GBD-DRP: Generalizable Bias-Disentangled Learning for Drug Response Prediction with Missing Modalities

Biyang Zeng , Shikui Tu*

School of Computer Science, Shanghai Jiao Tong University, China

{zengbiyang, tushikui}@sjtu.edu.cn

Abstract

Predicting the response of anticancer drugs for specific cancer cell lines remains an important challenge in cancer treatment, largely due to the diversity among cell lines. Existing machine learning methods usually perform well on cell lines seen during training, but their performance often drops when applied to new cell lines. Many approaches rely heavily on cell-line specific features, which limits their ability to generalize. At the same time, methods that focus mainly on generalization often fail to capture differences in drug response across distinct cellular environments. Some studies improve prediction accuracy by introducing additional and more detailed cell line features. However, these features are difficult to collect and are frequently missing, especially for newly studied cell lines. To address these issues, we propose a generalizable bias-disentangled framework for drug response prediction. Our method separates drug-related patterns that are shared across cell lines from representations that reflect cell-line specific environments. We further isolate real but non-generalizable biases that arise from noise or unobserved factors. This design helps preserve accurate predictions for specific cell lines while improving performance on unseen ones. In addition, the bias-related branch supports counterfactual consistency learning when some modalities are missing, allowing the model to handle incomplete data. Experimental results show that GBD-DRP achieves state-of-the-art performance, and additional tests on unseen cell lines and missing-modality settings indicate good generalization in practice.

1 Introduction

Cancer remains one of the most serious threats to human health, and drug-based therapy is currently the main approach for cancer treatment. Due to the genetic diversity of cancer cell lines, the same drug can show very different responses across different cell lines. As a result, accurately predicting drug response is an important problem and also a key component of future personalized medication. With the development of sequencing technologies and the construction of large-scale databases, machine learning methods have shown clear progress on this task. In particular, resources such as the Genomics of Drug Sensitivity in Cancer (GDSC) [Yang *et al.*, 2012] and the Cancer Cell Line Encyclopedia (CCLE) [Barretina *et al.*, 2012] have enabled data-driven modeling of drug response, and recent deep learning methods have achieved encouraging results [Baptista *et al.*, 2021; Partin *et al.*, 2023].

Due to the irregular structure of drug molecules and the heterogeneity of cell-line features, many approaches adopt graph-based models. For example, GraphDRP [Nguyen *et al.*, 2021] and DeepCDR [Liu *et al.*, 2020] represent drugs as molecular graphs and use Graph Convolutional Networks (GCN) to encode drug structures. Another line of work applies transformer-based models to drug SMILES representations. Methods such as DeepTTA [Jiang *et al.*, 2022] decompose canonical SMILES strings into functional groups, encode them as vocabularies, and use transformers to obtain drug representations. These models focus on capturing structural information from drugs in a flexible manner.

To further improve prediction accuracy, some studies pay more attention to modeling interactions between drugs and cell lines. Instead of encoding drugs and cell lines separately and simply concatenating their representations for prediction, these methods introduce interaction modeling during the encoding stage. FA^{2}GL [Xu *et al.*, 2025] constructs graphs for both drugs and cell lines and connects them through pipe nodes, allowing the heterogeneous graph to capture drug–cell interactions more explicitly. PASO [Wu *et al.*, 2025] use attention network to get drug representations and applies cross-attention with cell-line features to model the effects of drugs on different biological pathways. In addition, these approaches often incorporate multiple cell-line modalities, including gene expression (GEP), gene mutation profiles (MUT), and copy number variations (CNV), to better describe cell-line-specific responses.

Although these methods perform well under standard evaluation settings, they also introduce new challenges. As discussed by Codice et al. [Codicè *et al.*, 2025], many preclinical studies show limited generalization when tested on truly unseen cell lines, which are more relevant to real clinical

*Corresponding author

scenarios. This highlights the need to explicitly evaluate and improve generalization to unseen cell lines. Moreover, the use of additional modalities increases the risk of missing data. Compared with gene expression data, far fewer cell lines have complete MUT and CNV profiles. As a result, methods that rely on these features often exclude cell lines with missing modalities during testing. This can introduce bias and further worsen generalization to new cell lines, while also requiring additional data collection before predictions can be made.

To address these issues, we propose GBD-DRP, a generalizable bias-disentangled learning framework for drug response prediction. We decompose drug response into three channels. The first channel captures drug-invariant properties and represents the baseline efficacy of a drug. We model this component using an EdgeGAT-based graph encoder on drug molecular graphs. The second channel represents explicit cell-line-specific effects that can be modeled using available features such as GEP, MUT, and CNV. This branch captures drug responses under specific cellular environments and remains generalizable when features are available. The third channel models residual bias, which includes unobserved mechanisms and noise that cannot be explained by existing features. This component corresponds to domain bias in traditional domain generalization settings. By separating the residual bias from explicit cell-line-specific representations, our method preserves cell-line specificity while improving generalization to unseen cell lines.

For potentially missing modalities, we use the explicit cell-specific branch to process the modality when it is available, treating it as a supplement to the standard cell-line features. When the modality is not available, the residual branch is used to estimate its effect. By applying a counterfactual consistency constraint on the residual branch, the model can leverage existing samples of optional modalities to infer the influence of the missing samples. This allows the framework to utilize incomplete additional modality information effectively.

