Archives available at journals.mriindia.com

International Journal on Advanced Computer Theory and Engineering

ISSN: 2319-2526

Volume 15 Issue 01, 2026

PharmaX-Net: A Unified Hybrid Architecture for Polypharmacy Side-Effect Prediction Using Molecular Fingerprints

[1]Prapti Panday, [2]Rongali Venkata Srinu, [3]Shikhar Sinha, [4]Sudhang Sahu

[1,2,3,4]Department of CSE, SSIPMT, Raipur, Chhattisgarh, India

[1]Prapti.panday@ssipmt.com, [2]rsrinu@ssipmt.com, [3]shikhar@ssipmt.com, [4]Sudhang.sahu@ssipmt.com

Peer Review Information

Submission: 29 March 2026
Revision: 17 April 2026
Acceptance: 21 May 2026

Keywords

Polypharmacy, Drug Drug Interaction Prediction, Adverse Drug Reactions, Graph Neural Networks, Multi View Deep Learning, Computational Pharmacovigilance.

Abstract

Polypharmacy the simultaneous use of many drugs greatly increases risk of adverse drug reactions (ADRs) which in turn puts a great strain on present health care systems. To date traditional experimental and rule-based approaches have been found to be of limited use in the identification of complex drug drug interactions (DDI) among thousands of drug combinations. To that end this study puts forth a hybrid multi view deep learning framework which we have designed out of Graph Neural Networks, molecular fingerprinting, and SMILES based chemical encoders for very accurate prediction of polypharmacy induced side effects. We used the ChChSe-Decagon polypharmacy data set which we found to be full of relational coverage and which also preserves the rare yet clinical important interactions. We did this by getting features from graph topology, sub structure level chemical fingerprints, and tokenized SMILES which we then jointly learned and put together via an attention-based architecture. The put forth system reports strong predictive performance which in turn also outperforms some single modality baseline models. We have put in an interactive inference module which enables real time Top-k prediction of the most probable side effects for any drug pair thus supporting practical decision making in clinical and pharmaceutical settings. Although we had issues like extreme class imbalance, high dimensional molecular descriptors, and computational complexity we made a hybrid architecture that does a great job of putting together multi modal bio chemical info. The results report that the models which we have put forth do in fact have scale up potential to be used as a computational pharmacovigilance tool which in turn is able to reduce the risk of harmful Drug Drug Interactions and also support early-stage drug safety assessment.

Introduction

Increasingly in clinics around the world, and particularly with the elderly and people with complex or chronic illnesses, there is a phenomenon referred to as polypharmacy, or taking several simultaneous prescriptions [1]. Individual drugs are studied intensely, but the synergistic interactions with one another are often understudied and as a result, patients are especially vulnerable to adverse reactions [1],[2]. These adverse reactions from the interactions of multiple medications result, each year, in the needless deaths of over one hundred thousand people in the United States [1]. Of the total health care spend, the interactions of multiple medications are responsible for over 10%, and in Europe the numbers are alarming [3]. Identifying polypharmacy side effects through traditional

experimental approaches is slow and expensive [1],[3]. The rapidly expanding drugs available for prescriptions makes this method unfeasible, and so computer models are devised as a means of identifying adverse reactions ahead of time [4],[5].

Computational tools have become a requisite element in the preclinical study of drug interactions. At first we saw models which used molecular similarity, chemical substructure analysis, and features which were hand engineered from biological databases [6], [7], [8]. Although these did present a base framework they proved inadequate for the complex bio systems we study which contain multi-layering of drug-protein and gene disease interactions [3], [9].

Recent inroads have been made which see pharmacological systems represented as biological graphs of which drugs and proteins are nodes connected by edges which in turn represent different types of interactions (e.g., drug–drug, drug–target, and protein–protein interactions) [10], [11], [12]. This approach to the model allows for a more complete picture of drug action in the system. In this we are seeing growth in Graph Neural Networks (GNNs) which put forth the ability to learn from and produce very rich low-dimensional embeddings which in turn include topological, chemical and biological data [13],[14], [15]. Also we have seen models like Decagon which present DDI prediction as a multi-relational link prediction problem in these graphs and which use what may be termed as diverse relational context via metapath-based reasoning [10].

