



International Journal of Science, Architecture, Technology and Environment

Artificial Intelligence in Pharmacovigilance: Transforming Adverse Drug Reaction Detection, Signal Detection and Drug Safety Monitoring

Nikath Sulthana^{1*}

¹III Pharm D Department of Pharmacology Pharmacy Practice Hospital Pharmacy Division Nimra College of Pharmacy Vijayawada 521456 AP INDIA

*Corresponding author, nikhathahmedk786@gmail.com

DOI: <https://doi.org/10.63680/ijstate102603.01>

Abstract

Pharmacovigilance is an essential component of the medicine life cycle, concerned with the detection, assessment, understanding and prevention of adverse effects and other medicine-related problems. Conventional pharmacovigilance systems depend on spontaneous adverse drug reaction reporting, individual case safety reports, clinical databases, electronic health records, epidemiological studies and regulatory assessment. Although these systems have contributed substantially to medicine safety, the increasing volume, heterogeneity and complexity of health data have created significant challenges for timely and comprehensive safety monitoring. Artificial intelligence (AI), particularly machine learning (ML), deep learning, natural language processing (NLP) and large language models, has emerged as a potential means of augmenting pharmacovigilance activities. AI-based systems can process structured and unstructured data at a scale that is difficult to achieve through entirely manual workflows. Applications include adverse drug reaction and adverse drug event detection, automated case intake and triage, clinical-text mining, social-media surveillance, duplicate detection, case prioritization, safety signal detection, real-world data analysis and support for regulatory safety monitoring. Machine learning approaches can identify complex patterns in large datasets, while NLP techniques can extract drug-event relationships from clinical narratives and patient-generated content. More recent transformer-based models and large language models provide additional capabilities for contextual interpretation and information extraction, although their use introduces concerns regarding hallucination, reproducibility, transparency and validation. Despite substantial progress, AI should not be regarded as an autonomous replacement for pharmacovigilance professionals. Data incompleteness, reporting bias, confounding, class imbalance, poor generalizability, algorithmic bias, privacy concerns and limited interpretability can affect the reliability of AI-generated outputs. Governance, validation, human oversight and continuous performance monitoring are therefore essential. Recent regulatory initiatives, including the joint FDA and European Medicines Agency guiding principles for good AI practice in drug development, emphasize human-centered design, risk-based assessment, data governance, multidisciplinary expertise and life-cycle management. This review examines the role of AI throughout the pharmacovigilance process, from adverse-event identification and case processing to signal detection and real-world safety surveillance. It also discusses machine learning, NLP, deep learning, large language models, explainable AI, governance, regulatory

considerations, limitations and future directions. A human-centered model in which AI augments rather than replaces pharmacovigilance expertise may provide a practical framework for responsible implementation.

Keywords: Artificial intelligence; pharmacovigilance; machine learning; adverse drug reactions; natural language processing; safety signal detection; real-world data; large language models.

INTRODUCTION

Medicines provide major benefits in the prevention, diagnosis and treatment of disease, but their use is inevitably associated with the possibility of adverse effects. Some adverse reactions are predictable from pharmacological properties, whereas others may be uncommon, delayed, population-specific or difficult to identify during clinical development. Pharmacovigilance therefore extends beyond the initial demonstration of efficacy and safety in clinical trials and continues throughout the entire medicine life cycle. The World Health Organization defines pharmacovigilance as the science and activities concerned with the detection, assessment, understanding and prevention of adverse effects or other possible drug-related problems. Its broader purpose is to support safer medicine use and reduce preventable harm. [19]

Traditional pharmacovigilance has relied heavily on spontaneous reporting systems, individual case safety reports, medical literature, clinical investigations, regulatory databases and epidemiological studies. These sources remain fundamental. However, the modern healthcare environment generates an increasingly large quantity of digital information through electronic health records, insurance and claims databases, patient registries, social-media platforms, mobile applications, online patient communities and other real-world data sources. The scale of these data creates an opportunity for earlier identification of potential safety problems but simultaneously increases the burden on pharmacovigilance professionals who must review, interpret and prioritize information.

The increasing volume of information has encouraged researchers to investigate artificial intelligence as a means of supporting pharmacovigilance. AI is an umbrella term encompassing computational approaches capable of performing tasks that traditionally require aspects of human intelligence, including pattern recognition, classification, prediction, language interpretation and decision support. Within pharmacovigilance, machine learning has been investigated for data ingestion, adverse-event detection, signal detection, case prioritization and other tasks. A scoping review by Kompa and colleagues demonstrated that AI and machine-learning applications had already appeared across multiple pharmacovigilance activities, although the authors also noted that many studies did not reflect contemporary best practices in machine learning. [1]

The development of AI in pharmacovigilance has therefore progressed from relatively simple rule-based automation toward statistical machine learning, deep learning, transformer architectures and large language models. Systematic reviews have described the increasing use of machine-learning and deep-learning techniques for extracting adverse drug events from clinical text. [2] More recent methodological work has examined how machine-learning and deep-learning models perform tasks such as named-entity recognition and relation extraction in clinical adverse-event datasets. [3]

A major area of development is natural language processing. Pharmacovigilance information is frequently embedded in free-text documentation rather than presented in standardized structured fields. Clinical notes may contain medication names, doses, treatment indications, symptoms, laboratory abnormalities, temporal relationships and suspected adverse events within the same narrative. Extracting these relationships manually is time-consuming. NLP systems can potentially identify relevant concepts and relationships

automatically. A systematic review by Modi and colleagues showed that adverse drug event extraction from clinical notes has evolved through rule-based, machine-learning, deep-learning, hybrid and large-language-model approaches. [4]

AI can also extend pharmacovigilance beyond conventional clinical documentation. Patients increasingly discuss medicines and their experiences on social-media platforms and online communities. These sources can contain descriptions of suspected adverse reactions that may not immediately enter formal reporting systems. NLP and sequence-labeling methods have therefore been investigated for identifying adverse drug reaction mentions from social-media content. [8] Reviews have examined the broader potential and limitations of social-media pharmacovigilance. [9] Research has also demonstrated pipelines for early detection of adverse drug reactions in social-health networks. [10]

The development of deep linguistic features has further expanded automated adverse-event detection from social media. [11] Patient reviews have also been investigated as a source for identifying adverse drug reactions using NLP techniques. [12] However, converting social-media observations into reliable pharmacovigilance evidence remains challenging because informal language, sarcasm, spelling variation, incomplete clinical information and uncertain drug-event relationships can generate noise. Bousquet and colleagues identified several major challenges that must be addressed before social-media analysis can become efficiently integrated into operational pharmacovigilance. [13]

AI is not limited to unstructured information. Structured real-world data can also be analyzed using computational approaches. Electronic health records and large healthcare databases provide opportunities for active surveillance, outcome identification and epidemiological investigation. Research on signal detection from electronic healthcare databases has demonstrated the potential of data mining for identifying safety issues that may not be readily visible through spontaneous reports. [15] European initiatives such as the EU-ADR project demonstrated the feasibility of combining electronic healthcare databases for large-scale drug-safety monitoring. [16]

Large-scale surveillance infrastructures also illustrate how modern pharmacovigilance increasingly depends on distributed real-world data. The U.S. Food and Drug Administration Sentinel System was developed to monitor medical-product safety using information from disparate data sources and has supported numerous analyses informing regulatory decisions. [18] Such infrastructures provide a foundation on which AI can potentially be deployed for more efficient identification and prioritization of safety signals.

