ORIGINAL RESEARCH **Open access**

Graph Neural Networks for Drug-Drug Interaction Prediction in Polypharmacy Patients: A Conceptual Framework Using Prescription Sequences and Molecular Structures

Daniel Fischer[1*], Laura Meier[1], Thomas Braun[2], Stefan Koch[2], Felix Roth[1]

Abstract

Polypharmacy, defined as the concurrent use of five or more medications, is highly prevalent among older adults and patients with multiple chronic conditions and is associated with an increased risk of drug–drug interactions (DDIs), leading to adverse drug events, hospitalizations, and higher healthcare costs. Existing DDI databases are often incomplete and fail to capture higher-order interactions, while many machine learning approaches overlook temporal prescription patterns and molecular structure information, limiting their effectiveness in real-world clinical settings. To address these limitations, this study proposes a graph neural network (GNN)-based framework that integrates prescription sequence data with molecular representations to improve DDI prediction. The model constructs a unified graph where drug nodes encode both known interactions and learned similarities, while a prescription sequence encoder captures temporal co-prescribing patterns and a molecular encoder processes SMILES-based structures. These multimodal representations are fused within a patient–drug interaction graph and refined using GNN layers with attention mechanisms to enhance interpretability. By combining longitudinal clinical data with chemical structure information, the framework enables more accurate, context-aware, and patient-specific prediction of DDIs, supports the identification of novel interactions, and improves risk stratification in polypharmacy settings, offering a scalable and interpretable foundation for future clinical decision support systems.

Keywords Multi-modal fusion, Graph neural networks, Drug-drug interactions, Polypharmacy, Prescription sequences, Molecular graphs

*Correspondence:
Daniel Fischer
daniel.fischer@outlook.com
[1] Department of AI Healthcare Engineering, University of Freiburg, Freiburg, Germany
[2] Department of Intelligent Medical Systems, Karlsruhe Institute of Technology, Karlsruhe, Germany

Introduction

Polypharmacy, typically defined as the simultaneous use of five or more medications, affects a substantial proportion of older adults and patients managing chronic diseases. Drug-drug interactions arising in such regimens contribute significantly to adverse events, hospital admissions, and elevated healthcare expenditures. Addressing these risks requires sophisticated computational methods capable of modeling complex relational data inherent in clinical practice.

Existing DDI knowledge bases, such as those compiled in structured repositories, often prove incomplete and lag behind evolving clinical evidence. Conventional machine learning techniques frequently represent drugs independently, neglecting the sequential patterns observed in electronic health records where medications are

prescribed over time. This limitation hinders the detection of context-dependent interactions that manifest uniquely in polypharmacy patients.

Graph neural networks have demonstrated strong capabilities in learning from relational structures by propagating information across connected nodes. In pharmacological applications, drugs can be modeled as nodes within graphs, with edges informed by known interactions, structural similarities, or co-occurrence frequencies derived from real-world data sources. Such representations align well with the interconnected nature of medication effects.

This paper proposes a conceptual framework for DDI prediction using GNNs that integrates two complementary data sources: (1) prescription sequences from EHRs capturing temporal ordering and co-prescribing patterns, and (2) molecular structure graphs capturing chemical similarity. The framework enables prediction of novel DDIs and patient-specific risk assessment. Subsequent sections detail the background, framework architecture, modeling components, and evaluation considerations..

Background

Drug-drug interactions and polypharmacy

Drug-drug interactions encompass pharmacokinetic alterations affecting absorption, distribution, metabolism, or excretion, as well as pharmacodynamic synergies or antagonisms at target sites. In polypharmacy patients, these interactions frequently lead to serious clinical consequences, including toxicity, reduced efficacy, or unexpected side effects that complicate disease management. Prevalence is particularly high among older adults with multimorbidity, where cumulative medication burden amplifies vulnerability.

