The American Journal of Engineering and Technology
ISSN 2689-0984 Volume 07 - 2025

LLM-Powered Prescription Cart Intelligence: A Hybrid System for Real-Time Drug Interaction Detection in E-Commerce

[1]Deepanjan Mukherjee
[1]Independent Researcher, USA

Received: 13th Oct 2025 | Received Revised Version: 05th Nov 2025 | Accepted: 07th Dec 2025 | Published: 30th Dec 2025

Volume 07 Issue 12 2025 | Crossref DOI: 10.37547/tajet/v7i12-312

Abstract

The present online pharmacy market lacks real-time drug-drug interaction detection during the shopping experience. This paper presents a unique system to detect interactions directly in e-commerce pharmacy shopping carts, reducing the risk of adverse drug reactions that could lead to potential hospitalizations. The hybrid system combines the current rule-based checking using commercial databases (DrugBank, First DataBank) with Large Language Models (LLMs) to improve contextual analysis through Retrieval-Augmented Generation (RAG). A three-layer design comprising of interaction detection, LLM enhancement, and user experience layers is proposed, to achieve under 500ms response times through microservices architecture and multi-tier caching, while generating user-friendly natural language explanations. A confidence scoring mechanism flags uncertain outputs for further pharmacy review and intervention to ensure user safety. The system also addresses critical limitations of current similar tools requiring use of separate interaction checkers by providing seamless cart-level integration. The proposed evaluation methodology targets >90% sensitivity for major interactions and >80% specificity to minimize pharmacist fatigue due to false positives.

Keywords: LLM, Artificial Intelligence, Drug-drug interactions, E-commerce, Clinical decision support, Patient safety, Online pharmacy, Polypharmacy.

Cite This Article: Mukherjee, D. (2025). LLM-powered prescription cart intelligence: A hybrid system for real-time drug interaction detection in e-commerce. The American Journal of Engineering and Technology, 7(12), 144–151. https://doi.org/10.37547/tajet/v7i12-312.

1. Introduction

Every day, hundreds of thousands of consumers click on "Add to Cart" on online pharmacy platforms and inadvertently create combinations of medications that could send them to the emergency room. There are 6 emergency department visits for therapeutic and nontherapeutic medication harms per 1,000 patients, with about 38% of such visits subsequently requiring hospitalization [1]. The online pharmacy market reached $82.91 billion in 2023 and is projected to grow to $294.35 billion by 2030 [2], yet current online pharmacy platforms focus on convenience but lack proactive safety mechanisms.

Drug-drug interactions (DDIs) impose a massive public health and economic burden. Among patients taking any prescription drug, half are exposed to two or more drugs, and 5% take eight or more simultaneously [3]. Adverse drug reactions account for 4.2-30% of hospital admissions in the USA and Canada [4], with patients experiencing adverse effects staying hospitalized 1.2-3.8 days longer

and incurring additional costs of $2,284-$5,640 per patient [5].

Research reveals alarming gaps in current detection systems. An examination of 50 community pharmacies found potential DDIs in 39.2% of 1,000 prescriptions [6]. More concerning, a study of 64 Arizona pharmacies found only 28% correctly identified eligible interactions, with a median sensitivity of 0.85 for detecting well-established interactions [7]. Traditional clinical decision support systems generate excessive non-clinically relevant alerts, causing alert fatigue— pharmacists override 73.8% of DDI alerts on average [7]. This suggests that existing systems are poorly calibrated for both clinical and consumer contexts.

At present, most online pharmacy systems defer interaction checking until after a purchase is made, and back-end pharmacy staff processes hundreds of orders. However, while consumers are shopping, they do not receive any feedback on potential harmful combinations with new or existing medications. Recent advancements in Large Language Models (LLMs) present a paradigm shift in drug interaction detection. Although some of these models are found to accurately identify DDIs, the accuracy of such models remains low due to limited access to high-quality datasets, interpretability issues, and the inability to consider patient-specific factors.

