SEPARATING RESPONSE RANKING AND CLASSIFICATION FOR BULK-TO-SINGLE-CELL DRUG RESPONSE PREDICTION

Anonymous authors
Paper under double-blind review

ABSTRACT

Single-cell drug response prediction helps characterize heterogeneous responses to treatment, providing a finer-grained view of drug sensitivity than population-level measurements. Existing methods transfer drug response supervision from bulk data to single-cell data through domain adaptation. However, transferring drug response information requires both preserving the relative sensitivity of cells and determining an appropriate classification boundary without target response labels. We propose scRA-DRP, a framework that separately adapts response ranking and a classification reference. Starting from a shared domain-adversarial initialization, the ranking branch aligns target features with local source regions constructed using drug response labels. The classification branch uses weighted domain adaptation to adapt the classification rule to the target population and provide a reference number of sensitive cells. This count determines how many top-ranked cells are assigned to the sensitive class, preserving the response ordering. Experiments across 18 drug–single-cell dataset pairs show higher mean performance in both response ranking and classification compared with the evaluated baselines. The ablation studies further support the complementary roles of regional alignment and the separately adapted classification reference.

1 INTRODUCTION

Cells within the same cancer cell population can exhibit different responses to a drug. Characterizing this variation can help identify potentially resistant subpopulations and provide insights into drug action and resistance mechanisms (Shaffer et al., 2017). Single-cell RNA sequencing offers a view of the transcriptional states underlying this heterogeneity, but drug response measurements paired with individual cell transcriptomes remain limited (Ma et al., 2026). In contrast, bulk drug response resources such as GDSC (Yang et al., 2013) provide gene expression profiles and matched drug response measurements across cancer cell lines, offering extensive supervision for response prediction. A practical approach is therefore to transfer drug response knowledge learned from these bulk data to individual cells. Methods such as scDEAL (Chen et al., 2022) and SCAD (Zheng et al., 2023) use this supervision to predict single-cell drug responses.

Bulk-to-single-cell transfer requires adapting to differences in gene expression distributions while applying population-level drug response supervision to individual cells. Differences in cell states, population composition, and measurement technologies further complicate this transfer (Zheng et al., 2023). Existing methods address these differences through distribution alignment, adversarial learning, and structural constraints. For example, scDEAL (Chen et al., 2022) combines distribution matching with target-cluster regularization, whereas DAGFormer (Yan et al., 2025) uses cellular graphs to model within-domain relationships. Domain adaptation theory shows that aligning marginal feature distributions alone is insufficient to guarantee target-domain discrimination (Zhao et al., 2019). Moreover, ranking quality and probability calibration involve distinct statistical requirements (Menon et al., 2012), and discriminative response scores do not directly guarantee accurate classification decisions in the target domain. Even when the relative response ordering of cells is preserved, shifts in the score distribution can reduce the applicability of a source-domain decision threshold. Bulk-to-single-cell drug response prediction must therefore address both the transfer of response ranking and the determination of a target-domain decision boundary.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

Transferring response ranking requires preserving drug-response-related discriminative structure while adapting to expression distribution differences, so that the relative sensitivity of target cells can be distinguished. Determining a decision boundary additionally requires identifying where to divide the target response ranking into sensitive and resistant cells. Local discriminative structure supports comparisons between cells but does not directly determine the number of sensitive cells in the target population. A population-level classification reference can, in turn, constrain the size of the sensitive set, while its membership depends on response ordering. This distinction motivates adapting response scores and classification references separately and combining them at prediction time to make cell-level decisions.

Building on these considerations, we propose **scRA-DRP**, a framework that separately adapts response ranking and a classification reference for bulk-to-single-cell drug response prediction. Starting from a domain-adapted initialization and using the same source supervision, the ranking branch transfers local discriminative structure through regional alignment with source drug-response regions. The classification branch uses the original empirical source distribution as the reference for domain alignment and incorporates target-specific feature statistics to support classification-rule transfer to the target population. The former provides the response ordering of cells, and the latter supplies a population-level sensitive-cell count reference. This reference determines the cutoff in the response ranking, yielding cell-level sensitive or resistant predictions.

Our contributions are as follows:

- We propose scRA-DRP, which separately adapts response ranking and a classification reference, combining them via ranking-guided decisions for bulk-to-single-cell drug response prediction.
- We develop a bidirectional regional alignment objective that uses source response labels to guide target feature adaptation while allowing variation within local regions.
- Experiments on 18 drug–single-cell dataset pairs show higher mean performance on all four metrics, with ablations supporting the contributions of regional alignment and separate adaptation.

2 Related Work

2.1 Drug Response Transfer to Single Cells

Single-cell drug response prediction encompasses unsupervised domain adaptation, semi-supervised domain adaptation, few-shot learning, and domain generalization. For unsupervised bulk-to-single-cell adaptation, bulk response labels and unlabeled target expression profiles are used to transfer drug-response relationships across expression distributions. scDEAL (Chen et al., 2022) and SCAD (Zheng et al., 2023) address this transfer through distribution matching and domain-adversarial learning, respectively. Subsequent methods extend this approach through additional sources and richer representations: scAdaDrug (Liu et al., 2025) uses adaptive feature weighting for multi-source domain adaptation; DAGFormer (Yan et al., 2025), DTLCDR (Xu et al., 2025), and scGSDR (Huang et al., 2025) introduce cellular graphs, multilevel feature alignment, and gene-pathway semantics, respectively. Pretrained representations provide another basis for transfer: single-cell foundation models such as scGPT (Cui et al., 2024) and scFoundation (Hao et al., 2024) learn expression representations for downstream tasks, while scATD (Zhou et al., 2025) combines pretrained single-cell representations with style-based adaptation for drug response prediction.

When target drug response labels are available, SSDA4Drug (Huang & Liu, 2025) uses limited target labels for semi-supervised adaptation, while SCREP (Ge et al., 2025) combines bulk meta-pretraining with few-shot fine-tuning on labeled single cells. FourierDrug (Song et al., 2026) further explores domain generalization without access to target data during training. scXDR (Qi et al., 2026) instead transfers between single-cell datasets, using labeled source cells to predict responses in an unlabeled target dataset. DREEP (Pellecchia et al., 2023) predicts drug responses by performing enrichment analysis of gene sets derived from individual cells against gene lists ranked by their expression correlations with drug sensitivity. Our method retains unsupervised bulk-to-single-cell adaptation and focuses on separately adapting response ranking and a classification reference.

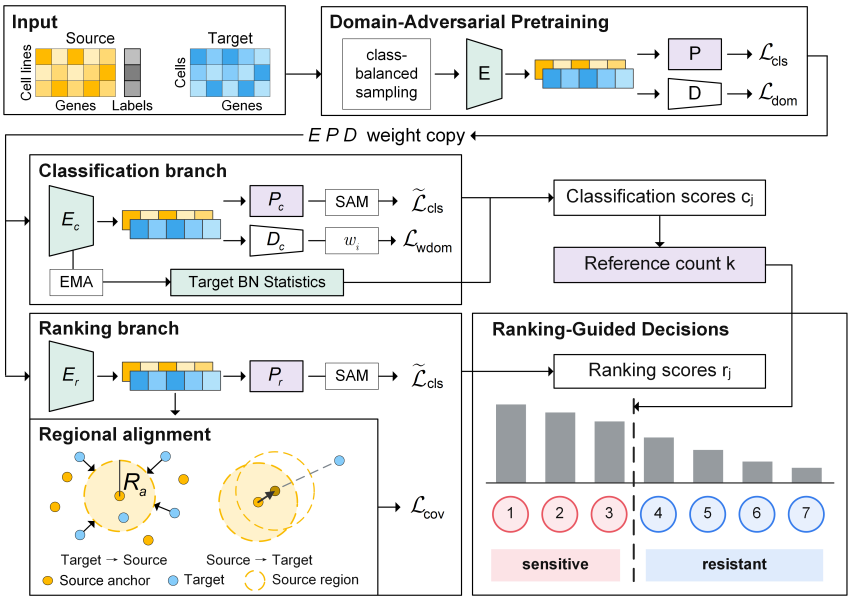

Figure 1: Overview of scRA-DRP.

