Epidemiology & Health Data Insights
(eISSN 3080-8111)

Review Article

Developing and Validating Predictive Models for Adverse Drug Reactions using Electronic Health Records (EHRs): A Narrative Review

Vivian Ukamaka Nwokedi[1], Michael Kayode Odedele[2], Taofeek Adeshina Yusuff[3], Irene Adjoa Anderson[4], Olukunle O. Akanbi[5], Chiamaka Pamela Agu[6], Amber Otibhor Omoike[7]

[1]*Department of Clinical Pharmacy, Faculty of Pharmacy, University of Benin, Benin City, Nigeria*

[2]*Department of Kinesiology, College of Health, Human Services & Nursing, California State University, Dominguez Hills, Carson, CA, USA*

[3]*Department of Business Analytics, Pompea College of Business, University of New Haven, West Haven, CT, USA*

[4]*Department of Pharmacology, Faculty of Pharmacy and Pharmaceutical Sciences, Kwame Nkrumah University of Science and Technology, Kumasi, Ghana*

[5]*Department of Psychology and Behavioral Sciences, National Louis University, Tampa, FL, USA*

[6]*Department of Public Health, University of New Haven, West Haven, CT, USA*

[7]*Department of Pharmacy, Irrua Specialist Teaching Hospital, Irrua, Edo State, Nigeria*

Received: Sep 24, 2025
Accepted: Oct 28, 2025

Corresponding author's email: viviannwokedi250@gmail.com

Abstract:

Adverse drug reactions (ADRs) remain a major global challenge, contributing substantially to patient morbidity, mortality, and healthcare costs. Traditional pharmacovigilance approaches—spontaneous reporting and post-marketing surveillance—are hampered by underreporting, delays, and limited contextual data. The growing availability of electronic health records (EHRs), which capture longitudinal structured and unstructured patient information, presents an unprecedented opportunity to advance ADR prediction. This narrative review synthesizes recent progress in developing and validating predictive models that leverage EHRs, highlighting methodological approaches, challenges, and future directions. Predictive strategies range from traditional regression models to advanced machine learning and deep learning architectures, with multimodal frameworks increasingly integrating structured fields (demographics, labs, prescriptions) and unstructured clinical text through natural language processing. While ensemble and deep learning methods demonstrate superior performance, issues of data quality, missingness, bias, and interpretability persist. Robust validation frameworks—spanning internal cross-validation to multi-center external testing—are critical to ensure generalizability and clinical trustworthiness. Ethical considerations, including fairness, privacy, and transparency, remain central to safe deployment. Looking forward, promising avenues include federated learning across institutions, integration of multi-omics and pharmacogenomic data, explainable AI tailored for clinical use, and real-time monitoring through digital twin frameworks. These trajectories, combined with robust governance and clinician–data scientist collaboration, have the potential to transform ADR detection from a reactive process to proactive, personalized prevention. By synthesizing the existing evidence, this review provides insights into the development of more effective predictive models for ADRs and informs strategies for improving pharmacovigilance. This study will contribute to the ongoing efforts to leverage EHRs and predictive models for improving patient outcomes and reducing the burden of ADRs.

Keywords: Adverse Drug Reactions (ADRs); Electronic Health Records (EHRs); Predictive Modeling; Machine Learning; Pharmacovigilance; Artificial Intelligence (AI)

Introduction

Adverse drug reactions (ADRs) is a persistent, challenging and life threatening event encountered in modern medicine which is of public health concern. These unintended, harmful responses to medications—occurring at doses normally used for prophylaxis, diagnosis, or therapy—are a major cause of patient morbidity and mortality worldwide [1]. The spectrum of ADRs ranges from mild discomforts to severe, life-threatening events, and their unpredictability complicates efforts to mitigate risk. Clinical trials and pre-marketing studies often fail to identify the full range of ADRs, leaving clinicians and health systems to depend heavily on post-marketing surveillance, spontaneous reporting systems, and case studies. Yet these traditional methods are hampered by underreporting, delayed detection, and lack of detailed clinical context [1, 2].

The burden of ADRs on patient safety and healthcare systems is substantial. It has been estimated that ADRs account for between 3.6% and 6.5% of hospital admissions [3], with as many as 10-20% of hospitalized patients developing ADRs during their stay [4]; whereas 16% mortality rate was recorded to be associated with ADRs in a cross-sectional survey conducted in South Africa [5]. In older adult populations, the prevalence of ADRs is 11% [6], and it is responsible for 3.3% of hospital admissions, especially amid polypharmacy and multiple comorbid conditions [7]. Beyond patient harm, ADRs impose a significant economic cost on healthcare systems: increased length of stay, extra diagnostic and therapeutic interventions, and in some cases long-term disability. From the perspective of pharmacovigilance, the challenge is heightened by system inefficiencies—including delays in signal detection, incomplete reporting, and limited data on patient-level risk factors.

In recent years, Electronic Health Records (EHRs) have emerged as a promising resource to ameliorate many of these limitations. Data contained in EHRs is broadly categorized into structured data and unstructured data. The structured data include the following information demographics, diagnoses, prescriptions, laboratory results and imaging. On the other hand, the unstructured data include data contained in clinical notes. This category of data provides longitudinal, real-world views of patient health trajectories [1, 8], as well as an opportunity to first hand information relating to drug administration over time, co-morbidities, concomitant drug administration as well as outcomes of therapy (including clinical, patient, economic and humanistic outcomes) [9, 10]. In addition, data from EHRs significantly support research (including both retrospective and prospective studies) and offer opportunities for improved statistical power as well as generalizability as EHR feature large sample size.

Complementing the availability of rich EHR data, predictive modeling using artificial intelligence (AI) and machine learning (ML) has become a dynamic and rapidly expanding strategy to identify patients at risk of ADRs before harm occurs. Machine learning models—such as random forests, gradient boosting, support-vector machines—and more recently, deep learning architectures, knowledge graph embeddings, and attention mechanisms have been applied both to structured EHR data and unstructured clinical text [11, 12]. These approaches have shown promise: for example, AI-based models have achieved high sensitivity and specificity in certain settings, such as oncology or neonatology, and risk scores have been proposed to classify high versus low ADR risk in vulnerable populations [13, 14]. Nonetheless, challenges remain around model validation, transportability, interpretability, and bias.

This narrative review aims to synthesize recent advances in the development and validation of predictive models for ADRs using EHR data. Specifically, our objectives are: (1) to map the landscape of existing ADR-prediction models built from EHRs; (2) to examine the methodologies used for model development—including feature engineering, model selection, and use of structured versus unstructured data; (3) to evaluate how these models have been validated (internal and external), and assess reported performance metrics such as discrimination, calibration, and clinical usefulness; (4) to discuss key challenges (e.g., data quality, missingness, bias, interpretability) and ethical considerations; and (5) to propose future directions and recommendations for research, clinical translation, and governance. By providing this synthesis, we expect this review will help researchers, clinicians, and policymakers to understand what has been achieved, what gaps remain, and how best to leverage EHR-driven predictive modeling in ADR prevention to improve patient safety.

Methodology

PubMed and Google Scholar were the principal sources of literature for this review. Additionally, the reference list of the studied literature was examined to pinpoint more relevant publications. We subsequently examined and integrated findings from the relevant literature thematically in relation to the topic.

