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

OPEN ACCESS

Review Article

Artificial Intelligence and Machine Learning in Pharmacovigilance: Transforming Adverse Drug Reaction Detection and Signal Management.

Aditya Ghavat[1*], Yash Ghodke[2], Yash Mulay[3], Raj Dohale[4], Shreyash Pardeshi[5]

[1,2,3,4] *Research Scholar Department Of Pharmacy, Siddhi's Institute Of Pharmacy, Nandgaon, Maharashtra, India.*

[5] *Assistant Professor, Siddhi's Institute Of Pharmacy, Nandgaon, Maharashtra, India.*

ARTICLE INFO

Published: 11 Nov 2026

Keywords:

Adverse drug reactions, Artificial intelligence, Drug safety, Machine learning, Natural language processing, Regulatory integration, Signal detection

DOI:

10.5281/zenodo.24143171

ABSTRACT

The science and collection of practices known as pharmacovigilance focus on the identification, evaluation, comprehension, and avoidance of side effects and other issues associated to medications. Conventional pharmacovigilance mostly relies on literature, clinical records, spontaneous reports, and organized safety databases. These sources include useful information but are impacted by under-reporting, diverse terminology, insufficient case information, reporting delays, duplicate records, and the huge volume of data requiring human assessment. As tools to assist these operations, artificial intelligence (AI), machine learning (ML), natural language processing (NLP), and similar computational techniques are being studied more and more. Pharmacovigilance (PV) experimental uses of artificial intelligence (AI) have quickly progressed to routine usage. The promise of AI to transform drug safety monitoring is critically examined in this study, which focuses on practical implementation concerns including assuring AI's consistent and transparent performance, eliminating many sources of bias, and resolving interpretability issues. It highlights the shift of PV from experimental usage to regular, scalable capabilities. It looks at the evidentiary foundation for AI in particular applications, how it might improve actionable insights, and how businesses can protect themselves from unforeseen effects in multi-AI system setups. Current legislative and international initiatives increasingly stress responsible AI, human supervision, transparency and lifecycle governance. Therefore, rather than taking the role of pharmacovigilance experts, AI should be largely seen as an enhancement technology. Properly verified technologies may decrease repetitious burden and allow professionals focus on clinically meaningful evaluation, although final safety judgments should remain subject to appropriate expert and regulatory review.

***Corresponding Author:** Aditya Ghavat
Address: Research Scholar Department Of Pharmacy, Siddhi's Institute Of Pharmacy, Nandgaon, Maharashtra, India
Email ✉: adityaghavat14369@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.

INTRODUCTION

The use of artificial intelligence (AI) in pharmacovigilance (PV) has grown dramatically, with the potential to increase the speed and precision of adverse event identification. The transformation has been pushed by the rising challenges in drug development and post-market surveillance, including the enormous volume of data, complexity of medication–drug interactions, and patient variability. However, there are new difficulties when moving from experimental use to regular, reliable application. These are unfair, unfair, and that AI, and that AI's, and that AI's, and that AI's, and that AI's, and that AI's, and that AI's, and that AI's, post-implementation.[1,2]

Patients might benefit greatly from medications, but using them can also lead to adverse drug reactions (ADRs), dosage mistakes, interactions, and other safety issues. Therefore, pharmacovigilance is a crucial part of the medication life cycle. The World Health Organization characterizes vigilance as efforts aimed at avoiding, identifying, analyzing and interpreting adverse effects and other medical-product-related concerns. Pharmacovigilance's goals go beyond spotting specific events; they also involve preserving a positive benefit-risk ratio, encouraging the use of safer medications, and guiding clinical and regulatory action.[3]

The development of data mining algorithms for signal detection in spontaneous reporting systems (SRS) in the early 2000s marked the beginning of the use of AI in PV. Pioneering techniques such as the Bayesian Confidence Propagation Neural Network (BCPNN) method and the Multi-item Gamma Poisson Shrinker (MGPS) established the framework for more complex AI applications in PV. Since then, AI has been used in a number of PV applications, including as signal identification, automated case processing, and real-world evidence analysis.[4,5] Artificial intelligence encompasses computational systems designed to execute functions that typically demand elements of human cognitive ability, whereas machine learning constitutes a collection of techniques that extract patterns from datasets. Within the field of pharmacovigilance, these technological approaches serve to categorize reports, recognize entities and connections within textual content, identify anomalous correlations, forecast potential hazards, and rank cases according to priority for specialist examination.[6]

Artificial intelligence encompasses a diverse range of methodologies rather than representing a singular pharmacovigilance approach. These methodologies span conventional machine learning techniques including logistic regression, random forests, and gradient boosting algorithms; sophisticated deep-learning frameworks; natural language processing models; graph-based analytical methods; and the increasingly prevalent application of large language models (LLMs). The selection of an appropriate methodology is contingent upon multiple factors: the specific task requirements, the nature of the data source, how outcomes are defined, and the degree of interpretability that is necessary.

Contemporary literature reviews demonstrate escalating enthusiasm for artificial intelligence-driven adverse drug reaction identification and forecasting. However, these reviews simultaneously underscore significant constraints, including insufficient external validation across diverse populations, variability in dataset composition, and disparities in how model performance metrics are communicated. As a result, any deployment of these technologies should be grounded in empirical evidence, calibrated according to risk considerations, and reinforced by specialized pharmacovigilance knowledge and oversight.[7]

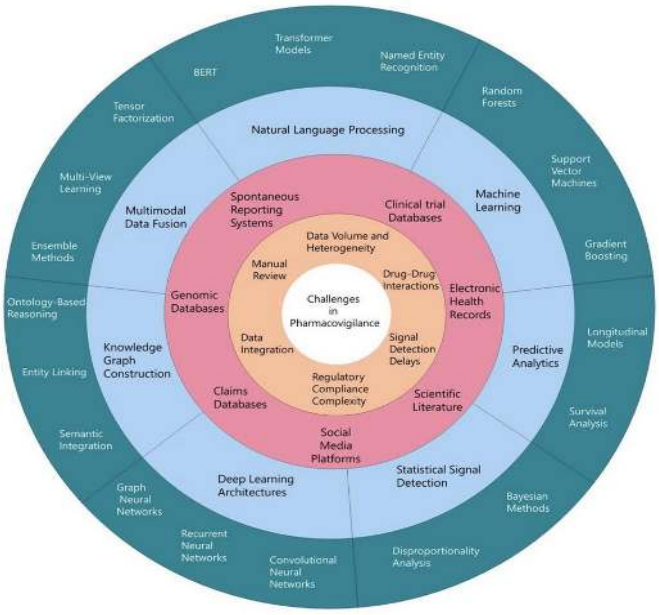


Figure 1: A Comprehensive Overview of State-Of-The-Art Challenges in PV, Sources of Data, and the Relevant AI Technologies to Address the Challenges.