2.2 Response Ranking and Classification Decisions

In score-based binary classification, continuous scores support both comparisons of relative class propensity and class assignment, but a good ranking does not itself determine an appropriate classification threshold. Menon et al. (2012) address general binary probability estimation by first learning scores with a ranking loss and then applying isotonic regression to estimate class probabilities. Related requirements also arise in search and retrieval. In video search, SBCR (Zhang et al., 2024) uses separate ranking and probability calibration networks, preserving the ranking through a monotonic map and gradient isolation. TMP Adapter (Posokhov et al., 2025) retains the retrieval ranking and calibrates scores to improve threshold consistency when filtering results. These studies provide practical precedents for treating score learning separately from subsequent calibration or truncation.

Applying this approach to prediction across domains also requires accounting for how distribution shift affects calibration. CPCS (Park et al., 2020) combines domain adaptation with importance weighting to calibrate predictions, while TransCal (Wang et al., 2020) fixes an adapted model and learns a transferable temperature using source labels and unlabeled target data. In bulk-to-single-cell drug response prediction, changes in expression distributions and the unit of supervision also require translating cell-line population response relationships into sensitive and resistant assignments for individual cells (Zheng et al., 2023). Target response labels are unavailable for directly adjusting these assignments on a target validation set. Motivated by this setting, we separately adapt response ranking and a population-level classification reference to guide cell-level drug response classification in the unlabeled target domain.

3 Method

3.1 Problem Formulation and Overview

For a given drug, the source domain consists of labeled bulk expression profiles, $\mathcal{D}_s = \{(x_i^s, y_i^s)\}_{i=1}^{n_s}$, and the target domain consists of unlabeled single-cell expression profiles, $\mathcal{X}_t = \{x_j^t\}_{j=1}^{n_t}$. Here, $x \in \mathbb{R}^p$ contains the p genes shared by both domains, with $y = 1$ indicating sensitivity and $y = 0$ indicating resistance. Our goal is to assign a response score r_j and a class label $\widehat{y}_j$ to each target cell. We use a transductive setting, adapting to the unlabeled expression profiles of the same target cells for which predictions are made.

To obtain response scores and class labels, scRA-DRP separately adapts response ranking and a classification reference, as shown in Figure 1. Two branches share a domain-adversarial initialization, then undergo separate adaptation. Ranking-guided decisions use the classification reference to determine how many top-ranked cells are labeled sensitive, with the remaining cells labeled resistant. For simplicity, the regional alignment illustration shows only one response class; the same procedure applies to both classes. The following subsections detail each component.

3.2 Domain-Adversarial Pretraining

To globally align source and target feature distributions while retaining drug response information, we use domain-adversarial pretraining to establish a common initialization for the ranking and classification branches. Specifically, we jointly train an expression encoder E, a drug response predictor P, and a domain discriminator D using labeled source and unlabeled target data. Let Q_s be the class-balanced source sampling distribution, and ℓ denote binary cross-entropy. The source classification loss is

$$\mathcal{L}_{\text{cls}} = \mathbb{E}_{(x,y)\sim Q_s}\big[\ell(P(E(x)), y)\big]. \tag{1}$$

For source and target batches B_s and B_t, the domain discrimination loss is

$$\mathcal{L}_{\text{dom}} = \frac{1}{|B_s| + |B_t|}\left[\sum_{i\in B_s} \ell(D(E(x_i^s)), 1) + \sum_{j\in B_t} \ell(D(E(x_j^t)), 0)\right]. \tag{2}$$

The pretraining loss is

$$\mathcal{L}_{\text{pre}} = \mathcal{L}_{\text{cls}} + \alpha\mathcal{L}_{\text{dom}}, \tag{3}$$

where α controls the domain loss weight. The predictor and discriminator minimize their respective losses, while a gradient reversal layer reverses the domain-loss gradient for the encoder (Ganin et al., 2016). The pretrained encoder E, predictor P, and discriminator D are copied to initialize the ranking and classification branches, which are subsequently adapted independently.

3.3 Source-Supervised Learning

The ranking and classification branches (r, c) both retain source drug response supervision during adaptation and draw source samples from the class-balanced distribution Q_s; to improve stability across random seeds, we optimize their source classification losses with SAM (Foret et al., 2021). For branch $b \in \{r, c\}$, let θ_b denote its encoder and predictor parameters, and δ_b a parameter perturbation computed from the source classification gradient. The source objective for each update is

$$\widetilde{\mathcal{L}}_{\text{cls},b}(\theta_b) = \mathcal{L}_{\text{cls}}\big(\theta_b + \text{sg}(\delta_b)\big). \tag{4}$$

Here, $\text{sg}(\cdot)$ treats the perturbation as constant during differentiation. SAM applies only to the source classification term; branch-specific transfer terms are evaluated at the original parameters. Details of the perturbation, gradient update, and batch-normalization handling are provided in Appendix A.1.

3.4 Response Ranking

The ranking branch aims to order target cells by drug sensitivity using response information from labeled source samples. Global domain alignment reduces distributional discrepancies between domains but does not explicitly model local drug-response structure. We therefore construct local regions around representative source samples from each response class and encourage target-cell features to align with these regions. This regional constraint allows feature differences within each region without requiring target features to coincide with individual source anchors.

Region construction. We construct regions by extracting normalized source features, selecting anchors within each response class, and setting each radius to a scaled distance from the anchor to its nearest opposite-class source sample.

Using the encoder E trained during domain-adversarial pretraining, we extract source-domain features and apply L_2 normalization:

$$z(x) = E(x)/\max(\|E(x)\|_2, \epsilon_z), \tag{5}$$

where ϵ_z is a small positive constant that prevents division by zero.

For each response class $y \in \{0, 1\}$, M anchors are selected in the normalized source feature space, giving $A = 2M$ anchors in total. The anchor closest to the normalized class centroid is chosen first, and the remaining anchors are obtained through iterative farthest-point sampling, which maximizes the distance to the nearest existing anchor.

Each selected anchor defines the center of a local region. For anchor a with expression x_a^s and label y_a, the distance d_a to the nearest opposite-class source sample is scaled by κ to determine the radius:

$$R_a = \kappa d_a = \kappa \min_{i:y_i^s \neq y_a} \|z(x_a^s) - z(x_i^s)\|_2, \tag{7}$$

where κ controls the region size relative to the nearest opposite-class distance. During subsequent training, the anchor samples and radii remain fixed, while their feature coordinates are recomputed with the current encoder.

Regional alignment. Alignment combines two complementary constraints: the target-to-source loss encourages target features to enter source regions, while the source-to-target loss encourages source region centers to move toward their nearest target features in the current batch.

Let z_a^s and z_j^t denote the current source anchor and target features, respectively. For each target feature, the target-to-source loss penalizes its minimum distance outside the source regions. For the current target batch B_t, this loss is

$$\mathcal{L}_{t\to s} = \frac{1}{|B_t|} \sum_{j \in B_t} \min_{a=1,\ldots,A} \left[\|z_j^t - \mathrm{sg}(z_a^s)\|_2 - R_a\right]_+^2, \tag{6}$$

where $[u]_+ = \max(u, 0)$. This penalty is zero when a target feature lies within any source region. In the reverse direction, each region is compared with its nearest target feature in the batch. The source-to-target loss penalizes the distance by which that feature falls outside the region:

$$\mathcal{L}_{s\to t} = \frac{1}{A} \sum_{a=1}^{A} \left[\min_{j \in B_t} \|z_a^s - \mathrm{sg}(z_j^t)\|_2 - R_a\right]_+^2. \tag{7}$$

The target-to-source loss therefore propagates gradients only through target features, while the source-to-target loss propagates gradients only through source anchor features. Both losses jointly update the shared encoder in the same training step. Averaging the two directional losses gives the regional alignment objective:

$$\mathcal{L}_{\mathrm{cov}} = \frac{1}{2}\left(\mathcal{L}_{t\to s} + \mathcal{L}_{s\to t}\right). \tag{8}$$

The target-to-source loss can be zero even when target features occupy only a few source regions. The reverse loss penalizes regions containing no target features from the current batch, encouraging alignment across the selected regions. This assumes that each selected source region represents a response pattern present in the target population.

