Knowledge-Guided Domain Adaptation Model for Transferring Drug Response Prediction from Cell Lines to Patients

Xuan Liu[1*], **Menglu Li**[2]

[1]School of Computer and Information Technology, Xinyang Normal University, Xinyang 464000, China
[2]College of Informatics, Huazhong Agricultural University, Wuhan 430070, China
lx666@xynu.edu.cn, mengluli@webmail.hzau.edu.cn

Abstract

Drug response prediction (DRP) is a longstanding challenge in modern oncology that underpins personalized treatment. Early DRP methods, trained on label-rich cell line samples, suffer from performance degradation when applied to label-scarce patient samples due to the distribution shift. Recently, a few transfer learning efforts have addressed this issue by aligning cell line (source domain) and patient (target domain) data via unsupervised domain adaptation (UDA). However, these efforts often treat each drug's response prediction as an isolated task, requiring model retraining when the drug changes; and focus only on aligning data distributions as a whole, neglecting the category (e.g., different cancers or tissues) confusion problem. To address these limitations, we propose a knowledge-guided domain adaptation model to transfer the DRP from cell lines to patients, named TransDRP. Specifically, TransDRP operates in two phases: pre-training and adaptation. In the first phase, we pre-train a multi-label graph neural network using molecular knowledge, to simultaneously predict responses for various drugs and capture their interdependencies. In the second phase, we implement a global-local domain adversarial strategy with clinical knowledge, to encourage representation alignment within same cancer categories and separation among different cancer categories across domains. Extensive experiments demonstrate that TransDRP outperforms state-of-the-art UDA methods in both transfer efficiency and precision for the patient DRP.

Code — https://github.com/liuxuan666/TransDRP

Introduction

Identification of drug response in cancer patients is pivotal for customizing therapy regimens to increase survival and reduce expenses. Since the financial limitation of testing multiple drugs on a patient (Adam et al. 2020), there is an urgent demand for computational models that can deliver drug response prediction (DRP) efficiently. Although the easy availability of genomic profiles for clinical patients (e.g., TCGA (Weinstein et al. 2013)), the corresponding drug response labels are insufficient to train a precise model. In this context, large-scale labeled data derived from preclinical cell lines (e.g., GDSC (Iorio et al. 2016)) have been considered as a viable alternative. Unfortunately, models trained on cell lines often fail to generalize in DRP for patients, due to the absence of an *in vivo* tumor micro-environment, resulting in significant discrepancies between genomic data of cell lines and patients (Seyhan 2019). In fact, such data discrepancy can be viewed as a domain-shift or out-of-distribution (OOD) (Liu et al. 2020a) issue, where cell lines/patients are referred to as the source/target domain. How to bridge the gap between the two domains holds great promise for achieving reliable patient-oriented DRP.

Recently, unsupervised domain adaptation (UDA) (Ganin and Lempitsky 2015), as a practical technique of transfer learning, has become the model of choice to tackle the OOD issue, with the assumption that labeled data are available only in a source domain. The core idea of UDA is to minimize discrepancies in genomic data distribution by learning invariant representations shared between source (cell line) and target (patient) domain, and then transfer the DRP model from cell lines to patients. For instance, discrepancy-based UDA methods (Warren et al. 2021; Chen et al. 2022) focused on aligning genomic data globally within a low-dimensional embedding space, using distance measures like maximum mean discrepancy (Borgwardt et al. 2006) or correlation alignment (Sun and Saenko 2017) functions. But these methods only considered domain alignment at the embedding level, without optimizing downstream DRP tasks, thus leading to suboptimal outcomes. In contrast, adversarial-based UDA methods (Peres da Silva et al. 2021; He et al. 2022; Shubham et al. 2024) offered a more powerful solution through an adversarial objective, inspired by generative adversarial networks (GANs) (Goodfellow et al. 2020). Technically, they worked by a genomic extractor and a domain discriminator competing with each other, wherein the former learns invariant representations that confuse the latter, which in turn tries to distinguish them from real domains. This principle enforces the trained DRP model not to differentiate between inputs from cell lines or patients.

Despite the commendable progress of efforts above, they typically treat the prediction of each drug's response as an individual task during the UDA process, so that the model needs to be retrained when drug changes. Obviously, this strategy is inefficient and lacks correlations across different drugs. A few methods have attempted to address this prob-

*Corresponding authors.

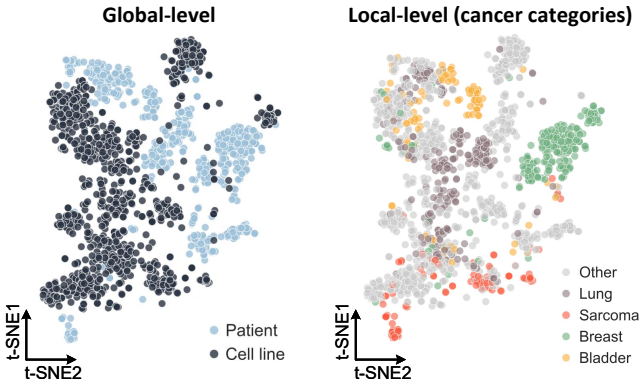

Figure 1: t-SNE visualization of cell line data (GDSC) and patient data (TCGA) on the gene expression spaces.

lem by integrating various drugs under a multi-task manner (Peres da Silva et al. 2021; Shubham et al. 2024), but they still face computational burdens induced by the number of drugs. Additionally, present UDA methods focused on aligning global data distributions between the source and target domains, but ignored the category (e.g., different cancers or tissues) confusion problem. In clinical, the distribution of different categories of genomic samples varies greatly, and confused alignment together will undoubtedly harm the performance of domain transfer. For instance, lung cancer samples of cell lines differ significantly from breast cancer samples of patients in the feature space (as shown in Figure 1), and aligning them would contradict the biology understanding. Consequently, it is imperative to call for a stronger UDA framework to jointly address these limitations.

Herein, we propose TransDRP, a knowledge-guided domain adaptation model for transferring DRP from cell lines (source) to patients (target). Specifically, TransDRP follows a two-phase fashion of pre-training and adaptation. In the first phase, we pre-train an encoder to extract shared representations from genomic profiles across both domains, and a multi-label graph neural network (GNN) decoder to predict various drug responses simultaneously on the source domain. One unique advantage of our decoder is treating each drug's response as a label, allowing the incorporation of molecular similarity knowledge to better capture dependencies among drugs (labels). In the second phase, we categorize data into different cancers based on clinical knowledge, and then implement a global-local domain adversarial strategy to generalize the multi-label decoder to target domain. The global one promotes entire alignment between two domains through a discriminator, while the local one encourages intra-category alignment and inter-category separation through prototypical contrastive learning. The contributions of this work are summarized as follows:

- We implement a multi-label GNN via molecular knowledge, capable of predicting various drug responses and capturing dependencies among them, whereas previous UDA methods focused only on a single-drug transfer.
- We devise a global-local domain adversarial strategy via clinical knowledge, to better align cross-domain representations within similar cancer categories while separating confusable ones during the DRP transferring.
- Extensive experiments show that TransDRP surpasses the state-of-the-art UDA baselines, particularly in the efficiency and precision of DRP for patients.