AI, artificial intelligence; PV, pharmacovigilance.

This illustration depicts the intricate obstacles within photovoltaic systems and the range of artificial intelligence solutions available to tackle them. The nested circular layers illustrate the PV ecosystem's structure, encompassing fundamental issues, information repositories, and the assortment of AI and machine learning methodologies applicable to optimize PV operations. This holistic perspective underscores AI's capacity to manage the intricacies of contemporary photovoltaic technology through the synthesis of multiple data streams and the implementation of advanced computational methods.

2.Methodology of the review

This narrative review was developed using published scientific literature and authoritative regulatory and international resources concerning artificial intelligence, machine learning, adverse drug events, and pharmacovigilance. Literature was identified primarily through PubMed-indexed publications and official resources from the World Health Organization (WHO) and European Medicines Agency (EMA). Priority was given to systematic reviews, recent reviews, methodological literature, and regulatory guidance. Key concepts included artificial intelligence, machine learning, deep learning, natural language processing, pharmacovigilance, adverse drug reaction detection, signal detection, real-world data, and large language models. The literature was synthesized thematically rather than through quantitative meta-analysis. Particular attention was given to distinguishing between AI systems used for information extraction or prioritization and systems designed to make high-stakes safety decisions. This distinction is critical because pharmacovigilance decisions require clinical interpretation, evaluation of alternative explanations, and comprehensive consideration of the totality of evidence.

3. Fundamentals of Pharmacovigilance

3.1 Adverse Drug Reactions and Adverse Drug Events

An adverse drug reaction is typically characterized as a detrimental and unforeseeable response resulting from the administration of a medicinal product at standard dosages. The more comprehensive term adverse drug event encompasses medication-related injury and may include occurrences stemming from medication errors. Consequently, pharmacovigilance systems assess both discrete cases and consolidated evidence.

3.2 Individual Case Safety Reports

An ICSR documents information pertaining to a suspected adverse occurrence linked to one or more

medicinal products. Contingent upon the reporting framework, information may encompass patient demographics, implicated product, concurrent medications, event specifics, timeframes, analytical findings, resolution, severity classification and source identification. Partial or missing information is prevalent, rendering case evaluation and medical judgment essential human responsibilities.

3.3 Signal Detection

A safety signal constitutes data indicating a novel and potentially causal relationship, or an emerging characteristic of an established relationship, between a pharmaceutical agent and an occurrence that necessitates additional scrutiny. Signal detection alone does not confirm causality. Analytical or algorithmic approaches can recognize correlations warranting medical and population-based assessment. A comprehensive examination of spontaneous-reporting signal methodologies revealed considerable inconsistency across approaches and underscored that the most suitable method hinges on the pharmaceutical agent, occurrence, information repository and intended application.

3.4 Signal Management

Signal management typically encompasses identification, substantiation, verification or prioritization, examination and appraisal, and suggestion of fitting interventions. Based on accessible information, interventions may encompass ongoing surveillance, supplementary examination, modifications to product documentation, harm-reduction protocols or additional regulatory measures. Artificial intelligence can facilitate multiple phases but should not autonomously dictate regulatory determinations.[8]

Table 01: Artificial Intelligence and Machine Learning: Concepts

Approach	Basic principle	Potential PV application
Supervised ML	Learns from labelled examples.	ADR classification, case triage, outcome prediction.
Unsupervised ML	Finds patterns without predefined labels.	Clustering, anomaly detection, emerging event patterns.
Semi-supervised ML	Combines labelled and unlabelled data.	Large-scale text and report classification.
Deep learning	Uses multi-layer neural networks to learn representations.	Clinical-text extraction, complex prediction tasks.
NLP	Processes human language in structured or free text.	Drug/event extraction, literature screening, social-media mining.

The selection of an algorithm should be guided by the specific pharmacovigilance challenge at hand, not by prevailing technological developments. While simpler models frequently provide superior interpretability and streamlined validation processes, more sophisticated models can effectively identify nonlinear patterns and process unstructured data. Comprehensive model development must encompass well-defined objectives, adequate training datasets, rigorous validation procedures, and continuous post-deployment monitoring.

Table 02: Major Data Sources for AI-Enabled Pharmacovigilance

Data source	Strength	Important limitation
Spontaneous reporting systems.	Large-scale-Real-world safety reporting.	Under-reporting, reporting bias, missing data.
EHRs	Clinical context, laboratory values and longitudinal record.	Data heterogeneity, coding differences, privacy issues.
Claims/administrative data	Large populations and utilization pattern.	Limited clinical detail and confounding.
Social media/patient forums	Patient-generated, near-real-time experiences.	Noise, unverifiable report, slang and demographic bias.

Data source	Strength	Important limitation
Registries	Disease/product- specific longitudinal data.	Coverage may be restricted to selected populations.

The World Health Organization indicates that artificial intelligence applications in pharmacovigilance encompass forecasting of prospective adverse events, sophisticated signal identification, and natural language processing for retrieving data from unstructured documentation and Individual Case Safety Reports. The variety of available data sources enables cross-validation, whereby signals detected in a single source can be assessed against corroborating independent evidence. [3]

4. Applications of AI in PV

The integration of AI technologies in PV is revolutionizing the way ADRs are detected, evaluated, and managed. This section explores the current applications of AI in PV, focusing on its role in enhancing ADR detection and evaluation

4.1 AI for ADR detection

Adverse drug reactions (ADRs) represent a significant public health challenge, affecting patient populations across all demographic groups. The prevalence is particularly elevated among elderly individuals aged 60 years and older, where 15%–35% encounter an ADR during hospitalization. This substantial incidence creates considerable strain on healthcare infrastructure, with estimated annual costs for ADR management reaching $30.1 billion. This situation underscores the urgent requirement for enhanced approaches to monitor, identify, and mitigate ADRs during both pre-market evaluation and post-market surveillance phases of pharmaceutical oversight.