The main contributions of this work are as follows:

- We introduce a Triple-Branch Disentanglement framework. It separates explicit biological specific features from spurious biases, which helps the model generalize better to new cell lines.
- We provide a solution for missing modalities using a Counterfactual Consistency mechanism. The model can use existing features to estimate the effects of missing modalities through the residual branch, without relying on imputation.
- Experiments shows GBD-DRP outperforms current state-of-the-art methods. And also maintains strong accuracy and robustness, especially for unseen cell lines and cases with missing modalities.

2 Related Works

2.1 Domain Generalization

Domain Generalization (DG) is a method for handling out-of-distribution (OOD) data [Zhou *et al.*, 2022]. It aims to train a model on source domains and allow it to generalize to target domains. Methods such as domain alignment, meta-learning, and data augmentation are commonly used. DG either models domain shifts directly or reduces the influence of domain-specific features through data or training techniques. One important branch is domain-agnostic learning [Peng *et al.*, 2019]. This approach disentangles input features into domain-invariant true signals and domain-specific bias. The domain-specific bias is then removed, allowing the model to make predictions that do not depend on the domain.

2.2 The Generalization Challenge in DRP

In principle, domain-agnostic learning seems suitable for predicting drug response in new cell lines, because it separates the interference caused by cell-line-specific features. However, drug response is highly specific to cell lines. Removing all domain-specific information often reduces prediction accuracy and can even cause model collapse. In this biochemical setting, it is more important to explicitly model cell-line-specific features to generate accurate predictions for each cell line. But at the same time, focusing too much on cell-line-specific features can increase bias in the model.

To address this dilemma, GBD-DRP further disentangles cell-line-specific features from residual biases that cannot be generalized. By separating these two parts, the model can accurately use cell-line features for new cell lines while removing non-generalizable bias. This design allows us to use domain generalization principles to improve generalization without losing important cell-line-specific information.

3 Methodology

3.1 Problem Formulation and Latent Decomposition

Let $\mathcal{D}$ and $\mathcal{C}$ denote the set of drugs and cell lines. For a pair (d, c), the input consists of drug molecular features x_d and cellular genomic features x_c^{bio} (composed of base and optional modalities). The ground truth response is $y \in \mathbb{R}$. The goal of Drug Response Prediction (DRP) is to learn a mapping $f : (x_d, x_c^{bio}) \to y$ that generalizes to unseen cell identities $c \notin \mathcal{C}_{train}$.

Standard regression often overfits to cell-specific spurious correlations (e.g., batch effects or unmeasured confounds). To address this, we do not treat the response as a monolithic target. Instead, we formulate the expected response $\hat{y}$ as a sum of three disentangled components, motivated by the separation of global pharmacological effects from local biological interactions and domain noise:

$$\hat{y} = \underbrace{\mu(d)}_{\text{Intrinsic}} + \underbrace{\tau(x_d, x_c^{bio})}_{\text{Explicit}} + \underbrace{\epsilon(x_d, c)}_{\text{Residual}} \tag{1}$$

This decomposition guides our **Triple-Branch Disentanglement** architecture (Figure 1), where each term corresponds to a distinct functional module:

- **Intrinsic Potency (μ):** The baseline effect of drug d (e.g., general toxicity) roughly invariant across cell lines. Modeled by $f_{inv}(z_{inv})$.

