MediTrust – Intelligent Hospital and Drug Interaction Platform

Manasa[1], Mithra Jayapal[2], Rithu Mary Varghese[3], Teresa Benny[4], A. S Anitha[5]
[1,2,3,4]*Student, Vidya Academy of Science and Technology, Thrissur*
[5]*Assistant Professor, Vidya Academy of Science and Technology, Thrissur*

***Abstract*—Drug–drug interactions (DDIs) are a major con- tributor to adverse drug reactions and pose a serious challenge to patient safety in modern healthcare systems, particularly in scenarios involving polypharmacy. Conventional hospital man- agement platforms primarily focus on appointments and record handling, while often lacking intelligent mechanisms for real- time drug safety analysis and personalized clinical support. As a result, potentially harmful medication combinations and abnormal patient conditions may go undetected until adverse outcomes occur.**

This paper presents MediTrust, an intelligent hospital and drug interaction platform that integrates machine learning–based DDI prediction with patient vital monitoring and automated consultation support. The proposed system enables users to log prescribed medicines and dosages, upon which trained machine learning models analyze possible interaction risks, associated side effects, and severity levels using publicly available drug–drug interaction datasets. In addition, patient-entered vital parameters such as blood pressure, blood sugar level, body temperature, and heart rate are incorporated to provide contextual and personalized safety insights.

To further enhance accessibility and decision support, an AI- driven conversational assistant is employed to deliver health guid- ance in simple and understandable language, while facilitating doctor–patient communication through appointment scheduling and secure chat. Experimental evaluation demonstrates that the integration of intelligent drug interaction analysis within a hospital workflow can significantly improve medication safety and clinical awareness. The proposed platform bridges the gap between drug safety research and real-world healthcare delivery, promoting proactive, data-driven, and patient-centric treatment.

***Index Terms*—Adverse Drug Reaction, Drug–Drug Interaction, Machine Learning, Graph-Based Drug Analysis, Intelligent Hos- pital Systems, Clinical Decision Support**

I. INTRODUCTION

Adverse Drug Reactions (ADRs) represent a major global healthcare concern, contributing significantly to patient mor- bidity, mortality, prolonged hospital stays, and increased med- ical costs. A substantial proportion of these reactions arise due to Drug–Drug Interactions (DDIs), particularly in patients un- dergoing polypharmacy for chronic or complex medical condi- tions. With the growing aging population and the widespread use of multiple medications, the risk associated with unsafe drug combinations has increased considerably, making early detection and prevention of ADRs a critical requirement in modern healthcare systems.

Despite advances in medical informatics, most conventional hospital management systems primarily focus on administra- tive functionalities such as appointment scheduling, electronic health records, and billing. These systems often lack intelligent mechanisms to analyze prescribed medications for potential interactions or to incorporate patient-specific physiological parameters into clinical decision-making. As a result, harmful drug interactions are frequently identified only after adverse symptoms are reported, limiting the opportunity for timely intervention and prevention.

Recent developments in Artificial Intelligence (AI) and Ma- chine Learning (ML) have demonstrated significant potential in addressing complex biomedical challenges, including drug safety analysis. In particular, graph-based learning methods have gained attention due to their ability to model relational and structural information. Drug molecules can be naturally represented as graphs, where atoms correspond to nodes and chemical bonds correspond to edges. Graph Neural Networks (GNNs) are well-

suited for learning from such molecular representations, enabling the capture of spatial, chemical, and topological properties that are difficult to model using tra- ditional feature-based approaches. Consequently, GNN-based models have shown promising performance in predicting drug interactions and associated side effects.

While several studies have explored DDI prediction us- ing deep learning models, many existing approaches remain confined to standalone prediction tasks and lack integration into real-world clinical environments. Moreover, most systems do not consider patient-specific contextual information such as vital signs, dosage levels, or real-time clinical conditions, which are crucial for personalized risk assessment. This gap between theoretical drug interaction research and practical hospital deployment limits the impact of AI-driven solutions in everyday healthcare practice.