Artificial intelligence applications for ADR identification have undergone significant transformation since the late 1990s, fundamentally altering pharmacovigilance methodologies. This progression can be segmented into three distinct phases, each marked by notable improvements in statistical modeling, natural language processing capabilities, and machine learning applications.[9,10]

Early pharmacovigilance applications concentrated on improving signal identification within spontaneous reporting frameworks. The BCPNN methodology, created by Bate et al., and the MGPS system, established by DuMouchel, represented substantial advancements compared to conventional statistical approaches.[11] These computational techniques enabled more streamlined analysis of extensive ADR report collections, with potential to detect safety signals that conventional manual examination might overlook. Nevertheless, these approaches encountered limitations regarding rare event management and polypharmacy interactions. For instance, while the BCPNN methodology exhibits strong sensitivity for ADR identification, it tends to produce elevated false positive frequencies, creating concerns regarding its specificity. This underscores the necessity for rigorous interpretation of algorithm-generated signals and reinforces the critical role of qualified expert assessment in signal validation.

As pharmacovigilance information sources diversified to encompass unstructured content from electronic health records and digital platforms, natural language processing methodologies became essential. Nikfarjam et al. illustrated the effectiveness of natural language processing in retrieving ADR information from digital communities, obtaining F-measures of 0.82 and 0.72 for ADR identification from platforms including Daily Strength and Twitter. These F-measure values demonstrate an appropriate equilibrium between precision and recall, indicating that the natural language processing approaches successfully recognized ADRs while reducing both false positives and false negatives. This investigation established novel opportunities for continuous ADR surveillance, potentially capturing patient-generated health information that might not surface in official reporting channels.[12]

However, these sophisticated models frequently operate as opaque systems, generating apprehension regarding their interpretability and practical utility in clinical settings. Rudin advocated for creating models with built-in interpretability rather than applying explanatory techniques after the fact, especially when addressing critical medical applications such as ADR

identification. This underscores the persistent tension between leveraging the capabilities of sophisticated AI systems and ensuring adequate transparency and explainability within healthcare contexts.[13]

To mitigate data privacy challenges in ADR identification, federated learning methodologies have been introduced. Federated learning enables multiple healthcare organizations to jointly develop AI models while safeguarding raw patient information, thereby protecting privacy. Choudhury et al. presented a federated learning framework for coordinated ADR identification across various healthcare institutions while maintaining data confidentiality, proving its capacity to forecast ADRs without exposing raw patient information. Conversely, research by Enthoven and Al-Ars revealed potential weaknesses to privacy breaches within federated learning frameworks, stressing the importance of implementing comprehensive security protocols.[14]

Prospective investigations should prioritize creating more dependable, transparent, and clinically validated AI systems for PV. This encompasses tackling the complexities of uncommon ADRs and underserved demographic groups, strengthening the consolidation of varied information sources while preserving confidentiality, and constructing AI frameworks capable of responding to shifts in pharmaceutical advancement and application patterns. As noted by Rieke et al., deliberate evaluation of the moral and legal ramifications of artificial intelligence in medical practice is essential, notably in critical domains including ADR identification.[15]

4.2 AI for signal detection and active surveillance

Signal detection represents a critical element of pharmacovigilance that focuses on identifying potential associations between pharmaceutical products and adverse health outcomes. Although conventional approaches utilizing disproportionality analysis of spontaneous reporting systems have provided valuable insights, they frequently encounter significant limitations including incomplete reporting, systematic reporting biases, and difficulties in recognizing multifaceted drug-event relationships. Artificial intelligence and machine learning technologies have become essential instruments for overcoming these obstacles, strengthening signal detection capabilities, enhancing detection precision, and facilitating the examination of multiple data sources. In addition to comprehensive data source evaluation, artificial intelligence is being leveraged within specialized regulatory intelligence systems to optimize safety surveillance and documentation workflows.[16]

Current developments in artificial intelligence-powered regulatory intelligence systems, commonly implemented as AI-driven conversational interfaces designed to retrieve or extract information from publicly accessible regulatory materials, are substantially transforming how pharmacovigilance professionals obtain, process, and implement extensive regulatory and safety information. Numerous machine learning models have been implemented for signal detection, demonstrating considerable enhancements relative to established methodologies. Bae and colleagues created a Gradient Boosting Machine algorithm to identify safety signals associated with nivolumab and docetaxel medications. This algorithm demonstrated a remarkable AUC score of 0.95, substantially exceeding the performance of conventional disproportionality analysis techniques, which generally produce AUCs of 0.55 or below.[17]

The application of artificial intelligence for signal detection in practical environments has demonstrated considerable promise. The FDA's Sentinel System, which incorporates the Active Risk Identification and Analysis component, facilitates the assessment of post-market safety signals through real-world data at remarkable speed and effectiveness. Machine learning-based technologies can strengthen critical aspects of safety evaluations, such as health outcome of interest validation, propensity score matching, and patient phenotyping. Since its comprehensive rollout, the Sentinel System has facilitated over 250 safety analyses. Artificial intelligence offers the capacity to streamline repetitive manual operations, improve uniformity, and minimize human bias while delivering critical insights for early signal identification. These systems can also retrieve information from adverse drug event documentation and determine case validity independently, facilitating expedited case handling.

Nevertheless, additional investigation and practical implementation evidence are necessary to establish the concrete advantages of employing these solutions relative to conventional approaches.[18]

While artificial intelligence methodologies have exhibited substantial promise in strengthening signal detection, obstacles persist regarding the necessity for extensive, superior-quality datasets for algorithm development, the transparency of sophisticated models, and the incorporation of AI-produced signals into established pharmacovigilance operations. Furthermore, the application of artificial intelligence in pharmacovigilance introduces significant ethical and regulatory questions. These encompass safeguarding information confidentiality and protection, mitigating potential algorithmic prejudices, and formulating explicit protocols for the verification and deployment of AI-based signal identification approaches. Governmental authorities including the Food and Drug Administration (FDA) and European Medicines Agency (EMA) are actively creating governance structures to manage artificial intelligence utilization in pharmaceutical safety oversight, highlighting the importance of openness, verifiability, and human participation in AI-supported analytical procedures.[19]

