

Zero-shot Learning for Preclinical Drug Screening

Kun Li¹, Weiwei Liu¹, Yong Luo¹, Xiantao Cai¹, Jia Wu² and Wenbin Hu^{1*}

¹School of Computer Science, Wuhan University, Wuhan, China

²Department of Computing, Macquarie University, Sydney, Australia

{li_kun, luoyong, caixiantao, hwb}@whu.edu.com, liuweimei863@gmail.com, jia.wu@mq.edu.au

Abstract

Conventional deep learning methods typically employ supervised learning for drug response prediction (DRP). This entails dependence on labeled response data from drugs for model training. However, practical applications in the preclinical drug screening phase demand that DRP models predict responses for novel compounds, often with unknown drug responses. This presents a challenge, rendering supervised deep learning methods unsuitable for such scenarios. In this paper, we propose a zero-shot learning solution for the DRP task in preclinical drug screening. Specifically, we propose a Multi-branch Multi-Source Domain Adaptation Test Enhancement Plug-in, called MSDA. MSDA can be seamlessly integrated with conventional DRP methods, learning invariant features from the prior response data of similar drugs to enhance real-time predictions of unlabeled compounds. The results of experiments on two large drug response datasets showed that MSDA efficiently predicts drug responses for novel compounds, leading to a general performance improvement of 5-10% in the preclinical drug screening phase. The significance of this solution resides in its potential to accelerate the drug discovery process, improve drug candidate assessment, and facilitate the success of drug discovery. The code is available at <https://github.com/DrugD/MSDA>.

1 Introduction

Improving the efficiency of preclinical drug candidate screening is a long-standing core challenge in the field of drug discovery [Li *et al.*, 2022]. It takes an average of 10 to 15 years and more than \$2 billion for a new drug to reach the pharmacy shelf [Siqueira-Neto *et al.*, 2023; Pandey *et al.*, 2022; Berdigaliyev and Aljofan, 2020]. Historically, natural products have been the main source of new drug entities; in recent years, however, there has been a shift towards high-throughput screening (HTS) techniques [Sadybekov *et al.*,

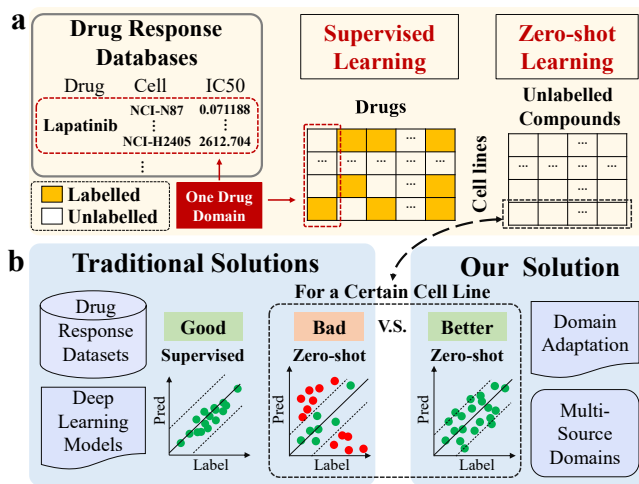


Figure 1: The difference between supervised learning and zero-shot learning for preclinical drug screening. (a) Schematic on the definitions of supervised learning and zero-shot learning in the DRP task. (b) The comparison of the effectiveness between the traditional solutions and our solution in preclinical trials. The traditional solutions perform well in supervised learning but poorly in zero-shot learning environments aimed at practical applications.

2022; Engels and Venkatarangan, 2001]. HTS drug screening methods can screen chemical libraries to identify the most promising compounds that have the desired effect on specific biological targets. Notably, the conventional libraries employed in HTS and virtual ligand screening [Irwin and Shoichet, 2016] (VLS) are constrained to less than 10 million accessible compounds, representing a mere fraction of the vast existing chemical space encompassing an estimated 10^{20} to 10^{60} novel compounds [Ertl, 2003]. This limitation of standard HTS and VLS decelerates the pace of drug discovery [Stein *et al.*, 2020; Lyu *et al.*, 2019], frequently yielding compounds with moderate affinity, limited selectivity, and initial hits displaying absorption, distribution, metabolism, excretion, and toxicity (ADMET) factors profiles [Wu *et al.*, 2022] that necessitate extensive testing.

In recent years, driven by the rapid advancement of AI technology, virtual screening based on deep learning (DL) is poised to emerge as a swifter and more cost-effective approach to drug discovery [Sadybekov and Katritch, 2023;

*Corresponding author

Drug	Condition	Data Division		GraphDRP		GratransDRP	
		Train / Test	Origin	+MSDA	Origin	+MSDA	
5-Fluorouracil	s	90.28% / 9.72%	0.91	-	0.89	-	
	z	0% / 100%	0.47	0.65 (40%)	0.58	0.65 (12%)	
Pelitinib	s	88.62% / 11.38%	0.89	-	0.89	-	
	z	0% / 100%	0.34	0.58 (73%)	0.45	0.57 (28%)	
Alectinib	s	89.58% / 10.42%	0.70	-	0.86	-	
	z	0% / 100%	0.14	0.42 (196%)	0.26	0.41 (60%)	
Global	s	88.89% / 11.11%	0.93	-	0.93	-	
	z	0% / 100%	0.44	0.48 (11%)	0.48	0.51 (5%)	

Table 1: Examples of performance comparison for drug response prediction under the conditions of supervised learning (s) and zero-shot learning (z). Drug response data for one drug with one cell line is recorded as one sample of this drug. The sample numbers of these drugs in the training dataset are set to 0 to simulate zero-shot learning. The evaluation metric is the *Pearson correlation coefficient*, the closer to 1 the better. The **Global average** represents the average performance of all types of drugs in the dataset.

Stärk *et al.*, 2022]. Exhaustive catalogs of somatic mutations in various cancer types have been created [Forbes *et al.*, 2017; Lawrence *et al.*, 2014] and major oncogenic mutations identified [Stratton *et al.*, 2009]. As a result, the establishment of cancer drug sensitivity databases, such as the Genomics of Drug Sensitivity in Cancer (GDSC) [Yang *et al.*, 2012], has made a large amount of drug response data newly available to researchers. Many approaches leverage these databases to design and validate a diverse array of models that aim to elucidate connections between genomic data and drug responsiveness in the context of drug discovery [Chu *et al.*, 2023; Nguyen *et al.*, 2022].