To address these limitations, this work proposes an intelli- gent hospital and drug interaction platform that integrates ma- chine learning-based DDI prediction within a comprehensive clinical workflow. By combining molecular graph-based drug analysis with patient-entered vital parameters and AI-assisted consultation support, the proposed system aims to enhance medication safety and clinical awareness. The integration of drug interaction prediction directly into hospital operations enables proactive decision support for both patients and health- care professionals, thereby promoting safer, data-driven, and patient-centric treatment.

II. RELATED WORK

The problem of predicting Adverse Drug Reactions (ADRs) caused by Drug–Drug Interactions (DDIs) has been widely studied due to its critical impact on patient safety and health- care outcomes. Early research efforts primarily relied on rule- based systems and curated pharmacological knowledge bases such as DrugBank and clinical guidelines. Although these approaches provided basic interaction alerts, they required ex- tensive manual updates and were unable to generalize to newly introduced drugs or previously unseen drug combinations.

With the availability of large-scale biomedical datasets, tra- ditional Machine Learning (ML) techniques were introduced for DDI prediction. Methods based on decision trees, support vector machines, logistic regression, and k-nearest neighbors utilized handcrafted features derived from chemical structures, therapeutic categories, and genomic information. While these models improved prediction performance compared to rule- based systems, their reliance on manual feature engineering limited scalability and reduced effectiveness when dealing with complex molecular relationships.

The emergence of deep learning significantly advanced drug interaction research by enabling automatic feature learning from raw data. Neural network models, including multilayer perceptrons and recurrent neural networks, were applied to drug interaction prediction using molecular fingerprints and sequence-based representations. However, these models often treated drugs as independent entities and failed to explicitly model the relational and structural properties inherent in molecular interactions.

Graph-based learning approaches addressed these limita- tions by representing drugs and their interactions as graphs. Graph Neural Networks (GNNs) have demonstrated strong performance in molecular property prediction, drug discovery, and interaction analysis. By modeling drugs as molecular graphs with atoms as nodes and bonds as edges, GNNs can capture spatial and chemical dependencies that are essential for accurate DDI prediction. Several studies have reported improved accuracy using graph convolutional networks, graph attention networks, and message-passing neural networks for predicting adverse drug interactions and side effects.

Recent research has further explored advanced techniques such as self-supervised learning, transfer learning, and en- semble models to enhance GNN performance. Pretraining molecular graphs using autoencoders or contrastive learning has been shown to improve representation quality and gener- alization, particularly when labeled data are limited. Ensemble strategies combining multiple graph-based models have also demonstrated increased robustness and predictive accuracy.

Despite these advances, most existing DDI prediction meth- ods remain limited to offline analysis and are rarely integrated into real-world clinical systems. Moreover, many approaches do not incorporate patient-specific contextual information such as dosage levels or vital signs, which are crucial for per- sonalized risk assessment. Addressing these

limitations, the present work focuses on integrating machine learning-based
DDI prediction into an intelligent hospital platform, bridging the gap between theoretical drug safety research and practical clinical deployment.

III. PROPOSED SYSTEM

This work proposes an intelligent healthcare platform, named MediTrust, which integrates drug–drug interaction pre- diction with hospital management functionalities to enhance medication safety and clinical decision support. The system combines machine learning techniques with real-time patient data and AI-assisted consultation to provide a comprehensive, data-driven healthcare solution. The overall architecture and functional components of the proposed system are described in the following subsections.

A. Overall System Architecture

The proposed system follows a modular architecture con- sisting of four primary user roles: Admin, Doctor, Patient (User), and an AI Assistant. The admin module is respon- sible for managing hospitals, doctors, users, and system- level configurations. The Doctor module enables healthcare professionals to view patient vitals, review predicted drug interaction risks, prescribe medications, and communicate with patients. The Patient module allows users to register, log prescribed medicines and dosages, monitor vital parameters, book appointments, and interact with doctors.
All modules are interconnected through a centralized back- end that processes medical data, executes drug interaction pre- dictions, and ensures secure data exchange. This architecture ensures scalability, role-based access control, and seamless integration of clinical workflows with intelligent drug safety analysis.

