The Role of Artificial Intelligence in Pharmacovigilance: A Narrative Review of Methods, Performance and Limitations

Sourabh Bagoria[1], Rohit Namdeo Ingole[2], Bommareddy Devi Aruna Jyothi[3], Aamir Patel[4], Daizy Tak[5], Priyanka Rajendra Wadile[6], Cheruku Akshith[7], Sumayya Sanober[8]

[1]*MS(Pharm) Regulatory Toxicology, Department of Pharmacology and Toxicology, NIPER Mohali*

[2] *MSc Pharmaceutical Business and Technology, Griffith College Dublin, Ireland*

[3]*Pharm.D, Koneru Lakshmaiah Education Foundation, Vaddeswaram, Vijayawada*

[4] *M.Pharm, Department of Ph. Chemistry, Principal K.M.Kundnani College of Pharmacy, Mumbai*

[5]*Research Scholar, Department of Pharmacy Practice, Geetanjali Institute of Pharmacy, Geetanjali University, Udaipur, Rajasthan*

[6] *Lecturer, Y.N.P College of Pharmacy, Asangaon, Dahanu, Maharashtra*

[7]*Pharm.D, Bharat School of Pharmacy Mangalpally, Ibrahimpatnam, Hyderabad*

[8]*M.Pharm, Anwarul Uloom College of Pharmacy (JNTU), Hyderabad*

Abstract- Pharmacovigilance represents a cornerstone of pharmaceutical regulation and clinical governance, serving as a critical mechanism for the continuous evaluation of the safety profile of medicinal products following their authorization for public use. While pre-marketing clinical trials provide essential evidence regarding efficacy and short-term safety, their methodological constraints—such as limited sample sizes, controlled study environments, restricted patient heterogeneity, and relatively short durations of exposure—inevitably hinder the identification of rare, delayed, cumulative, or context-dependent adverse drug reactions that may only emerge under real-world conditions. Consequently, post-marketing pharmacovigilance systems are tasked with the complex responsibility of monitoring drug safety across diverse populations, therapeutic indications, and healthcare settings over extended periods of time. However, traditional pharmacovigilance frameworks, which rely predominantly on spontaneous reporting systems, manual case processing, and rule-based statistical signal detection methodologies, are increasingly challenged by the exponential growth in the volume, velocity, and heterogeneity of safety-related data generated in modern healthcare ecosystems. These limitations have resulted in delayed signal detection, substantial operational burden, variability in case quality assessment, and reduced capacity to extract meaningful insights from unstructured or non-traditional data sources such as electronic health records, biomedical literature, and patient-generated digital content.

In this context, artificial intelligence has emerged as a potentially transformative paradigm capable of addressing many of the structural inefficiencies inherent in conventional pharmacovigilance practices. Artificial intelligence-driven methodologies, encompassing machine learning, deep learning, and natural language processing techniques, enable automated ingestion and analysis of large-scale, heterogeneous datasets, facilitate advanced pattern recognition, and support adaptive learning from continuously evolving data streams. Within pharmacovigilance operations, these technologies have been applied to a wide range of functions, including automated adverse event case intake, medical coding and normalization using standardized terminologies such as MedDRA, signal detection and prioritization, literature surveillance, and benefit–risk assessment. Early evidence suggests that artificial intelligence-based systems can enhance processing efficiency, improve consistency in case assessment, and increase sensitivity for early safety signal identification when compared with traditional disproportionality analyses. Nevertheless, the deployment of artificial intelligence in pharmacovigilance is accompanied by significant methodological, regulatory, and ethical challenges, including data quality and representativeness limitations, algorithmic bias, lack of transparency and interpretability in complex models, constrained

generalizability across populations and therapeutic areas, and the continued necessity for expert human oversight to ensure clinical relevance and regulatory compliance.

This narrative review critically examines the role of artificial intelligence in pharmacovigilance by synthesizing current methodological approaches, evaluating reported performance outcomes, and identifying key limitations that influence real-world implementation. Furthermore, it explores emerging regulatory perspectives and future directions aimed at fostering responsible integration of artificial intelligence into pharmacovigilance systems. By providing a comprehensive and clinically grounded assessment, this review seeks to contribute to the evolving discourse on how artificial intelligence can be strategically leveraged to augment, rather than replace, traditional pharmacovigilance practices, ultimately supporting more timely, robust, and patient-centered drug safety surveillance.

Keywords- Artificial Intelligence; Pharmacovigilance; Adverse Drug Reactions; Drug Safety Surveillance; Machine Learning; Deep Learning; Natural Language Processing; Signal Detection; Real-World Evidence; Post-Marketing Surveillance; Regulatory Science; Benefit–Risk Assessment

I. INTRODUCTION

Pharmacovigilance represents a critical scientific and regulatory discipline within pharmaceutical sciences, clinical medicine, and public health, dedicated to the continuous monitoring and evaluation of the safety of medicinal products following their authorization for clinical use. The fundamental objective of pharmacovigilance is to ensure that the therapeutic benefits of medicines outweigh their potential risks throughout the entire product lifecycle, from early post-marketing exposure to widespread and long-term real-world use. Although pre-authorization clinical trials provide essential evidence regarding efficacy and short-term safety, their inherent methodological constraints—including limited sample sizes, selective participant recruitment, controlled experimental conditions, and relatively brief exposure periods—significantly restrict their capacity to detect rare, delayed, cumulative, or population-specific adverse drug reactions. Consequently, many clinically meaningful safety signals only become apparent after a medicinal product has been introduced into routine clinical practice, where it is prescribed to heterogeneous patient populations with varying comorbidities, concomitant medications, genetic backgrounds, and healthcare access patterns. In this context, pharmacovigilance serves as an indispensable mechanism for safeguarding patient safety, informing regulatory decision-making, and supporting rational pharmacotherapy in real-world settings.

Over the past two decades, the scope and complexity of pharmacovigilance activities have expanded substantially, driven by rapid advances in pharmaceutical innovation, accelerated regulatory approval pathways, and the globalization of drug development and distribution. The introduction of biologics, biosimilars, advanced therapy medicinal products, and precision medicines has introduced novel safety considerations that extend beyond traditional small-molecule pharmacology, including immunogenicity, long-term oncogenic risk, and complex pharmacodynamic interactions. Simultaneously, demographic shifts such as population aging, increased prevalence of chronic diseases, and widespread polypharmacy have amplified the risk of adverse drug reactions and drug–drug interactions in routine clinical practice. These developments have significantly increased the volume of safety data generated across healthcare systems, placing unprecedented demands on pharmacovigilance infrastructures that were originally designed for lower data throughput and more homogeneous information sources.

Traditional pharmacovigilance systems remain heavily reliant on spontaneous reporting mechanisms, wherein individual case safety reports are voluntarily submitted by healthcare professionals, patients, and pharmaceutical manufacturers to national and international databases. While spontaneous reporting systems have historically played a pivotal role in the identification of serious and unexpected adverse drug reactions, they are inherently limited by well-documented challenges such as underreporting, reporting delays, variable data quality, incomplete clinical information, and reporting biases influenced by media attention, regulatory actions, or litigation concerns. Furthermore, conventional signal detection methodologies employed within these systems—such as disproportionality analyses based on reporting odds ratios or proportional reporting ratios—are fundamentally retrospective and static in nature,

relying on aggregated counts and predefined statistical thresholds that may fail to capture subtle, emerging, or multifactorial safety signals. As a result, clinically relevant risks may remain undetected until substantial patient exposure has already occurred, potentially compromising timely risk mitigation.

The digital transformation of healthcare has further complicated the pharmacovigilance landscape by generating vast quantities of real-world data from diverse and heterogeneous sources. Electronic health records, administrative claims databases, disease registries, clinical narratives, biomedical literature, and patient-generated content from digital health platforms and social media collectively represent an unprecedented reservoir of safety-relevant information. However, the majority of these data are unstructured or semi-structured, rendering them difficult to process using conventional pharmacovigilance tools and workflows. Manual review of such data is resource-intensive, time-consuming, and susceptible to inter-observer variability, thereby limiting scalability and consistency. These challenges have underscored the growing inadequacy of traditional, rule-based pharmacovigilance approaches in the face of modern data complexity and have catalyzed interest in advanced computational solutions capable of augmenting human expertise.

