# Looking at the BiG picture: Incorporating bipartite graphs in drug response prediction

David Earl Hostallero[1], Yihui Li[1], and Amin Emad[1*]

[1] Department of Electrical and Computer Engineering, McGill University, Montreal, QC, Canada

* Corresponding Author:

Amin Emad

755 McConnell Engineering Building

3480 University Street, Montreal, Quebec, Canada, H3A 0E9

Email: amin.emad@mcgill.ca

**ABSTRACT**

Motivation: The increasing number of publicly available databases containing drugs' chemical structures, their response in cell lines, and molecular profiles of the cell lines has garnered attention to the problem of drug response prediction. However, many existing methods do not fully leverage the information that is shared among cell lines and drugs with similar structure. As such, drug similarities in terms of cell line responses and chemical structures could prove to be useful in forming drug representations to improve drug response prediction accuracy.

Results: We present two deep learning approaches, BiG-DRP and BiG-DRP+, for drug response prediction. Our models take advantage of the drugs' chemical structure and the underlying relationships of drugs and cell lines through a bipartite graph and a heterogenous graph convolutional network that incorporate sensitive and resistant cell line information in forming drug representations. Evaluation of our methods and other state-of-the-art models in different scenarios show that incorporating this bipartite graph significantly improve the prediction performance. Additionally, genes that contribute significantly to the performance of our models also point to important biological processes and signaling pathways.

Availability and Implementation: An implementation of the algorithms in Python is provided in github.com/ddhostallero/BiG-DRP.

Contact: amin.emad@mcgill.ca

Supplementary Information: Online-only supplementary data is available at the journal's website.

## INTRODUCTION

Utilization of machine learning and statistical analyses in precision medicine has gained attention in the past decade. Prediction of drug response based on molecular profile of samples based on samples' molecular profiles is a major problem in this domain and various approaches have proposed for this purpose [1-5]. Gene expression profile of samples is widely used for this purpose due to their higher predictive ability compared to other molecular profiles [1]. The curation of large public databases of gene expression profiling of hundreds of cancer cell lines (CCLs) and their response to hundreds of different drugs (e.g., GDSC [6]) has development of novel methodologies in this domain.

Due to the similarity in molecular and chemical structure of different drugs and their mechanisms of action, machine learning (ML) methods that can take advantage of these similarities are of great interest. Instead of training a different ML model for each drug, one can formulate the drug response prediction as a paired prediction problem, such that a model takes in a (drug, CCL) pair as input and trains a single model for all drugs and CCLs [7-9]. This increases the number of samples, and enables information sharing across many drugs and drug families. Chemical structure data (e.g., PubChem [10], ChEMBL [11]) is particularly useful for representing the drugs and models have been developed to take advantage of these [12-14].

Some approaches [15, 16] have formulated this as a matrix factorization problem, forming a matrix of drugs and CCLs. One advantage of this is that these methods directly work with the "entities" (i.e. drugs and CCLs) and responses, and do not need to map feature representations

of the entities to their responses. However, this formulation is inherently transductive, since samples and drugs are expected to be present in the matrix. As a result, these models cannot be directly used to predict the response of a new CCL to a drug unless the CCL has drug response information in the training set for some other drugs.

Taking inspiration from the concept of "entity" from the matrix factorization approaches and to overcome their shortcoming due to their transductive nature, we propose to utilize the underlying matrix by transforming these entities into drug and CCL nodes and form a bipartite graph. We hypothesized that incorporating cell line information that are highly sensitive or resistant to a drug could improve the drug representation for drug response prediction. In our approach called Bipartite Graph-represented Drug Response Predictor (BiG-DRP and BiG-DRP+), we formed this graph by filtering the most sensitive and resistant CCLs for each drug, and linking them through an edge. Although drugs are not directly connected to each other through an edge, 2-hop message passing incorporates information on drug similarities. The model accepts drugs' descriptors and CCLs' gene expression profiles as input, and utilizes them as node attributes for the bipartite graph and as sample features. The output is a continuous drug response value pertaining to the predicted normalized log IC50.

To evaluate the performance of BiG-DRP and BiG-DRP+, we used 5-fold cross validation and compared these results across different baselines and other drug response prediction approaches, namely NRL2DRP [17] and PathDNN [7]. We tested on two data-splitting methods, leave-pairs-out and leave-cell lines-out, which represent two possible scenarios of data

availability. In both scenarios, we have shown significant improvement compared to other approaches. In addition, using a computational pipeline that we developed for identifying the most contributing features, we identified genes that pointed to biological processes and signaling pathways involved in drugs' mechanisms of action.

## METHODS

### Bipartite Graph-based Drug Response Prediction

We developed a novel deep learning-based drug response prediction model that takes advantage of a bipartite graph between drugs and cell lines, which we called Bipartite Graph-represented Drug Response Predictor (BiG-DRP). We also proposed an extension of BiG-DRP, called BiG-DRP+, which accounts for constantly changing drug representations in the former approach. An overview of the architecture of these models are provided in Figure 1.

The BiG-DRP pipeline first obtains latent embeddings CCLs and drugs and uses them in the drug response prediction task. To obtain drug embeddings, first a heterogeneous bipartite graph composed of CCL nodes and drug nodes is formed. The nodes of the bipartite graph are connected via two types of edges: sensitive edges or resistant edges. These edges are based on the log IC50 values of each CCL-drug pair. A sensitive edge implies that the CCL is likely to be sensitive to the drug, while a resistant edge implies that it is likely to be resistant to the drug. In addition, each CCL node is assigned attributes corresponding to its gene expression (GEx) profile and each drug node is assigned attributes corresponding to its drug descriptors. Then, a heterogenous graph convolutional network (H-GCN) generates embeddings of each drug

(denoted as $h_d^{(2)}$ in Figure 1) using this bipartite graph. For each drug of interest, the H-GCN obtains an embedding that not only captures the molecular characteristic of the drug itself, but also captures the characteristics of other drugs that induce a similar sensitive/resistant pattern in CCLs. Inclusion of the GEx profiles of CCLs as node attributes in the bipartite graph allows the model to define the "similar pattern" mentioned above in a broader sense: instead of requiring a similar pattern in the exact same CCLs, the model can identify such patterns in CCLs that have a similar GEx profile.

