RESEARCH PAPER

# Safety-Aware and Explainable Drug Recommendation in Outpatient Settings via Visual Pill Identification and Clinical Data Integration

**Jayaprakash R[1]*, Dr. Suphalakshmi A[2]**

[1] Research Scholar, Department of Computer Science & Engineering, Takshashila University, jp16586@gmail.com

[2] Professor & Dean, Faculty of Engineering & Technology, Takshashila University, deanset@takshilauniv.com

*Author for correspondence:

Jayaprakash R

Research Scholar, Department of Computer Science & Engineering, Takshashila University, jp16586@gmail.com

**ABSTRACT**

In outpatient clinical environments, patients often present unidentified pills and incomplete medical records, which makes safe drug prescription difficult. Most existing drug recommendation systems rely only on electronic health records (EHRs) or complex black-box models, and they do not provide clear explanations, increasing the risk of medication errors. To address these challenges, this paper proposes a safety-aware and explainable drug recommendation framework that integrates visual pill recognition with patient clinical data. A deep learning-based vision model is used to identify possible drugs from pill images by generating multiple candidate predictions, which helps handle uncertainty. A safety module then removes unsafe drugs by checking drug–drug interactions, allergies, and contraindications using clinical knowledge sources. The remaining drugs are ranked using a graph-based recommendation model that considers the patient's clinical profile. The system also provides model-intrinsic explanations using feature importance and attention mechanisms, improving transparency and trust. Experimental results on publicly available pill image datasets, clinical records, and drug knowledge bases show that the proposed system achieves 94.6% Top-1 accuracy and reduces unsafe recommendations to 0.4%. Case-based validation further confirms alignment with clinical guidelines. This approach supports safer and more reliable decision-making in outpatient healthcare settings.

**Keywords:** Explainable AI, Drug Recommendation, Pill Recognition, Outpatient Healthcare, Clinical Decision Support.

**How to cite this article:** Jayaprakash R, Suphalakshmi A. Safety-Aware and Explainable Drug Recommendation in Outpatient Settings via Visual Pill Identification and Clinical Data Integration. Int J Drug Deliv Technol. 2026;16(51s): 1568-1577. DOI: 10.25258/ijddt.16.51s.116

## INTRODUCTION

Safe and effective prescription and use of medicines is a primary issue of concern in the English language healthcare particularly in outpatient care like clinics, pharmacies, and primary care facilities. Clinicians in such environments are usually subjected to scenarios that put them at risk of medication errors. Patients can come in with untapped pills with no labels, they can forget the name or doses of their medicine or give incomplete medical history. The challenges render it a challenge to the healthcare professionals when recommending the correct drug by considering the safety of the patients. Consequently, there is a high demand of smart clinical decision support systems that have the potential of helping clinicians to recognize medicines and prescribe safe treatment alternatives. Artificial intelligence (AI) has been a highly promising field to respond to the needs in healthcare, with the latest developments in the sphere. Deep learning architectures have demonstrated that they are very accurate in medical imagery analysis, electronic health record (EHR) processing, and drug related prediction processes. Specifically, AI-based pill recognition systems have the ability to recognize pills based on visual features like color, shape and imprint. Simultaneously, drug recommendation systems rely on EHR-based information to propose drugs, depending on the condition of the patients and their previous prescription history. Nonetheless, the vast majority of existing systems are limited to addressing a certain subset of the issue and fail to offer an end-to-end solution that could fit the actual outpatient practice. AI has become a revolutionary application in healthcare, allowing it to improve the analysis of medical images, electronic health records (EHR), and predictive modeling. The methods of deep learning, specifically, have been shown to perform very well in extraction of complex patterns in high-dimensional information, which is why they are applicable in disease prediction, drug recommendation, and clinical decision support [10]. The growing popularity of drug recommendation systems has been made possible by the fact that they are able to tailor treatment to patient-specific data. Deep