Despite these challenges, specialized approaches have shown promise. For instance, smaller fine-tuned models like Phi-3.5 with 2.7 billion parameters achieved sensitivity of 0.978 and accuracy of 0.919 for DDI prediction [8]. This paper explores a hybrid system for DDI detection: using rule-based databases for reliable detection of known interactions, and LLMs providing contextual analysis and consumer-friendly explanations. It further introduces Prescription Cart Intelligence, an e-commerce system that leverages LLMs to detect DDIs directly within the shopping cart experience, transforming online pharmacies to proactively monitor patient safety.

2. Background and Related Work

Drug interactions occur when one medication alters the pharmacological effect of another, potentially leading to adverse events or treatment failures. These interactions are classified as pharmacokinetic or pharmacodynamic in nature [9]. Pharmacokinetic interactions affect how the body processes drugs through absorption, distribution, metabolism, or elimination pathways [10]. Pharmacodynamic interactions occur when one drug affects the pharmacological response to another without altering its concentration [11]. The interactions are also classified by clinical severity, which helps determine intervention decisions. This is critical for e-commerce applications, as different severity levels will lead to different user experiences – from informational to blocking alerts. The interactions also have a significant burden on healthcare systems. Given that prescription regimens contain an average of 6.58 medications with potential for 2.68 drug interactions per patient [12], the challenge scales exponentially with polypharmacy prevalence.

2.1 Current Drug Interaction Detection systems

Existing DDI detecting systems include standalone checkers, pharmacy management platforms, and clinical decision support systems.

1) **Standalone checkers** like CVS Drug Interaction Checker, Drugs.com, Medscape, and Micromedex that are currently available to consumers, provide database driven interaction checking. Studies have found that these systems detect DDIs at varied levels, revealing significant inconsistency. They are also designed for healthcare professionals with clinical knowledge, making them inappropriate for direct consumer use.
2) **Pharmacy Management Systems** (like PrimeRx) have an integrated process of interaction checking prior to dispensing drugs to consumers. However, a study of 64 Arizona pharmacies found only 28% correctly identified eligible interactions, with median sensitivity of 0.85 even for well-documented interactions [7]. This suggests that even professional grade systems miss a substantial percentage of known interactions.
3) **Clinical Decision Support Systems (CDSS)** integrated into electronic health records provide real-time alerts during prescribing workflows. Commercial databases such as Drug Therapy Monitoring System and National Drug Data File Plus power many EHR alerting systems. However, a study found that out of the 92,272 drug orders in the year 2000, 92.2% were overridden by the physicians at the point of prescribing [13].

2.2 Machine Learning for DDI Prediction

Traditional Machine Learning approaches employed classical methods (including random forests and support vector machines) requiring manual feature engineering based on domain expertise [14]. While these approaches achieved moderate success, they were limited by the need for extensive labeled training data and manually crafted molecular descriptors. Graph Neural Networks (GNN) is

another approach for DDI prediction by directly processing molecular graph structures [15]. Unlike traditional methods requiring handcrafted features, GNNs learn representations by aggregating information from neighboring nodes in molecular graphs [16]. However, there are still significant challenges in the above methods. Current methods have insufficient learning ability for massive multi-source data and struggle with zero-shot scenarios for novel drug combinations [17]. More importantly, these models lack the ability to explain predictions to non-experts, and cannot incorporate patient health data like age, comorbidities, or dosage to determine relevance.

2.3 Large Language Models in Healthcare and Drug Discovery

Recent studies demonstrate LLMs' true potential across healthcare domains. Indeed, LLMs show extraordinary performance in general tasks, but their performance in specialized application domains is grossly limited [17]. Nonetheless, LLMs remain superior to traditional approaches in consumer-facing applications. Natural language generation capabilities help give simple explanations of complex pharmacological concepts. These technologies are also especially beneficial for merging data across various sources – drug databases, literature, and patient history.