At the same time, AI introduces its own risks. Models can reproduce biases present in their training data, generate false-positive or false-negative results, fail when applied to new populations, or produce outputs that are difficult for human users to interpret. Governance is therefore central to successful AI implementation. Recent work on AI/ML governance in pharmacovigilance emphasizes the need for defined responsibilities, quality management, appropriate validation and continued human oversight. [6]

The regulatory environment is also evolving. In January 2026, the U.S. Food and Drug Administration and European Medicines Agency jointly released guiding principles for good AI practice in drug development. The principles address human-centered design, risk-based approaches, standards, context of use, multidisciplinary expertise, data governance, model development, performance assessment, life-cycle management and clear communication. [20]

Against this background, the present review examines how AI can transform pharmacovigilance while considering the scientific, operational, ethical and regulatory requirements necessary for responsible implementation.

PHARMACOVIGILANCE: CONCEPTS AND CONVENTIONAL WORKFLOW

Pharmacovigilance traditionally involves the continuous collection, assessment and interpretation of information relating to the safety of medicines. Although the precise organizational structure differs among countries and organizations, the underlying objective remains the same: identifying potential risks and generating evidence that can support appropriate action.

A conventional pharmacovigilance workflow may include:

1. Collection of safety information.
2. Identification of suspected adverse events.
3. Creation or receipt of individual case safety reports.
4. Data entry and coding.
5. Medical review.
6. Assessment of seriousness and expectedness.
7. Evaluation of causality and clinical context.
8. Duplicate detection.
9. Follow-up for missing information.
10. Aggregate analysis.
11. Signal detection.
12. Signal validation and prioritization.
13. Further epidemiological or clinical investigation.
14. Regulatory communication and risk-management action.

Spontaneous reporting systems are particularly important because they can identify previously unrecognized adverse reactions after medicines are used in larger and more diverse populations. However, spontaneous reporting is affected by under-reporting, stimulated reporting, reporting differences between countries, differences in awareness, incomplete clinical information and other forms of reporting bias.

Quantitative signal-detection methods have traditionally been developed to identify disproportionate reporting patterns. Such methods can provide statistical indicators that a particular drug-event combination is reported more frequently than expected. Bate and Evans described quantitative approaches to signal detection using spontaneous adverse-drug-reaction reports and highlighted the importance of statistical methods in systematic safety surveillance. [14]

Nevertheless, a statistical signal is not equivalent to a confirmed causal relationship. Signal detection is fundamentally a hypothesis-generating activity. A potential signal requires clinical assessment, evaluation of alternative explanations, consideration of biological plausibility and often additional epidemiological evidence.

Electronic healthcare databases offer another complementary source of evidence. Trifirò and colleagues examined the use of electronic health-record databases for signal detection and highlighted the importance of selecting appropriate outcomes and surveillance strategies. [15] The EU-ADR project subsequently demonstrated how several European electronic healthcare databases could be combined for large-scale drug-safety monitoring. [16]

The development of large-scale active surveillance systems has further changed the pharmacovigilance landscape. The Sentinel System represents a national infrastructure designed to use disparate healthcare data to monitor medical-product safety and support regulatory decision-making. [18]

AI should therefore be understood as an additional analytical layer rather than an independent pharmacovigilance system. Its value depends on the quality of the underlying data, the appropriateness of the analytical method, the clinical meaning of the output and the ability of pharmacovigilance experts to evaluate the resulting evidence.

ARTIFICIAL INTELLIGENCE AND ITS MAJOR COMPONENTS

Artificial intelligence encompasses multiple computational approaches rather than one specific technology. In pharmacovigilance, the most relevant categories include machine learning, deep learning, natural language processing, neural networks, transformer models and increasingly large language models.

AI technique	Main function	Pharmacovigilance application
Machine learning	Pattern recognition and classification	ADR detection and signal detection
Deep learning	Complex feature extraction	Clinical text and image/text analysis
NLP	Processing human language	ADR extraction from clinical notes and social media
Transformer models	Contextual language understanding	Drug–event relationship extraction
Generative AI/LLMs	Text generation and summarization	Case processing and literature review
Explainable AI	Interpretation of model outputs	Supporting expert review and accountability

TABLE-01: AI TECHNIQUES AND THEIR APPLICATIONS IN PHARMACOVIGILANCE.

Machine Learning

Machine learning allows computational systems to identify patterns from data and use those patterns to classify, predict or prioritize new observations. Depending on the task, machine learning can be supervised, unsupervised or semi-supervised.

Supervised learning uses labelled examples. In pharmacovigilance, a model may be trained using clinical reports labelled as containing or not containing an adverse event. The model then learns features associated with those labels.

Unsupervised learning does not require predefined labels and can be used for clustering, anomaly detection or discovering previously unrecognized patterns.

Semi-supervised approaches combine smaller amounts of labelled information with larger amounts of unlabelled data, potentially reducing the burden of manual annotation.

Machine learning has been explored across numerous pharmacovigilance tasks. The scoping review by Kompa and colleagues demonstrated that applications included data ingestion, information extraction and other pharmacovigilance activities. [1]

Deep Learning

Deep learning is a subset of machine learning based primarily on multilayer neural networks. These systems can learn complex representations directly from large datasets.

Convolutional neural networks, recurrent neural networks, long short-term memory networks and transformer-based architectures have all contributed to biomedical NLP and clinical information extraction.

The increasing availability of benchmark datasets has enabled systematic evaluation of deep-learning methods for adverse drug event extraction. Li and colleagues reviewed machine-learning and deep-learning approaches applied to clinical text and examined their performance in tasks including named-entity recognition and relation classification. [3]

Natural Language Processing

Natural language processing enables computers to process human language. It is particularly relevant to pharmacovigilance because important safety information often exists in unstructured text.

NLP can potentially identify:

- Drug names
- Adverse events
- Symptoms
- Laboratory abnormalities
- Dosage
- Route of administration
- Treatment indication
- Duration of exposure
- Temporal relationships
- Drug-event relationships
- Severity
- Patient characteristics

The extraction process frequently involves named-entity recognition followed by relation extraction. The systematic review by Modi and colleagues identified these as central tasks in automated adverse-drug-event extraction from clinical notes. [4]

DATA SOURCES FOR AI BASED PHARMACOVIGILANCE

AI systems can analyze multiple sources of safety information. Each source has distinct advantages and limitations.