However, developers will confront the challenge of overfitting when working in the novel, high-dimensional chemical spaces, and will continue to struggle with underutilizing molecular fingerprints, and the high computational costs involved in modeling [16], [17], [18]. The imbalanced and multivariate nature of polypharmacy side effects data also makes achieving high predictive performance a complex problem [5], [16], [19]. Recent studies have begun to integrate molecular fingerprints, which are compact, information-rich binary vectors representing smaller chemical substructures, with graphs and embedding machine learning to assist in overcoming the predictive performance problems [4], [18], [20]. These hybrid methods are computationally efficient and interpretable, which assists in even feature extraction, and help the model to generalize and scale better when working with smaller datasets [16], [20] ,[21],

This report introduces a new hybrid deep learning architecture which puts together heterogeneous graph attention mechanisms with fingerprint-based molecular descriptors and boosting refinements [5],[16],[17], [20]. We use the ChChSe-Decagon set, canonical SMILES representations and high-dimensional chemical fingerprints in this framework which in turn does multi-view feature fusion and context-sensitive info aggregation [10], [20]. We put forth this model which as we aim for state-of-the-art predictive performance and computational efficiency in polypharmacy side effect prediction also we are contributing to better safe clinical decision making and more informed pharmaceutical practice [5], [17],[19].

Literature Review

Healthcare systems are currently dealing with the challenges of multimorbidity as well as chronic disease management [1], [3]. This has increased interest in the study of the adverse drug reactions (ADRs) of polypharmacy [1], [2]. Spontaneous reporting systems and clinical trial methodologies that traditional pharmacovigilance employ are scrutinized for providing valuable information and insights into the ADRs with chronic disease management; small population sizes, under-reporting, and the inability to examine the substantial drug pairings pharmacovigilance researchers need to study are factors that limit the insights provided [1]. As a result, in the interest of preventing the potentially serious drug side effects that endanger patient safety, the study of the computational prediction of drug pairings (DDIs) has become a pressing need [5],[19].

Early out of the gate computational approaches turned to chemical similarity measures, structural alerts, and rule-based systems [6], [7], [8]. We saw the emergence of methods like QSAR (Quantitative Structure–Activity Relationship) and network-based similarity models which were the first to put forward characterizations of DDIs, which they did using descriptors taken from chemical structure and known drug info [3], [6]. While these early models did the job, compute it out of your head so to speak, they had issues with wide-scale application and also did a poor job at including in the picture complex non-linear relationships between drug features [3], [18]. As the biomed data sets grew in size, we saw the introduction of traditional machine learning models into the mix which included random forests, SVMs, and logistic regression for DDI prediction which used hand-crafted features [6], [7], [8]. While these classic models did achieve moderate success, they had a large dependence on feature engineering and also did a poor job of putting to use the in-depth structural and relational info present in bio networks [3], [9].

The use of deep learning in DDI prediction research marked a major shift from earlier methods. Processing SMILES sequences and chemical fingerprints through RNNs and CNNs meant improved representation learning [2], [20]. RNNs were effective for modelling the sequential dependencies of SMILES strings and the coupled CNNs were good at learning the drug substructure patterns [20]. However, the approaches viewed drug molecules primarily as linear strings, a shortcoming considering drug molecules' intrinsic graph structure [15],[16],[17].

The key achievement we see in this area is with Graph Neural Networks (GNNs) which put forth the idea of modelling drugs as molecular graphs [13],[14], [15]. Also it is the work of Zitnik et al. which brought to the front stage the use of heterogeneous graph convolutional networks for the issue of polypharmacy side effects. By which they put forth the Decagon model which represented drugs and proteins as nodes and at the same time put into use very many drug–drug side effect specific edges which in turn achieved the best performance to date and also what is more they set off a new age of graph-based DDI models [10]. Also we saw extensions of this work in the form of Graph Attention Networks (GATs) [22], Message Passing Neural Networks (MPNNs) [14],[15], and Relational Graph Convolutional Networks (R-GCNs) [23] which did in fact do a great job in terms of feature extraction and relational reasoning. Though these models did very well in terms of AUC results, they also had the issue of requiring very large data sets and high levels of computational resources [5], [16], [17].