Related Works

Drug response prediction

The past few years have witnessed a proliferation of computational DRP methods (Firoozbakht and Schwikowski 2022), which relate inputs (drug structures, cell line omics, or response associations) to desired outputs (sensitivity classifications or specific IC_{50} values). Early network-based methods constructed bio-networks with cell lines and drugs, then converted the problem to a link prediction task, undertaken on the matrix completion (Wang et al. 2017) or random walk (Stanfield, Coşkun, and Koyutürk 2017). Machine learning-based methods assembled handcrafted features to represent cell lines and drugs, so as to train a response predictor, such as support vector machine (Yang et al. 2019) and random forest (Lind and Anderson 2019). Deep learning-based methods exhibited stronger performance by providing an end-to-end solution from feature extraction to response prediction (Liu et al. 2020b; Hostallero, Li, and Emad 2022; Liu and Zhang 2023). Nowadays, there is a growing trend toward transfer learning frameworks, such as few-shot (Ma et al. 2021), fine-tuning (Hostallero et al. 2023), and UDA (Mourragui et al. 2019; Sharifi-Noghabi et al. 2021), aimed at deploying DRP model from cell lines (seen) to patients (unseen). Of them, adversarial-based UDA (Ganin and Lempitsky 2015) was wildly used for the alignment of cell line and patient data, attributed to their superiority of domain-invariant learning compared to prior techniques. However, they usually require building one model per drug during the UDA process, which leads to computational inefficiencies. Thus, we moved a step forward and implemented a multi-label paradigm into this field, aiming to accelerate prediction for various drugs and capture dependencies among them.

Unsupervised domain adaptation

Being popular in transfer learning, UDA allows models to transfer knowledge learned from source domains with abundant labeled training samples to target domains with unlabeled data only, the application of which ranges from computer vision to bioinformatics (Liu et al. 2022c; Bai et al. 2023; Cui et al. 2020). Learning domain-invariant representation is critical for UDA, and a line of works attempted to align the source and target distributions measured by discrepancy metrics or domain adversarial training, where the latter dominated in the top performance methods (Liu et al. 2022b). For example, the domain adversarial training presented by Ganin and Lempitsky (2015), introduced a discriminator to distinguish between target and source features, an extractor to fool this discriminator by generating domain-invariant features. On that basis, Bousmalis et al. (2016) explicitly learned image representations that are partitioned into two subspaces: one private to each domain and one shared across domains. (Yuan et al. 2022) further advanced

this with a category-level alignment mechanism, which encourages the cross-domain consistency between same categories and differentiation among diverse categories. Recently, the evolution of UDA has also sparked a rising interest in extending it to the DRP tasks (Shubham et al. 2024). Despite their success, the category-level alignment between cell lines and patients has not been explored yet. To fill this gap, we incorporate clinical cancer categories into UDA for category-divergence domain alignment.

Methodology

In this section, we first formulate the UDA task of cell line-to-patient DRP. Subsequently, we elaborate on our method TransDRP, encompassing the pre-training and adaptation two phases. The former initializes an extractor for invariant domain representations and a drug response classifier for the source domain, and the latter transfers the classifier to the target domain with these invariant representations. The overview of TransDRP is illustrated in Figure 2.

Problem formulation

Suppose that the source domain is a cell line-oriented dataset $\mathcal{D}^s = \{(\boldsymbol{x}_1^s, \boldsymbol{y}_1), ..., (\boldsymbol{x}_N^s, \boldsymbol{y}_N)\}$ composed of N sample-label pairs, and the target domain is a patient-oriented dataset $\mathcal{D}^t = \{\boldsymbol{x}_1^t, ..., \boldsymbol{x}_M^t\}$ composed of M samples without labels. Here, $\boldsymbol{x}_i^s$ and $\boldsymbol{x}_i^t$ stand for the input genomic feature of i-th cell line and patient, respectively. The label $\boldsymbol{y}_i = [y_{k,i}]_{k=1}^K$ denotes the response outcomes to K different drugs, where $y_{k,i} \in \{-1, 0, 1\}$ with '1'/'0' indicating a positive/negative response classification of cell line i to drug k, and '-1' indicating that the response label is missing. Our goal is to train a DRP model (including representation extraction and response prediction) on the source domain that can accurately infer sample labels for the target domain.

Pre-training phase

Domain representation extractor. Extracting the invariant representation of source and target domains is essential for UDA, as it reduces input noise and produces more consistent information for downstream domain alignment. To do so, we utilize the pre-trained autoencoder (composed of neural networks) to acquire domain-shared and domain-specific representations, following Bousmalis et al. (2016).

Specifically, each input $\boldsymbol{x}$ will be encoded into two separate representations: one by its corresponding cell line or patient private encoder $\boldsymbol{f}_p$, and another by a shared encoder $\boldsymbol{f}_e$. The concatenation of two representations is then used to reconstruct the proximate genomic feature $\hat{\boldsymbol{x}}$ via a shared decoder $\boldsymbol{g}_e$, just as follows:

$$\hat{\boldsymbol{x}} = \boldsymbol{g}_e\left(\boldsymbol{f}_p(\boldsymbol{x}') \oplus \boldsymbol{f}_e(\boldsymbol{x}')\right) \tag{1}$$

where $\boldsymbol{x}'$ is obtained by adding random Gaussian noise to $\boldsymbol{x}$, $\oplus$ stands for the vector concatenation operation. The loss of mean square error (MSE) is adopted to minimize the distance between $\boldsymbol{x}$ and $\hat{\boldsymbol{x}}$, ensuring consistency in genomic features before and after reconstruction:

$$\mathcal{L}_{\text{rec}} = \frac{1}{N}\sum_{i=1}^{N}\left(\boldsymbol{x}_i^s - \hat{\boldsymbol{x}}_i^s\right)^2 + \frac{1}{M}\sum_{j=1}^{M}\left(\boldsymbol{x}_j^t - \hat{\boldsymbol{x}}_j^t\right)^2 \tag{2}$$

Moreover, a difference loss is applied to encourage encoders to produce such separate representations. To this end, we introduce a soft subspace orthogonality constraint between private and shared representations in each domain:

$$\mathcal{L}_{\text{dif}} = \left\|\boldsymbol{H}^s \cdot \boldsymbol{O}^{s^T}\right\|_2 + \left\|\boldsymbol{H}^t \cdot \boldsymbol{O}^{t^T}\right\|_2 \tag{3}$$

where $\boldsymbol{H}^s$, $\boldsymbol{H}^t$ (or $\boldsymbol{O}^s$, $\boldsymbol{O}^t$) are matrices whose rows indicate the shared (or private) representations $\boldsymbol{h}^s = \boldsymbol{f}_e(\boldsymbol{x}^s)$, $\boldsymbol{h}^t = \boldsymbol{f}_e(\boldsymbol{x}^t)$ (or $\boldsymbol{o}^s = \boldsymbol{f}_p(\boldsymbol{x}^s)$, $\boldsymbol{o}^t = \boldsymbol{f}_p(\boldsymbol{x}^t)$) from samples of source and target domain, respectively; $\|\cdot\|_2$ means the squared Frobenius norm; T is the transpose operation.

