# Multi-Level Domain Adaptation and Contrastive Domain Isolation with Bilinear Fusion for Patient Drug Response Prediction

**Yuting Bai, Hanwen Lv, Wanwan Shi, Zhiyi Zou, Jiawei Luo***

College of Computer Science and Electronic Engineering, Hunan University, Changsha, 410082, China
{baiyuting, hanwenlv, shiwanwan, zouzhiyi31, luojiawei}@hnu.edu.cn

## Abstract

Accurate prediction of patient drug response is critical for precision cancer medicine but remains constrained by limited clinical data. While in vitro cell line data offer a scalable alternative, effective cross-domain transfer remains challenging. Many existing methods tend to overlook heterogeneous domain shifts across biological contexts, underrepresent the intrinsic differences between cell lines and patient tissues, and insufficiently capture high-order gene-drug interactions. To address these challenges, we propose MACB-DRP, a hierarchical transfer learning framework comprising three complementary stages that progressively coordinate adaptation across tissue, drug, and sample levels while enabling representation separation. The framework begins with tissue-aware domain adaptation, leveraging cancer-type classification and unsupervised alignment to preserve biologically meaningful structure across domains. It then incorporates drug-conditioned adversarial transfer for distribution alignment, coupled with bilinear fusion to model nonlinear and high-order gene-drug interactions. Finally, contrastive anchoring with feature-matched pairs enables fine-grained sample-level alignment, while feature-mismatched negatives preserve irreducible biological disparities. Experimental evaluation demonstrates that MACB-DRP achieves comprehensive predictive performance for patient drug responses, with robust results across multiple cancer types and nine drugs, and further reveals hierarchical structure across drugs and tissues in the visualization. These findings highlight the potential of biologically guided domain adaptation for improving translational pharmacogenomics.

**Code** — https://github.com/YutingBai99/MACB-DRP

## Introduction

Accurate prediction of patient-specific drug responses (DRP) is critical for realizing the promise of precision cancer medicine, as it directly informs treatment selection and profoundly influences therapeutic efficacy and patient survival. However, the development of robust predictive models remains fundamentally constrained by the scarcity of large-scale, high-quality clinical datasets. This limitation poses a major bottleneck, since supervised learning approaches typically require extensive labeled data to generalize effectively across diverse patient populations. Although TCGA (Weinstein et al. 2013) provides rich molecular profiles across a wide spectrum of cancer types, it contains only a limited number of matched drug response measurements, making it insufficient for direct supervised model training. In contrast, in vitro pharmacogenomic datasets derived from cancer cell lines, such as GDSC (Iorio et al. 2016), provide abundant molecular and drug response profiles, making them a scalable alternative for model training. To bridge the gap between preclinical and clinical domains, recent studies have explored domain adaptation techniques that transfer drug response knowledge from cell lines to patient tissues. Nevertheless, the biological discrepancies between these domains present substantial challenges for effective cross-domain generalization.

To mitigate the scarcity of clinical drug response data, unsupervised domain adaptation has been introduced into drug response prediction, aiming to transfer knowledge by aligning genomic distributions between cell lines (source domain) and patient samples (target domain). Representative work such as Celligner (Warren et al. 2021) performs global transcriptomic alignment to uncover similarities between cell lines and tumors, while Velodrome (Sharifi-Noghabi et al. 2021) formalizes the task as an out-of-distribution generalization problem, leveraging unlabeled data to learn domain-invariant representations. In terms of representation learning, VAEN (Jia et al. 2021) employs variational autoencoders to compress high-dimensional gene expression profiles and denoise predictive signals. CODE-AE (He et al. 2022) further introduces a shared-private encoder architecture with adversarial training to disentangle common biological features from domain-specific noise, enabling cross-domain prediction for five drugs. Building on CODE-AE, WISER (Shubham et al. 2024) incorporates weak supervision to enhance representation discriminability by utilizing large-scale unlabeled patient samples. Notably, TransDRP (Liu and Li 2025) is the first to explicitly account for tissue-type characteristics, integrating prototype contrastive learning to strengthen cancer-type-specific gene representations and demonstrating effectiveness on a curated clinical dataset covering nine drugs.

While recent efforts have yielded meaningful advances, some key challenges persist. Many existing methods tend to simplify complex domain shifts into a uniform abstrac-

*Corresponding author.

tion, overlooking the hierarchical heterogeneity across tissues, drugs, and individual samples. For example, Celligner (Warren et al. 2021) and Velodrome (Sharifi-Noghabi et al. 2021) rely on global alignment without explicitly modeling context-specific variations. TransDRP introduces tissue-level prototype contrastive learning to address cancer-type discrepancies but remains limited in capturing drug-conditioned and sample-specific shifts. In parallel, existing frameworks frequently enforcing rigid distribution alignment between cell lines and patient tissues, thereby overlooking biologically meaningful domain-specific signals during representation transfer. Although models such as CODE-AE (He et al. 2022) and WISER (Shubham et al. 2024) incorporate shared and private decomposition, they insufficiently retain the structural integrity of biologically unique features during adaptation. Furthermore, current models typically adopt shallow fusion strategies such as direct concatenation, which fail to capture high-order gene-drug interactions underlying response heterogeneity. These limitations collectively underscore the need for a unified framework capable of coordinating multi-level adaptation while preserving essential biological structure.