4.3 AI for real-time monitoring and active surveillance

Technological progress in artificial intelligence is facilitating a shift toward more anticipatory surveillance methodologies. This transition enables earlier identification of potential safety concerns and faster implementation of corrective measures. Real-time monitoring within pharmacovigilance encompasses ongoing examination of multiple data streams to identify potential adverse drug reactions as they emerge. Conversely, active surveillance entails deliberate and systematic gathering and evaluation of safety information across a product's entire operational period. Artificial intelligence technologies have substantially enhanced both dimensions of pharmacovigilance. The preceding discussion highlighted a significant application of artificial intelligence in real-time monitoring—examination of social media content and patient discussion boards for expedited signal identification. Beyond this, a vital application of artificial intelligence involves the systematic examination of electronic health records to identify adverse drug reactions within hospital environments. Van de Burgt and colleagues created and deployed a computational text analysis method for extracting adverse drug reactions from unstructured Dutch electronic health records. The method leveraged the Medical Dictionary for Regulatory Activities terminology and Systematized Nomenclature of Medicine Clinical Terms classification system to identify potential adverse drug reactions within an extensive collection of 35,000 electronic health record documents. Their results demonstrated a positive predictive value of 70% and a sensitivity rate of 73% for adverse drug reaction identification. Nevertheless, they emphasize the importance of independent validation and assessment before implementing these computational methods in clinical hospital environments. Regarding active surveillance approaches, artificial intelligence has achieved notable advancements. For example, the Centers for Disease Control and Prevention has established multiple active surveillance frameworks for immunization safety monitoring, including the Vaccine Safety Datalink initiative. Initiated in 1990, the Vaccine Safety Datalink project oversees immunization safety and performs investigations into post-vaccination adverse occurrences. In 2005, the Vaccine Safety Datalink team introduced an active surveillance framework termed Rapid Cycle Analysis, intended to track adverse occurrences following immunization with minimal time delay. The incorporation of artificial intelligence into comparable frameworks has the capacity to further optimize their operational effectiveness and precision. [20]

4.4 Artificial Intelligence for Adverse Drug Reaction Detection

Artificial intelligence facilitates ADR detection by transforming extensive volumes of both structured and unstructured data into potential safety evidence. Within clinical documentation, natural language processing models can recognize pharmaceutical names, adverse-event terminology, clinical symptoms, laboratory results, negation patterns, and associations between medications and adverse events. Within structured electronic health record systems, machine learning models can examine interactions among patient age,

medical diagnoses, laboratory measurements, prescribed medications, and existing health conditions to forecast risk levels.

4.4.1 Named-Entity Recognition

Named-entity recognition locates pertinent information including pharmaceutical names, medical conditions, clinical symptoms, adverse events, and relevant patient demographics. Contemporary deep-learning methodologies can surpass conventional feature-engineering approaches in certain extraction applications. A 2024 examination of clinical-text adverse drug event extraction demonstrated robust performance from deep learning for named-entity recognition and highlighted exceptional outcomes for transformer-based BERT architectures in comprehensive end-to-end assignments.

4.4.2 Relation Extraction

Identifying isolated terms proves inadequate because pharmacovigilance demands comprehension of interconnections. Relation-extraction algorithms work to establish whether a specific pharmaceutical is linked to a documented adverse event and may incorporate temporal or situational context. For instance, the algorithm may differentiate between a patient who experienced a rash following medication initiation and a hereditary predisposition to rash unconnected to therapeutic intervention.[21]

4.4.3 Classification and Triage

Machine-learning classification systems can organize incoming data based on severity level, predictability, medication-event correlation, or handling urgency. These approaches may minimize repetitive human effort; however, false negatives warrant particular attention because overlooking a critical safety incident carries substantially greater risk than a low-priority false positive.

4.4.4 Risk Estimation Through Predictive Analytics

Predictive analytical models utilize patient-specific characteristics to quantify the likelihood of unfavourable health outcomes. According to a comprehensive systematic review, machine learning approaches have been examined incorporating demographic factors, laboratory measurements, and existing medical conditions, though the review emphasized that external validation is insufficiently represented across most published research. Consequently, predictive outputs should be understood as risk quantification rather than evidence establishing causality between a therapeutic intervention and an adverse outcome.[22]

4.5 Natural Language Processing in Pharmacovigilance

NLP holds significant value in pharmacovigilance because a substantial portion of safety monitoring data exists in unstructured text format. Medical documentation, Individual Case Safety Report narratives, peer-reviewed publications, case studies, and user-contributed information frequently contain medically relevant information that proves challenging to extract through simple keyword-based search methods alone.

4.5.1 NLP Pipeline

- Data collection and removal of personally identifiable information where applicable.
- Text standardization and linguistic analysis.
- Identification of entities including medications, adverse events, and medical terminology.
- Detection of negations and contextual factors.
- Extraction of relationships and temporal sequences.
- Categorization, risk assessment, or adverse event signal identification.
- Expert evaluation and validation.

Advanced transformer-based language models have significantly enhanced the capacity to interpret context and semantic connections within text. However, pharmacovigilance implementations require specialized validation due to the intricacy of medical language, prevalence of acronyms, and the potential for identical terms to carry distinct meanings depending on usage context.[21]

4.6 Artificial Intelligence and Real-World Evidence

Real-world data sourced from electronic health records, insurance claims, disease registries, and comparable databases can offer insights into medication

utilization and clinical outcomes in settings beyond controlled clinical trials. These datasets enable evaluation of associations across larger, more heterogeneous populations. Artificial intelligence can facilitate the extraction of relevant variables, management of intricate patterns, and detection of vulnerable patient subgroups. Nevertheless, real-world data does not inherently eliminate bias. Confounding by indication, incomplete data, shifts in prescribing patterns, disparities in healthcare availability, and inconsistencies in coding standards can all influence findings. An algorithmic model may capture healthcare system practices rather than genuine biological medication risks. Consequently, establishing causal relationships demands rigorous epidemiological methodology and specialized clinical knowledge alongside strong predictive accuracy.

Current systematic reviews demonstrate that external and temporal validation of artificial intelligence models for adverse events remains infrequent, and transparency methodologies show considerable variability. These gaps must be resolved prior to widespread implementation in clinical practice or regulatory decision-making.[23]

4.7 Deep Learning and Large Language Models

Deep learning enables the processing of intricate, high-dimensional data through neural networks. Historically, convolutional and recurrent architectures have been employed for text and sequence analysis, whereas transformer architectures have gained prominence in language-based applications. BERT-type models demonstrate the capacity to capture contextual relationships within clinical text and have delivered encouraging outcomes in identifying adverse drug events. Large language models offer expanded capabilities, encompassing document summarization, information retrieval, structured data extraction, and generation of pharmacovigilance narratives. Nevertheless, generative systems may generate inaccurate or unsubstantiated claims. Consequently, pharmacovigilance applications should prioritize well-defined tasks, outputs anchored in source material, validation procedures, and human oversight.