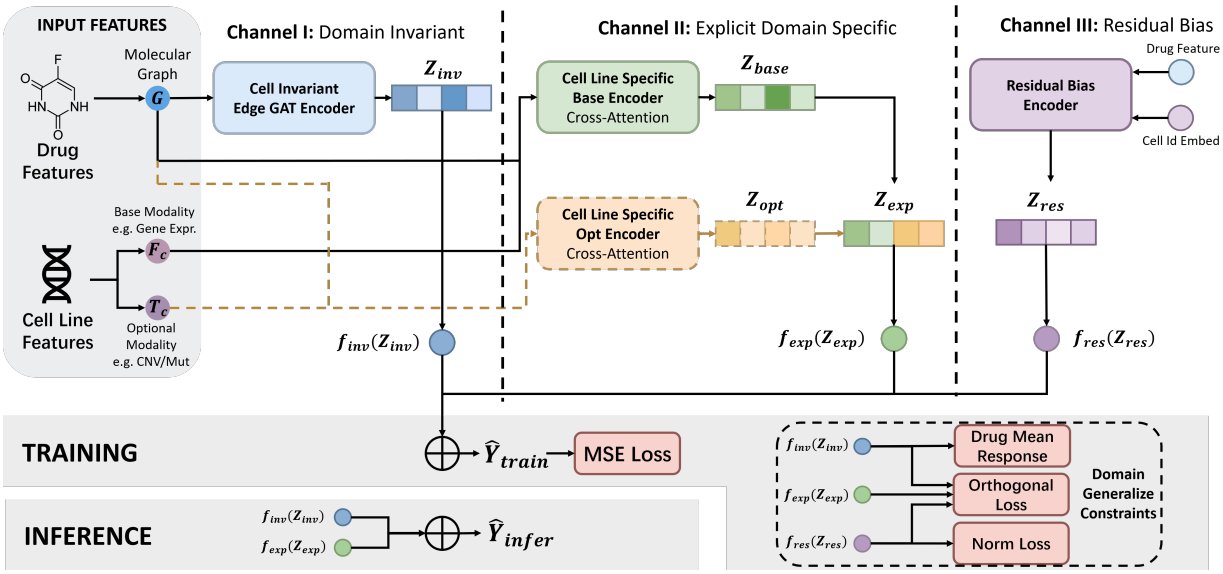

Figure 1: The overall architecture of GBD-DRP. Our framework consists of three functionally disentangled channels: (i) Channel I (Domain Invariant) captures the drug's intrinsic potency $\mu(d)$ through an Edge GAT encoder; (ii) Channel II (Explicit Domain Specific) models mechanistic drug-cell interactions via cross-attention, accommodating both base modalities (F_c) and optional modalities (T_c) through dynamic routing; (iii) Channel III (Residual Bias) isolates residual domain noise U_c using cell identity embeddings. During training, all channels contribute to $\hat{Y}_{train}$ to isolate biases. For inference on unseen cell lines, Channel III is deactivated to ensure generalization, with the final prediction $\hat{Y}_{infer}$ derived solely from invariant and explicit biological signals.

- **Explicit Specificity (τ):** The variation attributable to specific interactions between drug structure and *observed* genomic features x_c^{bio}. This represents the mechanistic, generalizable signal. Modeled by $f_{exp}(z_{exp})$.
- **Residual Bias (ϵ):** The discrepancy arising from unobserved factors coupled with cell identity c (e.g., experimental noise, unmeasured omics). Modeled by $f_{res}(z_{res})$.

GBD-DRP Overview

Guided by Eq. (1), GBD-DRP implements a **Triple-Branch Disentanglement** architecture to model these components explicitly. As illustrated in Figure 1, the model maps the input space into three orthogonal latent representations:

$$\hat{Y} = \underbrace{f_{inv}(z_{inv})}_{\approx\mu(d)} + \underbrace{f_{exp}(z_{exp})}_{\approx\tau(d,c_{bio})} + \underbrace{f_{res}(z_{res})}_{\approx\epsilon(d,c_{id})} \tag{2}$$

where z_{inv}, z_{exp}, and z_{res} encode the invariant drug features, mechanism-aware biological interactions, and residual cell-identity biases, respectively. During the training phase, we employ a **Conditional Orthogonality Constraint** to force these representations to be statistically independent. Crucially, during the inference phase, the spurious branch $f_{res}(z_{res})$ is discarded. This ensures that the final prediction relies solely on the invariant potency and explicit biological interactions, thereby achieving robust generalization.

3.2 Channel I: Invariant Drug Encoder (E_{inv})

To model the **intrinsic potency** $\mu(d)$ defined in Sec. 3.1, the Invariant Channel extracts structural features that remain consistent across different cellular environments. This context-independent baseline is captured via a dual-pathway architecture.

Graph-level Encoding. We utilize an **EdgeGAT** encoder to process the molecular graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. To capture nuanced chemical motifs, we explicitly integrate bond features $\mathbf{e}_{j,i}$ (e.g., bond type, conjugation) into the attention mechanism. The atom-level update for node i at layer ℓ is formulated as:

$$\mathbf{v}_i^{(\ell+1)} = \sigma\left(\sum_{j\in\mathcal{N}(i)} \alpha_{ij}^{(\ell)} \mathbf{W}_v^{(\ell)} [\mathbf{v}_j^{(\ell)} \parallel \mathbf{e}_{j,i}]\right) \tag{3}$$

where $\alpha_{ij}^{(\ell)}$ is the attention coefficient computed over concatenated node-edge features, and σ is an activation function. These local features are then aggregated via a self-attention pooling layer to yield a permutation-invariant graph embedding $\mathbf{g}_{emb}$.

Global Descriptor. In parallel, 2048-bit Morgan fingerprints are projected through an MLP to generate a descriptor $\mathbf{fp}_{emb}$, capturing long-range functional groups. The final invariant representation z_{inv} is obtained by:

$$z_{inv} = \text{MLP}_{inv}(\mathbf{g}_{emb} \parallel \mathbf{fp}_{emb}) \tag{4}$$

By anchoring z_{inv} to the drug's mean response across cell line (see Sec. 3.6), we ensure it captures the general pharmacological baseline rather than cell-specific variations.

3.3 Channel II: Cell Line-Aware Explicit Encoder (E_{exp})

To model the mechanistic interaction $\tau(x_d, x_c^{bio})$, the Explicit Channel aligns drug structural motifs with cell-line-specific

genomic features. Unlike the global bias isolated in Channel III, this channel captures generalizable patterns.

Cross-Modality Interaction. We utilize a multi-head cross-attention mechanism to identify dependencies between drug atoms and multi-omics features. For a specific omics modality $k \in \{base, opt_1, \ldots, opt_K\}$, let $\mathbf{V}$ be the drug atom embeddings and $\mathbf{x}_k$ be the genomic feature vector. The interaction is captured as:

$$\mathbf{z}_{att}^{(k)} = \text{Attention}(\mathbf{V}\mathbf{W}_Q, \mathbf{x}_k\mathbf{W}_K, \mathbf{x}_k\mathbf{W}_V) \tag{5}$$

where the queries are derived from chemical structures and keys/values from biological features. This identifies specific genomic sites sensitive to particular functional groups.

Modality-Specific Encoding. To facilitate the reconstruction requirement ($\mathcal{L}_{recon}$ in Sec. 3.6) and the counterfactual consistency, we generate a latent code for each observed modality:

$$z_{exp}^{(k)} = \text{MLP}_{exp}(\mathbf{z}_{att}^{(k)} \parallel \mathbf{g}_{emb} \parallel \mathbf{fp}_{emb}) \tag{6}$$

The total explicit contribution is then the masked summation: $f_{exp}(z_{exp}) = \sum_k m_k \cdot f_{exp}^{(k)}(z_{exp}^{(k)})$. By maintaining modality-specific latent paths, the model ensures that the biological semantics of each omics layer are preserved and can be reconstructed, addressing the ambiguity regarding feature fusion.