Conventional drug response prediction (DRP) methods typically rely on supervised learning using labeled response data from the drugs, as depicted in Figure 1(a). Predicting the responses of unlabeled novel compounds constitutes a zero-shot learning challenge [Pourpanah *et al.*, 2023] since these novel compounds must not have been encountered during training. Furthermore, the drug responses of novel compounds in practical preclinical drug screening applications are unknown. The distribution of data between known drug response data and that of novel compounds is often inconsistent, rendering supervised learning-based drug response prediction methods ineffective for novel compounds [Nguyen *et al.*, 2022], as illustrated in Figure 1(b).

To underscore the shortcomings of supervised learning methods in practical preclinical drug screening, we test two state-of-the-art (SOTA) DRP methods (GraphDRP [Nguyen *et al.*, 2022], GratransDRP [Chu *et al.*, 2023]) on three drugs (5-Fluorouracil, Pelitinib, Alectinib), as shown in Table 1. These methods typically achieve correlation metrics exceeding 90% in a supervised learning experimental setup but exhibit a significant drop to only 20-40% correlation in a zero-shot learning scenario. The experimental results indicate that these methods are unable to accurately predict novel compounds in the practical application of preclinical drug screening. How, then, can the drug responses of previously unseen novel compounds be effectively predicted? This zero-shot learning problem has become the focus of DRP tasks in preclinical drug screening.

As shown in Figure 1(a), each drug can interact with numerous known cell lines, yielding corresponding response data. We define the collection of response data for a single drug across multiple cell lines as a single drug domain. The response data for one drug interacting with one cell line is recorded as one sample for that drug. DA aims to maximize the performance of an unknown domain (the **target domain**) by leveraging knowledge from known domains (**source domains**). This approach aligns well with the requirements of zero-shot learning in the context of the DRP task. This is because DA is an adaptive approach, proposed with the assumption that it allows the model access to the samples in the target domain [Bai *et al.*, 2023]. This aligns with the conditions of zero-shot learning during the testing phase of the DRP task, wherein the DRP model must handle novel compounds without access to any prior samples.

In this paper, we propose a zero-shot learning solution for the DRP task in preclinical drug screening. Within this solution, we present the Multi-branch Multi-Source Domain Adaptation Test Enhancement Plug-in (MSDA), designed to enhance the effectiveness of DRP model predictions for novel compounds. During the inference phase, the MSDA identifies several drug domains as the source domains with the strongest correlation to the target domain. It guides the pre-trained model to acquire invariant features from these domains, facilitating adaptation to the target domain. The MSDA consists of two modules. The first module, a multi-source domain selector, employs the Wasserstein distance metric [Panaretos and Zemel, 2019] on drug features to identify the most relevant drug domains from a large training dataset, treating them as multi-source domains. Unlike conventional multi-source domain methods, the number of source domains in the MSDA is not the regular two or three, but several tens or hundreds. The second module is a multi-branch drug domain adaptor. This module comprises two distinct prediction branches: the first is the prediction branch of the pre-training model itself, and the second is the target domain adaptation branch responsible for transferring multi-source domains to the target domain. The latter aligns cell line types from various domains before computing the Maximum Mean Discrepancy (MMD) [Li *et al.*, 2020].

To showcase the substantial performance improvements attained with the MSDA, we conduct comparative studies with other SOTA methods on the GDSCv2 and CellMiner [Reinhold *et al.*, 2012] datasets. In the inference phase, the MSDA provides average improvements of 26.1%, 15.8%, 16.7%, 9.3%, 9.8% on GDSCv2 and 26.2%, 15.6%, 1.5%, 10.2%, 11.8% on CellMiner across four metrics for each of the five DRP methods, respectively. We further perform ablation studies on various aspects of the MSDA, including the strategy for the selection of source domains and the design of the target domain adaptation branch. Additionally, we conduct a series of hyperparameter experiments to investigate the influence of expanding the number of domain adaptation branches on performance. These experiments affirm that the MSDA plug-in in our solution is plug-and-play, highly versatile, significantly enhances generalization, and can potentially facilitate higher performance when given adequate computational resources. The significance of this zero-shot learning solution

resides in its capacity to improve drug candidate evaluation and facilitate successful drug discovery.

2 Related Work

There are many research areas related to zero-shot learning, such as domain generalization, meta-learning, transfer learning, covariate shifting, and so on. Recently, *Domain Adaptation* (DA) has gained popularity in the machine learning field due to its demonstrated effectiveness in enhancing model prediction performance, especially when dealing with unlabeled test samples. This applicability extends to domains like computer vision [Zhou *et al.*, 2023], as well as natural language processing [Han *et al.*, 2023]. Nevertheless, there are still challenges associated with effectively enhancing the zero-shot learning capabilities of DRP models for preclinical drug screening through the utilization of DA methods. These challenges include:

1. *Insufficient theoretical research on domain adaptation methods for regression tasks.* Numerous effective DA methods have been devised for classification tasks [Fang *et al.*, 2022; Liang *et al.*, 2021]. To the best of our knowledge, however, there is a paucity of methods specifically tailored to regression tasks [Ou *et al.*, 2023; Yang *et al.*, 2022]. One perspective [Jiang *et al.*, 2021] posits that domain alignment may widen the margins between classes in the target domain, thereby enhancing model generalization. However, it is essential to note that the regression space is typically continuous; this can be contrasted with the classification space, in which clear decision boundaries exist. The other perspective [Chen *et al.*, 2021] asserts that classification is robust to feature scaling but regression is not, and that aligning the distributions of deep representations would alter the feature scale and impede domain adaptation regression. To summarize, it can be concluded that domain adaptation methods pose greater challenges when applied to regression tasks compared to classification tasks.
2. *Numerous open-source domains, not limited source domains.* Considering a single drug domain as the source domain for data mining results in a significant decrease in model generalization accuracy due to data availability limitations and the demands of practical application [Liu *et al.*, 2018]. Each drug’s response data should be treated as an independent drug domain, which results in an exceptionally large number of drug domains. Furthermore, the conventional single-domain adaptation method will fail when confronted with substantial distribution differences between the source and target domains [Ben-David *et al.*, 2010]. Therefore, the efficient construction of source domains from a pool of over 20,000 known drug domains presents a challenge due to the presence of inter-domain shifts in these source domains [Peng *et al.*, 2019].
3. *Complex distribution patterns exist between the source and target domains.* Conventional DA methods require the alignment of only one type of input (e.g., feature maps of images [Jaritz *et al.*, 2022], sentence embeddings [Wang *et al.*, 2022]). However, the input data for

the DRP task comprises two distinct types: drugs and cell lines. The types of cell lines in the source domains may not perfectly correspond to those in the target domain [He *et al.*, 2022]. Additionally, the distributions of combined features from different drugs and cell lines are inconsistent. Hence, the effective alignment of complex source and target domains is a challenging task.