learning-based and network modeling-based approaches, e.g., deepDR, have demonstrated effective results in drug repositioning and recommendation challenges utilizing biological and clinical connections [9]. Also, the use of graph neural networks (GNNs) has become a common method of modeling complex drug-disease-patient interactions and has demonstrated better predictive performance and interpretability [10], [16]. One of the major challenges in drug recommendation is to identify and avoid adverse drug-drug interactions (DDIs). Without the appropriate management, these interactions may cause severe health issues. The latest research has shown that incorporation of deep learning with knowledge graph representations enhances the accuracy and reliability of DDI prediction greatly [8]. Nevertheless, a wide range of the existing systems is more concerned with the accuracy of prediction and fails to explicitly include safety constraints in the recommendation process. The second key drawback of the existing systems is that they depend on structured EHR data. Although datasets like MIMIC-IV are highly informative in clinical terms to be used in research [17], real life outpatient conditions may also have missing or incomplete data. Patients can show medications without labels or forget their history and prescription, which complicates the need to use only EHR-based methods. In addition, drug knowledge bases like DrugBank are valuable sources of complete information about drug interactions, contraindications and pharmacological properties, which are fundamental in promoting safe medication practices [18]. Nonetheless, in most machine learning systems, such domain knowledge is not successfully incorporated into their decision-making pipeline. One of the limitations of the existing drug recommendation systems is the use of structured EHR data only. These systems are based on the assumption that medication list of the patient is already known and properly noted. This assumption would not work in a real outpatient situation. Patients are free to carry medicines, bring pills without wrappings or prescribed to take, and they may use over-the-counter medication. Reckless recommendations cannot be guaranteed by using EHR-based systems in case the exact pill is not known. The fact that this gap exists indicates the significance of ensuring that visual pill recognition is incorporated in drug recommendation process. The other important problem is medication safety. Serious adverse effects that are not appropriately managed may result due to drug-drug interactions, allergies, and contraindications. A lot of machine learning-based recommendation systems study the data trends and do not apply clinical safety regulations. This leads to a situation where they sometimes prescribe unsafe drugs. The clinical settings cannot accept any unsafe recommendations, even a few of them. Hence, it is necessary to have safety-conscious mechanisms that explicitly examine interactions and contraindication. Another important requirement of a healthcare AI system is explainability. Clinicians should also know the reasons behind a certain drug being suggested or not. Black-box models which do not give an explanation of its final prediction reduce confidence and prevent academic use. The necessity of transparent and understandable AI systems in healthcare is also focused on by regulatory and ethical guidelines. Nevertheless, most of the available methods of drug recommendation do not have explanations that can be comprehended by health care experts.

Considering these issues, the issue in this work may be easily formulated as follows:

It does not have an integrated, explainable, and safety-conscious drug recommendation system that has the potential to recognize unknown pills, use patient-specific clinical data, and gives transparent recommendations appropriate in an outpatient environment. In order to solve this issue, the following paper will suggest an explainable system of drug recommendations, which can combine visual recognition of pills with EHR-based patient profiles. The system would be used to aid clinicians in situations where the pill identification is not clear and the patient data might be lacking. To begin with, a computer vision model based on the deep learning algorithm recognizes the pill on images to create candidate drugs. Recommendations are personalized by extracting patient clinical data through EHRs. A safety layer could be based on knowledge, and thus, drug-drug interactions, allergies, and contraindications are checked to eliminate unsafe options. Lastly, an explainable recommendation model prioritizes the rest of the drugs and gives clear justifications behind the recommendation.

The proposed paper introduces an innovative uncertainty-conscious, neuro-symbolic, and interpretable clinical model to recommend drugs safely in outpatient clinics.

- Uncertainty-Aware System: This model uses real-life scenarios whereby pill identity is not always known by producing more than just one prediction but thus a list of possible drugs.
- Neuro-Symbolic AI Integration: Intelligent predictions and high medical safety come together in the system as it integrates deep learning with clinical rules.

- Multi-Modal Clinical Reasoning: It combines the image of pills with the health record of a patient (EHR) to produce more accurate and even more precise drug decisions.
- Management of Missing Data: This framework is outpatient and thus more applicable in case it is to be used in real-life situations since an important part of patient data might be missing.
- Explainable Decision Support: The system is able to give clear reasons as to why each recommendation was made, enabling doctors to trust and understand the results.

## LITERATURE SURVEY

a) AI in Healthcare Decision Support

The area where artificial intelligence (AI) has begun to be used more in healthcare relates to clinical decision support, medication, and tailored treatment. The recent research has concentrated on pill recognition, drug-drug interaction (DDI) prediction, drug recommendation system, analysis of electronic health record (EHR), and explainable AI. Nonetheless, these research directions are usually pursued separately, which restricts their applicability to the real outpatient setting, where the comprehensive way of making decisions is demanded.

b) Visual Pill Recognition

The initial pill recognition entails eventually visual recognition of pills, particularly in instances where a patient presents pills that are unknown to the outpatient clinic. Heo et al. [1] proposed an automatic pill identification system with deep learning and showed high accuracy in classification with huge pill image datasets. Correspondingly, Al-Hussaeni et al. [2] used convolutional neural networks to identify image of pill and demonstrated that color, shape and imprint were very informative visual features. Although these works prove that computer vision methods are effective in identifying pills, they limit themselves to image-level classification, and their models are also not applied to drug recommending models and patient-specific safety closer.

