Drug decision support systems: A review of challenges, advances, and emerging trends

Jamilu BELLO[1*], Adepeju Abeke ADIGUN[1], Patrick OZOH[1], Michael Olugbenga ABOLARINWA[2], Ibrahim Olufemi ADIGUN[3]

[1] Department of Computer Science, Faculty of Computing and Information Technology, Osun State University, Osogbo, Osun State, Nigeria

jamilu.bello@pgc.uniosun.edu.ng (*Corresponding author), adepeju.adigun@uniosun.edu.ng, patrick.ozoh@uniosun.edu.ng.

[2] Department of Cybersecurity, Faculty of Computing and Information Technology, Osun State University, Osogbo, Osun State, Nigeria

michael.abolarinwa@uniosun.edu.ng,

[3] Department of Electrical Engineering, Faculty of Electrical Engineering, Technical University of Applied Sciences Würzburg-Schweinfurt, Schweinfurt, Germany

olufemiibrahim@gmail.com.

Abstract: Background: The integration of Drug Decision Support Systems (DDSS) is a foundational element of modern medical practice designed to mitigate the global burden of adverse drug events. Driven by rapid advances in genomics, artificial intelligence (AI), and computational modeling, DDSS architectures have transitioned from primitive, rule-based engines into adaptive, data-driven frameworks engineered to support precision medicine.

Objective: This scoping review systematically synthesizes the methodological and architectural milestones of this technological transition, mapping the evolution from deterministic rules to advanced probabilistic and neural hybrid systems.

Methodology: Surveying core literature sources across PubMed, Scopus, and IEEE Xplore for the period spanning January 2012 to March 2024, this review systematically appraised and categorized the computational architectures of relevant peer-reviewed studies.

Results: The review highlights a structural shift toward probabilistic inference methods, evaluating Viterbi decoding to reduce noise in clinical sequence data, alongside deep learning models like Graph Neural Networks (GNNs) that refine drug-drug interaction (DDI) predictions. However, primary findings indicate that widespread clinical adoption is stalled by a pervasive reliance on retrospective simulation studies rather than prospective real-world trials, severe multi-omics data heterogeneity, and stark equity gaps due to minority genome underrepresentation in training sets.

Conclusion: Ultimately, realizing the full potential of personalized pharmacotherapy relies on resolving these structural deficits through federated learning infrastructures and transparent algorithmic governance. Achieving this clinical integration requires the deployment of highly interpretable, Explainable Artificial Intelligence (XAI) models capable of seamless execution within existing electronic health records.

Keywords: Clinical Decision Support, Explainable AI, Drug–Drug interactions, Medication prescribing, Personalized medicine.

1. Introduction

Avoidable morbidity, mortality and unnecessary health care costs are still a large part of the problem of adverse drug events (ADEs) and medication errors worldwide. ADEs have been found to occur in both inpatient and outpatient settings in the literature through surveillance (Bright et al., 2012; Sutton et al., 2020), and a significant proportion of these are thought to be due to incorrect prescribing. The likelihood of unexpected adverse drug reactions (ADRs) and the increased complexity of polypharmacy, the genetic variability among patients, and the rise in chronic disease prevalence all increase the risk of encountering unexpected outcomes in healthcare, further emphasizing the need for the use of decision-support technologies to help clinicians at the point of care (Chen et al., 2023).

Drug Decision Support Systems (DDSS) are unique engines, originally rule-based, designed

https://doi.org/10.33436/v36i2y202605

to be integrated into Computerised Physician Order Entry Systems (CPOES). These early systems were focused on drug-drug interactions (DDIs), dosing issues, contraindications, and allergy alerts (Bright et al., 2012). While these systems are interpretable and relatively easy to deploy, they are limited in their adaptability to patient-specific features, their number of false positives, and their fixed rule sets (Sutton et al., 2020; Amann et al., 2022).

In the last 20 years, especially from 2021-2025, DDSS research has focused on patient-centred, data-driven models. An important facilitator, pharmacogenomics, has been shown to be associated with variations in drug metabolism and treatment outcomes (Wake et al., 2022; Blagec et al., 2022). With the advent of next-generation sequencing (NGS) and saliva-based genotyping, as well as the incorporation of pharmacogenomic databases (e.g., PharmGKB), DDSS that recommend drug use based on a patient's genetic makeup have been developed (Wishart et al., 2018; Chen et al., 2023). Concurrently, developments in artificial intelligence (AI), machine learning (ML) and computational modeling have enabled novel predictive approaches for DDIs, drug efficacy prediction, and the uncertainty of missing, incomplete, or noisy data (Ryu et al., 2018; Deng et al., 2020; Lin et al., 2023; Luo et al., 2024).