Training and prediction. The ranking branch combines source supervision with the regional alignment loss:

$$\mathcal{L}_r = \widetilde{\mathcal{L}}_{\mathrm{cls},r} + \mathcal{L}_{\mathrm{cov}}. \tag{9}$$

For prediction, the ranking branch assigns each target cell a drug response score:

$$r_j = P_r\big(E_r(x_j^t)\big). \tag{10}$$

3.5 Classification Reference

The classification branch predicts cell responses to provide a reference number of sensitive cells. It performs domain-adversarial adaptation with source sample weights in the domain loss.

Weighted domain adaptation. Class-balanced source sampling gives sensitive and resistant samples equal emphasis during source-supervised learning, but also changes the source distribution presented to the domain discriminator. To separate class-balanced supervision from the choice of source distribution for alignment, we weight source samples in the classification branch's domain loss by the ratio of their original class proportions to their sampling probabilities. This retains the original empirical source distribution as the alignment reference while maintaining balanced supervision for drug response prediction.

Specifically, let $\pi_s(y)$ denote the natural class proportion in the source training set, and $q_s(y)$ the class probability under the sampling distribution Q_s. For source sample i in a batch, we define

$$\widetilde{w}_i = \frac{\pi_s(y_i^s)}{q_s(y_i^s)}, \qquad w_i = \frac{|B_s|\widetilde{w}_i}{\sum_{l\in B_s}\widetilde{w}_l}. \tag{11}$$

Branch training uses the weighted domain loss:

$$\mathcal{L}_{\mathrm{wdom}} = \frac{1}{|B_s|+|B_t|}\left[\sum_{i\in B_s} w_i \ell(D_c(E_c(x_i^s)),1) + \sum_{j\in B_t}\ell(D_c(E_c(x_j^t)),0)\right]. \tag{12}$$

These weights are normalized within each batch and applied only to the source-domain term of the domain loss.

Training and prediction. The classification branch combines SAM source supervision with weighted domain adaptation:

$$\mathcal{L}_c = \widetilde{\mathcal{L}}_{\mathrm{cls},c} - \beta\mathcal{L}_{\mathrm{wdom}}. \tag{13}$$

Here, β controls the strength of domain adaptation relative to source supervision. This objective is minimized with respect to the classification encoder and response predictor, while the domain discriminator minimizes $\mathcal{L}_{\mathrm{wdom}}$. During training, the classification encoder maintains an auxiliary set of target-specific batch-normalization means and variances, updated by an exponential moving average over target batches. These statistics are used only for prediction, not for the forward passes used to optimize the classification branch:

$$c_j = P_c\big(E_c^{\mathrm{BN}}(x_j^t)\big). \tag{14}$$

Here, E_c^{BN} denotes the classification encoder evaluated with the auxiliary target-specific batch-normalization statistics.

3.6 Ranking-Guided Decisions

Ranking-guided decisions use the classification scores c_j to estimate the number of sensitive cells and the ranking scores r_j to select those cells. The classification branch provides the reference count

$$k_0 = \sum_{j=1}^{n_t}\mathbf{1}\{c_j \geq 0.5\}. \tag{15}$$

The selected count is $k = k_0$ unless the reference cutoff splits a group of cells with identical ranking scores. In that case, k is chosen to be as close as possible to k_0 while assigning the entire group to the same class. The complete selection rule is provided in Appendix A.2.

The top k cells ranked by r_j are then classified as sensitive, and the remaining cells as resistant:

$$\widehat{y}_j = \mathbf{1}\{\mathrm{rank}_{\downarrow}(r_j) \leq k\}. \tag{16}$$

Here, $\mathrm{rank}_{\downarrow}(r_j)$ denotes the position of cell j when ranking scores are sorted in descending order, starting at 1.

4 EXPERIMENTS

4.1 EXPERIMENTAL SETUP

Datasets. We evaluate bulk-to-single-cell drug response prediction using labeled GDSC expression profiles (Yang et al., 2013) and 18 drug–single-cell dataset pairs: 9 from GEO and 9 EpiSen pairs covering JHU006 and SCC47 (Kinker et al., 2020). Models access unlabeled target expression during adaptation, while target response labels are reserved for evaluation. Data sources and preprocessing procedures are provided in Appendix A.3.

Baselines. We compare with SCAD (Zheng et al., 2023), scDEAL (Chen et al., 2022), DAGFormer (Yan et al., 2025), scAdaDrug (Liu et al., 2025), FourierDrug (Song et al., 2026), and DREEP (Pellecchia et al., 2023). We also evaluate scFoundation (Hao et al., 2024) and scGPT (Cui et al., 2024), foundation models pretrained on large-scale single-cell transcriptomic data, by using their frozen representations as inputs to SCAD. This follows the integration in the original scFoundation study and applies the same downstream framework to scGPT. We include CatBoost (Prokhorenkova et al., 2018) as a source-only baseline to assess transfer performance without target-domain adaptation. Baseline configurations are documented in Appendix A.4.

Evaluation. We report AUROC, AUPR (average precision), Weighted-F1 (W-F1), and accuracy. W-F1 averages the F1 scores of the sensitive and resistant classes, weighted by their respective sample counts. Metrics are averaged over three seeds (42, 43, 44) within each dataset, then equally across datasets.

Implementation. Domain-adversarial pretraining runs for 5 epochs, followed by 15 epochs of independent adaptation for each branch. We use Adagrad with learning rate 0.001 and batch size 32, setting $M = 4$, $\kappa = 0.25$, $\alpha = \beta = 0.2$, and the SAM perturbation radius to 0.05. The same hyperparameter configuration is used across all datasets without dataset-specific tuning.

4.2 SINGLE-CELL DRUG RESPONSE PREDICTION RESULTS

Table 1 shows that scRA-DRP achieves the highest mean AUROC, AUPR, Weighted-F1, and accuracy on both GEO and EpiSen datasets. SCAD is the strongest baseline for overall ranking performance, whereas DAGFormer leads in overall classification performance; scRA-DRP improves both, attaining an overall AUPR of 0.8186 and Weighted-F1 of 0.7652. On GEO datasets, scRA-DRP achieves modest improvements in AUROC and AUPR over SCAD while substantially increasing Weighted-F1 and accuracy by 11.7 and 6.6 percentage points, respectively. The gains over SCAD are larger on EpiSen datasets across all four metrics. scDEAL has low mean performance under the evaluated configuration and substantial variation across random seeds. scAdaDrug is evaluated using its two-source configuration. DREEP uses GDSC-derived ranked gene lists and covers 9 GEO and 8 EpiSen datasets because no corresponding gene list is available for Cetuximab. On these same 17 datasets, scRA-DRP also outperforms DREEP on all four mean metrics.

Among the two foundation-model baselines, scFoundation outperforms scGPT on all four overall metrics, although scGPT has higher AUROC and AUPR on GEO datasets. Both underperform SCAD in each dataset group, which may reflect a mismatch between general expression patterns learned during pretraining and drug-specific response features.

Finally, CatBoost performs worse than most domain adaptation methods, providing a source-only reference for the transfer task. Per-dataset results and cross-seed variability for scRA-DRP, together with the common 17-dataset comparison across all methods, are provided in Appendix A.5.