In this study, we introduce MACB-DRP, a multi-level domain adaptation framework designed to transfer pharmacogenomic knowledge from cell lines to patient drug response prediction. MACB-DRP performs domain alignment at the tissue, drug, and sample levels through three stages. Specifically, tissue-level adaptation is facilitated through tissue type classification combined with unsupervised domain adaptation. At the drug level, adversarial transfer mitigates drug-conditioned domain shifts between cell lines and patients, while bilinear fusion captures high-order gene-drug interactions. Furthermore, sample-level adaptation is performed via contrastive anchoring, which simultaneously pulls together a small number of cross-domain sample pairs with similar features and pushes apart a large set of mismatched pairs, thereby simulating intrinsic cross-domain disparities and enforcing domain isolation. This design enables MACB-DRP to hierarchically extract and align biological, pharmacological, and representational information across cell line and patient domains, providing a robust solution for drug response prediction.

In brief, the contributions of MACB-DRP are as follows:

- We propose a hierarchical domain adaptation framework for patient drug response prediction, which aligns representations across tissue, drug, and sample levels, while preserving intrinsic biological characteristics via contrastive domain isolation.
- Bilinear fusion is incorporated to explicitly model high-order gene-drug interactions, enabling more effective capture of complex pharmacogenomic relationships compared to direct feature concatenation.
- Representation robustness is enhanced by integrating residual connections into MLP, facilitating multi-scale feature extraction from genomic profiles.
- The effectiveness of MACB-DRP is validated via comprehensive evaluations, demonstrating strong generalization and improved performance over baselines.

# Related Works

## Drug Response Prediction

Drug response prediction (DRP) has attracted increasing attention in recent years, particularly in the context of precision oncology. Early efforts predominantly focused on in vitro cell-line data, leveraging pharmacogenomic datasets such as GDSC to predict drug sensitivity using metrics like IC50. Representative models, including but not limited to GraphDRP (Nguyen et al. 2021), DeepDSC (Li et al. 2019), Uno (Xia et al. 2022), and SubCDR (Liu and Zhang 2023), utilize various deep learning architectures to capture molecular and transcriptomic patterns in cell lines.

With the growing demand for personalized treatment and the emergence of patient-level drug response measurements, researchers have shifted focus toward predicting clinical drug responses. Due to the scarcity of labeled patient data, many approaches adopt transfer learning strategies to bridge the gap between cell-line and patient domains. These strategies include few-shot learning (Ma et al. 2021) , fine-tuning (Hostallero et al. 2023), weakly supervised learning methods such as WISER (Shubham et al. 2024), as well as unsupervised domain adaptation techniques such as PRECISE (Mourragui et al. 2019), Velodrome (Sharifi-Noghabi et al. 2021), and TransDRP (Liu and Li 2025). Domain-adversarial training has also been widely used, with representative methods including CODE-AE (He et al. 2022), WISER, TransDRP, and AD-AE (Dincer, Janizek, and Lee 2020). Among these, domain adversarial training is primarily employed for global alignment between the source and target domains. However, to date, limited attention has been given to the application of adversarial variants such as conditional domain adversarial approaches in clinical drug response prediction.

In addition, several recent models have introduced auxiliary clinical information to enhance prediction. For instance, DruID (Jayagopal et al. 2025) integrates clinical-grade next-generation sequencing (cNGS) panels, while PREDICT-AI (Jayagopal et al. 2024) incorporates mutation profiles and patient survival data in a multi-task learning framework. Although these approaches improve biological fidelity, they impose stricter data requirements and are typically limited to a smaller set of drugs.

## Contrastive Learning

Contrastive learning has emerged as a powerful paradigm for representation learning, particularly effective in encouraging models to distinguish informative features by comparing semantically similar and dissimilar sample pairs. Building on this paradigm, approaches like SimCLR (Chen et al. 2020), MoCo (He et al. 2020), and InfoNCE-based models (Wu, Wu, and Huang 2021) have achieved promising results by refining the mechanisms of contrastive signal formulation, with applications spanning visual and biomedical domains.

In the context of drug response prediction, contrastive learning has been extensively applied to cell-line datasets. For instance, GraphCDR (Liu et al. 2022) integrates contrastive objectives as a regularizer within a multi-task learning framework to enhance generalization. ADGCL (Zhuang

et al. 2023) employs a dual graph contrastive strategy, combining self-supervised and micro-supervised learning to extract high-level representations. Similarly, GCFMCLT (Wei et al. 2023) incorporates contrastive signals into a graph-based collaborative filtering architecture to predict sensitivity relationships between miRNA and drugs.

However, applications of contrastive learning to patient-level drug response prediction remain limited. To the best of our knowledge, TransDRP (Liu and Li 2025) is the only method that leverages tissue-level prototype contrastive learning to jointly predict responses to nine drugs within a single model, achieving strong predictive performance in drug response prediction. Despite its demonstrated effectiveness, the potential of contrastive learning in clinical drug response modeling is far from fully explored, especially in capturing nuanced biological heterogeneity across patients.

## Methodology

This section begins by formally defining the problem and outlining the primary objective. We then present MACB-DRP, a three-stage framework designed to address the identified challenges, as illustrated in Figure 1. The first stage performs tissue-aware representation learning through cancer-type classification and unsupervised domain adaptation. The second stage introduces drug-conditioned adversarial transfer and bilinear fusion to align domain-specific distributions and capture gene-drug interactions. Finally, the third stage applies contrastive anchoring with domain isolation to achieve fine-grained sample-level alignment and enforce representation separation. The trained model is subsequently used to infer drug responses in patient samples.

