

## PERSPECTIVE

# Machine Learning for Cancer Drug Combination

Ziyan Wang<sup>1</sup>, Hongyang Li<sup>1</sup> and Yuanfang Guan<sup>1,\*</sup>

**When treating multiple complex diseases, such as cancer, polytherapy may demonstrate efficiency than monotherapy. However, due to the multiplicative relationship between the number of drugs and cell lines vs. the number of combinations, it is impractical to test all drug combinations using high-throughput preclinical approaches. An alternative to experimental tests is predicting drug synergy through computational models. Here, we summarize recent computational approaches for predicting drug synergy, discuss current limitations, and propose future directions.**

## BACKGROUND

Polytherapy, the combination of two or more drugs, often achieves better clinical outcomes than monotherapy through reducing drug toxicity and improving therapeutic efficacy. The degree to which multiple drugs synergize or antagonize is quantified by the drug synergy score. Traditionally, the synergy score is measured through *in vitro* or *in vivo* drug screening experiments and calculated by comparing it with the additive effect of individual drug treatments. However, because the number of drug combinations grows exponentially with the number of drugs, it is infeasible to experimentally test the synergistic effects of all possible drug combinations in multiple cell lines under consideration. Recently, there is a growing trend in leveraging computational techniques to extract relevant biological signals from available drug combination

datasets, build predictive machine learning models, and provide guidance for experimental design across diverse diseases and studies.

Artificial intelligence (AI) is a field of computer science that aims to bring human-like intelligence to the world of computing. Machine learning (ML) can be recognized as a type of AI, which aims to let machines learn and extract signals from data. There are two major types of ML algorithms: (i) supervised learning that builds predictive models based on existing training data with desired supervisory labels and predicts labels for new data, and (ii) unsupervised learning that infers characteristic patterns from data without predefined labels. For supervised learning, the predictive performance is evaluated by comparing predicted labels with experimentally measured labels on the held-out testing data, using scoring

metrics, such as Pearson's correlation coefficient. Currently, ML algorithms have been successfully applied to solve various biological and biomedical problems. In this paper, we will first introduce several ML approaches for predicting drug combination synergy.