Electronic Health Records as a Data Source for ADR Prediction

Electronic Health Records (EHRs) data that captures a patient's journey through healthcare thereby enabling the identification of adverse drug reactions. Hence, such data offer a wide range of opportunities owing to the fact that they provide patient follow-up information across interventions as well as medical and medication records, laboratory investigations and results and pertinent information on responses to various medication classes. This enables not only retrospective studies of ADRs but also potentially near-real-time monitoring. However, EHRs provide valuable insights into real-world clinical practice compared to the controlled environment of clinical trials. Therefore they feature different patient populations as well as multiple comorbidities, polypharmacies that are often underrepresented in trial data. With this real-world features of EHRs medication usage data including drug usage patterns, off-label uses, and rare ADRs may be more visible [14].

Notwithstanding, the adoption of EHRs data for ADR prediction models should be with caution to strengthen data validity and generalizability of findings as EHRs face numerous drawbacks. Amongst such obstacles is "data heterogeneity". This is because different healthcare systems use different coding methods, terminologies, and different EHR platforms. Laboratory test names or units may differ; medication names may be non-standardized; adverse events may be recorded with different vocabularies. Missing data is pervasive: structured fields may be blank, or lab tests may be ordered infrequently; crucial exposures or outcomes may be absent or mis-recorded. Coding errors occur, both in diagnoses and medications: misclassification or undercoding of ADRs is common. Privacy and governance concerns complicate data access and sharing, especially when unstructured clinical notes might contain identifiable or sensitive information.

Within this mixed data environment, Natural Language Processing (NLP) plays a crucial role in extracting ADR signals from unstructured clinical text. Clinical notes often contain descriptions of symptoms, side effects, drug tolerability, temporal relationships (e.g., when a reaction started relative to drug administration), and free-text contextual details that structured fields miss [15]. Recent work has shown that integrating structured and unstructured data via NLP can significantly increase sensitivity in detecting ADRs. For example, pediatric pharmacovigilance studies demonstrate that many adverse drug events are recorded only in clinical notes but not in diagnostic codes or billing data. NLP pipelines, including rule-based methods, traditional supervised learning models, and more recently transformer-based architectures, have been applied to identify drug-reaction pairs, to extract medication information from narrative text, and to normalize these into standard terminologies (e.g., RxNorm, MedXN) [16]. Tools that adopt interoperability standards, such as Health Level Seven Fast Healthcare Interoperability Resources (HL7 FHIR), to encode extracted substance, timing, and reaction attributes help ensure that output features are usable across institutions [17].

Overall, Figure 1 provides suitable illustration of how EHR data flows from raw form to a state usable for predictive modeling. First, multiple raw EHR sources provide structured data (demographics, prescriptions, labs, diagnoses, etc.) and unstructured data (clinical notes, narratives). This raw data must undergo cleaning and preprocessing: dealing with missing data, standardizing codes, normalizing units, detecting and removing duplicates or errors, ensuring privacy via de-identification. After cleaning, feature engineering constructs predictors from the data: drug exposure histories (dose, duration), patient covariates (comorbidities, age, sex, lab trends), temporal relationships, perhaps signal features derived from text (frequency of symptom mentions, co-occurrence). Then the modeling stage applies machine learning or deep learning methods, perhaps including multimodal models that combine structured and text-derived features. Finally, model outputs include risk prediction (probability of an ADR), risk stratification (low, moderate, high risk), and alerts or decision support suggestions. The pipeline ensures that the expressive content of EHRs is converted into usable features while maintaining validity, transparency, and reproducibility.

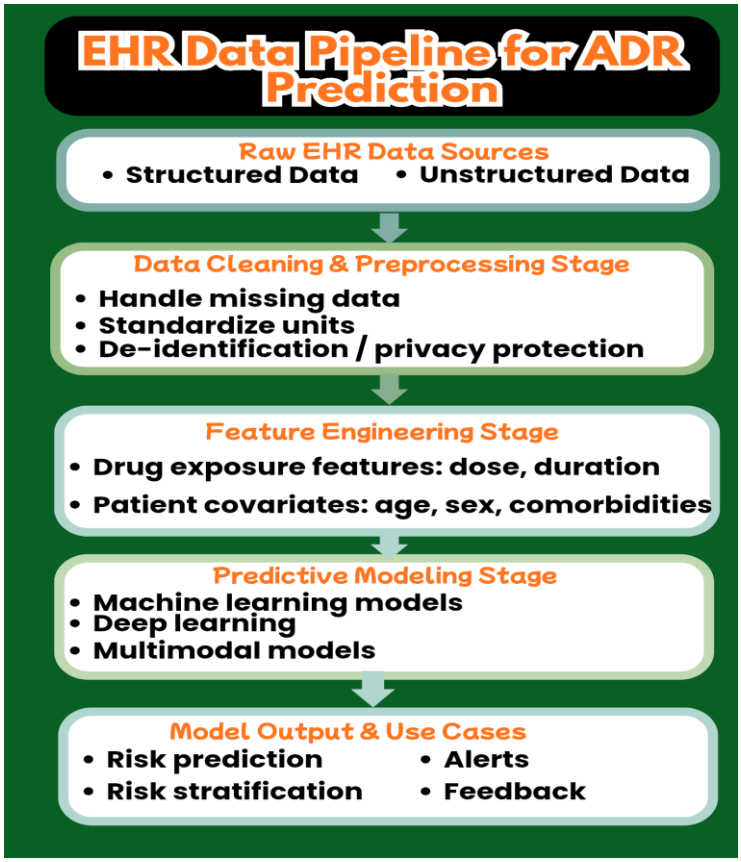

Figure 1: Pipeline for ADR Prediction Using EHRs Data

Predictive Modeling Approaches for ADRs

There are many approaches that have been found to be useful in the growing field of ADR prediction. Such as statistical modeling, machine learning, and deep learning approaches. Each is associated with different assumptions, strengths, and limitations. Whereas classical statistical modeling provides interpretability and simplicity, modern machine learning and deep networks seek to capture nonlinearities, interactions, and high-dimensional signals; hybrid models aim to harness both structured and unstructured data to optimize performance. We discussed these approaches in detail under the following headings:

Traditional Statistical Methods

These are the oldest mathematical models that have been found useful in biomedical prediction. They include logistic regression and Cox proportional hazards modeling and they have been the starting points in ADR research. For example, in logistic regression, the log-odds of an ADR event as a linear function of covariates (drug exposures, patient comorbidities, demographics) is modeled. Its appeal lies in interpretability: estimated coefficients can be directly understood as odds ratios, facilitating clinical insight. Cox proportional hazards models extend this to time-to-event frameworks, modeling the hazard of ADR occurrence as a function of covariates over follow-up time, accommodating censoring of patients who never experience the ADR.

Although these models have recorded great success in their application, they are faced with a lot of limitations in their application in ADR prediction. Usually they often assume linear effects (in logistic regression), proportional hazards (in Cox models), and require prior specification of interaction terms. In high-dimensional settings—such as when combining labs, medication histories, and text-derived features—the curse of dimensionality and multicollinearity may hinder performance. Furthermore, complex, nonlinear interactions or latent factors in EHRs may be missed. That said, logistic regression is still often used as a baseline or benchmark in comparative modeling studies [18, 19].