### Problem Definition

We address the task of predicting drug responses for patient samples by transferring pharmacogenomic knowledge from cell lines. The source domain is defined as $\mathcal{D}_s = \{(\mathbf{x}_s^i, t_s^i, y_s^{i,d})\}_{i=1}^{N_s}$, where $\mathbf{x}_s^i$ is the gene expression profile of the $i$-th cell line, $t_s^i$ is its cancer tissue type, and $y_s^{i,d} \in \{0,1\}$ represents the drug response label for drug $d$ (1 for sensitive, 0 for resistant), with $d \in \{1,\ldots,N_D\}$ indexing the set of $N_D$ distinct drugs. The target domain is $\mathcal{D}_t = \{(\mathbf{x}_t^j, t_t^j)\}_{j=1}^{N_t}$, consisting of $N_t$ patient samples with gene profiles $\mathbf{x}_t^j$ and tissue types $t_t^j$, but without drug response annotations. The goal is to learn a predictive function that generalizes to patient data and accurately estimates binary drug responses conditioned on both gene expression and drug features.

For clarity and conciseness, when operations are applied identically to both source and target domains, we omit domain-specific subscripts (e.g., s,t)and use unified notation throughout the formulation.

### Tissue-Aware Representation Learning

This stage performs unsupervised pretraining using all available gene expression and tissue-type data from both source ($\mathcal{D}_s$) and target ($\mathcal{D}_t$) domains, without incorporating drug response labels.

**Residual MLP.** In order to facilitate direct propagation of gene expression signals and enhance the capture of complex tissue-specific patterns, we employ a multilayer perceptron augmented with residual connections (Residual MLP) as our gene expression encoder, complemented by an attention mechanism for sample weighting. Each residual block consists of two consecutive Linear-BatchNorm-SELU-Dropout layers with a shortcut connection, enabling efficient information flow and improved feature refinement.

For each input sample $\mathbf{x} \in \mathcal{D}_s \cup \mathcal{D}_t$, the encoder first computes a latent representation $E_\phi(\mathbf{x})$ using a residual-enhanced MLP. A scalar attention weight $\alpha(E_\phi(\mathbf{x})) \in [0,1]$ is then generated to modulate the representation:

$$\mathbf{h} = \alpha(E_\phi(\mathbf{x})) \cdot E_\phi(\mathbf{x}). \tag{1}$$

This sample-wise scaling yields a tissue-aware, domain-invariant representation $\mathbf{h}$, which is subsequently used for classification and adversarial domain alignment.

**Supervised Tissue Classification.** To promote tissue-aware representation learning, we introduce a single-hidden-layer MLP as the cancer tissue-type classifier $C_{\text{tis}}$. Both the source domain $\mathcal{D}_s$ and the target domain $\mathcal{D}_t$ provide ground-truth tissue labels for all samples, denoted as $t_s^i$ and $t_t^j$, respectively. This enables fully supervised training using cross-entropy loss.

Due to distributional differences in tissue types across domains, we compute inverse-frequency weights $w^{(s)}$ and $w^{(t)}$ based on source and target sample counts, respectively. The weighted cross-entropy loss is then defined as:

$$\begin{aligned}\mathcal{L}_{\text{tis}} = &\frac{1}{N_s}\sum_{i=1}^{N_s} w_{t_s^i}^{(s)} \cdot \text{CE}\big(\mathbf{t}_s^i,\ C_{\text{tis}}(\mathbf{h}_s^i)\big) + \\ &\frac{1}{N_t}\sum_{j=1}^{N_t} w_{t_t^j}^{(t)} \cdot \text{CE}\big(\mathbf{t}_t^j,\ C_{\text{tis}}(\mathbf{h}_t^j)\big)\end{aligned} \tag{2}$$

**Unsupervised Domain Adaptation.** Domain adaptation is achieved via a discriminator $D_{\text{adv}}$, which encourages the encoder to learn domain-invariant representations by making features indistinguishable across source and target domains. The corresponding adversarial loss is:

$$\mathcal{L}_{\text{adv}} = -\frac{1}{N_s}\sum_{i=1}^{N_s}\log\ \hat{d}_s^i - \frac{1}{N_t}\sum_{j=1}^{N_t}\log\big(1-\hat{d}_t^j\big), \tag{3}$$

where $\hat{d}_s^i$ and $\hat{d}_t^j$ are the outputs of the domain discriminator $D_{\text{adv}}$ for source and target samples, respectively.

**Loss Function.** The total loss for Stage 1 is defined as the sum of the tissue-type classification loss and the adversarial domain adaptation loss:

$$\mathcal{L}_{\text{s1}} = \mathcal{L}_{\text{tis}} + \mathcal{L}_{\text{adv}}. \tag{4}$$

### Drug-Conditioned Adversarial Transfer and Bilinear Fusion

In Stage 2, we predict drug response for both source and target domains by combining tissue-aware patient features