2.4 E-commerce technology in Healthcare

The online pharmacy market is growing rapidly with multiple platforms increasingly incorporating AI technologies. AI-driven recommendation engines suggest personalized OTC products based on health profiles and purchase history. AI chatbots provide customer support and automated prescription verification (PharmbotAI). While platforms are still adopting AI for operational efficiency, online pharmacies have yet to successfully integrate safety interventions into their shopping workflows. Interaction checking is primarily a back-end pharmacy function which occurs after the consumer has already made a purchase, and no major platform provides real-time DDI alerts during cart assembly or prior to checkout. This represents a significant gap where e-commerce technology has not yet been applied to a critical safety concern. There are a few challenges in this approach:

1) Response time requirements are different – clinical systems tolerate multi-second delays while e-commerce users expect sub-500ms interactions.
2) Real-time state management is complex – interaction risk depends on the complete cart contents plus any medications users are already taking.
3) Liability concerns arise when consumer-facing systems provide medical guidance without pharmacist supervision.

3. System Design and Architecture

Here we walk through the proposed architecture of our LLM-powered prescription cart intelligence system. This section describes the high-level design, core components, data flow, and integration of touchpoints that enable real-time drug interaction detection within an e-commerce shopping experience.

This system implements a hybrid architecture that combines deterministic rule-based interaction checking with LLM-powered contextual analysis. The design balances the regulatory compliance and reliability of commercial drug databases while leveraging the natural language reasoning and explanation capabilities of large language models. This is a multi-tiered architecture consisting of 3 functional layers – the **Interaction Layer** (DDI checking using commercial drug databases), the **LLM layer** (contextual analysis, severity assessment, and natural language explanation), and the **User Experience layer** (present alerts through progressive disclosure mechanisms).

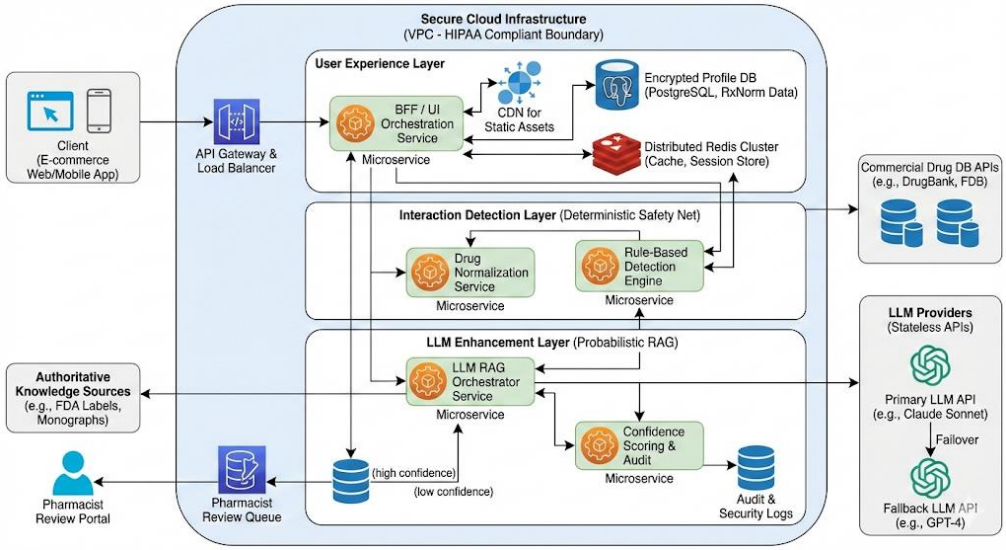

Figure 1: System Design and Architecture of LLM-Powered Prescription cart

3.1 Interaction Detection Layer (Rule-based)

This layer integrates commercial drug interaction databases that provide comprehensive documentation of DDIs to perform deterministic DDI checking.

1) Drug Database Integration – connect DrugBank's commercial API, which returns structures JSON responses with severity ratings (minor, moderate, major), interaction mechanisms, and management recommendations and can be batched up to 40 drugs [18].
2) Drug Normalization Pipeline – extract active ingredients via regular expressions, match them to RxNorm Concept Unique Identifiers, map RxCUI codes to DrugBank IDs, and normalize dosage information if required. This is required since e-commerce product catalogs have various naming conventions while interaction databases are standardized.
3) Response Time Optimization – cache drug identifier mappings, interaction results for drug in Redis or similar data stores to obtain 500ms response times. Modifications to the cart can be handled by parallel API calls, and mappings can have different TTLs (time-to-live) as required.