By combining the above two losses, we can ensure that the shared representation remains unaffected by the private representation of each domain, thereby yielding more invariant information that generalize well across domains:

$$\mathcal{L}_{\text{pre}} = \mathcal{L}_{\text{rec}} + \mathcal{L}_{\text{dif}} \tag{4}$$

After pre-training the autoencoder for τ_a epochs, the shared encoder $\boldsymbol{f}_e$ will serve as the domain representation extractor.

Multi-label response classifier. Using the shared representations, we further pre-train a drug response classifier on the labeled source domain, to provide prior knowledge for subsequent fine-tuning on the unlabeled target domain. Instead of modeling a single-drug classifier like previous UDA methods, here we centre on predicting responses between each cell line and different drugs in a multi-label fashion, where each label represents an individual drug's response outcome (sensitive or resistant). However, considering each label independently may diminish prediction accuracy when deploying multi-label learning (Chen et al. 2021); thus, we create a classifier based on graph neural networks, so as to capture dependencies within the label space.

For each cell line, we start by establishing an undirected label graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{X})$ whose structure is known. Here, $\mathcal{V} = \{v_1, ..., v_K\}$ stands for the set of drug (label) nodes, $\mathcal{E} \subset \mathcal{V} \times \mathcal{V}$ is the set of edges measured by the co-occurrence rate (Huynh and Elhamifar 2020) of pairwise nodes (labels) in the dataset, and $\mathcal{X} \in \mathbb{R}^{K \times F}$ consists of node attributes initialized by concatenating different drug molecular vectors $\{\boldsymbol{d}_1, ..., \boldsymbol{d}_K\}$ with the shared representations $\boldsymbol{h}$. Specifically, $\mathcal{X} = [\boldsymbol{h} \oplus \boldsymbol{d}_1, ..., \boldsymbol{h} \oplus \boldsymbol{d}_K]$, where $\boldsymbol{d}$ is derived from the extended connectivity fingerprint (ECFP) (Rogers and Hahn 2010). Next, we leverage a graph attention network (GAT) (Veličković et al. 2017) decoder $\boldsymbol{g}_a$ to graph $\mathcal{G}$, so that the associations/dependencies among all labels can be captured and aggregated. Concretely, $\boldsymbol{g}_a$ computes the embedding $\boldsymbol{z}$ of each node by iteratively convolving over neighbour nodes using the below propagation rule:

$$\boldsymbol{z}_v^l = \boldsymbol{\phi}\left(\boldsymbol{z}_v^{l-1}, \sum_{u \in \mathcal{N}(v)} \boldsymbol{a}\left(\boldsymbol{z}_v^{l-1}, \boldsymbol{z}_u^{l-1}\right) \boldsymbol{\psi}\left(\boldsymbol{z}_u^{l-1}\right)\right) \tag{5}$$

where $\boldsymbol{z}_v^l = \boldsymbol{Z}^l[v, :]$ is the hidden embedding of node v in the l-th layer with $\boldsymbol{Z}^0 = \mathcal{X}$; $\boldsymbol{\phi}$ and $\boldsymbol{\psi}$ are neural networks used for combining and extracting embeddings, respectively; $\mathcal{N}(v)$ denotes a set of nodes adjacent to v in $\mathcal{E}$; $\boldsymbol{a}$ calculates the attention weight from node v to u.

