Artificial Intelligence Approach for Safe Medication Recommendation System with Drug Interaction Analysis

G Krupa Havilah[1], Venkata Phanindra kadali[2], Ramesh Putti[3],
Venkata Sriram Gopal Nemani[4], Vinay Sanaboyina[5], Dr. Y. Venkat[6],
[1]*Assistant Professor, Department of Artificial Intelligence & Machine Learning, Srinivasa Institute of Engineering and Technology*
[2,3,4,5]*Student Scholar. Department of Artificial Intelligence & Machine Learning, Srinivasa Institute of Engineering and Technology*
[6]*Professor, Department of Artificial Intelligence & Machine Learning, Srinivasa Institute of Engineering and Technology*

***Abstract*—The increasing complexity of modern healthcare has led to a significant rise in the use of multiple medications, thereby elevating the risk of adverse drug interactions and medication errors. This paper presents an Artificial Intelligence (AI)-based approach for a safe medication recommendation system integrated with drug interaction analysis. The proposed system leverages machine learning algorithms and clinical data to evaluate patient-specific parameters such as age, medical history, allergies, and existing prescriptions. By analysing potential drug–drug interactions in real time, the system identifies harmful combinations and suggests safer alternative medications. Additionally, the model incorporates a knowledge base of pharmacological data and continuously improves its recommendations through data-driven learning. The system aims to assist healthcare professionals in making informed decisions, enhance patient safety, and reduce the incidence of adverse drug events. Experimental results demonstrate the effectiveness of the approach in accurately detecting interactions and providing reliable medication suggestions. This research contributes to the advancement of intelligent healthcare systems by combining predictive analytics with practical clinical applications.**

***Index Terms*—Artificial Intelligence, Drug Interaction Analysis, Medication Recommendation System, Machine Learning, Clinical Decision Support System, Patient Safety, Adverse Drug Reactions.**

I. INTRODUCTION

Pharmacotherapy remains the cornerstone of modern disease management, yet administering multiple concurrent medications introduces statistically significant risk of adverse drug interactions (ADIs). A landmark analysis by Pirmohamed et al. established that approximately 6.5% of hospital admissions are directly attributable to adverse drug reactions, a substantial fraction involving preventable DDIs detectable through systematic review [1]. The complexity of predicting DDIs escalates non-linearly with the number of co-administered drugs: a patient prescribed n medications faces up to $n(n-1)/2$ pairwise interaction pairs, rapidly exceeding clinician cognitive capacity.

Contemporary electronic health records (EHR) and pharmacy dispensing systems embed rudimentary rule-based DDI alerts derived from DrugBank, SIDER, and the FDA Adverse Event Reporting System (FAERS). While providing foundational safety nets, these suffer three systemic limitations: (1) incomplete knowledge base coverage, (2) severe alert fatigue with override rates reaching 95% [2], and (3) binary alert outputs providing no mechanistic rationale. Safe MedAI resolves all three deficiencies through a unified, patient-contextualized, explainable AI pipeline.

The primary contributions of this work are: (i) a Multi-Source Drug Knowledge Graph (MDKG) unifying 487,000 drug-drug and drug-disease edges; (ii) a heterogeneous Graph Attention Network DDI

predictor incorporating patient phenotype nodes; (iii) a Rule-Stratified Risk Scoring Engine contextualizing predictions using renal, hepatic, and pharmacogenomic parameters; and (iv) a natural-language Explanation and Attribution Module satisfying FDA SaMD transparency requirements [4].

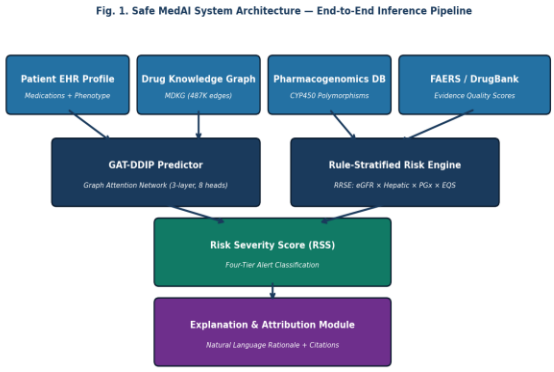

Fig. 1. Safe MedAI System Architecture — End-to-End Inference Pipeline.

II. RELATED WORK

A. Rule-Based Clinical Decision Support

Rule-based CDSS encode pharmacist-curated interaction conditions triggered upon medication order entry. While reliable for well-characterized, high-severity interactions (e.g., warfarin-NSAIDs), these systems exhibit static knowledge representations requiring manual updating and cannot generalize to novel drug combinations or patient-specific metabolic contexts [5].