3 Methods

3.1 Problem Definition

In the DRP task, the input data is the drug d and the cell line c , and the output is the half maximal inhibitory concentration (IC_{50} scores) of this drug on one cell line. A drug domain \mathcal{D} is defined as a set of response data for a specific drug d , across multiple cell lines, denoted as $\{c_1, c_2, \dots, c_n\}$. In Figure 1(a), one column of the matrix corresponds to a drug domain \mathcal{D} . The drug response datasets are divided into drug domains based on drug types. These drug domains are divided into a source domain (defined as S) and a target domain (defined as T) according to the drug domain as the smallest unit, regarded as the training set and the test set, respectively. The source domain S is denoted as $\{(\mathcal{D}_i, Y_i)\}_{i=1}^N$, and the target domain T is denoted as $\{(\mathcal{D}_j, Y_j)\}_{j=1}^M$. The set $\{(d_i, c_j)\}_{j=1}^{|\mathcal{D}_i|}$ denotes the inputs of the drug domain \mathcal{D}_i , and $Y_i = \{y_j\}_{j=1}^{|\mathcal{D}_i|}$ is the corresponding IC_{50} scores. N is the size of the source domain, while M is the size of the target domain. In general, M is a finite number in our experimental setting; in the real world, M would be enormous.

In the field of drug discovery, the core requirement of the DRP task is as follows: the model trained on S can achieve high accuracy on T to verify that the model already has high generalization ability, which ultimately serves the practical application of drug discovery.

3.2 Drug and Cell Line Feature Extractors

The drug feature representation branch and the cell line gene feature representation branch share weights derived from pre-trained DRP models, which are not incorporated into gradient calculations. This paper does not provide a detailed exposition of the computational methods employed in these two data representation branches. Broadly, we categorize feature extraction techniques for drugs and cell lines into Convolutional Neural Networks (CNNs), Multi-Layer Perceptrons (MLPs), Graph Neural Networks (GNNs), Transformers, and Transformer-based GNNs.

The input types for the drug extractor are the SMILES sequence or molecular graph. We uniformly describe the process of branching feature representation of a drug d_i as follows:

$$\mathbb{D}_i = \Phi_{drug}(d_i), \quad (1)$$

where, Φ_{drug} denotes the drug extractor with shared weights, while \mathbb{D}_i is the representation of the drug d_i . Similarly, c_j is the input of the cell line extractor which can be described as follows:

$$\mathbb{C}_j = \Phi_{cell}(c_j), \quad (2)$$

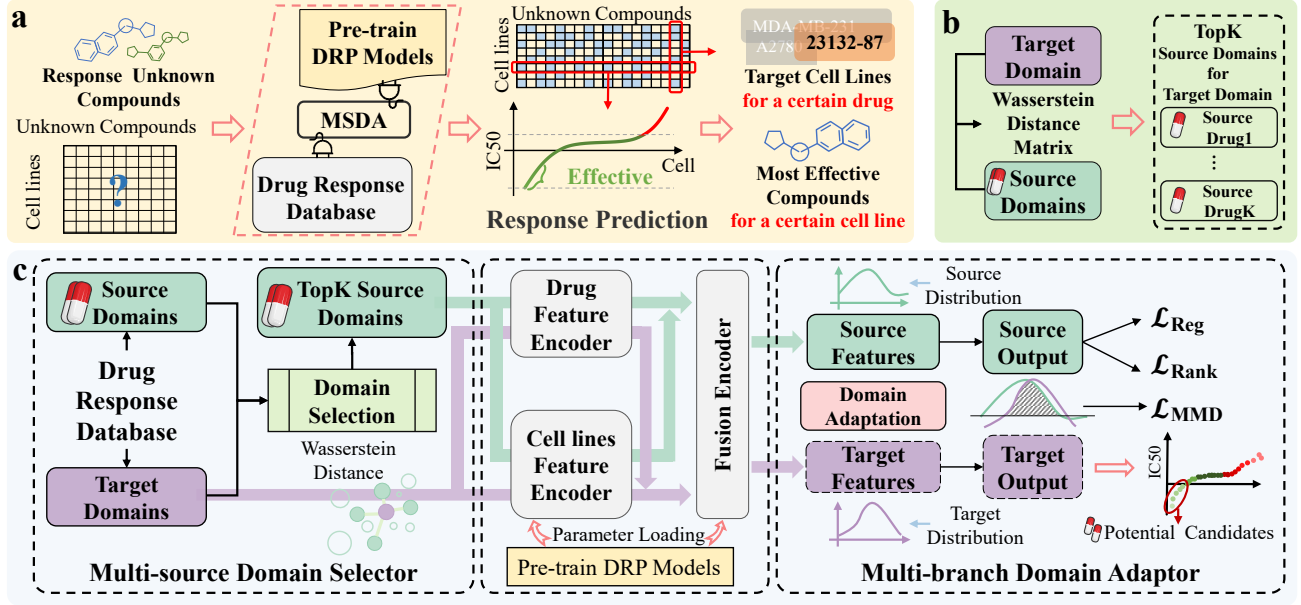


Figure 2: **Overview of the our solution.** (a)The main application of the MSDA as a plug-in for DRP tasks. (b)A flow illustration of the domain selector, where the input is a target domain and the output is the top K drug domains from the source domain with the highest similarity to the target domain. (c)The framework of the MSDA. The input is the source domains and target domains, the output is the drug response prediction of the target domains.

where, Φ_{cell} denotes the cell line extractor with shared weights, while \mathbb{C}_j is the representation of the cell line c_j , $j \in [1, |\mathcal{D}_i|]$.

3.3 Multi-source Domain Selector

Due to the vast number of drugs with known responses, utilizing the complete set of drug domains as source domains when implementing the MSDA for domain adaptation presents significant computational and time-related challenges. Therefore, a practical approach to mitigate computational resource constraints and enhance inference speed involves the selection of partially effective drug domains as source domains.