3.4 Channel III: Residual Bias Encoder (E_{res})

The final component of our decomposition addresses the unobserved cell-specific factors $\epsilon(x_d, c)$, such as experimental batch effects or unmeasured omics. To prevent the model from "hallucinating" these non-generalizable variations into the explicit branch, Channel III acts as a **residual buffer**. It isolates domain-specific noise during training to shield the mechanistic integrity of the z_{exp} representations.

Drug-Conditioned Bias Encoding. We model the residual bias as a joint interaction between drug properties and cell identity. Let $\mathbf{e}_c \in \mathbb{R}^{d_{emb}}$ be a learnable embedding representing the unique signature of cell line c.

Modality-Aware Conditioning. To address the ambiguity of mask incorporation, we explicitly condition the residual encoder on the modality availability mask $\mathbf{m} \in \{0, 1\}^K$. This allows the branch to dynamically adjust its bias estimation based on which omics data are missing. The residual latent code is synthesized as:

$$z_{res} = \text{MLP}_{res}(\mathbf{g}_{emb} \parallel \mathbf{fp}_{emb} \parallel \mathbf{e}_c \parallel \mathbf{m}) \tag{7}$$

where $\parallel$ denotes concatenation. By incorporating $\mathbf{m}$, z_{res} can learn to compensate for missing biological signals (detailed in Sec. 3.5).

To ensure z_{res} captures identity-specific noise rather than generalizable signals, we regularize it with an identity classifier $\mathcal{C}_{res}$ (Sec. 3.6). Crucially, this entire channel is **deactivated during inference**. This force-filters all residual domain biases, ensuring that predictions for unseen cell lines are driven exclusively by invariant and explicit biological mechanisms.

3.5 Missing-Aware Learning via Counterfactual Compensation

To handle incomplete multi-omics data without relying on noise-prone imputation, we propose a **Counterfactual Compensation** mechanism. We reframe modality absence as a dynamic signal migration: when a modality is missing, its biological effect is not "lost" but instead migrates from *explicit interaction features* (z_{exp}) to *implicit residual representations* (z_{res}).

Dynamic Routing Mechanism

We define cellular features as a set of $K + 1$ modalities, where $k = 0$ is the stable base modality (e.g., GEP) and $k \in \{1, \ldots, K\}$ are optional modalities (e.g., CNV, Mutation). Modality availability is controlled by a binary mask $\mathbf{m} \in \{0, 1\}^K$. The drug response prediction is dynamically routed:

$$\begin{aligned}\hat{y} = f_{inv}(z_{inv}) + f_{base}(z_{base}) \\ + \sum_{k=1}^{K} m_k \cdot f_{opt}^{(k)}(z_{exp}^{(k)}) + f_{res}(z_{res}, \mathbf{m})\end{aligned} \tag{8}$$

By conditioning f_{res} on $\mathbf{m}$ (via concatenation as described in Sec. 3.4), the model adaptively adjusts the residual contribution. This architecture ensures that when $m_k = 0$, the system does not zero out the underlying biological contribution but instead shifts its estimation to the residual latent space.

Counterfactual Consistency Learning

To formally train z_{res} to recover missing signals, we introduce the **Counterfactual Consistency** loss ($\mathcal{L}_{cons}$). For a sample where modality k is observed ($m_k = 1$), we artificially simulate a missing state by applying a mask $\mathbf{m}_{\setminus k}$.

We define the **target bias contribution** $\mathcal{T}_k$ under the full-information state as the sum of the implicit residual and the explicit interaction of modality k:

$$\mathcal{T}_k = f_{res}(z_{res}, \mathbf{m}) + f_{opt}^{(k)}(z_{exp}^{(k)}) \tag{9}$$

We then enforce a consistency constraint using Mean Squared Error:

$$\mathcal{L}_{cons} = \sum_{k=1}^{K} m_k \cdot \left\| f_{res}(z_{res}, \mathbf{m}_{\setminus k}) - \text{sg}(\mathcal{T}_k) \right\|_2^2 \tag{10}$$

where $\text{sg}(\cdot)$ denotes the **stop-gradient** operator. This treats $\mathcal{T}_k$ as a stable reference, compelling the residual branch to "learn to compensate" for the missing modality by aligning its partial-information output with the full-information target.

Through this mechanism, knowledge is transferred from the explicit mechanistic branch (f_{opt}) to the implicit residual branch (f_{res}) under simulated missing contexts. Consequently, during inference on incomplete data, GBD-DRP can recover missing biological signals through the residual latent space while shielding the explicit branch from zero-filling artifacts.

3.6 Training Objective and Statistical Constraints

To ensure that the decomposed latent representations ($z_{inv}, z_{exp}, z_{res}$) capture distinct functional aspects of the

drug response, we formulate a joint optimization objective. This composite loss integrates the predictive task with statistical constraints designed to enforce representational fidelity and geometric disentanglement.