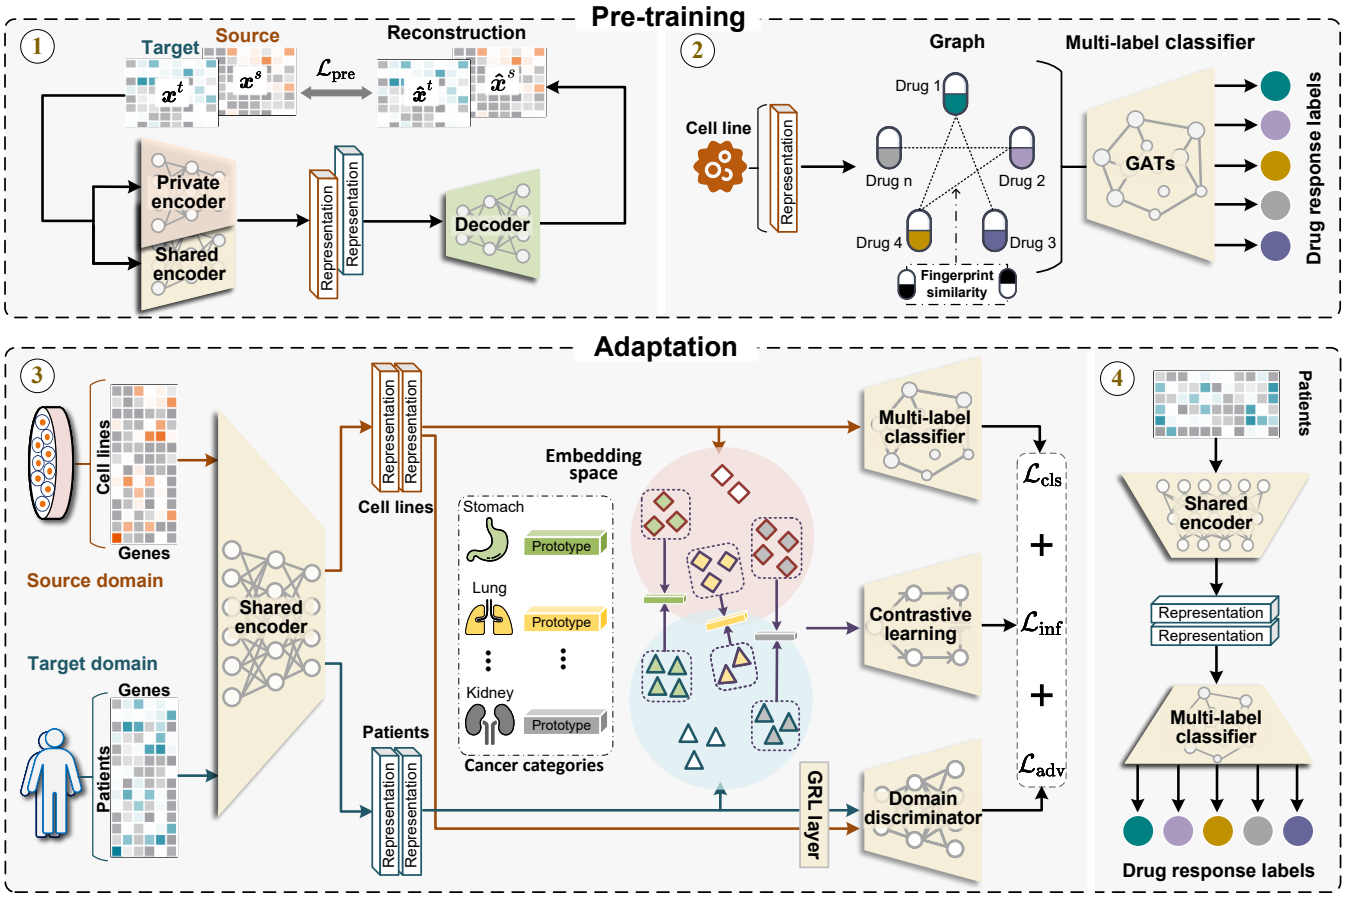

Figure 2: Overview of the TransDRP. TransDRP comprises a pre-training phase for building ① domain representation extractor and ② multi-label classifier on the source domain, as well as an adaptation phase that involves ③ domain adversarial training and ④ transferring drug response prediction to the target domain.

Lastly, we map the hidden embedding of each node to a one-dimensional scalar in the last layer of GAT: $\hat{y} \in \mathbb{R}^1 \leftarrow \mathrm{Sigmoid}(\boldsymbol{z}^l)$. And all nodes' outputs can be arranged as a vector $\hat{\boldsymbol{y}} \in \mathbb{R}^K$, representing predicted sensitivity probabilities to K drugs. Based on this, the objective of multi-label classification is to minimize the below cross-entropy loss:

$$\mathcal{L}_{\text{cls}} = -\frac{1}{N}\sum_{i=1}^{N} \boldsymbol{y}_i \log\left(\hat{\boldsymbol{y}}_i\right) + \left(1-\boldsymbol{y}_i\right)\log\left(1-\hat{\boldsymbol{y}}_i\right) \quad (6)$$

where $\boldsymbol{y}_i \in \mathbb{R}^K$ is the true label of cell line i to K drugs.

Adaptation phase

Domain adversarial training. For generalization of multi-label classifier $\boldsymbol{g}_a$ to the unlabeled target domain, the invariant representations learned by the extractor $\boldsymbol{f}_a$ should be tailored for the downstream DRP task. Domain adversarial training is a straightforward solution to fine-tune the extractor to produce target domain representations with a distribution closer to that of the source domain. In its procedure, samples from both domains need to align globally, without distinguishing diverse cancer categories. Intuitively, aligning category-unmatched samples would not reduce the domain gap and therefore brings no benefit for multi-label classification in the target domain. Instead, it is important to align representations within similar cancer categories while distinguishing between confusable ones. Along this line, we devise a novel domain adversarial strategy that aligns the two domains globally and cancer categories locally.

For the global one, we adopt a standard adversarial-based UDA, which fine-tunes the extractor $\boldsymbol{f}_e$ to learn domain-invariant features by fooling a domain discriminator $\boldsymbol{\varphi}$, a neural network designed to distinguish between source and target domain samples. Specifically, $\boldsymbol{\varphi}$ is optimized to minimize the discrimination loss $\mathcal{L}_{\text{dis}}$ (i.e., binary cross entropy)); meanwhile, $\boldsymbol{f}_e$ is optimized to maximize $\mathcal{L}_{\text{dis}}$ and minimize the multi-label loss $\mathcal{L}_{\text{cls}}$ with the purpose of balancing classifier $\boldsymbol{g}_a$ to adapt target domain. As a result, the global objective $\mathcal{L}_{\text{adv}}$ can be defined as follows:

$$\begin{aligned} \mathcal{L}_{\text{adv}} &= \underset{\boldsymbol{f}_e,\boldsymbol{\varphi}}{\operatorname{argmin}} \left\{\mathcal{L}_{\text{cls}} - \mathcal{L}_{\text{dis}}\right\} \\ \mathcal{L}_{\text{dis}} &= -\left(\mathbb{E}_{i\in\mathcal{D}^s}\left\{\log\left(\boldsymbol{\varphi}\left(\boldsymbol{f}_e(\boldsymbol{x}_i^s)\right)\right)\right\} + \right. \\ &\qquad \left. \mathbb{E}_{j\in\mathcal{D}^t}\left\{\log\left(1-\boldsymbol{\varphi}\left(\boldsymbol{f}_e(\boldsymbol{x}_j^t)\right)\right)\right\}\right) \end{aligned} \quad (7)$$

To achieve such 'mini-max' adversarial loss, gradient reversal layer (GRL) is used via multiplying the gradient from $\boldsymbol{\varphi}$

by a negative constant during the back-propagation to $\boldsymbol{f}_e$.

For the local one, we divide samples into different cancer categories based on clinical knowledge, and borrow the idea of contrastive learning to promote cross-domain representative consistency between same categories and discrepancy among different categories. Contrastive learning (Jaiswal et al. 2020) is a powerful self-supervised paradigm that brings an anchor (i.e., sample) closer to a positive/similar instance and away from many negative/dissimilar instances, by optimizing their mutual information (MI). However, it is unwise to yield positive/negative instances using regular data augmentation or sampling strategies, as they may not satisfy biological logic or fail to materialize in some mini-batches (explanations in **Appendix 1-1**). Hence, we consider constructing positive and negative instances through the prototypical network (Zhou et al. 2023).

Prototype is based on the concept that there exists an embedding in which samples cluster around a single representation for each category. First, we define each prototype $\boldsymbol{p}$ as the mean vector of samples belonging to its category:

$$\boldsymbol{p}_c^s = \frac{1}{|\mathcal{S}_c^s|} \sum_{i \in \mathcal{S}_c^s} \boldsymbol{h}_i^s, \quad \boldsymbol{p}_c^t = \frac{1}{|\mathcal{S}_c^t|} \sum_{j \in \mathcal{S}_c^t} \boldsymbol{h}_j^t \tag{8}$$

where $\mathcal{S}_c^s$/$\mathcal{S}_c^t$ represents the samples labeled with cancer category c in the source/target domain, $\boldsymbol{h}$ is the shared representation obtained during the pre-training phase.

Next, we designate each target domain sample $\boldsymbol{x}_i^t$ as the anchor, and source domain prototype $\boldsymbol{p}_i^s$ consistent with its category as the positive instance, while other target domain prototypes $\{\boldsymbol{p}_1^t, ..., \boldsymbol{p}_C^t\}$ are regarded as negative instances. Accordingly, the MI of positive pair (anchor and positive instance) should be maximized, and the MI of negative pairs (anchor and negative instances) should be minimized. At last, the local objective is to minimize the following InfoNCE-based (Parulekar et al. 2023) loss:

$$\mathcal{L}_{\text{inf}} = -\frac{1}{M} \sum_{i=1}^{M} \log \frac{\boldsymbol{\Gamma}\left(\boldsymbol{f}_e(\boldsymbol{x}_i^t), \boldsymbol{p}_i^s\right)}{\sum_{c=1}^{C} \boldsymbol{\Gamma}\left(\boldsymbol{f}_e(\boldsymbol{x}_i^t), \boldsymbol{p}_c^t\right)} \tag{9}$$

where C indicates the number of cancer categories, $\boldsymbol{\Gamma}$ is a similarity function measuring the MI score between two variables. Such a prototype design, built upon previously learned representations, could facilitate the construction of positive and negative instances for diverse cancer categories, to satisfy our category-divergence contrastive learning.