4.3 ABLATION STUDIES

We evaluate branch combination and regional alignment in Table 2, and the effect of SAM on cross-seed stability in Table 3. Source-domain classification and domain-adversarial learning are commonly used components in bulk-to-single-cell transfer and are not individually ablated here.

Table 1: Single-cell drug response prediction results. W-F1 denotes weighted F1. Higher is better.

	GEO				EpiSen				Overall			
Method	**AUROC**	**AUPR**	**W-F1**	**ACC**	**AUROC**	**AUPR**	**W-F1**	**ACC**	**AUROC**	**AUPR**	**W-F1**	**ACC**
scRA-DRP	**0.7654**	**0.7394**	**0.7152**	**0.7075**	**0.8962**	**0.8979**	**0.8152**	**0.8174**	**0.8308**	**0.8186**	**0.7652**	**0.7624**
SCAD	0.7527	0.7118	0.5984	0.6413	0.8126	0.7987	0.5041	0.5872	0.7826	0.7552	0.5513	0.6143
DAGFormer	0.6555	0.5826	0.6467	0.6384	0.7380	0.8090	0.7198	0.7203	0.6967	0.6958	0.6832	0.6794
scAdaDrug	0.6676	0.6095	0.5432	0.6120	0.6760	0.6859	0.4338	0.5369	0.6718	0.6477	0.4885	0.5744
scDEAL	0.4824	0.4515	0.4464	0.5695	0.4819	0.5218	0.3832	0.5075	0.4822	0.4866	0.4148	0.5385
FourierDrug	0.5393	0.4619	0.4290	0.5719	0.6174	0.6202	0.3369	0.5014	0.5783	0.5411	0.3829	0.5367
DREEP	0.7473	0.6827	0.5994	0.6213	0.7205	0.7646	0.6706	0.6788	0.7347	0.7212	0.6329	0.6484
scFoundation	0.5903	0.5490	0.5587	0.5782	0.6453	0.6600	0.4944	0.5835	0.6178	0.6045	0.5266	0.5808
scGPT	0.6409	0.5608	0.4775	0.5305	0.5744	0.5769	0.4121	0.5149	0.6076	0.5688	0.4448	0.5227
CatBoost	0.5872	0.5005	0.4628	0.5604	0.5267	0.5534	0.3699	0.4961	0.5569	0.5269	0.4164	0.5282

Table 2: Branch combination and regional alignment.

Setting	**AUROC**	**AUPR**	**W-F1**	**ACC**
scRA-DRP	**0.8308**	**0.8186**	**0.7652**	**0.7624**
w/o R	0.8059	0.7766	0.7356	0.7330
w/o C	**0.8308**	**0.8186**	0.5974	0.6563
SMI	**0.8308**	**0.8186**	0.7403	0.7461
Mean-Fusion	0.8282	0.8061	0.7215	0.7299
w/o $\mathcal{L}_{\text{cov}}$	0.7450	0.7113	0.6846	0.6819
w/o $\mathcal{L}_{t\to s}$	0.7378	0.7040	0.6804	0.6779
w/o $\mathcal{L}_{s\to t}$	0.8209	0.8054	0.7544	0.7516
PA	0.8134	0.7862	0.7592	0.7564
DA	0.8044	0.7781	0.7352	0.7326

Table 3: Effect of SAM. Values are mean ± sample SD across three seeds, with each seed's score first averaged over the 18 datasets.

Setting	**AUROC**	**AUPR**	**W-F1**	**ACC**
scRA-DRP	0.8308 ± **0.0063**	0.8186 ± **0.0142**	0.7652 ± **0.0049**	0.7624 ± **0.0054**
w/o SAM	0.8304 ± 0.0178	0.8150 ± 0.0206	0.7690 ± 0.0234	0.7661 ± 0.0234

Branch combination. We compare the full method with either branch alone (w/o R or w/o C), single-model inference (SMI), and probability averaging (Mean-Fusion). SMI uses the ranking model for both outputs, with ordinary BN statistics for ranking and target-specific statistics for count estimation. The full method retains the ranking branch's AUROC and AUPR while increasing W-F1 from 0.5974 to 0.7652, and also improves classification over the classification branch alone. Its higher W-F1 and accuracy than SMI and Mean-Fusion support separately adapting the classification reference and combining the branches through ranking-guided decisions.

Regional alignment. Keeping the classification branch and pretraining unchanged, we remove the complete regional loss or either directional term, replace regions with points by setting all radii to zero (PA), or replace regional alignment with domain-adversarial alignment (DA). Removing target-to-source alignment causes a larger performance decrease than removing source-to-target alignment. The lower mean AUROC and AUPR of PA suggest that allowing heterogeneity within local regions benefits prediction. The full method also outperforms DA, supporting regional alignment over global domain alignment.

Effect of SAM. Across the same three seeds, removing SAM changes each mean metric by less than 0.005 but increases all four cross-seed standard deviations. SAM thus improves predictive stability across the evaluated seeds while maintaining similar mean performance.

4.4 Hyperparameter Sensitivity

We examine the anchor count M, domain-adversarial weights for pretraining (α) and the classification branch (β), and region radius coefficient κ. Figure 2 shows overall mean scores, with dashed lines marking defaults.

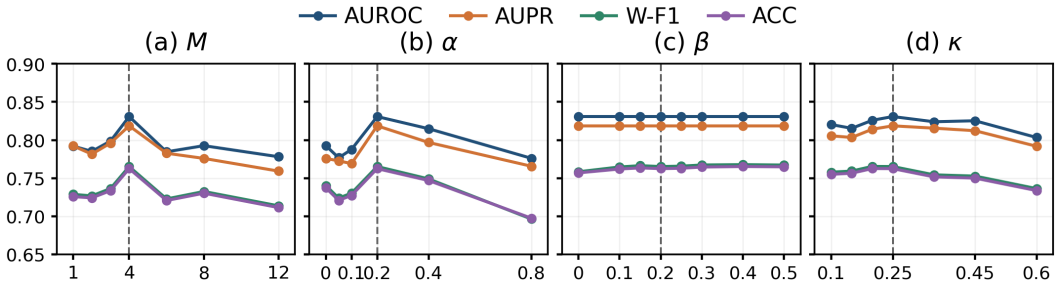

Figure 2: Hyperparameter sensitivity.

Among the tested values, $M = 4$ and $\alpha = 0.2$ achieve the highest mean scores on all four metrics. Varying β has little effect on W-F1 and accuracy and leaves AUROC and AUPR unchanged because the ranking branch is fixed. Performance remains similar near $\kappa = 0.20$–0.25 but declines at larger radii. A large SAM perturbation radius reduces W-F1 and accuracy (Appendix A.6).

4.5 Cell-Level Prediction Changes

To assess the effect of the dual-branch design across datasets, we quantify cell-level corrections and newly introduced errors relative to using the classification branch alone. Figure 3 shows that scRA-DRP improves accuracy on 10 datasets, reduces accuracy on four, and leaves four approximately unchanged ($|\Delta\text{ACC}| \leq 0.5\%$). Appendix A.7 illustrates these decision changes through prediction score distributions on PLX-4720/GSE108394.

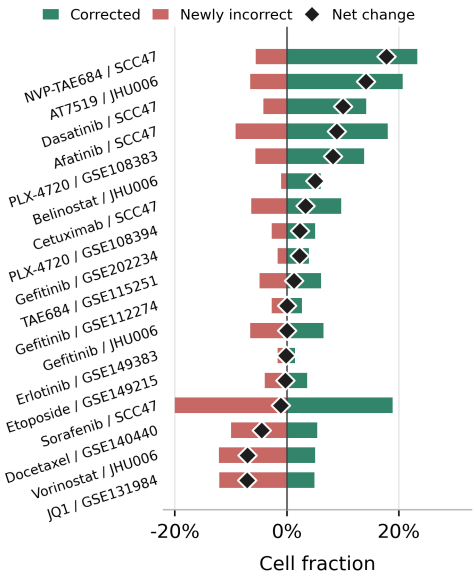

Figure 3: Cell-level prediction changes.