Several methodological directions have been identified in the recent literature. Graph Neural Network (GNN) and knowledge-graph methods are used to model the associations among drugs, targets and phenotypes using heterogenous biomedical data (Ma et al., 2023; Chen et al., 2024; Wang et al., 2024). It is now acknowledged that interpretability and predictive accuracy are increasingly considered trade-offs in the design of architectures, as one of the major hurdles to translating AI systems into clinical use (Amann et al., 2022; Labkoff et al., 2024).

2. Methodology

This review was undertaken using the PRISMA-ScR guidelines to ensure a rigorous and transparent literature synthesis and to ensure that the review is repeatable and reproducible.

A comprehensive literature search was conducted in PubMed, Scopus, and IEEE Xplore. The search terms used included Boolean logic that combined terms like ("Drug Decision Support Systems" OR "DDSS" OR "Clinical Decision Support") AND ("Pharmacogenomics" OR "Machine Learning" OR "Artificial Intelligence").

The literature search was conducted from January 2012 to March 2024. Studies published prior to 2012 were included to provide background for the research, and the main focus was on those published from 2021-24 to show recent developments in the field. For the purposes of this study, only papers published in peer-reviewed journals in the English language and that explicitly proposed, implemented, and/or evaluated a DDSS architecture were included. Studies not employing formal computational or clinical methodology, as well as editorials that were not drug-related, were not included.

Studies were critically appraised to determine their strength of evidence and ensure that only studies with strong evidence were included. The review actively classifies the cited works into sections of Randomized Controlled Trials (RCTs), retrospective observational studies, and *in silico* simulations, to appropriately evaluate the clinical applicability and comparative validation for each of the proposed technologies.

3. Literature review

3.1. Background of drug decision support systems

3.1.1. Rule-based systems

Traditionally, drug decision support systems (DDSS) are designed and developed using optimised academic principles and best practices in clinical care, with alerts being generated at the point of prescription, based on expert knowledge and rules. They are naturally clear and transparent, and can be traced back to authoritative clinical guidelines (Bright et al., 2012; Sutton et al., 2020),

but their comparative performance is less solid in more complex contexts. Such systems are evaluated primarily by retrospective observational studies as opposed to prospective Randomized Controlled Trials (RCTs) and are inadequate for learning genotype-phenotype associations or reacting to multifactorial drug-drug interactions (Amann et al., 2022). Another major drawback is that they often fail to perform well across different clinical populations, leading to a high false-positive rate (in observational datasets, this can be as high as 80% to 90%). This invariably results in "alert fatigue" on the part of the clinician, thereby compromising their clinical use (Dullabh et al., 2022; Labkoff et al., 2024).

3.1.2. Data-driven systems

As the volume of electronic health records (EHRs), pharmacogenomic databases, and multi-omics data grows, data-driven DDSS architectures have arisen. Clinical biomarkers are used to predict drug reactions, with statistical and machine-learning approaches, including single-nucleotide polymorphisms (SNPs) and multimodal clinical biomarkers (Ryu et al., 2018; Deng et al., 2020; Chen et al., 2023). However, an important limitation of these systems is that they rely on retrospective observational data or *in silico* simulations rather than RCTs. There are several impressive *in silico* examples of AUC scores exceeding 0.85 for graph-based neural networks (GNNs) and knowledge graph-based methods (KGs) modeling complex drug-target-disease relationships (Ma et al., 2023; Chen et al., 2024); however, the applicability of these methods in clinical settings has not yet been definitively validated. The main limitation of these models is their susceptibility to dataset bias and the lack of RCTs proving the effectiveness of bedside models.