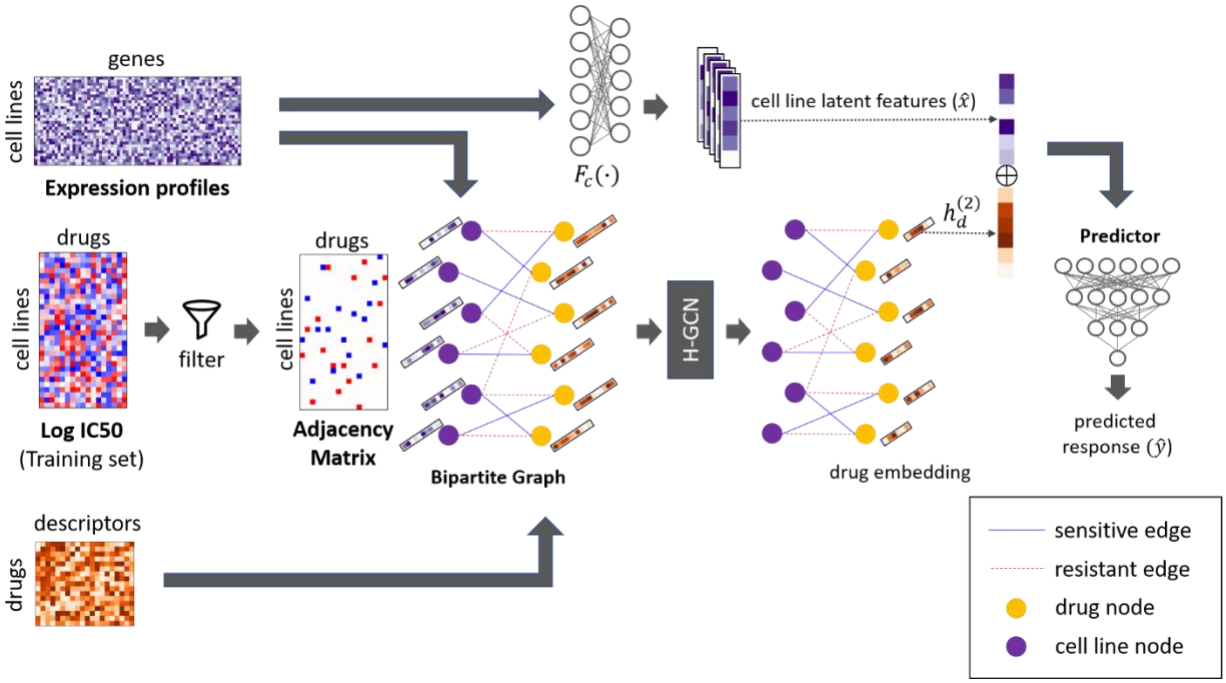


**Figure 1:** The computational pipeline and architecture of BiG-DRP and BiG-DRP+. Latent drug embeddings are generated using a heterogenous graph convolutional network based on a bipartite graph of drug-CCLs and drug descriptors. In parallel, CCL embeddings are generated using an encoder neural network based on their gene expression profile. These embeddings are then used by a predictor neural network to predict the drug response values.

To obtain embeddings of the CCLs based on their GEx profiles (denoted as $\hat{x}$ in Figure 1), the model uses a neural network that is separate from the H-GCN. While it is possible to use the bipartite graph and the H-GCN to obtain CCL embeddings, such a choice would limit the applicability of the pipeline to only CCLs that are already present in the training set. The reason is that a CCL that is not present in the training set will be in the form of a single disconnected node in the bipartite graph and no embedding can be found for it using the H-GCN. However, in many practical applications (e.g., prediction of clinical drug response of patients based on models trained on preclinical CCLs [4, 5]), a model must be able to predict drug response of samples that are not seen by the model during the training for any drug. To avoid this limitation, the CCL embeddings are obtained independent of the H-GCN network and the bipartite graph. The drug and CCL embeddings are then concatenated, representing each (drug, CCL) pair. Then, a series of neural network layers (called *predictor*) are used to predict the drug response of each such pair using the concatenated embeddings.

The BiG-DRP+ is an extension of BiG-DRP, which aims to stabilize the trained model. After the "last" training epoch of BiG-DRP, we train the model for one more epoch but with a smaller learning rate and "frozen" drug embeddings. The lower learning rate prevents the predictor from overfitting while the freezing of the embeddings allows the predictor to learn the finite set of drugs instead of constantly changing representations of the exact same drugs.

**Construction of the Heterogenous Bipartite Graph**

We denote the heterogenous bipartite graph as $G(V_C, V_D, E_r, E_s)$, where $V_C$ is the set of CCL nodes used to build the graph (a subset of all the CCLs in the study) and $V_D$ is the set of drug nodes. $E_r$ is the set of edges that connect drugs to their "most resistant" CCLs, while $E_s$ is the set of edges that connect drugs to their "most sensitive" CCLs. For a fixed value of $k$, a drug is connected via a resistant edge to CCLs whose log IC50 is among the top $k$ percent and is connected via a sensitive edge to CCLs whose log IC50 is among the bottom $k$ percent of the CCLs. The set $V_C$ is then the union of all such CCLs whose drug response are among the top $k$ or bottom $k$ percent of all cell lines for at least one drug. It is worth noting that the edges in this graph are unweighted and the log IC50 values are only used to determine whether a resistant (or sensitive) edge exists or not. We used $k = 1$ in our analysis, but the performance of BiG-DRP and BiG-DRP+ were not sensitive to the choice of $k$, as discussed in Results.

**Drug embedding using heterogenous graph convolutions**

We used a 2-layer heterogenous graph convolutional network (H-GCN) to find a network-based embedding of the drugs. An H-GCN is a variation of graph convolutional network [18], which allows multiple edge types. A forward pass of an H-GCN can be summarized using the following equation:

$$h_v^{(l+1)} = \sigma\left(\sum_{r\in\mathcal{R}}\left(b_r^{(l)} + \frac{1}{\sqrt{|\mathcal{N}(v,r)|}}\sum_{u\in\mathcal{N}(v,r)} h_u^{(l)} W_r^{(l)}\right) + \alpha h_v^{(l)}\right)$$

where $h_v^{(l)}$ is node $v$'s embedding at the $l$th layer, $\sigma$ is a non-linearity function, $\mathcal{N}(v,r)$ is node $v$'s set of neighbours connected using the edge type $r$. $W_r^{(l)}$ and $b_r^{(l)}$ are the weights and biases

at the $l$th H-GCN layer for edge type $r$, respectively. Intuitively, this allows a separation of GCN parameters for each edge type, and thus creates context during the message passing. The normalization factor $\sqrt{|\mathcal{N}(v,r)|}$ prevents the embedding values from exploding due to a large number of neighbours.

Although we constructed a bipartite graph, artificially adding self-loops to the graph is a common practice in GCNs to retain some information from the previous layer, avoiding the complete dependence of the node's embedding to its neighbors. However, in the case of H-GCN, self-loops increase the complexity of the model by adding another set of parameters. To avoid this, we injected a residual term $(\alpha h_v^{(l)})$ to the forward pass to simulate self-loops. Here, $\alpha$ is a hyperparameter (we fixed the value to $\alpha = 0.5$ ) pertaining to the amount of information to be retained for the next layer.

The bipartite graph and the H-GCN allow us to find a drug embedding that captures relevant information about the CCLs that are generally resistant/sensitive to it (its 1-hop neighbours), as well as information on other drugs to which these CCLs have a similar or inverse pattern of response (its 2-hop neighbours). These embeddings enable sharing of information across drugs that are connected to similar set of cell lines via similar edge types.

**Data Acquisition and Preprocessing**

We obtained the drug response data in the form of log IC50 values from the Genomics of Drug Sensitivity in Cancer (GDSC) database [6]. We only selected drugs with known log IC50 values as

well as binarized responses that allow us to calculate the key performance metrics used for evaluation of different methods. We also filtered out duplicate drugs that came from different batches, which are tagged with different drug IDs, named with synonyms, or labeled as "rescreens". In cases of such duplicates, we only kept the one for which the drug response in a larger number of cell lines was measured. We collected the Simplified Molecular Input Line Entry System (SMILES) encoding [19] of these drugs and used the RDKit software [12] to generate drug descriptors (e.g. molecular weight, number of aromatic rings) from these encodings. Descriptors with missing values were excluded from the analysis. At the end of these data cleaning steps, we were left with 237 unique drugs, each with feature vectors of length 198 (representing their drug descriptors).

