INTERNATIONAL JOURNAL OF PHARMACEUTICAL SCIENCES
[ISSN: 0975-4725; CODEN(USA): IJPS00]
Journal Homepage: https://www.ijpsjournal.com

OPEN ACCESS

**Review Article**

# Generative Artificial Intelligence and Large Language Models in Pharmacovigilance: Current Applications, Challenges and Future Perspectives.

**Samiksha Jadhav*, Appasaheb Kuhile, Yogita Lanke, Sujata Nagare**

*Department of pharmacy, JES's SND College Of Pharmacy Babhulgaon (Yeola), India.*

ARTICLE INFO

Published: 12 Nov 2026
Keywords:
Pharmacovigilance, Generative Artificial Intelligence, Large Language Models
DOI:
10.5281/zenodo.24155757

## ABSTRACT

Pharmacovigilance continuously monitors medicine safety by identifying adverse drug reactions and medication-related risks. The rapid growth of data from case reports, scientific literature, and electronic health records has significantly increased manual workloads. Consequently, Generative Artificial Intelligence (GenAI) and Large Language Models (LLMs) are emerging as pivotal tools to optimize information processing and interpretation. Current applications span data extraction, individual case safety report processing, case summarization, literature screening, signal and causality assessment, drug–drug interaction evaluation, and automated safety documentation. Evidence suggests these technologies enhance efficiency when handling vast datasets and repetitive tasks. However, critical challenges persist regarding hallucinated or inaccurate outputs, data biases, privacy concerns, lack of explainability, terminology errors, and insufficient real-world validation. Because pharmacovigilance drives safety-critical decisions, AI-generated outputs must never replace professional judgment. Human oversight, rigorous quality control, data protection, and regulatory framework integration remain vital. Future advancements in domain-specific language models, standardized evaluations, and human-supervised workflows will be essential to ensure the safe, prospective integration of GenAI into routine pharmacovigilance practice.

## INTRODUCTION

Patient safety is a major focus in healthcare, which is why pharmacovigilance (PV) is so critical. Its main job isn't just tracking bad side effects of drugs, but also understanding and preventing future risks. We need to monitor medicines closely even after they hit the market. Clinical trials are small, so many rare side effects only show up when thousands of different people start using the medicine out in the real world. This ongoing

***Corresponding Author:** Samiksha Jadhav
**Address:** Department of pharmacy, JES's SND College Of Pharmacy Babhulgaon (Yeola), India
**Email** ✉: jadhavsamiksha114@gmail.com
**Relevant conflicts of interest/financial disclosures**: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

watch helps doctors spot early warning signs and evaluate Keeping track of drug safety has become a massive task simply because we have too much healthcare data nowadays. Information about bad side effects comes from everywhere-like individual case safety reports (ICSRs), hospital records, research papers, and even what patients post on social media. The real headache is that this data is a total mess. It is mixed up in structured and unstructured files, uses different medical terms, and often misses key details. Checking all this by hand or using old computer programs is just too slow. Because of this lag, safety teams often fail to catch and understand important drug warnings before it is too late [1,2].

Recent reviews have highlighted the increasing use of artificial intelligence (AI), machine learning (ML), and natural language processing (NLP) to support pharmacovigilance activities such as adverse event detection, signal management, data classification, and analysis of real-world information. Although conventional AI and ML approaches have demonstrated potential in pharmacovigilance, many models are developed for specific analytical or predictive tasks. The emergence of Generative Artificial Intelligence (GenAI) has expanded the potential role of AI by enabling systems to generate, transform, summarize, and interpret information in natural language. Among these technologies, Large Language Models (LLMs) have attracted increasing attention because of their ability to process and generate human-like text and their potential application to complex biomedical and clinical information [3,4].

Recent studies and reviews have explored the use of GenAI and LLMs in several pharmacovigilance-related activities. These include extraction of drug and adverse-event information from unstructured text, summarization of safety reports, analysis of medication related harm, literature processing, signal management, and assistance with causality assessment. For example, recent investigations have evaluated LLM-based approaches for adverse-event extraction and ICSR causality assessment, demonstrating useful capabilities while also showing that performance can vary according to the model, task, prompting strategy, and evaluation criteria. Evidence from broader reviews also suggests that GenAI and LLMs may assist in reducing the burden associated with processing large volumes of medication-safety information [5,6].

Despite these opportunities, the application of GenAI and LLMs in pharmacovigilance remains an evolving field. Important concerns include hallucination, omission or generation of incorrect information, data quality, bias, limited explainability, terminology and coding errors, privacy and confidentiality, model validation, and integration into existing pharmacovigilance workflows. Because pharmacovigilance is a safety-critical field, generated information requires appropriate verification and expert oversight before it can be used for regulatory or clinical decision-making [7,8].