Transferring drug response prediction. The ultimate objective of our global-local domain adversarial training is summarized as below:

$$\mathcal{L} = \alpha \mathcal{L}_{\text{adv}} + (1 - \alpha) \mathcal{L}_{\text{inf}} \tag{10}$$

where α is a coefficient to balance the contributions of different losses. After fine-turning this adversarial objective for τ_b epochs, domain representation extractor $\boldsymbol{f}_e$ and multi-label classifier $\boldsymbol{g}_a$ can be isolated, and then directly applied to each genomic input of target domain to output the response probability of patient on multiple drugs: $\hat{\boldsymbol{y}}^t \in \mathbb{R}^K \leftarrow \boldsymbol{g}_a\left(\boldsymbol{f}_e(\boldsymbol{x}^t)\right)$. Hyperparameter settings and implementation details of TransDRP can be found in **Appendix 1-2**.

Experiments

Experimental setups

Dataset. In the pre-training phase, we collected 1,517 cell line samples (source domain) and 9,808 patient samples (target domain) from the CCLE (Barretina et al. 2012) and TCGA (Weinstein et al. 2013) databases for training the autoencoder. Following He et al. (2022)'s study, each cell line or patient sample was represented as its genomic feature (i.e., gene expression vector), and annotated with the cancer category. The labeled dataset used for pre-training the source domain's multi-label classifier was obtained from the GDSC database (Iorio et al. 2016), which contains IC_{50} values (labels) quantifying the drug efficacy in inhibiting cell line growth. Then, we binarized IC_{50} values to conduct classification tasks by using a z-score threshold of '0', where a drug response was deemed sensitivity at z-score(IC_{50})$<$0 and resistance otherwise. In the adaptation phase, 1,186 labeled patient samples, containing drug feedback records, were collected for testing. Each record details patient's medication progress, with complete and partial responses classified as sensitivity, and stable and progressive diseases classified as resistance. At last, 9 drugs (including 5-Fluorouracil, Cisplatin, Cyclophosphamide, Docetaxel, Doxorubicin, Etoposide, Gemcitabine, Temozolomide, and Paclitaxel) were screened for target domain testing, each with at least 90 labeled patient response records. Detailed dataset compilations are listed in **Appendix 2-1**.

Evaluation protocol. This work focused on examining the performance of DRP models in a zero-shot learning scenario, i.e., predicting drug responses in out-of-distribution data unseen during training. Initially, the 5-fold cross-validation (5-CV) was executed using labeled cell line data (source domain) and unlabeled patient data (target domain), to train the DRP model. Subsequently, the trained model was tested on the target domain, to predict drug response labels for patients. Notably, samples of the target domain were only utilized for unsupervised pre-training, and drug feedback records (labels of target domain) did not appear during the training phases. Four common metrics were utilized to measure the classification performance of models: area under the curve (AUC), area under the precision-recall curve (AUPR), accuracy (ACC), and f1-score (F1), with AUC and AUPR primarily displayed. Note that, Appendix is available in the **Code** link.

Baseline. We first evaluated TransDRP against several UDA methods for anticancer DRP, including three discrepancy-based methods (Celligner (Warren et al. 2021), Velodrome (Sharifi-Noghabi et al. 2021), scDEAL (Chen et al. 2022)) and three adversarial-based methods (AITL (Sharifi-Noghabi et al. 2020), CODE-AE (He et al. 2022), WISER (Shubham et al. 2024)). Moreover, we compared several advanced cell line-oriented methods that were not fine-tuned using the transfer learning strategy, including two convolutional neural network-based methods (tCNNs (Liu et al. 2019), DeepCDR (Liu et al. 2020b)), and two graph neural network-based methods (GraphCDR (Liu et al. 2022a), Bi-