1. Primary Prediction ($\mathcal{L}_{task}$) & Mean Anchoring ($\mathcal{L}_{mean}$): The total prediction follows our tri-channel summation: $\hat{y} = f_{inv}(z_{inv}) + f_{exp}(z_{exp}) + f_{res}(z_{res})$. The task loss is $\mathcal{L}_{task} = \frac{1}{N}\sum(y_i - \hat{y}_i)^2$. To anchor z_{inv} as the context-independent potency $\mu(d)$, we introduce $\mathcal{L}_{mean} = \frac{1}{N}\sum(f_{inv}(z_{inv}^{(i)}) - \bar{y}_{d^{(i)}})^2$, where $\bar{y}_{d^{(i)}}$ is the drug-specific average response in the training set. This prevents z_{inv} from absorbing cell-specific variations.

2. Semantic Reconstruction ($\mathcal{L}_{recon}$): To prevent z_{exp} from losing biological meaning during fusion, we require it to reconstruct the original omics features. For each modality k, we employ a decoder $Dec^{(k)}$ (a 2-layer symmetric MLP):

$$\mathcal{L}_{recon} = \sum_k m_k \cdot \|\mathbf{x}_c^{(k)} - Dec^{(k)}(z_{exp}^{(k)})\|_2^2 \quad (11)$$

Beyond semantic grounding, this constraint is critical for **preventing posterior collapse**. Specifically, the conditional orthogonality constraint ($\mathcal{L}_{orth}$) may inadvertently drive z_{exp} to become uninformative to minimize overlap with the residual space. $\mathcal{L}_{recon}$ acts as an information-preservation bottleneck, forcing the mechanistic branch to retain sufficient biological specificity even under strict disentanglement pressure.

3. Residual Identity Prior ($\mathcal{L}_{class}$): To force z_{res} to isolate domain-specific noise (e.g., identity-linked batch effects), we employ an auxiliary classifier $\mathcal{C}_{res}$ (a single linear layer with Softmax) to predict the cell identity c from the residual code:

$$\mathcal{L}_{class} = \text{CrossEntropy}(\mathcal{C}_{res}(z_{res}), c) \quad (12)$$

4. Conditional Orthogonality ($\mathcal{L}_{orth}$): To eliminate redundant information, we minimize the squared cosine similarity between branches. Crucially, the orthogonality between z_{exp} and z_{res} is *conditional* (gated by m_k):

$$\mathcal{L}_{orth} = \cos^2(z_{inv}, z_{res}) + \sum_k m_k \cdot \cos^2(z_{exp}^{(k)}, z_{res}) \quad (13)$$

This gating allows z_{res} to capture information from missing modalities ($m_k = 0$) while prohibiting information leakage when full data is observed.

5. Joint Optimization: The final objective is an end-to-end weighted sum:

$$\begin{aligned}\mathcal{L}_{total} = \mathcal{L}_{task} + \lambda_m\mathcal{L}_{mean} + \lambda_r\mathcal{L}_{recon} + \lambda_c\mathcal{L}_{class} \\ + \lambda_o\mathcal{L}_{orth} + \lambda_s\mathcal{L}_{cons} + \lambda_n\|f_{res}(z_{res})\|^2\end{aligned} \quad (14)$$

where $\mathcal{L}_{cons}$ is the counterfactual consistency loss from Sec. 3.5. The penalty term $\lambda_n\|f_{res}(z_{res})\|^2$ forcing the model to prioritize explicit channels. Hyperparameters λ were determined via grid search; a detailed sensitivity analysis is provided in Section 4.5.

4 Experiment

4.1 Datasets and Experimental Setup

Following the experimental framework of PASO [Wu *et al.*, 2025], we utilized a large-scale pharmacogenomic dataset consolidated from the GDSC and CCLE databases. The final curated dataset comprises 141,222 drug-response pairs, covering 233 unique drugs and 688 cancer cell lines.

Drug Representations
To capture both global and local molecular semantics, we extracted two complementary features for each compound: (1) Molecular Graphs for structural message passing, and (2) Morgan Fingerprints for fixed-length substructure encoding.

Cell Line Features
We integrated three distinct omics modalities to characterize cellular heterogeneity: Gene Expression Profiles (GEP), Copy Number Variations (CNV), and Somatic Mutations (MUT).

Evaluation Protocol
We adopted a 10-fold cross-validation strategy to benchmark performance. To comprehensively assess model robustness, we designed three distinct experimental scenarios:

- **Mixed Setting:** Drug-cell pairs are randomly split into training and testing sets.
- **Cell-Blind Setting:** The dataset is split based on cell line identities, ensuring no cell line in the test set appears during training. This simulates a "precision medicine" scenario for unseen patients.
- **Missing Modality Setting:** We randomly masked a certain ratio of the optional modalities (CNV and MUT), testing the model's ability to maintain performance relying primarily on the stable modality (GEP) and utilize unstable modalities.

4.2 Baseline Methods

To evaluate the performance of our proposed framework, we compare it against a comprehensive set of baselines ranging from classical machine learning algorithms to state-of-the-art deep learning models:

Random Forest (RF): Classical machine learning baselines that utilize concatenated drug features (e.g., fingerprints or encodings) and cell line genomic profiles to model non-linear response patterns via decision tree ensembles and kernel-based regression, respectively.

XGBoost [Chen, 2016]: An optimized gradient boosting decision tree algorithm that iteratively trains weak learners to capture complex, non-linear interactions between high-dimensional omics features and drug descriptors.

PathDSP [Tang and Gottlieb, 2021]: A deep neural network approach that integrates drug molecular fingerprints with pathway-level features derived from gene expression enrichment analysis to predict drug sensitivity.