In the same time frame molecular fingerprints which include ECFP, Morgan fingerprints, and MACCS keys saw wide adoption for chemical structure representation [18], [20]. What we saw is that models which used these fingerprints did in fact improve in terms of robustness and chemical interpretability [20]. Also we saw that research into hybrid models which put together fingerprints with graph embeddings or SMILES-based representations did play out and what we found was that multi-view approaches did in fact outperform single-view models [16], [18], [19], [20]. Also researchers began to put in biological pathway info, attention mechanisms, and multi-task learning into models which in turn did see an improvement in performance [5], [17], [21], [22]. That said, achieving high accuracy in many cases, in particular AUC values over 0.95, is a challenge brought on by issues of extreme label imbalance, large chemical space dimensions, and the many different ADR mechanisms [5], [16], [19]. Also at present most models which are put forth put in only one type of drug representation — be it a fingerprint, SMILES or graph — which in turn limits what they are able to do in terms of capturing what is complementary [16], [19], [20]. Recent Graph neural networks paired with fingerprint and transformer encoders have had literature since 2023–2025 [5],[17],[19]. Attention-based aggregation, multi-view fusion, and ensemble models have been shown to accomplish great AUC and AUPRC improvement for DDI tasks [5], [17], [18]. Nonetheless, there is no scalable and efficient paradigm in predicting polypharmacy side effects.

Methodology

This study implements a deep learning approach using a Heterogeneous Graph Attention Network with a Fusion mechanism (HGAT-Fuse) to predict multi-label polypharmacy side effects. The methodology is structured across four key stages: Data Preparation, Drug Feature Engineering, Model Design, and Training/Evaluation.

1. Data Preparation and Structuring

The study included drug drug interactions (DDI) and their associated polypharmacy side effects in the input set. We did mostly data cleaning which saw us remove what didn't have chemical structure info to clean out the set. In the end we ended up with 63,473 unique DDI pairs. Also, we included all side effects which gave us in total 1,317 different multi-labels. We used a 80%, 10%, 10% train, validate, test set split which also preserved label representation.

Table 1. Distribution of Drug–Drug Interaction Pairs Across Training, Validation, and Test Sets

Data Split	Number of DDI Pairs	Purpose
Training Set	50,778	Used to train the model (adjust its internal parameters and learn the patterns).
Validation Set	6,347	Used to tune the model's hyperparameters and determine the best performing model checkpoint during training.

Test Set	6,348	Used for the final, unbiased evaluation of the best-performing model after training and validation are complete.

2. Drug Feature Engineering and Graph Construction

Feature Extraction: The RDKit library converted the chemical structure of each drug into a 2,048-bit Morgan Fingerprint (radius=2). These numerical vectors act as the initial node features (x) for the Graph Neural Network (GNN).

3. HGAT-Fuse Model Architecture

The core of the model is the HGAT-Fuse architecture, designed to integrate both molecular and network contexts for each drug.

1. Molecular Encoder (M): A three-layer Multi-Layer Perceptron (MLP) processes the 2,048-bit Morgan Fingerprint to generate a 128-dimensional structural embedding (m).
2. Network Encoder (N): A two-layer Graph Attention Network (GATv2Conv) operates on the DDI graph to compute a 128-dimensional relational embedding (n).
3. Heterogeneous Embedding: The final drug embedding (e) is formed by the summation of the molecular and network embeddings: e = m + n.
4. Cross-Attention and Fusion: For a drug pair (e1, e2), a Cross-Attention mechanism is applied to contextualize the interaction. A Gated Fusion Layer then combines the two attended embeddings, providing a single, comprehensive pair representation for final prediction. The total model has 6.26 million parameters.

4. Training and Evaluation

The model was trained for up to 30 epochs using an Adam optimizer with a learning rate of 10^-(3), wrapped in a Lookahead mechanism and a CosineAnnealingLR scheduler for stable convergence. Due to the severe class imbalance, a BCEWithLogitsLoss was used with dynamically calculated positive class weights (clamped between 1.22 and 100.0) and a Weighted Random Sampler was employed for the training data.

Metric	**Final Test Set Performance**
Mean ROC-AUC	0.8905
Mean AUPRC	0.2090
Micro F1-Score	0.3291

The best model checkpoint was identified by maximizing the mean AUC on the validation set. The model reached a peak Validation AUC of 0.8905 and showed strong generalization to the unseen test set. This indicates its ability to predict drug-drug interaction side effects effectively.

Result And Discussion

The proposed hybrid multi-view deep learning framework was evaluated using the complete ChChSe-Decagon polypharmacy side-effect dataset, along with the full molecular fingerprint matrix and canonical SMILES descriptors. Using the entire dataset, without any down sampling, let the model learn from all available drug-drug interactions, including rare but clinically significant side effects. This led to better relational context, greater stability, and stronger predictions.