This is exactly why current research strongly points toward keeping a human-in-the-loop, setting up tight guardrails, and running strict, task-specific validation tests. Given how fast GenAI and LLM tools are growing-and how eager everyone is to use them for drug safety-we urgently need a clear-headed look at what they can actually do right now and where they fail. This review directly addresses that need. We take a close look at Generative AI and Large Language Models in pharmacovigilance, breaking down their current uses, major bottlenecks, and where the field is heading next. Throughout this analysis, our focus remains on why human expertise, thorough testing, and responsible deployment are non-negotiable in a high-stakes field like drug safety [5,8].

## 2. Generative Artificial Intelligence and Large Language Models

### 2.1 Artificial Intelligence and Machine Learning

Computers are now doing heavy jobs that once required real human thinking-like recognizing clinical patterns, learning from data, and helping doctors make decisions. This is what we call artificial intelligence. A specific tool under this is machine learning, where computer systems teach themselves from old data to predict future outcomes. In drug safety work, these machine learning tools have become very useful. Researchers use them to automatically categorize bad patient experiences, read safety reports, and catch risky drug signals much faster than before [3,5].

### 2.2 Deep Learning and Natural Language Processing

To spot highly complex clinical details hidden in massive patient datasets, researchers often turn to deep learning, which uses advanced multi-layer neural networks. Alongside this, natural language processing-or NLP-has become a game-changer for reading human language. It is incredibly useful for scanning messy, unstructured text like doctor's case narratives or published research papers, allowing computers to instantly pull out specific drug names and side effects without manual effort[5,14].

### 2.3 Generative Artificial Intelligence

Instead of just sorting data or guessing outcomes like traditional computer models, generative artificial intelligence actually builds entirely new content by learning from existing patterns. This means GenAI can rewrite, summarize, translate, and blend complex information together. In the world of drug safety, this flexibility opens up huge possibilities. It gives safety teams a faster way to handle massive workloads, especially when they need to summarize long clinical reports, pull out hidden details from patient files, or scan through piles of medical research papers [2,4].

### 2.4 Large Language Models

Think of large language models, or LLMs, as computer programs that read massive amounts of text to learn how humans communicate. Systems like GPT and other transformer models are great at multiple jobs, from summarizing reports and pulling out data to answering tricky questions. Right now, drug safety researchers are exploring what these tools can do. They are using them to spot adverse drug reactions, handle safety documents faster, analyse medical papers, and even assist in checking if a medicine is truly responsible for a patient's side effect [10,11].

### 2.5 Transformer Architecture and Prompt Engineering

Most new LLMs are built on what we call transformer systems, using attention mechanisms to figure out how words relate to one another in a sentence. If you want a good response from these models, you have to practice prompt engineering, which simply means writing clear instructions for the AI. What is interesting is that new research shows how these prompts are written heavily affects the results in drug safety jobs, like tracing if a pill actually caused a side effect. This proves that GenAI is moving past old systems that just group data, and is now actually understanding human speech. The future for this in drug safety looks great, but we still absolutely need human eyes to double-check and validate all the data it spits out [8,16].

## 3. PHARMACOVIGILANCE

### 3.1 Definition and Objectives of Pharmacovigilance

Let's talk about medicine safety. That's exactly where pharmacovigilance-or just PV-comes into the picture. Basically, it is the science behind spotting, tracking, and stopping bad side effects before they hurt anyone. This safety watch never actually stops; it begins way back during early clinical trials and stays active long after the drug hits pharmacy shelves.

The big goal? Keep patients safe. By catching nasty drug reactions early, experts can figure out if a medicine does more good than harm. To get this done, teams look at a massive mix of clues. We are talking about individual patient files, hospital papers, digital health records, and general real-world data [13,19].

### 3.2 Need for AI and Technological Advancement in Pharmacovigilance

Data is flooding into the drug safety sector at an unbelievable speed right now. Because of this massive wave, pharmacovigilance badly needs a tech upgrade. Honestly, looking at traditional PV, it is just way too exhausting for human teams. Doing manual stuff like sorting cases, finding adverse reactions, skimming journals, and spotting safety signals takes up hours of tedious work [3,13].

Enter modern computer solutions. Today, experts actively test artificial intelligence (AI), machine learning (ML), and natural language processing (NLP) to break down these huge datasets. They act like smart assistants that read unstructured medical text, map patterns, and speed up signal detection. On top of that, Generative AI and Large Language Models (LLMs) are opening completely fresh doors. They summarize and extract tough info using everyday conversational language.