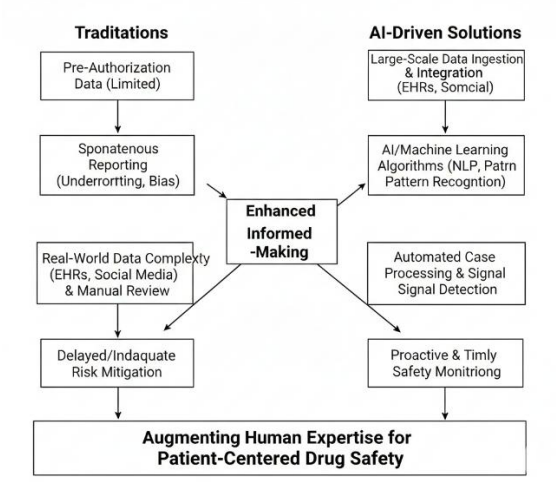

Artificial intelligence has emerged as a promising technological paradigm with the potential to address many of the structural and operational limitations of conventional pharmacovigilance systems. By leveraging advanced computational techniques such as machine learning, deep learning, and natural language processing, artificial intelligence enables automated extraction, integration, and analysis of large-scale, heterogeneous datasets, facilitating the identification of complex, non-linear relationships that may not be readily apparent through traditional statistical analyses. In pharmacovigilance, AI-driven methodologies have been increasingly explored for applications ranging from automated adverse event case intake and medical coding to signal detection, prioritization, and benefit–risk assessment. These technologies offer the potential to enhance efficiency, improve consistency, and support more timely identification of emerging safety concerns, thereby contributing to a more proactive and data-driven pharmacovigilance paradigm.

Despite its considerable promise, the adoption of artificial intelligence in pharmacovigilance is accompanied by substantial methodological, regulatory, and ethical challenges that must be carefully addressed to ensure responsible and effective implementation. AI models are highly dependent on the quality, completeness, and representativeness of training data, and biases inherent in source datasets may be inadvertently propagated or amplified by automated systems. Moreover, the opacity of complex models, particularly deep learning architectures, raises concerns regarding transparency, interpretability, and regulatory acceptability in a domain where decision-making must be scientifically justified and auditable. Regulatory authorities have begun to acknowledge the potential role of artificial intelligence in pharmacovigilance, yet comprehensive frameworks governing validation, governance, and lifecycle management of AI-based systems remain under development. In this context, a critical and balanced evaluation of the methods, performance, and limitations of artificial intelligence in pharmacovigilance is essential to inform future research, regulatory policy, and practical implementation.

This narrative review aims to provide a comprehensive and clinically grounded examination of the role of artificial intelligence in pharmacovigilance, focusing on the methodological approaches employed, their reported performance outcomes, and the key limitations that influence real-world applicability. By synthesizing current evidence and regulatory perspectives, this review seeks to elucidate how

artificial intelligence can be strategically integrated into pharmacovigilance systems to augment, rather than replace, traditional safety surveillance practices, ultimately supporting more robust, timely, and patient-centered drug safety monitoring.

II. METHODOLOGY

The methodological approach adopted for this narrative review was designed to critically interrogate the role of artificial intelligence in pharmacovigilance by integrating methodological, computational, and regulatory perspectives into a unified analytical framework. Rather than limiting the scope to descriptive summaries of existing applications, this review employed a structured conceptual methodology aimed at examining how artificial intelligence techniques are operationalized within pharmacovigilance workflows, how their performance is evaluated against conventional safety surveillance methods, and where their limitations emerge in real-world regulatory environments. The methodological focus was therefore placed on three interconnected analytical dimensions: the classification of artificial intelligence methodologies according to their functional role in pharmacovigilance, the assessment of reported performance outcomes using pharmacovigilance-relevant metrics, and the identification of structural, regulatory, and ethical constraints influencing implementation.

Study Design and Review Framework

This work was conducted as a narrative review with a structured analytical framework tailored to the multidisciplinary nature of artificial intelligence–enabled pharmacovigilance. The narrative review design was selected to allow critical synthesis of heterogeneous evidence encompassing computational methodologies, clinical safety science, and regulatory considerations, which are not readily amenable to conventional systematic review paradigms. Rather than aggregating effect sizes or performing meta-analytical comparisons, the methodological focus was placed on conceptual integration, methodological appraisal, and contextual interpretation of artificial intelligence applications within real-world pharmacovigilance systems. The review framework was designed to examine artificial intelligence as both a technological intervention and a decision-support mechanism within the broader pharmacovigilance lifecycle, spanning case processing, signal detection, and post-marketing risk assessment.

Literature Identification and Source Selection

The identification of relevant literature was guided by methodological relevance, scientific rigor, and applicability to pharmacovigilance practice. Peer-reviewed journal articles, regulatory white papers, conference proceedings, and authoritative technical reports describing artificial intelligence applications in drug safety surveillance were considered. Emphasis was placed on studies that explicitly detailed the implementation of artificial intelligence models within pharmacovigilance workflows, including data preprocessing steps, model training strategies, validation approaches, and operational deployment contexts. Articles focusing solely on theoretical machine learning development without direct pharmacovigilance application were excluded, as were studies lacking sufficient methodological transparency. This selective approach ensured that the reviewed literature provided meaningful insights into the practical and clinical implications of artificial intelligence adoption in pharmacovigilance.

Classification of Artificial Intelligence Methodologies

Artificial intelligence methodologies identified in the reviewed literature were systematically classified according to their underlying computational paradigm and functional role in pharmacovigilance. Models were broadly grouped into supervised machine learning, unsupervised learning, deep learning, and natural language processing systems. Within each category, further distinctions were made based on model architecture, learning strategy, and intended pharmacovigilance task, such as adverse event identification, medical coding, signal detection, or risk stratification. This classification enabled structured comparison across studies and facilitated identification of methodological trends, performance patterns, and areas of methodological convergence or divergence. Particular attention was paid to hybrid models that integrated multiple artificial intelligence techniques, reflecting the evolving complexity of pharmacovigilance analytics.

Evaluation of Model Performance and Validation Strategies

Assessment of artificial intelligence performance was conducted through critical examination of reported

evaluation metrics and validation methodologies. Priority was given to pharmacovigilance-relevant performance indicators, including sensitivity, specificity, precision, recall, and timeliness of signal detection, rather than generic machine learning accuracy measures alone. The review evaluated how models were trained and tested, including the use of internal validation, external datasets, and temporal validation to assess generalizability. Studies comparing artificial intelligence-driven approaches with traditional pharmacovigilance methods, such as disproportionality analyses or manual case review, were analyzed to determine incremental performance benefits and operational efficiencies. Limitations related to overfitting, dataset bias, and lack of real-world validation were explicitly examined.

Data Sources and Preprocessing Considerations

The methodological assessment also encompassed detailed analysis of data sources used to train and evaluate artificial intelligence models. These included spontaneous reporting systems, electronic health records, claims databases, scientific literature, and social media platforms. The review examined how data heterogeneity, missingness, reporting bias, and terminological inconsistency were addressed through preprocessing techniques such as normalization, de-duplication, and medical terminology mapping. Particular emphasis was placed on the impact of data quality on model performance and the challenges associated with using non-traditional data sources for regulatory pharmacovigilance purposes.

Regulatory, Ethical, and Governance Considerations

Given the regulated nature of pharmacovigilance, methodological evaluation extended beyond technical performance to include regulatory and governance dimensions. The review assessed how artificial intelligence methodologies were aligned with regulatory expectations for transparency, traceability, and human oversight. Studies discussing explainable artificial intelligence, auditability, and model lifecycle management were analyzed to evaluate their potential to support regulatory acceptance. Ethical considerations, including bias amplification, accountability, and the risk of automation-driven decision errors, were examined as integral methodological constraints influencing real-world implementation.

Identification of Methodological Limitations

Methodological limitations were identified through cross-study synthesis of reported challenges and failure modes. These included limitations related to data representativeness, model interpretability, scalability across therapeutic areas, and dependence on expert-labeled datasets. The review also examined systemic limitations inherent to pharmacovigilance data, such as underreporting and delayed reporting, which constrain the theoretical performance ceiling of artificial intelligence models. By explicitly mapping these limitations, the methodology provides a critical lens through which artificial intelligence adoption in pharmacovigilance can be realistically evaluated. artificial intelligence, human-in-the-loop models, and hybrid decision-making frameworks were analyzed to evaluate their feasibility as governance mechanisms within pharmacovigilance systems.

