HIGH SCHOOL EDITION
Journal of Student Research
Volume 14 Issue 1 (2025)

High-Throughput Prediction of Drug-Drug Interactions from Molecular Structures Using Convolutional Neural Networks with Multi-Resolution Filters

Aidan Lee[1] and Sherrie Lah[#]

[1]Crean Lutheran High School, USA
[#]Advisor

ABSTRACT

Drug-drug interaction (DDI) occurs when two or more drugs influence each other's performance which potentially alters their safety and efficacy. Predictions of these interactions are important in the medical field because they can lead to unexpected side effects, reduced effectiveness, increased blood levels, and other complications. There are four types of DDIs: mechanical, effect, advice, and no interaction. Identifying all possible DDIs in the current drug pool is challenging during the development of new drugs. Identifying potential interactions during clinical trials is time-consuming and resource-intensive, which contributes to the slowing pace of drug development. To address this problem, I proposed a convolutional neural network-based high-throughput system for predicting drug-drug interactions from molecular structures. The system takes the molecular structures of two drugs and predicts their possible interactions. It achieved an accuracy of 84.89% on a published dataset. Further analysis shows that the filter size of the convolutional neural network significantly impacts the system's accuracy.

Introduction

Drug-drug interaction (DDI) occurs when two or more drugs interact with each other which alters the effect of one or both of the drugs involved. These interactions can lead to changes in how a drug is absorbed, distributed, metabolized, or excreted by the body. The effects of a DDI can vary widely; they might enhance or diminish the intended therapeutic effects, or even cause unexpected and potentially harmful side effects. For example, one drug might increase the toxicity of another, or a drug might reduce the effectiveness of another by speeding up its metabolism.

Predicting DDIs is a critical step in the drug discovery pipeline as these interactions can cause significant issues during the later stages of clinical trials. Drug development is a lengthy process, often taking over 15 years to complete (Gowing Life 2024). For instance, the preclinical safety testing phase alone can take 12 to 15 months, and it may require 9 to 12 years to gather sufficient data to submit a new drug application (Kelder and Parker 2021).

To gain approval from the Food and Drug Administration (FDA), a drug must undergo numerous trials with various other drugs to identify all potential interactions. This process is resource-intensive, time-consuming, and labor-intensive. By identifying potential interactions early in the development process, significant time and resources can be saved. Therefore, the importance of understanding and predicting drug-drug interactions cannot be overstated, as they have critical implications for clinical outcomes, patient safety, and the overall success of pharmaceutical products.

In this research, I proposed a machine learning-based system for predicting drug-drug interactions. The system consists of two main modules: a preprocessing module and a drug-drug interaction prediction module. The preprocessing module converts the molecular structure of a drug into embedding vectors, which mathematically represent the drug compound's structure. The prediction module then takes a pair of these embedding vectors as input and predicts their potential interaction. Additionally, I investigated the relations between the network's filter size and its predictive performance.

The remaining chapters of this paper are structured as follows: Chapter 2 provides an overview of drug-drug interaction and object classification which are essential for understanding the proposed system. Chapter 3 details the proposed machine learning-based approach, while Chapter 4 evaluates its effectiveness through extensive experiments. Finally, Chapter 5 summarizes the findings of the paper.

Background Knowledge

Drug-Drug Interaction

In this research, I aim to classify drug-drug interactions (DDIs) into four categories: Mechanism, Effect, Advice, and No Interaction.

Mechanism

This type of interaction occurs during the metabolism, transport, or receptor binding of a drug. For example, one drug may inhibit an enzyme that metabolizes another drug, leading to reduced metabolism and causing diminished therapeutic effects. These mechanisms are well understood, allowing healthcare providers to manage such interactions by adjusting doses, selecting alternative medications, or closely monitoring the patient

Effect

This interaction directly influences the efficacy of a drug, potentially causing potentiation, attenuation, or modification of its therapeutic effects. These changes can result in either enhanced therapeutic benefits or adverse effects.

Advice

In this category, drugs indirectly affect each other. One drug may alter the pharmacokinetics of another without directly impacting its therapeutic effect, such as increasing the blood concentration of a specific drug. These interactions are managed by adjusting dosages and monitoring drug levels to prevent therapeutic failure or toxicity.