B. Molecular Graph Construction

To enable accurate prediction of drug–drug interactions, drugs are represented in the form of molecular graphs. Each drug is first obtained in the SMILES (Simplified Molecular Input Line Entry System) notation, which encodes its chem- ical structure. The SMILES strings are then converted into graph representations, where atoms are modeled as nodes and chemical bonds are modeled as edges.
Each node is associated with atom-level features such as atomic number, degree, valence, and aromaticity, while edges represent bond types and connectivity. This graph-based rep- resentation preserves the structural and chemical properties of drugs, enabling effective learning of interaction patterns by graph-based models.

C. Graph Neural Network Model

The constructed molecular graphs are used as inputs to a Graph Neural Network (GNN) model for predicting drug–drug interactions and potential adverse effects. The GNN aggregates information from neighboring nodes through multiple graph convolution layers, allowing the model to capture local and global molecular structures.
The learned embeddings from individual drug graphs are combined to model interaction dynamics between drug pairs.

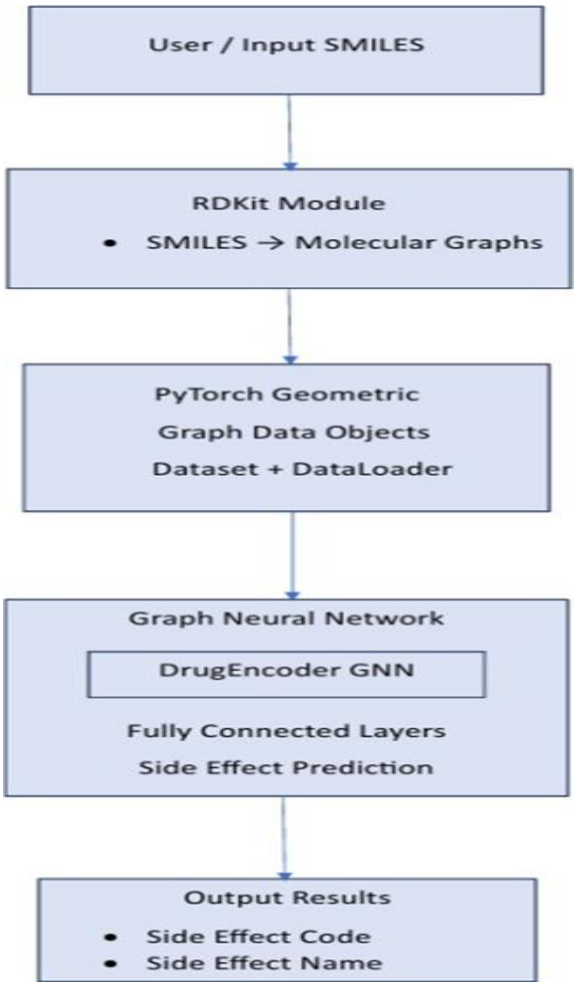

Fig. 1. System Architecture

The final representation is passed through fully connected layers to predict interaction risks, side effects, and severity levels. This approach enables the model to learn complex chemical relationships that are difficult to capture using tradi- tional feature-based methods.

D. Workflow of the Proposed System

The operational workflow of the proposed system begins when a patient enters prescribed medication details and dosage information through the user interface. The system retrieves the corresponding molecular representations of the drugs and performs drug–drug interaction analysis using the trained GNN model.

Simultaneously, patient vital parameters such as blood pres- sure, blood sugar level, body temperature, and heart rate are collected. These vitals are analyzed alongside the predicted interaction risks to provide context-aware safety insights. The prediction results are made available to both patients and doctors through an intuitive dashboard, supporting informed clinical decisions.