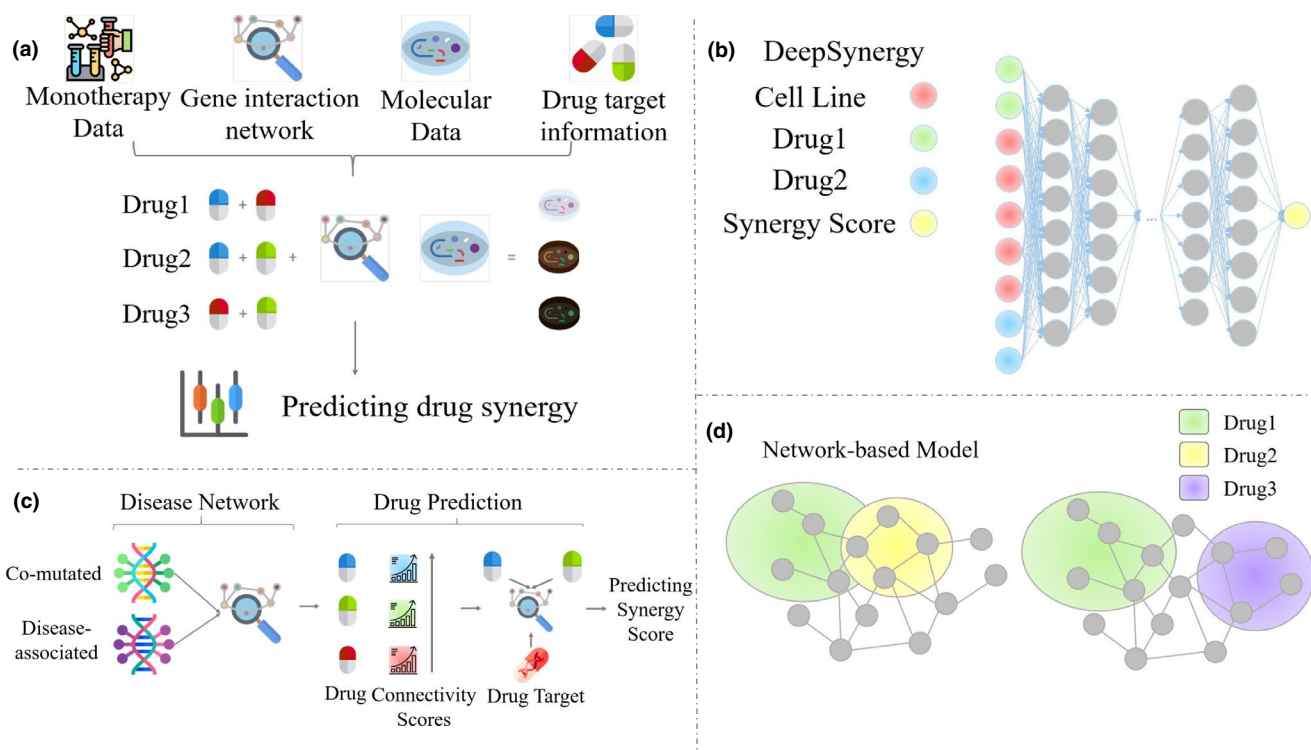
## CURRENT METHODS

In the AstraZeneca-Sanger Drug Combination Prediction DREAM Challenge, a large dataset<sup>1</sup> was used to systematically evaluate the predictive performance of computational models. The winning algorithm introduced a novel network propagation method to simulate the posttreatment genomic profile from the pretreatment profile of a cancer cell line based on drug target information and the gene-gene interaction network.<sup>2</sup> Together with the simulated genomic profiles, the monotherapy data were used to build tree-based conventional ML models for predicting drug synergy (Figure 1a). When tested on a sizeable held-out dataset, this method ranked first among 160 teams in the challenge and established a new state-of-the-field algorithm in the pharmacogenomics research community. Of note, this method is approaching the accuracy of experimental replicates and potentially guides experimental design for future candidate drug pairs.

Recently, a new class of ML algorithm, deep learning (DL), has shown great promise to address many biomedical problems (see **Supplementary Materials**). DL utilizes multiple, or "deep" nonlinear layers to progressively extract high-level features from the input. DeepSynergy is an example of how DL can be used for predicting drug combination synergy<sup>3</sup> (Figure 1b). DeepSynergy accepts both cell line-specific genomic profiles and compound-specific

<sup>1</sup>Department of Computational Medicine and Bioinformatics, University of Michigan, Ann Arbor, Michigan, USA. \*Correspondence: Yuanfang Guan (gyuanfan@umich.edu)

Received July 23, 2019; accepted December 3, 2019. doi:10.1002/cpt.1773



**Figure 1** Schematic illustrations of the machine learning approach for predicting drug synergy. **(a)** The machine learning model for predicting drug synergy by combining the simulated genomic profiles of cancer cell lines and the monotherapy data, which ranked first in the drug combination prediction DREAM challenge. **(b)** The DeepSynergy model using drug and cell line information to predict synergy. **(c)** The SynGeNet model for predicting drug synergy by using the co-mutated genes and disease-associated gene expression signature. **(d)** The network-based model to predict drug synergy.

chemo-informatic features as inputs—the latter input was not provided in the previous DREAM challenge. This model was trained on the O’Neil dataset<sup>4</sup> and surpassed other traditional ML algorithms.

Furthermore, Regan-Fendt *et al.* presented SynGeNet,<sup>5</sup> an innovative network that predicts drug combination effects by integrating transcriptomics-based connectivity mapping and network centrality analysis with clustering algorithms. They demonstrated that SynGeNet outperformed several other tools that used disease-associated and drug-associated gene expression data to predict drug combination synergy (**Figure 1c**).