3.1.3. Probabilistic frameworks

Hidden Markov Models (HMMs) have been used in many applications related to pharmacogenomics, such as probabilistic inference and variant-detection pipelines. Probabilistic inference methods have been used to augment the hypothesis space in the presence of uncertainty and noise. Enhancing sensitivity and generalization in genomic applications is continually explored with techniques like beam search decoding and neural network-augmented emission probabilities, which have shown quantitative sensitivity improvements of 10–15% over baseline algorithms in *in silico* simulations (Lin et al., 2023; Luo et al., 2024). The idea behind these hybrid systems is that they aim to benefit from the modelling capabilities of an estimate-based system while retaining the flexibility of a purely rule-based system, without giving up the rigid limitations of the latter. They are not, however, easy to compute at the bedside and, at present, lack prospective RCT validation to justify the costs of clinical infrastructure (Amann et al., 2022; Labkoff et al., 2024).

3.1.4. Network-based and knowledge graphs

To express complex biomedical relationships among medicines, targets, and genetic diversity, knowledge graph frameworks have emerged. Graph databases and GNNs are powerful because they can model higher-order relationships and interactions, rather than just two-way interactions, including those known as DDIs (Chen et al., 2023; Wang et al., 2024). Such systems are also shown to be quantitatively superior to relational systems, with precision improvements of between 12-20% in predicting polypharmacy risks in simulation experiments (Ma et al., 2023; Luo et al., 2024). Their major constraint, however, is the need to rely on large, carefully constructed biomedical ontologies that are not often found in diverse hospital EHRs. There is no real-world evidence from RCTs or any sufficiently strong observational studies to support the widespread clinical deployment of these models at the point of care.

3.1.5. Probabilistic neural hybrid architectures

Probabilistic inference is being increasingly combined with deep learning in modern applications. For instance, uncertainty in DDIs and predictions of adverse drug events are increasingly estimated using probabilistic graphical models and deep embeddings, evaluated primarily in retrospective observational studies (Luo et al., 2024; Lin et al., 2023). These hybrid methods aim to strike a balance between predictive capability and interpretability for regulatory purposes. However, a major limitation is the inherent trade-off: marginal AUC improvements (often <5% over simpler models) are frequently offset by drastic reductions in interpretability. There are

virtually no RCTs directly comparing these advanced approaches with established clinical baselines.

3.1.6. Scalable and cloud-based infrastructures

The ability to map through massive genomic and phenotypic data becomes a bottleneck as DDSS integrates the dataset. Distributed computation and real-time updates have been suggested to be enabled by cloud-based infrastructures and microservices (Ryu et al., 2018). Theoretically, federated learning can be used to train predictive models in decentralized nodes without sharing sensitive patient information. However, validated primarily through *in silico* networking simulations rather than clinical RCTs, these infrastructures face significant limitations, including network reliability variations and the logistical hurdle of integrating the learning models into fragmented, legacy health systems (Chen et al., 2024; Wang et al., 2024).