E. Vital-Based Contextual Safety Analysis

Unlike conventional DDI prediction systems, the proposed platform incorporates patient-specific vital parameters to en- hance risk assessment. Abnormal vital readings may amplify the severity of certain drug interactions, and therefore are considered as contextual factors in safety analysis.

By integrating vitals with drug interaction predictions, the system provides personalized alerts and recommendations, en- abling early identification of high-risk scenarios and promoting proactive medical intervention.

F. AI-Assisted Consultation Support

To improve accessibility and user engagement, an AI-driven conversational assistant is integrated into the system. The as- sistant provides medication guidance, explains interaction risks in simple language, and responds to health-related queries. It also supports appointment scheduling and directs users to seek medical attention when abnormal patterns are detected.

This AI-assisted interface enhances patient understanding, reduces dependence on manual interpretation of medical data, and supports efficient doctor–patient communication within the platform.

IV. DATASET DESCRIPTION

The performance and reliability of any drug–drug inter- action prediction model largely depend on the quality and diversity of the dataset used for training and evaluation. In this work, publicly available drug interaction datasets are utilized to ensure transparency, reproducibility, and relevance to real- world clinical scenarios. The datasets provide comprehensive information about interacting drug pairs and their associated adverse effects.

A. Dataset Source

The primary dataset used in this study is obtained from pub- licly accessible drug–drug interaction repositories, including the Kaggle Drug–Drug Interaction dataset and curated phar- macological databases such as DrugBank and TwoSides. These datasets contain verified interaction records derived from clinical reports, pharmacological studies, and post-marketing surveillance data. The use of well-established datasets ensures that the proposed model is trained on reliable and clinically meaningful interaction patterns.

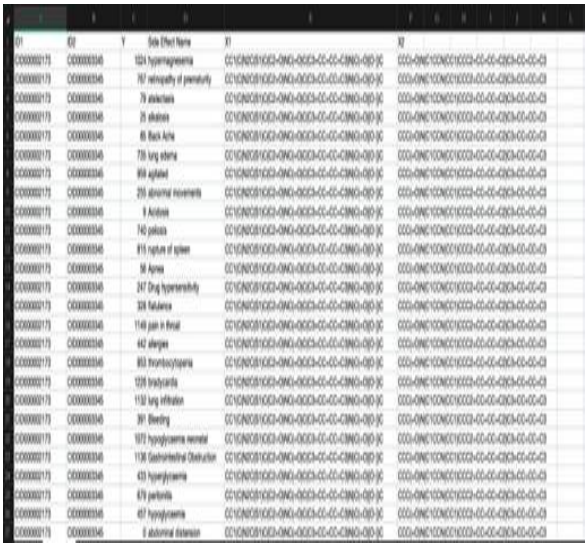

Fig. 2. Dataset

B. Drug Pair Representation

Each data sample in the dataset consists of a pair of drugs that may interact when administered together. Every drug is uniquely identified and represented using

its chemical structure in the form of SMILES (Simplified Molecular Input Line Entry System) notation. These SMILES strings serve as the basis for constructing molecular graphs, where atoms are represented as nodes and chemical bonds are represented as edges. By pairing two drug graphs, the dataset enables the model to learn interaction dynamics between different molecular structures.

C. Side-Effect Labels

The dataset is annotated with interaction outcomes in the form of side-effect labels. Each drug pair is associated with one or more adverse effects that may arise due to their com- bined usage, making the problem a multi-label classification task. Side-effect labels include clinically relevant outcomes such as dizziness, nausea, cardiac irregularities, and other physiological complications. These labels allow the model to predict not only the presence of an interaction but also the potential severity and nature of the adverse reaction.

D. Data Preprocessing

Several preprocessing steps are performed to prepare the dataset for model training. Initially, duplicate drug interaction records are removed to avoid data redundancy and bias. Incomplete entries with missing chemical structures or labels are filtered out to ensure data consistency. The SMILES strings are then converted into molecular graph representations using cheminformatics tools, enabling graph-based learning. To improve training stability and generalization, the dataset is further normalized and split into training, validation, and testing subsets. Class imbalance, which is common in medical datasets, is addressed through appropriate sampling strategies. These preprocessing steps ensure that the dataset is well- structured, balanced, and suitable for effective learning using Graph Neural Networks.