Machine Learning Techniques

To overcome the linearity constraints of traditional models, machine learning (ML) techniques have been widely adopted in ADR prediction. Core methods include decision forest and tree-based ensembles (e.g. random forests, gradient boosting machines), support vector machines (SVMs), and other kernel methods. These models are capable of capturing nonlinear relationships and interactions implicitly, and can usually

handle higher-dimensional feature spaces as long as overfitting is managed.

In a notable work, Hu et al. (2024) [19] systematically reviewed ADR prediction using machine learning and noted that ensemble models such as random forests, XGBoost, and AdaBoost often outperform simpler classifiers under appropriate regularization and tuning. Zhao et al. (2021) [20] applied a variety of ML methods (SVM, random forest, AdaBoost, XGBoost) on an imbalanced EHR-derived ADR dataset, combining structured and text features, and found that ensemble methods improved predictive performance and robustness compared to traditional logistic regression. Additionally, they utilized ensembling in order to reduce the effects of class imbalance, which is a persistent problem when the number of adverse drug reaction (ADR) occurrences is small in comparison to the number of non-events.

The adoption of gradient boosting, particularly XGBoost and LightGBM, has become a popular method for biomedical predictive tasks. This is because it constructs additive models by sequentially focusing on residuals, which enables excellent performance and the ability to handle missingness. The black-box aspect of support vector machines (SVMs), as well as the complexity of parameter adjustment, sometimes hinders interpretability and clinical acceptance. However, the potential of SVMs to work in high-dimensional feature spaces utilizing kernel tricks has led to their employment in adverse drug reaction (ADR) scenarios.

Overall, these models—even those that employ strong machine learning methods—generally necessitate meticulous feature building, as well as risk of overfitting, and may not exhibit internal transparency. Thus, many recent efforts have turned to deep learning to further push the frontier.

Deep Learning Approaches

Deep learning (DL) architectures are well-suited to capture complex hierarchical and latent patterns in large-scale data. In the ADR domain, neural networks—including feedforward deep neural networks, recurrent neural networks (RNNs), and transformer architectures—have been used to integrate heterogeneous features, model longitudinal trajectories, and process unstructured text.

Moreover, Mohsen et al. (2010) [21] combined drug-induced gene expression profiles from the Open TG-GATEs database with FAERS ADR labels to train deep neural networks for ADR prediction, achieving mean classification accuracy of 85.7% across modeled endpoints. This illustrates the power of DL in integrating multi-omics signals with pharmacovigilance data. In other work, convolutional neural networks have also been applied to encode drug and ADR embeddings for prediction in combination therapies [22].

RNNs, including long short-term memory (LSTM) networks, are particularly advantageous in modeling longitudinal EHR sequences—capturing temporal dependencies in lab values, drug changes, and patient state. Transformer architectures, recently adapted for healthcare timeseries and clinical text, offer even greater capacity for modeling long-range dependencies and attention-based weighting of features.

One recent model, ADRNet, presents a hybrid collaborative-filtering + representation-learning architecture that integrates both clinical ADR occurrence data and non-clinical drug descriptors (e.g. chemical, biological features). ADRNet uses a shallow collaborative filtering module to capture co-occurrence patterns and a deep drug representation network to guide latent embeddings, thus bridging structured and unstructured (descriptor) domains [23].

Graph neural networks (GNNs) have also gained traction. Kwak et al. (2020) [24]constructed a drug-disease graph from claims/EHR disease codes and applied GNNs to learn node embeddings predictive of drug-ADR associations; their model achieved AUROC ~0.795, demonstrating the power of relational learning.

Hybrid / Multimodal Models (Structured + Unstructured Data)

To exploit the full richness of EHRs, hybrid models that fuse structured features (labs, diagnoses, medications) and unstructured text features (symptom mentions, narrative context) have emerged. It is typical to use a two-branch architecture: one branch is an ML or DL pipeline for structured inputs, and the other is an NLP encoder (for example, BERT or CNN) for text, with fusion layers integrating latent representations later on. This enables the model to take use of both numerical precision and subtleties included in the actual narrative.

For instance, it is frequently discovered that research that employs a combination of structured and narrative features lead to improvements in sensitivity or positive predictive value when compared to structured-only models. The added information provided by clinical notes is especially useful in resolving ambiguous signals, determining the timing of adverse events, and distinguishing between incidental symptoms and those that are related to the medicine in question.

One example of a deep representation model that utilizes multimodal longitudinal EHRs is Patient2Vec. Although its major focus was on predicting hospitalizations rather than adverse drug reactions, its architecture

provides an illustration of how mixed temporal EHR data can be used to develop interpretable embeddings. Its attention mechanism allows one to gain an understanding of which visits or features have an impact on predictions. It is possible to make adjustments to alternative dispute resolution (ADR) models by using comparable design concepts [25].

Synthesis and Lessons Learned

From the diverse landscape of ADR prediction modeling, several key patterns emerge. First, there is no one-size-fits-all algorithm: logistic regression may suffice as a baseline, but more flexible ML and DL models are often required to capture EHR complexity. Second, as model complexity increases, so does the importance of rigorous validation, regularization, and interpretability methods (e.g. feature importance, attention weights). Third, hybrid and multimodal models that bring together structured, unstructured, and external drug knowledge (via embeddings or graphs) consistently show promise for improved discriminative ability. Finally, limitations in data quality, class imbalance, and generalizability remain major challenges across studies, which must be carefully addressed in any implementation or translation to clinical practice.

Model Validation and Evaluation Frameworks

In predictive modeling for adverse drug reactions (ADRs) using EHRs, building a model is only half the battle. Without rigorous validation, even a model with seemingly excellent performance in development can collapse under new data, overfit to quirks of the development set, or produce biased estimates when transported to other settings. Robust validation guards against overfitting, quantifies how well the model generalizes, and helps ensure that risk predictions remain dependable when deployed in real clinical environments [26, 27]. Indeed, the dictum "there is no such thing as a validated model" (i.e., models must continuously be tested and updated) underscores that validation is never a one-time check but a continuous scrutiny of performance across settings and time [28].

A central distinction in validation strategies lies between internal and external validation. Internal validation evaluates model performance using the same dataset (or subsets thereof) on which the model was developed, typically via cross-validation (e.g., k-fold), bootstrapping, or split-sample approaches. The aim is to correct for optimism—i.e., the tendency of a model to appear overly strong on training data due to overfitting [27, 29]. Bootstrapping is often preferred because it makes efficient use of all data and provides optimism-corrected estimates of performance. In contrast, cross-validation (e.g. 5- or 10-fold) partitions the data to iteratively test "held-out" folds, giving an estimate of reproducibility. Temporal cross-validation (e.g. training on earlier time periods, testing on later ones) is especially relevant in EHR-based studies, where data drift over time is common.