But wait, we cannot blindly rush into this. In a field where patient lives are on the line, these models need strict checking. A tiny AI error-like a random hallucination, missed data, or bad medical mapping-can completely ruin critical safety decisions.

Thus, technological advancement in pharmacovigilance is aimed not at replacing safety professionals, but at supporting efficient data processing and decision-making while maintaining appropriate human oversight and validation [2,8].

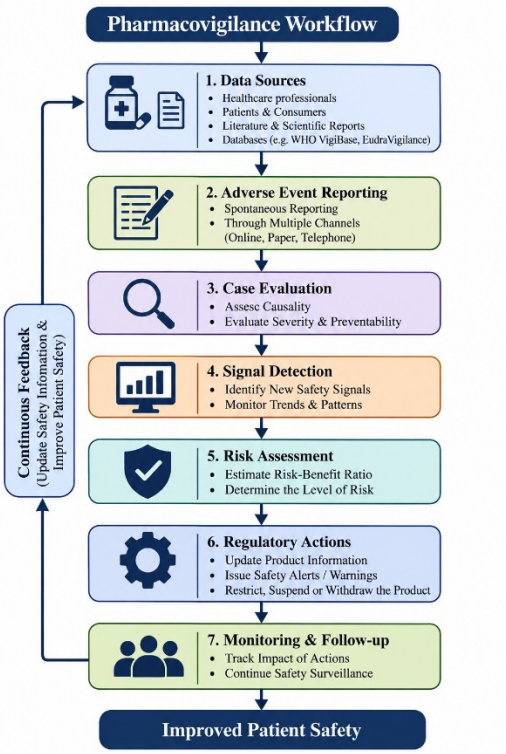



**Fig. General Workflow of Pharmacovigilance.**

## 4. APPLICATIONS OF GENERATIVE AI AND LARGE LANGUAGE MODELS IN PHARMACOVIGILANCE

Generative AI and Large Language Models (LLMs) are increasingly being explored for pharmacovigilance activities involving large volumes of structured and unstructured safety information. Their ability to process natural language, extract information, summarize text and generate structured outputs creates opportunities to support several stages of the pharmacovigilance workflow.

### 4.1 Individual Case Safety Report (ICSR) Processing

ICSR processing involves the identification, extraction, assessment and documentation of information related to suspected adverse events. GenAI and LLMs can assist in extracting relevant information such as patient characteristics, suspected drugs, adverse events, concomitant medications and clinical outcomes from unstructured narratives. They can also support case summarization and preparation of structured information for further pharmacovigilance processing [1,2].

### 4.2 Adverse Drug Reaction and Adverse Drug Event Extraction

LLMs can analyse free-text sources such as clinical narratives, medical literature and patient reported information to identify drug names, adverse events and associated clinical information. Recent studies have demonstrated the potential of LLMs for extracting drug and ADR information from text, although performance may vary according to the model, task and evaluation dataset [5,9].

### 4.3 Case Summarization and Medical Information Synthesis

Pharmacovigilance professionals frequently need to review lengthy case narratives and safety documents. GenAI can summarize relevant clinical and safety information into a more structured form, potentially reducing the time required for manual review. It may also assist in synthesizing information from multiple safety-data sources while preserving important clinical details [1,2].

### 4.4 Literature Screening and Safety Information Retrieval

Pharmacovigilance requires continuous monitoring of scientific literature for newly reported adverse events and drug-safety information. LLMs can support literature screening by classifying documents, identifying safety-relevant information, summarizing articles and prioritizing potentially relevant publications for expert review. These capabilities may help manage the increasing volume of biomedical literature [2,13].

### 4.5 Signal Detection and Signal Management

AI-based methods have been investigated for identifying and prioritizing potential safety signals from pharmacovigilance databases. Conventional machine-learning and NLP approaches have shown applications in signal detection and validation, while GenAI/LLMs may provide additional support through natural-language analysis and synthesis of evidence surrounding a potential signal. However, current evidence for LLMs in routine signal management remains limited. Therefore, LLM-generated outputs should be considered supportive information rather than independent evidence for regulatory decision-making [8,11].

### 4.6 Causality Assessment

Causality assessment evaluates the possible relationship between a suspected medicine and an adverse event. Recent research has investigated biomedical LLMs and prompting strategies for supporting ICSR causality assessment. Although some combinations of biomedical LLMs, causality algorithms and prompting approaches have shown agreement with human assessments, important disagreements remain, particularly for critical causality elements. Therefore, LLMs may assist experts by organizing relevant information and providing preliminary assessments, but final causality decisions require appropriate expert evaluation [8,10].