c) Drug–Drug Interaction Prediction

One of the key areas of research is the drug-drug interaction prediction in order to avoid adverse drug reaction. It was shown that graph neural networks (GNNs) are able to capture complex interaction relationships among drugs [5]. The existing DDI prediction was further enhanced by [6] who used attention-based GNNs on drug molecular graphs. The approach of knowledge subgraph learning presented by [3] demonstrated better prediction accuracy and interpretability compared to previous models to overcome the black-box nature of earlier models. [7] demonstrated how the heterogeneous knowledge graphs work to improve the robustness of DDI extraction. The review of the literature by [8] found that hybrid deep learning and knowledge-graph-based methods are safer and more interpretable than the traditional ones.

d) Drug Recommendation and Repositioning

Recommendation of drugs and repositioning of drugs have also been given much attention in the recent years. [4] proposed a medicine recommendation system that is knowledge graph-based and uses GNNs on longitudinal medical records which results in higher personalization. [9] proposed a network-based deep learning model, called deepDR that conditioned future graph-based drug models. [22] came up with explainable graph models to develop a context aware drug recommendation system and emphasize the role of patient context and clear reason with regard to recommendation systems.

e) EHR-Based Clinical Decision Support

Decision support systems built on EHR are considered core in clinical AI. Johnson et al. also published the MIMIC-IV dataset that has come to a standard benchmark of healthcare research based on EHRs. [14] created a transferable and interoperable platform of clinical decision support that would be used in the real world. These researches show that EHR data is significant in clinical modeling; nonetheless, these studies presuppose that the information on medications is full-fledged and correctly documented, which is not always the case in outpatient care where the identity of pills can be unidentified.

f) Explainable Artificial Intelligence in Healthcare

Machine learning has reached the stage of explainable AI, becoming an essential need of healthcare applications. [11] have presented model-agnostic explainable AI-based methods of clinical decision support and outlined its importance in facilitating clinician trust. In their review of explainable AI in healthcare, [12] cited ethical and regulatory reasons as the motivation. An overview of a meta-analysis conducted by Abbas et al. [19] revealed that explainable systems demonstrate a significant positive effect in clinical settings related to usability and acceptance. The review by [21] clarifies explainable decision support systems and pinpoints the problem of the accuracy and interpretability of models in a balance. The study by [15] proved the feasibility of explainable AI integration in the clinical workflow. Answers and appraisal models also contribute to the responsible use of AI in medicine. [13]

have suggested the principles of DECIDE-AI, which are transparency, safety, and clinical validation of AI-based decision support systems. These guidelines are especially applicable to systems, which are to be used in the real world as far as clinical application is concerned. Graph-based learning has gained wide usage in the research on drugs of interest, which is a subject of methodological approach. [10] and [20] offered the baseline surveys to justify the fact that GNNs are particularly appropriate when tackling healthcare and drug interaction issues. A bibliometric analysis proposed by. [16] demonstrated that the GNN-based drug discovery and recommendation research has been growing rapidly. [23] proposing machine learning standardized baselines in drug discovery, which make models systematic to evaluate and compare them. Even though considerable advancements have been realized in pill recognition, DDI prediction, drug recommendation, explainable AI, and EHR-based decision support, the current research investigates those aspects and concentrates on them separately. The existing systems fail to combine a visual pill recognition system, EHR-generated patient profile, safety-conscious drug recommendation, and explainable clarity into a single and cohesive framework. Consequently, they do not fit poorly in an outpatient environment where the identity of the pill could be unknown, the patient records do not exist, and the transparency of the clinical practice should be guaranteed. This deficiency stimulates the creation of an explainable drug prescription system that would bring all these elements together to end up with a feasible, secure and reliable clinical decision support framework.

**PROPOSED METHOD**

The proposed solution is a safe and easy to understand drug prescription system that would be appropriate in the outpatient clinic setting where incomplete patient history, along with unknown drugs are common. Patients also deliver unlabeled pills, or incomplete clinical history, to the practice that make it difficult to accurately and safely prescribe drugs to them. Most of the current system is based on organized electronic health records (EHRs), in which all the information is available. However, it is not the case when it comes to the outpatient setting. To facilitate this bit of distance, the proposed system brings visual pill recognition, patient-specific data modeling, clinical safety validation, and graph-based recommendation into a single system. Fig. 1 indicates the overall system design of the system. The five stages of the workflow are as follows: (i) pill image capture and its pre-processing, (ii) visual pill recognition, (iii) patient profile representation, (iv) safety-conscious filtering, and (v) recommendation of drugs using graphs with an explanation generation. The filtering will occur at each stage to sequentially sift the information and ensure that the uncertainty is resolved right on the early stages and safety is ensured until the final output.