Yet internal validation alone is insufficient when the ultimate goal is deploying a model in novel settings. External validation assesses model performance in datasets entirely independent from the development cohort—whether from a different hospital, geographic region, or time period (temporal validation). Its purpose is to evaluate transportability and generalizability [30]. Good external performance suggests that the model has captured more fundamental patterns rather than idiosyncrasies of a particular site. However, as Van Calster et al. (2023) [28] and la Roi-Teeuw et al. (2024) [31] caution, an external validation in a single new dataset is never proof of universal validity; models must be assessed across multiple contexts, and calibration drift over time must be monitored. Misconceptions about external validation—treating it as a singular "seal of approval"—remain widespread [31].

When assessing model performance, several metrics become key. Discrimination—the ability to distinguish patients who will develop an ADR from those who will not—is commonly quantified by the Area Under the Receiver Operating Characteristic (AUC or c-statistic). In time-to-event settings, concordance (C-index) may serve a similar role. However, good discrimination alone is insufficient. Calibration, which measures how closely predicted probabilities align with observed event frequencies, is often called the Achilles' heel of predictive models [32]. Calibration assessment includes calibration-in-the-large (systematic over- or under-prediction), calibration slope (overfitting/underfitting), and graphical calibration plots. A model with excellent discrimination but poor calibration may mislead clinicians if its probabilities are untrustworthy. Complementing discrimination and calibration, assessing clinical utility is essential: decision curve analysis, net benefit, and threshold analysis show whether a model's predictions would lead to better decision-making or patient outcomes rather than mere statistical performance [32].

Yet even with internal and external validation and a full set of metrics, the challenge of generalizability and transportability across health systems remains formidable. Differences in patient populations, case mix, measurement protocols, medication formularies, and

practice patterns can degrade performance in new settings. For example, if one hospital uses a different laboratory analyzer that shifts reference ranges, or prescribes different drug brands, predictor distributions might shift, inducing calibration drift or reduced discrimination [33]. In addition, heterogeneity in patient features or coding practices may exacerbate these issues [34]. Further, as newer interventions, changing disease prevalence, or updated protocols emerge, the model's relationships may become outdated, requiring recalibration or retraining.

A number of techniques have been suggested in order to address these difficulties. The ability of the model to generalize across centers can be evaluated through the use of internal-external cross-validation, which involves splitting by centers or clusters and leaving one center out in each fold. According to Riley et al. (2016) [35], meta-analyses of performance across external validations can provide a quantitative measure of heterogeneity. It is essential to sustain performance as practice advances by continuously monitoring and updating dynamic models, periodically recalibrating them, and employing model updating strategies [36]. Furthermore, the idea of targeted validation emphasizes the importance of validating in the population and setting where the model is expected to be utilized as opposed to verifying in arbitrary external cohorts [29]. In addition, the validation settings, participant selection, and overall performance heterogeneity need to be reported with complete transparency, such as through TRIPOD [26].

Key Challenges and Ethical Considerations

The development and deployment of ADR predictive models using EHRs, faces significant challenges. These problems are not only technological problems, they also relate to ethical issues and are discussed in the subsequent paragraphs. However Table 1 offers actionable insights that can be followed to address these challenges.

Amongst these problems are those pertaining to EHR data quality and missing data. Clinical systems were not designed for research, therefore data gaps, errors, and inconsistent recording are common. Opportunistic laboratory tests and vital signs may be recorded, and many patients may overlook critical factors. Algorithms struggle with numerical features with high missing fractions because imputation may introduce bias or variance [37]. A recent study of missing data management in EHRs indicated that while machine learning–based imputation algorithms including multiple imputation, KNN, and matrix factorization show promise, no single approach prevails across use scenarios [38]. Coding errors—mislabeling diagnoses, drug entries, or adverse events—are widespread and degrade model signals. Poor data quality can thus propagate to flawed predictions or overconfident models.

Closely intertwined is the issue of algorithmic bias and disparities in prediction across demographic groups. Because EHR data reflect historical care patterns—often marred by existing health inequities—models trained on them may reproduce or even magnify bias [39, 40]. A canonical example outside ADR modeling is the algorithm used for risk prediction in health systems, which assigned lower risk scores to Black patients with equivalent needs, leading to less access to care [41]. In the ADR context, this could mean that minority populations or underserved groups may systematically receive underpredicted risk, thereby undermining equity in clinical decision support. Reviews of bias in EHR-based AI emphasize multiple bias types—selection bias, measurement bias, confounding bias, temporal bias—and the need for structured mitigation strategies [42, 43]. Unless addressed, algorithmic bias can compromise both fairness and the ethical legitimacy of predictive models.

Turning to governance, data privacy, governance, and the ethics of secondary EHR use must be carefully negotiated. Clinical narratives, even when de-identified, may carry re-identification risks, especially when combined across variables or with external data sources. Legislation such as GDPR in Europe, and evolving biomedical data policies, emphasize the need for transparency, consent, minimal data use, and accountability in algorithmic deployment. Scholars argue that algorithmic ethics must be proactive, not an afterthought: embedding algorithmic audits, human-in-the-loop review, and dynamic compliance monitoring increases trust [44]. In contexts with weaker data protections, secondary use of EHRs can raise concerns of exploitation, patient autonomy, and community consent—particularly when certain populations are overrepresented in datasets.

Even if one overcomes data and governance challenges, interpretability and explainability of ML and deep learning models remain a central ethical and practical hurdle. Clinicians and patients may resist a "black box" recommendation whose logic cannot be scrutinized. Interpretable machine learning (IML) methods—such as attention mechanisms, feature importance, rule extraction, surrogate models—attempt to bridge this

gap. Yet many methods struggle in balancing fidelity (how closely explanations reflect true model logic) with simplicity (how understandable they are) [45]. The healthcare domain demands trust: systems like RETAIN illustrate how reverse-time attention models can yield predictions while preserving interpretability [46]. Tools such as VBridge have also been proposed to visually connect explanatory features to patient records to help clinicians interpret model decisions [42]. Nonetheless, poor or misleading explanations may lead to distrust, misunderstanding, or misuse.

Finally, even assuming model performance and interpretability are adequate, regulatory and clinical acceptance barriers can stymie adoption. Clinicians may fear accountability for following or rejecting algorithmic advice, especially if the model is unclear. Institutional governance, medical device regulation, and approval processes vary by jurisdiction, and many regulatory frameworks are not yet ready for clinical predictive algorithmic technologies. Unless AI models prove their value and explainability, physicians may favour simpler decision criteria or conventional risk scores. The "last mile" of implementation—integrating forecasts into clinical workflows, aligning incentives, retraining staff, and monitoring—is sometimes harder than model training.