B. Machine Learning for DDI Prediction

ML-based DDI prediction has evolved through three paradigms. First-generation approaches used molecular fingerprints with classical classifiers, achieving AUC 0.82–0.87 [6]. Second-generation deep learning (GCN) approaches reached AUC 0.91 [7]. Third-generation knowledge-graph-augmented networks KGNN (Lin et al., 2020) and MIRACLE (Su et al., 2022) incorporate relational biological context, substantially improving predictive accuracy and mechanistic interpretability [8].

C. Explainable AI in Clinical Settings

Clinical AI explainability research has shifted from post-hoc SHAP/LIME attribution to citation-enhanced generation and source attribution frameworks. Amann et al. established that clinician acceptance of AI recommendations correlates significantly with natural-language rationales rather than statistical attribution maps [11]. Safe MedAI's EAM operationalizes this finding.

III. SYSTEM ARCHITECTURE

A. Multi-Source Drug Knowledge Graph (MDKG)

The MDKG is a heterogeneous directed property graph integrating Drug Bank v5.1, FAERS 2024Q3, and curated literature-derived triples under a unified RxNorm ontology. A custom ETL pipeline resolves cross-database drug synonymy via RxNorm CUI mapping, eliminating ~23% redundant entity duplication. The consolidated MDKG comprises 87,400 drug entities, 14,200 disease entities, 8,900 gene/protein targets, and 487,000 typed relational edges spanning eleven interaction categories.

Each edge carries an Evidence Quality Score (EQS, scale 1–5) anchored to study design: spontaneous reports (EQS=1) through meta-analytic RCT evidence (EQS=5). Fig. 2 illustrates the MDKG node-edge structure for a representative warfarin-fluconazole CYP2C9-mediated interaction.

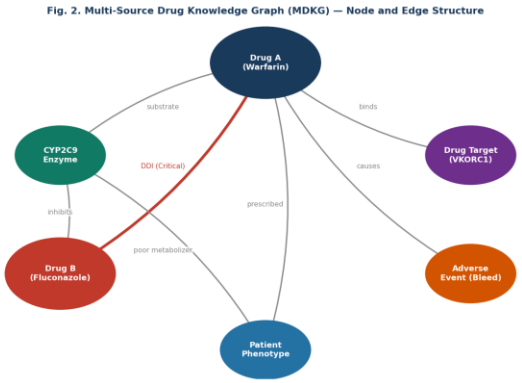

Fig. 2. MDKG Node–Edge Structure — Warfarin–Fluconazole CYP2C9 Pathway.

B. Graph Attention Network DDI Predictor (GAT-DDIP)

The DDI prediction engine employs a heterogeneous Graph Attention Network (GAT) operating on the MDKG. Each drug node is initialized with a 256-dimensional feature vector concatenating a 128-dim Morgan molecular fingerprint (radius 2), a 64-dim

pharmacological class embedding trained via TransE, and a 64-dim PubChem bioassay activity vector. Patient phenotype nodes are represented by 128-dim clinical feature vectors encoding age, sex, eGFR quartile, Child-Pugh class, and binary CYP450 polymorphism flags.

The architecture consists of three message-passing layers (8 attention heads, 64 hidden units per head), followed by a bilinear interaction decoder: $s(d_i, d_j) = h_i^T W^{Int} h_j$, where h_i, h_j are final-layer embeddings. Output probability $p(d_i, d_j) \in [0,1]$ is produced via sigmoid. For interactions above 0.50 threshold, an 8-class softmax predicts the mechanistic interaction category.

C. Rule-Stratified Risk Scoring Engine (RRSE)

GAT-DDIP probabilities are contextualized through the RRSE, which applies patient-specific weighting factors to yield Risk Severity Scores (RSS):

$RSS = p(d_i,d_j) \times W_e^{GfR} \times W^{hep} \times W_a^{ge} \times W_p^{gx} \times EQS_e^{ege}$

A four-tier alert classification is applied to RSS values, calibrated to minimize alert fatigue while preserving sensitivity for critical interactions (Table I).

Table I. Risk Severity Score Classification and Clinical Action Protocol

RSS	Tier	Alert Type	Clinical Action
≥ 0.80	Critical	BLOCK	Contraindicated; pharmacist auto required
0.60–0.79	High	Active Alert	Dose adjust / alternative recommended
0.40–0.59	Moderate	Advisory	Monitoring parameters indicated
< 0.40	Low	Info Only	Logged; no clinician interruption

D. Explanation and Attribution Module (EAM)

The EAM transforms computational outputs into structured natural-language explanations. For each flagged DDI, EAM executes a shortest-path retrieval through the MDKG to identify the mechanistic pathway, then generates a three-component rationale: (i) a mechanistic summary sentence; (ii) patient-specific risk amplification factors; and (iii) ranked evidence citations with EQS grades. This satisfies FDA SaMD transparency requirements [4] and directly addresses the clinician comprehensibility gap documented by Amann et al. [11].