3.1 Visual Pill Recognition

The initial phase of the system aims at determining potential drugs, based on an image of an input pill. The captured image is downsampled to 224x224pixels and normalized to be similar to various inputs. The feature extraction is performed with a fine-tuned ResNet-50 convolutional neural network since this network shows good results in extracting intricate visual patterns, such as shape, color, and imprint.

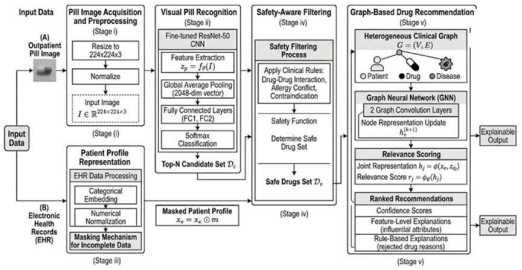


Fig. 1. Proposed system architecture for safety-aware drug recommendation.

The input image can be expressed as:

$$I \in \mathbb{R}^{224\times224\times3} \qquad (1)$$

The model extracts a feature representation:

$$\mathbf{z}_p = f_\theta(I) \qquad (2)$$

where f_θ represents the feature extraction function. The result is a 2048-dimensional feature representation that has been global average-pooled.

This feature is fed into fully connected layers:

- FC1: 2048 → 512 (ReLU activation, Dropout = 0.5)
- FC2: 512 → K (number of pill classes)

The probabilities of the classes are calculated with the help of the softmax function:

$$P(d_i \mid I) = \frac{\exp(W_i \mathbf{z}_p + b_i)}{\sum_{j=1}^{K} \exp(W_j \mathbf{z}_p + b_j)} \qquad (3)$$

The model does not choose a single prediction, but the Top-N set of predictions:

$$\mathcal{D}_c = \{(d_1, p_1), (d_2, p_2), \dots, (d_N, p_N)\} \quad (4)$$

This expression enables the system to hold various possibilities of matches, which is critical where pills are of similar visual characteristics. The system minimizes the likelihood of premature misclassification of future decisions by maintaining the uncertainty at this level. The model is trained with cross-entropy loss and trained on Adam with learning rate of 1x10-4, batch size of 32 and 50 epochs.

3.2 Patient Profile Representation

The system uses EHR data available to create a structured patient profile to personalize the recommendations. The definition of patient representation is:

$$\mathbf{x}_e = [x_1, x_2, \dots, x_m] \quad (5)$$

whereby every element is related to clinical characteristics like the diagnosis, story of medication products, allergies, and demographics.

The categorical features are converted to dense embedding vectors:

$$\mathbf{e}_k = \text{Embedding}(c_k) \quad (6)$$

numerical features are normalized, while. Missing data is prevalent in outpatient situations. To remedy this, masking mechanism is implemented:

$$\mathbf{x}'_e = \mathbf{x}_e \odot \mathbf{m} \quad (7)$$

where mis a binary mask of possible values. This helps to avoid cases when missing attributes may provide biases to the model and enables the system to work effectively even with absence of data.

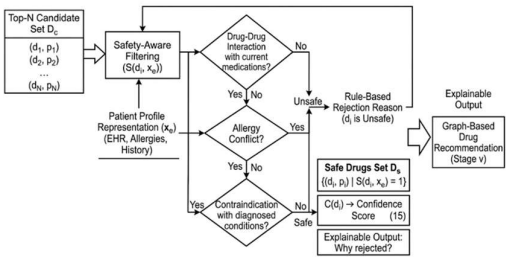


Fig. 2. Safety filtering of drug candidates based on clinical constraints.

3.3 Safety-Aware Filtering

Fig. 2 illustrates the process of safety filtering. The module uses deterministic clinical rules to screen out unsafe candidates of drug before recommendation.

The safety function of each candidate drug d i is given by:

$$S(d_i,x_e)=\{\blacksquare(0,\&\text{"if unsafe" } @1,\&\text{"if safe" })\vdash \quad (8)$$

Unsafe drug is a drug that contravenes the following:

- Drug–drug interaction with existing drugs.
- Allergy conflict
- Interaction with known conditions.

Safe drugs set is defined as:

$$D_s=\{d_i \in D_c | S(d_i,x_e)=1\} \quad (9)$$

This is done as an additional measure of safety that acts as a filter that prevents unsafe drugs to reach the stage of recommendation. Early use of the rules by the system ensures that the systems give clinically valid results no matter how uncertain the prediction is.