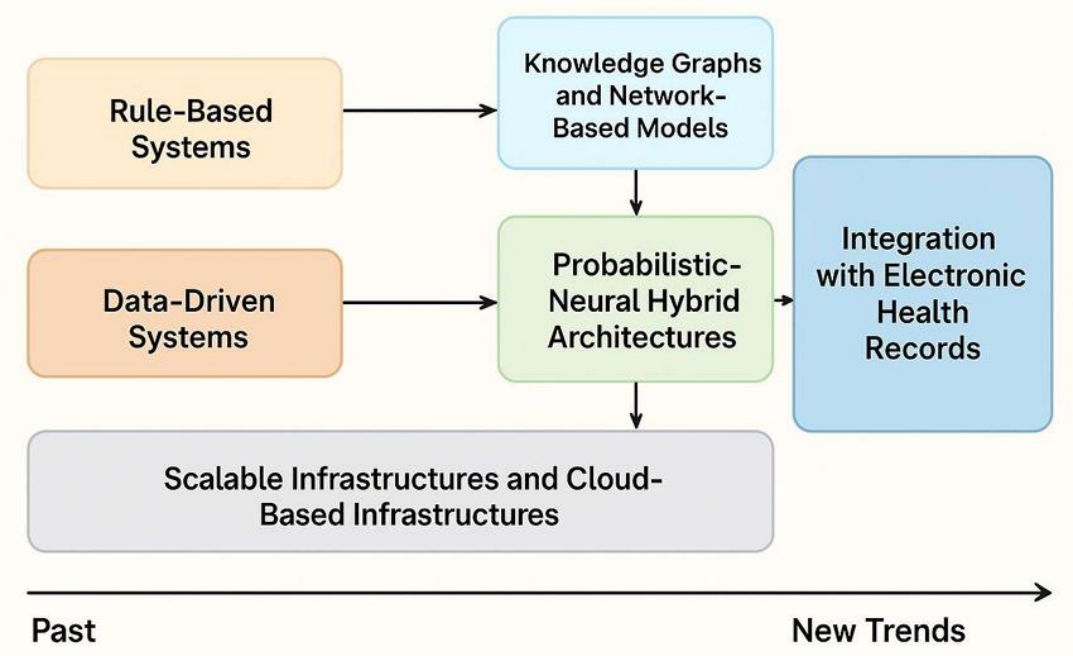

Figure 1. Drug decision support systems history (Source: Author's own).

3.2 The architecture trends and real-world applications

One of the most significant architectural trends is the extensive embedding of DDSS within electronic health records (EHRs). Theoretically, the decision support embedded in the prescribing workflow decreases cognitive load (Sutton et al., 2020; Dullabh et al., 2022). Embedded pharmacogenomic (PGx) modules have been developed for commercial vendors such as Epic and Cerner for dose and DDI alerts (Blagec et al., 2022; Labkoff et al., 2024). However, interoperability problems arise during deployment; the variety of EHR standards often makes it difficult to process them in real time for effective performance (Wake et al., 2022; Ma et al., 2023).

3.3 Reflections on practical implementation

Real-world deployments provide the best understanding of the capabilities and limitations of DDSS. In the case of some tightly managed clinical trials, there is a reduction in adverse drug events and guideline adherence (Bright et al., 2012; Dullabh et al., 2022), but there is limited evidence that this can be generalized. Workflow disruption, poor generalizability in the out-of-sample population, and alert fatigue driven by exhausted false-positive rates severely limit the large-scale use of these models (Amann et al., 2022; Labkoff et al., 2024). Socio-technical issues, such as clinicians' trust, usability, and system readiness, influence implementation success, which cannot be achieved by technical accuracy alone (Wake et al., 2022; Chen et al., 2024).

3.4 Comparison with previous reviews

The widespread change in Clinical Decision Support Systems (CDSS) has been well documented, but little research has considered Drug Decision Support Systems (DDSS) from a broad conceptual or monolithic perspective, or from a narrow, single-technology focus. Bright et al. (2012) and Sutton et al. (2020) conducted previous systematic reviews; these reviews concentrated on rule-based, deterministic systems, with a strong emphasis on simple integration into a basic electronic health record (EHR), usability measures, and the ubiquitous problem of clinician alert fatigue. For both, the discussion shifted to a more recent focus on the ethical ramifications of "black-box" AI systems and the utility of PGx as a standalone clinical tool, respectively (Amann et al., 2022; Wake et al., 2022).

However, there are also shortcomings in previous research: these studies often lack a solid, reproducible methodological approach (e.g., PRISMA-ScR), and they do not provide detailed analysis of the technologies, instead describing them in general terms without dissecting the underlying computational architectures. Moreover, many previous reviews combine theoretical and in silico computational achievements with clinical readiness at the point of care (PoC).

The present work has three main contributions:

i. This paper looks at the architectural evolution of DDSS, as opposed to the previous reviews that solely examined clinical results, explicitly outlining the development of the computational models, from the simple rule-based engines to the more complex probabilistic-neural hybrid models (such as HMMs with Viterbi decoding and Graph Neural Networks).

ii. The undertakes a serious examination of the evidence base. It does not assume that the algorithm is accurate but instead systematically differentiates between technologies validated using retrospective datasets and those tested in prospective Randomized Controlled Trials (RCTs).

iii. This combines concepts from structural biology (e.g., AlphaFold2 protein modeling) with the integration of multiple OMICS streams, bridging theoretical and clinical practices. It uniquely identifies data heterogeneity and genomic equity as the fundamental challenges to architectural scalability and provides a new taxonomy of emerging DDSS trends based on data characteristics.

