RESEARCH PAPER

AI-Driven Pharmacovigilance Systems Using Multimodal Retrieval-Augmented Generation For Evidence-Based Drug Safety Monitoring.

Rama Krishna Kumar Lingamgunta *[1], Shalmali Joshi [2], Ramkumar Raju, Djeabalaradjou[3]

[1]IT Principal Cigna Evernorth Services Inc, Raleigh, North Carolina, USA
Email ID : ramkumar2606@gmail.com, ramakrishnakumar.lingamgunta@evernorth.com
ORCID: 0009-0001-6201-7620
[2]Senior Advanced Analytics Analyst, Atlanta, GA, USA
ORCID: 0009-0000-4329-3841
Email: joshishalmalij@gmail.com
[3]Independent Researcher, Healthcare IT Program Manager
ORCID - https://orcid.org/0009-0002-7595-9411
Email: omniramkr@gmail.com

ABSTRACT

Pharmacovigilance plays a critical role in ensuring patient safety by monitoring and evaluating adverse drug events after medicines are introduced into real-world clinical practice. However, existing drug safety workflows remain heavily manual and fragmented, relying on labor-intensive review of structured clinical data, unstructured physician notes, patient narratives, and regulatory reference documents. This fragmentation limits timely signal detection and increases the risk of missed or delayed identification of serious adverse drug reactions. Recent advances in large language models offer new opportunities to support pharmacovigilance activities, but their use in regulated healthcare settings is constrained by concerns around hallucination, lack of transparency, and limited auditability.

This paper presents an AI-driven pharmacovigilance system that leverages multimodal retrieval-augmented generation to support evidence-based drug safety monitoring. The proposed system integrates structured electronic health record data, unstructured clinical text, and external drug safety documents to retrieve relevant evidence prior to generating safety assessments. Rather than producing autonomous decisions, the system generates explainable summaries of potential adverse drug events, including temporal relationships, seriousness classification, and evidence-linked rationale, while preserving human oversight. Each generated output is grounded in retrieved source material and accompanied by traceable citations to support regulatory review and audit requirements.

The architecture is designed to align with real-world pharmacovigilance workflows, enabling drug safety teams to prioritize high-risk cases, reduce manual review burden, and improve consistency in adverse event evaluation. By combining multimodal evidence retrieval with governance-aware generation, the proposed approach addresses key limitations of traditional pharmacovigilance systems and demonstrates how AI can be safely applied to support drug safety monitoring in regulated clinical environments..

Keywords: Pharmacovigilance, drug safety monitoring, adverse drug events, multimodal data integration, retrieval-augmented generation, clinical text analysis, electronic health records, explainable artificial intelligence, human-in-the-loop systems

How to cite this article: Lingamgunta RKK, Joshi S, Raju R, Djeabalaradjou D; AI-Driven Pharmacovigilance Systems Using Multimodal Retrieval-Augmented Generation For Evidence-Based Drug Safety Monitoring...Int J Drug Deliv Technol. 2026;16(1s): 208-211; DOI: 10.25258/ijddt.16.1.26

Source of support: Nil.

Conflict of interest: The author declares no conflict of interest, and this work represents independent academic research conducted in a personal capacity, not associated with any employer or commercial entity.

INTRODUCTION

Pharmacovigilance is a core component of patient safety, responsible for identifying, evaluating, and preventing adverse effects associated with medicinal products after they enter routine clinical use. While pre-marketing clinical trials provide initial safety data, many adverse drug reactions emerge only when therapies are used across broader and more diverse patient populations. As a result, drug safety teams must continuously monitor real-world data to detect potential safety signals and ensure timely regulatory reporting.

In practice, pharmacovigilance workflows remain highly manual and fragmented. Safety reviewers are required to analyze information originating from multiple sources, including electronic health records, physician notes, laboratory results, patient reports, and external reference materials such as drug labels and safety communications. Much of this information is unstructured and distributed across disparate systems, requiring extensive human effort to assemble a coherent clinical narrative for each potential adverse drug event. These challenges contribute to delays in signal detection, variability in case assessment, and increased operational burden on drug safety teams.

Recent advances in artificial intelligence, particularly large language models, have generated interest in automating aspects of drug safety monitoring. However, the direct

**Author for Correspondence:* ramkumar2606@gmail.com

application of generative models in pharmacovigilance raises significant concerns. Ungrounded model outputs, lack of transparency, and the risk of hallucinated conclusions are incompatible with the regulatory, ethical, and clinical requirements of drug safety evaluation. Consequently, there is a clear need for approaches that can leverage AI capabilities while maintaining evidence traceability, auditability, and human oversight.