Spontaneous Reporting Systems

Spontaneous reports remain one of the principal sources of post-marketing safety information. They provide direct reports of suspected adverse reactions and can reveal unexpected events.

AI can potentially assist with:

- Report classification
- Case prioritization

- Duplicate detection
- Data extraction
- Coding support
- Signal generation

However, models trained on spontaneous reports inherit limitations such as under-reporting, reporting bias and incomplete information.

Data source	Examples	Type of data	PV application
Spontaneous ADR reports	ICSR databases	Structured/semi-structured	ADR and signal detection
Electronic health records	Patient records, laboratory data	Structured + unstructured	ADE detection and monitoring
Clinical notes	Physician/nurse notes	Unstructured text	NLP-based ADR extraction
Social media	Patient posts, forums	Unstructured text	Early ADR identification
Scientific literature	Articles, case reports	Unstructured text	Literature surveillance
Patient reports	Patient-generated information	Structured + unstructured	Patient-centered safety monitoring
Real-world databases	Claims, registries	Structured	Active safety surveillance

TABLE-02: DATA SOURCES FOR AI-BASED PHARMACOVIGILANCE.

Electronic Health Records

Electronic health records contain information on prescriptions, diagnoses, laboratory results, clinical encounters and patient outcomes. They can therefore provide longitudinal information that is often unavailable in spontaneous reports.

AI can potentially identify temporal associations between medication exposure and clinical outcomes. However, electronic health records are not primarily designed for pharmacovigilance. Differences in documentation practices, coding systems, missing data and confounding by indication must be considered.

The use of electronic healthcare databases for active safety surveillance has been studied for many years, including through European collaborative initiatives. [15] The EU-ADR project illustrated the value of combining healthcare databases for large-scale drug-safety monitoring. [16]

Social Media

Social media provides patient-generated information in near-real-world settings. Patients may describe symptoms, treatment experiences and perceived relationships between medicines and adverse effects.

NLP methods have been used to identify adverse drug-reaction mentions in social-media posts. [8] Reviews have demonstrated that social-media pharmacovigilance can provide supplementary safety information but requires careful handling of data quality and clinical interpretation. [9]

Early detection pipelines have also been developed to identify possible adverse drug reactions from social-health networks. [10]

Patient Reviews and Online Communities

Patient reviews represent another form of patient-generated data. Unlike conventional spontaneous reports, reviews may describe symptoms in natural language and provide information about the patient's experience with a medicine.

NLP can be applied to extract potential adverse reactions from these reviews [12]. Nevertheless, online information may lack confirmation, standardized terminology and sufficient clinical context.

Biomedical Literature

Scientific publications contain detailed clinical and pharmacological information. Automated literature surveillance can assist pharmacovigilance teams in identifying new reports, case descriptions and safety evidence.

AI-based literature screening may reduce manual workload, but automated identification should be followed by appropriate human review because not every mention of a drug and clinical event represents an adverse drug reaction.

AI IN ADVERSE DRUG REACTION DETECTION

One of the most direct applications of AI in pharmacovigilance is automated adverse drug reaction detection.

A conventional manual process requires a pharmacovigilance professional to read a report, identify suspected adverse events and determine whether relevant information has been documented. When the number of reports becomes very large, manual review can become a substantial operational burden.

Machine-learning systems can classify reports according to the likelihood that they contain relevant safety information. NLP models can then extract the specific drug and event information.

For example, a clinical note may state:

"Following initiation of Drug X, the patient developed severe nausea and dizziness two days later. Drug X was discontinued and symptoms improved."

A human reviewer can easily interpret this narrative. A computational system, however, must identify:

- Drug X as the medication.
- Nausea as an adverse event.
- Dizziness as another adverse event.
- Two days as a temporal relationship.
- Discontinuation as a management action.
- Improvement after withdrawal as potentially relevant evidence.

Modern NLP systems are designed to identify these relationships.

The literature indicates that ADE extraction is commonly divided into named-entity recognition and relation extraction. [4] Deep-learning and contextual language models have demonstrated strong performance in benchmark clinical datasets, although performance varies according to dataset, task and relation type. [3]

Automated detection may therefore act as a screening mechanism. Rather than replacing clinical reviewers, AI can identify reports requiring attention and organize information for human assessment.

AI ASSISTED INDIVIDUAL CASE SAFETY REPORT PROCESSING

Individual case safety report processing is a labour-intensive component of pharmacovigilance. Reports can arrive through multiple channels and may contain inconsistent terminology or incomplete information.

AI can potentially support several steps.

Case Intake

AI systems can classify incoming documents and determine whether they potentially represent safety reports.

Information Extraction

NLP can extract information including:

- Patient demographics
- Suspected medicinal products
- Concomitant medicines
- Indications
- Adverse events
- Dates
- Dosages
- Routes
- Outcomes
- Laboratory findings

Medical Coding Support

Automated systems can suggest standardized terminology for reported events. Human review remains important because clinical terms may be ambiguous.

Duplicate Detection

Multiple reports may describe the same clinical case. Machine-learning approaches can compare demographic characteristics, dates, medicines and clinical events to identify potentially duplicate cases.

Case Prioritization

AI can assign cases to review queues according to predefined criteria, such as seriousness, novelty or the presence of particular drug-event combinations.

The advantage of automation is not simply speed. AI can potentially create a more consistent first-pass process across large numbers of reports. However, incorrect classification can result in missed cases or unnecessary workload, making validation essential.




FIGURE-01: AI-BASED ADR DETECTION

NATURAL LANGUAGE PROCESSING IN PHARMACOVIGILANCE

NLP is one of the most important AI technologies for pharmacovigilance because a large proportion of clinically meaningful information is expressed in natural language.

Named-Entity Recognition

Named-entity recognition identifies relevant concepts within text. In pharmacovigilance, these may include drugs, adverse events, symptoms, diseases and clinical attributes.

For example: "Patient developed acute kidney injury after starting Drug A."

The model should identify:

- Drug A → medication
- Acute kidney injury → adverse event

Relation Extraction

Identifying entities alone is insufficient. The system must determine which drug is associated with which event.

A sentence may mention several medications and multiple symptoms. Relation extraction therefore establishes connections between entities.

Modi and colleagues found that adverse-event extraction research has used pipeline, joint-task and multi-task approaches, with deep learning increasingly represented among published methods. [4]

Context Recognition

Clinical language is highly contextual.

For example: “Patient denies nausea.”

The word “nausea” is present, but it does not represent an adverse event experienced by the patient.

Similarly: “History of rash with Drug A.”

The model must interpret whether this represents a current event, historical event or allergy.

Therefore, pharmacovigilance NLP requires contextual understanding rather than simple keyword matching.