4. Computer systems of drug DSS

4.1. Deep Learning and Machine Learning

The predictive modeling capabilities of deep learning (DL) and machine learning (ML) have extended beyond simple rule-based logic. In response to the challenge of modeling therapeutic responses with multi-omics information, convolutional neural networks (Chen et al., 2023) and recurrent neural networks (Wang et al., 2024) are often suggested. Theoretically, severity should be stratified using integrated features, and retrospective observational studies frequently report impressive predictive accuracies (e.g., F1-scores > 0.88). However, most of these models lack strong, specific clinical validation via RCTs in target populations. A critical limitation is that ensemble algorithms remain "black-box" models. While they achieve high performance in simulations, they require explainable AI techniques to meet regulatory requirements and gain clinician acceptance in real-world deployment (Amann et al., 2022; Lin et al., 2023).

4.2. AI-Based variant detection

Incorporation of PGx requires the correct detection of genetic variations by DDSS. AI-powered variant callers (e.g., DeepVariant, Clair3) perform exceptionally well in silico (often >99% precision) at low sequencing coverage on benchmark datasets. These algorithms can be used within the theoretical pipeline of DDSS to adjust dosing parameters (Luo et al., 2024). However, a

major limitation is that they are heavily dependent on benchmark datasets such as Genome in a Bottle (GIAB); consequently, observational studies applying these models to underrepresented genetic groups demonstrate a marked reduction in sensitivity. They have yet to be validated by prospective RCTs for equitable, real-world PGx bedside pipelines (Wake et al., 2022; Chen et al., 2024).

4.3. High-performance workflows

Bedside utility heavily limits computational scalability. High-Performance Computing (HPC) workflows (e.g., HPC-GVCW) partition the task of detecting subunit variants, thereby reducing the time required by 30-50% in controlled simulations (Deng et al., 2020; Ryu et al., 2018). While these systems are highly efficient at variant discovery, their critical limitation is their intense computational resource consumption. Validations remain entirely in silico, with a notable gap in RCTs or observational evidence demonstrating their practical translation to mapping variants to corresponding drug responses, which severely restricts their day-to-day clinical utility (Blagec et al., 2022; Labkoff et al., 2024).

5. Pharmacogenomics Decision Support Systems

Pharmacogenomics (PGx) transforms DDSS from a rule-of-thumb approach to precision medicine by matching genetic variations with individual differences in drug efficacy and safety. There are platforms that offer curated guidelines, such as PharmGKB and the Clinical Pharmacogenetics Implementation Consortium (CPIC), which focus on actionable gene-drug interactions and allow the use of genomic knowledge in a formal way within the DDSS logic (Ryu et al., 2018; Blagec et al., 2022; Wake et al., 2022).

Recent literature shows that PGx-guided DDSS is useful. In certain patient populations, patient-centered profiles have been demonstrated to reduce adverse effects and improve adherence (Chen et al., 2024). Another area where PGx integration has potential is polypharmacy analysis, which involves simultaneously analyzing multiple medications that can interact in various ways (Labkoff et al., 2024; Wang et al., 2024).

The next frontier, and one that is coming closer, is the use of advances in the field of computational structural biology. Using AlphaFold2, proteins can be modeled to provide high-resolution information on protein-drug interactions (Chen et al., 2024). It is, however, important to emphasize that no peer-reviewed literature currently describes or validates the integration of AlphaFold2 into clinical DDSS pipelines; it remains theoretical and without clinical validation.

However, the current implementation of PGx-DDSS in clinical practice is hindered by significant systemic challenges, inconsistent data across populations, pronounced disparities in minority representation in PGx data, and integration issues with legacy EHRs (Amann et al., 2022; Dullabh et al., 2022).

6. Challenges and future forecasts

Retrospective observational data and simulated environments are a common drawback of the literature analyzed in this article. Many advanced DDSS architectures are difficult to evaluate based on the clinical impact due to the absence of prospective Randomized Controlled Trials (RCTs).

6.1. Explain ability and clinician trust

Although complex ML/Deep Learning models have improved predictive performance, they remain “black-box” models, hindering clinical uptake (Amann et al., 2022; Lin et al., 2023). Obscure recommendations are less likely to be accepted by clinicians, especially for high-risk drugs. While conceptually proposed techniques have been introduced for XAI (Chen et al., 2024; Wang et al., 2024), there are essentially no validated real-world prescribing environmental RCTs for these techniques.