### 4.7 Drug–Drug Interaction (DDI) and Medication-Related Harm Assessment

AI-based approaches can support the identification and assessment of potential drug–drug interactions and medication-related harm. GenAI and LLMs have been investigated in areas related to medication safety, while recent work has proposed human-augmented AI systems that combine pharmacological knowledge, real-world information and clinician feedback for more proactive DDI risk assessment [1,15].

### 4.8 Generation of Safety and Regulatory Documentation

Generative AI can assist in drafting structured case narratives, summaries and other safety related documents from extracted information. This may improve consistency and reduce repetitive documentation work. However, generated content requires verification to prevent inaccurate, incomplete or unsupported information from entering safety documentation [2,8].

### 4.9 Multilingual and Unstructured Data Processing

Pharmacovigilance data may originate from different languages and diverse unstructured sources. GenAI-based language processing may assist with translation, information extraction and summarization of multilingual safety information. This could support broader use of global safety data, although terminology accuracy and validation remain important considerations [6,13].

Overall, GenAI and LLMs have potential applications across multiple stages of the pharmacovigilance workflow. However, most current evidence supports their use as decision support and productivity-enhancing tools rather than autonomous systems. Human review, validation and appropriate safeguards remain essential, particularly for safety-critical activities.

**Pharmacovigilance.**

**Table 1: Application of Generative AI and Large Language Models In**

| SR.NO. | PV APPLICATIONS | ROLE OF GENERATIVE AI / LLMs. | POTENTIAL BENEFITS. |
|---|---|---|---|
| 1. | ICSR Processing. | Data extraction and Summarization. | Faster case processing. |
| 2. | ADR/ADE Extraction. | Identify adverse events from text. | Improved information Extraction. |
| 3. | Case Summarization. | Generate concise case summaries. | Reduce documentation Burden. |
| 4. | Literature Screening. | Screen and summarize safety literature. | Saves time. |
| 5. | Signal Detection. | Support signal identification And prioritization. | Efficient safety monitoring. |
| 6. | Causality Assessment. | Organize evidence and Support assessment. | Decision support. |
| 7. | DDI Assessment. | Identify and summarize Potential DDIs. | Improved medication safety. |
| 8. | Safety Documentation. | Draft safety narratives / reports. | Improved efficiency. |

## 5. EVIDENCE FROM RECENT STUDIES

Recent studies have explored the application of Artificial Intelligence (AI), Generative AI (GenAI), and Large Language Models (LLMs) in different pharmacovigilance activities. The available evidence includes studies on medication-related harm, adverse drug reaction extraction, ICSR processing, signal management, causality assessment, and safety information analysis. However, the evidence differs in terms of study design, datasets, AI models, and evaluation methods [1,13].

### 5.1 Generative AI and Medication-Related Harm

Ong et al. conducted a scoping review of Generative AI and LLM applications related to medication-related harm and adverse drug events. The identified applications included adverse event identification and classification, drug–drug interaction assessment, and clinical decision support. The review indicated that GenAI and LLMs have potential to support medication safety activities; however, prospective real-world implementation evidence remains limited.[1]

### 5.2 ICSR Processing and ADR Extraction

LLMs have been investigated for extracting drug and adverse-event information from unstructured text. Hirschman discussed the use of LLMs for extracting drug names and ADR information from free-text sources. GPT-4 demonstrated promising performance for drug mention extraction in an expert-annotated dataset. However, expert evaluation and validation remain necessary before LLM-generated information can be incorporated into safety-critical databases [11].

Generative AI has also been proposed for processing ICSR narratives. It can assist in identifying patient characteristics, drug information, adverse events, and clinical outcomes from unstructured narratives and converting them into structured information. Such approaches may support faster case processing, but the generated information requires verification.

### 5.3 AI in Signal Detection and Signal Management

AI and machine-learning techniques have been studied for pharmacovigilance signal detection and management. Warner et al. reviewed studies involving machine learning, natural language processing, and other AI approaches for signal detection, validation, and evaluation. Several studies reported promising performance of AI-based approaches compared with traditional methods.

However, most of the available evidence in this area concerns conventional AI/ML rather than Generative AI or LLMs. GenAI and LLMs may provide additional support through natural language analysis, evidence synthesis, and summarization of information associated with potential safety signals [14].

### 5.4 LLMs for Causality Assessment

Recent research has evaluated biomedical LLMs for causality assessment of Individual Case Safety Reports. Heckmann et al. investigated different biomedical LLMs, prompting strategies, and causality algorithms and compared their results with human expert assessments. The study reported that the combination of Medicine LLaMA-3 8B, the Naranjo algorithm, and chain-of-thought prompting achieved the highest agreement among the evaluated approaches. However, the overall agreement remained limited, particularly for important causality-related elements. This indicates that LLMs may assist in organizing and analysing case information, but their performance is not yet sufficient for independent causality assessment [10].