We performed z-score normalization on drug descriptors, one feature at a time across all drugs. We also z-score normalized the log IC50 values of each drug (one drug at a time) across cell lines. This is necessary since the log IC50 values of different drugs have significantly different means and standard deviations, which renders the calculated metrics incomparable across drugs and inflates the overall correlation coefficient. For example, a relatively small mean squared error for a certain drug, or a high overall spearman correlation do not necessarily indicate good performance without such a normalization. This drug-wise normalization allows us to compare results across different drugs, and prevents overestimation of the models' performance.

For the 237 drugs above, we obtained the RNA-seq GEx profile of 1001 CCLs from the Cell Model passports [20]. We performed $\log_2$(FPKM +1) transformation on the FPKM values. We excluded

genes that showed a small variability across the cell lines (genes with standard deviation <0.1) as well as genes with missing values in some cell lines. After these preprocessing steps, we ended up with 944 unique CCLs and their GEx values of 13,823 genes. This amounted to a total of 181,380 labeled CCL-drug pairs.

**Training Procedure**

As discussed earlier, to enable the model to generalize to completely new CCLs (those that are not seen by the model for any drug during training), we used a separate neural network, parallel to the H-GCN. As input, this network accepts the CCLs' gene expression vector $x$ and produces a latent representation $\hat{x} = F_c(x)$. We then concatenate $\hat{x}$ with the drug $d$'s embedding, $h_d^{(2)}$, and use it as input for our predictor, a 3-layer neural network that outputs the predicted drug response values ($\hat{y}$).

The model was trained end-to-end using the mean squared error $\mathcal{L} = (y - \hat{y})^2$ and Adam as the optimizer [21]. We also fixed the learning rate to 0.0001 and batch size to 128. We used Leaky ReLU for all non-linearity functions (i.e. $\sigma(x) = \max(0, x) + 0.01 \times \min(0, x)$). The number of training steps were decided by randomly selecting samples from the training data and using them as a validation set for early stopping. The model was then re-trained with the entire training set and the previously identified optimal number of training steps. For BiG-DRP+, the extra epoch's learning rate was set to 0.00001.

In our approach, elements of a batch are (drug, CCL) pairs, although all drug embeddings can be generated simultaneously for each forward pass. Embeddings generated using graph convolutional networks rely on the node connectivity. This generally means that a small perturbation of a node's embedding may affect the embeddings of its neighbors in the next GCN (or H-GCN) layer. Unlike regular dense neural networks, it is possible that a dramatic change would occur in the embeddings, even with a relatively small learning rate. In such cases, the predictor may not easily map the "new" embedding to the "known" ones, especially if the drug was not part of the batch during the previous training step. The predictor could see this as having an infinite number of drugs, increasing the level of complexity to the learning process. To address this "moving embedding problem", we developed BiG-DRP+, which slightly modifies the training of BiG-DRP.

The idea of BiG-DRP+ is to stop the training of the H-GCN component after several epochs but continue the training of the predictor using the "frozen" drug embeddings. In our BiG-DRP+ model, we froze the drug embeddings obtained by BiG-DRP (after the number of epochs determined by early stopping), but continued the training of other components of the architecture for one extra epoch (we used a lower learning rate for this epoch). This stabilizes the training of the predictor and enables it to identify CCLs that were treated by the same drug (since the half of the input to the predictor pertaining to the drug features are now fixed). The lower learning rate is a preventative measure to avoid overfitting.

**Evaluation and Cross-validation**

To evaluate the performance of our model we used 5-fold cross validation (CV), in which one fold was kept aside as the test set for evaluation and was not used during training of the model nor for the selection of hyperparameters. This process was repeated five times (each time with a different fold as the test set) to ensure that the specific choice of the test set does not bias the results. We adopted two types of data splitting techniques to form the folds, namely leave-pairs-out (LPO) and leave-cell lines-out (LCO).

In the LPO-CV, the disjoint folds were randomly selected from the set of all (CCL, drug) pairs, while in the LCO-CV the folds contained randomly selected sets of mutually exclusive CCLs. Prior to training, GEx values were z-score normalized per gene. We used only the training folds' unique CCLs to calculate the means and standard deviations to prevent data leakage between training and test sets. Imposing the uniqueness criterion above ensures that the models are not biased towards CCLs that exists in a larger number of (drug, CCL) pairs. To ensure a fair comparison, identical folds were used for all methods. For each drug, predictions of the five folds on their respective test sets were collected and were used to evaluate different methods.

To assess generalizability of our models to independent datasets, we obtained GEx profile (in the form of FPKM) and RECIST clinical drug response of patient tumours from The Cancer Genome Atlas (TCGA) [22] treated with cisplatin or gemcitabine. We selected these drugs since they were present in our training dataset and the clinical drug response was available for more than 150 patients treated with these drugs (n = 309 for cisplatin and n = 158 for gemcitabine) and there

were at least 50 sensitive and 50 resistant patients in each dataset. Similar to previous studies [23], we considered "stable disease" or "progressive disease" as resistant and "complete response" or "partial response" as sensitive. Similar to the preprocessing steps used for GDSC dataset, the expression of each gene in the testing set (in the form of $\log_2$(FPKM +1)) was z-score normalized across the samples, before being given to the models as input.

**Baseline Methods**

We compared our method against several baseline algorithms including both deep learning-based and traditional machine learning methods, detailed below. First, we used a multilayer perceptron (MLP) with a similar architecture and hyperparameters as BiG-DRP. Similar to BiG-DRP, the MLP also utilized the GEx and drug features. However, instead of an H-GCN, we replaced it with a dense neural network, which takes the drug features as input. We also used a support vector regressor (SVR) with a linear kernel as well as a SVR with a radial basis function (RBF) as traditional ML baselines. The concatenation of the GEx and drug features were used as the input to SVR models. Due to the large size of the data, we used Nystroem's transformation [24] to approximate the SVR's kernels to improve its efficiency. Hyperparameters, namely the number of Nystroem components, regularization factor, and gamma for RBF were selected using nested cross validation.

NRL2DRP [17] is a graph representation learning-based method that uses a graph composed of genes, drugs, and CCL nodes, connected by edges according to their sensitivity, mutation, and protein-protein interactions. However, NRL2DRP uses a topology-based graph embedding called

LINE [25], which is typically used for transductive learning. We slightly modified NRL2DRP to predict continuous values instead of binary values (so that it can be applied to our data). PathDNN [7] is another deep learning method that proposes to add some level of explainability to the drug response prediction problem by constraining the neural network connectivity using a pathway mask. This method uses drug targets and gene expressions, both of which should be in any of the Kyoto Encyclopedia of Genes and Genomes (KEGG) pathways [26]. We obtained the drug targets and pathway information from the PathDNN's repository. The drug targets were represented by their normalized STITCH [27] confidence score, which indicates a non-zero value for genes in the drug's targets. However, we removed three compounds because they did not have known targets in the KEGG pathways.

**Identification of genes that are most predictive of drug response**

To identify genes that are predictive of drug response, we used a neural network explainer called CXPlain [28] and a similar approach which we previously developed to aggregate contribution across CCLs and identify top contributors [5]. CXplain uses Granger's causality [29] as the basis of the feature attribution. Intuitively, for each of the features, it tries to predict the increase in the sample's loss if that specific feature is zeroed-out. We trained separate explainers for each of the drugs, since this eliminates the unnecessary complexity of learning attributions for multiple drugs, as well as the additional feature dimensions (i.e., drug features). We pooled the scores by calculating the mean attribution across all the CCLs for each of these drugs. The top genes were identified by a threshold calculated using kneedle [30], with sensitivity S=2.