AT7519/JHU006 shows a clear improvement, with accuracy increasing from 0.6818 to 0.8232. Further exploration finds a stronger overall correlation with a proliferation-program score based on published S and G2/M gene programs (Kinker et al., 2020) for the ranking branch (Spearman correlation 0.6620) than for the classification branch (0.5077) or the ranking branch without regional alignment (0.3956; Appendix A.8). On Sorafenib/SCC47, newly introduced errors slightly outweigh corrections. The scarcity of sensitive samples in the source training set (25 of 314 cell lines) may limit the transfer of features predictive of drug response. For the Docetaxel dataset, predicting responses across both DU145 and PC3 cells poses an additional challenge due to potential differences between the two cell lines.

5 Conclusion

We introduced scRA-DRP for bulk-to-single-cell drug response prediction, combining regional alignment for response ranking with a separately adapted classification reference. This reference determines how many top-ranked cells are classified as sensitive. Experiments on 18 drug–single-cell dataset pairs show improved mean ranking and classification performance, while ablations support the contributions of regional alignment and separate adaptation. Hyperparameter analysis shows that the classification weight has little effect on performance, while the anchor count and pretraining weight have a greater impact. Cell-level analysis reveals both corrected and newly introduced errors, with a noticeable net decrease in accuracy on only 4 of the 18 datasets. These results support separately adapting response ordering and the classification reference for prediction without target response supervision.

AI USE STATEMENT

We used generative AI tools to assist with methodological discussions, experimental design, code debugging, data processing, and interpretation of experimental results. We also used these tools to identify relevant literature, organize manuscript content, draft and revise text, and assist with translation. These tools supported the research and writing process, while the authors made the final methodological and editorial decisions and take responsibility for the accuracy, originality, and integrity of all content in this work.

REPRODUCIBILITY STATEMENT

The core implementation of scRA-DRP is available at `https://anonymous.4open.science/r/scRA-DRP-28E0/`. Dataset sources, preprocessing procedures, and baseline configurations are provided in Appendices A.3 and A.4, with training settings described in the experimental setup.

REFERENCES

Junyi Chen, Xiaoying Wang, Anjun Ma, Qi-En Wang, Bingqiang Liu, Lang Li, Dong Xu, and Qin Ma. Deep transfer learning of cancer drug responses by integrating bulk and single-cell RNA-seq data. *Nature Communications*, 13:6494, 2022. doi: 10.1038/s41467-022-34277-7. URL https://doi.org/10.1038/s41467-022-34277-7.

Haotian Cui, Chloe Wang, Hassaan Maan, Kuan Pang, Fengning Luo, Nan Duan, and Bo Wang. scGPT: toward building a foundation model for single-cell multi-omics using generative AI. *Nature Methods*, 21:1470–1480, 2024. doi: 10.1038/s41592-024-02201-0. URL https://doi.org/10.1038/s41592-024-02201-0.

Pierre Foret, Ariel Kleiner, Hossein Mobahi, and Behnam Neyshabur. Sharpness-aware minimization for efficiently improving generalization. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=6Tm1mposlrM.

Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *Journal of Machine Learning Research*, 17(59):1–35, 2016. URL https://jmlr.org/papers/v17/15-239.html.

Shuang Ge, Yiming Ren, Huan Xu, Qiang Cheng, Shuqing Sun, and Zhixiang Ren. Few-shot generalization to novel compounds in single-cell drug response via graph-infused meta-pretraining. *IEEE Transactions on Computational Biology and Bioinformatics*, 22(6):3438–3449, 2025. doi: 10.1109/TCBBIO.2025.3626801. URL https://doi.org/10.1109/TCBBIO.2025.3626801.

Minsheng Hao, Jing Gong, Xin Zeng, Chiming Liu, Yucheng Guo, Xingyi Cheng, Taifeng Wang, Jianzhu Ma, Xuegong Zhang, and Le Song. Large-scale foundation model on single-cell transcriptomics. *Nature Methods*, 21:1481–1491, 2024. doi: 10.1038/s41592-024-02305-7. URL https://doi.org/10.1038/s41592-024-02305-7.

Kaishun Huang and Hui Liu. Identification of drug-resistant individual cells within tumors by semi-supervised transfer learning from bulk to single-cell transcriptome. *Communications Biology*, 8:530, 2025. doi: 10.1038/s42003-025-07959-3. URL https://doi.org/10.1038/s42003-025-07959-3.

Yu-An Huang, Xiyue Cao, Zhu-Hong You, Yue-Chao Li, Xuequn Shang, and Zhi-An Huang. scGSDR: Harnessing gene semantics for single-cell pharmacological profiling. *Communications Biology*, 8:1508, 2025. doi: 10.1038/s42003-025-08788-0. URL https://doi.org/10.1038/s42003-025-08788-0.

Francesco Iorio, Theo A. Knijnenburg, Daniel J. Vis, et al. A landscape of pharmacogenomic interactions in cancer. *Cell*, 166(3):740–754, 2016. doi: 10.1016/j.cell.2016.06.017.

Gabriela S. Kinker, Alissa C. Greenwald, Rotem Tal, Zhanna Orlova, Michael S. Cuoco, James M. McFarland, Allison Warren, Christopher Rodman, Jennifer A. Roth, Samantha A. Bender, Bhavna Kumar, James W. Rocco, Pedro A. C. M. Fernandes, Christopher C. Mader, Hadas Keren-Shaul, Alexander Plotnikov, Haim Barr, Aviad Tsherniak, Orit Rozenblatt-Rosen, Valery Krizhanovsky, Sidharth V. Puram, Aviv Regev, and Itay Tirosh. Pan-cancer single-cell RNA-seq identifies recurring programs of cellular heterogeneity. *Nature Genetics*, 52(11):1208–1218, 2020. doi: 10.1038/s41588-020-00726-6. URL https://doi.org/10.1038/s41588-020-00726-6.

Hui Liu, Wei Duan, and Judong Luo. Predicting single-cell drug sensitivity utilizing adaptive weighted features for multi-source domain adaptation. *IEEE Journal of Biomedical and Health Informatics*, 29(8):6043–6054, 2025. doi: 10.1109/JBHI.2025.3553126. URL https://doi.org/10.1109/JBHI.2025.3553126.

Xiaofeng Ma, Zhuo Zuo, Wei Shi, and Yulong Sun. Algorithmically defined therapeutic targets: integrating single-cell transfer learning frameworks with small molecule drugs to reverse disease-associated cell fates. *Frontiers in Pharmacology*, 17:1786561, 2026. doi: 10.3389/fphar.2026.1786561. URL https://doi.org/10.3389/fphar.2026.1786561.

Aditya Krishna Menon, Xiaoqian Jiang, Shankar Vembu, Charles Elkan, and Lucila Ohno-Machado. Predicting accurate probabilities with a ranking loss. In *Proceedings of the 29th International Conference on Machine Learning*, 2012. URL https://icml.cc/2012/papers/372.pdf.

Sangdon Park, Osbert Bastani, James Weimer, and Insup Lee. Calibrated prediction with covariate shift via unsupervised domain adaptation. In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pp. 3219–3229. PMLR, 2020. URL https://proceedings.mlr.press/v108/park20b.html.

Simona Pellecchia, Gaetano Viscido, Melania Franchini, and Gennaro Gambardella. Predicting drug response from single-cell expression profiles of tumours. *BMC Medicine*, 21:476, 2023. doi: 10.1186/s12916-023-03182-1. URL https://doi.org/10.1186/s12916-023-03182-1.

Pavel Posokhov, Sergei Masliukhin, Stepan Skrylnikov, Danil Tirskikh, and Olesia Makhnytkina. Relevance scores calibration for ranked list truncation via TMP adapter. In *Findings of the Association for Computational Linguistics: ACL 2025*, pp. 7728–7734. Association for Computational Linguistics, 2025. doi: 10.18653/v1/2025.findings-acl.402. URL https://aclanthology.org/2025.findings-acl.402/.