3.2 LLM Layer

This layer augments rule-based detection with contextual analysis using Retrieval-Augmented Generation (RAG) to ground responses in authoritative drug information while leveraging LLMs' reasoning capabilities.

1) Retrieval-Augmented Generation – optimizes LLM output by referencing an authoritative knowledge base outside training data sources before generating responses. When the system **detects** an interaction through rule-based checking, it will retrieve relevant information from DrugBank and FDA drug labels. It then **formats** the retrieved information into a structured prompt for the LLM. Finally, it **generates** consumer-appropriate explanations, severity assessments, and recommendations.
2) Prompt Engineering Framework – uses refined prompts that structure LLM tasks for maximum accuracy and reliability. The LLM **identifies** the interaction type (pharmacokinetic vs. pharmacodynamic), then **extracts** relevant patient information from cart data, and finally **synthesizes** information into consumer-appropriate language.
3) LLM API Selection and Fallback Strategy – support multiple LLM providers including Anthropic Claude, OpenAI GPT-4, and Google Gemini through a unified adapter interface. If the primary LLM API fails, the system will fall back to secondary option to ensure continuous operation. If both fail, default to rule-based interaction warnings.
4) Confidence Scoring and Uncertainty Handling – the system implements a confidence scoring mechanism to evaluate LLM responses for uncertainty markers, hallucination indicators, and consistency with respect to retrieved context. Responses are scored on a 0-1 scale, and anything scoring below 0.7 is flagged for

pharmacist review rather than being shown directly to consumers.

3.3 User Experience Layer

The UX layer displays back-end interaction detection as actionable guidance for consumers through alert mechanisms and progressive disclosure.

1) Alert System – implements severity-based notifications that are displayed near cart items. **Minor** interactions (informational, requires simple monitoring like taking medications at different time of day), **Moderate** interactions (warning, recommend pharmacist or healthcare provider consultation with contact information), and **Major** interactions (blocking checkout, required to remove item or pharmacy intervention for override).
2) Natural Language Explanation – uses LLM generated simple explanations to complex drug interactions that consumers can understand. Medical terminology is substituted with consumer-friendly alternatives ("anticoagulation" → "blood thinning," "QT prolongation" → "heart rhythm changes").
3) Progressive Disclosure – interface enables controlled display of information to avoid overwhelming users while also ensuring detailed information is available if needed. Options like “Learn More” and “Why am I seeing this?” provide secondary level of information, references, and links to data sources that help build user trust.
4) Integration with E-commerce flow – empowers consumers to view interactions at multiple steps in their journey. **Product detail** pages (potential interactions based on items in cart), **Cart** page (aggregated interaction warnings and filter on drug combinations), and **Checkout** page (final validation and blocking if required).

3.4 Data Management and Privacy

The system also effectively maintains user medication lists while protecting sensitive health information. This is critical to ensure patient safety and high-quality interaction results.

1) Medication profile management – users are recommended to list prescription and OTC medications using RxNorm codes, with optional dosage/frequency fields. Data is encrypted at rest and in transit, and the system cross checks new cart items against stored medications.
2) Privacy and Compliance – all interactions are HIPAA compliant, and no health information is shared in LLM API calls. Users can control if the system can view/edit/delete their profile data and trigger deletion across associated logs if required.
3) Audit Logging and Monitoring – all interactions are logged with timestamps, drug identifiers, LLM responses, confidence scores, and user actions for quality assurance, performance monitoring, and regulatory compliance.

3.5 Scalability and Performance

The system must handle high-volume traffic patterns typical of e-commerce platforms while maintaining high availability, consistency, and response times.