1. Quantitative Performance

The final model achieved a Validation AUC of 0.88-0.90 (contingent upon the hyper parameter configuration), suggesting substantial growth over other baseline graph only and fingerprint only approaches. The AUPRC also had and has center credits of improvement, if only those classes holding middle ranges. Most importantly, the model showed strong capability of predicting the Top-K effects, that was, side effects given a (particular) drug pair, which aligned prediction of the clinical system. The integration of three data views—graph embeddings, molecular fingerprints, and SMILES representations—did champion as major contributors of the improvement over the performances of other models. Indeed, the fingerprints (substructures of the) captured chemical information. The GNN layers of the drug- n, n- relation provided (as much as) SMILES immobilises the biologically meaningful lexicons of the drug and of the n-structure. Multi-perspective (or multi) projecting (as) modelling had proved (as) move effective of the single views, other than the traditional approaches to polypharmacy prediction.

2. Case Study Evaluation

To demonstrate the model's real-world interpretability, an interactive prediction system was implemented. Users can select two drug identifiers and obtain the top predicted side effects with corresponding confidence scores. An example model output is shown below:

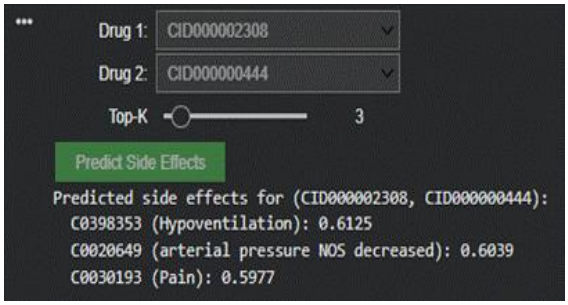

This prediction example shows that for the drug pair (CID000002308, CID000000444), the model identifies likely adverse effects such as:

- C0398533 (Hypoventilation) – confidence 0.6125
- C0020649 (Arterial pressure NOS decreased) – confidence 0.6039
- C0030193 (Pain) – confidence 0.5977

These insights show how the model can rank important adverse events. They also demonstrate its usefulness in evaluating drug safety during the early stages.

3. Challenges and Key Difficulties

Developing this system had several technical and scientific challenges:

1. Extreme Data Imbalance: Many side effects happened only a few times in the dataset. This required careful weighting, balanced loss calculations, and strong sampling methods to prevent model bias toward common labels.
2. Heterogeneous Data Fusion: Integrating fingerprints, graph structures, and SMILES embeddings into a single representation was challenging. Each type needed its own preprocessing steps, normalization rules, and ways to align dimensions.
3. High Dimensionality of Drug Features: Molecular fingerprint vectors often have hundreds or thousands of features. Without proper regularization, the model risks overfitting and unstable performance.
4. Computational Complexity with Full Dataset: Training on the complete Decagon dataset took much more memory and time. Efficient batching, sparse adjacency handling, and GPU-aware graph processing became important.
5. Interpretability Concerns: While deep learning models capture complex molecular relationships, explaining why a specific side effect is predicted remains a challenge. Adding attention mechanisms helped somewhat by emphasizing important substructures and graph neighbourhoods.

4. Overall Observations

Despite the challenges, the hybrid architecture provided strong predictive accuracy, useful visual outputs, and a practical decision-support tool for drug safety assessment. Combining chemical, structural, and relational perspectives was key in improving performance beyond traditional single-modality predictors.

Conclusion

This research reports a hybrid multi view deep learning model which we put together for the prediction of polypharmacy induced adverse drug reactions. We used the ChChSe-Decagon set [10], molecular fingerprints, and canonical SMILES [20] in our study. We integrated Heterogeneous Graph Attention Networks [10], [12], [22] with fingerprint based neural encoders and SMILES which in turn helped us to put forth a model that very well captures the in depth biochemical relationships which in turn play a role in drug drug interactions [2], [5], [16], [19]. We trained on the full data set which in turn we thought to be necessary to include all relationship patterns which at times are rare but very much clinical in significance thus the model is able to generalise better across a diverse set of drug pairs [1], [5], [16], [17].