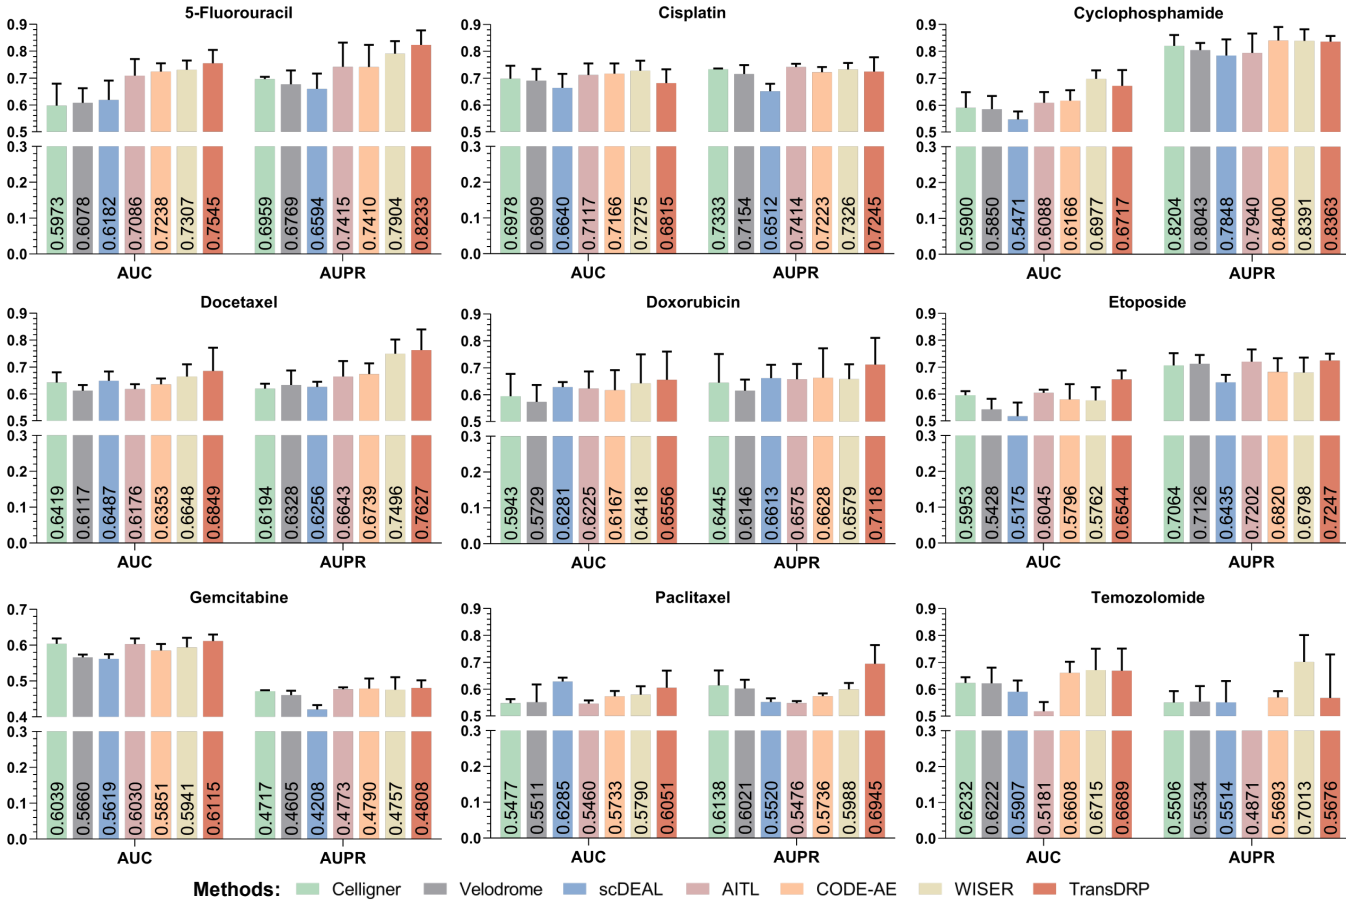

Figure 3: Performance (AUC and AUPR scores) comparison of all methods on the TCGA dataset for 9 clinical drugs.

GNN (Hostallero, Li, and Emad 2022)). All baseline methods were implemented using their publicly available source codes, with the best or default parameters applied. Detailed pipelines for each method are provided in **Appendix 2-2**.

Experimental results

Performance comparison. As depicted in Figure 3 and **Appendix 2-3**, TransDRP outperformed baselines with superior AUC and AUPR scores on 6 drugs (5-Fluorouracil, Docetaxel, Doxorubicin, Etoposide, Gemcitabine, and Paclitaxel), while maintaining competitive performance on the remaining 3 drugs (Cisplatin, Cyclophosphamide, and Temozolomide). All methods exhibited significant drug-specific performance variations. This phenomenon may arise from the imbalance ratio, i.e., the discrepancy between the number of sensitive and resistant samples across different drugs. In contrast, our global-local alignment demonstrates more stable and general performance across all drugs by accurately capturing domain-invariance knowledge. Regarding computational efficiency, we compared the running time of all methods at the same training epochs, as shown in **Appendix 2-4**. Obviously, TransDRP's multi-label design allows for the simultaneous prediction of multiple drugs (including but not limited to the 9 tested), thus consuming less running time compared to the baselines.

In addition to comparing with UDA baselines, we further tested cell line-oriented models when directly applied to TCGA patient data without transfer learning adjustment. As illustrated in **Appendix 2-5**, all cell line models performed suboptimally in predicting responses for 9 drug data, with AUC and AUPR scores declining by approximately 0.1 to 0.2 compared to UDAs. Of these results, cell line models approximate the performance of UDA methods for only one drug (Gemcitabine), but even the best-performing Bi-GNN falls behind our TransDRP. These findings confirmed the domain-shift between cell lines and patients, highlighting the strength of UDA strategy in bridging their DRP gap.

Ablation study. To investigate the necessity of each module in our model architecture, we conducted several comparisons between TransDRP with its variants:

- TransDRP (w/o AE) that removes the autoencoder in pre-training phase, without performing shared representation extraction across domains.
- TransDRP (w/o GD) that removes the GNN decoder and label graph for drug dependency learning, using only MLP to train the multi-label classifier.
- TransDRP (w/o ML) that removes the multi-label DRP

tasks, taking the response prediction of each drug as an individual task for separate training.

- TransDRP (w/o LA) that removes the category contrastive loss $\mathcal{L}_{\text{inf}}$ used for local domain alignment.
- TransDRP (w/o GA) that removes the generative adversarial loss $\mathcal{L}_{\text{adv}}$ used for global domain alignment.

Methods	AUC	AUPR	ACC	F1
TransDRP	0.6654	0.7029	0.6418	0.6287
TransDRP(w/o AE)	0.6473	0.6908	0.6174	0.6021
TransDRP(w/o GD)	0.6455	0.7001	0.6350	0.5930
TransDRP(w/o ML)	0.6602	0.6813	0.5986	0.5749
TransDRP(w/o LA)	0.6396	0.6871	0.6343	0.6017
TransDRP(w/o GA)	0.5944	0.6330	0.5450	0.5637

Table 1: Ablation results (average of 9 drugs)

Ablation results are shown in Table 1. When the autoencoder (w/o AE) was removed, the AUC and AUPR scores dropped from 0.6654 to 0.6473 and 0.7029 to 0.6908 respectively, implying the usefulness of our shared representation extraction for the UDA task. The result of variants (w/o GD) and (w/o ML) emphasized that using the multi-label GNN as a response classifier contributes to the downstream DRP transfer. Moreover, after removing the local or global alignment loss, the performance of their variants (w/o LA) and (w/o GA) were significantly degraded, highlighting that such a local-global combination strategy does facilitates DRP transfer. Overall, TransDRP with the above modules together delivered superior performance, and removing any modules will compromise its power.

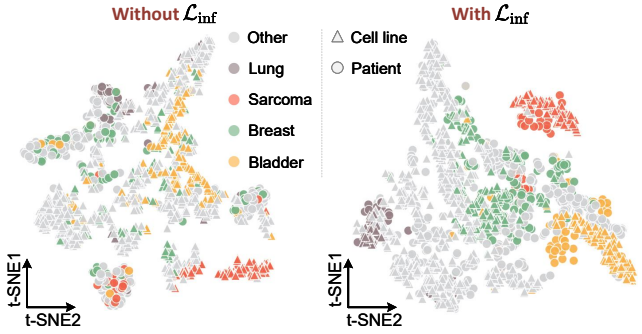

Figure 4: t-SNE visualization of representations produced by TransDRP with and without local loss $\mathcal{L}_{\text{inf}}$, where triangles/-circles indicate source/target domain samples and different colours indicate different cancer categories.

Visualization analysis. Since the ablation study emphasized the importance of our local alignment, we tended to investigate how the contrastive loss $\mathcal{L}_{\text{inf}}$ affects domain alignment in TransDRP. For this purpose, the shared representations from adaptation phase were visualized using t-SNE, with or without $\mathcal{L}_{\text{inf}}$ loss (Figure 4). Without $\mathcal{L}_{\text{inf}}$, two domains exhibited indistinguishable overlap compared to the original gene expression distribution (Figure 1), yet an inter-domain gap still persisted within cancer categories. With $\mathcal{L}_{\text{inf}}$, similar cancer categories between the two domains were closer, and different cancer categories remained distinct. This finding demonstrates that $\mathcal{L}_{\text{inf}}$ indeed promotes category-divergence domain alignment.