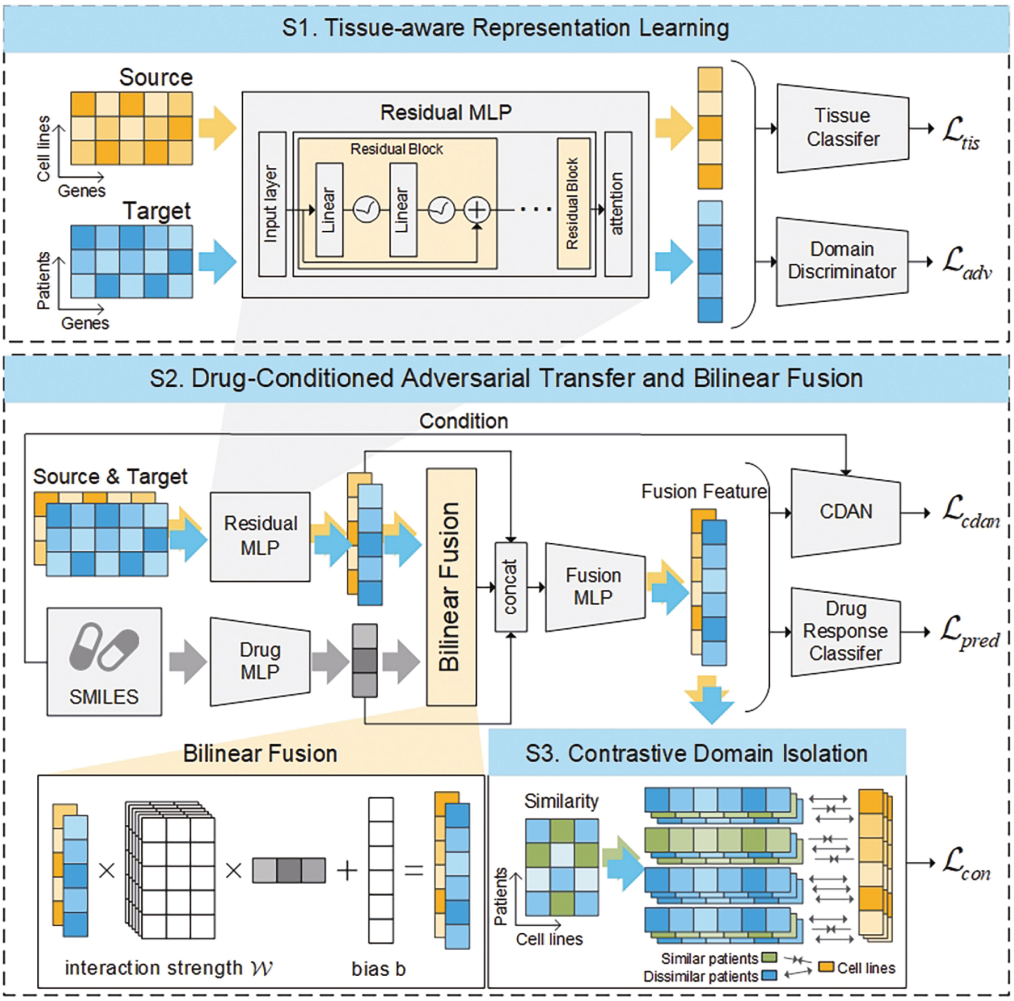


Figure 1: Overview of the MACB-DRP framework. The framework consists of three stages: S1. Tissue-aware representation learning, S2. Drug-Conditioned Adversarial Transfer and Bilinear Fusion, and S3. Contrastive domain isolation. In the main text, these components are referred to as Stage 1, Stage 2, and Stage 3, respectively.

with molecular representations. These are bilinearly fused with representations $\mathbf{h}$ to capture gene-drug interactions. To enhance generalization across domains with respect to drug-specific variations, a conditional adversarial discriminator aligns fused representations across domains under drug-specific supervision. This stage leverages labeled source-domain data $D_s$ and unlabeled target-domain data $D_t$.

**Drug Representation.** Each drug $d$ is initially represented by its SMILES string, which is subsequently converted into an extended connectivity fingerprint (ECFP), a fixed-length binary vector capturing molecular substructures. The ECFP feature $\mathbf{r}_d$ is projected into a latent representation via a two-hidden-layer MLP: $\mathbf{h}_d = E_d(\mathbf{r}_d)$.

**Bilinear Fusion.** Bilinear fusion is favored for its ability to explicitly and granularly model targeted interactions between feature pairs. Unlike most existing methods that rely on direct concatenation of patient and compound features, we adopt a tensorized bilinear fusion mechanism to capture high-order dependencies, which refer to interaction patterns involving pairwise multiplicative terms between gene and drug features. These patterns are essential for modeling synergistic or antagonistic effects that cannot be linearly separated across biological modalities.

Formally, let $\mathbf{H} \in \mathbf{R}^{N\times d}$ and $\mathbf{H}_D \in \mathbf{R}^{N_D\times d_D}$ denote the gene and drug feature matrices for the source or target domain, where each row corresponds to a sample's feature vector $\mathbf{h} \in \mathbf{R}^d$ or a drug representation $\mathbf{h}_D \in \mathbf{R}^{d_D}$. The fusion output $\mathbf{Z}_b \in \mathbf{R}^{N\times k\times N_D}$ is computed as:

$$\mathbf{Z}_b = \mathbf{H}\mathcal{W}\mathbf{H}_D^\top + \mathbf{b}, \tag{5}$$

Here, $\mathcal{W} \in \mathbf{R}^{k\times d\times d_D}$ is a learnable three-dimensional parameter that encodes pairwise interaction strengths between gene and drug features across $k$ output dimensions, and $\mathbf{b} \in \mathbf{R}^k$ is the bias vector added to each fusion channel.