V. METHODOLOGY

This section describes the complete methodology adopted for predicting adverse drug reactions caused by drug–drug interactions using Graph Neural Networks (GNNs). The pro- posed approach follows a structured pipeline starting from dataset preparation and molecular graph construction to model

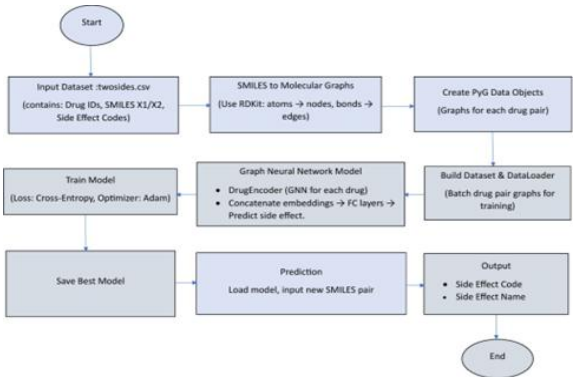

Fig. 3. Work Flow Diagram

training and prediction. Fig 3 illustrates the overall workflow of the system.

A. Input Dataset

The primary input to the system is the TwoSides drug–drug interaction dataset, provided in CSV format. The dataset contains drug identifiers, SMILES representations of two inter- acting drugs, and corresponding side-effect codes. Each record represents a potential interaction between a pair of drugs along with the observed adverse reaction. This dataset serves as a reliable benchmark for learning drug interaction patterns and associated side effects.

B. Feature Extraction

Feature extraction is performed at the molecular level using chemical structure information derived from SMILES notation. Each SMILES string is processed using the RDKit library to extract atom-level and bond-level features. Atom features include atomic number, degree, valence, aromaticity, and hybridization state, while bond features include bond type and bond order. These features provide a rich chemical representation required for effective learning of molecular interactions.

C. Graph Representation

Each drug molecule is represented as a molecular graph, where atoms are modeled as nodes and chemical bonds are modeled as edges. Using RDKit, SMILES strings are con- verted into graph structures. The resulting graphs preserve the topological and chemical properties of the molecules.

To facilitate deep learning, the molecular graphs are con- verted into PyTorch Geometric (PyG) data objects. For every drug pair, two individual molecular graphs are created and grouped together as a single

training sample, enabling the model to learn interaction dynamics between the two drugs.

D. Dataset Construction and DataLoader

After graph construction, the dataset is organized into a format suitable for batch processing. A custom dataset class is implemented to load pairs of molecular graphs along with their corresponding side-effect labels. PyTorch Geometric's DataLoader is used to batch multiple drug-pair graphs during training, enabling efficient memory usage and faster convergence

E. Graph Neural Network Model

The core of the proposed system is a Graph Neural Network model designed to encode molecular information from both drugs. Each drug molecule is processed independently using a shared GNN-based drug encoder. The encoder consists of multiple graph convolution layers that aggregate neighborhood information to generate node-level embeddings, which are subsequently pooled to obtain graph-level embeddings.

The embeddings obtained from both drug graphs are con- catenated and passed through fully connected layers. These layers learn the interaction patterns between the two drugs and produce a final output representing the probability of each side effect. This design allows the model to capture both individual drug characteristics and interaction-specific effects.

F. Model Training

The model is trained using supervised learning. During training, batches of drug-pair graphs are forwarded through the network, and predictions are compared with ground truth side- effect labels. The Adam optimizer is employed for parameter optimization due to its efficiency and adaptive learning rate properties. Model training is conducted over multiple epochs until convergence is achieved.