Negation Detection

Negation is particularly important because phrases such as “no rash,” “denies dizziness” or “without evidence of hepatotoxicity” should not be interpreted as positive events.

Temporal Information

Temporal relationships can determine whether an event plausibly occurred after drug exposure.

“Three days after starting Drug B, the patient developed diarrhoea.”

This provides information that may be important for subsequent clinical assessment.

AI IN SOCIAL-MEDIA PHARMACOVIGILANCE

Patients increasingly communicate about medicines through online platforms. These discussions provide a potentially valuable supplementary source of safety information.




FIGURE-02: NLP PIPELINE IN PHARMACOVIGILANCE.

Nikfarjam and colleagues demonstrated the use of sequence-labelling approaches with word-embedding cluster features to identify adverse drug-reaction mentions from social media. [8]

A broader review by Sarker and colleagues examined the use of social-media data for pharmacovigilance and described both opportunities and methodological challenges. [9]

One advantage of social-media surveillance is the possibility of identifying patient experiences that may not immediately appear in conventional reporting systems. Online discussions may also contain information about treatment discontinuation, symptom severity and patient-perceived relationships between a medicine and an event. However, social-media information presents substantial challenges.

Patients may use:

- Slang
- Abbreviations
- Misspellings
- Informal descriptions
- Sarcasm
- Hyperbole
- Product nicknames

Furthermore, an online statement that “this medicine caused my headache” represents a patient's perception rather than established causality.

Nikfarjam and colleagues developed a natural-language-processing pipeline for early detection of adverse drug reactions in social-health networks, illustrating the potential for automated monitoring. [10]

Deep linguistic features have also been investigated for social-media adverse-drug-reaction detection. [11] Patient reviews have similarly been analyzed using NLP approaches. [12]

Despite this progress, operationalizing social-media pharmacovigilance remains difficult. Bousquet and colleagues described major challenges involving data access, identification of relevant information, validation, interpretation and integration into pharmacovigilance workflows. [13]

Therefore, social-media AI should be considered a supplementary surveillance mechanism rather than a replacement for validated pharmacovigilance systems.

MACHINE LEARNING AND SAFETY SIGNAL DETECTION

Safety signal detection aims to identify previously unknown or insufficiently characterized associations between medicines and adverse events.

Traditional quantitative signal detection methods can identify disproportionate reporting patterns. Bate and Evans described quantitative approaches using spontaneous ADR reports, emphasizing the role of statistical signal detection in pharmacovigilance. [14]

Machine learning provides an additional approach because it can incorporate multiple variables and potentially identify nonlinear relationships.

Potential applications include:

- Classification of drug-event pairs

- Ranking of potential signals
- Anomaly detection
- Temporal pattern recognition
- Multivariable risk prediction
- Integration of heterogeneous data sources

However, machine learning does not eliminate the need for statistical reasoning.

A model may identify a pattern because of confounding rather than a true drug effect. For example, a medicine may appear associated with a particular adverse outcome because it is preferentially prescribed to patients who already have a higher baseline risk of that outcome.

CONVENTIONAL STATISTICAL METHODS VERSUS AI-BASED APPROACHES

Conventional statistical signal detection has several important advantages. Established methods can be mathematically interpretable, reproducible and relatively straightforward to monitor.

AI models can potentially capture more complex patterns but may be harder to interpret.

The two approaches should therefore not necessarily be regarded as competitors.

A practical pharmacovigilance system could use statistical disproportionality methods for broad signal screening and AI models for additional prioritization or contextual analysis.




FIGURE-03: AI-ASSISTED SAFETY SIGNAL DETECTION.

For example:

Stage 1: Conventional statistical signal detection identifies drug-event combinations with disproportionate reporting.

Stage 2: Machine learning prioritizes signals using additional characteristics.

Stage 3: NLP extracts clinical context from free-text reports.

Stage 4: Epidemiological analysis evaluates the association in real-world data.

Stage 5: Pharmacovigilance experts review the evidence.

Such a layered approach can combine established pharmacovigilance principles with computational methods.

AI AND A REAL-WORLD DATA

Real-world data can originate from electronic health records, insurance claims, registries, pharmacy systems and other healthcare sources. AI can potentially identify patterns across these large datasets more rapidly than manual review.

Trifirò and colleagues investigated electronic health-record databases for pharmacovigilance signal detection and demonstrated the importance of selecting suitable outcomes for monitoring. [15]

The EU-ADR project showed how multiple European electronic healthcare databases could be combined to support large-scale drug-safety monitoring. [16]

Subsequent work examined the extent to which electronic healthcare databases could provide leverage for active drug-safety surveillance. [17]

More recent reviews have specifically considered AI applied to structured real-world data for pharmacovigilance. [5]

Real-world data are particularly valuable because they may contain information about populations that are underrepresented in clinical trials. They can also provide longitudinal information about exposure and outcomes. However, real-world data are not automatically high-quality evidence.

Potential limitations include:

- Missing data
- Inconsistent coding
- Confounding
- Misclassification
- Differences in healthcare utilization
- Changes in clinical practice
- Incomplete medication exposure information

AI models cannot automatically correct these limitations.

THE ROLE OF LARGE-SCALE SAFETY SURVEILLANCE SYSTEM

Modern pharmacovigilance increasingly depends on infrastructure capable of integrating information from multiple healthcare sources.

The FDA Sentinel System represents an important example. Developed in response to legislative requirements for medical-product safety monitoring, Sentinel uses data from disparate sources and has supported hundreds of analyses, including analyses informing regulatory decisions. [18]

Such systems create an important environment for AI because machine-learning models can potentially analyze large-scale real-world datasets while standardized infrastructure supports data governance and reproducibility.

AI could potentially assist Sentinel-like systems by:

- Prioritizing drug-outcome combinations.
- Detecting unusual temporal patterns.
- Identifying patient subgroups.
- Screening large numbers of potential outcomes.
- Supporting automated data-quality checks.
- Integrating structured and unstructured information.

However, the use of AI in large surveillance infrastructures must be carefully governed. A model that performs well in one database may perform differently in another because patient populations, coding practices and healthcare systems differ.

DEEP LEARNING OF TRANSFORMER MODELS

Deep-learning systems have become increasingly important for clinical text processing. Traditional machine-learning systems often require manually engineered features. Deep-learning models can learn representations directly from data. Recurrent neural networks and long short-term memory networks were previously widely used for sequential clinical text. Transformer-based architectures subsequently changed NLP by enabling more effective contextual representation. Transformer models can consider relationships among words across long passages. This is particularly useful in pharmacovigilance because clinical relationships may be expressed across multiple sentences.

Li and colleagues reviewed machine-learning and deep-learning methods for clinical text-based adverse drug event detection and emphasized the importance of evaluating models across benchmark datasets and extraction tasks. [3]

Large language models represent a further development. These models are trained on enormous quantities of text and can perform multiple language tasks with limited task-specific training.

