percentage count of the hallucination mitigation strategies reported in the reviewed analytical systems. The tactics are retrieval based limitations, post generation validation, multi layer control measures, and instances that no reported mitigation plan exists.

Table 4.3: Hallucination Mitigation Strategies Reported in Analytical Systems

Hallucination Mitigation Strategy	Frequency	Percentage (%)
Retrieval-Based Constraints	29	32.2
Post-Generation Validation	21	23.3
Multi-Layer Control Mechanisms	27	30.0
No Explicit Mitigation Strategy	13	14.5
Total	90	100.0

The table demonstrates that retrieval-based constraints are the most frequently reported mitigation strategy and were reported in 29 studies (32.2%). Multi-layer control mechanisms are described in 27 studies (30.0%), which means that the necessity to achieve a set of system-level protection is taken into consideration. In 23.3% of the studies post-generation validation is used with 14.5% of the literature not mentioning any hallucination mitigation strategy. These findings indicate that despite the fact that the majority of the studies recognize the risk of hallucinations, the level of mitigation and rigor of these methods differs significantly.

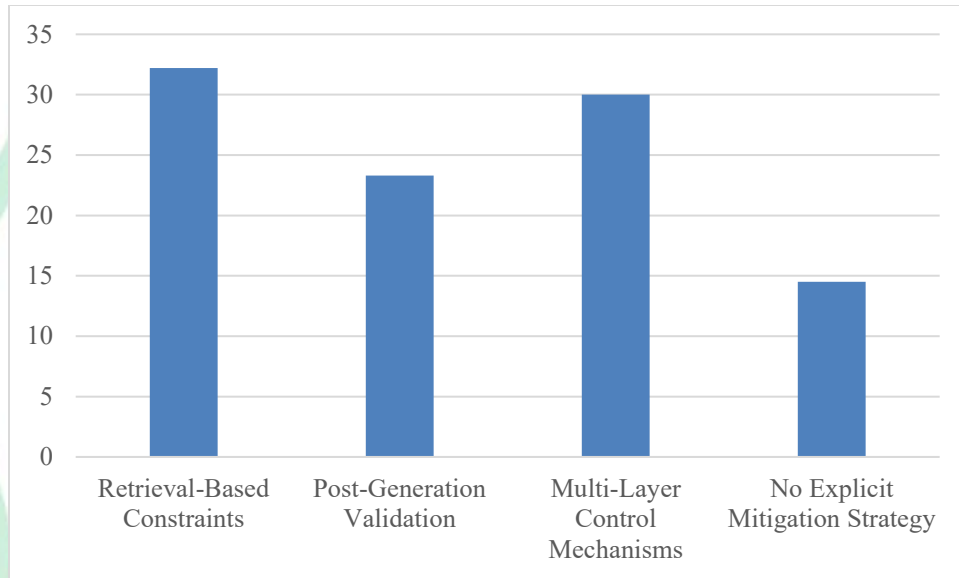


Figure 3: Graphical Representation of the Percentage of Hallucination Mitigation Strategies Reported in Analytical Systems

The distribution of the percentage of strategies of mitigating hallucinations is graphically illustrated in figure 3. The figure shows the preeminence of retrieval-based and multi-layer mechanisms and also attracts attention to the percentage of studies that do not have definite mitigation initiatives. This diagram supports the need to have integrated architectural solutions to mitigate hallucination in the analytical systems.

Table 4.4 classifies the reviewed studies on the main realm of applications of retrieval-augmented analysis systems. These domains are healthcare analytics, financial and business intelligence, scientific research and knowledge discovery, governance and policy analysis, and general-purpose big data analytics.

Table 4.4: Application Domains of Retrieval-Augmented Analytical Systems (N = 90)

Application Domain	Frequency	Percentage (%)
Healthcare Analytics	22	24.4
Financial and Business Intelligence	18	20.0
Scientific Research and Knowledge Discovery	20	22.2
Governance and Policy Analysis	15	16.7
General-Purpose Big Data Analytics	15	16.7
Total	90	100.0

The table shows that the largest percentage of applications is in healthcare analytics (24.4%), then scientific research and knowledge discovery (22.2%). Financial and business intelligence is 20.0% of the studies with governance and general-purpose analytics making 16.7 each. This spread is an indication that retrieval-enhanced analytical systems have been specifically appreciated in areas where evidence foundation, accuracy and explainability hold importance.