### 5.5 Guardrails for LLM-Based Pharmacovigilance

The use of LLMs in safety-critical areas introduces the possibility of hallucinated, incomplete, or incorrect information. Hakim et al. highlighted the need for guardrails when applying LLMs to pharmacovigilance.

Fixing these software errors requires basic, practical action. Right now, the best move is double-checking medical terms and drug names to stop bad mix-ups. Teams should also look for broken or weird documents while forcing the AI to show whenever it feels uncertain about a finding. Testing the actual meaning of the text adds another layer of safety. Most of all, real human experts must stay in charge to review the final data. Doing these things is exactly how we make language models safe to use in drug tracking[8].

### 5.6 Evidence from AI-Based DDI and Real-World Data Applications.

Right now, people are testing how smart software tracks dangerous drug mixes using everyday patient data. The process is pretty straightforward. These systems combine basic pharmacy facts with actual hospital checks to find medicine risks early on. So far, the latest studies prove that letting computers scan this complicated data works wonders. Just keep one important detail in mind. These tools are completely different from modern chatbots like ChatGPT. The truth is, most of these older frameworks run on standard, basic machine learning tricks rather than creative generative tech.

Once the software spits out this information, a trained expert must look it over carefully. A human eyes check has to happen first. Only after this professional review can teams safely use the data to process medical cases or judge drug safety[15].

### 5.7 Overall Evidence

The reviewed studies indicate that AI, GenAI, and LLMs are being investigated across several pharmacovigilance activities, particularly information extraction, ICSR processing, literature and safety-information analysis, signal management, causality assessment, and medication safety applications.

The evidence is promising but heterogeneous. Differences in datasets, models, prompting methods, evaluation criteria, and study designs make direct comparison difficult. In addition, several studies are reviews, perspectives, or early-stage evaluations rather than prospective real-world implementation studies.

Overall, current evidence suggests that GenAI and LLMs have considerable potential to support pharmacovigilance workflows, but further standardized benchmarking, prospective validation, and real-world evaluation are required before their wider implementation in safety critical processes [8,10].

## 6. HUMAN–AI COLLABORATION IN PHARMACOVIGILANCE

Generative AI and Large Language Models have the potential to support pharmacovigilance professionals in processing and analysing large volumes of safety information. However, pharmacovigilance is a safety-critical field in which AI-generated outputs may contain errors, omissions, hallucinations, or incorrect terminology. Therefore, the effective use of GenAI and LLMs requires collaboration between AI systems and human experts rather than complete replacement of human decision-making [8,15].

### 6.1 Human-in-the-Loop Approach

The human-in-the-loop approach involves integrating human experts at important stages of an AI-assisted pharmacovigilance workflow. In this approach, the AI

system performs tasks such as information extraction, classification, summarization, or prioritization, while a pharmacovigilance professional reviews and validates the output.

For example, an LLM may extract suspected drugs and adverse events from an ICSR narrative. Once the software spits out this information, a trained expert must look it over carefully. A human eyes check has to happen first. Only after this professional review can teams safely use the data to process medical cases or judge drug safety.

**6.2 Role of Human Experts.**

At the end of the day, computers just cannot match actual human brains for reading deep medical context. We completely depend on real experts to tell if software data is truly reliable, sort out confusing cases, and make the big safety calls. Having a specialist look things over is vital for proving what caused a bad reaction, judging safety signals, signing official legal documents, and decoding highly complex patient stories. On top of that, sharp human eyes easily spot wrong medical words, blank spaces in reports, wild claims, and fake hallucinations that chatbots randomly make up [10].

**6.3 Human-Augmented AI**

Basically, human-augmented AI is just a team setup where smart software and actual human brains back each other up. Computers tackle the exhausting work, digging through massive, messy data piles in a flash. On the flip side, real doctors bring the essential clinical know-how, drug logic, background context, and final expert choices to the table. Look at dangerous drug-mix checks for instance. New setups combine pharmacy facts and lab metrics with real hospital tracking and direct doctor feedback. Mixing these forces makes it way easier to figure out highly confusing medicine safety details[15].

**6.4 Verification and Quality Control.**

You just cannot skip double-checking facts when using software to track drug safety. It is that simple. Before anyone relies on computer-generated notes for actual safety choices, that info must be matched up with trusted, raw source data. When setting up these audits, here are the main spots you need to look at:

- Basic details about the patients and the specific medicines given
- The exact words used to describe bad side effects
- Standard MedDRA terms and official medical codes
- Actual clinical results and how the patients recovered
- Clues showing if the drug actually caused the reaction
- Clean references and real supporting proof
- Checking if the final software summaries missed anything important

Having strict quality control steps in place is the only way to slash the risk of trusting broken or completely baseless AI data [8,10].