Potential pharmacovigilance applications include:

- Case summarization
- Adverse-event extraction
- Drug-event relation extraction
- Literature screening
- Medical-text classification

- Query assistance
- Signal documentation
- Regulatory-document support

GENERATIVE AI AND LARGE LANGUAGE MODELS

Generative AI differs from conventional classification models because it can generate new text rather than simply assign a label.

A large language model may be asked:

“Summarize the suspected adverse reaction described in this case.”

The system could potentially transform a long clinical narrative into a structured summary.

Another possible task is:

“Identify all medicines and adverse events mentioned in this report.”

The model could produce structured information that can subsequently be reviewed.

A 2026 systematic review by Kim and colleagues examined applications of large language models in adverse drug reaction detection and pharmacovigilance and highlighted both their potential and the challenges associated with their use. [7]

One major risk is hallucination. A generative model may produce information that was not present in the source document. In pharmacovigilance, this could be particularly serious because invented clinical information could influence downstream assessment.

Another concern is reproducibility. The output of a generative system may vary according to model version, prompt, temperature or system configuration.

For this reason, generative AI should be treated as an assistive technology rather than an unquestioned source of clinical truth.

A safer implementation could require the model to:

1. Extract only information present in the source.
2. Identify the exact source passage.
3. Distinguish reported facts from interpretation.
4. Flag uncertainty.
5. Avoid generating unsupported causal conclusions.
6. Require human verification before final submission.

EXPLAINABLE ARTIFICIAL INTELLIGENCE

Explainability is particularly important in pharmacovigilance because AI outputs can influence patient-safety activities.

A pharmacovigilance professional should ideally understand why a system prioritized a particular case or signal.

For example, instead of simply reporting:

“Drug X–Event Y: high priority,” a system could indicate:

- Number of relevant reports.
- Time trend.
- Patient characteristics.
- Statistical evidence.
- Similarity to previous cases.
- Clinical-text features.
- Presence of seriousness criteria.
- Relevant real-world evidence.

Explainable AI methods can provide information about which variables contributed to a model's prediction.

However, explanation should not be confused with proof of causality.

A model may identify important predictive features without demonstrating that those features have a biological causal role.

Therefore, explainability should support human assessment rather than replace it.

BIAS AND DATA QUALITY

AI systems learn from data. If the data contain systematic biases, the resulting model may reproduce or amplify them.

Reporting Bias

Spontaneous reporting systems are affected by differences in reporting behaviour. Certain medicines or adverse events may attract greater attention than others.

Demographic Bias

If training datasets underrepresent certain demographic groups, models may perform poorly in those populations.

Language Bias

NLP systems developed primarily using English-language data may perform less effectively with other languages or multilingual clinical documentation.

Dataset Shift

A model developed using one hospital system may perform differently when implemented in another healthcare environment.

Labelling Errors

Machine-learning models depend heavily on the quality of their training labels. Incorrectly labelled adverse events can reduce performance. These issues demonstrate why AI performance should be evaluated in the specific context in which the system will be used.

FALSE POSITIVES AND FALSE NEGATIVES

No AI model is perfect.

A false positive occurs when the system identifies a potential adverse event or safety signal that is not genuinely relevant. A false negative occurs when a real event is missed. Both can create problems.

Excessive false positives may overwhelm pharmacovigilance teams with low-value cases. This can increase workload and delay review of genuinely important signals. These may be more concerning when the missed information relates to a serious or unexpected adverse event. The appropriate balance depends on the intended use.

For a preliminary screening system, higher sensitivity may be desirable because human reviewers can subsequently filter results.

For automated decision-making, however, higher requirements for specificity, validation and control would be necessary.

Therefore, model evaluation should consider more than overall accuracy. Relevant measures can include:

- Sensitivity
- Specificity
- Precision
- Recall
- F1 score
- Area under the receiver operating characteristic curve
- Calibration
- Positive predictive value
- Negative predictive value

The selected metrics should correspond to the intended clinical and pharmacovigilance application.

DATA PRIVACY AND SECURITY

Pharmacovigilance data frequently contain sensitive patient information.

Electronic health records may include:

- Names
- Dates
- Diagnoses
- Laboratory values
- Medication histories
- Clinical narratives
- Demographic information

AI systems must therefore operate within appropriate privacy and security frameworks.

Data governance should include:

- Controlled access
- De-identification where appropriate

- Secure storage
- Encryption
- Audit trails
- Data-use agreements
- Appropriate retention policies
- Monitoring of unauthorized access

Privacy is particularly important when large language models are used. Sensitive clinical information should not be submitted to external systems without appropriate safeguards and authorization.

The FDA and EMA's 2026 principles explicitly emphasize data governance, documentation and protection of sensitive information throughout the AI technology life cycle. [20]

GOVERNANCE OF AI IN PHARMACOVIGILANCE

AI implementation cannot be separated from governance.

Governance determines:

- Who is responsible for the model?
- Who validates it?
- Who approves its use?
- Who monitors performance?
- Who investigates failures?
- Who is responsible for final pharmacovigilance decisions?

Glaser and Littlebury emphasized that AI/ML governance in pharmacovigilance requires clearly defined responsibilities and appropriate oversight. [6]

A practical governance structure should involve multidisciplinary expertise.

Relevant professionals may include:

- Pharmacovigilance specialists
- Clinical pharmacologists
- Pharmacists
- Physicians
- Epidemiologists
- Biostatisticians
- Data scientists
- AI specialists
- Information-security professionals
- Regulatory specialists
- Quality-assurance professionals

This multidisciplinary structure is important because AI expertise alone does not provide sufficient understanding of pharmacovigilance requirements.

Governance area	Key consideration	Importance in pharmacovigilance
Model validation	Demonstration of accuracy and reliability	Ensures trustworthy outputs
Transparency	Documentation of model development and use	Supports accountability
Explainability	Ability to understand relevant AI outputs	Facilitates expert review
Data protection	Secure handling of patient information	Protects confidentiality
Bias monitoring	Evaluation of performance across datasets/populations	Reduces systematic errors
Human oversight	Expert review of important AI outputs	Supports responsible decision-making
Auditability	Maintaining records of model performance and decisions	Enables traceability
Continuous monitoring	Monitoring model performance after deployment	Detects performance deterioration
Change management	Controlled updating of AI models	Maintains validated performance
Regulatory compliance	Alignment with applicable regulatory requirements	Supports lawful implementation

TABLE-03: REGULATORY AND GOVERNANCE CONSIDERATIONS FOR AI IN PHARMACOVIGILANCE.

REGULATORY CONSIDERATIONS

Pharmacovigilance is a highly regulated field because its primary objective is patient safety. AI systems used in pharmacovigilance should therefore be evaluated according to their intended purpose and potential risk.