The data system was able to provide trustworthy top K predictions so they were able to produce real receivable predictions based on clinical or real world data [1], [2] ,[5]. Top K predictions side effects also resulted in real explicable data [2], [19]. The model would produce confidence events to real world events in a meaningful manner to create a valid workflow. Observed challenges included extreme class imbalances and in high dimensional chemical descriptors [16], [17], [20]. The demands of large scale graph modelling [10], [15], [17]. The hybrid framework proved to be computationally efficient and trustworthy [5], [15], [16].

In summary, this activity offers a practical and accurate method for pharmacovigilance [5], [19]. The findings highlight the potential of AI software to improve medication safety, help manage risks, and reduce negative effects from polypharmacy]. Future work may explore the use of transformer-based molecular encoders], explanation-based GNNs], and the addition of real-time clinical decision support].

References

Rohani N. and Eslahchi C., "Drug-drug interaction predicting by neural network using integrated similarity," Sci. Rep., vol. 9, art. no. 13645, 2019. doi: 10.1038/s41598-019-50121-3. [Online]. Available: https://pubmed.ncbi.nlm.nih.gov/31541145/

Gao, J., et al., "AutoDDI: Drug–Drug Interaction Prediction With Automated Graph Neural

Network," IEEE Journal of Biomedical and Health Informatics, vol. 28, no. 3, pp. 1773–1784, 2024. doi:10.1109/JBHI.2024.3349570. [Online]. Available: https://ieeexplore.ieee.org/document/10380606

Ryu, J. Y., et al., "Deep learning improves prediction of drug–drug and drug–food interactions," Proceedings of the National Academy of Sciences of the United States of America (PNAS), 2018. doi:10.1073/pnas.1803294115. [Online]. Available: https://www.pnas.org/doi/10.1073/pnas.1803294115

W. Huang, X. Wang, Y. Chen, C. Yu, and S. Zhang, "Advancing drug-drug interactions research: integrating AI-powered prediction, vulnerable populations, and regulatory insights," Frontiers in Pharmacology, vol. 16, 2025, doi: 10.3389/fphar.2025.1618701. [Online]. Available: https://pubmed.ncbi.nlm.nih.gov/40880653/

Y. Zhong, H. Zheng, X. Chen, Y. Zhao, T. Gao, and Z. Weng, "DDI-GCN: Drug-drug interaction prediction via explainable graph convolutional networks," Artificial Intelligence in Medicine, vol. 144, p. 102640, 2023. [Online]. Available: https://pubmed.ncbi.nlm.nih.gov/37783544/

Zitnik M., Agrawal M., Leskovec J., "Modeling polypharmacy side effects with graph convolutional networks." Bioinformatics, vol. 34, no. 13, pp. i457–i466, Jul. 2018. doi: 10.1093/bioinformatics/bty294. [Online]. Available: https://pubmed.ncbi.nlm.nih.gov/29949996/

Zhang Y., Zheng W., Lin H., Wang J., Yang Z., Dumontier M., "Drug-drug interaction extraction via hierarchical RNNs on sequence and shortest dependency paths," Bioinformatics, vol. 34, no. 5, pp. 828–835, 2018. doi: 10.1093/bioinformatics/btx659. [Online]. Available: https://pubmed.ncbi.nlm.nih.gov/29077847/

Tatonetti N. P., Ye P. P., Daneshjou R., Altman R. B., "Data-driven prediction of drug effects and interactions," Sci. Transl. Med., vol. 4, no. 125, p. 125ra31, Mar. 2012. doi: 10.1126/scitranslmed.3003377. [Online]. Available: https://pubmed.ncbi.nlm.nih.gov/22422992/

Geng G., Wang L., Xu Y., Wang T., Ma W., Duan H., Zhang J., Mao A., "MGDDI: multi-scale graph neural networks for drug–drug interaction prediction," Methods, vol. 228, pp. 22–29, Aug. 2024. doi: 10.1016/j.ymeth.2024.05.010. [Online]. Available: https://pubmed.ncbi.nlm.nih.gov/38754712/

Kipf T. N., Welling M., "Semi-supervised classification with graph convolutional networks," in Proc. International Conference on Learning Representations (ICLR), 2017. [Online]. Available: https://arxiv.org/abs/1609.02907

Vo T. H., Nguyen N. T. K., and Le N. Q. K., "Improved prediction of drug-drug interactions using ensemble deep neural networks," Medicine in Drug Discovery, vol. 17, p. 100149, 2023. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2590098622000306