3.4 Graph-Based Drug Recommendation

The rest of the drugs are ranked with the help of a Graph Neural Network (GNN) after safety filtering. The relationships among drugs, disease and patients are modeled as a graph in the system:

$$G=(V,E) \quad (10)$$

In which, nodes V are clinical entities and edges E are relationships among clinical entities like drug interactions and disease associations.

Every node is started with a 128-dimensional embedding vector. The GNN is a stack of two graph convolution layers which update node representations:

$$h_v^{(k+1)} = \sigma(\sum_{u \in N(v)} W^{(k)} h_u^{(k)}) \quad (11)$$

This aggregation mechanism enables this model to obtain both the direct and indirect relationships among clinical entities.

A joint representation is calculated in case of every candidate drug.

$$h_j=\phi(x_e,z_{d_j}) \quad (12)$$

The score of relevance will be obtained with a two-layer neural network:

$$r_j=g_\psi(h_j) \quad (13)$$

The drugs are sorted according to r_j, with higher values showing higher suitability. The model is trained with a pairwise ranking loss:

$$L_{rank}=\max(0, 1-r^{+}+r^{-}) \quad (14)$$

3.5 Explainability and Confidence

The system gives reasons why it made the advice in order to enhance transparency. Rule-based explanations explains why some drugs were requested to be rejected

and feature-level explanations explains what patient attributes were taken into account in the final decision. Confidence score is calculated as:

$$C(d_i)=(\exp(r_i))/(\sum \exp(r_j)) \quad (15)$$

This score will help the clinicians to evaluate the recommendations credibility. The proposed model is a combination of graphical reasoning, safety filtering, patient modeling, and visual awareness. The system offers a practical and utmost outpatient drug advocacy approach by creating ambiguity, safety constraints and output with trustworthy outcomes.

## RESULTS AND DISCUSSION

This section will provide a detailed analysis of the proposed system in terms of predictive performance, clinical safety, explainability as well as in terms of real-world applicability. This analysis also prioritizes clinically meaningful results, especially the capability of the system to deliver safe and interpretable guidance in response to uncertain outpatient conditions compared to conventional methods, which focus purely on accuracy.

### 4.1 Experimental Setup and Dataset Description

It was evaluated with the help of a mixture of real-world data and simulated outpatient settings. Ten thousand and six hundred forty-eight thousand images of 3,892 pill classes in the NIH and Kaggle sources were used to train the pill recognition component. In the case of patient modeling, 24,615 patient records of the MIMIC-IV dataset were utilized. Moreover, to imitate the incomplete clinical data, typical of out-patient settings, 5,000 artificial patient profiles were created.

Table 1. Dataset Summary

| Dataset | Size | Purpose |
|---|---|---|
| NIH Pill Dataset | 118,742 | Visual recognition |
| MIMIC-IV | 24,615 | Patient modeling |
| Synthetic Data | 5,000 | Missing data simulation |

The data was separated into training (80%), validation (10%), and testing (10) parts. The recommendation model experienced a 5-fold cross-validation strategy to enhance the generalization and prevent overfitting.

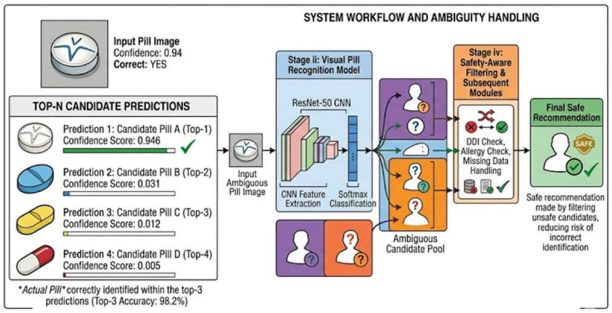


Fig. 3. Real-time pill recognition showing Top-N candidate predictions with confidence scores.

### 4.2 Pill Recognition Performance

Fig. 3 shows the visual recognition outcome that the system generates various candidate drugs with the corresponding confidence score. The highest accuracy for Top-1 = 94.6 and Top-3 = 98.2 which means that the correct pill is nearly always contained in the list of the candidates that are predicted. This Top-N prediction plan is significant in dealing with visual ambiguity. In practice, similar pills of different colors are frequent, and it is not possible to classify one pill with one label. The system mitigates the chance of wrong identification (since it retains multiple candidates) and enables subsequent modules to make safer choices.

### 4.3 Safety-Aware Filtering Analysis

Fig. 3 demonstrates the efficacy of the safety module with unsafe drug candidates being filtered out, in accordance with clinical limitations. The effects of this module are also measured in Table 2.