Table 2: Challenges Associated with Developing and Deploying ADR Prediction Models using EHRs

Challenge	Description	Example	Potential Mitigation Strategy
Incomplete data entry	High proportions of missing entries or variable absence	Lab test values missing in many patient records	Use robust imputation (multiple imputation, ML methods), auxiliary data, or missingness modeling
Coding errors	Misclassification of diagnoses, drugs, adverse events	ADR event coded incorrectly or omitted in ICD codes	Cross-validation against narrative notes, audit samples, NLP to detect contradictions
Population bias / disparities	Training data unrepresentative of minority or marginalized groups, leading to skewed performance	Underprediction in racial minorities	Stratified modeling, fairness-aware training, oversampling or reweighting, bias audits
Data privacy & governance	Risks to patient confidentiality and consent in secondary EHR use	Re-identification risk from combinations of variables	De-identification, federated learning, differential privacy, ethics review, transparency policies
Interpretability / explainability	Black box models whose decisions are opaque to clinicians	Deep NN giving ADR risk probability without insight	Use inherently interpretable architectures, post-hoc explanation methods, attention-based models
Clinical integration / acceptance	Resistance or logistical barriers to embedding models into practice	Clinicians ignoring alerts or refusing to trust AI output	Engage end-users early, iterative usability design, clinician–AI collaboration frameworks, education

Clinical and Research Implications

The implementation of prediction models for adverse drug reactions (ADRs) is expected to have a significant impact on the fields of pharmacovigilance, clinical decision support, personalized medicine, and the life cycle of drug development. If such models are carefully deployed and validated, they have the potential to transform the current paradigm of reactive detection of ADRs into proactive prevention, thereby improving patient safety, decreasing the burden on healthcare systems, and catalyzing innovation in the field of drug safety research.

Pharmacovigilance systems can be greatly improved by predictive ADR models. Spontaneous reporting, human case review, and statistical signal detection can lag behind real-world events in traditional pharmacovigilance. Using predictive models in surveillance pipelines can speed up ADR identification, prioritize high-risk drug–patient combinations for further assessment, and filter false positives by ranking likely true signals [47, 48]. Using real-world data, a preliminary predictive surveillance model identified product safety signals [49]. By continuously scanning EHRs, claims, and integrated sources, ADR prediction systems may catch safety signals earlier than traditional methods, thus reducing patient harm and informing regulators more swiftly.

Second, the integration of predictive models into clinical decision support systems (CDSS) is a critical path to translation. Embedding ADR risk scores or alerts into the electronic prescribing workflow can help clinicians anticipate and avoid harmful drug–patient interactions before they occur. However, to be accepted, these alerts must be clinically credible, actionable, and minimally disruptive to workflow. Models should provide not only a risk probability but also contextual justification (e.g., contributing risk factors) to support clinician understanding and trust. Furthermore, continuous feedback loops—where clinician responses feed back into model refinement—can help maintain performance in real practice settings [50].

Beyond prevention of harm, ADR prediction models can enhance personalized medicine and risk stratification. Because patients differ in genetics, comorbidities, drug metabolism, and polypharmacy, a one-size-fits-all approach to drug safety is often inadequate. Predictive models trained on large, heterogeneous EHR datasets can uncover individualized risk profiles, enabling tailored prescribing (e.g., selecting safer alternatives, adjusting doses, monitoring more intensively) for patients at high predicted risk. In this way, ADR prediction becomes a component of precision therapeutics, aligning safety measures with individual vulnerability.

Moreover, ADR prediction models have an important role across the drug development and post-marketing surveillance continuum. During drug development, models trained on preclinical, mechanistic, and early-phase human data may help anticipate safety liabilities before large trials. For instance, deep learning models combining gene expression and toxicity data have shown promise in flagging potential adverse outcomes early [21]. In post-marketing surveillance, predictive models may complement spontaneous reporting by continuously reassessing risk as new real-world data accumulate, and by identifying patient subgroups at disproportionate risk. In regulatory contexts, predictive models may support adaptive safety monitoring, dynamic labeling, or conditional risk mitigation plans.

On the other hand, these aspirations related to translation necessitate meticulous coordination. Model performance needs to be consistent across different environments, interpretability needs to meet the requirements of physicians and regulators, governance frameworks need to protect patient privacy and responsibility, and deployment needs to conform to existing clinical infrastructure. As a matter of fact, the most significant obstacles frequently encountered in practice do not originate in the building of the model but in the processes of integration, trust, maintenance, and ongoing validation.

Future Directions in ADR Prediction Research

As the concept of ADR prediction grows, the frontier is moving towards systems that are scalable, privacy-aware, multimodal, and interpretable systems that can be implemented across institutions and feature deeper biological insights. There are a number of new trajectories that show a great deal of potential and merit consideration.

Multi-center and federated learning for ADR modelling is one of the promising breakthroughs. Hospitals and health systems routinely silo adverse event data, reducing sample diversity and generalizability. Federated learning facilitates model training across sites without exchanging patient-level data, protecting privacy while learning from more populations [51]. A federated approach that pooled local models to generate a global ADR predictor by Choudhury et al. (2020)

[52] showed that decentralized learning is possible without losing data sovereignty. Recently, Guo and Choo (2025) [51] examined how federated learning and large language models might improve ADR prediction across institutional borders, particularly for unstructured clinical material. Nevertheless, federated learning in healthcare still faces thorny challenges—statistical heterogeneity (non-IID data), communication overhead, model aggregation strategies, and privacy attacks [53]. Advances in meta-learning federated frameworks (e.g., FedMetaMed) may help models better adapt to local idiosyncrasies while benefiting from shared knowledge [54].

Beyond distributed training, another rich frontier is the integration of genomics, pharmacogenomics, and multi-omics data with EHR-derived features. ADR susceptibility is often mediated by variation in drug metabolism genes, transporters, or molecular pathways; thus, combining molecular-level data with clinical trajectories can refine risk models. Recent reviews in pharmacogenomics highlight that AI and multi-omics integration can improve precision in predicting drug toxicity across contexts [55]. The fusion of longitudinal EHR time series with genotype or expression data could capture the interplay of patient biology and therapy exposure. In cancer pharmacovigilance and cardiology, hybrid models that blend omics and clinical features have begun to show incremental gains over clinical-only predictors. As sequencing becomes more routine, this direction will only strengthen.

Closely tied is the drive to design explainable AI (XAI) techniques tailored for clinical ADR prediction and adoption. While powerful predictive models are already feasible, their opaque nature poses a barrier to use. Recent surveys of XAI in medical settings emphasize that methods such as SHAP, LIME, counterfactual explanations, and attention visualization have promise but often falter in clinical fidelity and interpretability tradeoffs [56, 57]. The challenge is to tie explanations back to domain concepts—drug mechanisms, temporal relationships, lab trends—so clinicians can verify or contest predictions meaningfully. Emerging hybrid strategies that embed mechanistic knowledge or causal constraints, or that extract human-interpretable rules from deep models, may help bridge this gap. Systematic benchmarks of XAI methods in real-world ADR datasets remain scant; thus, future work should compare interpretability techniques under realistic clinical settings, not just synthetic tasks.

Another exciting horizon lies in digital twins, real-time ADR monitoring, and global EHR collaborations. The notion of a medical digital twin—a continuously updated in silico replica of a patient's physiological and therapeutic state—offers a paradigm in which ADR risk is simulated and forecast before it emerges [58]. In such frameworks, the twin may ingest EHR changes, wearable signals, pharmacokinetic/pharmacodynamic models, and molecular inputs to predict future events in silico. The twin could be used to "stress test" alternate drug regimens and flag high-risk trajectories. In practice, early efforts are underway in cardiology and metabolic disease, though ADR-specific instantiations remain nascent [59]. Closely allied is the idea of real-time ADR monitoring, where streaming EHR data or device inputs feed prediction engines to detect adverse signals as they unfold, triggering alerts or adaptive monitoring. At a broader scale, global EHR collaborations (e.g., federations of health systems sharing model weights or summary statistics) can amplify population diversity, improve external validity, and support comparative safety meta-analyses.