The structure of the domain selector is illustrated in Figure 2(b). In the initial stage of the MSDA, we design the multi-source domain selector to select multiple drug domains that are similar to the target domain from the training dataset; these are defined as the source domains. The set of drugs in the target domain T is denoted by $D^t = \{d_i\}_{i=1}^M$. Next, for a drug d_i from D^t , the set of drugs $D_i^{s \rightarrow t}$ from S are selected as shown below:

$$D_i^{s \rightarrow t} = \text{Top}_K(W_{ij}), \quad \text{where } j \in [1, N], \quad (3)$$

where $\text{Top}_K(\cdot)$ denotes the operation of sorting from smallest to largest order and fetching the first K elements. The Wasserstein distance W_{ij} is a distance function defined between probability distributions on a given metric space, and its 1-order form is formulated as follows:

$$W_{ij} = W(\mathbb{D}_j^s, \mathbb{D}_i^t) = \inf_{\gamma \in \Gamma} \mathbb{E}_{(x,y) \sim \gamma} [\|x - y\|], \quad (4)$$

where $\Gamma = \Pi(\mathbb{D}_j^s, \mathbb{D}_i^t)$ denotes the set of all joint distributions $\gamma(x, y)$ whose marginals are respectively $\mathbb{D}_j^s = \Phi_{drug}(d_j)$

and $\mathbb{D}_i^t = \Phi_{drug}(d_i)$; here, $\Phi_{drug}(\cdot)$ is the shared drug feature extractor referred to in Section 3.2. The distributions of various drug features exhibit significant dissimilarity, often with minimal overlap.

3.4 Multi-branch Drug Domain Adaptor

The multi-branch drug domain adaptor includes three functions: the fusion of drugs and cell line features, the adaptation from the multi-source domains to the target domain, and the predictions of drug responses in the target domain.

The inputs of the drug and cell line feature fusion branches are the outputs of the shared drug feature extractor Φ_{drug} and the shared cell line feature extractor Φ_{cell} , respectively (refer to Section 3.2 for details). Specifically, this module has a fusion branch for multi-source domains and several fusion branches for the target domain. Denote the fusion branch for the target domain as the target domain adaptation branch. The general fusion branch can be expressed as follows:

$$Y_{s \rightarrow t}^{pred} = \Phi_{fusion}(\mathbb{F}^{s \rightarrow t}), \quad (5)$$

where, $\mathbb{F}^{s \rightarrow t} = [\mathbb{D}^{s \rightarrow t}, \mathbb{C}^{s \rightarrow t}]$, and Φ_{fusion} denotes the fusion branch that can be replaced by different methods. The fusion branch of the multi-source domains uses pre-trained model parameters and is not involved in the gradient computation.

For example, in the case of the target domain D_i^t , each target domain fusion branch is constrained by three conditions, which are regression loss and ranking loss in multi-source domains $D_i^{s \rightarrow t}$ and feature-invariant consistency based on the MMD distance between $D_i^{s \rightarrow t}$ and D_i^t . The objects of the feature-invariant constraint are $FC(\mathbb{F}_i^{s \rightarrow t})$ and $FC(\mathbb{F}_i^t)$,

Methods		GDSCv2				CellMiner			
		PCC ↑	RMSE ↓	SPC ↑	Rank ↓	PCC ↑	RMSE ↓	SPC ↑	Rank ↓
tCNNs [Liu <i>et al.</i> , 2019]	Original	0.2073	0.0421	0.2032	0.0023	0.3320	0.00331	0.3627	0.0212
	+MSDA	0.3709	0.0080	0.3648	0.0054	0.3309	0.00146	0.3604	0.0106
	(Improv.)	79.0%	80.9%	79.53%	-134.78%	-0.3%	55.9%	-0.63%	50.00%
DeepCDR [Liu <i>et al.</i> , 2020]	Original	0.4442	0.0048	0.4391	0.0358	0.3593	0.00028	0.3766	0.0064
	+MSDA	0.4946	0.0039	0.4954	0.0281	0.3752	0.00024	0.3936	0.0045
	(Improv.)	11.3%	17.9%	12.82%	21.51%	4.4%	14.3%	14.30%	29.69%
GraphDRP [Nguyen <i>et al.</i> , 2022]	Original	0.4402	0.0056	0.4447	0.0337	0.4393	0.00027	0.4616	0.0054
	+MSDA	0.4898	0.0042	0.4888	0.0263	0.4398	0.00027	0.4624	0.0051
	(Improv.)	11.3%	24.0%	9.92%	21.96%	0.1%	0.2%	0.20%	5.56%
GratransDRP [Chu <i>et al.</i> , 2023]	Original	0.4848	0.0050	0.4851	0.0241	0.4324	0.00026	0.4562	0.0058
	+MSDA	0.5103	0.0039	0.5073	0.0224	0.4380	0.00025	0.4611	0.0042
	(Improv.)	5.3%	20.6%	4.58%	7.05%	1.3%	6.1%	6.10%	27.59%
TransEDRP [Li and Hu, 2022]	Original	0.5060	0.0052	0.5039	0.0295	0.4380	0.00025	0.4578	0.0055
	+MSDA	0.5316	0.0044	0.5300	0.0254	0.4422	0.00021	0.4608	0.0038
	(Improv.)	5.1%	15.2%	5.2%	13.9%	1.0%	14.7%	0.7%	30.8%

Table 2: Overall experiment. The table shows the model performance of five representative DRP methods on both the GDSCv2 and CellMiner datasets before and after test-time domain adaptation with the MSDA. The RMSE and Rank metrics are better when they are closer to 0, while the PCC and SPC metrics, which indicate the correlation between the predicted results and the true results, are better when they are closer to 1. Positive improvements are highlighted in red, while negative improvements are highlighted in green.