1) Microservices Architecture – each component (drug normalization, rule-based checking, LLM layer, UI API) deploys as an independent microservice (with Kubernetes) and scales automatically based on CPU utilization or queue depth. The system scales from baseline 100 requests/second to peak 1000+ requests/second within 2-3 minutes during traffic spikes.
2) Caching – enables minimizing external API calls. In-memory (within services for frequently access data), Redis cluster (to store interaction results from drug databases), and CDN caching (for static assets). The system also employs mechanisms to invalidate the caches based on schedules or event-based triggers (drug databases update).
3) Performance monitoring – use real-time dashboards to track API response time distributions (p50, p75, p95, p99), cache hit rates, LLM API latency, and user action distributions, with alerts triggering when p95 times exceed 500ms or error rates exceed 1%.

4. Proposed Evaluation Methodology

To validate the proposed system's effectiveness, a comprehensive evaluation framework must assess detection accuracy, system performance, and user experience.

The evaluation would employ standard metrics used in DDI detection research – sensitivity (ability to detect true interactions), specificity (ability to correctly identify non-interactions), positive predictive value (PPV), and negative predictive value (NPV) which are calculated from true positives (TPs), true negatives (TNs), false positives (FPs), and false negatives (FNs) [19], using the formulas in Fig. 2 below. Our hybrid system should target >90% sensitivity

for major/contraindicated interactions to surpass existing tools, while maintaining >85% specificity to minimize false positives that contribute to alert fatigue.

$$Sensitivity = TP/(TP + FN)$$

$$Specificity = TN/(FP + TN)$$

$$PPV = TP/(TP + FP)$$

$$NPV = TN/(TN + FN)$$

Figure 2: Formula to evaluate Sensitivity, Specificity, Post Predictive Value, and Negative Predictive Value [19]

System performance benchmarks must include response time distributions (p50, p95, p99) under varying loads (100-1000+ requests/second), with success defined as <500ms p95 latency and 99.9% availability. User acceptance would be assessed through controlled trials in alert comprehension rates, action distributions (removed items, contacted pharmacists, acknowledged) and cart abandonment rates during transactions with the LLM-enriched system versus rule-based baseline.

5. Discussion

This hybrid architecture overcomes major shortcomings of current drug interaction detection systems by integrating deterministic rule-based verification with LLM-driven contextual analysis, leading to advantages over the current methods.

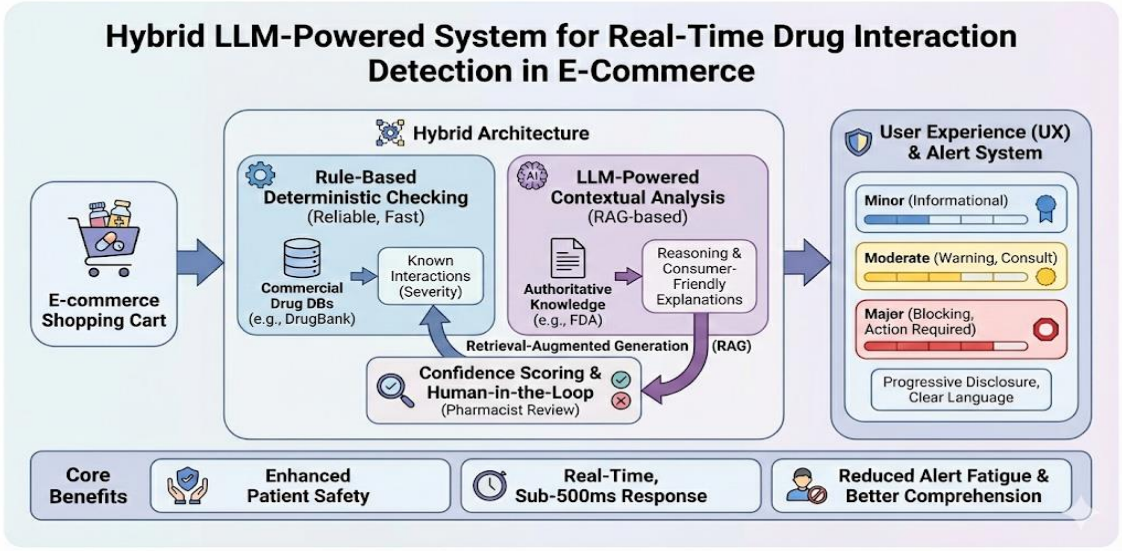

Figure 3: Core components of a hybris LLM-powered system for real-time drug interaction detection in e-commerce