Epidemiological evidence underscores the public health burden of unmanaged DDIs in polypharmacy contexts [1]. Healthcare systems incur considerable costs from associated adverse events, prompting the need for proactive prediction strategies. Frameworks that account for both individual drug properties and patient-level usage patterns offer a promising avenue to mitigate these risks through informed prescribing guidance [2].

Existing DDI prediction approaches

Knowledge-based approaches draw primarily from curated repositories that catalog documented interactions, yet these resources remain limited in scope and fail to capture emerging patterns in diverse patient cohorts. Similarity-based techniques leverage chemical structure comparisons to infer potential DDIs, providing a foundation for hypothesis generation when direct evidence is absent. Machine learning methods, including matrix factorization and early deep learning variants, have extended these capabilities by learning latent representations from large-scale interaction data [3].

Recent advances incorporate textual embeddings and knowledge graphs to enrich traditional similarity metrics, addressing gaps in static databases. Nevertheless, many models overlook the temporal dynamics inherent in prescription sequences, limiting their applicability to real-world polypharmacy scenarios. Hybrid strategies that combine multiple data modalities show potential for improving predictive robustness across clinical settings [4].

Table 1 provides a comparative analytical perspective highlighting how the proposed framework advances beyond existing DDI prediction paradigms.

Table 1. Comparative Analytical Framework of DDI Prediction Paradigms

Dimension	Knowledge-Based Approaches	Similarity-Based Models	Conventiona ML Models
Data Dependency	Curated DDI databases	Molecular similarity metrics	Interaction matrices + features
Temporal Awareness	None	None	Limited
Structural Understanding	Implicit	Strong	Moderate
Patient-Specific Context	Absent	Absent	Weak
Generalization to Novel Drugs	Poor	Moderate	Moderate

Higher-Order Interaction Modeling	Limited	Limited	Partial
Explainability	High (rule-based)	Moderate	Low
Scalability	Moderate	High	Moderate
Clinical Applicability	Static	Hypothesis-generating	Predictive

Graph neural networks for relational data

Graph convolutional networks (GCNs) enable efficient propagation of features across graph-structured data through localized message-passing operations, making them suitable for modeling complex drug interaction networks. Graph attention networks (GATs) further refine this process by assigning learnable weights to neighboring nodes, capturing varying degrees of relational importance in pharmacological graphs. GraphSAGE extends these ideas with inductive learning, allowing embeddings to generalize to unseen nodes and supporting scalable applications in large drug libraries [5].

Node embeddings generated through iterative updates encode both local structural context and global graph properties, facilitating downstream tasks such as link prediction. In pharmacological domains, these mechanisms have been adapted to integrate heterogeneous information sources beyond simple adjacency matrices. The resulting representations offer a flexible foundation for multi-modal fusion in DDI prediction frameworks [6].

Framework Overview

High-level architecture

The proposed architecture begins with two distinct input modalities comprising longitudinal prescription sequences from electronic health records and molecular structure graphs derived from SMILES notations. These inputs feed into a unified graph construction module that assembles nodes and edges reflecting both temporal co-prescribing patterns and chemical similarities. Subsequent GNN layers process the integrated graph to produce enriched node embeddings that encode patient-specific and structural information for final DDI prediction [7].

Output from the GNN component undergoes a pairwise interaction head that computes probabilities and types of potential DDIs for every relevant drug pair within a patient's regimen. The end-to-end design maintains differentiability throughout, enabling seamless optimization of sequence encoding, molecular representation, and graph propagation modules. This modular structure ensures that temporal and structural signals reinforce each other during inference [8].

Figure 1 illustrates the hierarchical multi-modal graph neural network architecture integrating prescription sequences and molecular structures for patient-specific DDI prediction.