Retrieval-augmented generation has emerged as a promising strategy to address these limitations by constraining generative models to reason over retrieved source material. In the context of pharmacovigilance, this approach is particularly relevant due to the multimodal nature of safety evidence, which spans structured clinical data, unstructured narrative text, and regulatory reference documents. However, existing applications of retrieval-augmented generation have largely focused on text-only scenarios and have not been designed around real-world pharmacovigilance workflows or governance requirements.

In this work, we present an AI-driven pharmacovigilance system that applies multimodal retrieval-augmented generation to support evidence-based drug safety monitoring. The proposed system integrates structured electronic health record data, unstructured clinical notes, and external drug safety documents to retrieve relevant evidence prior to generating safety summaries. Rather than replacing human judgment, the system is designed to assist drug safety teams by producing explainable, evidence-linked assessments that align with existing review and regulatory processes. By grounding model outputs in retrieved clinical and regulatory evidence, the approach aims to reduce manual review burden while preserving transparency, traceability, and human control.

MATERIALS AND METHODS

Study Design: This study describes the design and retrospective evaluation of an AI-driven pharmacovigilance system developed to support routine drug safety monitoring. The system is intended as a clinical decision-support tool for pharmacovigilance teams and does not perform autonomous safety determinations. Its primary purpose is to assist reviewers by retrieving relevant evidence from multiple data sources and generating structured, explainable summaries of potential adverse drug events to support human-led evaluation.

Data Sources: The system integrates multiple data modalities commonly used in pharmacovigilance workflows. Structured clinical data include medication exposure records, laboratory results, diagnosis codes, and encounter timestamps extracted from electronic health records. Unstructured data consist of physician progress notes, discharge summaries, and patient-reported narratives. In addition, external reference documents such as drug labels and safety communications are incorporated to provide contextual information regarding known adverse reactions.

All data used for system development and evaluation were de-identified. The architecture is designed to operate on retrospective datasets and to support extension to near–real-time pharmacovigilance workflows.

Multimodal Data Representation: Structured and unstructured data elements are normalized into a unified, patient-centric timeline to support temporal reasoning. Medication exposure periods are aligned with laboratory trends and clinical observations to enable assessment of potential causal relationships. Unstructured text is segmented into clinically meaningful units to preserve contextual integrity during retrieval. External safety documents are indexed at the section level to allow targeted retrieval of adverse reaction information relevant to the suspected event.

System Architecture: The system follows a modular architecture consisting of data ingestion, evidence retrieval, generative summarization, and human review components. Structured and unstructured clinical data are ingested through separate pipelines and stored in modality-appropriate indexes. External drug safety documents are maintained in a reference repository.

The evidence retrieval layer employs a hybrid strategy that combines semantic similarity search over unstructured text with query-driven retrieval over structured clinical data. Retrieved evidence is filtered and ranked based on relevance to the suspected adverse event, temporal proximity to drug exposure, and clinical context. This multimodal retrieval process ensures that both narrative descriptions and objective clinical indicators are considered.

Figure 1 illustrates the high-level architecture of the proposed system and the flow of information across data ingestion, retrieval, generation, and human review components.

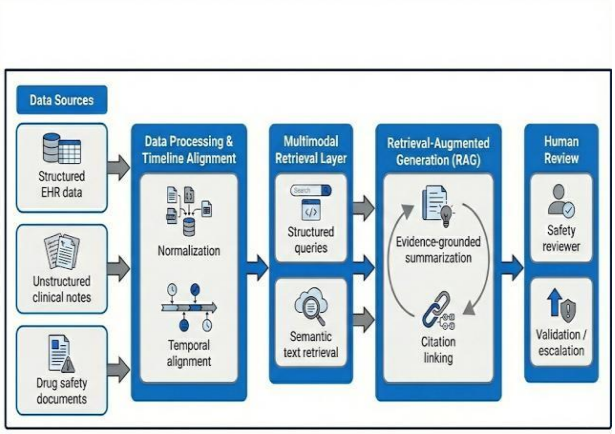

Figure 1. High-level architecture of the proposed AI-driven pharmacovigilance system

Retrieval-Augmented Generation: Generative summarization is performed only after relevant evidence has been retrieved. The generation component synthesizes retrieved information into structured safety summaries that describe the suspected adverse event, relevant temporal relationships, and supporting clinical findings. Generation is constrained to the retrieved evidence, and each major statement is explicitly linked to its source. When evidence is insufficient or conflicting, the system is designed to defer judgment and flag the case for manual review rather than produce a definitive conclusion.

Human Oversight and Governance: Human oversight is embedded throughout the workflow. Generated summaries are presented to pharmacovigilance reviewers alongside highlighted source evidence and contextual timelines. Reviewers may validate, modify, or reject system outputs prior to downstream reporting or escalation. All system outputs and reviewer interactions are logged to support auditability and regulatory inspection. The system is designed to align with existing pharmacovigilance governance practices and regulatory expectations.