which are activated by the fully connected layer $FC(\cdot)$ once. The formulas for these constraints are as follows:

$$\mathcal{L}_{reg} = \sum_{i=0}^{|D_i^{s \rightarrow t}|} \text{MSELoss}(Y_i^{pred}, Y_i), \quad (6)$$

where, $\text{MSELoss}(\cdot)$ denotes the L2 norm loss. The label value y_i is subtracted from the model output (estimate) $f(x_i)$, after which the square is calculated to obtain the L2 norm loss, which is expressed as follows:

$$\text{MSELoss}(x, y) = \frac{1}{n} \sum_{i=1}^n (y_i - f(x_i))^2, \quad (7)$$

The ranking loss \mathcal{L}_{rank} is computed using MarginRankingLoss (MRLoss). For data (x_1, x_2, r) containing N samples, x_1 and x_2 are the two inputs given to be ranked, and r represents the true ranking labels. The ranking loss \mathcal{L}_{rank} is computed as shown below:

$$\mathcal{L}_{rank} = \sum_{i=0}^{|D_i^{s \rightarrow t}|} \text{MRLoss}(Y_i^{pred}, Y_i), \quad (8)$$

$$\text{MRLoss} = \max(0, -r(x_1 - x_2) + \text{margin}), \quad (9)$$

where margin denotes the margin value. If this value is larger, it means that the expectation x_1 is further away from x_2 (i.e. the margin is larger). The MMD is one of the most widely used loss functions in transfer learning, especially in domain adaptation, and is mainly used to measure the distance between two different but related distributions. The distance between two distributions is defined as follows:

$$\text{MMD}(X, Y) = \left\| \frac{\sum_{i=1}^n \phi(x_i)}{n} - \frac{\sum_{j=1}^m \phi(y_j)}{m} \right\|_H^2, \quad (10)$$

where, $\phi(\cdot)$ denotes a mapping from the original space to Hilbert space H . Hilbert space is the extension of Euclidean

space that is no longer restricted to the finite-dimensional situation. The dimensions of the drug and cell line features encoded by the different methods are different. The fact that the calculation of the distance between two distributions is not limited by the dimension of the sampled features is one of the keys to the functioning of the MSDA as a general plug-in. The distribution distances of invariant features from between multi-source domains $D_i^{s \rightarrow t}$ and the target domain D_i^t are constrained by MMD distances as follows:

$$\mathcal{L}_{mmd} = \sum_{i=0}^{|D_i^{s \rightarrow t}|} \text{MMD}(\varphi_{\text{align}}(\mathbb{F}_i^{s \rightarrow t}, \mathbb{F}_i^t)), \quad (11)$$

where $\varphi_{\text{align}}(\cdot)$ denotes the operation in which the feature vectors of two distributions are first fused through a fully connected layer, after which a union set is taken according to the cell line types, and eventually, the features are aligned according to their cell line types.

Finally, the overall loss of each target domain adaptation branch is obtained by weighted summation, as follows:

$$\mathcal{L} = \mathcal{L}_{reg} + \alpha \mathcal{L}_{rank} + \beta \mathcal{L}_{mmd}. \quad (12)$$

where, $\alpha \in [0, 1]$ and $\beta \in [0, 1]$ are the weights of the loss function. The computation of the loss of each target domain adaptation branch, the back-propagation, and the updating of the gradient are independent and serial.

4 Experiments

4.1 Evaluation Strategies and Metrics

We evaluate the performance of the MSDA plug-in on two publicly available datasets: GDSCv2 [Yang *et al.*, 2012] and CellMiner [Reinhold *et al.*, 2012]. The drug domain is then randomly partitioned into source and target domains in an 8:2 ratio. The DRP methods without the MSDA are trained and validated on the source domains and tested on the target domains. The response data is clustered by drug type and then

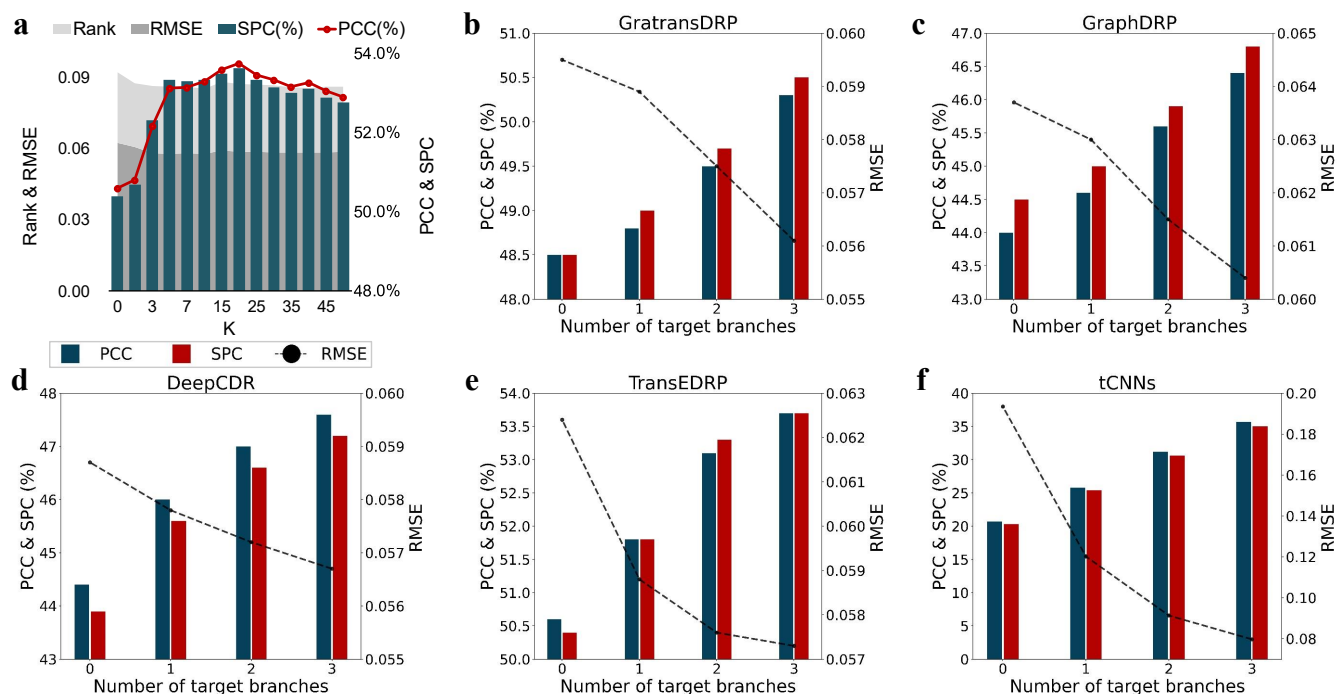


Figure 3: Hyperparameter experiment. (a)The impact of the source domain consisting of the number of drug domains K on the performance of the MSDA when the number of target domain adaptation branches is 2. (b-f)The impact of the number of target domain adaptation branches n on the performance of the MSDA when the number of drug domains corresponding to each target domain branch is fixed at 5. The experiments are conducted on the GDSCv2 public dataset, where (a) is the MSDA loaded on the TransEDRP method, and (b-f) are five representative DRP methods with MSDA loaded on.

randomly partitioned into source and target domains, with the type of drug serving as the criterion for division. As illustrated in Figure 2(c), it should be noted here that a single drug domain represents the smallest unit of division.