No Interaction

Some drugs show no interaction when used together, meaning they do not influence each other's metabolism, efficacy, or safety.

Object Classification

Object classification is a fundamental task in the field of computer vision and machine learning. The goal of this system is to assign the inputted object data to one of categories or classes. These systems can be developed using various machine learning algorithms, with convolutional neural networks being particularly effective due to their strong performance. During training, the model's parameters are updated to minimize a predefined loss function using gradient descent algorithms. In this research, DDI prediction is approached as an object classification problem. Specifically, the input consists of pairs of molecular structures of two drugs, and the output is one of four interaction types described in Chapter 2.1.

Drug-Drug Interaction Prediction System

The overall process of this experiment is to input two molecular structures, preprocess the molecules into a transcript that a machine learning model can process, then classifying the input into four categories of interactions. This research can save up time and resources as it can prevent future clinical failures. The process is largely divided into 2 parts.

During the preprocessing step, molecular structures are adjusted into a one dimensional line entry. Then the modified string is once again changed to a one-hot vector that can finally be imputed into the Convolutional Neural Network (CNN). In the DDI predicting step, inputs are assigned to predefined groups.

Preprocessing

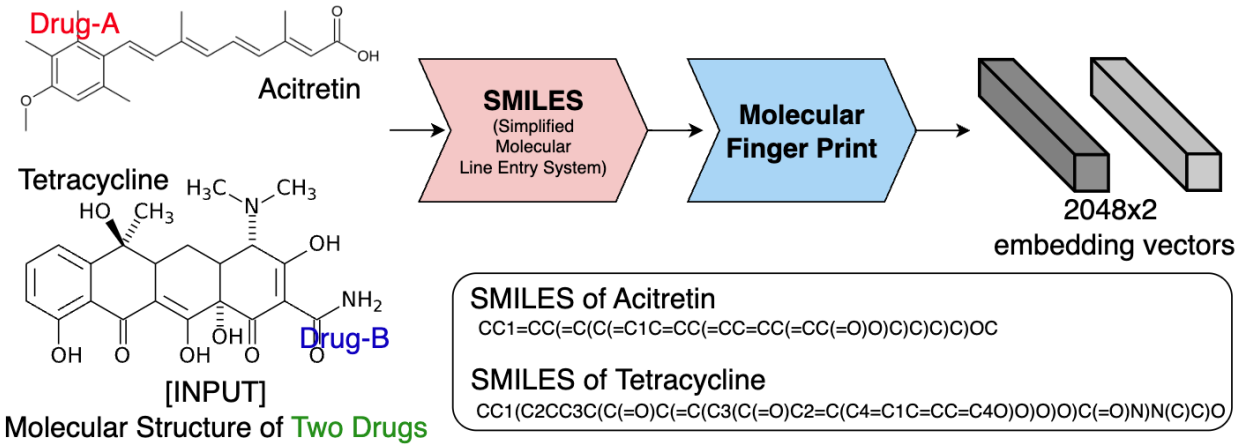

Figure 1. Flow chart of the proposed preprocessing

CNN is not able to understand the molecular structures given as a 2d or 3d model. To prepare or preprocess the molecule for the CNN, the Simplified molecular-input line-entry system (SMILES) (Toropov 2005) is first used. This algorithm converts a molecular structure into a one-dimensional line entry. The string then is represented through a one-hot vector. The string is then modified to a matrix consisting of 0s and 1s, which the AI can interpret. From this process, the molecule is ready to be inputted to the CNN.