Liudmila Prokhorenkova, Gleb Gusev, Aleksandr Vorobev, Anna Veronika Dorogush, and Andrey Gulin. CatBoost: unbiased boosting with categorical features. In *Advances in Neural Information Processing Systems*, volume 31, 2018. URL https://proceedings.neurips.cc/paper/2018/hash/14491b756b3a51daac41c24863285549-Abstract.html.

Guanpeng Qi, Liugen Wang, Mengdi Nan, Yuhan Fu, Qing Ren, Yuan Zhang, Ziyan Sun, Zhixin Shi, Jiayi Lu, and Jie Gao. scXDR: drug response prediction across single-cell datasets via heterogeneous network transfer learning. *Communications Biology*, 9:139, 2026. doi: 10.1038/s42003-025-09418-5. URL https://doi.org/10.1038/s42003-025-09418-5.

Sydney M. Shaffer, Margaret C. Dunagin, Stefan R. Torborg, Eduardo A. Torre, Benjamin Emert, Clemens Krepler, Marilda Beqiri, Katrin Sproesser, Patricia A. Brafford, Min Xiao, Elliott Eggan, Ioannis N. Anastopoulos, Cesar A. Vargas-Garcia, Abhyudai Singh, Katherine L. Nathanson, Meenhard Herlyn, and Arjun Raj. Rare cell variability and drug-induced reprogramming as a mode of cancer drug resistance. *Nature*, 546:431–435, 2017. doi: 10.1038/nature22794. URL https://doi.org/10.1038/nature22794.

Ran Song, Yinpu Bai, Xuejun Liu, and Hui Liu. FourierDrug: a domain generalization framework for robust drug response prediction via frequency-space asymmetric attention. *Bioinformatics*, 42(6):btag276, 2026. doi: 10.1093/bioinformatics/btag276. URL https://doi.org/10.1093/bioinformatics/btag276.

Ximei Wang, Mingsheng Long, Jianmin Wang, and Michael I. Jordan. Transferable calibration with lower bias and variance in domain adaptation. In *Advances in Neural Information Processing Systems*, volume 33, 2020. URL https://proceedings.neurips.cc/paper/2020/hash/df12ecd077efc8c23881028604dbb8cc-Abstract.html.

Dechen Xu, Jie Li, and Li Zhou. A deep transfer learning method based on multi-perspective feature alignment for cancer drug response prediction. In *2025 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*. IEEE, 2025. doi: 10.1109/BIBM66473.2025.11357178. URL https://doi.org/10.1109/BIBM66473.2025.11357178.

Fen Yan, ZhiHua Du, and Yu-An Huang. DAGFormer: A graph-based domain adaptation approach for single-cell cancer drug response prediction. *PLOS Computational Biology*, 21(12):e1013832, 2025. doi: 10.1371/journal.pcbi.1013832. URL https://doi.org/10.1371/journal.pcbi.1013832.

Wanjuan Yang, Jorge Soares, Patricia Greninger, Elena J. Edelman, Howard Lightfoot, Simon Forbes, Nidhi Bindal, Dave Beare, James A. Smith, I. Richard Thompson, Sridhar Ramaswamy, P. Andrew Futreal, Daniel A. Haber, Michael R. Stratton, Cyril Benes, Ultan McDermott, and Mathew J. Garnett. Genomics of drug sensitivity in cancer (GDSC): a resource for therapeutic biomarker discovery in cancer cells. *Nucleic Acids Research*, 41(D1):D955–D961, 2013. doi: 10.1093/nar/gks1111. URL https://doi.org/10.1093/nar/gks1111.

Shunyu Zhang, Hu Liu, Wentian Bao, Enyun Yu, and Yang Song. A self-boosted framework for calibrated ranking. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 6226–6235. ACM, 2024. doi: 10.1145/3637528.3671570. URL https://doi.org/10.1145/3637528.3671570.

Han Zhao, Remi Tachet Des Combes, Kun Zhang, and Geoffrey Gordon. On learning invariant representations for domain adaptation. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 7523–7532. PMLR, 2019. URL https://proceedings.mlr.press/v97/zhao19a.html.

Zetian Zheng, Junyi Chen, Xingjian Chen, Lei Huang, Weidun Xie, Qiuzhen Lin, Xiangtao Li, and Ka-Chun Wong. Enabling single-cell drug response annotations from bulk RNA-Seq using SCAD. *Advanced Science*, 10(11):2204113, 2023. doi: 10.1002/advs.202204113. URL https://doi.org/10.1002/advs.202204113.

Murong Zhou, Zeyu Luo, Yu-Hang Yin, Qiaoming Liu, Guohua Wang, and Yuming Zhao. scATD: a high-throughput and interpretable framework for single-cell cancer drug resistance prediction and biomarker identification. *Briefings in Bioinformatics*, 26(3):bbaf268, 2025. doi: 10.1093/bib/bbaf268. URL https://doi.org/10.1093/bib/bbaf268.

A Appendix

A.1 SAM Optimization Details

During branch adaptation, SAM is applied to the source classification loss of each branch $b \in \{r, c\}$. Let θ_b denote the encoder and predictor parameters. At each training step, we compute the source classification gradient and its normalized perturbation:

$$g_b^{\mathrm{src}} = \nabla_{\theta_b} \mathcal{L}_{\mathrm{cls},b}(\theta_b), \qquad \delta_b = \rho \frac{g_b^{\mathrm{src}}}{\|g_b^{\mathrm{src}}\|_2 + 10^{-12}}, \tag{17}$$

where $\rho = 0.05$. Sensitivity to the perturbation radius is reported in Appendix A.6. The perturbation is recomputed at each step and treated as constant during differentiation. It is applied to the encoder and predictor parameters, excluding the domain discriminator.

The parameter update combines the source classification gradient evaluated at the perturbed parameters with the branch-specific transfer gradient evaluated at the original parameters:

$$g_b^{\mathrm{update}} = \nabla_{\vartheta} \mathcal{L}_{\mathrm{cls},b}(\vartheta)|_{\vartheta=\theta_b+\mathrm{sg}(\delta_b)} + g_b^{\mathrm{transfer}}(\theta_b). \tag{18}$$

Here, g_b^{transfer} is the regional alignment gradient for the ranking branch and the weighted domain-adversarial gradient for the classification branch. These transfer terms update the encoder, while the predictor receives the source classification gradient. The classification branch's domain discriminator is updated separately using the unperturbed weighted domain loss.

The perturbed forward pass reuses the same source batch and the random states of the original source forward pass, preserving the corresponding dropout masks. After computing the perturbed source gradient, we restore the original parameters and the batch-normalization buffers saved immediately before this extra forward pass. We also restore the random state from that point. The combined gradient in Eq. 18 is then applied using Adagrad. Thus, the extra SAM forward pass does not introduce an additional update to the running normalization statistics.

A.2 Count Selection with Tied Ranking Scores

Let $r_{(1)} \geq \cdots \geq r_{(n_t)}$ denote the ranking scores sorted in descending order. Cells with exactly equal scores form a tie group and are assigned the same response label. If the L distinct score groups have sizes $n_1, \ldots, n_L$ in descending score order, the admissible sensitive-cell counts are

$$\mathcal{K}_r = \left\{ 0,\ n_1,\ n_1 + n_2,\ \ldots,\ \sum_{\ell=1}^{L} n_\ell \right\}. \tag{19}$$

Given the classification branch's reference count $k_0 = \sum_{j=1}^{n_t} \mathbf{1}\{c_j \geq 0.5\}$, we select the admissible count closest to k_0, choosing the smaller count when two candidates are equally close:

$$k = \min \left\{ u \in \mathcal{K}_r :\ |u - k_0| = \min_{v \in \mathcal{K}_r} |v - k_0| \right\}. \tag{20}$$

The top k cells are classified as sensitive, and the remaining cells as resistant. When k_0 is an admissible count, this rule gives $k = k_0$. Otherwise, the cutoff moves to the nearest boundary between distinct score groups. The admissible counts include both 0 and n_t, allowing all cells to be classified as resistant or sensitive. Ties are determined by exact equality of the computed ranking scores, without rounding or a tolerance threshold.