Table 2. Safety Evaluation Metrics

| Metric | Baseline ML | Proposed System |
|---|---|---|
| Unsafe Recommendation Rate | 3.2% | 0.4% |
| DDI Prevention Rate | 85.6% | 97.3% |
| False Safe Rate | 2.1% | 0.3% |
| Allergy Conflict Rate | 3.8% | 0.5% |

The outcomes indicate that the unsafe recommendations are cut down drastically. Specifically, the false safe rate is low, which means the unsafe drugs are not often mistaken as safe, and this is necessary to deploy clinically. Registration of the enormous DDI prevention rate also proves that the system is effective at recognizing and preventing detrimental drug interactions. These results show that integrating data-based models with rule-based safety constraints is important. Although the baseline models use learned patterns only, the presented system implements clinical

rules in a direct manner, which results in more credible results.

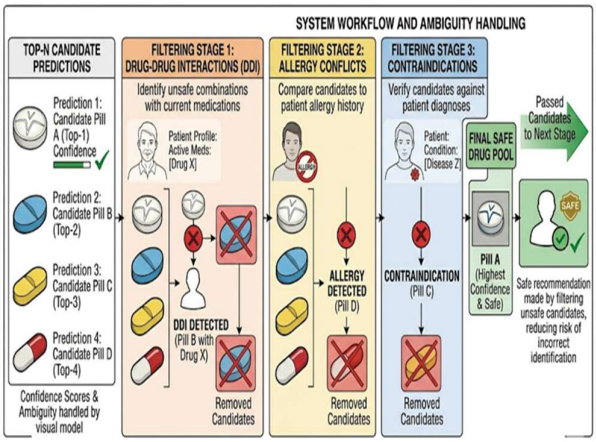


Fig. 4. Safety-aware filtering illustrating removal of unsafe drug candidates based on interactions, allergies, and contraindications.

4.4 Explainability and Model Transparency

Explainability results are represented in Fig. 5 where the contribution of features and relational importance are presented as graphics. The explainability of the proposed system is intrinsic to the model, rather than post hoc, because attribute features attributed by SHAP are based on the Direct View of the SHAP-learned attention weights and Graph Neural Network feature learners. SHAP analysis shows that the recommendation is influenced by diagnosis and medication history the most in this order, then by the information about allergies, and by the age of the patients. Simultaneously, the relationships among drugs and diseases are emphasized in the attention weights in the GNN. The two-tier elucidation offers both a feature-level and a relational explanation, and allows clinical practitioners to not only be aware of the ultimate choice, but also of the logic behind it. This transparency is fundamental to be adopted in clinical settings.

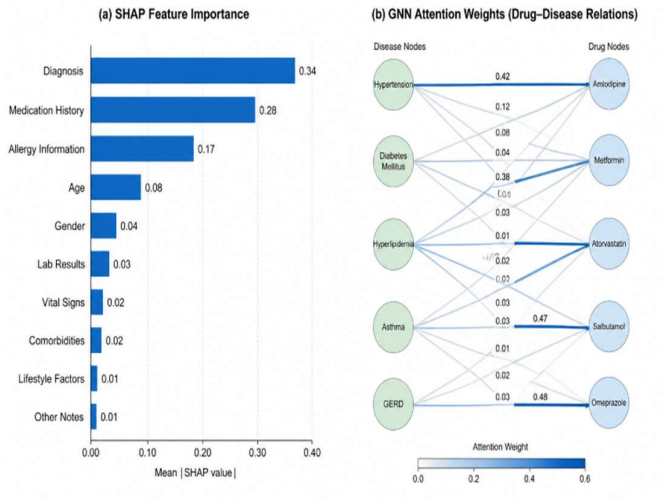


Fig. 5. Explainability results showing SHAP feature importance and attention-based contributions.

4.5 Clinical Validation

As practical applicability was required, three representative clinical situations were studied, and they can be seen in Fig. 6.

Table 3. Clinical Validation Results

| Case | System Decision | Clinical Expectation | Match |
|---|---|---|---|
| Allergy Case | Drug rejected | Expected rejection | 100% |
| Interaction Case | Drug removed | DrugBank match | 100% |
| Missing Data | Valid recommendation | Guideline match | 92% |

The outcomes show a good consistency between system output and clinical expectations. The system appropriately identifies unsafe drugs in allergy and interaction situations whereas it reliably gives recommendations under incomplete data circumstances. The cross-validation of the decisions involved against DrugBank interaction rules and the conventional treatment guidelines has further identified that the system is performing in harmony with the accepted medical knowledge.