6.2. Data heterogeneity and integration

Multi-source integration (Dullabh et al., 2022; Wake et al., 2022) is complicated by the structural variability of EHRs, the pipelines used to annotate genes, and incomplete datasets. Theoretically, multi-omics data can provide more comprehensive models, but it is technically challenging to apply them on at a scalable and standardized level (Ryu et al., 2018; Deng et al., 2020).

6.3. Computational scalability and workflow

The influx of whole-genome data outpaces the capacity of standard hospital computing infrastructure. Furthermore, technical functionality does not guarantee clinical success. Poorly integrated systems generate alert fatigue, repetitive recommendations, and workflow distractions that actively harm clinical efficiency (Bright et al., 2012; Labkoff et al., 2024). Mandating interoperable standards (FHIR, HL7) is a strict prerequisite for future adoption (Sutton et al., 2020; Blagec et al., 2022).

6.4. Ethical, equity, and governance

Data bias can greatly affect the fairness of DDSS algorithms. The current evidence for PGx largely comes from populations of European ancestry; therefore, algorithms derived from this evidence may be less effective, or even harmful, for other populations worldwide (Amann et al., 2022). Patient privacy, data sovereignty, and algorithmic accountability are largely unaddressed regulatory challenges (Dullabh et al., 2022; Chen et al., 2024).

7. Emerging directions

Future progress depends on establishing a link between theory and proven clinical evidence. Hybrid architectures that combine probabilistic inference with neural networks hold promise for achieving a balance between interpretability and accuracy. Structural biology tools represent new avenues for understanding drug-protein binding (Wang et al., 2024), but they need to move from *in silico* experiments to clinically validated pipelines. To achieve the vision of precision medicine at scale, cloud-based, real-time DDSS must be integrated into EHRs (Blagec et al., 2022; Luo et al., 2024) and guided by population-inclusive data and clear governance. Table 1 presents an integrated analytical synthesis of these emerging trends, summarizing their maturity levels, evidence bases, and inherent limitations.

Table 1. Analytical synthesis of emerging trends in Drug Decision Support Systems (DDSS)

Trend	Core Features	Maturity Level	Evidence Base	Known Limitations & Challenges	Representative References
Integration with EHRs	Embedding DDSS into workflows (Epic, Cerner); specific dose/ interaction alerts	Clinically Deployed (in select systems)	Observational & limited RCTs	Interoperability issues, non-standard EHR data, workflow disruption, alert fatigue	Sutton et al. (2020); Dullabh et al. (2022); Blagec et al. (2022); Labkoff et al. (2024)
Knowledge Graphs & Network Models	Representing drug–target–disease relationships using GNNs	Research Prototype	Computational Simulations	Requires massive, high-quality ontologies; severely lacks prospective clinical validation	Ma et al. (2023); Chen et al. (2024); Wang et al. (2024); Luo et al. (2024)
Probabilistic–Neural	Blending HMM princi-	Research Prototype	Computational Simulations	Computationally intensive;	Amann et al. (2022)

Hybrids	ples with neural networks for robust sequence analysis			ongoing trade-off between interpretability and accuracy	
Scalable Cloud & Federated Learning	Distributed computation, privacy-preserving model training across nodes	Pilot / Research	Simulations & Retrospective Data	High implementation costs, variable network reliability, fragmented legacy systems	Ryu et al. (2018); Deng et al. (2020); Chen et al. (2024); Wang et al. (2024)
AI-Based Variant Detection	Deep learning variant callers (DeepVariant) intended for PGx pipelines	Mature Tool (Standalone), Prototype (DDSS Integration)	Retrospective Data	Heavily dependent on specific benchmark datasets; poor generalization to diverse genomes	Wake et al. (2022)
High-Performance Computing (HPC)	Parallelized genotype calling optimized for whole-genome data scaling	Research Prototype	Simulations	Highly resource-intensive; focuses on variant discovery rather than clinical drug response	Deng et al. (2020); Ryu et al. (2018)
Explainable AI (XAI)	Attention-based, surrogate, and causal inference methods for transparency	Research Prototype	Simulations	Extremely limited validation in real-world prescribing contexts; "black-box" stigma remains	Amann et al. (2022); Lin et al. (2023); Chen et al. (2024); Wang et al. (2024)
Equity & Governance Frameworks	Fairness-aware algorithms, inclusive PGx data curation, open governance	Conceptual / Early Policy	N/A (Policy/ Ethics Reviews)	Severe underrepresentation of minority genomes limits global clinical application; governance complexity	Dullabh et al. (2022); Chen et al. (2024)
Protein Structure–Informed DDSS	Theoretical integration of structural prediction with PGx variant data	Purely Theoretical	In silico concepts	No concrete integration demonstrated; requires immense computational power and data fusion	Wang et al. (2024); Chen et al. (2024)