Besides, Cheng *et al.* proposed a network-based approach to identify drug combinations<sup>6</sup> (**Figure 1d**). They used the bootstrapping algorithm, z-score, and separation measurement to analyze the distance between drug targets and disease proteins. They demonstrated that the network-based approach outperformed traditional approaches, and the network proximity offered an effective strategy to identify new drug combinations.

## LIMITATIONS

Although pioneering efforts have been made to develop ML models for predicting drug combination synergy, as we mentioned above, there are still some limitations.

First, owing to the lack of sufficient data size, current models may not have a comparable prediction when applied to external datasets. This is because most public datasets contain a relatively small portion of all possible drug combinations. When trained on a small dataset, ML models may overfit to the training data, resulting in higher errors during testing. Menden *et al.*<sup>1</sup> also reported that when applying the top-performing models in the DREAM challenge to the O’Neil dataset,<sup>4</sup> the results became worse than those on their original dataset. This is especially true when two datasets only have a few common drugs or cancer cell lines. The cohort and batch effects further introduce noise and differences among datasets. Therefore, a model may perform well on one dataset, but poorly on other datasets.

Second, it is also vital to improve the biological interpretation of ML models.

Hidden prior knowledge is often crucial for developing a powerful prediction model and a better understanding of the mechanism underlying drug synergism. For instance, Li *et al.*<sup>2</sup> leveraged prior information of the gene-gene interaction network and drug target genes to improve prediction accuracy, and, in DeepSynergy,<sup>3</sup> both genomic profiles and chemical compounds were considered. However, it remains unknown to what extent the hidden biological information is needed to perfectly predict drug synergy. Gilvary *et al.*<sup>7</sup> mentioned that AI approaches made significant contributions to several research fields in medicine, but predictive performance can be further improved by developments of biological understanding. In recent years, a surging number of papers were published, which aimed to improve the representation of hidden biological information. For instance, Wang *et al.*<sup>8</sup> proposed a gene set embedding method (Grep) to learn the representation of biological information. They described that the GRep framework could also be applied to other biological

analyses like drug networks. If we can exploit hidden biological information and integrate it into ML models, we will achieve better predictive performance in the future.

Last, different evaluation metrics have been used in different studies. Without a common evaluation metric, it is difficult to compare different models objectively. Data science challenges, such as DREAM, are examples of unbiasedly comparing computational methods using the same scoring metric on the same held-out testing datasets. We recommend that a set of standard and official evaluation metrics should be used for the research community.