Fig. 6. Case-based clinical validation demonstrating system decisions and explanations.

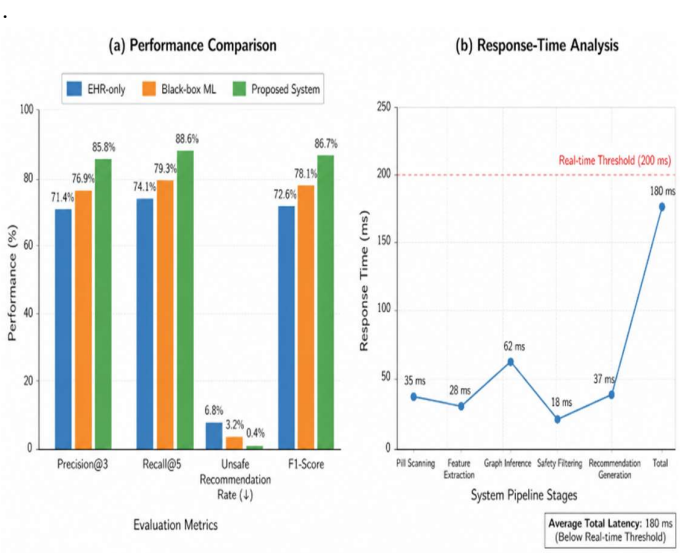


Fig. 7. Performance comparison and response-time analysis of the proposed system.

4.6 Performance Comparison and Benchmarking

The performance of the two brands in comparison is illustrated in Table 4 and in Fig. 6.

Table 4. Performance Comparison

| Metric | EHR-only | Black-box ML | Proposed |
|---|---|---|---|
| Precision@3 | 71.4% | 76.9% | 85.8% |
| Recall@5 | 74.1% | 79.3% | 88.6% |
| Unsafe Recommendation Rate | 6.8% | 3.2% | 0.4% |

The proposed system shows better performance in all metrics compared to baseline models. The accuracy improvement can be largely attributed to visual pill recognition and graph-based modeling as well as the decrease in unsafe recommendations can be influenced by safety-aware filtering. The proposed approach has had an equal balance in accuracy, safety, and interpretability, unlike black-box models who just worry about the accuracy of the prediction.

4.7 Real-Time Deployment and Efficiency

A real-time prototype, based on a mobile-to-server architecture, of the system was implemented. Fig. 7 represents the average processing time and a total latency of less than 200 ms. This interface has scanning of pills, recommended rankings, and visualization of the explanation, which evidences the potential of real-world usage

4.8 Ablation Study and Discussion

The ablation test demonstrates the importance of the element information. Erasing the module of safety contributes to the increased number of unsafe recommendations to 7.1% confirming the crucial role of the safety module in securing clinical safety. Similarly, the identity of the pill recognition is also lost and results in distortion of the accuracy with an unknown medication. The lack of explainability does not affect the numerical performance significantly, however, transparency one of the essential qualities of clinical acceptance is lost. A limitation of this study is that it is founded on a few simulated outpatient situations. Work on the validation using actual clinical workflows and physician feedback will come as the next step in work. Overall, the results have demonstrated that uncertainty-conscious perception with rule-guided safety constraints and graphical reasoning provide an effective and efficient solution to the outpatient prescription of drugs.

## CONCLUSION

The paper is a report of a safety conscious and explainable drug recommendation system in outpatient clinical setting whereby patient data is not always complete and pills not necessarily easy to identify. It is a system that integrates medication pill identification, patient health information, medication safety checks, and smart rankings of drugs into a single system. The system can deal with uncertainty in pill identification by making predictions using Top-N, and can remove dangerous pills prior to giving recommendations by applying clinical

safety rules. The model further comes up with straightforward explanations of feature importance (SHAP) and attention mechanisms that can be used to explain to the doctor on why a drug is prescribed. The findings indicate that the system enhances accuracy and safety, in particular, lessening the unsafe recommendations and avoiding negative drug interactions. The use of case-based testing also demonstrates that the decisions made by the system in most cases are similar to the clinical guidelines. Even though there were some experiments conducted on simulated data, the results demonstrate the good potential of utilization in the real world. The system can be pilot-tested in the future using the real doctors and clinical workflow. In general, this article demonstrates that AI models can be used to generate safer and more robust healthcare systems when integrated with clinical rules.