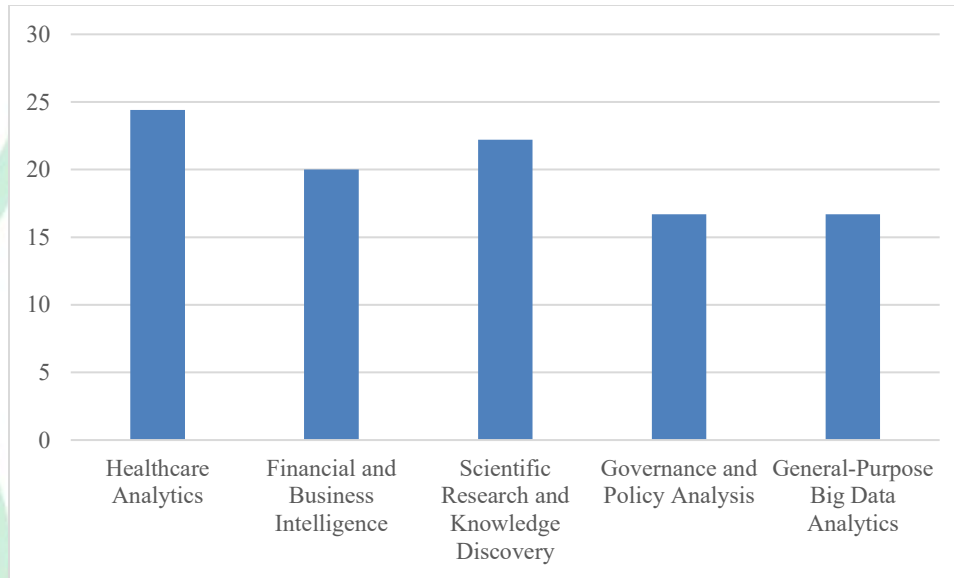


Figure 3: Graphical Representation of the Percentage of Application Domains of Retrieval-Augmented Analytical Systems

The graphical illustration of the application domains visually shows the comparatively equal distribution of retrieval-augmented systems in a variety of sectors. The figure indicates high usage in healthcare and scientific research with the validation of the widespread applicability of retrieval-augmented analytical systems in a wide range of big data settings.

5. CONCLUSION

The paper has conceptually discussed retrieval-augmented analytical systems as a productive method to explore big data having solid evidence foundation and low chances of hallucination. Such systems reduce the major shortcomings of traditional generative analytics, especially the factual accuracy and transparency problems, through the combination of external data retrieval with controlled generative reasoning.

It is concluded that retrieval-augmented architectures are an important step towards reliable AI-based analytics. They allow verifiable, interpretable and contextual insights and can be used in areas of critical relevance where the reliability of analytics is required. It is possible that future studies can be aimed at the empirical validation, performance benchmarking and development of the standardized evaluation metrics of evidence grounding and hallucination mitigation.

REFERENCES

1. Zhang, W., & Zhang, J. (2025). Hallucination mitigation for retrieval-augmented large language models: a review. *Mathematics*, 13(5), 856.
2. Abdelghafour, M. A. M., Mabrouk, M., & Taha, Z. (2024). Hallucination mitigation techniques in large language models. *International Journal of Intelligent Computing and Information Sciences*, 24(4), 73-81.
3. Bhattacharya, R. (2024). Strategies to mitigate hallucinations in large language models. *Applied Marketing Analytics*, 10(1), 62-67.
4. Kazlaris, I., Antoniou, E., Diamantaras, K., & Bratsas, C. (2025). From Illusion to Insight: A Taxonomic Survey of Hallucination Mitigation Techniques in LLMs. *AI*, 6(10), 260.
5. Wu, W., Wang, H., Li, B., Huang, P., Zhao, X., & Liang, L. (2025, May). Multirag: a knowledge-guided framework for mitigating hallucination in multi-source retrieval augmented generation. In *2025 IEEE 41st International Conference on Data Engineering (ICDE)* (pp. 3070-3083). IEEE.
6. Joshi, S. (2025). Comprehensive Review of AI Hallucinations: Impacts and Mitigation Strategies for Financial and Business Applications.

7. Sharma, C. (2025). Retrieval-Augmented Generation: A Comprehensive Survey of Architectures, Enhancements, and Robustness Frontiers. *arXiv preprint arXiv:2506.00054*.
8. Huang, L., Yu, W., Ma, W., Zhong, W., Feng, Z., Wang, H., ... & Liu, T. (2025). A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions. *ACM Transactions on Information Systems*, 43(2), 1-55.
9. Tonmoy, S. M. T. I., Zaman, S. M., Jain, V., Rani, A., Rawte, V., Chadha, A., & Das, A. (2024). A comprehensive survey of hallucination mitigation techniques in large language models. *arXiv preprint arXiv:2401.01313*, 6.
10. Anjum, S., Zhang, H., Zhou, W., Paek, E. J., Zhao, X., & Feng, Y. (2025, June). Halo: Hallucination analysis and learning optimization to empower llms with retrieval-augmented context for guided clinical decision making. In *Proceedings of the ACM/IEEE International Conference on Connected Health: Applications, Systems and Engineering Technologies* (pp. 187-198).
11. Ding, H., Pang, L., Wei, Z., Shen, H., & Cheng, X. (2024). Retrieve only when it needs: Adaptive retrieval augmentation for hallucination mitigation in large language models. *arXiv preprint arXiv:2402.10612*.
12. Amatriain, X. (2024, March). *Measuring and mitigating hallucinations in large language models: a multifaceted approach*.
13. Niu, M., Li, H., Shi, J., Haddadi, H., & Mo, F. (2024). Mitigating hallucinations in large language models via self-refinement-enhanced knowledge retrieval. *arXiv preprint arXiv:2405.06545*.
14. Joshi, S. (2025). Mitigating LLM Hallucinations: A Comprehensive Review of Techniques and Architectures. *Available at SSRN 5267540*.
15. Alansari, A., & Luqman, H. (2025). Large Language Models Hallucination: A Comprehensive Survey. *arXiv preprint arXiv:2510.06265*.