Novel drug screening capability. To verify the practical value of TransDRP in screening novel drugs, we performed a case study using samples with testicular and pancreatic cancer (derived from TCGA database), to predict their responses on seven unseen drugs: Tamoxifen, Afatinib, Pemetrexed, Fulvestrant, Methotrexate, Vinorelbine, and Vinblastine (as shown in Figure 5). Among 40 testicular cancer patients, Vinorelbine and Vinblastine exhibited higher therapeutic probabilities than other drugs, both of which are alkaloids approved for clinical treatment of metastatic testicular cancer (Barrales-Cureño 2015). Among 16 pancreatic cancer patients, Pemetrexed and Tamoxifen were identified as two drugs with the highest predicted probability. Pemetrexed has shown improved overall survival in advanced pancreatic cancer patients during a phase II trial (Oettle et al. 2006), and Tamoxifen has also been found to weaken the microenvironment of solid tumors in mice and may be repurposed for pancreatic cancer treatment (Cortes et al. 2019). These results reveal the great potential of TransDRP in discovering new clinically anticancer drugs.

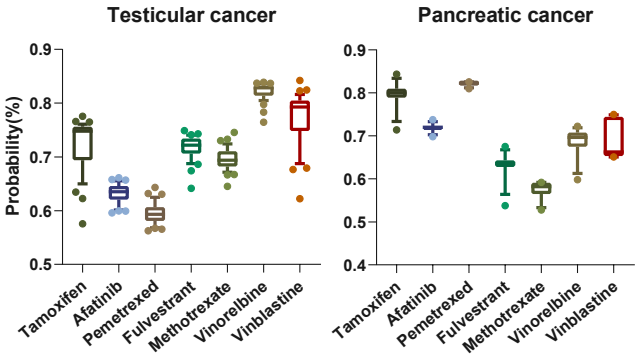

Figure 5: Response predictions of TransDRP for testicular and pancreatic cancer patients with seven unseen drugs.

Conclusion

In this work, we leverage bio-knowledge within the UDA framework to present TransDRP, a robust deep learning model for transferring DRP from pre-clinical cell lines (source domain) to clinical patients (target domain). On one hand, it fuses molecular similarity information and graph neural networks for multi-label (drug) response prediction, improving domain transfer efficiency. On the other hand, it incorporates a novel domain adversarial training with category-divergence alignment, improving domain transfer accuracy. In transfer experiments from the GDSC dataset (cell lines) to the TCGA dataset (patients), TransDRP surpasses other state-of-the-art methods in terms of prediction performance and computational efficiency. Subsequent clinical drug screening analyses further highlight TransDRP's excellent generalization capability in DRP tasks.

Acknowledgments

This work was supported by the Key Scientifc Research Proiect of Colleges and Universities of Henan Province in China (25A170003). The funders have no role in study design, data collection, data analysis, data interpretation, or writing of the manuscript.

References

Adam, G.; Rampášek, L.; Safikhani, Z.; Smirnov, P.; Haibe-Kains, B.; and Goldenberg, A. 2020. Machine learning approaches to drug response prediction: challenges and recent progress. *NPJ precision oncology*, 4(1): 19.

Bai, P.; Miljković, F.; John, B.; and Lu, H. 2023. Interpretable bilinear attention network with domain adaptation improves drug–target prediction. *Nature Machine Intelligence*, 5(2): 126–136.

Barrales-Cureño, H. J. 2015. Pharmacological applications and in vitro biotechnological production of anticancer alkaloids of Catharanthus roseus. *Biotecnol Apl*, 32(1): 1101–10.

Barretina, J.; Caponigro, G.; Stransky, N.; Venkatesan, K.; Margolin, A. A.; Kim, S.; Wilson, C. J.; Lehár, J.; Kryukov, G. V.; Sonkin, D.; et al. 2012. The Cancer Cell Line Encyclopedia enables predictive modelling of anticancer drug sensitivity. *Nature*, 483(7391): 603–607.

Borgwardt, K. M.; Gretton, A.; Rasch, M. J.; Kriegel, H.-P.; Schölkopf, B.; and Smola, A. J. 2006. Integrating structured biological data by kernel maximum mean discrepancy. *Bioinformatics*, 22(14): e49–e57.

Bousmalis, K.; Trigeorgis, G.; Silberman, N.; Krishnan, D.; and Erhan, D. 2016. Domain separation networks. *Advances in neural information processing systems*, 29.

Chen, J.; Wang, X.; Ma, A.; Wang, Q.-E.; Liu, B.; Li, L.; Xu, D.; and Ma, Q. 2022. Deep transfer learning of cancer drug responses by integrating bulk and single-cell RNA-seq data. *Nature Communications*, 13(1): 6494.

Chen, Z.-M.; Wei, X.-S.; Wang, P.; and Guo, Y. 2021. Learning graph convolutional networks for multi-label recognition and applications. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(6): 6969–6983.

Cortes, E.; Lachowski, D.; Robinson, B.; Sarper, M.; Teppo, J. S.; Thorpe, S. D.; Lieberthal, T. J.; Iwamoto, K.; Lee, D. A.; Okada-Hatakeyama, M.; et al. 2019. Tamoxifen mechanically reprograms the tumor microenvironment via HIF-1A and reduces cancer cell survival. *EMBO reports*, 20(1): e46557.

Cui, S.; Jin, X.; Wang, S.; He, Y.; and Huang, Q. 2020. Heuristic domain adaptation. *Advances in Neural Information Processing Systems*, 33: 7571–7583.

Firoozbakht, F.; and Schwikowski, B. 2022. An overview of machine learning methods for monotherapy drug response prediction. *Briefings in bioinformatics*, 23(1): bbab408.

Ganin, Y.; and Lempitsky, V. 2015. Unsupervised domain adaptation by backpropagation. In *International conference on machine learning*, 1180–1189. PMLR.

Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2020. Generative adversarial networks. *Communications of the ACM*, 63(11): 139–144.

He, D.; Liu, Q.; Wu, Y.; and Xie, L. 2022. A context-aware deconfounding autoencoder for robust prediction of personalized clinical drug response from cell-line compound screening. *Nature Machine Intelligence*, 4(10): 879–892.

Hostallero, D. E.; Li, Y.; and Emad, A. 2022. Looking at the BiG picture: incorporating bipartite graphs in drug response prediction. *Bioinformatics*, 38(14): 3609–3620.

Hostallero, D. E.; Wei, L.; Wang, L.; Cairns, J.; and Emad, A. 2023. Preclinical-to-clinical anti-cancer drug response prediction and biomarker identification using TINDL. *Genomics, Proteomics & Bioinformatics*, 21(3): 535–550.

Huynh, D.; and Elhamifar, E. 2020. Interactive multi-label cnn learning with partial labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9423–9432.

Iorio, F.; Knijnenburg, T. A.; Vis, D. J.; Bignell, G. R.; Menden, M. P.; Schubert, M.; Aben, N.; Gonçalves, E.; Barthorpe, S.; Lightfoot, H.; et al. 2016. A landscape of pharmacogenomic interactions in cancer. *Cell*, 166(3): 740–754.