In contrast to other existing Drug Interaction checkers, which require users to navigate away from shopping to manually fill out medications, our system also integrates easily with the cart-level feature, enhancing the likelihood that they will check for interaction before a purchase is made. By creating a RAG-based grounding mechanism, the hallucination problem seen in standalone LLMs can be alleviated. It anchors LLM outputs to authoritative drug databases and generates explanations that are consumer-appropriate. In addition, confidence scoring mechanics and pharmacist review queue create a human-in-the-loop safety net to alleviate the concerns of deploying AI machines in healthcare applications.

Nevertheless, the implementation still has some problems, including dependency upon commercial API availability and costs, and restricted individual context for patients without their complete medical histories. Despite these limitations, the system is an important leap for consumer-facing pharmaceutical safety by illustrating that proactive clinical decision support can be seamlessly infused into e-commerce workflows with no adverse impact on user experience or system performance.

6. Conclusion and Future Work

This paper presents a new LLM-powered drug-drug interaction detection system which is designed to be used within e-commerce pharmacy shopping carts. It helps address a critical safety gap in the online pharmacy market. The proposed hybrid architecture combines deterministic rule-based checking using commercial databases (DrugBank, First Databank) with LLM-enhanced contextual analysis through Retrieval-Augmented Generation, which helps achieve more reliable detection. The three-layer system design – comprising interaction detection, LLM enhancement, and user experience layers – demonstrates that proactive clinical decision support can be integrated into consumer-facing applications. By implementing severity-based progressive disclosure, confidence scoring mechanisms, and HIPAA-compliant data management, the system balances patient safety with user experience.

Future work should focus on empirical validation through real-world data to measure detection, accuracy, user acceptance, and potential healthcare cost savings by preventing adverse drug reactions and hospitalizations. Technical enhancements can include usage of more fine-tuned models (like Phi-3.5) that can extend to drug-food interactions. The system can also be extended to integrate with electronic health records to get access to complete medication histories and patient-specific risk factors and adding multi-language support would facilitate reach to a larger user base.