G. Loss Function

Since a drug pair may be associated with one or more side effects, the prediction task is formulated as a multi-label classification problem. Binary Cross-Entropy loss is used to measure the discrepancy between predicted probabilities and true labels. This loss function is well-suited for handling mul- tiple independent labels and ensures stable gradient updates during training.

H. Model Evaluation Metrics

To assess the performance of the proposed model, multiple evaluation metrics are employed. These include accuracy, precision, recall, and F1-score. Accuracy measures the overall correctness of predictions, while precision and recall evaluate the model's ability to correctly identify relevant side effects. The F1-score provides a balanced measure by considering both precision and recall. These metrics collectively offer a comprehensive evaluation of the model's predictive capability.

I. Prediction Phase

After training, the best-performing model is saved and used for inference. During prediction, new drug pairs are provided in the form of SMILES strings. These inputs are converted into molecular graphs and passed through the trained model to predict possible side effects. The output includes both side- effect codes and their corresponding names, enabling clear interpretation and clinical relevance.

VI. CONCLUSION AND FUTURE WORK

This paper presented an intelligent healthcare framework for predicting adverse drug reactions caused by drug–drug interactions using Graph Neural Networks. By representing drug molecules as molecular graphs and learning interaction patterns through graph-based deep learning, the proposed system effectively captures complex chemical and structural relationships that traditional feature-based approaches fail to model.

Unlike conventional drug interaction detection systems, the proposed platform integrates prediction capabilities within a hospital management environment, enabling real-time clini- cal decision support. The inclusion of patient-specific vital parameters further enhances the contextual understanding of drug safety, allowing more personalized and proactive risk assessment. Experimental evaluation on publicly available drug–drug interaction datasets demonstrates that the proposed approach achieves reliable predictive performance and offers significant potential for improving medication safety.

Overall, this work bridges the gap between theoretical drug interaction research and practical healthcare deployment, con- tributing toward safer, data-driven,

and patient-centric medical treatment.

Future work will focus on integrating electronic health records to enable automated prescription analysis, incorporat- ing explainable AI techniques to improve model interpretabil- ity and clinical trust, and conducting large-scale clinical vali- dation to assess real-world effectiveness and scalability.

REFERENCES

[1] A. P. Rodrigues and M. Rocha, "Drug–drug interaction prediction using deep learning," Briefings in Bioinformatics, vol. 21, no. 2, pp. 454–466, 2020.

[2] O. ChandraUmakantham, S. Srinivasan, and V. Pathak, "Detecting side effects of adverse drug reactions through drug–drug interactions using graph neural networks," IEEE Access, vol. 12, pp. 93823–93840, 2024.

[3] Y. Deng, X. Xu, Y. Qiu, J. Xia, W. Zhang, and S. Liu, "A multimodal deep learning framework for predicting drug–drug interaction events," Bioinformatics, vol. 36, no. 15, pp. 4316–4322, 2020.

[4] K. K. C. Lee, R. Li, and A. R. Haake, "Predicting biomedical inter- actions with higher-order graph convolutional networks," IEEE/ACM Transactions on Computational Biology and Bioinformatics, vol. 19, no. 2, pp. 987–998, 2022.

[5] M. Zitnik, M. Agrawal, and J. Leskovec, "Modeling polypharmacy side effects with graph convolutional networks," Bioinformatics, vol. 34, no. 13, pp. i457–i466, 2018.

[6] T. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," in Proc. ICLR, 2017.

[7] P. Velic˘kovic´ et al., "Graph attention networks," in Proc. ICLR, 2018.

[8] W. Hamilton, Z. Ying, and J. Leskovec, "Inductive representation learning on large graphs," in Proc. NeurIPS, pp. 1024–1034, 2017.

[9] D. K. Duvenaud et al., "Convolutional networks on graphs for learning molecular fingerprints," in Proc. NeurIPS, pp. 2224–2232, 2015.

[10] S. Kearnes et al., "Molecular graph convolutions: moving beyond fingerprints," Journal of Computer-Aided Molecular Design, vol. 30, pp. 595–608, 2016.