### Pathway characterization analysis

We used KnowEnG's gene set characterization pipeline [31] to perform pathway enrichment analysis (using Reactome pathways [32]). The p-values of Fisher's exact test were corrected for multiple tests (i.e., multiple pathways) using Benjamini-Hochberg false discovery rate (FDR).

## RESULTS

### Performance of BiG-DRP and BiG-DRP+ based on leave-pair-out cross validation

First, we evaluated BiG-DRP, BiG-DRP+, and other baseline algorithms using a five-fold LPO-CV, in which the folds were randomly selected among the set of all possible (CCL, drug) pairs. Table 1 shows a summary of the performance results using area under the receiver operating characteristic curve (AUROC) and Spearman's correlation coefficient (SCC). To calculate these metrics across all drugs, we first calculated them separately for each drug (Supplementary Table S1) and then obtained mean and standard deviation across the drugs. BiG-DRP+ outperforms all other methods according to both metrics, and BiG-DRP outperforms all baselines but has a slightly worse performance compared to BiG-DRP+. BiG-DRP+ has a ~5% higher AUROC and a ~11% higher SCC compared to that of MLP, which utilizes a similar architecture to BiG-DRP+ (except for the usage of the bipartite graph and the H-GCN). This highlights the importance of this novel aspect of the algorithm.

**Table 1:** The performance of BiG-DRP, BiG-DRP+ and baseline methods using five-fold LPO-CV evaluation. Best performance values are underlined. *Since PathDNN requires availability of at least one drug target in any of the signaling pathways, we could only apply it to 234 drugs.

| | Drug Attributes | Other inputs | Number of Drugs | AUROC mean (± std.) across drugs | SCC mean (± std.) across drugs |
|---|---|---|---|---|---|
| **BiG-DRP+** | Descriptors | GEx | 237 | <u>0.878</u> (±0.068) | <u>0.748</u> (±0.100) |
| **BiG-DRP** | Descriptors | GEx | 237 | 0.875 (±0.068) | 0.742 (±0.099) |
| **MLP** | Descriptors | GEx | 237 | 0.835 (±0.083) | 0.675 (±0.120) |
| **NRL2DRP** | None | Drug-CCL-Gene network | 237 | 0.804 (±0.085) | 0.516 (±0.119) |
| **PathDNN** | Drug Targets | GEx, pathway information | 234* | 0.766 (±0.083) | 0.516 (±0.115) |
| **SVR-RBF** | Descriptors | GEx | 237 | 0.737 (±0.100) | 0.502 (±0.123) |
| **SVR-Linear** | Descriptors | GEx | 237 | 0.736 (±0.101) | 0.494 (±0.129) |

Figure 2 compares the performance of BiG-DRP+ against other methods for individual drugs (measured based on SCC). Each circle in the scatter plots reflects a drug, and the color of the circles reflect the density of other circles in their vicinity. Comparing BiG-DRP+ and BiG-DRP shows that the drug-specific SCC values are generally close to each other (concentrated around the diagonal line); however, the one-sided Wilcoxon signed rank test (p=2.26E-36) suggests that the performance for the majority of the drugs have improved in BiG-DRP+, albeit a small amount. Comparing to other baselines, the figure shows that the majority (and in many cases all) of the circles are above the diagonal line, suggesting a substantial improvement of their response prediction by BiG-DRP+. One-sided Wilcoxon signed rank tests also confirmed this observation, resulting in statistically significant p-values (Figure 2 and Supplementary Table S2).

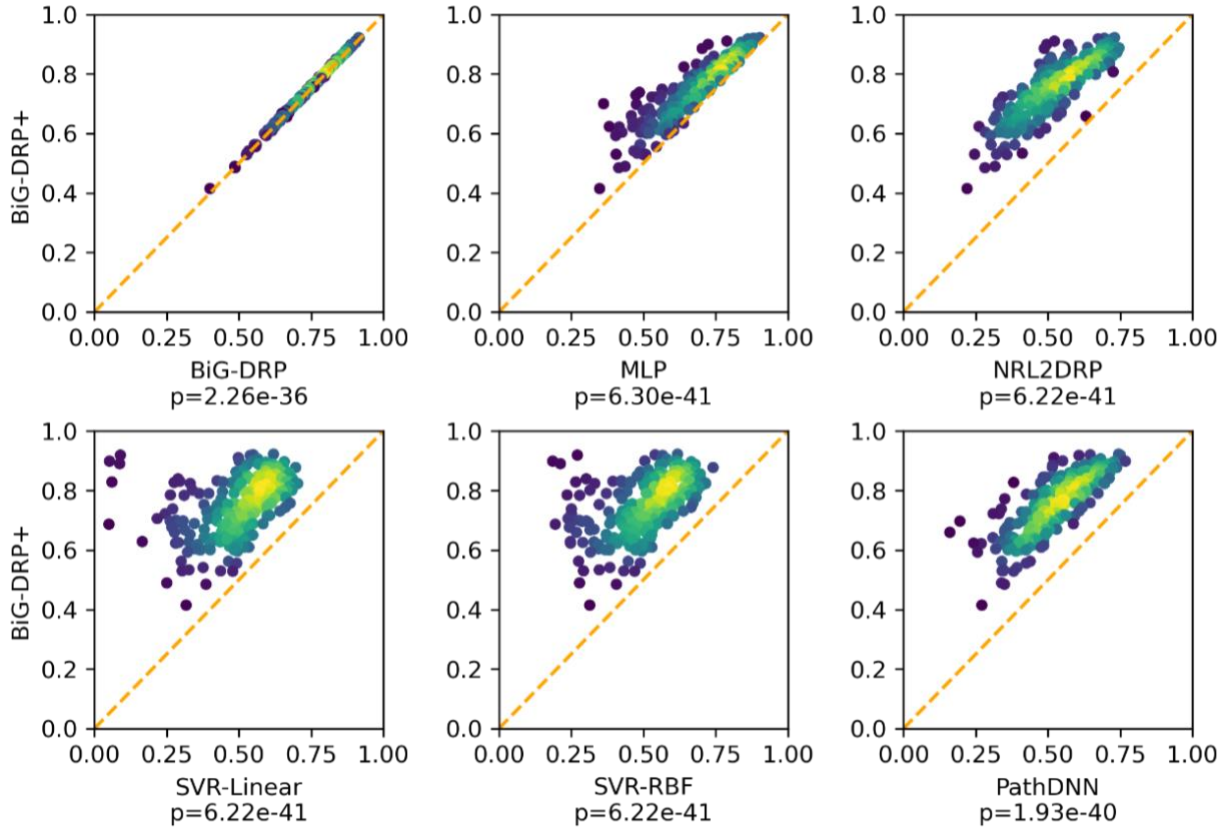


**Figure 2:** The drug-wise performance of BiG-DRP+ compared to baseline methods using five-fold LPO-CV evaluation. Each circle represents a drug, and the color of the circles reflect the density of other circles in their vicinity (yellow shows that there are many circles concentrated in that area). The coordinates reflect SCC for BiG-DRP+ (y-axis) and baseline methods (x-axis). The p-values are obtained using a one-sided Wilcoxon signed rank test, comparing the SCC of BiG-DRP+ and the baselines across drugs.