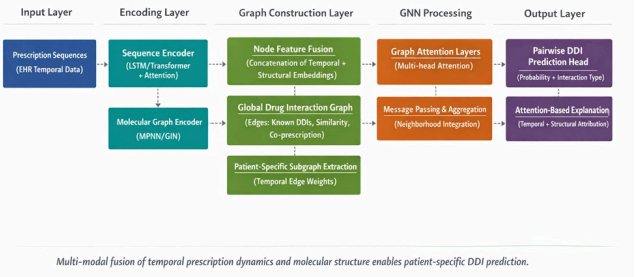

Figure 1. Hierarchical Multi-Modal GNN Architecture for Patient-Specific Drug–Drug Interaction Prediction

Core assumptions

Longitudinal EHR data containing timestamped medication orders, start and stop dates, and dosage adjustments are assumed to be available in structured format for each patient. Molecular structures in SMILES format are presumed accessible for all drugs under consideration, allowing consistent graph conversion across the dataset. Sufficient historical co-prescribing patterns exist within the EHR corpus to support meaningful learning of sequence-based embeddings [9].

The framework further assumes that known DDIs from established databases can serve as supervisory signals during training while permitting generalization to novel combinations. Patient populations exhibit polypharmacy levels typical of older adults and chronic disease cohorts, ensuring the model encounters realistic combinatorial complexity. These assumptions align with contemporary healthcare data infrastructures and enable practical deployment scenarios [1].

Design principles

Multi-modal integration stands as a foundational principle, ensuring that prescription sequence embeddings and molecular graph encodings are fused at the node level rather than handled in isolation. Temporal awareness is embedded through sequence modeling components that preserve ordering and proximity information within patient histories. Scalability to large drug sets is achieved via inductive GNN architectures that avoid reliance on fixed graph sizes or exhaustive pairwise computations [10].

Explainability is prioritized through attention mechanisms that highlight influential structural features or temporal patterns contributing to each prediction. The design also emphasizes transferability across institutions by relying on standardized data representations such as SMILES and common EHR fields. Collectively, these principles guide the framework toward clinically actionable, interpretable, and generalizable DDI risk assessment [11].

Prescription Sequence Modeling

Temporal prescription data from EHR

Medication order timestamps, start and stop dates, and dose change records form the raw input for capturing real-world prescribing dynamics in polypharmacy patients. Sequential patterns emerge naturally from these data, revealing co-occurrence frequencies, ordering dependencies, and gaps between successive prescriptions that reflect clinical decision-making. Such temporal granularity distinguishes the approach from static interaction matrices by incorporating the evolving nature of patient regimens over time [9].

Data preprocessing normalizes timestamps into relative intervals and encodes dosage trajectories to enrich the sequence representation. Patient-specific trajectories are segmented into observation windows that balance historical context with computational feasibility. The resulting structured sequences serve as direct input to downstream embedding layers within the overall framework [12].

Sequence encoding methods

Recurrent architectures such as LSTM networks process prescription sequences to generate contextual drug embeddings that reflect temporal dependencies and co-prescribing regularities. Transformer-based alternatives leverage self-attention to capture long-range interactions across extended medication histories without sequential processing bottlenecks. Hybrid attention mechanisms further refine these embeddings by weighting the contribution of specific historical prescriptions according to their relevance to current regimens [13].

Learned sequence embeddings are produced at both the individual drug and patient-graph levels, providing complementary views of temporal context. These representations are designed to be concatenated with molecular features during node initialization, ensuring multi-modal synergy. The encoding stage therefore transforms raw temporal data into rich, fixed-dimensional vectors suitable for graph-based propagation [14].

Molecular Structure Encoding

Molecular graph representation

Atoms serve as nodes and covalent bonds as edges in the molecular graph representation, with each node initialized using feature vectors that encode atom type, hybridization state, formal charge, and other chemical attributes. Conversion from SMILES strings to graph format follows standardized cheminformatics pipelines that preserve stereochemical and topological information. This graph view captures intrinsic structural properties that influence pharmacokinetic and pharmacodynamic behavior beyond simple fingerprint representations [10].