[11] J. Gilmer et al., "Neural message passing for quantum chemistry," in Proc. ICML, pp. 1263–1272, 2017.

[12] M. Wishart et al., "DrugBank 5.0: a major update to the DrugBank database," Nucleic Acids Research, vol. 46, no. D1, pp. D1074–D1082, 2018.

[13] A. Tatonetti, P. Ye, R. Daneshjou, and R. Altman, "Data-driven predic- tion of drug effects and interactions," Science Translational Medicine, vol. 4, no. 125, 2012.

[14] F. Wu et al., "MoleculeNet: a benchmark for molecular machine learning," Chemical Science, vol. 9, pp. 513–530, 2018.

[15] S. Zhang et al., "Predicting drug–drug interactions based on integrated similarity measures," PLoS One, vol. 12, no. 2, 2017.

[16] Y. Liu et al., "Deep learning for predicting adverse drug reactions," IEEE Journal of Biomedical and Health Informatics, vol. 23, no. 6, pp. 2336–2345, 2019.

[17] Z. Wang et al., "Explainable artificial intelligence for healthcare: a review," IEEE Access, vol. 9, pp. 109998–110020, 2021.

[18] J. Esteva et al., "A guide to deep learning in healthcare," Nature Medicine, vol. 25, pp. 24–29, 2019.

[19] K. He et al., "Deep residual learning for image recognition," in Proc. CVPR, pp. 770–778, 2016.

[20] A. Rajkomar, J. Dean, and I. Kohane, "Machine learning in medicine,"New England Journal of Medicine, vol. 380, pp. 1347–1358, 2019.

[21] M. Zitnik and J. Leskovec, "Predicting multicellular function through multi-layer tissue networks," Bioinformatics, vol. 33, no. 14, pp. i190– i198, 2017.

[22] R. Xu et al., "Drug–drug interaction prediction based on graph convolu- tional networks," IEEE Journal of Biomedical and Health Informatics, vol. 25, no. 6, pp. 2113–2124, 2021.

[23] Y. Luo et al., "A deep learning approach for predicting adverse drug reactions," BMC Bioinformatics, vol. 19, no. 1, pp. 1–12, 2018.

[24] J. Shang et al., "End-to-end structure-aware convolutional networks for drug–drug interaction prediction," Bioinformatics, vol. 36, no. 2, pp. 364–372, 2020.
[25] C. Yang et al., "Molecular graph convolutional networks for drug discovery," Journal of Chemical Information and Modeling, vol. 59, no. 9, pp. 3817–3828, 2019.
[26] A. Holzinger et al., "What do we need to build explainable AI systems for the medical domain?" arXiv preprint arXiv:1712.09923, 2017.
[27] Y. Wang et al., "SMILES-based deep generative model for drug discov- ery," Journal of Cheminformatics, vol. 11, no. 1, pp. 1–18, 2019.
[28] J. Ramsundar et al., "Massively multitask networks for drug discovery,"arXiv preprint arXiv:1502.02072, 2015.
[29] S. Ji et al., "A survey on knowledge graphs: representation, acquisition, and applications," IEEE Transactions on Neural Networks and Learning Systems, vol. 33, no. 2, pp. 494–514, 2022.
[30] A. E. Hoerl and R. W. Kennard, "Ridge regression: biased estimation for nonorthogonal problems," Technometrics, vol. 12, no. 1, pp. 55–67, 1970.
[31] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," Nature, vol. 521, pp. 436–444, 2015.
[32] S. Min et al., "Deep learning in bioinformatics," Briefings in Bioinfor- matics, vol. 18, no. 5, pp. 851–869, 2017.
[33] M. Miotto et al., "Deep learning for healthcare: review, opportunities and challenges," Briefings in Bioinformatics, vol. 19, no. 6, pp. 1236–1246, 2018.
[34] T. Chen et al., "A simple framework for contrastive learning of visual representations," in Proc. ICML, pp. 1597–1607, 2020.