Drug-Drug Interaction Prediction

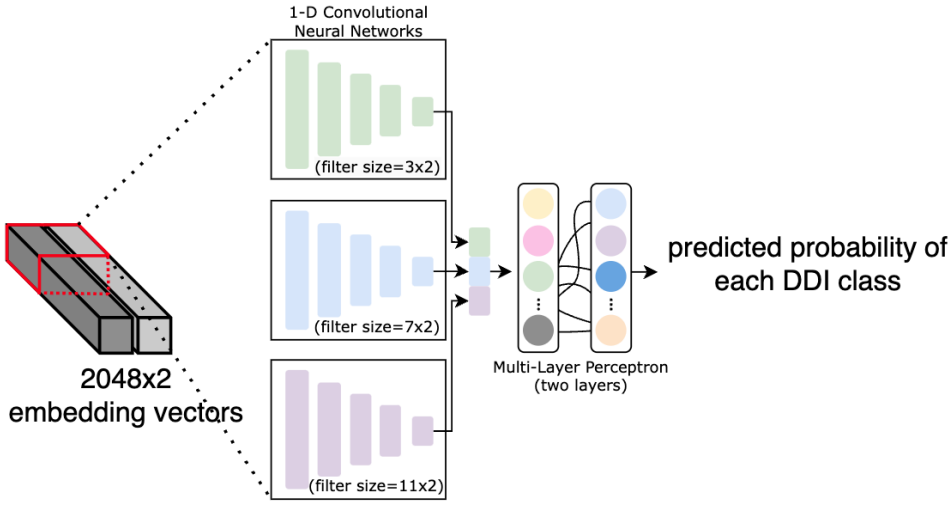

Figure 2. Architecture of the proposed drug-drug interaction prediction network

As the result of preprocessing, the molecular data is represented as two matrices. The two-dimensional matrices are converted to a three dimensional matrix by stacking the two matrices together. The resulting matrix is inserted into the CNN, which creates a feature map. The map is flattened out to a one dimensional matrix then inputted again into a neural network. The neural network produces four possibilities that correspond to the four prefixed categories.

In this context they would be: mechanism DDI, information that we already know of; effect DDI, which affects the drug directly; advice DDI, which indirectly affects the drug; and lastly no interaction indicating that there are no interactions. These outputs are again evaluated through the cross-entropy function to minimize errors.

Equation 1. Cross entropy loss function

$$L = -\sum_{c=1}^{C} GT_c \times log_e(P_c)$$

The function above is the cross entropy loss function used in system training. Here, L represents the error or loss value or that should be minimized during learning. P stands for the probability of an event c and G represents the four values in a vector formed by 0s and 1s. For example, when GT has a value of [1, 0, 0], and the predicted possibility for the first value is 80%, we get L = -(1×ln(0.8)) = 0.2231. In this case, the cross entropy loss is 0.223.

Experimental Results

Dataset

Three types of datasets were used in the evaluation. The DS1 (Zhang et3 al. 2017), containing 548 drugs, DS2 (Wan et al. 2019), containing 707 drugs, and DS3 (Gottlieb et al. 2012) containing 807 drugs. The number of interactions for each dataset were: 97,168, 34,412, 10,078, and the number of non-interactions were: 203,136, 465,437, 641,171. The data was split into 80% training and 20% test.

Dataset Type	Number of Drugs	Number of Interactions	Number of Non-Interactions
DS1 (Zhang et al. 2017)	548	97,168	203,136
DS2 (Wan et al. 2019)	707	34,412	465,437
DS3 (Gottlieb et al. 2012)	807	10,078	641,171

Figure 3. Dataset summary

Evaluation

For the experiment, four kinds of CNN were utilized for evaluation. Starting with VGG-19, MobileNet-V2, Inception-V3, and lastly ResNet-50. These four AIs had various layer depths. The resulting matrix is composed of four kinds of results. Accuracy, Recall, Precision, and F1-Score. The Accuracy can be found in Equation 2-5.

Equation 2. Accuracy

$$Accuracy = \frac{True\ Positive + True\ Negative}{Total}$$

Equation 3. Recall

$$Recall = \frac{True\ Positive}{True\ Positive + False\ Negative}$$

Equation 4. Precision

$$Precision = \frac{True\ Positive}{True\ Positive + False\ Positive}$$

Equation 5. F1-Score

$$F1score = 2 \times \frac{Precision \times Recall}{Precision + Recall}$$

The resulting table consists of the four CNN architectures, and the four methods of evaluation. The evaluation comparison graph shows that ResNet-50 is the most effective of the four CNN architectures. We can infer that the ResNet-50 model was most effective in identifying the four interactions with the accuracy of 84.89% and F1 score of 84.82%. The confusion matrix shows the percentile value on how the trained model classified the drug interactions into four categories. It shows that advice, effect, and mechanism were classified thoroughly, but finding no interaction was not as effective as the classifying other interactions. The following table consists of the accuracy of the four different CNN architectures with various numbers of filters. The accuracy table shows the resulting accuracy from the models with different numbers of filters. Systems with different numbers of filters, from local to global were included to observe the results. The corresponding graph represents each of the network's depths, and their accuracies. The accuracy increases as the number of layers increases in the filter, but reaches saturation in the (11, 15, 19) filter. We can infer that the ResNet-50 system with the filters (7, 11, 15) had the highest accuracy, which was 88.23%.