Node and edge features are augmented with additional descriptors such as aromaticity indicators and ring membership to enhance expressiveness. The resulting molecular graphs remain compact yet information-rich, enabling efficient processing by subsequent GNN encoders. Uniform representation across all drugs ensures consistent feature spaces for multi-modal integration within the patient-drug graph [15].

Molecular graph encoding

Message-passing neural networks (MPNNs) and graph isomorphism networks (GINs) iteratively update atom-level representations by aggregating information from neighboring atoms and bonds. The final readout function pools these atom embeddings into a fixed-size drug-level vector that encodes global chemical properties relevant to interaction potential. This encoding step produces

embeddings that are inherently inductive and transferable to novel molecular structures not encountered during training [11].

Learned molecular embeddings capture substructural motifs and functional groups known to mediate specific DDI mechanisms, providing a complementary signal to sequence-derived context. The output vectors are aligned dimensionally with sequence embeddings to facilitate straightforward concatenation at the node initialization stage. Consequently, the molecular encoder contributes chemically grounded features that ground temporal patterns in underlying structural reality [7].

Graph Construction

Drug interaction graph

Drugs function as nodes within the global interaction graph, where each node aggregates features from both prescription sequence embeddings and molecular structure encodings. Edges are instantiated using multiple criteria, including documented interactions drawn from established pharmacological databases, quantitative measures of structural similarity computed from molecular graphs, and empirical co-prescribing frequencies extracted from aggregated electronic health record data. This multi-source edge construction creates a heterogeneous graph that reflects both explicit clinical knowledge and latent relational signals inherent to polypharmacy regimens [16].

The resulting graph topology captures higher-order connectivity patterns that extend beyond simple pairwise links, enabling the model to propagate information across clusters of frequently co-prescribed medications. Structural similarity edges derived from molecular fingerprints further enrich the adjacency matrix by linking drugs that share substructures associated with similar interaction mechanisms. Consequently, the drug interaction graph serves as a foundational scaffold for subsequent GNN processing that integrates temporal and chemical perspectives [17].

Patient-specific subgraph

For each individual patient, a subgraph is induced by extracting only those nodes and edges corresponding to the current active medication set identified in the longitudinal prescription record. Temporal weights are assigned to edges based on the recency and duration of overlapping prescriptions, thereby emphasizing interactions that are clinically proximate within the patient's treatment trajectory. This patient-centric view transforms the global graph into a personalized representation that accounts for the unique combinatorial context of polypharmacy [12].

Subgraph induction is performed dynamically at inference time to accommodate evolving regimens without retraining the entire model. Edge weights derived from sequence data modulate message passing so that recent co-prescribing patterns exert stronger influence on node updates. The approach ensures that predictions remain anchored in the patient's actual medication history while still benefiting from the broader structural knowledge encoded in the global graph [18-20].

Edge feature engineering

Edge features are engineered as multi-type vectors that differentiate known DDIs, structural similarity scores, and temporal proximity metrics, allowing the model to distinguish between evidence-based and inferred relationships. Attention coefficients are computed over these heterogeneous edge types during graph propagation, enabling the network to learn the relative importance of each relation category for a given prediction task. This explicit encoding of edge semantics prevents information loss that occurs when graphs rely solely on binary adjacency [21].

Additional edge attributes, such as dosage overlap indicators or metabolic pathway co-occurrence flags, are concatenated to further enrich the feature space for message passing. The engineering process maintains compatibility with standard GNN layers by projecting diverse edge features into a unified embedding dimension. Ultimately, these engineered edges facilitate more nuanced interaction modeling that aligns with the multi-modal design principles of the framework [22].

Graph neural network architecture

Node feature initialization

Node features are initialized by concatenating the sequence-derived embedding, the molecular graph embedding, and static drug properties such as therapeutic class, half-life, and primary metabolic pathways. This initialization strategy ensures that each drug node encodes

complementary signals from temporal usage patterns and intrinsic chemical composition before any propagation occurs. The concatenated vector is passed through a linear projection layer to produce a unified initial representation suitable for subsequent GNN layers [23].