To obtain high-level representations for downstream prediction and alignment, the fused features $\mathbf{Z}_b$ are concatenated with the original gene and drug features to form an

enriched representation $\tilde{\mathbf{Z}} = [\mathbf{Z}_b; \mathbf{H}; \mathbf{H}_D]$. This composite vector is then transformed by a three-hidden-layer MLP, producing the final embedding matrix $\mathbf{Z}$.

**Drug Response Prediction.** Fused representations $\mathbf{Z}$ are passed through a drug response predictor comprising a linear output layer with sigmoid activation. The response scores are computed as: $\hat{\mathbf{y}} = \sigma(W_r\,\mathbf{Z} + b_r) \in (0, 1)$. Since drug response labels are only available for source-domain samples, the prediction loss is computed using binary cross-entropy over source samples:

$$\mathcal{L}_{\text{pred}}^{(s)} = -\frac{1}{N_s}\sum_{i=1}^{N_s}\left[y_i \log\,\hat{y}_i + (1-y_i)\log(1-\hat{y}_i)\right] \quad (6)$$

**Drug-Conditioned Domain Adversarial Network.** Building on the prior unsupervised domain adaptation stage that achieved tissue-aware alignment of gene expression features, we apply Conditional Domain Adversarial Network (CDAN) to perform drug-conditioned alignment of fused representations.

To construct the conditioning signal, we first apply a simple drug classifier $C_D(\cdot)$ to the fusion representations $\mathbf{Z}$, followed by softmax activation, yielding predicted drug-level predictions $\hat{\mathbf{P}}_D = \text{Softmax}(C_D(\mathbf{Z}))$. These predictions are then concatenated with the fusion features and fed into the domain discriminator $D_{\text{C}}$ for drug-conditioned domain alignment, denoted as $\mathbf{F} = [\mathbf{Z}, \hat{\mathbf{P}}_D]$. The corresponding loss function is defined as:

$$\begin{aligned}\mathcal{L}_{\text{cdan}} = &-\frac{1}{N_s}\sum_{i=1}^{N_s}\log\,D_{\mathbf{C}}(\mathbf{F}_i^{(s)}) \\ &-\frac{1}{N_t}\sum_{j=1}^{N_t}\log\left(1 - D_{\mathbf{C}}(\mathbf{F}_j^{(t)})\right)\end{aligned} \quad (7)$$

Here, $\mathbf{F}_i^{(s)}$ and $\mathbf{F}_j^{(t)}$ are the concatenated representations from source and target domains, respectively.

**Loss Function.** During Stage 2, the model is optimized by combining drug response prediction loss on source-domain samples with drug-conditioned adversarial loss:

$$\mathcal{L}_{\text{s2}} = \mathcal{L}_{\text{pred}}^{(s)} + \mathcal{L}_{\text{cdan}}. \quad (8)$$

## Contrastive Anchoring with Domain Isolation

Stage 1 and Stage 2 enable global alignment at the tissue and drug levels, respectively, allowing coarse-grained consistency between source and target domains. However, samples derived from cell lines and patient tissues inherently differ in biological distribution, and samples rarely to simultaneously match across tissue type, drug identity, and response outcome. Appendix 1 provides detailed statistics supporting this observation.

To address this, we designed Contrastive Anchoring with Domain Isolation. In each epoch, cosine similarity is computed between all fused representations from the source and target domains. For each source sample, the Top-K most similar target samples are selected as positive examples, while all remaining target samples are treated as negatives. This construction yields a compact set of highly relevant positives that serve as anchors pulling source samples toward biologically similar targets, along with a large pool of dissimilar negatives that enforce domain-level isolation by explicitly pushing apart incompatible representations. We adopt $K = 7$ based on results from the hyperparameter sensitivity analysis. The contrastive loss over source samples is defined as:

$$\mathcal{L}_{\text{con}} = -\frac{1}{N_s}\sum_{i=1}^{N_s}\log\left(\frac{S_i^+}{S_i^+ + S_i^-}\right) \quad (9)$$

with $S_i^+ = \sum_{j\in\mathcal{P}(i)}\exp\left(\cos(\mathbf{z}_i^s, \mathbf{z}_j^t)/\tau\right)$, and $S_i^- = \sum_{j\in\mathcal{N}(i)}\exp\left(\cos(\mathbf{z}_i^s, \mathbf{z}_j^t)/\tau\right)$, where $\mathcal{P}(i)$ and $\mathcal{N}(i)$ denote the Top-K most similar and remaining target samples, respectively. The temperature $\tau$ controls the sharpness of the similarity distribution and is empirically set to 0.1.

The total training objective for Stage 3 combines the same source-supervised drug response prediction used in Stage 2 with contrastive alignment:

$$\mathcal{L}_{\text{s3}} = \mathcal{L}_{\text{pred}}^{(s)} + \alpha\,\mathcal{L}_{\text{con}} \quad (10)$$

where $\alpha$ denotes the strength of contrastive alignment. We set $\alpha = 0.1$ in our experiments, as further examined in the sensitivity analysis.