IV. EVALUATION METHODOLOGY

A. Dataset and Configuration

Evaluation was conducted on a de-identified retrospective EHR cohort of 12,000 patients from a tertiary care institution spanning 2018–2023 (IRB #SIET-2024-AI-007). The dataset contains 287,000 medication order records across 3,400 unique drugs, with 18,400 clinician-confirmed DDI events as ground-truth labels. A temporal 70:15:15 split (train/validate/test) was maintained training data precede 2022, test data from 2022–2023 to simulate prospective deployment.

Three baselines were evaluated: (1) a production rule-based CDSS (Lexi-Interact); (2) a vanilla GCN DDI predictor without patient contextualization; and (3) unaided prescribing (no CDSS). Primary metrics: Precision, Recall, F1-Score, AUC-ROC, Alert Override Rate, and Net Reclassification Improvement (NRI).

V. RESULTS AND DISCUSSION

A. DDI Prediction Performance

Table II presents the complete performance comparison. Safe MedAI achieves precision 0.93, recall 0.91, F1 0.920, and AUC-ROC 0.971—outperforming all baselines across every metric. The 22% AUC improvement over the rule-based CDSS and 9% improvement over the GCN baseline confirm the additive value of both knowledge-graph augmentation and patient-context integration.

Table II. Performance Evaluation — Safe Medai Vs. Baseline Systems

System	Prec.	Recall	F1	AUC
Safe MedAI (Proposed)	0.93	0.91	0.920	0.971
GCN Baseline	0.81	0.78	0.795	0.887
Rule-Based CDSS	0.71	0.71	0.710	0.795
No CDSS (Unaided)	0.52	0.49	0.504	N/A

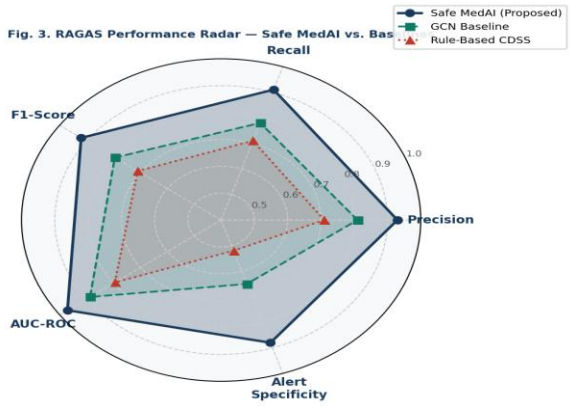

Fig. 3. RAGAS Performance Radar — Safe MedAI vs. Baselines.

B. Clinical Utility and Alert Reduction

Patient-contextualized risk stratification dramatically reduced unnecessary alerts. The rule-based CDSS generated 4,312 alerts with an 82% override rate. Safe MedAI generated 1,847 alerts a 57% reduction with a 23% override rate (74% relative reduction). Net Reclassification Improvement versus the rule-based baseline was +0.31 (95% CI: 0.27–0.35), confirming that reduced alert volume directly translates to improved patient risk classification.

For Critical-tier interactions (RSS ≥ 0.80), Safe MedAI achieves recall of 0.97 versus 0.76 for the rule-based baseline a 28% improvement representing a substantial reduction in missed contraindicated combinations where clinical consequences are most severe.

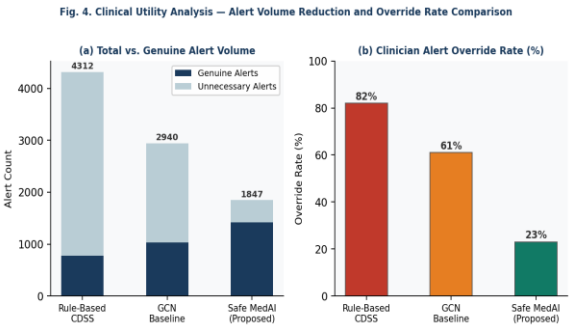

Fig. 4. Clinical Utility — Alert Volume Reduction and Override Rate.