Dimensional alignment across modalities is achieved through learned projection matrices that preserve information while reducing redundancy. Static properties serve as auxiliary features that anchor the learned embeddings in established pharmacological knowledge. The resulting node initialization therefore provides a rich starting point that reflects both patient context and molecular reality [24].

Graph attention layers

Multi-head attention mechanisms operate over the neighboring nodes within each layer, assigning learnable importance scores that reflect the varying relevance of different interaction types encoded in the edge features. Graph attention network variants allow the model to focus selectively on structurally similar neighbors or those linked by recent co-prescribing patterns, thereby capturing context-dependent influences. The attention heads are aggregated through concatenation or averaging to produce updated node representations that integrate diverse relational signals [6].

Layer-wise attention coefficients are normalized across all incoming edges to ensure stable training dynamics across graphs of varying density. This design choice enables the architecture to adaptively weigh contributions from molecular similarity versus temporal proximity depending on the specific drug pair under consideration. The outcome is a set of refined node embeddings that encode both local neighborhood context and global graph topology [25].

Message passing and aggregation

Iterative message passing updates each node by aggregating transformed features from its neighbors according to the learned attention weights, allowing information to flow across the patient-specific subgraph. Mean, max, and attention-based pooling operations are applied at selected layers to derive higher-level representations that summarize subgraph-level properties relevant to DDI risk. The multi-layer design progressively refines embeddings by propagating increasingly abstract interaction patterns through the graph [8].

Residual connections between layers preserve initial node features while permitting deep propagation without gradient degradation. Aggregation functions are chosen to balance expressiveness with computational efficiency, ensuring scalability to large polypharmacy regimens. The final aggregated node states therefore encapsulate both fine-grained structural details and coarse-grained relational summaries essential for accurate prediction [10].

Table 2 delineates the functional interplay between framework components and their respective contributions to DDI prediction.

Table 2. Multi-Modal Component Interactions and Functional Contributions in the Proposed Framework

Component	Input Data Type	Representation Learned	Functiona Role
Prescription Sequence Encoder	Temporal EHR data	Contextual sequence embeddings	Captures ordering, c prescribin patterns
Molecular Graph Encoder	SMILES-derived graphs	Structural embeddings	Encodes chemical properties and substructur
Node Feature Fusion	Multi-modal vectors	Unified drug representation	Aligns temporal ar structural signals
Drug Interaction Graph	Multi-source edges	Heterogeneous graph topology	Represent known + inferred relationshi
Patient Subgraph Module	Patient-specific prescriptions	Contextualized subgraph	Personalize global grap
Graph Attention Layers	Node + edge features	Weighted relational embeddings	Prioritizes relevant neighbors
Message Passing Layers	Graph structure	Aggregated node embeddings	Propagate higher-orde interaction

Prediction Head	Node embeddings	Interaction probabilities	Computes pairwise DD
Explanation Module	Attention weights	Attribution maps	Provides interpretabil

Pairwise DDI prediction

For every drug pair within the patient subgraph, node embeddings are combined through a bilinear interaction module or a multi-layer perceptron that operates on their concatenated representations augmented by edge features. The prediction head outputs both a probability score for interaction occurrence and a categorical label indicating the expected interaction type, such as pharmacokinetic or pharmacodynamic. This pairwise formulation transforms graph-level representations into clinically actionable DDI forecasts [11].

Softmax normalization across possible interaction categories provides calibrated confidence estimates that can inform clinical prioritization. The architecture supports end-to-end differentiability, allowing gradients to flow back through the attention layers and modality encoders during conceptual training. Consequently, the pairwise prediction stage completes the mapping from raw inputs to interpretable DDI outputs [26].