**Performance of BiG-DRP and BiG-DRP+ based on leave-cell line-out cross validation**

Next, we evaluated the performance of different models using a five-fold LCO-CV. This is a stricter evaluation, since unlike LPO-CV, a CCL in the test set is never seen by the models during training, since folds are randomly selected based on the CCLs and not based on (CCL, drug) pairs. Table 2 shows the summary of the results using our performance metrics. Note that due to the transductive nature of NRL2DRP's embedding method (LINE [25]), this method could not be applied to the LCO-CV evaluation and hence is not included in this table.

Based on these evaluations, BiG-DRP+ has the best performance using AUROC and SCC, while BiG-DRP has the second-best performance. The BiG-DRP+ clearly outperforms MLP using both metrics, further highlighting the importance of the bipartite graph and H-GCN in the drug response prediction task. Similar to LPO-CV evaluation, a drug-wise analysis using SCC for each drug showed a significantly superior performance of BiG-DRP+ compared to all baseline methods (one-sided Wilcoxon Signed-Rank test, Supplementary Table S2). Supplementary Table S1 provides the drug-specific performance metrics for all drugs.

**Table 2:** The performance of BiG-DRP+, BiG-DRP and baseline methods using five-fold LCO-CV evaluation. Best performance values are underlined. *Since PathDNN requires availability of at least one drug target in any of the signaling pathways, we could only apply it to 234 drugs.

| | **Drug Attributes** | **Other input features** | **Number of Drugs** | **AUROC mean (± std.) across drugs** | **SCC mean (± std.) across drugs** |
|---|---|---|---|---|---|
| **BiG-DRP+** | Descriptors | GEx | 237 | <u>0.746</u> (±0.077) | <u>0.431</u> (±0.094) |
| **BiG-DRP** | Descriptors | GEx | 237 | 0.743 (±0.077) | 0.426 (±0.095) |
| **MLP** | Descriptors | GEx | 237 | 0.730 (±0.086) | 0.413 (±0.100) |
| **SVR-RBF** | Descriptors | GEx | 237 | 0.680 (±0.110) | 0.348 (±0.120) |
| **SVR-Linear** | Descriptors | GEx | 237 | 0.666 (±0.102) | 0.324 (±0.119) |
| **PathDNN** | Drug Targets | GEx, pathway information | 234* | 0.612 (±0.074) | 0.193 (±0.061) |

To assess the generalizability of BiG-DRP+ and BiG-DRP to independent datasets, we used them to predict the drug response of patient tumours from the TCGA dataset treated with cisplatin or gemcitabine. Given the predicted log IC50 values, we used a one-sided Mann-Whitney U test to determine if our models can distinguish between the patients that are resistant from those that

377 are sensitive to these two drugs. Our statistical analysis showed significant p-values for both
378 drugs using BiG-DRP+ (p = 2.66E-6 for cisplatin and p = 2.01E-2 for gemcitabine) and using BiG-
379 DRP (p = 2.25E-6 for cisplatin and p = 1.58E-2 for gemcitabine).

380

381 **Detailed Evaluation of BiG-DRP+**

382 Since one major component of the BiG-DRP and BiG-DRP+ pipeline is the bipartite graph of the
383 CCLs and drugs, we sought to evaluate the effect of different thresholds for forming this graph.
384 As explained in Methods, a drug is connected to a CCL with a sensitive (resistant) edge if the log
385 IC50 of the CCL is among the bottom (top) k% of all the CCLs. In our analysis, we fixed this value
386 to be k = 1. To assess the robustness of the results to this parameter, we formed different
387 bipartite graphs with different choices of k = 0.5, 1, 2, 5, 10 and repeated the LPO-CV and LCO-
388 CV. Supplementary Table S3 provides the SCC and AUROC of BiG-DRP and BiG-DRP+ for these
389 evaluations for different values of k. These results suggest that the performance of our proposed
390 methods remain stable for these different choices of k, with a slight deterioration as k increases
391 (less than 1% in all evaluations when comparing k=10 to k=1). This deterioration is expected, since
392 an increase in k increases potentially erroneous edges in the bipartite graph.

393

394 Next, we asked whether the choice of drug features as attributes in the bipartite graph has a
395 significant effect on the performance of BiG-DRP+. To address this question, we used Morgan
396 fingerprints [13] of the drugs, alone or in addition to the drug descriptors, as the attributes of the
397 drug nodes in the bipartite graph. The results (provided in Table 3) revealed that there is not a

398 substantial difference between any of these choices, but simultaneously using both types of drug
399 features slightly improves the results.

400

401 **Table 3:** The performance of BiG-DRP+ with different drug attributes. The rows show the results
402 of BiG-DRP+ when drug descriptors (vectors of length 198), Morgan fingerprints (vectors of length
403 512), or the combination of both (vectors of length 710) are used as node attributes.
404

| Method | Drug Attribute | LPO-CV | | LCO-CV | |
|---|---|---|---|---|---|
| | | AUROC mean (± std.) | SCC mean (± std.) | AUROC mean (± std.) | SCC mean (± std.) |
| BiG-DRP+ | Descriptors | 0.878 (±0.068) | 0.748 (±0.100) | 0.746 (±0.077) | 0.431 (±0.094) |
| | Morgan | 0.878 (±0.068) | 0.748 (±0.100) | 0.743 (±0.080) | 0.426 (±0.098) |
| | Both | 0.879 (±0.068) | 0.748 (±0.099) | 0.746 (±0.077) | 0.433 (±0.095) |

405

406 Next, we sought to better characterize the bipartite graph and the drugs that have most
407 benefited from using this graph in the drug response prediction task. For this purpose, we first
408 formed a single bipartite graph by aggregating the bipartite graphs corresponding to each of the
409 five folds in our LCO-CV evaluation (i.e., by finding the union of edges). Then, we used a nested
410 stochastic block model [33] to infer the modular substructure of the graph, while taking into
411 account the edge type (i.e., resistant and sensitive) connecting each two nodes. Figure 3A
412 illustrates the bipartite graph and clusters identified using this method (see Supplementary Table
413 S4 for the cluster assignment of drugs and CCLs). In particular, five drug clusters were identified.
414 Comparing the performance of BiG-DRP+ compared to MLP (SCC-LCO), revealed that all these
415 clusters significantly benefit from the use of the bipartite graph (one-sided Wilcoxon signed rank
416 test, Figure 3B). In particular, Cluster 3 had the highest median improvement in SCC (8.4%) and

had a significant improvement p-value (p = 5.25 E-5). The majority of the drugs in this cluster (13 out of 20) are protein kinase inhibitors, with 8 of them targeting members of serine/threonine protein kinase family and 5 of them targeting members of tyrosine kinase family. These observations suggest that information sharing across the bipartite graph used in our methods benefit certain groups of drugs more than others and this may be dependent on the similarity between drugs' mechanisms of action.