FA^{2}GL [Xu *et al.*, 2025]: A feature aggregation and graph learning network that captures fine-grained drug-cell interactions via cross-attention and models network-level associations using a Graph Neural Network (GNN) with auxiliary pipe nodes.

PASO [Wu *et al.*, 2025]: A hybrid deep learning model that combines pathway-based multi-omics difference features with drug SMILES sequences, utilizing multi-scale CNNs, Transformers, and attention mechanisms to capture molecular interactions.

4.3 Results

Method	RMSE ↓	Pearson (%) ↑	R^2 (%) ↑
RF	1.2171 ± 0.0103	90.06 ± 0.25	81.10 ± 0.44
XGB	1.1611 ± 0.0066	91.00 ± 0.16	82.80 ± 0.29
PathDSP	1.0499 ± 0.0154	92.82 ± 0.15	83.65 ± 1.16
FA^{2}GL	1.0037 ± 0.0103	93.35 ± 0.21	86.01 ± 0.26
PASO	0.9400 ± 0.0081	94.25 ± 0.11	88.38 ± 0.21
GBD-DRP	**0.9288 ± 0.0123**	**94.38 ± 0.17**	**88.99 ± 0.30**

Table 1: Performance comparison with state-of-the-art baselines in the Mixed Setting.

Mixed Cross Validation Results

Table 1 presents the quantitative comparison of GBD-DRP against baseline methods. GBD-DRP consistently outperforms all competing methods across all three evaluation metrics. Specifically, our model achieves the lowest RMSE of 0.9288 and the highest R^2 of 88.99%, surpassing the second-best method, PASO (RMSE: 0.9400, R^2: 88.38%). While PASO already establishes a strong baseline by utilizing pathway-based difference features and transformer encoders, GBD-DRP further reduces the prediction error by 1.2% in terms of RMSE. This performance gain is attributed to our explicit bias disentanglement strategy, which effectively isolates spurious domain noise from true biological signals.

Cell Blind Cross Validation Results

Method	RMSE ↓	Pearson (%) ↑	R^2 (%) ↑
RF	1.4826±0.0774	85.74±0.98	73.21±1.16
XGB	1.3721±0.0527	87.33±0.76	76.99±1.55
PathDSP	1.3513±0.1178	87.68±1.42	76.90±1.31
FA^{2}GL	1.3462±0.0795	87.74±1.14	77.83±0.98
PASO	1.3457±0.1778	87.75±1.21	77.25±1.77
GBD-DRP	**1.2872±0.0631**	**88.06±0.91**	**77.78±1.59**

Table 2: Performance comparison in the **Cell-Blind Setting**. GBD-DRP maintains superior generalization compared to baselines when facing unseen domains.

As summarized in Table 2, although the absolute improvement in the Cell-Blind setting is more modest than in the Mixed setting, GBD-DRP's advantages lie in its **predictive stability and noise resilience**.

Compared to PASO, GBD-DRP reduces the RMSE from 1.3166 to 1.2872 while significantly lowering the prediction variance (±0.0631 vs. ±0.1778). Since PASO lacks an explicit mechanism to isolate domain-specific biases, its attention mechanism inadvertently captures **non-generalizable experimental noise** (e.g., batch effects or lab-specific artifacts) and treats them as predictive biological signals.

In contrast, GBD-DRP addresses this through its **Residual Disentanglement** strategy. By delegating cell-specific residual variations to the dedicated buffer (z_{res}), our model "filters" the information entering the explicit branch (z_{exp}). This ensures that the learned "lock-and-key" interactions are grounded in invariant biological mechanisms rather than spurious correlations. Consequently, when transitioning to unseen cell lines, GBD-DRP maintains a more robust performance by relying on these filtered, generalizable patterns, whereas PASO's performance fluctuates due to its entanglement with the unobserved biases of the training domains.

Missing Modality Cross Validation Results

In clinical practice, multi-omics profiles are frequently incomplete due to cost or sample degradation. To evaluate GBD-DRP's robustness under such data-constrained scenarios, we simulate a challenging environment where 50% of the optional modalities (CNV and MUT) are randomly masked during inference.

To ensure a fair and rigorous comparison, we implemented three representative strategies for the baseline models that lack inherent missing-aware mechanisms: (1) **Mean Imputation**: Missing features are replaced by the global mean of the respective modality in the training set. (2) **Zero Filling**: Missing entries are set to zero, a common practice in deep learning to neutralize missing inputs. (3) **Feature Exclusion**: The models are restricted to using only the stable modality (GEP), effectively bypassing the unobserved features.