DDI Prediction and Interpretation

Predicting novel interactions

Zero-shot prediction for previously unseen drug combinations is enabled by relying on the inductive capabilities of the molecular graph encoder, which generalizes structural embeddings to novel compounds not present in the training set. Transfer learning from molecular features allows the model to infer potential interactions even when prescription sequence data for the new drug are sparse. This capacity addresses a critical limitation of knowledge-based systems that cannot extrapolate beyond curated databases [7].

The framework leverages shared substructures across molecular graphs to propagate learned interaction patterns to chemically related but untested drugs. Patient-specific subgraph induction further contextualizes these structural inferences within the individual's existing regimen. As a result, the model can generate plausible hypotheses for novel DDIs that emerge in evolving polypharmacy scenarios [15].

Attention-based explanation

Attention weights computed during graph propagation highlight which structural features or temporal patterns most strongly influenced each predicted interaction, offering direct insight into the model's reasoning process. Visualization tools map these weights back to specific atoms in molecular graphs or to particular prescription events in the patient sequence, producing human-interpretable explanations. Such transparency supports clinician trust and facilitates regulatory review of AI-assisted decision tools [13].

Post-hoc analysis of attention distributions across multiple patient cases can reveal recurring motifs, such as specific functional groups or sequencing patterns associated with elevated risk. The explanation module therefore bridges the gap between black-box graph predictions and actionable clinical knowledge. This interpretability feature distinguishes the framework from opaque deep learning alternatives commonly applied to DDI tasks [14].

Table 3 synthesizes how different data modalities and learning mechanisms jointly contribute to clinical value in the proposed framework.

Table 3. Analytical Mapping between Data Modalities, Learning Mechanisms, and Clinical Value

Data Modality	Learning Mechanism	Type of Signal Captured	Unique Analytica Advantag
Prescription Sequences	LSTM / Transformer	Temporal ordering, co-prescription patterns	Captures dynamic treatment evolution
Molecular Structures	Graph Neural Networks (MPNN/GIN)	Chemical substructures, functional groups	Enables mechanisti reasoning
Co-prescription Networks	Graph construction	Empirical association patterns	Reflects rea world

			prescribing behavior
Known DDI Databases	Supervised learning signals	Verified interaction knowledge	Anchors mo in clinical evidence
Attention Mechanisms	Weighted aggregation	Feature importance attribution	Enhances interpretabil
Multi-modal Fusion	Feature concatenation + projection	Integrated representation	Combines complement signals
Patient Subgraphs	Dynamic graph extraction	Individualized drug context	Personalize predictions

Evaluation Strategy

Metrics for DDI prediction

Area under the receiver operating characteristic curve and area under the precision-recall curve serve as primary metrics for assessing binary interaction detection performance across varying decision thresholds. Hits@k evaluation quantifies the model's ability to rank novel DDIs among the top candidates for clinical review, emphasizing discovery of previously undocumented interactions. Stratified accuracy metrics by interaction type, including contraindication, dose adjustment, or enhanced monitoring recommendations, provide granular insight into clinical utility [3].

Macro-averaged F1 scores across interaction categories account for class imbalance inherent in polypharmacy data, where most drug pairs exhibit no clinically significant interaction. These metrics are computed at both the global graph level and the patient-specific subgraph level to validate multi-scale performance. Collectively, the evaluation suite ensures comprehensive assessment of predictive accuracy and practical relevance [4].

Validation protocols

Time-based splitting partitions the longitudinal EHR data such that models are trained on earlier prescription sequences and evaluated on more recent records, simulating prospective deployment in evolving clinical environments. Cross-database validation leverages independent repositories to confirm generalizability beyond any single institution's prescribing patterns. Prospective silent mode testing within live EHR systems allows the framework to generate predictions without influencing care while accumulating real-world outcome labels for subsequent refinement [9].

Hold-out cohorts stratified by age, comorbidity burden, and polypharmacy intensity further stress-test robustness across diverse patient subpopulations. External validation on public benchmark datasets complements internal time-split experiments by providing standardized comparison points. These protocols collectively establish the framework's readiness for integration into clinical decision support workflows [1].