8. Conclusion

Drug Decision Support Systems (DDSS) have been revolutionized with the incorporation of cutting-edge computing, biomedical informatics, and deep-learning technologies. Modern DDSS architectures enable highly accurate data-driven prescribing, thanks to genomics, clinical, and multi-omics data. Despite these advances, there are substantial challenges to widespread clinical use, including problems with algorithms interpretability, data heterogeneity, computational scalability, workflow interoperability, and ethical governance. New AI paradigms such as Explainable Artificial Intelligence (XAI), federated learning architectures, and the incorporation of structural biology and protein modeling hold great potential for addressing these limitations. Moreover, future studies need to focus on population-based datasets, fairness-aware algorithms and transparent governance mechanisms to guarantee fair clinical outcomes. In conclusion, the effective implementation of DDSS will require careful consideration of computational complexity, clinical relevance, and ethical considerations. Achieving this balance will set the foundation for precision medicine, and revolutionize the delivery of safe, effective and truly personalized pharmacotherapy around the world.

Author contributions

Conceptualization: J. B. and A. A. A.; Methodology: J. B., A. A. A. and P. O.; Software: J. B., A. A. A. and M. O. A.; Validation: P. O, M. O. A. and I. O. A.; Formal analysis: J. B., A. A. A. and I. O. A.; Investigation: M. O. A. and I. O. A.; Writing, original draft preparation: J. B. and A. A. A.; Writing, review and editing: P. O., M. O. A. and I. O. A.; Visualization: J. B., A. A. A. and M. O. A. All authors have read and agreed to the published version of the manuscript.

Acknowledgements

We would like to thank the Molecular and Genetic Laboratory, Fountain University, Osogbo, Osun State, Nigeria.

Submission Received: 9 March 2026; Revised: 20 May 2026; Accepted: 29 May 2026; Published: 30 June 2026.

REFERENCES

Amann, J., Vetter, D., Blomberg, S. N. et al. (2022) To explain or not to explain? — Artificial Intelligence explainability in clinical decision support systems. *PLOS Digital Healt.* 1(2), e0000016. https://doi.org/10.1371/journal.pdig.0000016.

Blagec, K., Swen, J. J., Koopmann, R. et al. (2022) Pharmacogenomics decision support in the U-PGx project: Results and advice from clinical implementation across seven European countries. *PLOS ONE*. 17(6), e0268534. https://doi.org/10.1371/journal.pone.0268534.

Bright, T. J., Wong, A., Dhurjati, R. et al. (2012) Effect of clinical decision-support systems: A systematic review. *Annals of Internal Medicine*. 157(1), 29–43. https://doi.org/10.7326/0003-4819-157-1-201207030-00450.

Chen, Y., Chen, J., Zou, Q. et al. (2024) DrugDAGT: A dual-attention graph transformer with contrastive learning improves drug–drug interaction prediction. *BMC Biology*. 22, 233. https://doi.org/10.1186/s12915-024-02030-9.

Chen, Z., Liang, N., Zhang, H. et al. (2023). Harnessing the power of clinical decision support systems: challenges and opportunities. *Open Heart*. 10(2), e002432. https://doi.org/10.1136/openhrt-2023-002432.

Deng, Y., Xu, X., Qiu, Y. et al. (2020) A multimodal deep learning framework for predicting drug-drug interaction events. *Bioinformatics*. 36(15), 4316–4322. https://doi.org/10.1093/bioinformatics/btaa501.

Dullabh, P., Heaney-Huls, K., Lobach, D. F., et al. (2022) The technical landscape for patient-centered CDS: progress, gaps, and challenges. *Journal of the American Medical Informatics Association*. 29, 1101–1105. https://doi.org/10.1093/jamia/ocac029.

Labkoff, S., Oladimeji, B., Kannry, J. et al. (2024) Toward a responsible future: recommendations for AI-enabled clinical decision support. *Journal of the American Medical Informatics Association (JAMIA)*. 31(11), 2730–2739. https://doi.org/10.1093/jamia/ocae209.