EMA guidance underscores the importance of secure data input, rigorous output validation, ongoing development, and expert consultation when uncertainties emerge regarding large language model implementation. These considerations are especially critical for safety-sensitive operations.[24]

5. Benefits of Artificial Intelligence and Machine Learning in Pharmacovigilance

- Management of extensive volumes of both structured and unstructured information.
- Accelerated detection and ranking of emerging safety concerns.
- Streamlining of routine case-handling procedures through automation.
- Enhanced identification of pharmaceutical-adverse event correlations within narrative text.
- Consolidation of diverse patient and pharmaceutical attributes.
- Capacity to uncover nuanced or multifaceted trends.
- Facilitation of continuous real-time monitoring across designated information repositories.
- Capability to establish consistent and prioritized operational procedures when properly verified.

These benefits represent achievable possibilities rather than assured outcomes. Results are contingent upon information integrity, algorithmic architecture, and operational deployment circumstances.

6. Obstacles and Constraints in Artificial Intelligence Applications for Pharmacovigilance

Although artificial intelligence demonstrates considerable potential for advancing pharmacovigilance operations, multiple obstacles and constraints must be resolved to facilitate broad implementation and achieve peak effectiveness. This discussion addresses four primary areas of concern: representation bias across heterogeneous patient demographics, temporal variability in medication safety assessments, inference of causality within intricate multi-drug treatment regimens, and consolidation of heterogeneous data formats.

6.1 Representation Bias Across Heterogeneous Patient Demographics

Artificial intelligence systems, particularly those deployed in pharmacovigilance, frequently rely on training datasets that inadequately capture the demographic composition of actual patient populations. This limitation frequently produces representation bias, wherein the system demonstrates strong predictive accuracy for dominant demographic groups while exhibiting diminished performance for marginalized or underrepresented cohorts, including ethnic minorities, geriatric populations, or individuals with multiple concurrent health conditions. As an illustration, systems developed using predominantly European or North American patient data may demonstrate reduced effectiveness in identifying adverse drug reactions among populations from Asian or African regions.

This challenge intensifies due to the tendency of clinical research studies—a fundamental source of training information—to systematically exclude or underrepresent these demographic groups. Such exclusion compromises the capacity of artificial intelligence systems to maintain consistent predictive reliability across all demographic segments, potentially resulting in overlooked safety concerns affecting particular populations. Addressing these representation challenges requires implementing data expansion methodologies and statistical rebalancing approaches to rectify disparities within training datasets, supplemented by sophisticated quantitative assessment and remediation techniques for identifying and correcting bias.41 Methodologies including distributed machine learning frameworks facilitate system development using geographically dispersed datasets originating from multiple geographic regions while preserving data confidentiality, thereby offering a promising strategy for reducing demographic bias.[25]

6.2 Establishing Causality in Complex Polypharmacy Scenarios

Polypharmacy, defined as the concurrent use of multiple medications by a single patient—particularly prevalent among elderly populations—represents a substantial obstacle for artificial intelligence applications in pharmacovigilance. Although AI algorithms excel at identifying correlations between drug combinations and adverse events, determining true causal mechanisms remains exceptionally difficult, especially when multiple drugs interact simultaneously. Resolving this challenge is critical for advancing AI from proof-of-concept applications to standard clinical practice in pharmacovigilance, where the objective must be delivering reliable, unbiased, and actionable clinical insights.

Conventional machine learning approaches can detect associations between specific drug combinations and elevated adverse drug reaction rates; however, determining whether a single agent or the synergistic interaction among multiple agents drives these outcomes requires substantially more sophisticated analysis. Causal inference methodologies, including causal graph neural networks, are currently under investigation to address this analytical complexity.[26]

Specialized computational frameworks have been developed to tackle causality assessment in drug interaction contexts. The InferBERT framework exemplifies this progress by combining transformer-based natural language processing with do-calculus methodology to identify causal relationships in pharmacovigilance datasets, achieving notable performance in distinguishing medication-induced adverse outcomes. Computational strategies including graph-based causal networks, propensity score matching, and instrumental variable analysis provide robust mechanisms for controlling confounding factors in observational research. Bayesian networks constructed with expert clinical input have demonstrated considerable effectiveness, enabling clinical expertise to guide causal model architecture while allowing empirical data to optimize parameter estimation. Incorporating these specialized causal inference approaches alongside conventional machine learning necessitates fundamental modifications to data collection frameworks, as existing pharmacovigilance databases currently lack the structural characteristics required for rigorous causal inference analysis.[27]

6.3 Explainable AI and Regulatory Considerations

Explainable AI (XAI) serves as a cornerstone for building trust and ensuring transparency within AI-powered photovoltaic systems. To meet regulatory requirements, AI algorithms must deliver transparent rationales for their predictions, allowing stakeholders to comprehend, assess, and leverage AI-generated insights effectively. This capability is essential for obtaining regulatory clearance and positioning AI as a dependable, standard feature in photovoltaic applications.

Multiple established methodologies have emerged to overcome the opacity inherent in AI systems used for photovoltaic analysis. Local Interpretable Model-Agnostic Explanations (LIME) and SHapley Additive exPlanations (SHAP) represent leading XAI frameworks that effectively pinpoint both significant and insignificant variables in photovoltaic models, with SHAP demonstrating marginally superior performance in quantifying feature importance. These approaches have proven successful in photovoltaic contexts, such as forecasting undesirable outcomes with accuracy rates reaching 72% while furnishing clear documentation of how individual features contribute to predictions. Despite their effectiveness, these techniques encounter limitations when dealing with correlated features, as most contemporary XAI methods either overlook feature interdependencies or incorrectly presume feature independence.[28]

Additional methodologies encompass Gradient-weighted Class Activation Mapping (Grad-CAM) for generating visual interpretations, layer-wise relevance propagation for attributing predictions through neural network layers, and Anchors for producing rule-based explanations that enhance stakeholder comprehension of model reasoning. Deploying these methodologies within photovoltaic systems would substantially improve model interpretability, enabling regulators and industry professionals to discern how AI systems formulate specific safety determinations and identify circumstances requiring human oversight.[29]

The illustration presents the fundamental regulatory and ethical dimensions essential for AI-integrated photovoltaic systems. It encompasses four pivotal components: the existing regulatory environment, data protection and cybersecurity protocols, ethical governance structures, and the trajectory of regulatory development. These interconnected elements demonstrate the integrated strategy necessary for successfully implementing AI within photovoltaic technology. The regulatory environment delineates current standards established by organizations such as the FDA and EMA, whereas data protection and cybersecurity underscore vital considerations including GDPR adherence and sophisticated encryption technologies. Ethical governance frameworks tackle significant challenges such as algorithmic fairness and system transparency. Furthermore, regulatory development reflects the evolving character of this domain, highlighting the importance of flexible approval mechanisms and coordinated international data standardization initiatives. This holistic framework illustrates the delicate equilibrium required between harnessing AI's capabilities in pharmaceutical safety surveillance and maintaining compliance with rigorous regulatory and ethical requirements.