Conclusion

The conceptual framework synthesizes graph neural networks with prescription sequence modeling and molecular structure encoding to address DDI prediction in polypharmacy patients. By constructing a unified graph that fuses temporal usage patterns from electronic health records with chemically grounded embeddings, the architecture provides a comprehensive view of interaction risk. This integrated approach moves beyond isolated drug representations to capture the combinatorial complexity inherent in real-world medication regimens.

Key advantages include the incorporation of temporal context that reflects actual prescribing dynamics, the exploitation of structural similarity for generalization to novel compounds, and the provision of attention-based explanations that enhance clinical interpretability. The multi-modal design principles ensure that sequence-derived and molecular signals reinforce each other, yielding more robust and context-aware predictions than unimodal alternatives. Scalability through inductive GNN layers further positions the framework for deployment across large healthcare systems.

Limitations arise from the requirement for structured longitudinal EHR data, the computational demands of deep graph propagation on dense patient subgraphs, and the challenges associated with rare drugs that lack sufficient historical co-prescribing information. Data quality issues, such as incomplete timestamp records or missing molecular annotations, could also affect embedding fidelity in certain settings. Future extensions may incorporate

additional modalities or federated learning strategies to mitigate these constraints while preserving privacy.

Implementation on public datasets such as MIMIC, SIDER, and DrugBank, followed by integration into clinical decision support systems, represents the logical next step toward real-world impact. Collaborative efforts between AI researchers, pharmacologists, and health informaticians will be essential to refine and validate the framework across diverse populations. Ultimately, this GNN-based conceptual approach offers a pathway to more personalized and proactive DDI risk assessment that could meaningfully improve medication safety in polypharmacy care.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 07 Aug 2021 Revised: 25 Oct 2021 Accepted: 04 Jan 2022
Published online: 20 July 2022

References

Ouchi K, Lindvall C, Chai PR, Boyer EW. Machine learning to predict, detect, and intervene older adults vulnerable for adverse drug events in the emergency department. J Med Toxicol. 2018;14(3):248-52. https://doi.org/10.1007/s13181-018-0674-4.

Al-Rabeah MH, Lakizadeh A. Prediction of drug-drug interaction events using graph neural networks based feature extraction. Sci Rep. 2022;12(1):15590. https://doi.org/10.1038/s41598-022-19895-8.

Wang M, Wang H, Liu X, Ma X, Wang B. Drug-drug interaction predictions via knowledge graph and text embedding: instrument validation study. JMIR Med Inform. 2021;9(6):e28277. https://doi.org/10.2196/28277.

Karim MR, Cochez M, Jares JB, Uddin M, Beyan O, Decker S. Drug-drug interaction prediction based on knowledge graph embeddings and convolutional-LSTM network. In: Proc 10th ACM Int Conf Bioinform Comput Biol Health Inform. 2019. p. 113-23. https://doi.org/10.1145/3307339.3342161.

Lin X, Quan Z, Wang ZJ, Ma T, Zeng X. KGNN: knowledge graph neural network for drug-drug interaction prediction. In: Proc IJCAI. 2020. p. 2739-45.

Hong Y, Luo P, Jin S, Liu X. LaGAT: link-aware graph attention network for drug-drug interaction prediction. Bioinformatics. 2022;38(24):5406-12.

Deng Y, Xu X, Qiu Y, Xia J, Zhang W, Liu S. A multimodal deep learning framework for predicting drug-drug interaction events. Bioinformatics. 2020;36(15):4316-22.

Yu Y, Huang K, Zhang C, Glass LM, Sun J, Xiao C. SumGNN: multi-typed drug interaction prediction via efficient knowledge graph summarization. Bioinformatics. 2021;37(18):2988-95.

Zhang W, Kuang Z, Peissig P, Page D. Adverse drug reaction discovery from electronic health records with deep neural networks. In: Proc ACM Conf Health Inference Learn. 2020. p. 30-39. https://doi.org/10.1145/3368555.3384453.