### FUTURE DIRECTIONS

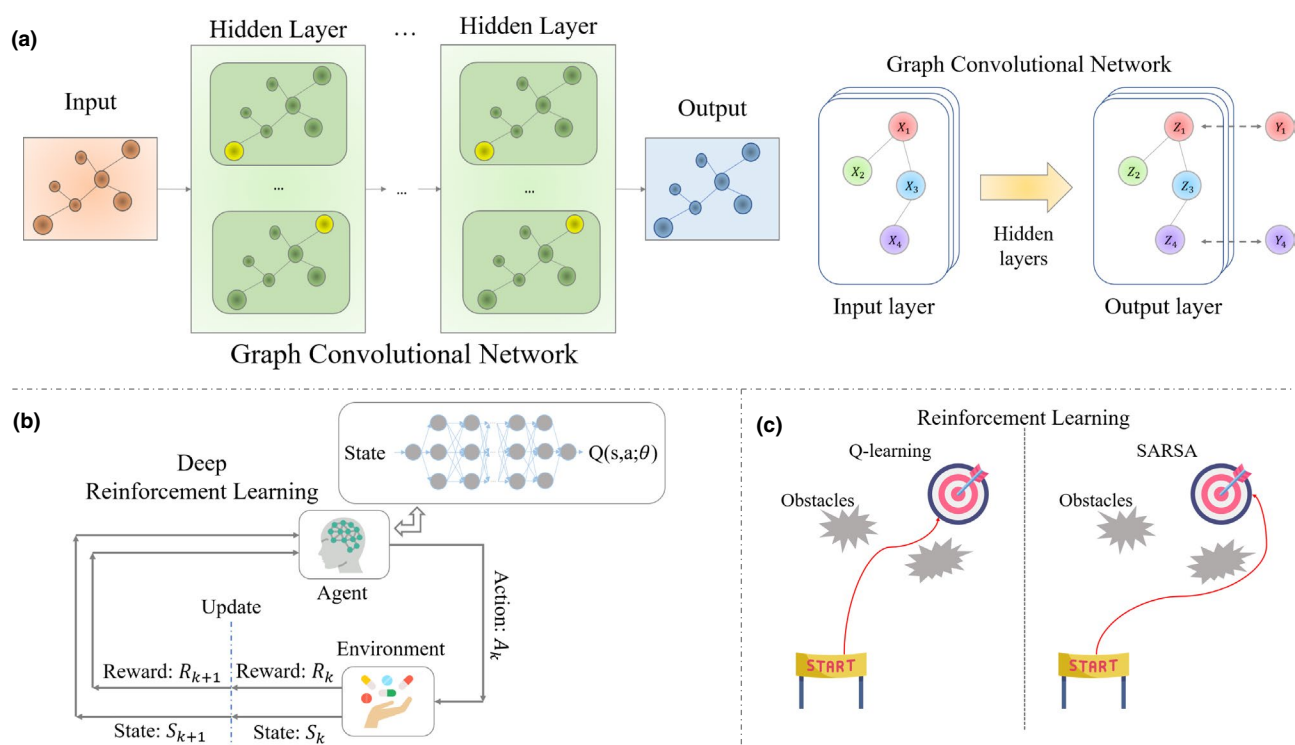
Despite those limitations discussed in the previous section, the development of computational models may benefit from the development of current research. Here, we propose the Graph Convolutional Network (GCN)<sup>9</sup> and Reinforcement Learning (RL),<sup>10</sup> which are rapidly evolving research approaches recently in the AI field, with the overall aim of better predicting drug synergy.

Although DL methods show a great success on Euclidean data, those non-Euclidean data, like chemical data represented as graphs, need to be effectively analyzed. Inspired by the ideas of Convolutional Neural Network and Recurrent Neural Network from DL methods, AI researchers have developed the architecture of GCN to analyze and solve the complexity of graph data (Figure 2a). GCN has already had great applications on medical predicting tasks because it aims to learn a representation of each node with respect to its neighbors. Torng and Altman also use GCN to predict drug-target interactions.<sup>9</sup> We believe that this approach may have a significant effect when predicting synergy scores of drug combinations if we can adequately define the “graph” for drug combinations.

Besides, before introducing the RL approach, we would like to introduce the basic idea of the Markov decision process (MDP). Different from the Hidden Markov Model, in which the next state is only affected by the current state, MDP takes the actions of the agent into

consideration (i.e., the next state is affected by the current state and the current action). An MDP consists of four parts: States (S), Actions (A), Transition Probability ( $P_{sa}$ ), and Reward function (R). The transition probability means the probability distribution from the current state to the next state, and the reward function means the reward that the agent will gain if it takes action from the current state (i.e.,  $S \times A \rightarrow R$ ). The solution of the MDP is a policy ( $\pi$ ) that describes the series of best actions that the agent takes in each state, as well as the real values (V) for it.

Under the basic definitions of the MDP, if the transition probabilities and rewards are unknown, then the MDP problem will become an RL problem. RL is an area of ML and it describes an agent that optimally makes decisions from trials in a fixed environment with reward functions. In Figure 2b, we present the flowchart of the Deep Reinforcement Learning (DRL) algorithm, which is derived from RL and adds the idea of Deep Neural Network when estimating the Q-value. Then, the agent will select the next action by maximizing the



**Figure 2** Schematic illustrations of the Graph Convolutional Network and Reinforcement Learning approach. (a) The flowcharts of the Graph Convolutional Network, where the convolution operation is done on paths that connect the atoms of the chemicals. (b) The overview of the Deep Reinforcement Learning approach. (c) The resulting path generated by two classic Reinforcement Learning algorithms. SARSA, State-action-reward-state-action.