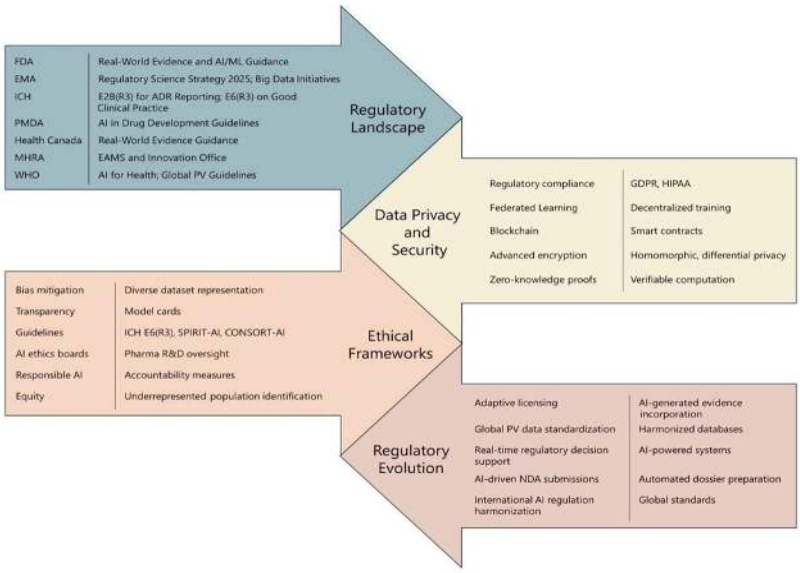


Figure 2: Regulatory and Ethical Considerations for AI-enhanced PV, Highlighting Current Guidelines and Expected Course of Evolution.

Artificial intelligence and machine learning technologies are fundamentally transforming the entire value chain within pharmacovigilance, spanning from initial data collection through information extraction, comprehensive analytics, actionable insights, and regulatory reporting to strategic intelligence generation. AI facilitates the identification, classification, and prioritization of pertinent information from multiple sources, including electronic health records, scientific literature, and digital platforms, thereby streamlining

and strengthening the data collection framework. AI-driven algorithms uncover concealed safety signals embedded within extensive datasets, empowering pharmacovigilance specialists to detect and evaluate emerging risks with greater speed and precision compared to conventional approaches. Furthermore, AI continuously refines its models as additional data becomes accessible, thereby enhancing drug safety detection and analytical capabilities. AI facilitates collaborative human-centered decision-making by delivering data-driven insights and forecasts derived from underlying patterns and trends, while simultaneously ensuring adherence to dynamic regulatory standards and documentation obligations through process automation and continuous surveillance of data integrity and sufficiency. Notwithstanding rapid technological advancement, significant obstacles persist. The efficacy of AI models remains contingent upon the availability of high-quality and sufficient data, particularly within resource-constrained environments. Given the diversity inherent in data sources, sophisticated AI models capable of synthesizing multiple data formats while guaranteeing precise and dependable results are essential. The integration of AI systems within the pharmacovigilance community continues to hinge on successfully resolving these core obstacles.

7. Ethical, Legal and Regulatory Considerations

Pharmacovigilance AI systems process sensitive health information and can shape critical safety determinations. Responsible deployment demands robust privacy safeguards, equitable treatment, clear communication, accountability mechanisms and meaningful human involvement. Comprehensive data governance frameworks must define access permissions, processing procedures, retention periods and mechanisms for validating model outputs. Regulatory frameworks continue to develop. The EMA's reflection paper on artificial intelligence in the medicinal-product life cycle establishes foundational principles for AI and machine learning applications throughout development and post-market surveillance phases. The EMA and FDA have initiated coordinated efforts addressing AI applications in pharmacovigilance, underscoring the importance of harmonized regulatory approaches. The

WHO recognizes artificial intelligence as a valuable tool for enhancing pharmacovigilance systems while acknowledging associated risks and governance requirements. The WHO's 2026 global smart pharmacovigilance strategy reinforces the importance of strengthening safety-monitoring infrastructure and promoting responsible deployment of digital technologies and real-world evidence sources.[30] Artificial intelligence should be integrated into established quality management systems rather than implemented as a standalone solution. Essential components include validation protocols, version management, comprehensive audit documentation, structured change procedures, systematic incident tracking and continuous performance assessment for applications with significant impact.[31]

7.1 Human-in-the-Loop Pharmacovigilance

Human-in-the-loop frameworks position qualified pharmacovigilance specialists at critical decision points throughout the workflow. The AI system may perform identification, information extraction, prioritization or synthesis functions, while qualified professionals evaluate the underlying evidence and render final clinical judgments. This methodology acknowledges that pharmaceutical safety assessment demands nuanced contextual analysis and cannot be adequately addressed through statistical correlation alone. An effective implementation strategy involves allocating more intensive human review to higher-risk activities. For instance, automated document screening may accommodate a higher false-positive threshold, whereas systems designed to filter out critical adverse events demand substantially more rigorous validation standards.

8. LIMITATIONS AND CHALLENGES

A major lesson from the literature is that technological novelty should not substitute for methodological quality. A 2022 scoping review found that many pharmacovigilance ML studies did not fully meet recommended good-practice characteristics. Therefore, reproducibility, validation and transparent reporting should be treated as core requirements.

Table 03: Limitations and Challenges

Challenge	Why it matters	Possible mitigation
Data quality	Incomplete or inconsistent data can produce unreliable outputs.	Data validation, cleaning and quality monitoring.
False positives	Can increase workload and distract reviewers.	Threshold optimization and human review.
False negatives	May miss clinically important safety information.	Safety-oriented validation and conservative thresholds.
Bias	Training data may not represent all populations.	Diverse datasets and subgroup evaluation.
Model drift	Data and clinical practice change over time.	Continuous performance monitoring.
Privacy	Health data are sensitive.	De-identification, access controls and governance.
External validation	Good performance on one dataset may not generalize.	Independent, temporal and multicenter validation.
Explainability	Complex models may be difficult to interpret.	Training, oversight and clear accountability.
Generative errors	LLMs may produce unsupported information	Grounding, source checking and restricted use.