**6.5 Benefits of Human–AI Collaboration.**

Teamwork between humans and AI brings some massive upgrades to the whole pharmacovigilance scene. Instead of doing everything the hard way, this combination offers several clear perks:

- It completely wipes out the exhausting, repetitive manual chore of typing in data.
- You get to tear through mountain-sized piles of drug safety files in a fraction of the time.
- It works like a charm for skimming medical journals and extracting key facts instantly.
- The software acts as a quick assistant to help pull together clean case summaries.
- It does an amazing job organizing messy, scattered safety metrics into actual order.
- It helps human specialists rank the most dangerous signals so they know what to fix first.
- It perfectly blends super-fast computer numbers with vital real-world doctor logic.

These practical perks are a lifesaver when you are completely buried under massive heaps of messy, unstructured safety paperwork that needs a thorough review [1,2].

**6.6 Limitations of Human–AI Collaboration**

Human–AI collaboration also presents challenges. Excessive reliance on AI-generated outputs may lead to automation bias, in which users accept incorrect

outputs without adequate verification. In addition, differences in user expertise, inconsistent review practices, and lack of standardized validation methods may affect the reliability of AI-assisted workflows. Therefore, human involvement should be meaningful and appropriately positioned within the workflow rather than being limited to a final superficial check [13].

### 6.7 Future Direction of Human–AI Collaboration

Future pharmacovigilance systems may increasingly use collaborative workflows in which AI performs high-volume information-processing tasks while human experts provide interpretation, validation, and final decision-making. Such systems may incorporate continuous monitoring, feedback mechanisms, domain-specific knowledge, and standardized quality control procedures. The objective should be to develop AI-assisted rather than AI-independent pharmacovigilance, where technological efficiency is combined with human clinical and regulatory expertise. Overall, human–AI collaboration provides a framework for using the processing capabilities of GenAI and LLMs while retaining human expertise and accountability in pharmacovigilance [15].

## 7. CHALLENGES OF GENERATIVE AI AND LLMs IN PHARMACOVIGILANCE.

Even with all the hype, smart software and LLMs run into heavy trouble when tracking real-world medicine safety. Here are the biggest headaches:

### 7.1 Generating Fake Facts

Language tools often just make up wrong data out of nowhere. Having these random glitches in your reports can easily ruin critical safety choices [8].

### 7.2 Messy Data Feeds

If the raw health files you feed the computer are broken or full of bias, the machine will spit out shaky and useless results [20].

### 7.3 Leaking Patient Secrets

Safety databases are packed with private medical details. Because of this, teams must use airtight security setups to stop data leaks [13].

### 7.4 No Clear Explanations

Complex networks rarely show their homework. Since they do not explain how they found an answer, human experts face a tough time checking the logic [8].

### 7.5 Erratic Test Results

An AI changes its behavior based on the specific prompts or datasets you use. To stop these wild shifts, we badly need basic benchmarks and stable testing rule [10].

### 7.6 Ruining Medical Codes

Computers easily mix up standard drug names and side-effect terms. They even make silly mistakes with official MedDRA mapping tags.

### 7.7 System Integration Hassles

Hooking new tech into old hospital databases is tough. Plus, keeping regulators happy requires massive amounts of paperwork and strict audit logs.

### 7.8 Missing Live Evidence

Most studies right now are just early tests looking back at old files. We still lack actual proof showing how these tools survive daily clinic tasks.

In short, strict checking, tight security, and constant human eyes are the only ways to safely use this software for tracking drugs.

## 8. REGULATORY AND ETHICAL CONSIDERATIONS.

Dropping GenAI and LLMs into the drug safety field brings up some major ethical and legal questions that we simply cannot ignore. The reason is pretty straightforward. These automated systems end up handling deeply personal patient details and high-stakes safety profiles every single day. Their use should therefore focus on improving efficiency while maintaining patient safety and professional oversight [8,13].

### 8.1 Patient Privacy and Data Security

Pharmacovigilance data may contain confidential patient information. Appropriate data protection, de-identification, access control and secure data handling are important when using AI systems.

### 8.2 Bias and Fairness

AI models may reflect biases present in their training data. Biased or incomplete data can affect adverse event identification and safety assessment. Regular monitoring and evaluation are therefore necessary to reduce potential bias [13,20].

### 8.3 Accuracy and Hallucination

LLMs may generate incorrect or unsupported information. In pharmacovigilance, such errors could affect safety assessment or documentation. AI-generated outputs should therefore be verified against original data by qualified experts [11].