Method	RMSE ↓	Pearson (%) ↑	R^2 (%) ↑
Strategy 1: Mean Imputation			
RF	1.2042 ± 0.0115	90.03 ± 0.31	80.84 ± 0.52
XGB	1.1605 ± 0.0084	91.05 ± 0.22	82.80 ± 0.35
PathDSP	1.1210 ± 0.0182	91.54 ± 0.20	82.11 ± 1.05
FA^{2}GL	1.0931 ± 0.0122	92.21 ± 0.25	83.23 ± 0.41
PASO	1.0876 ± 0.0165	92.10 ± 0.21	84.05 ± 0.39
Strategy 2: Zero Filling			
RF	1.2191 ± 0.0132	90.11 ± 0.42	81.10 ± 0.61
XGB	1.2243 ± 0.0091	89.21 ± 0.28	79.88 ± 0.44
PathDSP	1.1567 ± 0.0211	90.88 ± 0.25	80.56 ± 1.20
FA^{2}GL	1.1023 ± 0.0151	92.46 ± 0.22	83.88 ± 0.48
PASO	1.0542 ± 0.0143	92.88 ± 0.18	85.12 ± 0.33
Strategy 3: Feature Exclusion (GEP only)			
RF	1.2171 ± 0.0103	90.06 ± 0.25	81.10 ± 0.44
XGB	1.1611 ± 0.0066	91.00 ± 0.16	82.80 ± 0.29
PathDSP	1.0499 ± 0.0154	92.82 ± 0.15	83.65 ± 1.16
FA^{2}GL	1.0037 ± 0.0103	93.35 ± 0.21	86.01 ± 0.26
PASO	0.9882 ± 0.0120	93.63 ± 0.12	87.09 ± 0.25
GBD-DRP	**0.9322 ± 0.0179**	**94.22 ± 0.22**	**88.72 ± 0.35**

Table 3: Performance comparison under 50% missing modality setting. GBD-DRP demonstrates superior resilience by adaptively compensating for unobserved signals.

As shown in Table 5, GBD-DRP demonstrates a decisive advantage, achieving an RMSE of 0.9322 and a Pearson correlation of 94.20%, significantly outperforming all baseline methods. Notably, the performance of GBD-DRP under 50% modality missingness remains remarkably close to its full-modality performance, whereas baselines experience a sharper decline.

Intriguingly, baselines perform better with *Feature Exclusion* than naive *Mean/Zero Imputation*, suggesting that input-level infilling introduces detrimental "phantom signals" and

overfitting to imputation artifacts. In contrast, GBD-DRP leverages **Counterfactual Consistency** to adaptively compensate for missing explicit signals in the latent space. By dynamically reallocating unobserved variance to the residual channel (z_{res}), GBD-DRP "recovers" the biological essence of missing modalities without noise injection, ensuring superior resilience in incomplete data environments.

4.4 Ablation Study and Sensitivity Analysis

To rigorously justify the necessity of GBD-DRP's components, we conduct a fine-grained ablation study across both *Mixed* and *Cell-Blind* settings. We evaluate four variants: (1) **w.o. Attn**: Replacing cross-attention in Channel II with concatenation; (2) **w.o.** $\mathcal{L}_{orth}$: Removing orthogonality constraints; (3) **w.o.** $\mathcal{L}_{cons}$: Disabling the counterfactual consistency mechanism; (4) **w.o. Res**: Removing the entire Residual Channel and its associated losses ($\mathcal{L}_{class}, \mathcal{L}_{orth}$).

	Mixed Setting			Cell-Blind Setting		
Method	**RMSE** ↓	**Prsn**% ↑	R^2% ↑	**RMSE** ↓	**Prsn**% ↑	R^2% ↑
w.o. Attn	1.0037	93.16	86.01	1.4050	86.48	74.06
w.o. $\mathcal{L}_{orth}$	0.9652	93.84	87.52	1.3630	87.13	75.58
w.o. Res	0.9882	93.63	87.09	1.4488	86.08	72.42
GBD-DRP	**0.9288**	**94.38**	**88.99**	**1.2872**	**88.06**	**77.78**

Table 4: Ablation results across different generalization scenarios.

Method	**RMSE** ↓	**Pearson** (%) ↑	R^2 (%) ↑
w.o. $\mathcal{L}_{cons}$	0.9475	93.81	88.23
GBD-DRP	**0.9322**	**94.22**	**88.72**

Table 5: Ablation results under 50% missing modality setting.

Mechanistic Interaction via Attention: The *w.o. Attn* variant yields inferior results, particularly in R^2. This suggests that simple concatenation fails to capture the "lock-and-key" dependencies between chemical motifs and genomic sites, whereas attention provides the necessary granularity for modeling drug sensitivity.

Generalization to Unseen Domains (Cell-Blind): As shown in Table 4, the performance of the *w.o. Res* variant collapses significantly in the Cell-Blind setting, with RMSE increasing from. This confirms that without a dedicated residual channel to sequester domain-specific noise, non-generalizable biases leak into the predictive branches, severely hindering zero-shot transfer.

Fine-grained Analysis of the Residual Branch: The degradation in the *w.o.* $\mathcal{L}_{orth}$ variant indicates that representational disentanglement is not achieved by structural separation alone; explicit geometric constraints are required to prevent information redundancy.

Counterfactual Consistency ($\mathcal{L}_{cons}$): Removing $\mathcal{L}_{cons}$ under 50% missingness increases RMSE from 0.9337 to 0.9475 (Table 5). This confirms $\mathcal{L}_{cons}$ is vital for **signal compensation**: it compels the residual branch (z_{res}) to actively recover functional information lost from omitted modalities. These results validate that our mechanism successfully reallocates predictive logic from the explicit branch (z_{exp}) to the residual buffer during data depletion, ensuring a robust "information safety net."

4.5 Implementation Details

Our model is implemented using PyTorch and Deep Graph Library (DGL). The choice of hyper-parameter are based on grid search and shown in Table 6

Category	**Hyperparameter**	**Symbol**	**Value**
Optimization	Learning Rate	η	1×10^{-3}
	Batch Size	B	1024
Architecture	Embedding Dimension	d_{emb}	512
	Dropout Rate	p	0.2
Loss Weights	Task Loss	λ_{task}	1.0
	Orthogonality Loss	λ_{orth}	1.0
	Residual Norm Loss	λ_{res}	0.5
	Classification Loss	λ_{class}	0.1
	Drug Mean Constraint	λ_{mean}	0.1
	Reconstruction Loss	λ_{recon}	0.01

Table 6: Hyperparameter settings used in our experiments.