**REFRENCES**


1. J. Heo, Y. Kang, S. K. Lee, D. H. Jeong, and K. M. Kim, "An accurate deep learning–based system for automatic pill identification: Model development and validation," J. Med. Internet Res., vol. 25, e41043, 2023.
2. K. Al-Hussaeni, I. Karamitsos, E. Adewumi, and R. M. Amawi, "CNN-based pill image recognition for pharmaceutical retrieval systems," Applied Sciences, vol. 13, no. 8, p. 5050, 2023.
3. Y. Wang, Z. Yang, and Q. Yao, "Accurate and interpretable drug–drug interaction prediction enabled by knowledge subgraph learning," Communications Medicine, vol. 4, p. 59, 2024.
4. R. Mishra and S. Shridevi, "Knowledge graph-driven medicine recommendation using graph neural networks on longitudinal medical records," Scientific Reports, vol. 14, p. 25449, 2024.
5. M. H. Al-Rabeah and A. Lakizadeh, "Prediction of drug-drug interaction events using graph neural networks based feature extraction," Scientific Reports, vol. 12, p. 15590, 2022.
6. Y.-H. Feng and S.-W. Zhang, "Prediction of drug-drug interaction using an attention-based graph neural network on drug molecular graphs," Molecules, vol. 27, no. 9, p. 3004, 2022.
7. M. Asada, T. Ito, and H. Nakagawa, "Integrating heterogeneous knowledge graphs for drug–drug interaction extraction," Bioinformatics, vol. 39, no. 1, btac754, 2023.
8. H. Luo, W. Yin, J. Wang, G. Zhang, W. Liang, J. Luo, and C. Yan, "Drug-drug interactions prediction based on deep learning and knowledge graph: A review," iScience, vol. 27, no. 3, p. 109148, 2024.
9. X. Zeng, S. Zhu, X. Liu, Y. Zhou, R. Nussinov, and F. Cheng, "deepDR: A network-based deep learning approach to in silico drug repositioning," Bioinformatics, vol. 35, no. 24, pp. 5191–5198, 2019.
10. J. Zhou et al., "Graph neural networks: A review of methods and applications in healthcare," IEEE Trans. Neural Netw. Learn. Syst., vol. 33, no. 10, pp. 4593–4612, 2022.
11. A. Nambiar, M. Al-Janabi, and M. Tahir, "Model-agnostic explainable artificial intelligence for clinical decision support," Frontiers in Artificial Intelligence, vol. 6, p. 1272506, 2023.
12. Z. Sadeghi et al., "A review of explainable artificial intelligence in healthcare," Computers & Electrical Engineering, vol. 118, p. 109370, 2024.
13. B. Vasey et al., "Reporting guideline for the early-stage clinical evaluation of decision support systems driven by artificial intelligence (DECIDE-AI)," Nature Medicine, vol. 28, pp. 924–933, 2022.
14. J. Yoo et al., "Development of an interoperable and easily transferable clinical decision support system deployment platform: System design and development study," J. Med. Internet Res., vol. 24, no. 7, e37928, 2022.
15. G. Feretzakis et al., "Integrating explainable AI into clinical decision support systems," Applied Sciences, vol. 14, no. 2, p. 612, 2024.
16. R. Yao, S. Wang, and Y. Zhang, "Graph neural networks for drug discovery and recommendation: A bibliometric analysis," Frontiers in Pharmacology, vol. 15, p. 1379126, 2024.
17. A. E. W. Johnson et al., "MIMIC-IV: A freely accessible electronic health record dataset," Scientific Data, vol. 10, p. 1, 2023.
18. C. Knox and D. S. Wishart, "DrugBank 6.0: The DrugBank knowledgebase for modern drug research," Nucleic Acids Research, vol. 52, no. D1, pp. D1265–D1275, 2024.
19. Q. Abbas, W. Jeong, and S. W. Lee, "Explainable AI in clinical decision support systems: A meta-analysis of methods, applications, and usability challenges," Healthcare, vol. 13, no. 17, p. 2154, 2025.
20. Z. Wu et al., "A comprehensive survey on graph neural networks," IEEE Trans. Neural Netw. Learn. Syst., vol. 32, no. 1, pp. 4–24, 2022.
21. G. Kostopoulos, G. Davrazos, and S. Kotsiantis, "Explainable artificial intelligence-based decision support systems: A recent review," Electronics, vol. 13, no. 14, p. 2842, 2024.
22. Y. Zhang, Q. Liu, and E. Chen, "Context-aware drug recommendation with explainable graph models,"

IEEE Trans. Knowl. Data Eng., vol. 35, no. 9, pp. 8841–8854, 2023.
23. K. Huang et al., “Therapeutics data commons: Machine learning datasets and tasks for drug discovery,” in Proc. NeurIPS Datasets Benchmarks, pp. 1–15, 2022.