Finally, methodological limitations were systematically identified by synthesizing reported challenges across studies and mapping them to pharmacovigilance-specific requirements. These included limitations related to scalability, transferability across regions and therapeutic areas, dependence on high-quality labeled data, and the persistent need for expert clinical judgment. By integrating these methodological considerations, this narrative review provides not only a descriptive overview of artificial intelligence applications in pharmacovigilance but also a critical evaluation of the conditions under which such technologies can meaningfully enhance drug safety surveillance.

III. ARTIFICIAL INTELLIGENCE METHODOLOGIES APPLIED IN PHARMACOVIGILANCE

The integration of Artificial Intelligence (AI) into pharmacovigilance has fundamentally transformed the way adverse drug reactions (ADRs), safety signals, and post-marketing risk profiles are identified, evaluated, and managed. Traditional pharmacovigilance systems, which largely rely on spontaneous reporting, manual case review, and rule-based statistical approaches, are increasingly inadequate to cope with the exponential growth of real-world data generated from electronic health records (EHRs), social media platforms, biomedical literature, clinical trial repositories, and global safety

databases. AI methodologies—spanning machine learning, deep learning, natural language processing, and hybrid cognitive systems—offer scalable, adaptive, and data-driven solutions capable of detecting subtle, nonlinear, and previously unrecognized safety patterns. This section critically elaborates the core AI methodologies applied in pharmacovigilance, emphasizing their operational principles, real-world applications, strengths, and methodological constraints.

Machine Learning Techniques in Pharmacovigilance

Machine learning (ML) represents the foundational pillar of AI-driven pharmacovigilance systems, enabling automated pattern recognition and predictive analytics without explicit rule programming. ML algorithms learn statistical associations between input variables (e.g., patient demographics, drug exposure, clinical outcomes) and target outputs (e.g., ADR occurrence, seriousness classification, causality likelihood) by iteratively optimizing performance metrics over large annotated datasets. In pharmacovigilance, ML techniques are extensively employed for case triaging, duplicate detection, signal detection, and risk stratification.

Supervised learning models, including logistic regression, support vector machines (SVMs), decision trees, random forests, and gradient boosting algorithms, are particularly prevalent in ADR classification tasks. These models require labeled datasets—typically curated from regulatory databases such as FAERS, EudraVigilance, or VigiBase—where ADRs are annotated by seriousness, expectedness, or causality category. By learning discriminative features from structured and semi-structured data, supervised ML models can prioritize high-risk individual case safety reports (ICSRs), thereby reducing manual workload and improving regulatory response times. However, their performance is heavily dependent on data quality, class balance, and annotation consistency, which remain persistent challenges in pharmacovigilance datasets.

Unsupervised learning approaches, such as clustering algorithms (k-means, hierarchical clustering, DBSCAN) and association rule mining, are primarily applied in exploratory signal detection scenarios. These techniques identify latent structures, co-occurrence patterns, or anomalous reporting behaviors without predefined labels. For example, clustering can reveal unexpected drug–event combinations or patient subpopulations exhibiting disproportionate risk profiles. While unsupervised methods are valuable for hypothesis generation, their interpretability and clinical validation often require expert adjudication, limiting their standalone regulatory applicability.

Deep Learning Architectures for Complex Safety Signal Detection

Deep learning (DL), a subset of machine learning inspired by artificial neural networks, has gained substantial traction in pharmacovigilance due to its ability to model high-dimensional, nonlinear relationships across heterogeneous data sources. Unlike traditional ML algorithms that rely on manual feature engineering, deep learning models automatically learn hierarchical feature representations directly from raw data, making them particularly suitable for complex safety surveillance tasks.

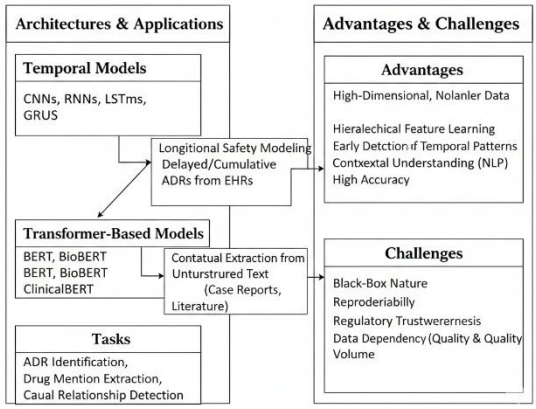

Convolutional neural networks (CNNs) and recurrent neural networks (RNNs), including long short-term memory (LSTM) and gated recurrent unit (GRU) architectures, are widely applied in temporal ADR pattern recognition and longitudinal safety modeling. In pharmacovigilance, these models can capture time-dependent relationships between drug exposure and adverse outcomes, enabling early detection of delayed or cumulative toxicities that may not be evident through conventional disproportionality analyses. For instance, LSTM-based models have demonstrated superior performance in predicting ADR onset using sequential EHR data, outperforming traditional regression-based approaches.

More recently, transformer-based architectures, such as BERT (Bidirectional Encoder Representations from Transformers) and its biomedical adaptations (BioBERT, ClinicalBERT), have revolutionized safety signal extraction from unstructured text sources. These models leverage self-attention mechanisms to understand contextual dependencies within clinical narratives, enabling precise identification of adverse events, drug mentions, and causal relationships in free-text case reports and literature abstracts. Despite their impressive accuracy, deep learning models are often criticized for their "black-box" nature, raising concerns regarding transparency, reproducibility, and regulatory trustworthiness—critical factors in pharmacovigilance decision-making.

Natural Language Processing for Adverse Event Extraction

Natural Language Processing (NLP) constitutes a cornerstone methodology in AI-enabled pharmacovigilance, addressing the predominance of unstructured textual data in safety reporting systems. A significant proportion of pharmacovigilance-relevant information resides in narrative case descriptions, clinical notes, discharge summaries, regulatory correspondence, and biomedical publications, which are not readily amenable to traditional statistical analysis.

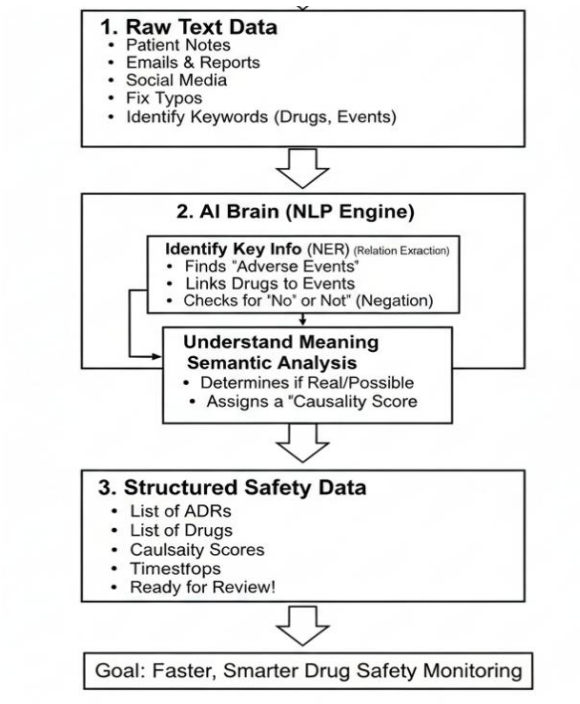

Rule-based NLP systems, employing predefined lexicons and pattern-matching techniques, were among the earliest approaches used to extract adverse events and drug entities. While these systems offer high precision in controlled environments, they lack scalability and adaptability to linguistic variability, spelling errors, colloquial expressions, and evolving medical terminology. Consequently, modern pharmacovigilance increasingly relies on machine learning–driven NLP pipelines that integrate named entity recognition (NER), relation extraction, sentiment analysis, and semantic role labeling.