Given these trajectories, we propose several recommendations for the research and policy community. First, researchers should prioritize multi-site, prospective validation studies: models should be tested in live EHR settings across geographies and care systems to assess robustness and drift. Second, standardized reporting frameworks—extending TRIPOD to cover federated, multimodal, and XAI models—should be adopted, ensuring transparency of data provenance, fairness audits, and explanation methodology. Third, funding agencies and institutional review boards should support privacy-preserving data enclaves, federated infrastructure, and collaborative networks, especially in low-resource settings. Fourth, policy makers and regulators must evolve frameworks for algorithmic safety, liability, and continuous post-deployment monitoring; regulatory paradigms should treat ADR prediction models as evolving medical devices requiring lifecycle oversight. Finally, bridging the academic–clinical interface is crucial: clinicians, pharmacologists, data scientists, and informaticians should co-design models, evaluation protocols, and interpretability interfaces so that algorithmic insights are actionable, trustworthy, and aligned with workflow.

In sum, the future of ADR prediction is not simply more complex models, but smarter architectures—privacy-aware, interpretable, biologically informed, and globally generalizable. The confluence of federated learning, multi-omics integration, explainable AI, digital twin architectures, and real-time monitoring offers a roadmap to transform predictive pharmacovigilance from a hopeful concept into a standard of clinical safety.

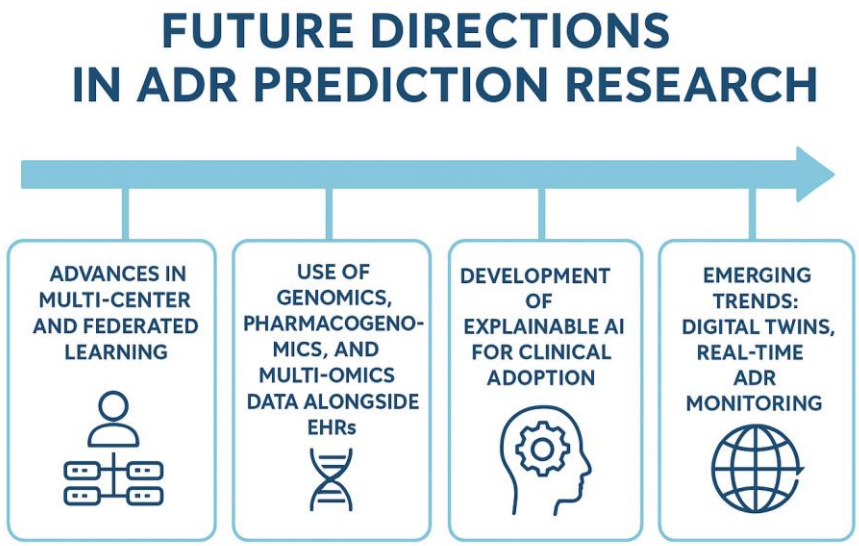

Figure 3: Future Research Roadmap for ADR Prediction Models

Conclusion

This review demonstrates that predictive modeling using electronic health records is redefining how adverse drug reactions are identified and prevented. By leveraging machine learning, deep learning, and multimodal integration, these models promise earlier detection and improved patient safety compared to traditional pharmacovigilance methods, which are hindered by underreporting and delayed signal capture. However, their potential is constrained by challenges of data quality, algorithmic bias, interpretability, and limited clinical uptake. Robust validation—across internal, external, and multi-center settings—remains essential to ensure models are reliable and generalizable in real-world use. Looking ahead, advances such as federated learning, multi-omics integration, explainable AI, and digital twin frameworks present transformative opportunities. Yet, progress must be accompanied by strong governance, transparency, and interdisciplinary collaboration. For clinicians, researchers, and policymakers alike, the path forward lies in co-designing systems that are technically sound, ethically responsible, and clinically meaningful. If achieved, predictive ADR models can evolve from experimental tools into trusted components of modern pharmacovigilance and personalized medicine.

Acknowledgments

Competing Interests: All the authors declare that they have no conflict of interest.

Ethics approval and consent to participate: Not applicable.

Consent for publication: Not applicable.

Availability of data and material: All the resources consulted in the review are provided in the reference section

Funding: No funding of any kind was received.

Authors' Contributions: This work was carried out in collaboration among all authors. All authors read and approved the final manuscript.

References

1. Wasylewicz A, van de Burgt B, Weterings A, Jessurun N, Korsten E, Egberts T, et al. Identifying adverse drug reactions from free-text electronic hospital health record notes. Br J Clin Pharmacol. 2022 Mar;88(3):1235-45. https://doi.org/10.1111/bcp.15068
2. Salas M, Petracek J, Yalamanchili P, Aimer O, Kasthuril D, Dhingra S, et al. The use of artificial intelligence in pharmacovigilance: a systematic review of the literature. Pharm Med. 2022 Oct;36(5):295-306. https://doi.org/10.1007/s40290-022-00441-z
3. Ayalew MB, Tegegn HG, Abdela OA. Drug related hospital admissions; a systematic review of the recent literatures. Bull Emerg Trauma. 2019