### 8.4 Human Oversight and Accountability

AI should act as a decision-support tool rather than an autonomous decision-maker in safety critical pharmacovigilance activities. Human experts should review, validate and approve AI generated information before final use[15].

### 8.5 Validation and Regulatory Governance

AI systems require task-specific validation, performance monitoring and proper documentation before their implementation in pharmacovigilance workflows. Clear governance is required for model changes, data handling and accountability.

Overall Consideration Responsible use of GenAI in pharmacovigilance requires privacy protection, bias control, output verification, human oversight, validation and regulatory governance. Current evidence supports a controlled human–AI collaboration model rather than completely autonomous pharmacovigilance.

## 9. FUTURE PERSPECTIVES

Generative AI and Large Language Models are expected to play an increasing role in pharmacovigilance as AI technologies, data resources, and validation methods continue to develop. Future applications may focus on improving the efficiency, accuracy, and integration of pharmacovigilance activities [2,13].

### 9.1 Advanced LLM-Based PV Systems

Future LLMs may be developed specifically for pharmacovigilance and biomedical applications. Domain-specific models may improve the understanding of medical terminology, adverse events, drug information, and safety narratives [16].

### 9.2 Integration with Real-World Data

Integration of LLMs with electronic health records, clinical data, literature, and other real world data sources may support more comprehensive safety monitoring and identification of medication-related risks [19].

### 9.3 Improved Signal Detection and Safety Monitoring

Future AI systems may combine conventional statistical methods, machine learning, and GenAI to support continuous analysis of safety information and prioritization of potential safety signals [18].

### 9.4 Improved Causality Assessment

Further development of biomedical LLMs and standardized prompting strategies may improve their ability to assist in ICSR causality assessment. However, reliable validation against expert assessments will remain important [12].

### 9.5 Automated but Controlled PV Workflows

Future systems may automate repetitive activities such as literature screening, information extraction, case summarization, and safety-document drafting while maintaining appropriate verification and quality-control mechanisms.

### 9.6 Standardized Validation and Regulatory Frameworks

Development of standardized benchmarks, validation procedures, documentation requirements, and regulatory guidance will be important for the responsible implementation of GenAI and LLMs in pharmacovigilance.

### 9.7 Personalized and Multilingual Safety Monitoring

Advances in language models may improve the processing of multilingual and diverse safety information, potentially supporting broader and more inclusive pharmacovigilance.

### 9.8 Overall Perspective.

Looking ahead, the integration of Generative AI into routine pharmacovigilance mandates the development of specialized, domain-specific architectures that are fully validated and operate under strict human oversight. Realizing the practical utility of these systems requires ongoing empirical research and longitudinal, real-world assessments to firmly establish their operational reliability and define their precise boundaries within drug safety pipelines [8,13].

## CONCLUSION.

We are watching a major shift unfold as Large Language Models and generative tools step out of tech labs and directly into everyday drug safety work. Right now, teams are deploying these systems across a wide mix of operational duties. They are taking on the heavy lifting for tasks like managing Individual Case Safety Reports-specifically by cleaning up ICSR files-along with isolating adverse events, condensing long patient histories, and sorting through endless medical journals. On top of that, researchers are actively trying to push these models into complex analytical territory to map out causal links, track safety signals, and evaluate drug-to-drug interactions.

Looking at the evidence, the main upside is clear: these algorithms are incredibly good at breaking up massive data bottlenecks. They can digest huge piles of cluttered safety data in a single heartbeat, taking a massive load of mindless, repetitive clerical work off human shoulders. Even so, nobody is ready to let them run completely on autopilot. The technology is still highly prone to inventing facts out of thin air, inheriting systemic human biases, leaking private details, and throwing unexpected coding glitches. Worse, they offer very little transparent logic, and we still lack the long-term field data needed to fully trust them.

That is why AI belongs in the passenger seat. For now, it must remain a supportive, human-in-the-loop tool rather than a standalone decision-maker. Getting these systems safely into routine workflows means setting up bulletproof testing protocols, tight quality checks, ironclad data walls, and strict oversight from regulatory experts. The path forward is all about building hyper-focused, domain-specific models. Only when these tools prove themselves in day-to-day operations can we build a faster, high-tech safety network without risking patient well-being or clinical accuracy.