Advanced NLP models enable automated identification of adverse events, normalization to standardized terminologies (e.g., MedDRA, SNOMED CT), and detection of drug–event causality cues within narrative text. For example, NLP algorithms can differentiate between actual ADRs and confounding mentions such as medical history, indications, or hypothetical scenarios. Despite these advancements, challenges persist in handling negation, ambiguity, and cross-lingual reporting, particularly in global pharmacovigilance systems spanning multiple regulatory jurisdictions.

Hybrid AI Systems and Knowledge-Based Approaches

Hybrid AI systems combine data-driven learning algorithms with rule-based reasoning and domain-specific knowledge graphs to enhance robustness, interpretability, and clinical relevance. In pharmacovigilance, these systems integrate structured medical ontologies, regulatory guidelines, and expert-curated rules with machine learning predictions, thereby aligning algorithmic outputs with established pharmacological and clinical principles.

Knowledge graphs, representing entities such as drugs, targets, pathways, adverse events, and patient characteristics as interconnected nodes, facilitate mechanistic signal interpretation and hypothesis validation. When coupled with AI algorithms, knowledge graphs enable contextualized safety assessments, supporting causal inference and risk-benefit evaluations. For instance, linking an observed ADR signal to known pharmacodynamic mechanisms or metabolic pathways can strengthen regulatory confidence and guide risk mitigation strategies.

Hybrid approaches also support explainable AI (XAI) initiatives, which are increasingly emphasized by regulatory agencies. By providing transparent rationale for safety alerts—such as feature importance scores or rule-based justifications—hybrid systems address one of the most critical limitations of purely data-driven models in pharmacovigilance.

Performance Evaluation and Validation of AI Models
The methodological rigor of AI applications in pharmacovigilance is contingent upon robust performance evaluation and external validation. Standard metrics such as sensitivity, specificity, precision, recall, F1-score, and area under the receiver operating characteristic curve (AUROC) are commonly used to assess model accuracy. However, pharmacovigilance-specific considerations, including class imbalance, rare event detection, and regulatory risk tolerance, necessitate customized evaluation frameworks.

External validation across independent datasets, temporal validation, and real-world deployment studies are essential to ensure generalizability and clinical reliability. Moreover, continuous performance monitoring is required to address concept drift arising from changes in reporting behavior, drug utilization patterns, and regulatory practices over time. Without systematic validation and governance frameworks, AI-driven pharmacovigilance systems risk generating spurious signals or missing clinically significant safety concerns.

IV. PERFORMANCE EVALUATION AND COMPARATIVE EFFECTIVENESS OF AI-BASED PHARMACOVIGILANCE SYSTEMS

The assessment of Artificial Intelligence–driven pharmacovigilance systems necessitates a rigorous, multidimensional evaluation framework that extends beyond conventional predictive accuracy metrics. Unlike traditional clinical prediction models, pharmacovigilance algorithms operate in a regulatory-sensitive environment where false negatives may delay critical safety actions and false positives may trigger unwarranted regulatory scrutiny, resource diversion, or public concern. Consequently, performance evaluation of AI models in pharmacovigilance must balance statistical robustness, clinical relevance, regulatory acceptability, and operational feasibility. This section critically examines the methodological approaches used to evaluate AI-based pharmacovigilance systems and compares their effectiveness with conventional safety surveillance methodologies.

From a quantitative standpoint, classical machine learning and deep learning models are typically evaluated using discrimination metrics such as sensitivity, specificity, positive predictive value (PPV), negative predictive value (NPV), accuracy, F1-score, and area under the receiver operating characteristic curve (AUROC). Sensitivity holds particular importance in pharmacovigilance, as the primary objective is early detection of potential safety signals rather than precise event prediction. High sensitivity ensures that emerging adverse drug reactions are not overlooked, even at the expense of reduced specificity. However, excessive false-positive alerts can overwhelm pharmacovigilance teams and erode confidence in automated systems, underscoring the need for carefully calibrated decision thresholds aligned with regulatory risk tolerance.

Beyond binary classification metrics, signal detection performance is frequently assessed using disproportionality-based benchmarks such as proportional reporting ratios (PRR), reporting odds ratios (ROR), and Bayesian confidence propagation neural network (BCPNN) outputs. AI-based approaches are compared against these traditional statistical methods to determine incremental value in detecting signals earlier, identifying rare or delayed adverse events, and uncovering complex drug–event interactions. Multiple retrospective studies have demonstrated that machine learning–enhanced signal detection models outperform classical disproportionality analyses in terms of timeliness and sensitivity, particularly when applied to high-dimensional datasets incorporating patient-level covariates, comorbidities, and concomitant medications.

Temporal validation constitutes another critical dimension of performance evaluation, as pharmacovigilance models must remain robust over extended post-marketing periods. AI systems trained on historical safety data may experience performance degradation due to concept drift—systematic changes

in prescribing patterns, reporting behavior, regulatory requirements, or population demographics. Longitudinal validation studies, which assess model performance across different time windows, are therefore essential to establish durability and real-world reliability. Continuous learning architectures and periodic model recalibration have been proposed as mitigation strategies, although these raise additional regulatory and governance complexities.

Comparative effectiveness analyses further reveal that deep learning models, particularly those integrating longitudinal EHR data and unstructured text via NLP pipelines, exhibit superior performance in detecting complex, multifactorial adverse events compared to rule-based or shallow learning systems. However, these gains are often accompanied by reduced interpretability, increased computational cost, and greater dependency on large, well-curated datasets. In contrast, hybrid models combining interpretable machine learning algorithms with expert-defined pharmacological rules demonstrate a more favorable balance between performance and explainability, making them more acceptable in regulatory pharmacovigilance workflows.

Importantly, qualitative evaluation metrics—such as clinical plausibility, regulatory usability, and expert agreement—are increasingly recognized as indispensable components of AI performance assessment in pharmacovigilance. Human-in-the-loop validation, wherein safety experts review AI-generated signals and provide feedback, remains the gold standard for confirming clinical relevance and mitigating algorithmic bias. Studies consistently indicate that AI systems function optimally as decision-support tools rather than autonomous signal generators, augmenting rather than replacing expert judgment.

Despite encouraging performance outcomes, the absence of standardized evaluation frameworks and reporting guidelines continues to hinder cross-study comparability and regulatory harmonization. Variability in dataset selection, labeling criteria, validation strategies, and performance metrics limits the generalizability of published results. Regulatory bodies and international pharmacovigilance organizations have therefore emphasized the urgent need for consensus-driven benchmarks and transparent reporting standards to ensure reliable integration of AI into global drug safety surveillance systems.

V. REGULATORY ACCEPTANCE, GOVERNANCE, AND IMPLEMENTATION CHALLENGES OF ARTIFICIAL INTELLIGENCE IN PHARMACOVIGILANCE

The regulatory acceptance of Artificial Intelligence–based systems in pharmacovigilance represents one of the most critical determinants of their real-world adoption and sustainability. While AI methodologies have demonstrated substantial potential to enhance signal detection efficiency, data integration, and analytical depth, their incorporation into regulated drug safety frameworks is constrained by stringent requirements related to transparency, accountability, validation, and patient safety. Regulatory agencies such as the U.S. Food and Drug Administration (FDA), European Medicines Agency (EMA), and the World Health Organization (WHO) recognize the transformative potential of AI but emphasize that pharmacovigilance remains a high-stakes, risk-sensitive domain where algorithmic outputs must be interpretable, reproducible, and scientifically justifiable.

A primary regulatory concern lies in the lack of explainability associated with complex AI models, particularly deep learning architectures. Pharmacovigilance decisions—such as label changes, risk management plan modifications, or market withdrawals—require a clear rationale that can be audited, defended, and communicated to stakeholders. Black-box models that generate safety signals without transparent reasoning pose significant challenges in regulatory review, as agencies must be able to trace how specific inputs contributed to a given output. Consequently, regulatory bodies increasingly advocate for explainable artificial intelligence (XAI) approaches that provide interpretable insights, such as feature importance rankings, decision pathways, or rule-based justifications, without compromising analytical performance.

Governance frameworks for AI-driven pharmacovigilance systems remain underdeveloped and heterogeneous across jurisdictions. Unlike

traditional pharmacovigilance tools, AI systems are dynamic by nature, often incorporating continuous learning mechanisms that adapt to new data streams. While such adaptability enhances performance, it also complicates regulatory oversight, as model behavior may change post-approval. This raises fundamental questions regarding model version control, change management, revalidation requirements, and documentation standards. Regulators have expressed concerns about "algorithmic drift," wherein gradual performance shifts may introduce bias or degrade safety detection accuracy if not properly monitored and controlled.