Oct;7(4):339. https://doi.org/10.29252/beat-070401

4. Davies EC, Green CF, Taylor S, Williamson PR, Mottram DR, Pirmohamed M. Adverse drug reactions in hospital in-patients: a prospective analysis of 3695 patient-episodes. PLoS One. 2009 Feb 11;4(2):e4439. https://doi.org/10.1371/journal.pone.0004439
5. Mouton JP, Mehta U, Parrish AG, Wilson DP, Stewart A, Njuguna CW, et al. Mortality from adverse drug reactions in adult medical inpatients at four hospitals in South Africa: a cross-sectional survey. Br J Clin Pharmacol. 2015 Oct;80(4):818-26. https://doi.org/10.1111/bcp.12567
6. Alhawassi TM, Krass I, Bajorek BV, Pont LG. A systematic review of the prevalence and risk factors for adverse drug reactions in the elderly in the acute care setting. Clin Interv Aging. 2014 Dec 1:2079-86. https://doi.org/10.2147/CIA.S71178
7. Pedrós C, Formiga F, Corbella X, Arnau JM. Adverse drug reactions leading to urgent hospital admission in an elderly population: prevalence and main features. Eur J Clin Pharmacol. 2016 Feb;72(2):219-26. https://doi.org/10.1007/s00228-015-1974-0
8. Choudhury R. Integrating EHRs for proactive pharmacovigilance. J Pharmacovigil. 2025 Apr 2;13(1):502. https://doi.org/10.35248/2329-6887.25.13.502
9. Davis SE, Zabotka L, Desai RJ, Wang SV, Maro JC, Coughlin K, et al. Use of electronic health record data for drug safety signal identification: a scoping review. Drug Saf. 2023 Aug;46(8):725-42. https://doi.org/10.1007/s40264-023-01325-0
10. Zhang F, Sun B, Diao X, Zhao W, Shu T. Prediction of adverse drug reactions based on knowledge graph embedding. BMC Med Inform Decis Mak. 2021 Feb 4;21(1):38. https://doi.org/10.1186/s12911-021-01402-3
11. Yang J, Hu Z, Zhang L, Peng B. Predicting Drugs Suspected of Causing Adverse Drug Reactions Using Graph Features and Attention Mechanisms. Pharmaceuticals. 2024 Jun 22;17(7):822. https://doi.org/10.3390/ph17070822
12. Abdeldjouad FZ, Brahami M, Sabri M. Evaluating the effectiveness of artificial intelligence in predicting adverse drug reactions among cancer patients: a systematic review and meta-analysis. In: 2023 IEEE International Conference Challenges and Innovations on TIC (I2CIT); 2023 Dec 21. p. 1-8. https://doi.org/10.48550/arXiv.2404.05762
13. Yalçın N, Kaşıkcı M, Çelik HT, Allegaert K, Demirkan K, Yiğit Ş, et al. An artificial intelligence approach to support detection of neonatal adverse drug reactions based on severity and probability scores: a new risk score as web-tool. Children. 2022 Nov 26;9(12):1826. https://doi.org/10.3390/children9121826
14. Wang W, Ferrari D, Haddon-Hill G, Curcin V. Electronic health records as source of research data. In: Machine Learning for Brain Disorders. 2023 Jul 23. p. 331-54. https://doi.org/10.1007/978-1-0716-3195-9_11
15. Hossain E, Rana R, Higgins N, Soar J, Barua PD, Pisani AR, et al. Natural language processing in electronic health records in relation to healthcare decision-making: a systematic review. Comput Biol Med. 2023 Mar 1;155:106649. https://doi.org/10.1016/j.compbiomed.2023.106649
16. Golder S, O'Connor K, Lopez-Garcia G, Tatonetti N, Gonzalez-Hernandez G. Leveraging Unstructured Data in Electronic Health Records to Detect Adverse Events from Pediatric Drug Use-A Scoping Review. medRxiv. 2025 Mar 20. https://doi.org/10.1101/2025.03.20.25324320
17. Tabari P, Costagliola G, De Rosa M, Boeker M. State-of-the-art fast healthcare interoperability resources (fhir)–based data model and structure implementations: Systematic scoping review. JMIR Med Inform. 2024 Sep 24;12(1):e58445. https://doi.org/10.2196/58445
18. Zhang F, Sun B, Diao X, Zhao W, Shu T. Prediction of adverse drug reactions based on knowledge graph embedding. BMC Med Inform Decis Mak. 2021 Feb 4;21(1):38. https://doi.org/10.1186/s12911-021-01402-3

19. Hu Q, Chen Y, Zou D, He Z, Xu T. Predicting adverse drug event using machine learning based on electronic health records: a systematic review and meta-analysis. Front Pharmacol. 2024 Nov 13;15:1497397. https://doi.org/10.3389/fphar.2024.1497397
20. Zhao YX, Yuan H, Wu Y. Prediction of adverse drug reaction using machine learning and deep learning based on an imbalanced electronic medical records dataset. In: Proceedings of the 5th international conference on medical and health informatics; 2021 May 14. p. 17-21. http://dx.doi.org/10.1145/3472813.3472817
21. Mohsen A, Tripathi LP, Mizuguchi K. Deep learning prediction of adverse drug reactions in drug discovery using open TG–GATEs and FAERS databases. Front Drug Discov. 2020 Oct 27;1:768792. https://doi.org/10.48550/arXiv.2010.05411
22. He M, Shi Y, Han F, Cai Y. Prediction of adverse drug reactions based on pharmacogenomics combination features: a preliminary study. Front Pharmacol. 2025 Mar 10;16:1448106. https://doi.org/10.3389/fphar.2025.1448106
23. Li H, Hu T, Xiong Z, Zheng C, Feng F, He X, et al. ADRNet: A generalized collaborative filtering framework combining clinical and non-clinical data for adverse drug reaction prediction. In: Proceedings of the 17th ACM Conference on Recommender Systems; 2023 Sep 14. p. 682-7. https://doi.org/10.1145/3604915.3608813
24. Kwak H, Lee M, Yoon S, Chang J, Park S, Jung K. Drug-disease graph: predicting adverse drug reaction signals via graph neural network with clinical data. In: Pacific-Asia Conference on Knowledge Discovery and Data Mining; 2020 May 6. p. 633-44. https://doi.org/10.48550/arXiv.2004.00407
25. Zhang J, Kowsari K, Harrison JH, Lobo JM, Barnes LE. Patient2vec: A personalized interpretable deep representation of the longitudinal electronic health record. IEEE Access. 2018 Oct 12;6:65333-46. https://doi.org/10.48550/arXiv.1810.04793
26. Collins GS, Dhiman P, Ma J, Schlussel MM, Archer L, Van Calster B, et al. Evaluation of clinical prediction models (part 1): from development to external validation. BMJ. 2024 Jan 8;384. https://doi.org/10.1136/bmj-2023-074819
27. Steyerberg EW, Harrell Jr FE. Prediction models need appropriate internal, internal-external, and external validation. J Clin Epidemiol. 2015 Apr 18;69:245. https://doi.org/10.1016/j.jclinepi.2015.04.005
28. Van Calster B, Steyerberg EW, Wynants L, Van Smeden M. There is no such thing as a validated prediction model. BMC Med. 2023 Feb 24;21(1):70. https://doi.org/10.1186/s12916-023-02779-w
29. Sperrin M, Riley RD, Collins GS, Martin GP. Targeted validation: validating clinical prediction models in their intended population and setting. Diagn Progn Res. 2022 Dec 22;6(1):24. https://doi.org/10.1186/s41512-022-00136-8
30. Ramspek CL, Jager KJ, Dekker FW, Zoccali C, van Diepen M. External validation of prognostic models: what, why, how, when and where?. Clin Kidney J. 2021 Jan 1;14(1):49-58. https://doi.org/10.1093/ckj/sfaa188
31. la Roi-Teeuw HM, van Royen FS, de Hond A, Zahra A, de Vries S, Bartels R, et al. Don't be misled: 3 misconceptions about external validation of clinical prediction models. J Clin Epidemiol. 2024 Aug 1;172:111387. https://doi.org/10.1016/j.jclinepi.2024.111387
32. Van Calster B, McLernon DJ, Van Smeden M, Wynants L, Steyerberg EW, Bossuyt P, et al. Calibration: the Achilles heel of predictive analytics. BMC Med. 2019 Dec 16;17(1):230. https://doi.org/10.1186/s12916-019-1466-7
33. Luijken K, Groenwold RH, Van Calster B, Steyerberg EW, van Smeden M. Impact of predictor measurement heterogeneity across settings on the performance of prediction models: A measurement error perspective. Stat Med. 2019 Aug 15;38(18):3444-59. https://doi.org/10.48550/arXiv.1806.10495
34. Navarro CL, Damen JA, van Smeden M, Takada T, Nijman SW, Dhiman P, et al.