# Experiments

## Experiment Settings

**Dataset.** We evaluate MACB-DRP on the same dataset used in TransDRP, comprising 1,517 cell line samples from CCLE (Barretina et al. 2019) as the source domain and 9,808 patient samples from TCGA (Weinstein et al. 2013) as the target domain. Each sample is annotated with gene expression profiles and cancer tissue types (He et al. 2022). Drug response labels for source-domain samples are derived from GDSC IC50 measurements (Iorio et al. 2016), binarized using a standard z-score threshold of 0: samples with $z$-score$(\text{IC}_{50}) < 0$ are labeled as sensitive, and those with $\geq 0$ as resistant. For the target domain, 1,186 labeled patient samples are used for testing. Following TransDRP, complete and partial responses are treated as sensitivity, while stable and progressive diseases are treated as resistance. We select 9 drugs with at least 90 labeled patient samples for evaluation: 5-Fluorouracil, Cisplatin, Cyclophosphamide, Docetaxel, Doxorubicin, Etoposide, Gemcitabine, Temozolomide, and Paclitaxel. The SMILES representations of these compounds are retrieved from PubChem (Kim et al. 2025). More dataset details are provided in the Appendix 1.

**Model Configuration.** In Stage 1, gene expression profiles are encoded using a residual MLP composed of an input layer, four ResNet blocks with hidden dimensions of 512, 256, 128, and 64, followed by an attention layer. The output is a 64-dimensional gene representation. In Stage 2, drug SMILES strings are first converted into molecular objects and then transformed into 128-bit fingerprints. These are passed through an MLP to obtain 64-dimensional

| Methods | MACB-DRP | TransDRP | WISER | CODE-AE | DruID | VAEN |
|---|---|---|---|---|---|---|
| **AUC** | **0.6317**/0.0892 | 0.6188/0.0867 | 0.5816/**0.0829** | 0.4912/0.1061 | 0.4923/0.0828 | 0.5049/0.0881 |
| **AUPR** | **0.7521**/0.1396 | 0.7511/0.1507 | 0.7106/0.1629 | 0.6844/0.1923 | 0.6526/0.1907 | 0.6794/**0.1346** |
| **F1** | 0.6519/**0.1110** | 0.5840/0.1269 | **0.6888**/0.2458 | 0.6103/0.2517 | 0.4037/0.1511 | 0.6163/0.2202 |
| **Average** | **0.6786/0.1133** | 0.6513/0.1214 | 0.6603/0.4126 | 0.5953/0.1834 | 0.5162/0.1415 | 0.6002/0.1476 |

Table 1: Performance comparison of baseline methods on patient-specific drug response prediction across 9 drugs, evaluated using AUC, AUPR, F1-score, and their average values. Values after / denote the standard deviation for each metric, or the average of the three standard deviations for the average row. The best performer is highlighted in **bold**.

drug embeddings. The bilinear fusion module takes the 64-dimensional gene and drug features as input and produces a 128-dimensional interaction representation. This is concatenated with the original gene and drug features to form a 256-dimensional fused vector, which is further refined by a three-hidden-layer MLP with hidden dimensions of 256, 128 and 64, yielding the final 64-dimensional representation.

**Baselines.** To demonstrate the effectiveness of MACB-DRP, we compare it against three latest methods, namely TransDRP (Liu and Li 2025), WISER (Shubham et al. 2024), and DruID (Jayagopal et al. 2025), as well as two classical baselines, CODE-AE (He et al. 2022) and VAEN (Jia et al. 2021). Among these, MACB-DRP, TransDRP, and DruID perform multi-drug prediction within a unified model, whereas WISER, CODE-AE, and VAEN train separate models for each drug. For fair comparison, all drugs are evaluated using consistent parameter settings within each method, based on configurations reported in the original papers and publicly available code. Detailed parameter specifications are provided in the Appendix 3.

**Evaluation Protocol.** We evaluate model robustness via five-fold cross-validation, using AUC, AUPR, and F1-score as primary metrics. Their average serves as a composite indicator, and the standard deviations are calculated across all drugs and folds for each metric, reflecting overall performance consistency. For reproducibility, refer to Appendix 2 for detailed system configurations.

## Experimental Results

**Drug Response Prediction.** As shown in Table 1, MACB-DRP attains the highest AUC and AUPR scores, with competitive standard deviations. While WISER reports a higher F1-score, its substantially larger standard deviation suggests poor consistency across folds and drugs. It is worth noting that DruID could not be trained on our dataset due to missing preprocessed data. Instead, we used the training set provided by DruID for evaluation, which likely contributed to its inferior performance. For the same reason, the reported standard deviation for DruID reflects variation only across drugs, without fold-level averaging, and is therefore excluded from cross-method stability comparisons. Overall, MACB-DRP achieves the highest average performance and the lowest average standard deviation across all metrics. These results demonstrate strong predictive accuracy and cross-drug robustness under five-fold cross-validation across 9 drugs. Appendix 4 lists experimental results for each drug across several key methods.

**Ablation Study.** To assess the contribution of each component in MACB-DRP, we design ablation experiments targeting five key modules: the tissue-aware loss $L_{\text{tis}}$, adversarial loss $L_{\text{adv}}$, conditional domain alignment loss $L_{\text{cdan}}$, bilinear fusion, and the contrastive loss $L_{\text{con}}$. All experiments are conducted using five-fold cross-validation. The average performance and the corresponding standard deviations across all drugs is reported in Table 2.