To comprehensively evaluate the impact of the MSDA, we employed several evaluation metrics: Root Mean Square Error (RMSE) to gauge deviation, Pearson correlation coefficient (PCC) for assessing linear correlation, Spearman’s Rank Correlation Coefficient (SRC) to measure monotonicity, and Margin Ranking Loss (Rank) to evaluate ranking performance.

4.2 Overall Experiment

Our proposed zero-shot learning solution aims to mitigate data limitations in drug response prediction, thereby enhancing the efficiency and accuracy of drug discovery and evaluation. Our goal is to validate the effectiveness of MSDA as a test-time enhancement plug-in for implementing this solution using various datasets and diverse DRP methods. In the overall experiment, we selected five publicly available DRP methods and utilized two drug response databases. The prediction results, including those that both do and do not integrate our proposed MSDA plug-in, are presented in Table 2. The prediction results of the original models are denoted by **Original**, while results after integrating MSDA are labeled **+MSDA**. The comparative analysis demonstrates that MSDA significantly enhances the performance of DRP methods in predicting responses for unknown drug domains. Across the

drug response databases GDSCv2 and CellMiner, the average PCC metrics exhibit improvements of 22.4% and 1.3%, respectively, while the average RMSE metrics indicate enhancements of 31.7% and 18.2%, respectively. This indicates that MSDA can bolster the assessment of preclinical drug candidates by offering more dependable predictions, a crucial factor for progressing potential drugs through the development pipeline.

4.3 Ablation Study

In this study, we introduce the MSDA plug-in as a key component of our proposed solution, which aims to enhance the performance of drug response prediction in the context of zero-shot learning. This plug-in can be seamlessly integrated with publicly available DRP methods. The MSDA creates multiple target domain adaptation branches by replicating the fusion prediction branches from the pre-trained DRP model while preserving the parameters of the original model prediction branches. Finally, the MSDA combines the predictions from the target domain adaptation branches with those of the original branches, using summation and averaging, to generate the fine-tuning results specific to the target domain. Therefore, in the design of the structure of the MSDA, two crucial aspects necessitate verification through ablation experiments conducted on the GDSCv2 public dataset:

A1 Whether the average adaptive fine-tuning of multiple target domain adaptation branches is better than that of

	K	n	Fusion Raw Branch	tCNNs	Results on GDSCv2 (RMSE ↓ / PCC ↑)			
					DeepCDR	GraphDRP	GratransDRP	TransEDRP
base	0	0	✗	0.1935/.2073	0.0587/.4442	0.0637/.4402	0.0595/.4848	0.0624/.5060
(A)	70	2	✓	0.0813/.3709	0.0549/.4946	<u>0.0579/.4898</u>	0.0561/.5103	0.0584/.5316
	70	2	✗	0.0544/.4852	0.0558/.4676	0.0577/.4788	0.0579/.4841	0.0590/.5143
(B)	15	3	✓	<u>0.0797/.3572</u>	0.0567/.4760	0.0604/.4643	<u>0.0561/.5034</u>	0.0573/.5371
	10	2	✓	0.0913/.3118	0.0572/.4699	0.0615/.4555	0.0575/.4950	<u>0.0576/.5311</u>
(C)	5	1	✓	0.1203/.2585	0.0578/.4597	0.0630/.4456	0.0595/.4848	0.0588/.5177
	10	1	✓	0.1134/.2773	0.0573/.4676	0.0621/.4524	0.0580/.4924	0.0580/.5273
	15	1	✓	0.1115/.2927	0.0569/.4740	0.0614/.4581	0.0570/.4979	<u>0.0591/.5323</u>

Table 3: Ablation study of the MSDA on GDSCv2 on five representative DRP methods. Base indicates the zero-shot learning performance of the DRP methods. (A) is used to compare the effect of merging the raw prediction branch from the original methods. (B) and (C) are used to compare the effect of K and n on the performance of the MSDA. The best performer and the second-best performer are highlighted in bold and underlined.

just one single branch for K drug domains as source domains filtered by the domain selector.

A2 Whether the results of the target domain adaptation branches n need to be merged with the results of the original branch predictions.

For **A1**, to investigate the structural performance disparities between multi-branching and single-branching within the MSDA framework while maintaining an equal number of source drug domains, we fixed the number of source domains K at 10 and 15, representing two and three times the original number of units (5), respectively. In the single-branching experiments, we set K to 10 and 15, respectively. In the multi-branching experiments, the number of source domains in each branch is fixed at 5, while n is set to 2 and 3, respectively. Table 3 (B) and (C) show the results of the experiments with an overall number of source domains of 15 and 10, respectively. These results indicate that a rational division of the K domains into $\frac{K}{n}$ sub-source domains, followed by fine-tuning using n target domain adaptation branches, is a more effective strategy.

For **A2**, we discuss the effectiveness of integrating the prediction branch of the original method with the target adaptation branch. The K is set to 70 and n is 2. We again select five representative DRP methods. As indicated in Table 3(A), the performance is enhanced when the original prediction branches of each DRP model are integrated. This is because, during the inference stage of domain adaptation, the fine-tuned branch lacks access to real results from the target domain, making it challenging to gauge the extent of domain adaptation adjustments. Integrating the domain adaptation branch with the original branch stabilizes the prediction results for the target domain. Furthermore, it is observed that when the original method’s effectiveness (e.g., tCNNs) is notably subpar, the mandatory fusion of the original branch and the domain adaptation branch diminishes the improvements brought about by the MSDA.

4.4 Hyperparameter Experiment

The MSDA acts as a multi-branch multi-source domain adaptation test-time plug-in that can be integrated into general DRP methods, incorporating two critical hyperparameters:

B1 The number of drug domains K selected by the domain selector while maintaining a fixed model structure.

B2 The number of target domain adaptation branches n with the number of drug domains remaining fixed for each branch adapted to the target domain.

For **B1**, we perform a sensitivity analysis with K values, considering 14 parameters selected at uneven intervals from 0 to 50, utilizing the TransEDRP method within the GDSCv2 dataset. The trend in Figure 3(a) illustrates the substantial impact of varying the value of K on the four metrics. Beyond $K = 20$, improvements either remain stable or exhibit a decreasing trend. This indicates that if the similarity between drugs in the source domains and those in the target domains is relatively low, confusion may be introduced into the pertinent information.