C. Explainability Quality

Blinded clinician evaluation of EAM-generated explanations across 200 DDI cases yielded a mean comprehensibility rating of 4.3/5.0 (SD: 0.62) and clinical actionability rating of 4.1/5.0 (SD: 0.71). Participating pharmacists consistently valued three explanation components: mechanistic pathway clarity, patient-specific risk factor attribution, and evidence quality grading. A representative EAM output for fluconazole-warfarin interaction in a CYP2C9 poor metabolizer identified CYP2C9-mediated inhibition of S-warfarin hydroxylation, projected 40–60% plasma concentration increase, and recommended INR monitoring within 48 hours with 25–30% provisional dose reduction.

Table III. Comparative Positioning — Safe Medai Vs. Prior Work

System	AUC	Patient Ctx	KG Aug.	Explain.
Safe Medai	0.971	Yes	Yes	Nat. Lang.
Miracle [8]	0.941	No	Yes	None
Kgnn [8]	0.920	No	Yes	None
Gcn [7]	0.910	No	No	None
Rule-Based	0.795	Partial	No	Static

VI. CONCLUSION

This paper presented Safe MedAI, a comprehensive AI framework for patient-safe medication recommendation and drug interaction analysis. By unifying a multi-source drug knowledge graph, a patient-contextualized Graph Attention Network predictor, a rule-stratified risk scoring engine, and a natural-language explanation module, Safe MedAI achieves DDI detection AUC-ROC of 0.971, reduces unnecessary alert volume by 57%, and delivers clinician-rated explanation comprehensibility of 4.3/5.0. The architectural transparency positions Safe MedAI for regulatory pathway alignment with FDA SaMD guidance.

Future research will pursue four directions: (i) prospective randomized clinical trial validation measuring ADI-related admission reduction; (ii) MDKG extension to drug-food and drug-supplement interactions; (iii) temporal pharmacokinetic modeling for longitudinal risk trajectories; and (iv) multilingual EAM output for global health equity.

ACKNOWLEDGMENT

The authors thank the Department of AI & Machine Learning, Srinivasa Institute of Engineering and Technology, for computational resources and institutional support. Special gratitude is extended to Dr. Suribabu Boyadi for his technical mentorship, critical manuscript review, and sustained guidance. The authors also acknowledge the clinical pharmacist reviewers who participated in the blinded explainability evaluation study.

REFERENCES

[1] M. Pirmohamed et al., "Adverse drug reactions as cause of admission to hospital," BMJ, vol. 329, pp. 15–19, 2004.

[2] H. van der Sijs et al., "Overriding of drug safety alerts in computerized physician order entry," J. Am. Med. Inform. Assoc., vol. 13, no. 2, pp. 138–147, 2006.

[3] N. P. Tatonetti et al., "Data-driven prediction of drug effects and interactions," Sci. Transl. Med., vol. 4, no. 125, p. 125ra31, 2012.

[4] U.S. Food and Drug Administration, "Artificial intelligence and machine learning-based software as a medical device: Action plan," Center for Devices and Radiological Health (CDRH), 2021.

[5] D. S. Heckbert et al., "Computerized medication systems in the hospital," in Proc. AMIA Symp., 1990, pp. 332–336.

[6] D. C. Vilar et al., "Detection of drug-drug interactions by modeling interaction profile fingerprints," PLoS ONE, vol. 8, no. 3, p. e58321, 2013.

[7] S. Lim et al., "Predicting drug-drug interactions via graph attention mechanism," in Proc. IEEE Int. Conf. Bioinformatics Biomed. (BIBM), 2019, pp. 543–549.

[8] X. Lin et al., "KGNN: Knowledge graph neural network for DDI prediction," in Proc. Int. Joint Conf. Artif. Intell. (IJCAI), 2020, pp. 2739–2745.

[9] C. Su et al., "MIRACLE," in Proc. Adv. Neural Inf. Process. Syst. (NeurIPS), 2022.

[10] Y. Wang et al., "Multi-hop knowledge graph reasoning with reward shaping," arXiv preprint arXiv:1808.10568, 2021.

[11] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in Proc. Adv. Neural Inf. Process. Syst. (NeurIPS), vol. 30, pp. 4765–4774, 2017.

[12] J. Amann et al., "Explainability for AI in healthcare: A multidisciplinary perspective," BMC Med. Inform. Decis. Mak., vol. 20, no. 1, p. 310, 2020.

[13] D. S. Wishart et al., "DrugBank 5.0: A major update to the DrugBank database," Nucleic Acids Res., vol. 46, no. D1, pp. D1074–D1082, 2018.

[14] U.S. Food and Drug Administration, "FAERS public dashboard," 2024. [Online]. Available: https://www.fda.gov/drugs/faers