In January 2026, FDA and EMA jointly released ten guiding principles for good AI practice in drug development. These principles apply across the medicine life cycle and emphasize responsible development and use of AI. [20]

The principles include:

1. Human-centered design.
2. Risk-based approach.
3. Adherence to standards.
4. Clear context of use.
5. Multidisciplinary expertise.
6. Data governance and documentation.
7. Appropriate model-development practices.
8. Risk-based performance assessment.
9. Life-cycle management.
10. Clear and essential information.

These principles are highly relevant to pharmacovigilance because safety-monitoring systems may directly influence decisions about medicines.

Regulatory acceptance of AI should therefore depend not merely on whether an algorithm has high predictive performance in a research dataset but on whether the entire system is appropriately designed, validated, documented and monitored.

HUMAN OVERSIGHT IN AI-BASED PHARMACOVIGILANCE

The concept of “human-in-the-loop” is particularly appropriate for pharmacovigilance. In a human-in-the-loop system, AI performs computational tasks while human experts retain responsibility for interpretation and final decisions.

For example:

AI: identifies 500 potentially relevant reports.

Human reviewer: evaluates the prioritized cases.

AI: summarizes relevant clinical information.

Human reviewer: verifies the summary against the original source.

AI: identifies a possible drug-event signal.

Human experts: evaluate clinical plausibility, confounding, epidemiological evidence and other information.

This model recognizes that pharmacovigilance involves judgment and contextual interpretation. Human oversight is also important when models encounter unusual cases outside their training distribution.

AI ASSISTED PHARMACOVIGILANCE WORKFLOW

A future AI-assisted pharmacovigilance workflow could be organized into several layers.

Layer 1: Data Acquisition

Data are collected from:

- Spontaneous reports
- Electronic health records
- Claims
- Registries
- Scientific literature
- Social media
- Patient reviews

Layer 2: Data Preprocessing

AI systems identify duplicates, standardize terminology and detect missing or inconsistent information.

Layer 3: Information Extraction

NLP extracts:

- Drugs
- Events
- Symptoms
- Dates
- Outcomes
- Dosages
- Patient characteristics

Layer 4: Case Prioritization

Machine-learning models rank cases according to predefined criteria.

Layer 5: Signal Detection

Statistical and AI-based systems evaluate drug-event relationships.

Layer 6: Signal Characterization

Real-world data and clinical evidence are examined.

Layer 7: Expert Review

Pharmacovigilance professionals interpret the evidence.

Layer 8: Regulatory Action

If appropriate, findings may contribute to:

- Label updates
- Risk-minimization measures
- Safety communications
- Additional monitoring
- Further studies

This architecture allows AI to contribute throughout the pharmacovigilance lifecycle without making the technology solely responsible for safety decisions.




FIGURE-04: INTEGRATED WORKFLOW OF AI-ASSISTED PHARMACOVIGILANCE FROM DATA ACQUISITION TO REGULATORY DECISION MAKING.

ADVANTAGES OF AI IN PHARMACOVIGILANCE

Increased Processing Capacity

AI can process large volumes of information more rapidly than manual review.

Automation of Repetitive Tasks

Tasks such as document classification, information extraction and preliminary prioritization may be automated.

Earlier Identification of Potential Signals

Automated surveillance can continuously analyze incoming data and identify unusual patterns.

Integration of Multiple Data Sources

AI can potentially combine structured and unstructured information.

Consistency

A validated algorithm applies predefined rules consistently.

Reduction of Manual Burden

Pharmacovigilance professionals can devote more time to complex medical evaluation.

Support for Real-World Surveillance

AI can facilitate analysis of large electronic healthcare databases.

Patient-Generated Data Analysis

NLP can make large amounts of patient-generated information more accessible to pharmacovigilance teams.

These advantages explain why AI has attracted increasing attention in pharmacovigilance research. However, potential efficiency should not be interpreted as evidence that AI automatically improves patient safety. Implementation quality remains critical.

LIMITATIONS OF AI IN PHARMACOVIGILANCE

AI implementation has several important limitations.

Poor Data Quality

If the input data are incomplete or inaccurate, AI output may also be unreliable.

Lack of Interpretability

Complex models may be difficult for users to understand.

Bias

Training data may contain demographic, reporting or institutional biases.

Generalizability

A model may perform well on a benchmark dataset but poorly in another environment.

False Results

Both false-positive and false-negative outputs remain possible.

Dependence on Training Data

New drugs, new adverse events or emerging clinical terminology may not be adequately represented.

Privacy Concerns

Large datasets contain sensitive information.

Regulatory Uncertainty

AI technologies evolve faster than many existing regulatory frameworks.

Model Drift

Performance may change over time as clinical practice, populations and data characteristics change.

Human Overreliance

Users may place excessive confidence in AI outputs.

These limitations support a controlled and monitored implementation strategy rather than unrestricted automation.

Challenge	Description	Potential consequence
Data quality	Missing, incomplete or inconsistent data	Reduced model performance
Data heterogeneity	Different formats and coding systems	Difficult data integration
Bias	Under-representation or systematic differences in datasets	Unequal detection performance
False positives	Incorrect identification of safety signals	Increased workload
False negatives	Missed ADRs or safety signals	Potential safety concerns
Explainability	Complex models may be difficult to interpret	Reduced transparency
Privacy	Sensitive patient information requires protection	Ethical and regulatory concerns
Model drift	Performance may change as data patterns change	Reduced reliability over time
External validity	Model performance may differ between datasets or populations	Limited generalizability
Human oversight	AI outputs require expert interpretation	Need for appropriate governance

TABLE-04: MAJOR CHALLENGES AND LIMITATIONS OF AI IN PHARMACOVIGILANCE.

COMPARISSION OF CONVENTIONAL AND AI-ASSISTED PHARMACOVIGILANCE

Conventional pharmacovigilance remains essential and provides the foundation for modern safety monitoring.

Traditional approaches generally depend more heavily on manual review and established statistical techniques, whereas AI-assisted systems can add automated information extraction, machine-learning classification and large-scale pattern recognition.

However, AI does not make conventional methods obsolete. Instead, the two approaches can complement each other.

For example, disproportionality analysis can identify potential signals from spontaneous reports, while machine learning can help prioritize those signals. NLP can extract additional clinical information, while epidemiological studies can evaluate associations using real-world data.

This combination provides a more realistic model of AI implementation than the idea of replacing pharmacovigilance professionals with autonomous algorithms.

Feature	Conventional pharmacovigilance	AI-assisted pharmacovigilance
Data processing	Mainly manual/semi-automated	Highly automated
Data volume	More limited	Large-scale
Unstructured text	Difficult to process manually	NLP can process efficiently
ADR identification	Manual review + statistical methods	ML/NLP-assisted identification
Signal detection	Statistical algorithms	ML + statistical approaches
Processing speed	Relatively slower	Potentially faster
Human involvement	High	Human oversight remains essential
Explainability	Generally straightforward	Can be challenging
Bias	Human/data-related bias	Algorithmic + data bias
Validation	Established methods	Requires model validation and monitoring

TABLE-05: CONVENTIONAL VERSUS AI-ASSISTED PHARMACOVIGILANCE.