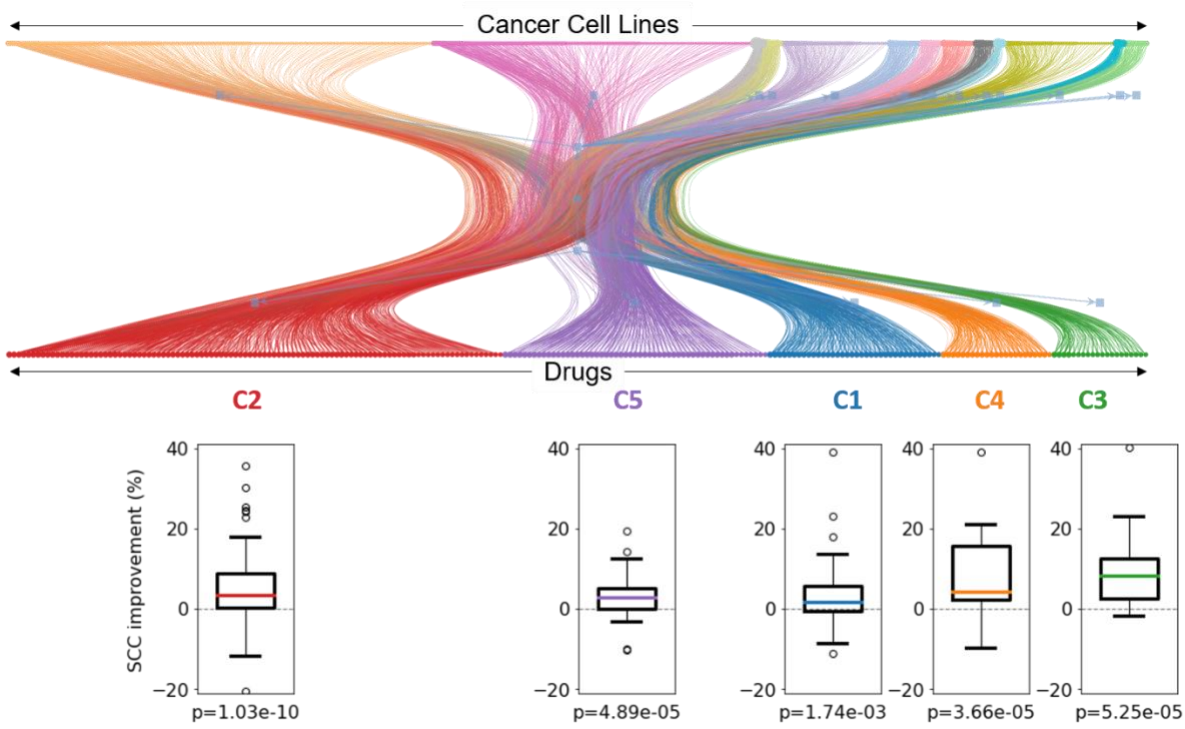


**Figure 3:** The CCL-drug bipartite graph and its clusters. The boxplots show the distribution of SCC improvements obtained for each drug using BiG-DRP+ compared to MLP in the LCO evaluation. The p-values are obtained using a one-sided Wilcoxon signed rank test.

## Identification of biomarkers of drug sensitivity

To identify genes whose expression substantially contribute to the predictive model, we used a pipeline similar to the one we proposed in a previous study [5]. This approach provides an aggregate contribution score for each gene in the model and uses these scores to systematically

identify the set of top contributing genes in each model. We focused on 15 drugs for which BiG-DRP+ provided the highest SCC values in the LCO-CV evaluation. Supplementary Table S5 provides the ranked list of genes that were implicated for each of the 15 drugs. We clustered the drugs based on the contribution scores of all implicated genes (Figure 4). Interestingly, four drugs formed a clear cluster, separate from the others: trametinib, refametinib, selumetinib, and pd0325901. Further investigation revealed that these drugs all are MEK inhibitors (i.e., inhibit the mitogen-activated protein kinase kinase enzymes) and involve some similar mechanisms of action [10].

Next, we focused on genes implicated for trametinib, a MEK-inhibitor for which BiG-DRP+ had the best performance (SCC in LCO-CV). For this drug, *ETV5* had the highest prediction contribution. ETV5 and ETV4 (the fourth highest contributor) are among the ETS family of oncogenic transcription factors. The expression of this family has been shown to be upregulated in solid tumours and they have been shown to be involved in tumour's progression, metastasis and chemoresistance [34]. Previous studies have shown ETV5 to be regulated by ALK, a receptor tyrosine kinase, in a MEK/ERK-dependent manner in neuroblastoma cell lines [35]. In addition, treatment of various cancer cell lines with trametinib has been shown to downregulate ETV5 [35-37]. Moreover, the overexpression of ETV4 and ETV5 have been shown to reduce the sensitivity different cancer cell lines to this drug [37].

To obtain a better functional characteristic of the genes implicated for trametinib, we also performed pathway enrichment analysis on genes implicated for this drug (see Supplementary

Table S6 for results of pathway enrichment analysis of all 15 drugs). Several important pathways related to MAPK signaling, EGFR signaling, and IL2 signaling were identified (Fisher's exact test, FDR<0.05). Taken together, these results suggest that genes that contribute to the predictive ability of BiG-DRP+ for trametinib point to important genes and signaling pathways involved in its mechanism of action.

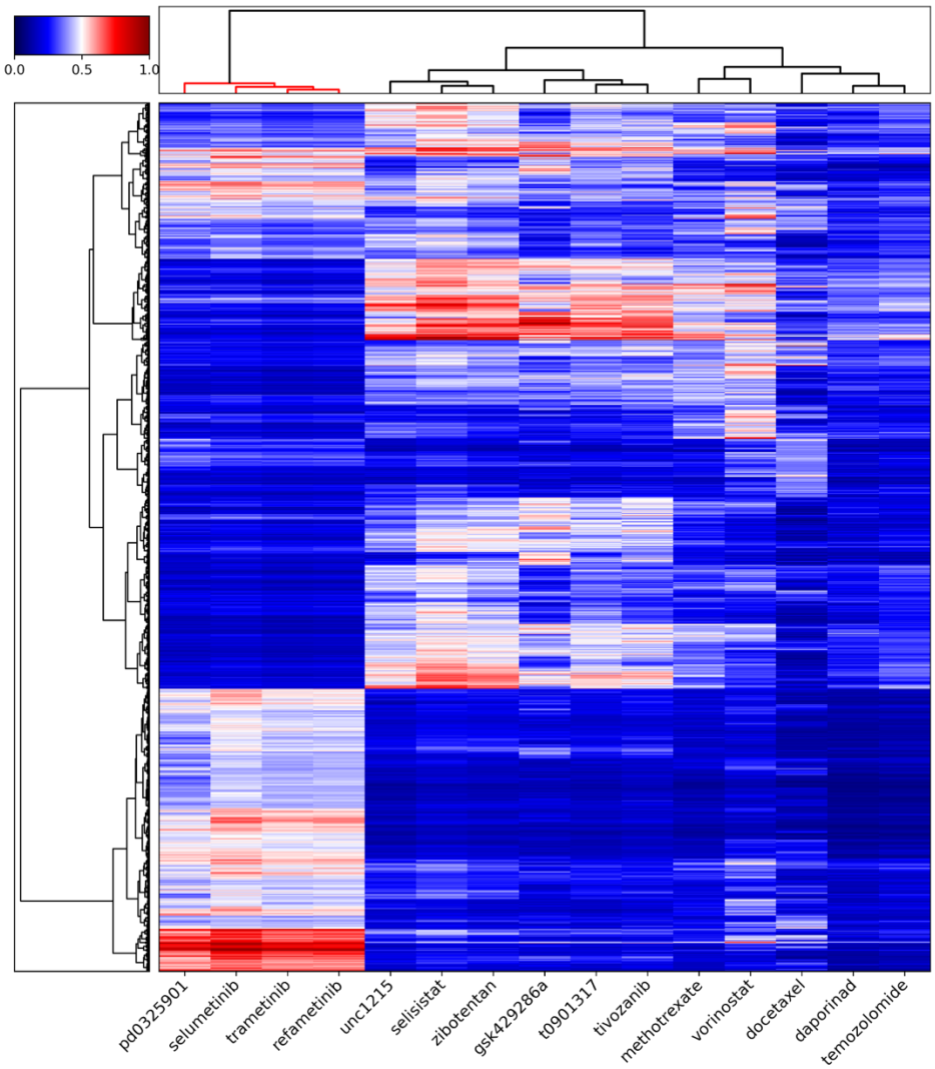


**Figure 4:** The clustering of 15 drugs based on contribution scores of their genes. The contribution scores of the union of genes implicated for these drugs is used to cluster drugs using hierarchical clustering. The heatmap shows the contribution scores.