Han X, Xie R, Li X, Li J. SmileGNN: drug-drug interaction prediction based on the SMILES and graph neural network. Life (Basel). 2022;12(2):319. https://doi.org/10.3390/life12020319.

Feng YH, Zhang SW. Prediction of drug-drug interaction using an attention-based graph neural network on drug molecular graphs. Molecules. 2022;27(9):3004. https://doi.org/10.3390/molecules27093004.

Olekhnovitch R, Hoertel N, Limosin F, Blanco C, Olfson M, Fagot-Campagna A, et al. Using filled prescription sequences to rank antidepressants according to their acceptability in the general population: the Constances cohort. J Psychiatr Res. 2020;123:72-80. https://doi.org/10.1016/j.jpsychires.2020.01.019.

Asada M, Miwa M, Sasaki Y. Extracting drug-drug interactions with attention CNNs. In: Proc BioNLP 2017. 2017. p. 9-18.

Liu S, Zhang Y, Cui Y, Qiu Y, Deng Y, Zhang Z, et al. Enhancing drug-drug interaction prediction using deep attention neural networks. IEEE/ACM Trans Comput Biol Bioinform. 2023;20(2):976-85. https://doi.org/10.1109/TCBB.2022.3172215.

Yang Z, Zhong W, Lv Q, Chen CY. Learning size-adaptive molecular substructures for explainable drug-drug interaction prediction by substructure-aware graph neural network. Chem Sci. 2022;13(29):8693-703. https://doi.org/10.1039/D2SC02218D.

Rodrigues AD, editor. Drug-Drug Interactions. Boca Raton (FL): CRC Press; 2019.

Niu J, Straubinger RM, Mager DE. Pharmacodynamic drug-drug interactions. Clin Pharmacol Ther. 2019;105(6):1395-406. https://doi.org/10.1002/cpt.1434.

Li Z, Liu F, Yang W, Peng S, Zhou J. A survey of convolutional neural networks: analysis, applications, and prospects. IEEE Trans Neural Netw Learn Syst. 2022;33(12):6999-7019. https://doi.org/10.1109/TNNLS.2021.3084827.

Patton K, Borshoff DC. Adverse drug reactions. Anaesthesia. 2018;73(Suppl 1):76-84. https://doi.org/10.1111/anae.14137.

Vazquez SR. Drug-drug interactions in an era of multiple anticoagulants: a focus on clinically relevant drug interactions. Blood. 2018;132(21):2230-9. https://doi.org/10.1182/blood-2018-06-848747.

He C, Liu Y, Li H, Zhang H, Mao Y, Qin X, et al. Multi-type feature fusion based on graph neural network for drug-drug interaction prediction. BMC Bioinformatics. 2022;23(1):224. https://doi.org/10.1186/s12859-022-04761-z.

Carpenter M, Berry H, Pelletier AL. Clinically relevant drug-drug interactions in primary care. Am Fam Physician. 2019;99(9):558-64.

Zhang T, Leng J, Liu Y. Deep learning for drug-drug interaction extraction from the literature: a review. Brief Bioinform. 2021;21(5):1609-27.

Lyu T, Gao J, Tian L, Li Z, Zhang P, Zhang J. MDNN: a multimodal deep neural network for predicting drug-drug interaction events. In: Proc IJCAI. 2021. p. 3536-42.

Jin W, Barzilay R, Jaakkola T. Hierarchical generation of molecular graphs using structural motifs. In: Proc Int Conf Mach Learn. 2020;119:4839-48.

Sudsakorn S, Bahadduri P, Fretland J, Lu C. 2020 FDA drug-drug interaction guidance: a comparison analysis and action plan by pharmaceutical industrial scientists. Curr Drug Metab. 2020;21(6):403-26. https://doi.org/10.2174/1389200221666200622094105.