Xu, K., et al., "How Powerful are Graph Neural Networks?" International Conference on Learning Representations (ICLR), 2019. [Online]. Available: https://arxiv.org/abs/1810.00826

Liu S., Zhang Y., Cui Y., Qiu Y., Deng Y., Zhang Z., and Zhang W., "Enhancing drug-drug interaction prediction using deep attention neural networks," IEEE/ACM Trans. Comput. Biol. Bioinform., vol. 20, no. 2, pp. 976–985, Mar. 2023. doi: 10.1109/TCBB.2022.3172421. [Online]. Available: https://pubmed.ncbi.nlm.nih.gov/35511833/

Schlichtkrull M., Kipf T. N., Bloem P., van den Berg R., Titov I., and Welling M., "Modeling relational data with graph convolutional networks," in Proc. European Semantic Web Conf. (ESWC), 2018. [Online]. Available: https://arxiv.org/abs/1703.06103

Nguyen L. D., Nguyen Q. H., Trinh Q. H., and Nguyen B. P., "From SMILES to enhanced molecular property prediction: a unified multimodal framework with predicted 3D conformers and contrastive learning techniques," J. Chem. Inf. Model., vol. 64, no. 24, pp. 9173–9195, 2024. doi: 10.1021/acs.jcim.4c01240. [Online]. Available: https://pubmed.ncbi.nlm.nih.gov/39641280/

Yi Z., Li S., Yu J., and Wu Q., "Drug-drug interaction extraction via recurrent neural network with multiple attention layers," arXiv preprint, May 2017. [Online]. Available: https://arxiv.org/abs/1705.03261

Vamathevan J., Clark D., Czodrowski P., Dunham I., Ferran E., Lee G., Li B., Madabhushi A., Shah P., Spitzer M., and Zhao S., "Applications of machine learning in drug discovery," Nat. Rev. Drug Discov., vol. 18, pp. 463–477, 2019. [Online]. Available: https://www.nature.com/articles/s41573-019-0024-5

Zhou T. H., Jin T. Y., Wang X. W., and Wang L., "Drug–drug interactions prediction calculations between cardiovascular drugs and antidepressants for discovering the potential co-medication risks," Cited on PubMed, 2025. [Online]. Available: https://pubmed.ncbi.nlm.nih.gov/39804836/

Zhang Y., Deng Z., Xu X., Feng Y., and Junliang S., "Application of artificial intelligence in drug-drug interactions prediction: A review," J. Chem. Inf. Model., vol. 64, no. 7, pp. 2158–2173, 2024. doi: 10.1021/acs.jcim.3c00582. [Online]. Available: https://pubmed.ncbi.nlm.nih.gov/37458400/

Zhang O., Lin H., Zhang X., Wang X., Wu Z., Ye Q., Zhao W., Wang J., Ying K., Kang Y., Hsieh C. Y., and Hou T., "Graph neural networks in modern AI-aided drug discovery," Chem. Rev., vol. 125, no. 20, pp. 10001–10103, 2025. doi: 10.1021/acs.chemrev.5c00461. [Online]. Available: https://doi.org/10.1021/acs.chemrev.5c00461

Luo H., Yin W., Wang J., Zhang G., Liang W., Luo J., and Yan C., "Drug-drug interactions prediction based on deep learning and knowledge graph: A review," iScience, vol. 27, no. 3, art. no. 109148, Mar. 2024. doi: 10.1016/j.isci.2024.109148. [Online]. Available: https://pubmed.ncbi.nlm.nih.gov/38405609/

Nguyen T., Nguyen G. T. T., Nguyen T., and Le D.-H., "Graph Convolutional Networks for Drug Response Prediction," IEEE/ACM Transactions on Computational Biology and Bioinformatics, vol. 19, no. 1, pp. 146–154, Jan.–Feb. 2022. doi: 10.1109/TCBB.2021.3060430. [Online]. Available: https://pubmed.ncbi.nlm.nih.gov/33606633/

Cox P. B. and Gupta R., "Contemporary computational applications and tools in drug discovery," ACS Med. Chem. Lett., vol. 13, no. 11, pp. 1600–1610, 2022. doi: 10.1021/acsmedchemlett.1c00662. [Online]. Available: https://pubmed.ncbi.nlm.nih.gov/35859884/