For **B2**, we conduct three sets of experiments, each with a varying number of target domain adaptation branches ($n \in [1, 2, 3]$), while using five drug domains as the source domains on each branch. Additionally, a blank control group is included. The experimental results on GDSCv2 for five existing published methods in Figure 3 (b-f) illustrate that as the number of target domain adaptation branches increases, the performance improvement of the models becomes more substantial for a specific number of drug domains that can be learned within each target domain adaptation branch.

5 Conclusion

In this paper, we introduce a zero-shot learning solution tailored to address the DRP task within preclinical drug screening. Our proposed solution MSDA is seamlessly integrated into conventional DRP methods, enabling the learning of invariant features from previous response label information and thereby augmenting the model’s real-time prediction capabilities. We conduct a series of experiments on GDSCv2 and CellMiner with the same division standard. Among the five most representative and high-performing DRP methods, the incorporation of the MSDA significantly enhances their predictive performance in the zero-shot learning scenario. To deal with the actual needs of the industry in real-world applications, the MSDA can adaptively select different parameters to balance the performance and inference speed requirements.

Acknowledgments

The work of Wenbin Hu was supported by the National Key Research and Development Program of China (2023YFC2705700). This work was supported in part by the Natural Science Foundation of China (No. 82174230), Artificial Intelligence Innovation Project of Wuhan Science and Technology Bureau (No. 2022010702040070), Natural Science Foundation of Shenzhen City (No. JCYJ20230807090211021).

References

- [Bai *et al.*, 2023] Peizhen Bai, Filip Miljković, Bino John, and Haiping Lu. Interpretable bilinear attention network with domain adaptation improves drug–target prediction. *Nature Machine Intelligence*, 5(2):126–136, 2023.
- [Ben-David *et al.*, 2010] Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman Vaughan. A theory of learning from different domains. *Machine learning*, 79:151–175, 2010.
- [Berdigaliyev and Aljofan, 2020] Nurken Berdigaliyev and Mohamad Aljofan. An overview of drug discovery and development. *Future medicinal chemistry*, 12(10):939–947, 2020.
- [Chen *et al.*, 2021] Xinyang Chen, Sinan Wang, Jianmin Wang, and Mingsheng Long. Representation subspace distance for domain adaptation regression. In *International Conference on Machine Learning*, pages 1749–1759. PMLR, 2021.
- [Chu *et al.*, 2023] Thang Chu, Thuy Trang Nguyen, Bui Duong Hai, Quang Huy Nguyen, and Tuan Nguyen. Graph transformer for drug response prediction. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 20(2):1065–1072, 2023.
- [Engels and Venkatarangan, 2001] MF Engels and P Venkatarangan. Smart screening: approaches to efficient hts. *Current opinion in drug discovery & development*, 4(3):275–283, 2001.
- [Ertl, 2003] Peter Ertl. Cheminformatics analysis of organic substituents: identification of the most common substituents, calculation of substituent properties, and automatic identification of drug-like bioisosteric groups. *Journal of chemical information and computer sciences*, 43(2):374–380, 2003.
- [Fang *et al.*, 2022] Zhen Fang, Jie Lu, Feng Liu, and Guangquan Zhang. Semi-supervised heterogeneous domain adaptation: Theory and algorithms. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(1):1087–1105, 2022.
- [Forbes *et al.*, 2017] Simon A Forbes, David Beare, Harry Boutselakis, Sally Bamford, Nidhi Bindal, John Tate, Charlotte G Cole, Sari Ward, Elisabeth Dawson, Laura Ponting, et al. Cosmic: somatic cancer genetics at high-resolution. *Nucleic acids research*, 45(D1):D777–D783, 2017.
- [Han *et al.*, 2023] Xiaochuang Han, Daniel Simig, Todor Mihaylov, Yulia Tsvetkov, Asli Celikyilmaz, and Tianlu Wang. Understanding in-context learning via supportive pretraining data. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12660–12673, 2023.
- [He *et al.*, 2022] Di He, Qiao Liu, You Wu, and Lei Xie. 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, 2022.
- [Irwin and Shoichet, 2016] John J Irwin and Brian K Shoichet. Docking screens for novel ligands conferring new biology: Miniperspective. *Journal of medicinal chemistry*, 59(9):4103–4120, 2016.
- [Jaritz *et al.*, 2022] Maximilian Jaritz, Tuan-Hung Vu, Raoul De Charette, Émilie Wirbel, and Patrick Pérez. Cross-modal learning for domain adaptation in 3d semantic segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(2):1533–1544, 2022.
- [Jiang *et al.*, 2021] Junguang Jiang, Yifei Ji, Ximei Wang, Yufeng Liu, Jianmin Wang, and Mingsheng Long. Regressive domain adaptation for unsupervised keypoint detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6780–6789, June 2021.
- [Lawrence *et al.*, 2014] Michael S Lawrence, Petar Stojanov, Craig H Mermel, James T Robinson, Levi A Garraway, Todd R Golub, Matthew Meyerson, Stacey B Gabriel, Eric S Lander, and Gad Getz. Discovery and saturation analysis of cancer genes across 21 tumour types. *Nature*, 505(7484):495–501, 2014.
- [Li and Hu, 2022] Kun Li and Wenbin Hu. Transedrp: Dual transformer model with edge emdedded for drug respond prediction, 2022.
- [Li *et al.*, 2020] Jingjing Li, Erpeng Chen, Zhengming Ding, Lei Zhu, Ke Lu, and Heng Tao Shen. Maximum density divergence for domain adaptation. *IEEE transactions on pattern analysis and machine intelligence*, 43(11):3918–3930, 2020.
- [Li *et al.*, 2022] Yuquan Li, Chang-Yu Hsieh, Ruiqiang Lu, Xiaoqing Gong, Xiaorui Wang, Pengyong Li, Shuo Liu, Yanan Tian, Dejun Jiang, Jiaxian Yan, et al. An adaptive graph learning method for automated molecular interactions and properties predictions. *Nature Machine Intelligence*, 4(7):645–651, 2022.
- [Liang *et al.*, 2021] Jian Liang, Dapeng Hu, Yunbo Wang, Ran He, and Jiashi Feng. Source data-absent unsupervised domain adaptation through hypothesis transfer and labeling transfer. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(11):8602–8617, 2021.
- [Liu *et al.*, 2018] Weiwei Liu, Donna Xu, Ivor W Tsang, and Wenjie Zhang. Metric learning for multi-output tasks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(2):408–422, 2018.