## DISCUSSION AND CONCLUSION

In this study, we proposed two novel graph representation deep learning methods to incorporate information regarding the sensitivity and resistance of cell lines, their gene expression profiles and chemical drug attributes to obtain better drug representations. Using cross-validation and different data splitting methods we showed significant improvement compared to traditional and state-of-the-art methods. Using a computational pipeline to make neural networks explainable, we identified a set of genes that substantially contribute to the predictive model. These genes implicated important signaling pathways and pointed to shared and unique mechanisms of action in the drugs.

Moreover, detailed evaluation of our methods showed a high degree of robustness towards changes in the threshold used to form the bipartite graph. This further supports the importance of different techniques we used to ensure stability of our proposed architecture: the normalization factor and the injected self-loop in our H-GCN's forward pass. More specifically, due to the injected self-loop, the nodes retain a portion of their own information, which forces the embeddings to have some level of separation. The normalization factor also helps by preventing the received messages from becoming too large and overpowering the self-loop. It is important to note that this robustness may not be applicable to some corner cases. For example, when a drug's connected CCLs are not connected to any other drug (i.e., it forms a disconnected star subgraph), this drug's embedding would not benefit from the existence of the second H-GCN layer. As another example, the second H-GCN layer will be obsolete if all the drugs happen to form disconnected stars, and thus no information sharing will take place across drugs. Another

example is when we add a new drug that results in a disconnected node. A disconnected node will not be able to incorporate CCL information into the drug embedding, which defeats the purpose of the H-GCN.

Unlike many previous models (e.g., NRL2DRP [25]) that require both cell lines and drugs to be present in the training set, BiG-DRP is designed to enable prediction of unseen cell lines (those that are not present in the training set). However, the drug embedding part of the model (the H-GCN) requires the drugs to be part of the bipartite graph. This constraint implies that the drugs present in the test set must be also present in the training set. As a result, this model generally is not applicable to predict the response of CCLs to unseen new drugs. Although this could be naively remedied by assuming known edges involving the unseen drug in the bipartite graph, this kind of solution is impractical and would be difficult to enact without reducing the test set. However, in most practical applications (e.g., prediction of drug response of cancer patients [4] and [5]), it is more crucial for the model to generalize to unseen samples (CCLs or patients). The reason is that before a drug enters clinical trial or enters clinical usage, many *in-vitro* studies on CCLs are first performed. Consequently, one can expect to have access to molecular description and drug response of a drug for which the drug responses of a new set of samples (CCLs or patients) are to be predicted.

In summary, this study proposes new techniques to improve drug response prediction and to identify biomarkers of drug sensitivity.

**Funding:**

This work was supported by the Government of Canada's New Frontiers in Research Fund (NFRF) [NFRFE-2019-01290] (AE), by Natural Sciences and Engineering Research Council of Canada (NSERC) grant RGPIN-2019-04460 (AE), and by McGill Initiative in Computational Medicine (MiCM) (AE). This work was also funded by Génome Québec, the Ministère de l'Économie et de l'Innovation du Québec, IVADO, the Canada First Research Excellence Fund and Oncopole, which receives funding from Merck Canada Inc. and the Fonds de Recherche du Québec – Santé (AE).

**Authors' contributions:**

AE and DEH conceived the study and designed the project. DEH designed the algorithm and implemented the pipeline. DEH and YL ran the baseline methods and performed the statistical analyses of the results. All authors read and approved the final manuscript.

## SUPPLEMENTARY TABLES

**Supplementary Table S1:** The performance of BiG-DRP+, BiG-DRP and baseline methods (columns) using five-fold CV for each drug (rows). Values presented are the mean across five folds. Each tab corresponds to a performance metric (Spearman's correlation coefficient and area under the receiver operating characteristic) for each data-splitting scenario: leave-cell lines-out (LCO) and leave-pairs-out (LPO).

**Supplementary Table S2:** The results of the one-sided Wilcoxon signed rank test, represented by the p-values. The test compares the per-drug Spearman's correlation coefficient (SCC) of the BiG-DRP+ and the other methods, where the alternative is that BiG-DRP+'s SCC is significantly larger.

**Supplementary Table S3:** The performance of BiG-DRP and BiG-DRP+ using for $k \in \{0.5, 1, 2, 5, 10\}$. The performance metrics were calculated independently per drug and presented as the mean and standard deviation.

**Supplementary Table S4:** The cluster assignments of the aggregated bipartite graph using the nested stochastic block model. The first tab shows that cluster indices for the drugs and the second tab shows the cluster indices for the cell lines. Although the graph partitioning was performed on the entire graph, drug and cell line clusters were mutually exclusive so similar indices in different tabs do not pertain to the same cluster.

**Supplementary Table S5:** The top genes and their normalized contribution scores for top-performing drugs in the leave-cell lines-out scenario. Each tab corresponds to the top genes for a specific drug.

**Supplementary Table S6:** The results of the pathway enrichment analysis of the top genes (Supplementary Table S5) on the Reactome pathways using Fisher's exact test. Each tab corresponds to a specific drug. The corrected p-values are indicated in the *pvalue_cor* column.

## REFERENCES


1. Costello JC, Heiser LM, Georgii E, Gonen M, Menden MP, Wang NJ, Bansal M, Ammad-ud-din M, Hintsanen P, Khan SA, et al: **A community effort to assess and improve drug sensitivity prediction algorithms.** *Nat Biotechnol* 2014, **32:**1202-1212.
2. Sharifi-Noghabi H, Zolotareva O, Collins CC, Ester M: **MOLI: multi-omics late integration with deep neural networks for drug response prediction.** *Bioinformatics* 2019, **35:**i501-i509.
3. Geeleher P, Zhang Z, Wang F, Gruener RF, Nath A, Morrison G, Bhutra S, Grossman RL, Huang RS: **Discovering novel pharmacogenomic biomarkers by imputing drug response in cancer patients from large genomics studies.** *Genome research* 2017, **27:**1743-1751.
4. Huang EW, Bhope A, Lim J, Sinha S, Emad A: **Tissue-guided LASSO for prediction of clinical drug response using preclinical samples.** *PLoS computational biology* 2020, **16:**e1007607.
5. Hostallero DE, Wei L, Wang L, Cairns J, Emad A: **A Deep Learning Framework for Prediction of Clinical Drug Response of Cancer Patients and Identification of Drug Sensitivity Biomarkers using Preclinical Samples.** *bioRxiv* 2021, **2021.07.06.451273**.
6. Yang W, Soares J, Greninger P, Edelman EJ, Lightfoot H, Forbes S, Bindal N, Beare D, Smith JA, Thompson IR, others: **Genomics of Drug Sensitivity in Cancer (GDSC): a resource for therapeutic biomarker discovery in cancer cells.** *Nucleic acids research* 2012, **41:**D955-D961.
7. Deng L, Cai Y, Zhang W, Yang W, Gao B, Liu H: **Pathway-Guided Deep Neural Network toward Interpretable and Predictive Modeling of Drug Sensitivity.** *J Chem Inf Model* 2020, **60:**4497-4505.
8. Nguyen T-T, Nguyen GTT, Nguyen T, Le D-H: **Graph convolutional networks for drug response prediction.** *IEEE/ACM Transactions on Computational Biology and Bioinformatics* 2021.
9. Liu P, Li H, Li S, Leung KS: **Improving prediction of phenotypic drug response on cancer cell lines using deep convolutional network.** *BMC Bioinformatics* 2019, **20:**408.
10. Kim S, Chen J, Cheng T, Gindulyte A, He J, He S, Li Q, Shoemaker BA, Thiessen PA, Yu B, et al: **PubChem in 2021: new data content and improved web interfaces.** *Nucleic Acids Res* 2021, **49:**D1388-D1395.
11. Gaulton A, Hersey A, Nowotka M, Bento AP, Chambers J, Mendez D, Mutowo P, Atkinson F, Bellis LJ, Cibrian-Uhalte E, et al: **The ChEMBL database in 2017.** *Nucleic Acids Res* 2017, **45:**D945-D954.
12. Landrum G: **RDKit: Open-source cheminformatics.** 2010.