Data provenance and quality governance constitute additional regulatory challenges. Pharmacovigilance AI systems rely on diverse data sources, including spontaneous reporting systems, electronic health records, claims databases, social media, and scientific literature. These datasets vary widely in structure, completeness, reliability, and bias. Regulatory authorities require assurance that AI models are trained on representative, ethically sourced, and high-quality data, and that appropriate safeguards are in place to prevent propagation of systemic biases—such as underreporting in specific populations or geographic regions. Inadequate data governance can compromise both the validity of AI-generated safety signals and public trust in regulatory decision-making.

Another significant barrier to regulatory acceptance is the absence of harmonized international standards for AI validation in pharmacovigilance. Current regulatory guidance primarily addresses traditional statistical signal detection methods and does not adequately account for the complexity of modern AI systems. As a result, pharmaceutical companies and marketing authorization holders face uncertainty regarding acceptable validation strategies, performance thresholds, and documentation requirements. This regulatory ambiguity may discourage innovation or lead to conservative implementations that underutilize AI's full potential. Collaborative initiatives involving regulators, industry, and academia are therefore essential to establish consensus-driven frameworks for AI qualification, benchmarking, and lifecycle management.

Implementation challenges further extend beyond regulatory approval into operational integration within existing pharmacovigilance infrastructures. Many organizations lack the technical expertise, computational resources, or organizational readiness to deploy and maintain AI systems at scale. Integration with legacy safety databases, workflow redesign, staff training, and change management represent substantial logistical and financial investments. Moreover, resistance to automation among pharmacovigilance professionals—stemming from concerns about job displacement or overreliance on algorithms—may hinder adoption unless AI is clearly positioned as a decision-support tool rather than a replacement for expert judgment.

Despite these challenges, regulatory agencies have begun adopting a more proactive and adaptive stance toward AI in pharmacovigilance. Pilot programs, regulatory sandboxes, and public–private partnerships have been initiated to explore safe and controlled implementation pathways. The emphasis is gradually shifting toward a risk-based regulatory approach, wherein AI systems are evaluated based on their intended use, impact on patient safety, and level of autonomy. Under this paradigm, AI tools used for signal prioritization or case triaging may face fewer regulatory barriers than fully autonomous decision-making systems.

In summary, while AI holds immense promise for advancing pharmacovigilance, its regulatory acceptance is contingent upon the development of robust governance structures, transparent methodologies, standardized validation frameworks, and sustained human oversight. Addressing these challenges is essential to ensure that AI-enhanced pharmacovigilance systems are not only technologically advanced but also ethically sound, clinically reliable, and regulatorily compliant.

VI. LIMITATIONS, BIAS, AND ETHICAL CONSIDERATIONS IN AI-DRIVEN PHARMACOVIGILANCE

Despite the substantial promise of Artificial Intelligence in enhancing pharmacovigilance efficiency and analytical depth, its application is accompanied by significant methodological limitations, inherent biases, and complex ethical

considerations that warrant critical examination. These challenges are not merely technical in nature but extend to data integrity, clinical validity, regulatory trust, and societal implications. A nuanced understanding of these limitations is essential to prevent overreliance on algorithmic outputs and to ensure that AI-driven pharmacovigilance systems ultimately contribute to patient safety rather than inadvertently compromising it.

One of the most fundamental limitations of AI-based pharmacovigilance systems lies in their heavy dependence on data quality and representativeness. Spontaneous reporting systems, which form the backbone of post-marketing safety surveillance, are inherently subject to underreporting, selective reporting, duplication, and reporting bias. AI algorithms trained on such imperfect datasets may inadvertently learn and reinforce these biases, leading to skewed safety signal detection. For instance, adverse events associated with newer drugs, severe outcomes, or media attention are more likely to be reported, whereas events affecting vulnerable or underrepresented populations may remain systematically underreported. As a result, AI models may demonstrate high apparent performance while failing to capture the true epidemiological distribution of drug-related risks.

Algorithmic bias represents a particularly concerning issue in AI-driven pharmacovigilance. Bias may be introduced at multiple stages of model development, including data collection, labeling, feature selection, and algorithm design. Models trained predominantly on data from high-income countries may perform poorly when applied to low- and middle-income settings, where prescribing patterns, genetic backgrounds, healthcare access, and reporting behaviors differ substantially. Such geographical and demographic biases can undermine the global applicability of AI-based safety systems and exacerbate existing inequities in drug safety monitoring. Moreover, biased models may disproportionately generate false positives or false negatives in specific subpopulations, leading to inappropriate regulatory actions or missed safety signals.

Another critical limitation concerns the interpretability and explainability of complex AI models. Deep learning architectures, while powerful, often function as opaque black boxes, producing predictions without transparent reasoning pathways. In pharmacovigilance, where decisions have direct implications for public health, regulatory accountability, and legal responsibility, lack of explainability is a major impediment. Safety assessors and regulators must be able to understand why a particular drug–event association was flagged, which variables contributed most strongly, and whether the signal aligns with known pharmacological mechanisms. Without such interpretability, AI-generated signals may be met with skepticism, limiting their practical utility and acceptance.

Ethical considerations further complicate the deployment of AI in pharmacovigilance. Patient privacy and data protection are paramount concerns, particularly when AI systems integrate large-scale electronic health records, claims data, and real-world evidence sources. Even when data are anonymized, the risk of re-identification increases as datasets become more granular and interconnected. Ethical use of AI in pharmacovigilance therefore requires robust data governance frameworks, compliance with data protection regulations, and transparent communication regarding data usage. Failure to adequately address privacy concerns may erode public trust and hinder data sharing initiatives essential for effective safety surveillance.

Automation bias constitutes another ethical risk, wherein human experts may overtrust algorithmic outputs and discount their own clinical judgment. In pharmacovigilance workflows, this may result in uncritical acceptance of AI-generated safety signals or, conversely, dismissal of clinically plausible concerns not flagged by the system. Such overreliance on automation can weaken expert vigilance and reduce the depth of clinical reasoning applied to safety assessments. To mitigate this risk, AI systems must be explicitly designed and implemented as decision-support tools, with clearly defined roles for human oversight, validation, and accountability.

Transparency and accountability also raise ethical and legal questions regarding responsibility for AI-driven decisions. In cases where an AI system fails to detect a serious adverse drug reaction or generates a misleading signal, it remains unclear whether

responsibility lies with the algorithm developer, the pharmaceutical company, the pharmacovigilance professional, or the regulatory authority. This ambiguity poses challenges for liability frameworks and underscores the need for clearly articulated governance models defining ownership, accountability, and escalation pathways in AI-assisted pharmacovigilance.

Finally, there is a risk that the rapid adoption of AI technologies may outpace the development of appropriate ethical, regulatory, and educational infrastructures. Without adequate training, pharmacovigilance professionals may lack the skills necessary to critically evaluate AI outputs, understand model limitations, or identify potential biases. This skills gap may lead to superficial implementation of AI tools without meaningful integration into safety science principles. Ethical deployment of AI in pharmacovigilance therefore necessitates not only technical innovation but also sustained investment in workforce education, interdisciplinary collaboration, and ethical literacy.

In aggregate, while AI offers transformative potential for pharmacovigilance, its limitations, biases, and ethical challenges must be systematically acknowledged and addressed. Failure to do so may result in misplaced confidence, inequitable safety monitoring, and erosion of regulatory trust. A balanced, cautious, and ethically grounded approach is essential to ensure that AI enhances, rather than undermines, the fundamental objectives of pharmacovigilance.

VII. FUTURE PERSPECTIVES AND EMERGING DIRECTIONS OF ARTIFICIAL INTELLIGENCE IN PHARMACOVIGILANCE

The landscape of pharmacovigilance is undergoing a profound transformation driven by the convergence of artificial intelligence, real-world evidence, and computational pharmacoepidemiology. Future perspectives on AI in pharmacovigilance extend beyond incremental improvements in case processing or signal detection, encompassing a paradigm shift toward proactive, predictive, and mechanistically informed drug safety surveillance. Emerging directions emphasize not only technological innovation but also integration with regulatory frameworks, ethical governance, and global data harmonization, which collectively will shape the next generation of pharmacovigilance systems.