- [Liu *et al.*, 2019] Pengfei Liu, Hongjian Li, Shuai Li, and Kwong Sak Leung. Improving prediction of phenotypic drug response on cancer cell lines using deep convolutional network. *BMC Bioinformatics*, 20(1):1–14, 2019.
- [Liu *et al.*, 2020] Qiao Liu, Zhiqiang Hu, Rui Jiang, and Mu Zhou. DeepCDR: a hybrid graph convolutional network for predicting cancer drug response. *Bioinformatics*, 36(26):i911–i918, 2020.
- [Lyu *et al.*, 2019] Jiankun Lyu, Sheng Wang, Trent E Balius, Isha Singh, Anat Levit, Yurii S Moroz, Matthew J O’Meara, Tao Che, Enkhjargal Algaa, Kateryna Tolmachova, et al. Ultra-large library docking for discovering new chemotypes. *Nature*, 566(7743):224–229, 2019.
- [Nguyen *et al.*, 2022] Tuan Nguyen, Giang T. T. Nguyen, Thin Nguyen, and Duc-Hau Le. Graph convolutional networks for drug response prediction. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 19(1):146–154, 2022.
- [Ou *et al.*, 2023] Fu-Zhao Ou, Baoliang Chen, Chongyi Li, Shiqi Wang, and Sam Kwong. Troubleshooting ethnic quality bias with curriculum domain adaptation for face image quality assessment. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 20718–20729, 2023.
- [Panaretos and Zemel, 2019] Victor M Panaretos and Yoav Zemel. Statistical aspects of wasserstein distances. *Annual review of statistics and its application*, 6:405–431, 2019.
- [Pandey *et al.*, 2022] Mohit Pandey, Michael Fernandez, Francesco Gentile, Olexandr Isayev, Alexander Tropsha, Abraham C Stern, and Artem Cherkasov. The transformational role of gpu computing and deep learning in drug discovery. *Nature Machine Intelligence*, 4(3):211–221, 2022.
- [Peng *et al.*, 2019] Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source domain adaptation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1406–1415, 2019.
- [Pourpanah *et al.*, 2023] Farhad Pourpanah, Moloud Abdar, Yuxuan Luo, Xinlei Zhou, Ran Wang, Chee Peng Lim, Xi-Zhao Wang, and Q. M. Jonathan Wu. A review of generalized zero-shot learning methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(4):4051–4070, 2023.
- [Reinhold *et al.*, 2012] William C Reinhold, Margot Sunshine, Hongfang Liu, Sudhir Varma, Kurt W Kohn, Joel Morris, James Doroshow, and Yves Pommier. Cellminer: a web-based suite of genomic and pharmacologic tools to explore transcript and drug patterns in the nci-60 cell line set. *Cancer research*, 72(14):3499–3511, 2012.
- [Sadybekov and Katritch, 2023] Anastasiia V Sadybekov and Vsevolod Katritch. Computational approaches streamlining drug discovery. *Nature*, 616(7958):673–685, 2023.
- [Sadybekov *et al.*, 2022] Arman A Sadybekov, Anastasiia V Sadybekov, Yongfeng Liu, Christos Iliopoulos-Tsoutsouvas, Xi-Ping Huang, Julie Pickett, Blake Houser, Nilkanth Patel, Ngan K Tran, Fei Tong, et al. Synthone-based ligand discovery in virtual libraries of over 11 billion compounds. *Nature*, 601(7893):452–459, 2022.
- [Siqueira-Neto *et al.*, 2023] Jair L Siqueira-Neto, Kathryn J Wicht, Kelly Chibale, Jeremy N Burrows, David A Fidock, and Elizabeth A Winzeler. Antimalarial drug discovery: Progress and approaches. *Nature Reviews Drug Discovery*, 22(10):807–826, 2023.
- [Stärk *et al.*, 2022] Hannes Stärk, Octavian Ganea, Lagnajit Pattanaik, Regina Barzilay, and Tommi Jaakkola. Equibind: Geometric deep learning for drug binding structure prediction. In *International conference on machine learning*, pages 20503–20521. PMLR, 2022.
- [Stein *et al.*, 2020] Reed M Stein, Hye Jin Kang, John D McCorvy, Grant C Glatfelter, Anthony J Jones, Tao Che, Samuel Slocum, Xi-Ping Huang, Olena Savych, Yurii S Moroz, et al. Virtual discovery of melatonin receptor ligands to modulate circadian rhythms. *Nature*, 579(7800):609–614, 2020.
- [Stratton *et al.*, 2009] Michael R Stratton, Peter J Campbell, and P Andrew Futreal. The cancer genome. *Nature*, 458(7239):719–724, 2009.
- [Wang *et al.*, 2022] Boxin Wang, Wei Ping, Chaowei Xiao, Peng Xu, Mostofa Patwary, Mohammad Shoeybi, Bo Li, Anima Anandkumar, and Bryan Catanzaro. Exploring the limits of domain-adaptive training for detoxifying large-scale language models. *Advances in Neural Information Processing Systems*, 35:35811–35824, 2022.
- [Wu *et al.*, 2022] Pengfei Wu, Siyi Lin, Guodong Cao, Jia-abin Wu, Hangbiao Jin, Chen Wang, Ming Hung Wong, Zhu Yang, and Zongwei Cai. Absorption, distribution, metabolism, excretion and toxicity of microplastics in the human body and health implications. *Journal of Hazardous Materials*, 437:129361, 2022.
- [Yang *et al.*, 2012] Wanjuan Yang, Jorge Soares, Patricia Greninger, et al. Genomics of drug sensitivity in cancer (gdsc): a resource for therapeutic biomarker discovery in cancer cells. *Nucleic acids research*, 41:D955–D961, 2012.
- [Yang *et al.*, 2022] Kai Yang, Jie Lu, Wanggen Wan, Guangquan Zhang, and Li Hou. Transfer learning based on sparse gaussian process for regression. *Information Sciences*, 605:286–300, 2022.
- [Zhou *et al.*, 2023] Kaiyang Zhou, Ziwei Liu, Yu Qiao, Tao Xiang, and Chen Change Loy. Domain generalization: A survey. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, 45(04):4396–4415, 2023.