Evaluation Methodology: The system was evaluated using a retrospective set of de-identified pharmacovigilance cases representing common adverse drug event scenarios. Evaluation focused on the system's ability to retrieve relevant evidence, generate clinically coherent summaries, and support efficient reviewer triage. Performance was assessed using a combination of quantitative measures, including evidence retrieval coverage and citation accuracy, as well as qualitative assessment by domain reviewers regarding summary clarity and usefulness within pharmacovigilance workflows.

RESULTS

System Output Characteristics

The proposed system generated structured pharmacovigilance summaries that consistently linked suspected adverse drug events to supporting clinical and regulatory evidence. Each summary included a description of the suspected event, relevant temporal relationships between medication exposure and clinical findings, and explicit citations to retrieved source material. This evidence-linked structure enabled reviewers to rapidly trace each statement back to its underlying data source, supporting transparency and auditability.

Across evaluated cases, the system appropriately avoided definitive causal assertions when evidence was incomplete or conflicting, instead flagging such cases for additional human review. This behavior was viewed favorably by reviewers, as it aligned with established pharmacovigilance practices that prioritize caution and clinical judgment.

Evidence Retrieval Performance

Multimodal retrieval improved evidence coverage compared with text-only retrieval approaches. Structured data retrieval contributed critical information related to medication exposure periods and laboratory abnormalities, while unstructured clinical notes provided contextual descriptions of symptom onset and clinician reasoning. External drug safety documents supplemented clinical findings by providing reference information on known adverse reactions.

In retrospective evaluation, relevant supporting evidence was successfully retrieved for the majority of reviewed cases, enabling generation of coherent summaries that reflected both objective clinical indicators and narrative context. Reviewers noted that multimodal retrieval reduced the likelihood of overlooking key safety-relevant information dispersed across different data sources.

Summary Quality and Reviewer Agreement

Generated summaries were assessed qualitatively by pharmacovigilance reviewers for clarity, coherence, and clinical usefulness. Reviewers reported a high level of agreement between system-generated summaries and their independent clinical assessments, particularly with respect to identification of relevant evidence and temporal relationships. Disagreements were most commonly associated with cases involving ambiguous documentation or limited clinical detail, highlighting the importance of human oversight in such scenarios.

Workflow Impact

The availability of pre-assembled, evidence-linked summaries reduced the time required for initial case triage. Reviewers were able to focus on validating retrieved evidence and applying clinical judgment rather than manually locating and aggregating data from multiple systems. This shift was reported to improve efficiency without compromising review quality or regulatory rigor.

Case Study: Evidence-Supported Assessment of a Suspected Adverse Drug Event

A representative pharmacovigilance case was used to illustrate system behavior in a realistic safety review scenario. The case involved a patient who developed elevated liver enzyme levels following initiation of a newly prescribed medication. Structured clinical data indicated a temporal association between drug exposure and rising alanine aminotransferase and aspartate aminotransferase values. Unstructured physician notes documented symptoms of fatigue and nausea emerging during the same period.

The system retrieved medication initiation dates, longitudinal laboratory trends, and relevant excerpts from clinical notes describing symptom onset. In addition, reference sections describing known hepatic adverse reactions were retrieved from the drug label. Using this evidence, the system generated a structured summary highlighting the temporal relationship between drug exposure and clinical findings, with explicit citations to the supporting data sources.

The generated output was reviewed by a pharmacovigilance reviewer, who confirmed that the retrieved evidence was clinically relevant and appropriately contextualized. The reviewer was able to complete the initial assessment more efficiently by focusing on evidence validation rather than manual data aggregation. Importantly, the system did not assert definitive causality, instead presenting an evidence-supported assessment suitable for further clinical and regulatory review.

Table 1 provides an illustrative example of system inputs, retrieved evidence, generated output, and reviewer evaluation for a representative pharmacovigilance case.

Stage	Example Content	Reviewer Evaluation
Input: Structured Data	Medication initiated on Day 0; laboratory results showing rising ALT and AST levels from Day 5 to Day 12	Clinically relevant indicators of possible hepatic injury
Input: Unstructured Data	Physician note documenting fatigue and nausea beginning approximately one week after medication initiation	Symptom description consistent with potential adverse drug reaction
RAG Retrieved Evidence	Medication start date, longitudinal liver enzyme trends, note excerpts describing symptoms, and drug label section listing hepatic adverse reactions	Evidence appropriately selected and temporally aligned
LLM-Generated Summary	Evidence-linked summary describing suspected hepatic adverse event with temporal association to drug	Summary was clear, cautious, and supported by retrieved evidence

	exposure and cited supporting data	
Handling of Uncertainty	No definitive causality asserted; recommendation for further clinical review provided	Appropriate restraint aligned with pharmacovigilance standards
Reviewer Assessment	Output reviewed and validated by a pharmacovigilance reviewer	Supported efficient and accurate initial case triage