Q-value of each action. In recently published papers, Popova *et al.* have used the DRL algorithms on drug discovery.<sup>10</sup> If we properly define the state space and the actions, DRL may also be applied to predict drug synergy. In addition, other reinforcement learning approaches, like State-action-reward-state-action (SARSA) or Q-learning (Figure 2c), may provide contributions to this field in the future as well.

To sum up, this paper summarizes current ML models for predicting synergy scores of drug combinations. How to improve the predictive performance remains to be explored because of (i) insufficient data, (ii) nondeep biological understanding, and (iii) a missing standard metric or benchmark. Although these limitations may not be solved in a short time, we believe that the performance could be enhanced using current datasets through proper novel AI algorithms, like GCN and DRL.

#### SUPPORTING INFORMATION

Supplementary information accompanies this paper on the *Clinical Pharmacology & Therapeutics* website ([www.cpt-journal.com](http://www.cpt-journal.com)).

#### Supplementary Text.

#### FUNDING

This work is supported by NSF-US14-PAF07599 (#1452656; CAREER: On-line Service for Predicting Protein Phosphorylation Dynamics Under Unseen Perturbations NSF), R35-GM133346-01 (Machine Learning for Drug Response Prediction) to Y.G.

#### CONFLICT OF INTEREST

Y.G. receives personal payment from Eli Lilly and Company, Genentech Inc., F. Hoffmann-La Roche AG, and Cleerly Inc.; holds equity shares at Cleerly Inc., and Ann Arbor Algorithms Inc.; receives research support from Merck KGaA as research contracts and Ryss Tech as unrestricted donations. All other authors declared no competing interests for this work.

© 2020 The Authors *Clinical Pharmacology & Therapeutics* © 2020 American Society for Clinical Pharmacology and Therapeutics

- Menden, M.P. *et al.* Community assessment to advance computational prediction of cancer drug combinations in a pharmacogenomic screen. *Nat. Commun.* <https://doi.org/10.1038/s41467-019-09799-2>. [epub ahead of print].
- Li, H., Li, T., Quang, D. & Guan, Y. Network propagation predicts drug synergy in cancers. *Cancer Res.* **78**, 5446–5457 (2018).
- Preuer, K. *et al.* DeepSynergy: predicting anti-cancer drug synergy with deep learning. *Bioinformatics* **34**, 1538–1546 (2018).
- O’Neil, J. *et al.* An unbiased oncology compound screen to identify novel combination strategies. *Mol. Cancer Ther.* **15**, 1155–1162 (2016).
- Regan-Fendt, K.E. *et al.* Synergy from gene expression and network mining (SynGeNet) method predicts synergistic drug combinations for diverse melanoma genomic subtypes. *NPJ Syst. Biol. Appl.* **5**, 6 (2019).
- Cheng, F., Kovács, I.A. & Barabási, A.-L. Network-based prediction of drug combinations. *Nat. Commun.* **10**, 1197 (2019).
- Gilvary, C., Madhukar, N., Elkhader, J. & Elemento, O. The missing pieces of artificial intelligence in medicine. *Trends Pharmacol. Sci.* **40**, 555564 (2019).
- Wang, S., Flynn, E. & Altman, R. GRep: Gene Set Representation via Gaussian Embedding. <https://doi.org/10.1101/519033>.
- Torng, W. & Altman, R.B. Graph convolutional neural networks for predicting drug-target interactions. *J. Chem. Inf. Model.* <https://doi.org/10.1021/acs.jcim.9b00628>. [e-pub ahead of print].
- Popova, M., Isayev, O. & Tropsha, A. Deep reinforcement learning for de novo drug design. *Sci. Adv.* **4**, eaap7885 (2018).