9. FUTURE PROSPECTIVE

Upcoming priorities should emphasize enhancing dataset quality and consistency, progressing NLP methodologies to strengthen clinical narrative comprehension, and constructing explainable AI frameworks. Regulatory structures must adapt to facilitate AI integration within pharmacovigilance, guaranteeing the establishment of standardized protocols for AI adoption and the development of extensive, open-

access datasets for training purposes. Moving past simple correlation-based methodologies through the incorporation of causal inference approaches will facilitate deeper comprehension of drug-adverse drug reaction relationships. The convergence of artificial intelligence and large-scale data analytics in pharmacovigilance holds transformative potential for drug safety surveillance, mitigating numerous obstacles stemming from escalating data volume and the imperative for instantaneous monitoring. As these innovations mature, they are positioned to enhance pharmacovigilance effectiveness, dependability, and scope, ultimately strengthening patient protection. Although human validation of AI-generated insights remains essential, continuous advancements in AI system dependability will diminish reliance on human-driven processes and broaden AI implementation across pharmacovigilance operations.

The evolution of AI-driven pharmacovigilance is expected to incorporate integrated systems that merge organized clinical datasets, unstructured narrative text, diagnostic results, and additional real-world information sources. Enhanced natural language processing capabilities may facilitate cross-language surveillance and retrieval of insights from intricate clinical narratives. Distributed learning approaches could potentially enable organizations to develop algorithms using decentralized information repositories while preserving individual patient confidentiality, though implementation obstacles and regulatory considerations persist.

Explainable artificial intelligence represents a critical priority. Pharmacovigilance specialists require transparency into the reasoning behind case prioritization and signal identification. Explanations must demonstrate clinical relevance rather than technical complexity alone. Rigorous prospective assessment, uniform evaluation criteria, and independent verification are equally essential. Generative AI technologies may progressively facilitate knowledge synthesis and record management. Nevertheless, current research supports controlled, purpose-built implementations rather than fully autonomous pharmacovigilance operations. Emerging studies examining large language models for adverse event identification have consistently documented concerns including model fabrication, insufficient precision, parameter dependency, restricted training data, and inadequate external verification.[33]

Cross-border cooperation will prove vital given that pharmaceutical products and safety concerns transcend geographic limits. Standardized nomenclature, compatible data frameworks, consistent validation protocols, and unified governance structures could enhance system compatibility and effectiveness. The World Health Organization and regulatory bodies are actively promoting technology-enabled and analytics-focused strategies to advance pharmacovigilance capabilities. [32]

DISCUSSION

Artificial intelligence and machine learning present significant opportunities to transform pharmacovigilance from a primarily reactive approach to information management into a more comprehensive, evidence-based surveillance system. Their most immediate benefit lies in alleviating the manual workload associated with routine, repetitive activities and enabling specialists to efficiently locate pertinent information across extensive datasets. These technologies can additionally facilitate integrated surveillance approaches, where potential safety signals are cross-referenced and validated using multiple data sources including spontaneous adverse event reports, electronic health records, published literature, and other relevant information streams.

However, AI should not be assumed to inherently enhance pharmacovigilance outcomes. Machine learning algorithms can perpetuate existing biases present in training data, overlook uncommon adverse events, produce spurious safety signals, and demonstrate reduced effectiveness when deployed in populations or clinical settings that differ from their training environments. Standard performance indicators such as accuracy or F1 scores do not comprehensively capture clinical applicability. It is essential to evaluate the potential impact of prediction errors, the frequency of the monitored event, and the operational burden created by false positive alerts. The scientific literature in this field is expanding substantially. Current systematic reviews demonstrate growing adoption of deep

learning and natural language processing techniques; however, they also highlight gaps in external validation and variability in research methodology. Consequently, upcoming research should document dataset specifications, precise outcome measurements, model design details, reference comparison methods, validation approaches, performance across different patient subgroups, and clinically meaningful error assessment.[34] [35]

For pharmacy students and emerging pharmacovigilance specialists, competency in artificial intelligence is becoming increasingly important. While pharmacists do not require expertise equivalent to software developers, familiarity with data integrity principles, fundamental machine learning concepts, algorithm constraints, adverse event signal identification, and quality assurance oversight can strengthen interdisciplinary partnerships with computational specialists and regulatory authorities.

CONCLUSION

Artificial intelligence and machine learning are increasingly recognized as valuable supportive technologies within pharmacovigilance systems. These technologies facilitate adverse drug reaction identification, data extraction from clinical records, case prioritization, ongoing literature monitoring, signal assessment and ranking, predictive risk modeling, and examination of real-world evidence. Natural language processing and transformer-based architectures prove especially effective for extracting pharmaceutical-adverse event correlations from narrative text sources, whereas machine learning algorithms excel at recognizing trends within organized clinical datasets. The effectiveness of artificial intelligence extends beyond mere algorithmic accuracy. Successful deployment demands superior data quality, rigorous validation protocols, operational transparency, safeguarding of confidential information, persistent system evaluation, and competent professional supervision. Existing research indicates that artificial intelligence functions optimally as a complementary instrument to enhance pharmacovigilance specialists rather than as an independent substitute for their expertise.

As regulatory structures and global standards advance and strengthen, thoughtful incorporation of artificial intelligence may enable pharmacovigilance operations to handle expanding quantities of safety data while maintaining the fundamental objective of patient protection. Subsequent investigations should emphasize independent validation across different populations, model interpretability, uniform evaluation criteria, representative information sources, and demonstrable enhancements in actual clinical safety performance. As artificial intelligence progresses from developmental phases toward mainstream integration in pharmacovigilance, its successful incorporation into daily operations depends on resolving critical obstacles surrounding data integrity, explainability, and the reduction of algorithmic prejudice. A comprehensive strategy encompassing regulatory compliance, operational clarity, and continuous assessment is essential to guarantee that AI applications in pharmacovigilance not only meet performance standards but also remain responsive to the dynamic nature of pharmaceutical safety monitoring. Resolving these issues will establish the foundation for increased confidence and dependable implementation of AI technologies within the pharmacovigilance sector.

REFERENCES

1. US Food and Drug Administration. *Good pharmacovigilance practices and pharmacoepidemiologic assessment*. Rockville, MD: US Department of Health and Human Services, 2005.
2. Price J. Drug–drug interactions: a pharmacovigilance road less traveled. *Clin Ther* 2023; 45 (2): 94–98. Crossref. PubMed. Web of Science.
3. World Health Organization. (n.d.). Vigilance: Health products regulation and prequalification learning catalogue. World Health Organization.
4. Hauben M, Hartford CG. Artificial intelligence in pharmacovigilance: scoping points to consider. *Clin Ther* 2021; 43 (2): 372–379. Crossref. PubMed. Web of Science.
5. Hauben M. Artificial intelligence and data mining for the pharmacovigilance of drug–drug interactions. *Clin Ther*2023; 45 (2): 117–133. Crossref. PubMed. Web of Science.