5 Conclusion

In this study, we presented **GBD-DRP**, a novel tri-channel disentanglement framework designed to enhance the generalizability and robustness of drug response prediction. By explicitly decomposing the latent space into invariant pharmacological effects, explicit biological interactions, and residual domain-specific biases, our model effectively mitigates the "shortcut learning" common in conventional DRP methods. Through the integration of a counterfactual consistency mechanism, GBD-DRP not only isolates spurious experimental noise but also provides a principled solution for the missing modality problem, enabling robust "signal compensation" without relying on data imputation. Extensive experiments demonstrate that GBD-DRP consistently achieves state-of-the-art performance in both Cell-Blind generalization and missing modality data scenarios. Future research can focus on extending the current domain generalization (DG) paradigm. While this work primarily addresses cellular-level biases, we aim to investigate a composite DG framework that treats both cell lines and drugs as distinct yet interacting domains. By modeling "Drug-Domain" shifts—such as those arising from chemical scaffold variations or novel drug classes—we anticipate developing a more holistic system capable of achieving simultaneous zero-shot generalization for both unseen patients and unseen therapeutic compounds. Additionally, integrating single-cell multi-omics data into this disentangled architecture could further refine the precision of mechanistic explanations at the sub-clonal level.

Ethical Statement

There are no ethical issues.

Acknowledgements

This work was supported by the Science and Technology Commission of Shanghai Municipality (grants No. 24510714300), and the Shanghai Municipal Science and Technology Major Project, China (Grant No. 2021SHZDZX0102).

References

[Baptista *et al.*, 2021] Delora Baptista, Pedro G Ferreira, and Miguel Rocha. Deep learning for drug response prediction in cancer. *Briefings in bioinformatics*, 22(1):360–379, 2021.

[Barretina *et al.*, 2012] Jordi Barretina, Giordano Caponigro, Nicolas Stransky, Kavitha Venkatesan, Adam A Margolin, Sungjoon Kim, Christopher J Wilson, Joseph Lehár, Gregory V Kryukov, Dmitriy Sonkin, et al. The cancer cell line encyclopedia enables predictive modelling of anticancer drug sensitivity. *Nature*, 483(7391):603–607, 2012.

[Chen, 2016] Tianqi Chen. Xgboost: A scalable tree boosting system. *Cornell University*, 2016.

[Codicè *et al.*, 2025] Francesco Codicè, Corrado Pancotti, Cesare Rollo, Yves Moreau, Piero Fariselli, and Daniele Raimondi. The specification game: rethinking the evaluation of drug response prediction for precision oncology. *Journal of Cheminformatics*, 17(1):33, 2025.

[Jiang *et al.*, 2022] Likun Jiang, Changzhi Jiang, Xinyu Yu, Rao Fu, Shuting Jin, and Xiangrong Liu. Deeptta: a transformer-based model for predicting cancer drug response. *Briefings in bioinformatics*, 23(3), 2022.

[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(Supplement_2):i911–i918, 2020.

[Nguyen *et al.*, 2021] Tuan Nguyen, Giang TT 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, 2021.

[Partin *et al.*, 2023] Alexander Partin, Thomas S Brettin, Yitan Zhu, Oleksandr Narykov, Austin Clyde, Jamie Overbeek, and Rick L Stevens. Deep learning methods for drug response prediction in cancer: predominant and emerging trends. *Frontiers in medicine*, 10:1086097, 2023.

[Peng *et al.*, 2019] Xingchao Peng, Zijun Huang, Ximeng Sun, and Kate Saenko. Domain agnostic learning with disentangled representations. In *International conference on machine learning*, pages 5102–5112. PMLR, 2019.

[Tang and Gottlieb, 2021] Yi-Ching Tang and Assaf Gottlieb. Explainable drug sensitivity prediction through cancer pathway enrichment. *Scientific reports*, 11(1):3128, 2021.

[Wu *et al.*, 2025] Yang Wu, Ming Chen, and Yufang Qin. Anticancer drug response prediction integrating multiomics pathway-based difference features and multiple deep learning techniques. *PLOS Computational Biology*, 21(3):e1012905, 2025.

[Xu *et al.*, 2025] Kaiyi Xu, Minhui Wang, Xin Zou, Chengfu Ji, Hua Zhou, and Chang Tang. Cancer–drug response prediction via feature aggregation and association graph learning. *Engineering Applications of Artificial Intelligence*, 151:110671, 2025.

[Yang *et al.*, 2012] Wanjuan Yang, Jorge Soares, Patricia Greninger, Elena J Edelman, Howard Lightfoot, Simon Forbes, Nidhi Bindal, Dave Beare, James A Smith, I Richard Thompson, et al. Genomics of drug sensitivity in cancer (gdsc): a resource for therapeutic biomarker discovery in cancer cells. *Nucleic acids research*, 41(D1):D955–D961, 2012.

[Zhou *et al.*, 2022] Kaiyang Zhou, Ziwei Liu, Yu Qiao, Tao Xiang, and Chen Change Loy. Domain generalization: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 45(4):4396–4415, 2022.