## REFERENCES


1. Sudha TYS, Babu R, Sasanka KSBS, Mahato SK, Singh H. Advancing pharmacoepidemiology and pharmacovigilance: A scoping review of challenges, technological innovations, and global strategies for drug safety. J Appl Pharm Sci 2026;16 (5):71–82.
2. Dimitsaki S, Natsiavas P, Jaulent MC. Applying AI to Structured Real-World Data for Pharmacovigilance Purposes: Scoping Review. J Med Internet Res 2024;26:e57824.
3. Kumar RKS, Velusamy S. Harnessing Artificial Intelligence for Enhanced Pharmacovigilance: A Comprehensive Review. Indian J Pharm Pract 2025;18 (2):171–179.
4. Khemani B, Malave S, Shinde S, Shukla M, Shikalgar R, Talwar H. AI-driven pharmacovigilance: Enhancing adverse drug reaction detection with deep learning and NLP. MethodsX 2025;15:103460.
5. Hakim JB, Painter JL, Ramcharran D, Kara V, Powell G, Sobczak P, Satō C, Bate A, Beam A. The need for guardrails with large language models in pharmacovigilance and other medical safety critical settings. Sci Reports 2025;15:27858.
6. Salvo F, Crupi L, Cholle C. Language models in pharmacovigilance: Applications, promises and limits. Therapies 2026;81 (2):140–146.
7. Algarvio RC, Conceição J, Rodrigues PP, Ribeiro I, Ferreira-da-Silva R. Artificial intelligence in pharmacovigilance: A narrative review and practical experience with an expert-defined

Bayesian network tool. Int J Clin Pharm 2025;47 (4):932–944.

8. Ramcharran D, Painter JL, Kara V, Glaser M, Vanini M, Chalamalasetti VR, Golds C, Abdelkarim A, Bate A, Stegmann JU. Orchestrating generative AI in pharmacovigilance: predicting and preempting the unpredictable. Ther Adv Drug Saf 2025;16:20420986251396023.
9. Praveen J. Empowering pharmacovigilance: Unleashing the potential of generative AI in drug safety monitoring. J Innov Appl Pharm Sci 2023;8 (2):464.
10. Hirschman L. The Promise and Challenge of Large Language Models for Pharmacovigilance. Drug Safety 2026;49 (2):145–146.
11. Heckmann NS, Papoutsi DG, Barbieri MA, Battini V, Mølgaard SN, Schmidt SØ, Melskens L, Sessa M. Biomedical Large Language Models and Prompt Engineering for Causality Assessment of Individual Case Safety Reports in Pharmacovigilance. Pharm Res 2026;43 (8):2761–2774.
12. Ren J, Carroll H, McCarthy K, Allen J, Tam J, Philip J, Mehta M, Clark D. A KNN based machine learning approach to automating causality assessment of adverse events. Sci Reports 2026;16:9140.
13. Zitu MM, Owen D, Manne A, Wei P, Li L. Large Language Models for Adverse Drug Events: A Clinical Perspective. J Clin Med 2025;14 (15):5490.
14. Warner J, Prada Jardim A, Albera C. Artificial Intelligence: Applications in Pharmacovigilance Signal Management. Pharm Med 2025;39 (3):183 –198.
15. Agarwal A, Singh G, Jain S, Mittal P. Beyond boundaries: Charting the frontier of healthcare with big data and AI advancements in pharmacovigilance. Health Sci Rev 2025;14:100214.
16. Spanakis M, De Pauw A, Brumer M, Symvoulakis EK, De Loof H. Closing the loop: human-augmented, mechanistically enhanced AI for proactive management of drug-drug interactions. Front Pharmacol 2026;17:1767646.
17. Ong JCL, Chen MH, Ng N, Elangovan K, Tan NYT, Jin L, Xie Q, Ting DSW, Rodriguez-Monguio R, Bates DW, Liu N. A scoping review on generative AI and large language models in mitigating medication related harm. npj Digit Med 2025;8 (1):182.
18. Jamil S, Saha S, Misra R. Enhancing adverse drug event extraction and summarization for cancer drugs through large language models. J Biomed Inform 2026;178:105009.
19. Yadav R, Gupta AK, Mishra AK, Singh P, Kumar A, Pandey A, Pandey C. Artificial Intelligence for Next-Gen Pharmacovigilance and Post Marketing Surveillance in Evolving the Healthcare Environment. Int J Aquat Res Environ Std 2025;6 (S4):157–175.
20. Zou H, Wang Y, Xiang K, Huang A. Knowledge-augmented large language model-based framework for adverse drug reaction analysis. Appl Soft Comput 2025;185:114025.

HOW TO CITE: Samiksha Jadhav, Appasaheb Kuhile, Yogita Lanke, Sujata Nagare, Generative Artificial Intelligence and Large Language Models in Pharmacovigilance: Current Applications, Challenges and Future Perspectives., Int. J. of Pharm. Sci., 2026, Vol 4, Issue 11, 1973-1984. https://doi.org/10.5281/zenodo.24155757