Our ablation studies confirm the synergistic design of MACB-DRP. Removing tissue-aware adaptation and adversarial transfer causes the largest performance drops, with the average score decreasing from 0.6786 to 0.6182 and 0.6187 respectively, validating their roles in encoding cancer-type priors and aligning global distributions. Excluding CDAN lowers the score to 0.6633, indicating its contribution to drug-conditional alignment. Bilinear fusion reduces the average score to 0.6631 and the F1 score from 0.6519 to 0.6201, further highlighting its role in modeling complex gene-drug interactions. Removing contrastive anchoring results in a score of 0.6491, reflecting its importance in sample-level alignment and the preservation of biologically meaningful disparities. Standard deviations also increase when each module is removed, especially for tissue-aware adaptation, adversarial transfer, and contrastive anchoring, further supporting their contributions to robustness and overall synergy of MACB-DRP.

In particular, since the positive and negative sample pairs in MACB-DRP's contrastive anchoring serve distinct purposes, namely precise cross-domain feature alignment and domain separation respectively, we additionally conducted an ablation study focused on their individual roles, as summarized in Table 3. Specifically, w/o con-p denotes contrastive learning using only negative pairs, while w/o con-n refers to the use of positive pairs exclusively. The results show that removing positive pairs leads to slight decreases in all three performance metrics, accompanied by a minor increase in standard deviation. In contrast, eliminating negative pairs causes substantial performance degradation and a pronounced increase in standard deviation. These observations demonstrate the indispensable role of domain separation in MACB-DRP.

**Visualization Analysis.** To enhance interpretability, we visualize the fused gene features for four representative drugs and illustrate the tissue-type distribution within the

| Methods | MACB-DRP | w/o tis | w/o adv | w/o cdan | w/o bilinear | w/o con |
|---|---|---|---|---|---|---|
| **AUC** | **0.6317/0.0892** | 0.5731/0.1098 | 0.5675/0.1192 | 0.6171/0.0987 | 0.6221/0.1024 | 0.5909/0.0934 |
| **AUPR** | **0.7521/0.1396** | 0.7258/0.1471 | 0.7075/0.1580 | 0.7437/0.1466 | 0.7471/0.1455 | 0.7137/0.1505 |
| **F1** | **0.6519/0.1110** | 0.5557/0.1445 | 0.5811/0.1488 | 0.6292/0.1223 | 0.6201/0.1265 | 0.6428/0.1443 |
| **Average** | **0.6786/0.1133** | 0.6182/0.1338 | 0.6187/0.142 | 0.6633/0.1225 | 0.6631/0.1248 | 0.6491/0.1294 |

Table 2: Predictive performance and standard deviations from five-fold CV across 9 drugs for ablation variants of MACB-DRP.

| Methods | MACB-DRP | w/o con-p | w/o con-n |
|---|---|---|---|
| **AUC** | **0.6317/0.0892** | 0.6271/0.0949 | 0.5861/0.1073 |
| **AUPR** | **0.7521/0.1396** | 0.7472/0.1406 | 0.7160/0.1472 |
| **F1** | **0.6519/0.1110** | 0.6391/0.1196 | 0.6399/0.1489 |
| **Average** | **0.6786/0.1133** | 0.6711/0.1184 | 0.6473/0.1345 |

Table 3: Ablation study results concerning the positive and negative samples used in contrastive anchoring.

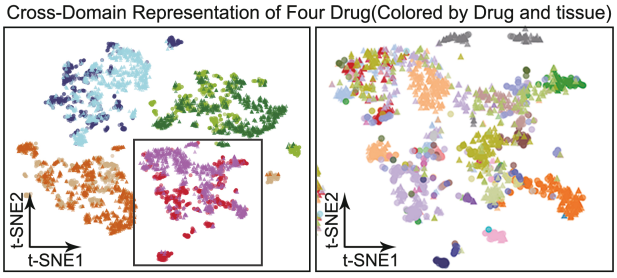


Figure 2: t-SNE visualization of cross-domain fused gene-drug representations for four drugs: 5-Fluorouracil, Cisplatin, Cyclophosphamide, and Docetaxel. Each point represents a sample, with circles denoting cell line samples and triangles denoting patient samples. The left panel uses color to indicate drug identity and sample domain (source vs. target), while the right panel illustrates tissue-type coloring for Docetaxel as a representative example.

fused features of Docetaxel, as shown in the Figure 2. In the left panel, the features of all four drugs exhibit clear aggregation, with source and target domain samples well integrated in the latent space. In the right panel, the fused features of Docetaxel display a distinct tissue-type-specific distribution, indicating that MACB-DRP effectively captures biologically meaningful hierarchically structured representations. This tissue-level stratification is not unique to Docetaxel, but consistently observed across other drugs, as documented in the Appendix 5. These results highlight the dual capability of MACB-DRP to achieve robust cross-domain alignment while preserving biological structure across both drug-level and tissue-level hierarchies.

**Sensitivity Analysis.** In the contrastive anchoring stage, we introduce two hyperparameters: the number of selected positive samples $K$, and the weight of the contrastive loss $\alpha$, as illustrated in the Figure 3.

The value of $K$ ranges from 1 to 15, considering only

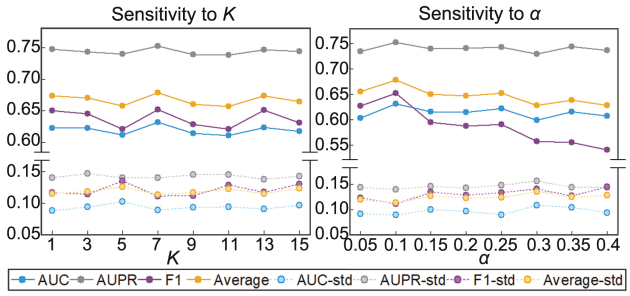


Figure 3: Sensitivity analysis of hyperparameters $K$ and $\alpha$ in the MACB-DRP framework.

odd integers. Overall, the model exhibits stable performance across different $K$ values, indicating robustness of the anchoring mechanism. Notably, when $K = 7$, the model achieves the highest average score (0.6786) and the lowest average standard deviation (0.1133).