Table 1. Evaluation result of the proposed drug-drug interaction prediction system

	Accuracy	**Recall**	**Precision**	**F1-Score**
VGG-19 (Simonyan et al. 2014)	0.8191	0.8181	0.8165	0.8173
MobileNet-V2 (Sandler et al. 2018)	0.8252	0.8246	0.8258	0.8252
Inception-V3 (Szegedy et al. 2015)	0.8371	0.8376	0.8347	0.8361
ResNet-50 (He et al. 2016)	0.8489	0.8492	0.8473	**0.8482**

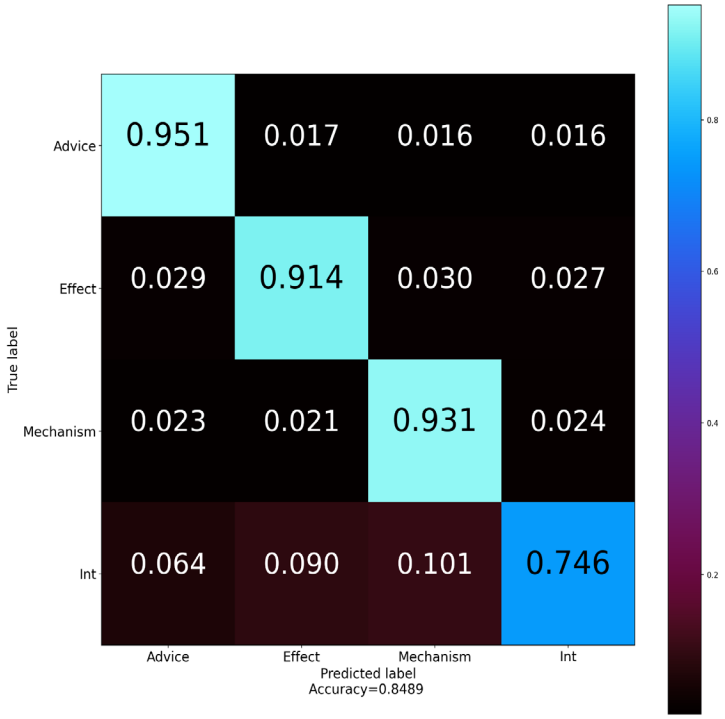

Figure 4. Confusion matrix result

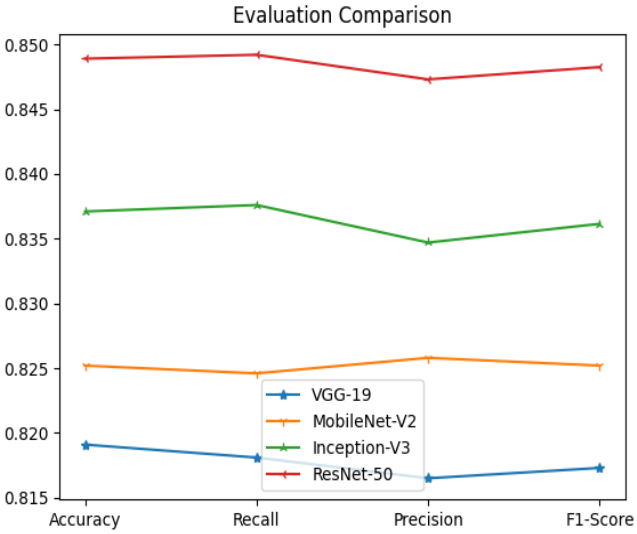

Figure 5. Evaluation result (four CNN architectures)

Table 2. Ablation result (filter size modification experiment)

	Accuracy +(3, 7, 11)	**Accuracy (7, 11, 15)**	**Accuracy (11, 15, 19)**	**Accuracy (baseline)**

VGG-19	0.8191	0.8286	0.8288	0.7699
MobileNet-V2	0.8252	0.8373	0.8382	0.7776
Inception-V3	0.8371	0.8635	0.8648	0.7907
ResNet-50	0.8489	0.8823	0.8822	0.8018

"+(a, b, c) indicates the model trained with filter sizes a, b, and c (see Figure 2 for reference)."