A.3 Datasets and Preprocessing

Data sources. The source domain consists of GDSC cell-line expression profiles (Yang et al., 2013), using the RMA-normalized expression matrix distributed through PaccMann and used by SCAD (Zheng et al., 2023), with binary drug response annotations from Table S5C of Iorio et al. (2016). Sensitive and resistant samples are encoded as 1 and 0, respectively. The target evaluation comprises 9 GEO pairs, available under the accession numbers listed in Table 4, and 9 EpiSen pairs involving JHU006 and SCC47. For EpiSen, we use the processed expression matrices provided in the SCAD repository[1], with drug-specific labels obtained from the released SCAD annotations and the supplementary drug sensitivity data of Kinker et al. (2020). The EpiSen pairs reuse the same 66 JHU006 cells or 120 SCC47 cells and therefore do not represent independent cell cohorts.

Expression preprocessing. We use processed GDSC bulk and single-cell expression matrices, align source and target genes by Entrez identifiers, and average duplicate gene mappings. Only shared genes are retained as input features, excluding response labels and IC50 metadata. Expression is first standardized gene-wise within each domain, using all available drug-specific source samples before splitting and all target cells. A second standardizer is fitted jointly to source training and unlabeled target expression and applied to the source training, source validation, and target sets.

Source partition. We use a fixed stratified source split with approximately 80% of samples for training and 20% for validation, obtained from the first fold of a shuffled five-fold partition with random state 42. The same split is used for training seeds 42, 43, and 44, so variation across seeds reflects training randomness under a fixed source partition.

[1] https://github.com/CompBioT/SCAD

Table 4: The 18 drug–single-cell dataset pairs used in the reported analysis. Cell counts refer to the target evaluation sets. EpiSen pairs from the same cell line reuse the same cells.

Drug	Target dataset	Source	Cells	Sensitive	Resistant
Gefitinib	GSE112274	GEO	507	37	470
Docetaxel	GSE140440	GEO	324	162	162
Erlotinib	GSE149383	GEO	2730	1201	1529
Etoposide	GSE149215	GEO	2271	1204	1067
PLX-4720	GSE108383	GEO	130	78	52
Gefitinib	GSE202234	GEO	800	400	400
TAE684	GSE115251	GEO	22779	6868	15911
JQ1	GSE131984	GEO	2626	1061	1565
PLX-4720	GSE108394	GEO	6545	3298	3247
Gefitinib	JHU006	EpiSen	66	33	33
Vorinostat	JHU006	EpiSen	66	33	33
Afatinib	SCC47	EpiSen	120	60	60
Sorafenib	SCC47	EpiSen	120	60	60
NVP-TAE684	SCC47	EpiSen	120	60	60
Belinostat	JHU006	EpiSen	66	33	33
AT7519	JHU006	EpiSen	66	33	33
Cetuximab	SCC47	EpiSen	120	60	60
Dasatinib	SCC47	EpiSen	120	60	60

Target access. The same target expression set is used for adaptation and prediction, corresponding to a transductive setting. Target response labels are stored separately from the training inputs and are used for evaluation. Each drug–target pair is modeled separately, using its aligned source and target expression features.

A.4 Baseline Configurations

Domain adaptation and generalization baselines. We use the available implementations of SCAD, scDEAL, DAGFormer, scAdaDrug, and FourierDrug with the fixed source partition. Target response labels are reserved for evaluation. For scAdaDrug, the two source inputs are formed by splitting each class-balanced GDSC training batch into two subsets. DAGFormer uses float64 computation for numerical stability. FourierDrug uses binary cross-entropy on its sigmoid outputs to correct the released loss interface; target expression is used for prediction-time standardization but not parameter optimization.

DREEP. We use GDSC-derived ranked gene lists and average the available drug-specific enrichment scores for each cell. The response score is $q = (1 - \overline{\mathrm{ES}})/2$, with $q \geq 0.5$ indicating sensitivity, without an enrichment-significance filter. DREEP lacks a matching Cetuximab signature, so comparisons involving DREEP are restricted to the same 17 pairs for all methods.

scFoundation and scGPT. We freeze the pretrained encoders and use their expression representations as inputs to SCAD. The downstream network uses hidden widths of 1024 and 256, batch size 32, Adagrad with a learning rate of 0.001, and a domain-adversarial coefficient of 0.2. Predictions are obtained from the final adaptation checkpoint using a threshold of 0.5.

CatBoost. We train a source-only classifier with up to 1000 iterations, tree depth 6, learning rate 0.03, and positive-class weighting equal to the resistant-to-sensitive sample ratio in the source training set. Training stops early if the source validation Matthews correlation coefficient (MCC) does not improve for 50 consecutive boosting iterations. The prediction threshold is selected from $\{0.1, 0.2, \ldots, 0.9\}$ by maximizing source validation MCC.

A.5 Additional Prediction Results

Table 5 reports scRA-DRP results for each drug–dataset pair as the mean and sample standard deviation across seeds 42, 43, and 44. Metrics are computed separately for each seed, without prediction ensembling.

Table 5: Per-dataset results of scRA-DRP. Values are mean ± standard deviation across three seeds. W-F1 denotes weighted F1.

Drug	Dataset	AUROC	AUPR	W-F1	ACC
Gefitinib	GSE112274	0.9949 ± 0.0020	0.9528 ± 0.0162	0.7919 ± 0.0157	0.7265 ± 0.0209
Docetaxel	GSE140440	0.4643 ± 0.1479	0.5248 ± 0.1910	0.5089 ± 0.0575	0.5093 ± 0.0572
Erlotinib	GSE149383	0.4960 ± 0.0083	0.4484 ± 0.0061	0.5051 ± 0.0040	0.5037 ± 0.0042
Etoposide	GSE149215	0.5662 ± 0.0070	0.5689 ± 0.0064	0.5322 ± 0.0031	0.5390 ± 0.0040
PLX-4720	GSE108383	0.9594 ± 0.0583	0.9686 ± 0.0469	0.8834 ± 0.0269	0.8846 ± 0.0266
Gefitinib	GSE202234	0.9914 ± 0.0131	0.9907 ± 0.0142	0.9466 ± 0.0273	0.9467 ± 0.0272
TAE684	GSE115251	0.7888 ± 0.1839	0.6466 ± 0.2766	0.7035 ± 0.1032	0.6942 ± 0.1019
JQ1	GSE131984	0.7428 ± 0.1652	0.6270 ± 0.1810	0.6952 ± 0.1289	0.6932 ± 0.1299
PLX-4720	GSE108394	0.8846 ± 0.0130	0.9267 ± 0.0078	0.8698 ± 0.0084	0.8701 ± 0.0084
Gefitinib	JHU006	0.9663 ± 0.0160	0.9677 ± 0.0139	0.8609 ± 0.0977	0.8636 ± 0.0946
Vorinostat	JHU006	0.8913 ± 0.0862	0.8927 ± 0.0917	0.7392 ± 0.0310	0.7475 ± 0.0315
Afatinib	SCC47	0.8522 ± 0.0319	0.8595 ± 0.0576	0.7831 ± 0.0581	0.7833 ± 0.0577
Sorafenib	SCC47	0.7168 ± 0.1984	0.7220 ± 0.1896	0.6574 ± 0.1842	0.6611 ± 0.1819
NVP-TAE684	SCC47	0.9564 ± 0.0180	0.9462 ± 0.0335	0.8597 ± 0.0198	0.8611 ± 0.0192
Belinostat	JHU006	0.9813 ± 0.0188	0.9735 ± 0.0299	0.9443 ± 0.0233	0.9444 ± 0.0231
AT7519	JHU006	0.8968 ± 0.0429	0.8958 ± 0.0156	0.8231 ± 0.0974	0.8232 ± 0.0974
Cetuximab	SCC47	0.9581 ± 0.0282	0.9512 ± 0.0248	0.8944 ± 0.0679	0.8944 ± 0.0679
Dasatinib	SCC47	0.8466 ± 0.1313	0.8722 ± 0.0672	0.7746 ± 0.0580	0.7778 ± 0.0603

Table 6: Comparison on the 17 drug–dataset pairs covered by DREEP. All methods use the same evaluation pairs. Higher is better; the highest mean in each column is bold.