For $\alpha$, we vary its value from 0.05 to 0.4 in increments of 0.05. When $\alpha < 0.1$, all metrics decline and standard deviations increase, suggesting insufficient contrastive strength. As $\alpha > 0.1$, AUPR remains relatively stable, whereas both AUC and F1-score begin to decrease, with F1 showing a particularly sharp drop. Consequently, the overall average score declines. Meanwhile, the standard deviation gradually rises, indicating reduced robustness in both five-fold cross-validation and cross-drug generalization. The best performance is observed at $\alpha = 0.1$, where the model achieves optimal balance between effectiveness and stability.

## Conclusion

In this work, we propose MACB-DRP, a domain adaptation framework for transferring knowledge from cell lines to patient-specific drug response prediction. To effectively extract individual-level representations across diverse tissues and drugs, MACB-DRP performs hierarchical cross-domain alignment at the tissue, drug, and sample levels while preserving intrinsic biological signals. It further models the complex interactions between genes and drugs through bilinear fusion. Experimental results demonstrate that MACB-DRP achieves strong predictive performance and cross-drug stability in terms of AUC, AUPR, and F1 score. Visualization analysis reveals that the learned gene-drug representations exhibit a hierarchical structure from drug identity to tissue specificity, underscoring their biological relevance and interpretability.

## Acknowledgments

This work was supported by the National Natural Science Foundation of China (Grant No.62372165 and 62032007)

## References

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. 2019. Addendum: The Cancer Cell Line Encyclopedia enables predictive modelling of anticancer drug sensitivity. *Nature*, 565(7738): E5–E6.

Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, 1597–1607. PmLR.

Dincer, A. B.; Janizek, J. D.; and Lee, S.-I. 2020. Adversarial deconfounding autoencoder for learning robust gene expression embeddings. *Bioinformatics*, 36(Supplement_2): i573–i582.

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.

He, K.; Fan, H.; Wu, Y.; Xie, S.; and Girshick, R. 2020. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9729–9738.

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.

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.

Jayagopal, A.; Walsh, R. J.; Hariprasannan, K. K.; Mariappan, R.; Mahapatra, D.; Jaynes, P. W.; Lim, D.; Tan, D. S. P.; Tan, T. Z.; Pitt, J. J.; et al. 2025. A multi-task domain-adapted model to predict chemotherapy response from mutations in recurrently altered cancer genes. *iScience*, 28(3).

Jayagopal, A.; Xue, H.; He, Z.; Walsh, R. J.; Hariprasannan, K. K.; Tan, D. S. P.; Tan, T. Z.; Pitt, J. J.; Jeyasekharan, A. D.; and Rajan, V. 2024. Personalised drug identifier for cancer treatment with transformers using auxiliary information. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 5138–5149.

Jia, P.; Hu, R.; Pei, G.; Dai, Y.; Wang, Y.-Y.; and Zhao, Z. 2021. Deep generative neural network for accurate drug response imputation. *Nature communications*, 12(1): 1740.

Kim, S.; Chen, J.; Cheng, T.; Gindulyte, A.; He, J.; He, S.; Li, Q.; Shoemaker, B. A.; Thiessen, P. A.; Yu, B.; et al. 2025. PubChem 2025 update. *Nucleic acids research*, 53(D1): D1516–D1525.

Li, M.; Wang, Y.; Zheng, R.; Shi, X.; Li, Y.; Wu, F.-X.; and Wang, J. 2019. DeepDSC: a deep learning method to predict drug sensitivity of cancer cell lines. *IEEE/ACM transactions on computational biology and bioinformatics*, 18(2): 575–582.

Liu, X.; and Li, M. 2025. Knowledge-Guided Domain Adaptation Model for Transferring Drug Response Prediction from Cell Lines to Patients. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 523–531.

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

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

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.

Nguyen, T.; Nguyen, G. T.; Nguyen, T.; and Le, D.-H. 2021. Graph convolutional networks for drug response prediction. *IEEE/ACM transactions on computational biology and bioinformatics*, 19(1): 146–154.

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.

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*.

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

Wei, J.; Zhuo, L.; Zhou, Z.; Lian, X.; Fu, X.; and Yao, X. 2023. GCFMCL: predicting miRNA-drug sensitivity using graph collaborative filtering and multi-view contrastive learning. *Briefings in Bioinformatics*, 24(4): bbad247.

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.

Wu, C.; Wu, F.; and Huang, Y. 2021. Rethinking infonce: How many negative samples do you need? *arXiv preprint arXiv:2105.13003*.

Xia, F.; Allen, J.; Balaprakash, P.; Brettin, T.; Garcia-Cardona, C.; Clyde, A.; Cohn, J.; Doroshow, J.; Duan, X.; Dubinkina, V.; et al. 2022. A cross-study analysis of drug response prediction in cancer cell lines. *Briefings in bioinformatics*, 23(1): bbab356.

Zhuang, L.; Wang, H.; Zhao, J.; and Sun, Y. 2023. Adaptive dual graph contrastive learning based on heterogeneous signed network for predicting adverse drug reaction. *Information Sciences*, 642: 119139.