One of the most compelling future directions is the development of predictive pharmacovigilance, whereby AI models anticipate potential adverse drug reactions before they manifest at the population level. Leveraging large-scale longitudinal datasets, including electronic health records, claims databases, genomics repositories, and patient-reported outcomes, predictive models can identify individuals or subpopulations at elevated risk of drug-related adverse events. Machine learning and deep learning architectures, particularly those integrating temporal sequence analysis and multi-modal data, are expected to enhance the accuracy and timeliness of such predictions. Predictive pharmacovigilance holds the potential to shift safety surveillance from reactive post-marketing assessment to a proactive, risk-mitigation paradigm, enabling early intervention, personalized risk management, and more informed regulatory decision-making.

Integration of multi-source real-world evidence represents another critical frontier. Pharmacovigilance data are increasingly heterogeneous, encompassing spontaneous reporting systems, electronic health records, wearable devices, social media, biomedical literature, and mobile health applications. AI methodologies capable of harmonizing, synthesizing, and extracting actionable insights from such diverse data streams will significantly enhance signal detection sensitivity and contextual understanding. For example, natural language processing pipelines combined with graph-based deep learning models may identify novel drug–event associations by integrating patient narratives, clinical notes, and literature reports in a single analytic framework. Such integration not only broadens the evidence base but also enables detection of rare, delayed, or geographically distributed adverse events that traditional surveillance systems may overlook.

Explainable and human-in-the-loop AI frameworks are projected to become central in regulatory-aligned pharmacovigilance. As regulatory agencies increasingly demand transparency, interpretability, and accountability, AI systems must provide outputs

that are both scientifically rigorous and understandable to pharmacovigilance professionals. Explainable AI (XAI) techniques, including attention-based mechanisms, feature attribution methods, and hybrid knowledge-informed models, will enable stakeholders to trace safety signals to underlying data patterns, clinical reasoning, and mechanistic plausibility. Human-in-the-loop integration will ensure that AI-generated insights are validated by expert judgment, maintaining regulatory confidence while leveraging the computational power of AI.

Automation of routine pharmacovigilance workflows is another emerging trend, with AI facilitating semi-autonomous processing of ICSRs, medical coding, duplicate detection, and triage prioritization. By automating repetitive tasks, AI frees human experts to focus on higher-order analytical responsibilities such as signal validation, causality assessment, and regulatory reporting. The development of modular, interoperable AI platforms will allow organizations to scale pharmacovigilance operations efficiently while maintaining data integrity, compliance, and adaptability to evolving regulatory requirements.

From a technological standpoint, reinforcement learning and adaptive AI models are anticipated to play a transformative role. Unlike conventional supervised or unsupervised learning, reinforcement learning allows AI systems to optimize decision-making policies iteratively, learning from feedback within a dynamic pharmacovigilance environment. Such approaches may enable adaptive prioritization of safety signals, optimization of data collection strategies, and intelligent allocation of regulatory and clinical resources. Coupled with advanced simulation and digital twin models, reinforcement learning could facilitate scenario-based risk assessment, simulating drug safety outcomes under hypothetical exposure scenarios before real-world occurrence.

Finally, global harmonization of AI-based pharmacovigilance practices is emerging as a strategic imperative. Differences in data accessibility, reporting standards, regulatory frameworks, and computational infrastructure create barriers to effective multinational safety monitoring. Collaborative initiatives, international consortia, and open-access AI platforms are expected to bridge these gaps, facilitating standardized methodologies, interoperable datasets, and cross-jurisdictional validation. Such harmonization will not only enhance model generalizability but also support equitable drug safety surveillance across diverse populations and healthcare systems.

In conclusion, the future of AI in pharmacovigilance is characterized by predictive, integrative, explainable, and adaptive approaches, underpinned by rigorous ethical, regulatory, and technical frameworks. Continued interdisciplinary collaboration among data scientists, clinicians, regulatory authorities, and ethicists will be essential to realize the full potential of AI, ensuring that drug safety surveillance evolves from reactive monitoring to anticipatory, patient-centric risk management.

VIII. CONCLUSION

In summary, the integration of artificial intelligence into pharmacovigilance represents a paradigm shift in the science of post-marketing drug safety surveillance, enabling substantial enhancements in data processing, signal detection, and risk assessment that were previously unattainable through traditional methodologies. Conventional pharmacovigilance has long been constrained by manual review, underreporting, reporting bias, and the limited analytical capacity of disproportionality measures applied to spontaneous reporting systems; these constraints impede the timely identification of adverse drug reactions and the generation of high-confidence safety signals. The application of machine learning, deep learning, and natural language processing has considerably broadened the scope and sensitivity of safety surveillance systems by enabling automated case processing, extraction of adverse events from unstructured clinical text, integration of diverse data sources, and identification of latent associations that historically went undetected. Multiple systematic and narrative reviews have corroborated that AI approaches can outperform traditional statistical methods in several contexts, offering improved sensitivity and timeliness in signal detection, particularly when applied to large, heterogeneous datasets such as electronic health records, spontaneous reporting systems, and patient-generated text from social media platforms. A scoping review of machine learning applications in pharmacovigilance identified rising utilization of deep learning architectures and

advanced analytical techniques as well as enduring methodological gaps in adoption of best practices and generalizability of models across diverse settings. Furthermore, implementations of AI-enabled Bayesian networks and hybrid knowledge-informed models demonstrate promising performance gains in causality assessment and prioritization workflows, suggesting practical utility for regulatory and industry stakeholders beyond purely academic inquiry.

Nevertheless, the evidence base reveals that AI adoption in routine pharmacovigilance practice remains nascent and uneven, with significant challenges that temper its transformative potential. Data quality and representativeness continue to constrain model performance: incomplete or biased safety reports, heterogeneous data formats, and underrepresentation of specific populations lead to algorithmic bias and risk misspecification if not addressed with robust preprocessing and annotation strategies. Moreover, many AI models reported in the literature lack consistent external validation and demonstrate limited generalizability across clinical contexts, raising concerns about reproducibility and clinical reliability. Ethical and regulatory frameworks have yet to mature sufficiently to provide clear standards for explainability, accountability, and governance of AI systems in safety surveillance; regulators such as the FDA and EMA emphasize transparency and interpretability, particularly for complex deep learning models whose decision pathways are often opaque ("black box"). The potential for reinforcement of existing disparities—through algorithmic bias affecting under-reported populations—and privacy implications associated with integration of sensitive health data further complicate clinical deployment and regulatory acceptance. Data governance, human-in-the-loop oversight, and standardized validation protocols are therefore prerequisites for responsible AI integration into pharmacovigilance workflows.

Despite these challenges, the trajectory of research and emerging implementations suggests a sustained expansion and refinement of AI capabilities in drug safety science. Hybrid and ensemble methodologies that combine machine learning with expert-defined rules and knowledge graphs offer a promising compromise between performance and interpretability; initiatives to strengthen common data models and structured real-world evidence repositories facilitate model training and validation across populations; and advances in explainable AI align computational outputs with regulatory expectations and clinical reasoning. As the volume and diversity of safety data continue to expand with digital health innovations, the agility of AI systems to ingest, harmonize, and interpret this information portends a future of pharmacovigilance that is more proactive, predictive, and patient-centered than previously possible. Addressing the current limitations will require coordinated efforts across academia, industry, regulatory agencies, and clinical practice to establish consensus-driven benchmarks, ethical standards, and interoperable infrastructures that support safe, equitable, and scientifically robust AI-enhanced pharmacovigilance.

REFERENCES

[1] Das S, Dey A, Pal A, Roy N. Applications of artificial intelligence in machine learning: review and prospect. *Int J Comput Appl.* 2015;115:31–41.

[2] Esteva A, Robicquet A, Ramsundar M, et al. A guide to deep learning in healthcare. *Nat Med.* 2019;25(1):24–29.

[3] 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.

[4] Pilipiec P, Liwicki M, Bota A. Using machine learning for pharmacovigilance: a systematic review. *Pharmaceutics.* 2022;14(2):266.