References

1. Sharath Kommu, C. Carter, and P. Whitfield, "Adverse Drug Reactions," *Nih.gov*, Jan. 10, 2024. http://ncbi.nlm.nih.gov/books/NBK599521/.
2. Research and Markets, "Online Pharmacy Market Report 2024, with Forecasts to 2030, Featuring Profiles of CVS Health, The Kroger Co., Walmart, Cigna Healthcare, DocMorris, UnitedHealth, Giant Eagle and Apollo Pharmacy," *GlobeNewswire News Room*, Nov. 12, 2024. https://www.globenewswire.com/news-release/2024/11/12/2979234/28124/en/Online-Pharmacy-Market-Report-2024.
3. K. J. Quinn and N. H. Shah, "A dataset quantifying polypharmacy in the United States," *Scientific Data*, vol. 4, Oct. 2017, doi: https://doi.org/10.1038/sdata.2017.167.
4. J. Sultana, P. Cutroneo, and G. Trifirò, "Clinical and economic burden of adverse drug reactions," Journal of Pharmacology and Pharmacotherapeutics, vol. 4, no. 5, p. 73, 2013, doi: https://doi.org/10.4103/0976-500x.120957.
5. R. Rodríguez-Monguió, M. J. Otero, and J. Rovira, "Assessing the Economic Impact of Adverse Drug Effects," *PharmacoEconomics*, vol. 21, no. 9, pp. 623–650, Jun. 2003, doi: https://doi.org/10.2165/00019053-200321090-00002.
6. R. Aparasu, R. Baer, and A. Aparasu, "Clinically important potential drug-drug interactions in outpatient settings," Research in Social and Administrative Pharmacy, vol. 3, no. 4, pp. 426–437, Dec. 2007, doi: https://doi.org/10.1016/j.sapharm.2006.12.002.
7. K. R. Saverno et al., "Ability of pharmacy clinical decision-support software to alert users about clinically important drug—drug interactions," *Journal of the American Medical Informatics Association*, vol. 18, no. 1, pp. 32–37, Jan. 2011, doi: https://doi.org/10.1136/jamia.2010.007609.
8. D. Vito, F. Ferrucci, and A. Angelakis, "LLMs for Drug-Drug Interaction Prediction: A Comprehensive Comparison," *arXiv.org, 2025*. https://arxiv.org/abs/2502.06890.
9. K. M. Giacomini, R. M. Krauss, D. M. Roden, M. Eichelbaum, M. R. Hayden, and Y. Nakamura, "When good drugs go bad," *Nature*, vol. 446, no. 7139, pp. 975–977, Apr. 2007, doi: https://doi.org/10.1038/446975a.
10. I. Cascorbi, "Drug Interactions," *Deutsches Aerzteblatt Online*, vol. 109, no. 33–34, Aug. 2012, doi: https://doi.org/10.3238/arztebl.2012.0546.
11. M. Rowland, "INTRODUCING PHARMACOKINETIC AND PHARMACODYNAMIC CONCEPTS*," Journal of Liposome Research, vol. 11, no. 4, pp. 395–422, Jan. 2001, doi: https://doi.org/10.1081/lpr-100108614.
12. J. Niu, R. M. Straubinger, and D. E. Mager, "Pharmacodynamic Drug–Drug Interactions," *Clinical Pharmacology & Therapeutics*, vol. 105, no. 6, pp. 1395–1406, Apr. 2019, doi: https://doi.org/10.1002/cpt.1434.
13. A. Alanazi, W. Alalawi, B. Aldosari, A. Alanazi, W. Alalawi, and B. Aldosari, "An Evaluation of Drug-Drug Interaction Alerts Produced by Clinical Decision Support Systems in a Tertiary Hospital," *Cureus*, vol. 15, no. 8, Aug. 2023, doi: https://doi.org/10.7759/cureus.43141.

14. Z. Fang, X. Zhang, A. Zhao, X. Li, H. Chen, and J. Li, "Recent Developments in GNNs for Drug Discovery," *arXiv.org*, 2025. https://arxiv.org/abs/2506.01302.
15. K. Abbas, C. Hao, X. Yong, M. K. Hasan, S. Islam, and A. Hadi, "Graph neural network-based drug-drug interaction prediction," *Scientific Reports*, vol. 15, no. 1, pp. 30340–30340, Aug. 2025, doi: https://doi.org/10.1038/s41598-025-12936-1.
16. Y. Chen et al., "DrugDAGT: a dual-attention graph transformer with contrastive learning improves drug-drug interaction prediction," *BMC Biology*, vol. 22, no. 1, Oct. 2024, doi: https://doi.org/10.1186/s12915-024-02030-9.
17. H. Qi, X. Li, C. Zhang, and T. Zhao, "Improving drug-drug interaction prediction via in-context learning and judging with large language models," *Frontiers in Pharmacology*, vol. 16, Jun. 2025, doi: https://doi.org/10.3389/fphar.2025.1589788.
18. "DrugBank Online," *Drugbank.com*, 2025. https://go.drugbank.com/clinical/drug_drug_interaction_checker.
19. L. A. Marcath, J. Xi, E. K. Hoylman, K. M. Kidwell, S. L. Kraft, and D. L. Hertz, "Comparison of Nine Tools for Screening Drug-Drug Interactions of Oral Oncolytics," *Journal of Oncology Practice*, vol. 14, no. 6, pp. e368–e374, Jun. 2018, doi: https://doi.org/10.1200/JOP.18.00086.