13. Rogers D, Hahn M: **Extended-connectivity fingerprints.** *Journal of chemical information and modeling* 2010, **50:**742-754.
14. Duvenaud D, Maclaurin D, Aguilera-Iparraguirre J, Gómez-Bombarelli R, Hirzel T, Aspuru-Guzik A, Adams RP: **Convolutional networks on graphs for learning molecular fingerprints.** In *Advances in Neural Information Processing Systems (NeurIPS)*. pp. 2224–2232. Montreal, Canada: MIT Press; 2015:2224–2232.
15. Wang L, Li X, Zhang L, Gao Q: **Improved anticancer drug response prediction in cell lines using matrix factorization with similarity regularization.** *BMC Cancer* 2017, **17:**513.
16. Suphavilai C, Bertrand D, Nagarajan N: **Predicting cancer drug response using a recommender system.** *Bioinformatics* 2018, **34:**3907-3914.
17. Yang J, Li A, Li Y, Guo X, Wang M: **A novel approach for drug response prediction in cancer cell lines via network representation learning.** *Bioinformatics* 2019, **35:**1527-1535.
18. Kipf TN, Welling M: **Semi-supervised classification with graph convolutional networks.** In *International Conference on Learning Representations (ICLR)*. 2017
19. Weininger D: **SMILES, a chemical language and information system. 1. Introduction to methodology and encoding rules.** *Journal of Chemical Information and Modeling* 1988, **28:**31-36.
20. van der Meer D, Barthorpe S, Yang W, Lightfoot H, Hall C, Gilbert J, Francies HE, Garnett MJ: **Cell Model Passports-a hub for clinical, genetic and functional datasets of preclinical cancer models.** *Nucleic Acids Res* 2019, **47:**D923-D929.
21. Kingma DP, Ba J: **Adam: A method for stochastic optimization.** In *International Conference on Learning Representations (ICLR)*. 2015
22. Cancer Genome Atlas Research N, Weinstein JN, Collisson EA, Mills GB, Shaw KR, Ozenberger BA, Ellrott K, Shmulevich I, Sander C, Stuart JM: **The Cancer Genome Atlas Pan-Cancer analysis project.** *Nat Genet* 2013, **45:**1113-1120.
23. Geeleher P, Cox NJ, Huang RS: **Clinical drug response can be predicted using baseline gene expression levels and in vitro drug sensitivity in cell lines.** *Genome Biol* 2014, **15:**R47.
24. Williams C, Seeger M: **Using the Nyström method to speed up kernel machines.** In *Advances in Neural Information Processing Systems (NeurIPS)*. 2001: 682-688.
25. Tang J, Qu M, Wang M, Zhang M, Yan J, Mei Q: **Line: Large-scale information network embedding.** In *International Conference on World Wide Web (WWW)*. 2015: 1067-1077.
26. Kanehisa M, Goto S: **KEGG: kyoto encyclopedia of genes and genomes.** *Nucleic Acids Res* 2000, **28:**27-30.
27. Szklarczyk D, Santos A, von Mering C, Jensen LJ, Bork P, Kuhn M: **STITCH 5: augmenting protein-chemical interaction networks with tissue and affinity data.** *Nucleic Acids Res* 2016, **44:**D380-384.
28. Schwab P, Karlen W: **CXPlain: Causal Explanations for Model Interpretation under Uncertainty.** In *Advances in Neural Information Processing Systems (NeurIPS)*. 2019
29. Granger CW: **Investigating causal relations by econometric models and cross-spectral methods.** *Econometrica: journal of the Econometric Society* 1969**:**424-438.
30. Satopaa V, Albrecht J, Irwin D, Raghavan B: **Finding a "Kneedle" in a Haystack: Detecting Knee Points in System Behavior.** In *2011 31st International Conference on Distributed Computing Systems Workshops*. 2011: 166-171.
31. Blatti C, 3rd, Emad A, Berry MJ, Gatzke L, Epstein M, Lanier D, Rizal P, Ge J, Liao X, Sobh O, et al: **Knowledge-guided analysis of "omics" data using the KnowEnG cloud platform.** *PLoS Biol* 2020, **18:**e3000583.
32. Jassal B, Matthews L, Viteri G, Gong C, Lorente P, Fabregat A, Sidiropoulos K, Cook J, Gillespie M, Haw R, et al: **The reactome pathway knowledgebase.** *Nucleic Acids Res* 2020, **48:**D498-D503.

626 33. Peixoto TP: **Hierarchical block structures and high-resolution model selection in large networks.**
627 *Physical Review X* 2014, **4:**011047.
628 34. Sizemore GM, Pitarresi JR, Balakrishnan S, Ostrowski MC: **The ETS family of oncogenic**
629 **transcription factors in solid tumours.** *Nat Rev Cancer* 2017, **17:**337-351.
630 35. Lopez-Delisle L, Pierre-Eugene C, Louis-Brennetot C, Surdez D, Raynal V, Baulande S, Boeva V,
631 Grossetete-Lalami S, Combaret V, Peuchmaur M, et al: **Activated ALK signals through the ERK-**
632 **ETV5-RET pathway to drive neuroblastoma oncogenesis.** *Oncogene* 2018, **37:**1417-1429.
633 36. Ranzani M, Alifrangis C, Perna D, Dutton-Regester K, Pritchard A, Wong K, Rashid M, Robles-
634 Espinoza CD, Hayward NK, McDermott U, et al: **BRAF/NRAS wild-type melanoma, NF1 status and**
635 **sensitivity to trametinib.** *Pigment Cell Melanoma Res* 2015, **28:**117-119.
636 37. Wang B, Krall EB, Aguirre AJ, Kim M, Widlund HR, Doshi MB, Sicinska E, Sulahian R, Goodale A,
637 Cowley GS, et al: **ATXN1L, CIC, and ETS Transcription Factors Modulate Sensitivity to MAPK**
638 **Pathway Inhibition.** *Cell Rep* 2017, **18:**1543-1557.

639