APPLICATIONS ACROSS THE PHARMACOVIGILANCE LIFECYCLE

AI can potentially contribute to multiple stages of medicine safety monitoring.

Pre-Marketing Development

AI may assist with analysis of safety information emerging during clinical development.

Clinical Trial Safety Monitoring

Automated analysis can identify adverse-event patterns across trial documentation.

Regulatory Submission

AI may support document review and information organization.

Post-Marketing Surveillance

AI can process large volumes of spontaneous reports and real-world data.

Signal Management

Machine learning can help rank potential drug-event combinations.

Literature Monitoring

NLP can identify safety-relevant publications.

Social-Media Monitoring

NLP can identify patient-reported adverse reactions.

Risk Management

AI-derived evidence can contribute to the broader evidence base used for risk assessment, provided that outputs are appropriately validated and interpreted.

Pharmacovigilance activity	AI application	Potential benefit
Case identification	NLP and machine learning	Rapid identification of potential ADR cases
Case intake	Automated information extraction	Reduced manual workload
Case processing	AI-assisted classification and coding	Faster processing of safety reports
Duplicate detection	Machine-learning algorithms	Identification of duplicate reports
ADR identification	NLP and deep learning	Extraction of drug–event relationships
Signal detection	ML and statistical methods	Identification of potential safety signals
Signal validation	AI-assisted pattern analysis	Support for expert evaluation
Literature surveillance	NLP and LLMs	Automated screening of publications
Risk assessment	Real-world data analytics	Identification of populations at increased risk
Post-marketing surveillance	Continuous AI monitoring	Earlier identification of emerging safety concerns

TABLE-06: APPLICATION OF AI ACROSS THE PHARMACOVIGILANCE LIFECYCLE.

CHALLENGES IN CLINICAL IMPLEMENTATIONS

Moving an AI system from research to operational pharmacovigilance presents significant challenges.

First, models must be integrated with existing pharmacovigilance infrastructure. A technically impressive algorithm may provide little benefit if it cannot communicate effectively with existing databases or case-processing systems.

Second, validation must be appropriate to the intended use. A model designed to assist case prioritization should not automatically be treated as validated for causal assessment.

Third, organizations need appropriately trained personnel. Pharmacovigilance professionals do not necessarily need to become data scientists, but they should understand the strengths and limitations of AI systems they use.

Fourth, quality systems must accommodate AI-specific processes.

Documentation should address:

- Model version
- Training data
- Validation dataset
- Performance metrics
- Intended use
- Known limitations

- Update procedures
- Monitoring requirements
- Change-control procedures

Glaser and Littlebury emphasized the importance of governance structures and clearly assigned responsibilities for AI/ML in pharmacovigilance. [6]

MODEL VALIDATION AND PERFORMANCE MONITORING

Validation should occur before operational deployment and continue after implementation.

Internal Validation

The model should be evaluated using data independent from the training dataset.

External Validation

Performance should be assessed using data from different populations or institutions.

Prospective Validation

Where possible, models should be evaluated under real operational conditions.

Continuous Monitoring

AI performance can change over time.

Changes may occur because of:

- New medicines
- New adverse-event terminology
- Changes in clinical practice
- Changes in reporting behaviour
- Changes in patient populations
- Changes in healthcare documentation

The 2026 FDA/EMA principles specifically emphasize life-cycle management and periodic reassessment of AI systems. [20]

Therefore, validation should not be treated as a one-time event.

FUTURE PERSPECTIVES

The future of AI in pharmacovigilance is likely to involve greater integration rather than complete automation.

One important direction is multimodal pharmacovigilance. Instead of analysing one data source at a time, future systems may combine spontaneous reports, electronic health records, literature, registries and patient-generated data.

Another direction is the development of domain-specific language models. General-purpose language models may be useful, but pharmacovigilance-specific models could potentially provide better performance for specialized terminology.

A third area is federated or privacy-preserving learning. Instead of transferring sensitive patient-level information to a central database, models may potentially learn across distributed institutions while reducing direct movement of patient data.

Another important development is explainable AI. Pharmacovigilance users will increasingly require systems that can communicate why a case or signal was prioritized.

Real-time safety monitoring is also likely to receive increasing attention. AI systems could potentially monitor continuously updated data streams rather than relying exclusively on periodic analysis.

However, future development should focus on clinically meaningful outcomes rather than technological novelty.

RESEARCH GAPS

Despite substantial research, important evidence gaps remain.

First, many AI pharmacovigilance studies are retrospective and use benchmark datasets. Performance in controlled research environments may not reflect operational performance.

Second, there is a need for standardized evaluation frameworks. Different studies use different datasets, outcome definitions and metrics, making direct comparison difficult.

Third, more work is required to evaluate AI across different populations, languages and healthcare systems.

Fourth, prospective studies should determine whether AI actually improves pharmacovigilance outcomes rather than simply achieving high machine-learning performance.

Fifth, the safety implications of generative AI require systematic evaluation. Hallucination, unsupported inference and reproducibility remain important concerns.

Sixth, explainability should be evaluated from the perspective of actual pharmacovigilance users.

Seventh, governance models require continued development as AI becomes more deeply integrated into regulated safety processes.

PROPOSED HUMAN CENTERED AI PHARMACOVIGILANCE MODEL

A practical future model can be described as Human-Centered AI Pharmacovigilance.

The model consists of five interconnected principles.




FIGURE-05: HUMAN-CENTERED AI PHARMACOVIGILANCE MODEL.

Principle 1: AI Performs High-Volume Computational Tasks

AI should be used where computational scale provides a meaningful advantage.

Examples include:

- Screening
- Classification
- Extraction
- Ranking
- Pattern detection

Principle 2: Pharmacovigilance Professionals Provide Clinical Interpretation

Human experts should evaluate the clinical meaning of AI-generated outputs.

Principle 3: Statistical and Epidemiological Methods Remain Essential

AI should complement, rather than replace, established pharmacovigilance methods.

Principle 4: Every Model Requires Governance

Models should have:

- Defined owners
- Documented intended use

- Validation
- Monitoring
- Change control
- Risk assessment

Principle 5: Patient Safety Remains the Primary Objective

The goal of AI adoption should not simply be reducing workload or cost.

The ultimate objective should remain the identification and management of medicine-related risks.