Jaiswal, A.; Babu, A. R.; Zadeh, M. Z.; Banerjee, D.; and Makedon, F. 2020. A survey on contrastive self-supervised learning. *Technologies*, 9(1): 2.

Lind, A. P.; and Anderson, P. C. 2019. Predicting drug activity against cancer cells by random forest models based on minimal genomic information and chemical properties. *PloS one*, 14(7): e0219774.

Liu, J.; Shen, Z.; He, Y.; Zhang, X.; Xu, R.; Yu, H.; and Cui, P. 2020a. Towards out-of-distribution generalization: A survey. *arXiv preprint arXiv:2108.13624*.

Liu, P.; Li, H.; Li, S.; and Leung, K.-S. 2019. Improving prediction of phenotypic drug response on cancer cell lines using deep convolutional network. *BMC bioinformatics*, 20: 1–14.

Liu, Q.; Hu, Z.; Jiang, R.; and Zhou, M. 2020b. DeepCDR: a hybrid graph convolutional network for predicting cancer drug response. *Bioinformatics*, 36(Supplement_2): i911–i918.

Liu, X.; Song, C.; Huang, F.; Fu, H.; and Zhang, W. 2022a. GraphCDR: a graph neural network method with contrastive learning for cancer drug response prediction. *Briefings in Bioinformatics*, 23(1): bbab457.

Liu, X.; Yoo, C.; Xing, F.; Oh, H.; El Fakhri, G.; Kang, J.-W.; Woo, J.; et al. 2022b. Deep unsupervised domain adaptation: A review of recent advances and perspectives. *APSIPA Transactions on Signal and Information Processing*, 11(1).

Liu, X.; and Zhang, W. 2023. A subcomponent-guided deep learning method for interpretable cancer drug response prediction. *PLOS Computational Biology*, 19(8): e1011382.

Liu, Y.; Li, R.; Luo, J.; and Zhang, Z. 2022c. Inferring RNA-binding protein target preferences using adversarial domain adaptation. *PLoS computational biology*, 18(2): e1009863.

Ma, J.; Fong, S. H.; Luo, Y.; Bakkenist, C. J.; Shen, J. P.; Mourragui, S.; Wessels, L. F.; Hafner, M.; Sharan, R.; Peng, J.; et al. 2021. Few-shot learning creates predictive models of drug response that translate from high-throughput screens to individual patients. *Nature Cancer*, 2(2): 233–244.

Mourragui, S.; Loog, M.; Van De Wiel, M. A.; Reinders, M. J.; and Wessels, L. F. 2019. PRECISE: a domain adaptation approach to transfer predictors of drug response from pre-clinical models to tumors. *Bioinformatics*, 35(14): i510–i519.

Oettle, H.; Richards, D.; Ramanathan, R.; van Laethem, J.; Peeters, M.; Fuchs, M.; Zimmermann, A.; John, W.; Von Hoff, D.; Arning, M.; et al. 2006. A phase III trial of pemetrexed plus gemcitabine versus gemcitabine in patients with unresectable or metastatic pancreatic cancer. Ann Oncol 2005; 16: 1639–1645. *Annals of Oncology*, 17(3): 535.

Parulekar, A.; Collins, L.; Shanmugam, K.; Mokhtari, A.; and Shakkottai, S. 2023. Infonce loss provably learns cluster-preserving representations. In *The Thirty Sixth Annual Conference on Learning Theory*, 1914–1961. PMLR.

Peres da Silva, R.; Suphavilai, C.; Nagarajan, N.; and Nagarajan, N. 2021. TUGDA: task uncertainty guided domain adaptation for robust generalization of cancer drug response prediction from in vitro to in vivo settings. *Bioinformatics*, 37(Supplement_1): i76–i83.

Rogers, D.; and Hahn, M. 2010. Extended-connectivity fingerprints. *Journal of chemical information and modeling*, 50(5): 742–754.

Seyhan, A. A. 2019. Lost in translation: the valley of death across preclinical and clinical divide–identification of problems and overcoming obstacles. *Translational Medicine Communications*, 4(1): 1–19.

Sharifi-Noghabi, H.; Harjandi, P. A.; Zolotareva, O.; Collins, C. C.; and Ester, M. 2021. Out-of-distribution generalization from labelled and unlabelled gene expression data for drug response prediction. *Nature Machine Intelligence*, 3(11): 962–972.

Sharifi-Noghabi, H.; Peng, S.; Zolotareva, O.; Collins, C. C.; and Ester, M. 2020. AITL: adversarial inductive transfer learning with input and output space adaptation for pharmacogenomics. *Bioinformatics*, 36(Supplement_1): i380–i388.

Shubham, K.; Jayagopal, A.; Danish, S. M.; AP, P.; and Rajan, V. 2024. WISER: Weak supervISion and supErvised Representation learning to improve drug response prediction in cancer. *arXiv preprint arXiv:2405.04078*.

Stanfield, Z.; Coşkun, M.; and Koyutürk, M. 2017. Drug response prediction as a link prediction problem. *Scientific reports*, 7(1): 40321.

Sun, B.; and Saenko, K. 2017. Correlation alignment for unsupervised domain adaptation. *Domain adaptation in computer vision applications*, 153–171.

Veličković, P.; Cucurull, G.; Casanova, A.; Romero, A.; Lio, P.; and Bengio, Y. 2017. Graph attention networks. *arXiv preprint arXiv:1710.10903*.

Wang, L.; Li, X.; Zhang, L.; and Gao, Q. 2017. Improved anticancer drug response prediction in cell lines using matrix factorization with similarity regularization. *BMC cancer*, 17: 1–12.

Warren, A.; Chen, Y.; Jones, A.; Shibue, T.; Hahn, W.; Boehm, J.; Vazquez, F.; Tsherniak, A.; and McFarland, J. 2021. Global computational alignment of tumor and cell line transcriptional profiles. *Nature Communications*, 12(1): 22–22.

Weinstein, J. N.; Collisson, E. A.; Mills, G. B.; Shaw, K. R.; Ozenberger, B. A.; Ellrott, K.; Shmulevich, I.; Sander, C.; and Stuart, J. M. 2013. The cancer genome atlas pan-cancer analysis project. *Nature genetics*, 45(10): 1113–1120.

Yang, J.; Li, A.; Li, Y.; Guo, X.; and Wang, M. 2019. A novel approach for drug response prediction in cancer cell lines via network representation learning. *Bioinformatics*, 35(9): 1527–1535.

Yuan, B.; Zhao, D.; Shao, S.; Yuan, Z.; and Wang, C. 2022. Birds of a feather flock together: Category-divergence guidance for domain adaptive segmentation. *IEEE Transactions on Image Processing*, 31: 2878–2892.

Zhou, F.; Wang, P.; Zhang, L.; Wei, W.; and Zhang, Y. 2023. Revisiting prototypical network for cross domain few-shot learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 20061–20070.