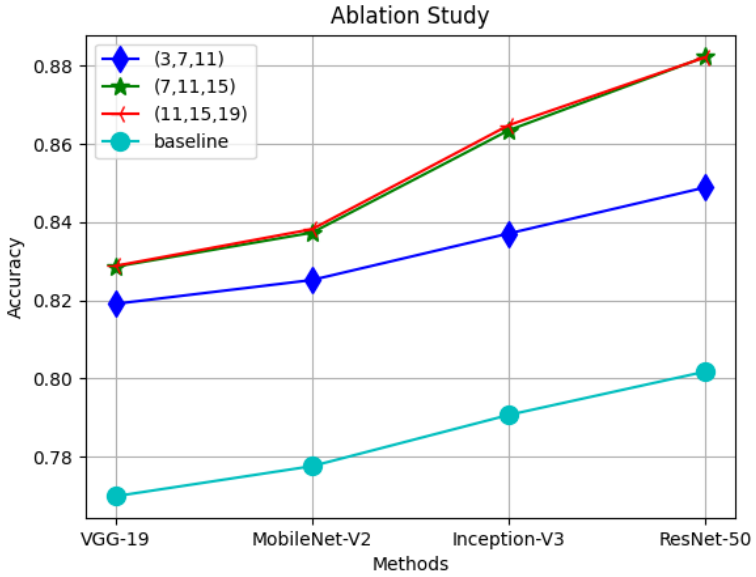

Figure 6. Ablation result (filter size modification experiment)

Conclusion

In this research, I conducted a test on classifying four different kinds of drug drug interactions, advice, effect, mechanism, and no effect, by using a machine learning model. The research was carried out by four kinds of CNN and inputting two kinds of drugs to classify their interactions. The experiment consisted of four different AIs using four types of evaluation methods and four different filter sizes. After conducting some experiments, the model ResNet-50 demonstrated an accuracy of 84.89%. I also found in additional research that the ResNet-50 with filter size (7, 11, 15) showed an accuracy of 88.23. I underline the accuracy of the research as it will take a huge role in reducing time and resources during the arduous task of drug development.

Acknowledgments

I would like to thank my advisor for the valuable insight provided to me on this topic.

References

Gottlieb, A., Stein, G. Y., Oron, Y., Ruppin, E., & Sharan, R. (2012). INDI: a computational framework for inferring drug interactions and their associated recommendations. Molecular systems biology, 8(1), 592.

Gowing Life. (2024, Mar 16). "*How Did Drug Discovery Become So Slow?*": Gowing Life.

https://www.gowinglife.com/how-did-drug-discovery-become-so-slow/

He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In Proceedings of the IEEE conference on computer vision and pattern recognition (pp. 770-778). https://doi.org/10.48550/arXiv.1512.03385

Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., & Chen, L. C. (2018). Mobilenetv2: Inverted residuals and linear bottlenecks. In Proceedings of the IEEE conference on computer vision and pattern recognition (pp. 4510-4520). https://doi.org/10.48550/arXiv.1801.04381

Simonyan, K., & Zisserman, A. (2014). Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556. https://doi.org/10.48550/arXiv.1409.1556

Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., ... & Rabinovich, A. (2015). Going deeper with convolutions. In Proceedings of the IEEE conference on computer vision and pattern recognition (pp. 1-9).

Toropov, A. A., Toropova, A. P., Mukhamedzhanoval, D. V., & Gutman, I. (2005). Simplified molecular input line entry system (SMILES) as an alternative for constructing quantitative structure-property relationships (QSPR).

Wan, F., Hong, L., Xiao, A., Jiang, T., & Zeng, J. (2019). NeoDTI: neural integration of neighbor information from a heterogeneous network for discovering new drug–target interactions. Bioinformatics, 35(1), 104-111.

Zhang, W., Chen, Y., Liu, F., Luo, F., Tian, G., & Li, X. (2017). Predicting potential drug-drug interactions by integrating chemical, biological, phenotypic and network data. BMC bioinformatics, 18, 1-12.