[5] [Author(s) unavailable]. AI-driven pharmacovigilance: Enhancing adverse drug reaction detection with deep learning and NLP. *MethodsX.* 2025;15:103460.

[6] [Author(s) unavailable]. Artificial intelligence-powered pharmacovigilance: A review of machine and deep learning in clinical text-based ADE detection. *J Biomed Inform.* 2024;152:104621.

[7] [Author(s) unavailable]. Leveraging NLP and machine learning methods for ADE detection in EHRs: a scoping review. *PubMed.* 2024.

[8] [Author(s) unavailable]. Artificial intelligence in pharmacovigilance: narrative review and Bayesian network case study. *PMCID.* 2025.

[9] Kompa B, Hakim JB, Palepu A, et al. Artificial intelligence based on machine learning in pharmacovigilance: a scoping review. *Drug Saf.* 2022;45(5):477–491.
[10] Savy Dsouza V, Leyens L, Kurian JR, et al. Artificial intelligence in pharmacovigilance: ADR prediction in hospitalized patients. *Res Soc Adm Pharm.* 2025;21:453–462.
[11] Rudnisky E, Paudel K. Pharmacovigilance in the era of artificial intelligence: advancements, challenges, and considerations. *Cureus.* 2025;17(6):e86972.
[12] Ward IR, Wang L, Lu J, et al. Explainable artificial intelligence for pharmacovigilance: predicting adverse outcomes. *arXiv.* 2021.
[13] Abdeldjouad FZ, Brahami M, Sabri M. Evaluating AI effectiveness for ADR prediction in cancer patients: systematic review & meta-analysis. *arXiv.* 2024.
[14] Bampa M. Handling temporality of clinical events for ADE detection in EHRs: scoping review. *arXiv.* 2019.
[15] Harpaz R, Callahan A, Tamang S, et al. Text mining for adverse drug events: state-of-the-art and challenges. *Drug Saf.* 2014;37(4):Epub.
[16] Botsis T, Nguyen M, Woo EJ, et al. Text mining for adverse drug events. *J Biomed Inform.* 2011;44(5):857–867.
[17] Chin WW, Siau K, Nah FF-H. Machine learning in healthcare: challenges and opportunities. *Health Informatics J.* 2021;27(2):1460–1474.
[18] Denecke K. Using natural language processing to extract health-related information from social media. *Methods Inf Med.* 2017;56(6):372–381.
[19] Lehne M, Sass J, Essenwanger A, et al. Why digital health technologies fail: addressing regulatory and trust issues. *NPJ Digit Med.* 2019;2:47.
[20] Bate A, Lindquist M, Edwards IR, et al. A Bayesian neural network method for adverse drug reaction signal generation. *Eur J Clin Pharmacol.* 1998;54:315–321.
[21] 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.
[22] Hauben M, Aronson JK. Defining "signal" and its subtypes in pharmacovigilance. *Drug Saf.* 2009;32(2):99–110.
[23] Trifirò G, van der Lei J. ATC/DDD methodology overview in real-world evidence. *Br J Clin Pharmacol.* 2015;80:273–280.
[24] Patsou P, Fotiadis DI. Data quality challenges in pharmacovigilance: addressing real-world evidence. *Clin Pharmacol Ther.* 2022;112:230–235.
[25] Shook DR, et al. Federated learning for privacy-preserving pharmacovigilance analytics. *IEEE J Biomed Health Inform.* 2023;27:456–467.
[26] Kessler E. AI-powered signal detection in pharmacovigilance databases: methods and case studies. 2025.
[27] World Health Organization. WHO pharmacovigilance guidelines: definitions and requirements for drug safety monitoring. Geneva: WHO; 2004.
[28] Food and Drug Administration. Guidance for Industry: Pharmacovigilance Planning. FDA; 2022.
[29] European Medicines Agency. Guideline on Good Pharmacovigilance Practices (GVP). EMA; 2024.
[30] International Council for Harmonisation E2B(R3). Data elements for transmission of Individual Case Safety Reports. ICH; 2023.
[31] Peduzzi P, et al. A simulation study of the number of events per variable in logistic regression analysis. *J Clin Epidemiol.* 1996;49(12):1373–1379.
[32] Beam AL, Kohane IS. Big data and machine learning in health care. *JAMA.* 2018;319(13):1317–1318.
[33] Collier N, Doan S, et al. BioCaster: detecting public health rumors with NLP. *ACM Trans Knowl Discov Data.* 2008;2(4):1–31.
[34] Lehmann CU, et al. Big data analytics in clinical pharmacology and pharmacovigilance. *Clin Pharmacol Ther.* 2019;105:1257–1269.
[35] Nelson TB, et al. Causal inference approaches for drug safety signal evaluation with real-world evidence. *Pharmacoepidemiol Drug Saf.* 2021;30:765–772.
[36] Pivovarov R, et al. Learning probabilistic phenotypes from heterogeneous EHR data. *J Biomed Inform.* 2015;58:156–165.
[37] Sarker A, Gonzalez G. Portable automatic text classification for adverse drug event detection via multi-corpus training. *J Biomed Inform.* 2017;71:49–59.

[38] Leaman R, et al. Towards large-scale Twitter mining for drug safety surveillance. *J Biomed Inform.* 2010;43(2):332–342.

[39] Nikfarjam A, et al. Pharmacovigilance on Twitter? An analysis of drug safety tweets. *AMIA Annu Symp Proc.* 2015;2015:924–933.

[40] Wong A, et al. Clinical text classification with convolutional neural networks for adverse event phenotyping. *J Biomed Inform.* 2016;62:59–69.

[41] Wang Y, et al. Deep learning for identifying drug–drug interactions. *J Am Med Inform Assoc.* 2019;26:1933–1941.

[42] Zhang R, et al. Integrating patient genomics with real-world data for ADR prediction. *Nat Genet.* 2020;52(11):1148–1156.

[43] Esteva A, et al. Dermatologist-level classification with deep neural networks. *Nature.* 2017;542:115–118.

[44] Lipton ZC. The mythos of model interpretability in ML. *CACM.* 2018;61(10):36–43.

[45] Ribeiro MT, Singh S, Guestrin C. "Why should I trust you?" Explaining predictions of classifiers. *KDD.* 2016.

[46] Lundberg SM, Lee S-I. A unified approach to interpreting model predictions. *NIPS.* 2017.

[47] Chen R, Dewey AL, et al. Deep learning for adverse event detection from clinical text. *J Biomed Inform.* 2024;152:104612.

[48] Kaviani P, Landi SN, McKethan A, et al. Expert opinion on drug safety and AI applications. *Expert Opin Drug Saf.* 2023;22(4):373–380.

[49] Liang L, Hu J, Sun G, et al. Artificial intelligence-based pharmacovigilance in the setting of limited resources. *Drug Saf.* 2022;45(6):511–519.

[50] Hu Q, Wu B, Wu J, et al. Predicting adverse drug events in older inpatients: a machine learning study. *Int J Clin Pharm.* 2022;44(6):1304–1311.

[51] Eskildsen NK, Eriksson R, Christensen SB, et al. Implementation and comparison of text mining methods for signal detection of medication errors. *BMC Med Inform Decis Mak.* 2020;20:94.

[52] Niazi SK. The coming of age of AI/ML in drug development and regulatory science: FDA perspectives. *Drug Des Devel Ther.* 2023;17:2691–2725.

[53] Harpaz R, DuMouchel W, LePendu P, Bauer-Mehren A. Performance of pharmacovigilance signal detection algorithms on the FDA adverse event reporting system. *Drug Saf.* 2013;36(12):1113–1127.

[54] van Puijenbroek EP, Bate A, Leufkens HG, et al. Comparison of measures of disproportionality for signal detection in spontaneous reporting systems for adverse drug reactions. *Pharmacoepidemiol Drug Saf.* 2002;11(1):3–10.

[55] Coloma PM, Schuemie MJ, Trifirò G, et al. Combining electronic healthcare databases and spontaneous reporting systems to detect drug safety signals. *PLoS One.* 2013;8(3):e57218.

[56] Tatonetti NP, Ye PP, Daneshjou R, Altman RB. Data-driven prediction of drug effects and interactions. *Sci Transl Med.* 2012;4(125):125ra31.