This human-centered approach is consistent with the direction of contemporary AI governance, including the FDA and EMA emphasis on human-centered design, risk-based approaches, multidisciplinary expertise, data governance and life-cycle management. [20]

CONCLUSION

Artificial intelligence is becoming an increasingly important technological component of modern pharmacovigilance. Machine learning, deep learning, natural language processing and large language models provide new opportunities to process large volumes of structured and unstructured safety information. [1]

AI can potentially support adverse drug reaction detection, individual case safety report processing, clinical-text mining, social-media surveillance, safety signal detection and real-world data analysis. Research has demonstrated the feasibility of extracting adverse drug events from clinical notes and identifying adverse-reaction information from patient-generated and social-media data. [3]

Natural language processing can improve the identification of drug-event relationships from unstructured clinical narratives, while automated approaches can reduce the workload associated with manual information extraction. [4]

Social-media and patient-generated information may provide supplementary sources of safety information, although these sources require careful validation because of incomplete clinical context, informal language and uncertain causality. [9]

Large electronic healthcare databases and active surveillance infrastructures provide additional opportunities for AI-supported medicine safety monitoring. [15]

Collaborative initiatives such as the EU-ADR project have demonstrated the value of combining electronic healthcare databases for large-scale drug-safety surveillance. [16]

Established quantitative signal-detection methods remain an important component of pharmacovigilance and can complement machine-learning approaches rather than being completely replaced by them. [14]

Large-scale surveillance infrastructures such as the U.S. FDA Sentinel System provide an important foundation for the continued development of data-driven approaches to medicine safety monitoring. [18]

At the same time, AI implementation presents challenges involving data quality, reporting bias, generalizability, privacy, interpretability, false-positive and false-negative results, and model performance over time. [6]

Large language models provide new possibilities for adverse-event extraction, summarization and pharmacovigilance information processing, but their outputs require appropriate validation and human verification because generative systems may produce unsupported information. [7]

The integration of AI into pharmacovigilance should therefore follow a human-centered approach in which computational systems assist with high-volume tasks while qualified pharmacovigilance professionals retain responsibility for clinical interpretation and safety decisions. [20]

Future development should focus not only on improving algorithmic performance but also on establishing robust governance, validation, transparency, data protection and continuous monitoring throughout the AI life cycle. [5]

Overall, AI should be regarded as an augmentation technology rather than an autonomous replacement for pharmacovigilance expertise. The combination of computational scalability with clinical and regulatory judgment may support more efficient, systematic and responsive medicine-safety monitoring while maintaining appropriate human oversight.

Declaration of Conflicting Interests

The authors declare no potential conflicts of interest with respect to the research, authorship and publication of this article.

Funding

The author received no financial support for the research, authorship and publication of this article.

References

1. Kompa B, Hakim JB, Palepu A, Kompa KG, Smith M, Bain PA, et al. Artificial Intelligence Based on Machine Learning in Pharmacovigilance: A Scoping Review. *Drug Saf.* 2022;45(5):477-491.
2. Pilipiec P, Liwicki M, Bota A. Using Machine Learning for Pharmacovigilance: A Systematic Review. *Pharmaceutics.* 2022;14(2):266.
3. Li Y, Tao W, Li Z, Sun Z, Li F, Fenton S, et al. Artificial intelligence-powered pharmacovigilance: A review of machine and deep learning in clinical text-based adverse drug event detection for benchmark datasets. *J Biomed Inform.* 2024; 152:104621.
4. Modi S, Kasmiran KA, Sharef NM, Sharum MY. Extracting adverse drug events from clinical notes: A systematic review of approaches used. *J Biomed Inform.* 2024; 151:104603.
5. 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.
6. Glaser M, Littlebury R. Governance of artificial intelligence and machine learning in pharmacovigilance: what works today and what more is needed? *Ther Adv Drug Saf.* 2024; 15:20420986241293303.
7. Kim TY, Oh WS, Jeong DH. Large Language Models in Adverse Drug Reaction Detection and Pharmacovigilance: A Systematic Review of Current Applications, Challenges, and Future Directions. *Diagnostics (Basel).* 2026;16(15):2435.

8. Nikfarjam A, Sarker A, O'Connor K, Ginn R, Gonzalez G. Pharmacovigilance from social media: mining adverse drug reaction mentions using sequence labeling with word embedding cluster features. *J Am Med Inform Assoc.* 2015;22(3):671-681.
9. Sarker A, Ginn R, Nikfarjam A, O'Connor K, Smith K, Jayaraman S, et al. Utilizing social media data for pharmacovigilance: A review. *J Biomed Inform.* 2015; 54:202-212.
10. Nikfarjam A, Ransohoff JD, Callahan A, Jones E, Loew B, Kwong BY, et al. Early Detection of Adverse Drug Reactions in Social Health Networks: A Natural Language Processing Pipeline for Signal Detection. *JMIR Public Health Surveill.* 2019;5(2): e11264.
11. Zhang Y, Cui S, Gao H. Adverse drug reaction detection on social media with deep linguistic features. *J Biomed Inform.* 2020; 106:103437.
12. Oyeboode O, Orji R. Identifying adverse drug reactions from patient reviews on social media using natural language processing. *Health Informatics J.* 2023.
13. Bousquet C, Dahamna B, Guillemin-Lanne S, Darmoni SJ, Faviez C, Huot C, et al. The Adverse Drug Reactions from Patient Reports in Social Media Project: Five Major Challenges to Overcome to Operationalize Analysis and Efficiently Support Pharmacovigilance Process. *JMIR Res Protoc.* 2017;6(9): e179.
14. Bate A, Evans SJW. Quantitative signal detection using spontaneous ADR reporting. *Pharmacoepidemiol Drug Saf.* 2009;18(6):427-436.
15. Trifirò G, Pariente A, Coloma PM, Kors JA, Polimeni G, Miremont-Salamé G, et al. Data mining on electronic health record databases for signal detection in pharmacovigilance: which events to monitor? *Pharmacoepidemiol Drug Saf.* 2009;18(12):1176-1184.
16. Coloma PM, Schuemie MJ, Trifirò G, Gini R, Herings R, Hippisley-Cox J, et al. Combining electronic healthcare databases in Europe to allow for large-scale drug safety monitoring: the EU-ADR Project. *Pharmacoepidemiol Drug Saf.* 2011;20(1):1-11.
17. Coloma PM, Trifirò G, Schuemie MJ, Gini R, Herings R, Hippisley-Cox J, et al. Electronic healthcare databases for active drug safety surveillance: is there enough leverage? *Pharmacoepidemiol Drug Saf.* 2012;21(6):611-621.
18. Brown JS, Mendelsohn AB, Nam YH, Maro JC, Cocoros NM, Rodriguez-Watson C, et al. The US Food and Drug Administration Sentinel System: a national resource for a learning health system. *J Am Med Inform Assoc.* 2022;29(12):2191-2200.
19. World Health Organization. Pharmacovigilance: ensuring the safe use of medicines. Geneva: World Health Organization; 2004. WHO Reference No. WHO/EDM/2004.8.
20. U.S. Food and Drug Administration; European Medicines Agency. Guiding Principles of Good AI Practice in Drug Development. January 2026.