Systematic review identifies the design and methodological conduct of studies on machine learning-based prediction models. J Clin Epidemiol. 2023 Feb 1;154:8-22. https://doi.org/10.1016/j.jclinepi.2022.11.015

35. Riley RD, Ensor J, Snell KI, Debray TP, Altman DG, Moons KG, et al. External validation of clinical prediction models using big datasets from e-health records or IPD meta-analysis: opportunities and challenges. BMJ. 2016 Jun 22;353. https://doi.org/10.1136/bmj.i3140
36. Saelmans A, Seinen T, Pera V, Markus AF, Fridgeirsson E, John LH, et al. Implementation and Updating of Clinical Prediction Models: A Systematic Review. Mayo Clin Proc Digit Health. 2025 May 23:100228. https://doi.org/10.1016/j.mcpdig.2025.100228
37. Ferri P, Romero-Garcia N, Badenes R, Lora-Pablos D, Morales TG, De La Camara AG, et al. Extremely missing numerical data in Electronic Health Records for machine learning can be managed through simple imputation methods considering informative missingness: A comparative of solutions in a COVID-19 mortality case study. Comput Methods Programs Biomed. 2023 Dec 1;242:107803. https://doi.org/10.1016/j.cmpb.2023.107803
38. Ren W, Liu Z, Wu Y, Missing Data in Electronic health Record s (MINDER) Group. Moving beyond medical statistics: a systematic review on missing data handling in electronic health records. Health Data Sci. 2024;4:0176. https://doi.org/10.34133/hds.0176
39. Cross JL, Choma MA, Onofrey JA. Bias in medical AI: Implications for clinical decision-making. PLOS Digit Health. 2024 Nov 7;3(11):e0000651. https://doi.org/10.1371/journal.pdig.0000651
40. Hanna MG, Pantanowitz L, Jackson B, Palmer O, Visweswaran S, Pantanowitz J, et al. Ethical and bias considerations in artificial intelligence/machine learning. Mod Pathol. 2025 Mar 1;38(3):100686. https://doi.org/10.1016/j.modpat.2024.100686
41. Obermeyer Z, Powers B, Vogeli C, Mullainathan S. Dissecting racial bias in an algorithm used to manage the health of populations. Science. 2019 Oct 25;366(6464):447-53. https://doi.org/10.1126/science.aax2342
42. Cheng F, Liu D, Du F, Lin Y, Zytek A, Li H, et al. Vbridge: Connecting the dots between features and data to explain healthcare models. IEEE Trans Vis Comput Graph. 2021 Oct 1;28(1):378-88. https://doi.org/10.48550/arXiv.2108.02550
43. Hofmann B. Biases in AI: acknowledging and addressing the inevitable ethical issues. Front Digit Health. 2025 Aug 20;7:1614105. https://doi.org/10.3389/fdgth.2025.1614105
44. Liaw ST, Liyanage H, Kuziemsky C, Terry AL, Schreiber R, Jonnagaddala J, et al. Ethical use of electronic health record data and artificial intelligence: recommendations of the primary care informatics working group of the international medical informatics association. Yearb Med Inform. 2020 Aug;29(01):051-7. https://doi.org/10.1055/s-0040-1701980
45. Abdullah TA, Zahid MS, Ali W. A review of interpretable ML in healthcare: taxonomy, applications, challenges, and future directions. Symmetry. 2021 Dec 17;13(12):2439. https://doi.org/10.3390/sym13122439
46. Choi E, Bahadori MT, Sun J, Kulas J, Schuetz A, Stewart W. Retain: An interpretable predictive model for healthcare using reverse time attention mechanism. Adv Neural Inf Process Syst. 2016;29. https://doi.org/10.48550/arXiv.1608.05745
47. Algarvio RC, Conceição J, Rodrigues PP, Ribeiro I, Ferreira-da-Silva R. Artificial intelligence in pharmacovigilance: a narrative review and practical experience with an expert-defined Bayesian network tool. Int J Clin Pharm. 2025 Jul 30:1-3. https://doi.org/10.1007/s11096-025-01975-3
48. Dsouza VS, Leyens L, Kurian JR, Brand A, Brand H. Artificial Intelligence in Pharmacovigilance: A Systematic Review on Predicting Adverse Drug Reactions in Hospitalized Patients. Res Social Adm Pharm. 2025 Feb 12. https://doi.org/10.1016/j.sapharm.2025.02.008

49. De Abreu Ferreira R, Zhong S, Moureaud C, Le MT, Rothstein A, Li X, et al. A pilot, predictive surveillance model in pharmacovigilance using machine learning approaches. Adv Ther. 2024 Jun;41(6):2435-45. https://doi.org/10.1007/s12325-024-02870-5
50. Dimitsaki S, Natsiavas P, Jaulent MC. Applying AI to structured real-world data for pharmacovigilance purposes: Scoping review. J Med Internet Res. 2024 Dec 30;26:e57824. https://doi.org/10.2196/57824
51. Guo D, Choo KK. Applications of Federated Large Language Model for Adverse Drug Reactions Prediction: Scoping Review. J Med Internet Res. 2025 Sep 8;27:e68291. https://doi.org/10.2196/68291
52. Choudhury O, Park Y, Salonidis T, Gkoulalas-Divanis A, Sylla I, k Das A. Predicting adverse drug reactions on distributed health data using federated learning. AMIA Annu Symp Proc. 2020 Mar 4;2019:313.
53. Ali MS, Ahsan MM, Tasnim L, Afrin S, Biswas K, Hossain MM, et al. Federated Learning in Healthcare: Model Misconducts, Security, Challenges, Applications, and Future Research Directions--A Systematic Review. arXiv preprint arXiv:2405.13832. 2024 May 22. https://doi.org/10.48550/arXiv.2405.13832
54. Gao J, Li Y. Fedmetamed: Federated meta-learning for personalized medication in distributed healthcare systems. In: 2024 IEEE International Conference on Bioinformatics and Biomedicine (BIBM); 2024 Dec 3. p. 6384-91. https://doi.org/10.48550/arXiv.2412.03851
55. Zack M, Stupichev DN, Moore AJ, Slobodchikov ID, Sokolov DG, Trifonov IF, et al. AI and Multi-Omics in Pharmacogenomics: A New Era of Precision Medicine. Mayo Clin Proc Digit Health. 2025 Jun 26:100246. https://doi.org/10.1016/j.mcpdig.2025.100246
56. Alkhanbouli R, Matar Abdulla Almadhaani H, Alhosani F, Simsekler MC. The role of explainable artificial intelligence in disease prediction: a systematic literature review and future research directions. BMC Med Inform Decis Mak. 2025 Mar 4;25(1):110. https://doi.org/10.1186/s12911-025-02944-6
57. Sun Q, Akman A, Schuller BW. Explainable artificial intelligence for medical applications: A review. ACM Trans Comput Healthc. 2025 Feb 17;6(2):1-31. https://doi.org/10.48550/arXiv.2412.01829
58. Li X, Loscalzo J, Mahmud AF, Aly DM, Rzhetsky A, Zitnik M, et al. Digital twins as global learning health and disease models for preventive and personalized medicine. Genome Med. 2025 Feb 7;17(1):11. https://doi.org/10.1186/s13073-025-01435-7
59. Sadée C, Testa S, Barba T, Hartmann K, Schuessler M, Thieme A, et al. Medical digital twins: enabling precision medicine and medical artificial intelligence. Lancet Digit Health. 2025 Jun 14. https://doi.org/10.1016/j.landig.2025.02.004