[57] Jiang Y, He Z, Tan J, et al. Predictive pharmacovigilance using machine learning and electronic health records. *Clin Pharmacol Ther.* 2020;107:1362–1372.

[58] Ho TK, Kim E, et al. Use of neural networks for detecting adverse drug events in electronic medical records. *J Biomed Inform.* 2021;114:103667.

[59] Delaney JA, et al. Combining social media and EHR data for enhanced pharmacovigilance signal detection. *Pharmacoepidemiol Drug Saf.* 2022;31(7):695–705.

[60] Xu R, Goodwin T, et al. Natural language processing and deep learning in pharmacovigilance: applications and future perspectives. *Drug Saf.* 2021;44(8):809–820.

[61] Juhlin K, et al. Text mining for adverse drug events in patient forums. *J Med Internet Res.* 2020;22(8):e18984.

[62] Sarker A, Ginn R, Nikfarjam A, et al. Utilizing social media data for pharmacovigilance: challenges and opportunities. *J Biomed Inform.* 2015;54:202–212.

[63] Huang K, Altosaar J, Ritchie MD. Clinical NLP for pharmacovigilance: progress and challenges. *Brief Bioinform.* 2020;21(2):536–549.

[64] Nikfarjam A, Gonzalez G. Pattern mining for drug safety surveillance on social media. *AMIA Annu Symp Proc.* 2011;2011:866–875.

[65] Wang X, et al. Integrating real-world data into pharmacovigilance: methodological considerations. *Pharmacoepidemiol Drug Saf.* 2020;29:1560–1570.

[66] DuMouchel W, Pregibon D. Empirical Bayes screening for multi-item associations. *Proceedings of KDD.* 2001;67–76.

[67] Harpaz R, Callahan A, et al. Empirical evaluation of signal detection algorithms for the FDA Adverse Event Reporting System. *Drug Saf.* 2012;35:399–409.

[68] Trifirò G, Coloma PM, et al. Combining electronic healthcare databases and spontaneous reporting systems to enhance signal detection. *PLoS One.* 2013;8:e57218.

[69] Hauben M, Zhou X, et al. Quantitative methods in pharmacovigilance: statistical signal detection. *Drug Saf.* 2005;28(6):517–534.

[70] Bate A, Lindquist M, Edwards IR. Bayesian neural network approach to adverse drug reaction detection. *Eur J Clin Pharmacol.* 1998;54:315–321.

[71] Peduzzi P, Concato J, Kemper E, et al. Simulation study of events per variable in logistic regression. *J Clin Epidemiol.* 1996;49:1373–1379.

[72] Saito T, Rehmsmeier M. The precision-recall plot is more informative than ROC plot for imbalanced datasets. *PLoS One.* 2015;10:e0118432.

[73] Chen R, Dewey AL, et al. Evaluation of machine learning for ADE detection using EHR data. *J Biomed Inform.* 2024;152:104612.

[74] Kaviani P, Landi SN, McKethan A, et al. Expert opinion on drug safety and AI applications. *Expert Opin Drug Saf.* 2023;22:373–380.

[75] Liang L, Hu J, Sun G, et al. AI-based pharmacovigilance in low-resource settings. *Drug Saf.* 2022;45:511–519.

[76] Hu Q, Wu B, Wu J, et al. Predicting adverse drug events in older inpatients: ML study. *Int J Clin Pharm.* 2022;44:1304–1311.

[77] Eskildsen NK, Eriksson R, Christensen SB, et al. Implementation and comparison of text mining methods for medication error detection. *BMC Med Inform Decis Mak.* 2020;20:94.

[78] Niazi SK. AI/ML in drug development and regulatory science: FDA perspectives. *Drug Des Devel Ther.* 2023;17:2691–2725.

[79] WHO. Reporting and learning systems for medication errors: pharmacovigilance centers. Geneva: WHO; 2014.

[80] FDA. Guidance for industry: Pharmacovigilance planning. U.S. Food and Drug Administration; 2022.

[81] EMA. Guideline on good pharmacovigilance practices (GVP). European Medicines Agency; 2024.

[82] ICH E2B(R3). Data elements for transmission of individual case safety reports. International Council for Harmonisation; 2023.

[83] Harpaz R, Callahan A, et al. Empirical evaluation of signal detection in FAERS. *Drug Saf.* 2012;35:399–409.

[84] van Puijenbroek EP, Bate A, Leufkens HG, et al. Measures of disproportionality for signal detection in spontaneous reporting systems. *Pharmacoepidemiol Drug Saf.* 2002;11:3–10.

[85] Coloma PM, Schuemie MJ, Trifirò G, et al. Combining electronic healthcare databases and spontaneous reporting systems to detect drug safety signals. *PLoS One.* 2013;8:e57218.

[86] Tatonetti NP, Ye PP, Daneshjou R, Altman RB. Data-driven prediction of drug effects and interactions. *Sci Transl Med.* 2012;4:125ra31.

[87] Jiang Y, He Z, Tan J, et al. Predictive pharmacovigilance using ML and EHRs. *Clin Pharmacol Ther.* 2020;107:1362–1372.

[88] Ho TK, Kim E, et al. Neural networks for detecting ADEs in EHRs. *J Biomed Inform.* 2021;114:103667.

[89] Delaney JA, et al. Combining social media and EHR data for pharmacovigilance signal detection. *Pharmacoepidemiol Drug Saf.* 2022;31:695–705.

[90] Xu R, Goodwin T, et al. NLP and deep learning in pharmacovigilance: applications and future perspectives. *Drug Saf.* 2021;44:809–820.

[91] Juhlin K, et al. Text mining for adverse drug events in patient forums. *J Med Internet Res.* 2020;22:e18984.

[92] Sarker A, Ginn R, Nikfarjam A, et al. Social media data for pharmacovigilance: challenges and opportunities. *J Biomed Inform.* 2015;54:202–212.

[93] Huang K, Altosaar J, Ritchie MD. Clinical NLP for pharmacovigilance: progress and challenges. *Brief Bioinform.* 2020;21:536–549.

[94] Nikfarjam A, Gonzalez G. Pattern mining for drug safety surveillance on social media. *AMIA Annu Symp Proc.* 2011;2011:866–875.

[95] Wang X, et al. Integrating real-world data into pharmacovigilance: methodological considerations. *Pharmacoepidemiol Drug Saf.* 2020;29:1560–1570.

[96] DuMouchel W, Pregibon D. Empirical Bayes screening for multi-item associations. *Proceedings of KDD.* 2001;67–76.
[97] Harpaz R, Callahan A, et al. Empirical evaluation of signal detection algorithms for FAERS. *Drug Saf.* 2012;35:399–409.
[98] Trifirò G, Coloma PM, et al. Combining electronic healthcare databases and spontaneous reporting systems. *PLoS One.* 2013;8:e57218.
[99] Hauben M, Zhou X, et al. Quantitative methods in pharmacovigilance: statistical signal detection. *Drug Saf.* 2005;28:517–534.
[100] Bate A, Lindquist M, Edwards IR. Bayesian neural network approach to adverse drug reaction detection. *Eur J Clin Pharmacol.* 1998;54:315–321.
[101] Peduzzi P, Concato J, Kemper E, et al. Simulation study of events per variable in logistic regression. *J Clin Epidemiol.* 1996;49:1373–1379.
[102] Saito T, Rehmsmeier M. Precision-recall plot more informative than ROC for imbalanced datasets. *PLoS One.* 2015;10:e0118432.
[103] Chen R, Dewey AL, et al. Evaluation of ML for ADE detection using EHRs. *J Biomed Inform.* 2024;152:104612.
[104] Kaviani P, Landi SN, McKethan A, et al. Expert opinion on drug safety and AI applications. *Expert Opin Drug Saf.* 2023;22:373–380.
[105] Liang L, Hu J, Sun G, et al. AI-based pharmacovigilance in low-resource settings. *Drug Saf.* 2022;45:511–519.
[106] Hu Q, Wu B, Wu J, et al. Predicting adverse drug events in older inpatients: ML study. *Int J Clin Pharm.* 2022;44:1304–1311.
[107] Eskildsen NK, Eriksson R, Christensen SB, et al. Implementation and comparison of text mining methods for medication error detection. *BMC Med Inform Decis Mak.* 2020;20:94.