6. European Medicines Agency. (2024). Reflection paper on the use of artificial intelligence (AI) in the medicinal product lifecycle. EMA/CHMP/CVMP/83833/2023.
7. Li, Y., Tao, W., Li, Z., Sun, Z., Li, F., Fenton, S., Xu, H., & Tao, C. (2024). Artificial intelligence-powered pharmacovigilance: A review of machine and deep learning in clinical text-based adverse drug event detection for benchmark datasets. Journal of Biomedical Informatics, 152, 104621.
8. Adverse drug reaction signal detection methods in spontaneous reporting system: A systematic review. (2024). Pharmacoepidemiology and Drug Safety.
9. Bate A, Evans SJW. Quantitative signal detection using spontaneous ADR reporting. *Pharmacoepidemiol Drug Saf* 2009; 18 (6): 427–436. Crossref. PubMed. Web of Science.
10. Harpaz R, DuMouchel W, Shah NH, et al. Novel data-mining methodologies for adverse drug event discovery and analysis. *Clin Pharmacol Ther* 2012; 91 (6): 1010–1021. Crossref. PubMed. Web of Science.
11. DuMouchel W. Bayesian data mining in large frequency tables, with an application to the FDA spontaneous reporting system. *Am Stat* 1999; 53 (3): 177–190. Crossref. Web of Science.
12. Nikfarjam A, Sarker A, O'connor K, et al. Pharmacovigilance from social media: mining adverse drug reaction mentions using sequence labeling with word embedding cluster features. *J Am Med Inf Assoc* 2015; 22 (3): 671–681. Crossref. PubMed. Web of Science.
13. Choudhury O, Park Y, Salonidis T, et al. Predicting adverse drug reactions on distributed health data using federated learning. In: AMIA annual symposium proceedings, Online, 2019, p. 313. American Medical Informatics Association.
14. Enthoven D, Al-Ars Z. An overview of federated deep learning privacy attacks and defensive strategies. In: Ur Rehman MH, Gaber MM (eds) *Federated Learning Systems*. 1st ed. Switzerland: Springer Cham, 2021, pp. 173–196. Crossref.
15. Rieke N, Hancox J, Li W, et al. The future of digital health with federated learning. *NPJ Digit Med* 2020; 3 (1): 1–7. Crossref. PubMed. Web of Science.
16. Harpaz R, DuMouchel W, LePendu P, et al. Performance of pharmacovigilance signal-detection algorithms for the FDA adverse event reporting system. *Clin Pharmacol Ther* 2013; 93 (6): 539–546. Crossref. PubMed. Web of Science.
17. Yang X, Bian J, Gong Y, et al. MADEx: a system for detecting medications, adverse drug events, and their relations from clinical notes. *Drug Saf* 2019; 42: 123–133. Crossref. PubMed. Web of Science.
18. US Food and Drug Administration. Sentinel system: five-year strategy 2019–2023. Silver Spring, MD: US Food and Drug Administration (FDA), 2019.
19. Salas M, Petracek J, Yalamanchili P, et al. The use of artificial intelligence in pharmacovigilance: a systematic review of the literature. *Pharmaceut Med* 2022; 36 (5): 295–306. Crossref. PubMed. Web of Science.
20. Van de Burgt BWM, Wasylewicz ATM, Dullemond B, et al. Development of a text mining algorithm for identifying adverse drug reactions in electronic health records. *JAMIA Open* 2024; 7 (3): ooae070. Crossref. PubMed. Web of Science.
21. Li, Y., Tao, W., Li, Z., Sun, Z., Li, F., Fenton, S., Xu, H., & Tao, C. (2024). Artificial intelligence-powered pharmacovigilance: A review of machine and deep learning in clinical text-based adverse drug event detection for benchmark datasets. Journal of Biomedical Informatics, 152, 104621.
22. Artificial intelligence (AI) in pharmacovigilance: A systematic review on predicting adverse drug reactions (ADR) in hospitalized patients. (2025). PubMed-indexed systematic review.
23. Systematic review of AI-based models in pharmacoepidemiology for adverse drug event prediction and detection. (2026). PubMed-indexed systematic review.
24. European Medicines Agency. (2026). Artificial intelligence: Large language model guiding

principles and principles for AI in medicine development. EMA.

25. Rieke N, Hancox J, Li W, et al. The future of digital health with federated learning. *NPJ Digit Med* 2020; 3 (1): 1–7. Crossref. PubMed. Web of Science.
26. Wang X, Xu X, Tong W, et al. InferBERT: a transformer-based causal inference framework for enhancing pharmacovigilance. *Front Artif Intell* 2021; 4: 659622. Crossref. PubMed. Web of Science.
27. Rodrigues PP, Ferreira-Santos D, Silva A, et al. Causality assessment of adverse drug reaction reports using an expert-defined Bayesian network. *Artif Intell Med* 2018; 91: 12–22. Crossref. PubMed. Web of Science.
28. Ward IR, Wang L, Lu J, et al. Explainable artificial intelligence for pharmacovigilance: what features are important when predicting adverse outcomes? *Comput Methods Programs Biomed* 2021; 212: 106415. Crossref. PubMed.
29. Salih AM, Raisi-Estabragh Z, Galazzo IB, et al. A perspective on explainable artificial intelligence methods: SHAP and LIME. *Adv Intell Syst* 2025; 7 (1): 2400304. Crossref. Web of Science.
30. European Medicines Agency. (2026). Artificial intelligence (AI) in pharmacovigilance cluster. EMA.
31. World Health Organization. (2025). The global smart pharmacovigilance strategy. WHO
32. World Health Organization. (2025). The global smart pharmacovigilance strategy. WHO.
33. Large language models in adverse drug reaction detection and pharmacovigilance: A systematic review of current applications, challenges, and future directions. (2026). PubMed-indexed systematic review.
34. Artificial intelligence (AI) in pharmacovigilance: A systematic review on predicting adverse drug reactions (ADR) in hospitalized patients. (2025). PubMed-indexed systematic review.
35. Artificial intelligence based on machine learning in pharmacovigilance: A scoping review. (2022). PubMed-indexed scoping review.

HOW TO CITE: Aditya Ghavat, Yash Ghodke, Yash Mulay, Raj Dohale, Shreyash Pardeshi, Artificial Intelligence and Machine Learning in Pharmacovigilance: Transforming Adverse Drug Reaction Detection and Signal Management., Int. J. of Pharm. Sci., 2026, Vol 4, Issue 11, 1625-1641. https://doi.org/10.5281/zenodo.24143171