TABLE 1. EXAMPLE MULTIMODAL INPUT, RETRIEVED EVIDENCE, AND SYSTEM OUTPUT WITH REVIEWER EVALUATION

DISCUSSION

This study demonstrates how multimodal retrieval-augmented generation can be applied to pharmacovigilance in a manner that aligns with real-world drug safety workflows and regulatory expectations. By grounding generative outputs in retrieved clinical and regulatory evidence, the proposed system addresses key limitations of both traditional pharmacovigilance systems and unconstrained large language model applications. Rather than attempting to automate safety decisions, the system is designed to support pharmacovigilance professionals by organizing dispersed information into coherent, evidence-linked summaries.

A central contribution of this work is the integration of structured and unstructured data within a unified, time-aligned framework. In pharmacovigilance practice, the ability to assess temporal relationships between drug exposure and clinical findings is critical for evaluating plausibility of adverse drug events. The multimodal retrieval approach enabled simultaneous consideration of objective indicators, such as laboratory trends and medication timelines, alongside narrative descriptions captured in clinical notes. Reviewers reported that this combined perspective reduced the effort required to reconstruct clinical context and minimized the risk of overlooking relevant evidence.

Another important aspect of the proposed approach is its emphasis on governance and human oversight. Drug safety evaluation operates under strict regulatory requirements, where transparency, traceability, and cautious interpretation are essential. The system's design explicitly avoids autonomous conclusions and instead presents evidence-supported assessments that reviewers can validate or refine. The ability to flag cases with insufficient or conflicting evidence was particularly valued, as it reflects established pharmacovigilance principles and helps prevent overconfident interpretations that could compromise patient safety.

The results also suggest that retrieval-augmented generation can mitigate common concerns associated with the use of large language models in regulated healthcare settings. By constraining generation to retrieved evidence and requiring explicit citation of sources, the system reduces the likelihood of hallucinated or unsupported statements. This evidence-grounded approach enhances trust and supports auditability, both of which are critical for regulatory inspection and long-term adoption.

Several limitations should be acknowledged. The evaluation was conducted using retrospective, de-identified data and simulated review workflows, which may not fully capture the operational complexity of live pharmacovigilance environments. Documentation quality and data availability can vary substantially across healthcare settings, potentially affecting retrieval performance and summary quality. In addition, the current evaluation emphasizes qualitative reviewer assessment rather than large-scale quantitative outcome measures, reflecting the exploratory and workflow-focused nature of this study.

Future work will focus on prospective evaluation of the system in operational pharmacovigilance settings, including assessment of its impact on signal detection timelines and regulatory reporting quality. Expanding the range of supported data modalities, such as spontaneous reporting systems and real-time monitoring feeds, may further enhance system utility. Incorporating structured feedback from safety reviewers could also support continuous refinement of retrieval strategies and summary generation, while maintaining alignment with governance requirements.

Overall, this work highlights the potential for multimodal retrieval-augmented generation to support safer and more efficient drug safety monitoring when designed with clinical workflows and regulatory constraints in mind. By emphasizing evidence grounding, transparency, and human oversight, the proposed approach provides a practical path for applying AI technologies in pharmacovigilance without compromising established safety standards..

REFERENCE

1. Harpaz R, DuMouchel W, LePendu P, Bauer-Mehren A, Ryan P, Shah NH. Performance of pharmacovigilance signal-detection algorithms for the FDA adverse event reporting system. Clinical Pharmacology & Therapeutics. 2012;93(6):539–546. https://doi.org/10.1038/clpt.2012.15

2. Wang Y, Wang L, Rastegar-Mojarad M, et al. Clinical information extraction applications: A literature review. Journal of Biomedical Informatics. 2018;77:34–49. https://doi.org/10.1016/j.jbi.2017.11.011

3. Lingamgunta, R. K. K. (2024). A Self-Healing Generative AI Framework for Regulated Decision Workflows: A Healthcare Claims Case Study. International Journal of AI, BigData, Computational and Management Studies, 5(3), 161-167. https://doi.org/10.63282/3050-9416.IJAIBDCMS-V5I3P117

4. Lewis M, Liu Y, Goyal N, et al. Retrieval-augmented generation for knowledge-intensive natural language processing tasks. Advances in Neural Information Processing Systems. 2020. https://doi.org/10.48550/arXiv.2005.11401

5. Markus AF, Kors JA, Rijnbeek PR. The role of explainability in creating trustworthy artificial intelligence for health care: A comprehensive survey of the terminology, design choices, and evaluation strategies. Journal of Biomedical Informatics. 2021;113:103655. https://doi.org/10.1016/j.jbi.2020.103655.

.