Method	AUROC	AUPR	W-F1	ACC
scRA-DRP	**0.8233**	**0.8108**	**0.7576**	**0.7547**
SCAD	0.7769	0.7497	0.5523	0.6153
DAGFormer	0.6980	0.6931	0.6862	0.6821
scAdaDrug	0.6681	0.6439	0.4939	0.5775
scDEAL	0.4900	0.4903	0.4184	0.5417
FourierDrug	0.5869	0.5459	0.3858	0.5388
DREEP	0.7347	0.7212	0.6329	0.6484
scFoundation	0.6009	0.5881	0.5086	0.5660
scGPT	0.6099	0.5660	0.4436	0.5224
CatBoost	0.5593	0.5253	0.4202	0.5299

Table 6 compares all methods on the same 17 drug–dataset pairs, excluding Cetuximab/SCC47 because DREEP lacks a matching GDSC signature. Scores are averaged across seeds within each pair and then equally across pairs. scRA-DRP achieves the highest mean scores on all four metrics in this common evaluation set.

A.6 Additional Hyperparameter Sensitivity

We vary the SAM perturbation radius ρ while keeping the remaining settings fixed. Results are averaged across three seeds within each of the 18 drug–dataset pairs and then equally across pairs. The bold parameter value indicates the default setting.

SAM perturbation radius. Table 7 reports results for $\rho \in \{0, 0.025, 0.05, 0.1, 0.2\}$, where $\rho = 0$ disables the SAM perturbation. Mean performance changes modestly between $\rho = 0$ and $\rho = 0.1$. Increasing ρ to 0.2 reduces W-F1 and accuracy to 0.7092 and 0.7188, respectively, while AUROC and AUPR remain close to their values at the default radius. The default $\rho = 0.05$ does not maximize every mean metric; the SAM ablation in the main text additionally evaluates its effect on cross-seed variability.

A.7 Prediction Score Distributions

Figure 4 shows prediction score distributions for PLX-4720/GSE108394, comprising 6,545 cells (3,298 sensitive and 3,247 resistant). Across three seeds, scRA-DRP achieves mean AUROC,

Table 7: Sensitivity to the SAM perturbation radius. The bold parameter value denotes the default configuration.

ρ	AUROC	AUPR	W-F1	ACC
0	0.8304	0.8150	0.7690	0.7661
0.025	0.8301	0.8169	0.7655	0.7625
0.05	0.8308	0.8186	0.7652	0.7624
0.1	0.8338	0.8223	0.7620	0.7600
0.2	0.8342	0.8162	0.7092	0.7188

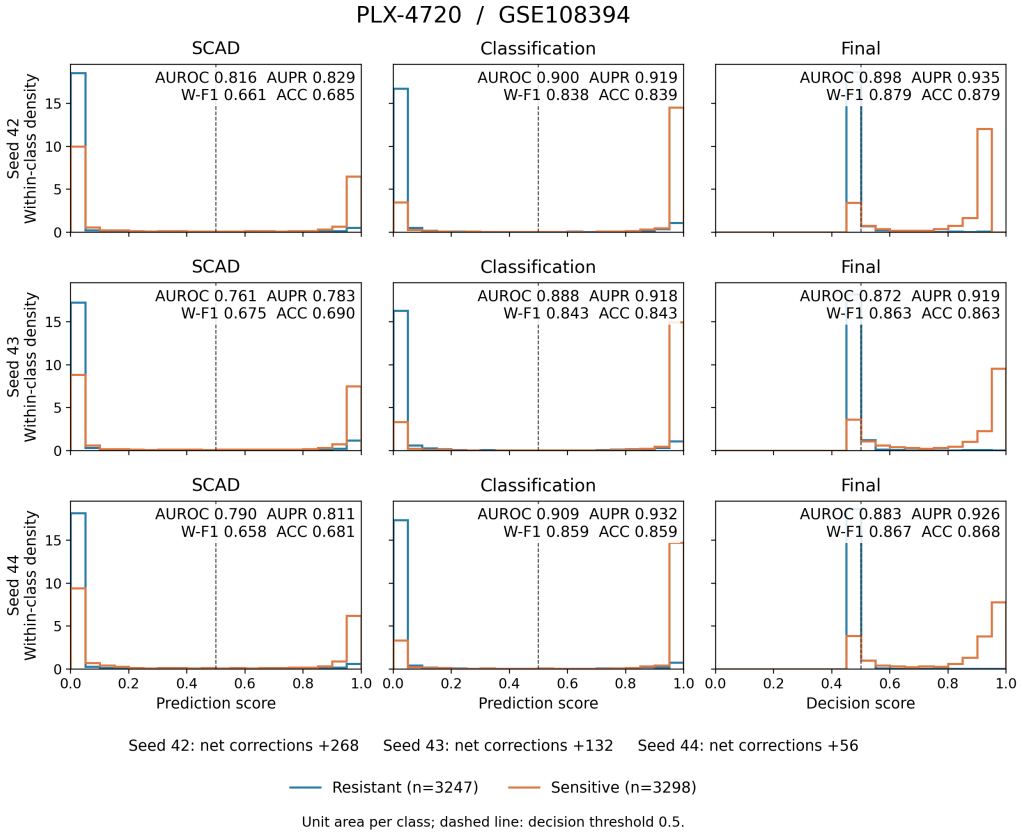

Figure 4: Prediction score distributions on PLX-4720/GSE108394. Rows correspond to seeds 42, 43, and 44; columns show SCAD, the classification branch, and scRA-DRP. Blue and orange denote resistant and sensitive cells, respectively. Each class histogram is independently normalized to unit area. Dashed lines mark the decision threshold of 0.5; final scores are decision scores rather than calibrated probabilities.

AUPR, accuracy, and weighted F1 of 0.8846, 0.9267, 0.8701, and 0.8698, respectively. Relative to the classification branch alone, the final predictions improve both sensitive- and resistant-class recall in every seed, with net corrections of 268, 132, and 56 cells. Although AUROC is slightly lower than that of the classification branch in all three seeds, accuracy and weighted F1 consistently improve. The predicted sensitive-cell count remains unchanged, while the ranking branch changes which cells are assigned to that count. This example illustrates how the two branches can improve classification without uniformly improving ranking metrics.

A.8 Proliferation-Program Analysis in AT7519/JHU006

We assess the association between predicted response scores and proliferation programs in the 66 evaluated AT7519/JHU006 cells. We use the S-phase and G2/M-phase gene sets from Supplemen-

tary Table S4 of Kinker et al. (2020), with 35 of 46 and 51 of 62 genes available, respectively. Within each cell, we rank all genes in the standardized target expression matrix, assign average ranks to ties, and divide ranks by the number of genes. Each program score is the mean normalized rank of its available genes, and the proliferation score is the arithmetic mean of the two program scores. Spearman correlations are computed across cells separately for seeds 42, 43, and 44 and then averaged. The resulting correlations are 0.6620 for the ranking branch, 0.5077 for the classification branch, and 0.3956 for the ranking branch without regional alignment. These results describe an exploratory association between response predictions and proliferation-related expression.