Lin, X., Dai, L., Zhou, Y. et al. (2023) Comprehensive evaluation of deep and graph learning on drug–drug interactions prediction. *Briefings in Bioinformatics*. 24(4), bbad235. https://doi.org/10.1093/bib/bbad235.

Luo, H., Yin, W., Wang, J. et al. (2024) Drug–drug interactions prediction based on deep learning and knowledge graph: A review. *iScience*, 27, 109148. https://doi.org/10.1016/j.isci.2024.109148.

Ma, M., Lei, X., et al. (2023) A dual graph neural network for drug–drug interactions prediction based on molecular structure and interactions. *PLOS Computational Biology* 19(1), e1010812. https://doi.org/10.1371/journal.pcbi.1010812.

Ryu, J. Y., Kim, H. U. & Lee, S. Y. (2018) Deep learning improves prediction of drug–drug and drug–food interactions. *Proceedings of the National Academy of Sciences of the United States of America*. 115(18), E4304–E4311. https://doi.org/10.1073/pnas.1803294115.

Sutton, R. T., Pincock, D., Baumgart, D. C. et al. (2020) An overview of clinical decision support systems: benefits, risks, and strategies for success. npj Digital Medicine, 3(17). https://doi.org/10.1038/s41746-020-0221-y.

Wake, D. T., Dunnenberger, H. et al. (2022) Pharmacogenomic clinical decision support: A review, how-to guide, and future vision. *Clinical Pharmacology & Therapeutics*. 113(4), 803–815. https://doi.org/10.1002/cpt.2387.

Wang, Y., Yang, Z. & Yao, Q. (2024). Accurate and interpretable drug–drug interaction prediction enabled by knowledge subgraph learning (KnowDDI). *Communications Medicine*. 4(1). https://doi.org/10.1038/s43856-024-00486-y.

Wishart, D. S., Feunang, Y. D., Guo, A. C. et al. (2018) DrugBank 5.0: A major update to the DrugBank database for 2018. *Nucleic Acids Research*. 46(D1), D1074–D1082. https://doi.org/10.1093/nar/gkx1037.

Jamilu BELLO is a Ph.D. student in the Department of Computer Science, Faculty of Computing and Information Technology, Osun State University, Osogbo, Osun State, Nigeria. He holds a Bachelor of Science and a Master of Science in Computer Science. His primary research interests include health informatics, probabilistic sequence modeling, clinical decision support systems, bioinformatics, soft computing, data science, and artificial intelligence (AI). He has published peer-reviewed papers on algorithmic optimizations and data-driven healthcare frameworks.

***

Adepeju Abeke ADIGUN is a Professor in the Department of Computer Science, Faculty of Computing and Information Technology, Osun State University, Osogbo, Osun State, Nigeria, and Director of the Centre for Entrepreneurship Studies and Innovations. She holds Bachelor's, Master's, and Ph.D. degrees in Computer Science. Her research interests include human-computer interaction (HCI), computer graphics, and operating systems. With more than 15 years of teaching and research experience in higher education, she has authored numerous publications in high-impact international journals. In recognition of her contributions to computer science in Nigeria, she has been awarded Fellowship of the Nigeria Computer Society (FNCS).

***

Patrick OZOH is an Associate Professor in the Department of Computer Science, Faculty of Computing and Information Technology, Osun State University, Osogbo, Osun State, Nigeria. He holds a Ph.D. in Computer Science and has extensive academic experience in advanced software engineering paradigms. His research interests include distributed computing, high-performance architectures, and systems design. He has contributed to teaching, research, and academic development at Osun State University.

***

Michael Olugbenga ABOLARINWA is a Lecturer in the Department of Cybersecurity, Faculty of Computing and Information Technology, Osun State University, Osogbo, Osun State, Nigeria. He holds Master's and Ph.D. degrees in Computer Science, with specialized training in software architectures and database security. His research interests include secure point-of-care interface design and the engineering of optimized clinical graphical user interfaces.

***

Ibrahim Olufemi ADIGUN is a Research Assistant in the Department of Electrical Engineering, Faculty of Electrical Engineering, Technical University of Applied Sciences Würzburg-Schweinfurt, Schweinfurt, Germany. His primary research fields include machine learning algorithms and intelligent systems architecture.