scBRIDGE: STOCHASTIC DISTRIBUTION TRANSPORT FOR SINGLE-CELL DRUG RESPONSE PREDICTION

Anonymous authors
Paper under double-blind review

ABSTRACT

Single-cell drug response prediction is limited by the scarcity of labeled scRNA-seq data, motivating knowledge transfer from large-scale bulk pharmacogenomic datasets. However, substantial distribution shifts between bulk and single-cell transcriptomes make such transfer challenging. Existing approaches are largely representation-centric, aligning the two domains in a shared latent space without explicitly modeling the transition between their distributions. To address this limitation, we formulate bulk-to-single-cell adaptation as an unpaired stochastic distribution transport problem and propose scBridge, a diffusion-based framework for single-cell drug response prediction. scBridge first learns compact latent representations and performs coarse covariance alignment to provide a stable cross-domain initialization. Drawing on Diffusion Schrödinger Bridge, scBridge then models a continuous stochastic transition between bulk and single-cell latent distributions through coupled forward and backward dynamics, optimized alternately via iterative proportional fitting with cycle-consistency regularization. Confidence-aware pseudo-label refinement further adapts the transported representations for downstream prediction. Across 12 scRNA-seq datasets, scBridge achieves a mean AUROC of 0.87, with a 9.33% relative improvement over the second-best model. The learned transport reduces domain separability by 56.8%–77.9% and reveals heterogeneous cell-level transport trajectories with distinct temporal correction dynamics across datasets. Model attribution further identifies biologically relevant mechanisms associated with BRAF inhibitor resistance. These results support stochastic distribution transport as an effective perspective for single-cell drug response prediction across modalities.

1 INTRODUCTION

Precision oncology seeks to match treatments to the molecular characteristics of individual tumors, but intratumoral heterogeneity complicates this goal (Qiao et al., 2025). Cells within the same tumor can occupy distinct transcriptional states and exhibit substantially different drug sensitivities, motivating the need to resolve drug response at single-cell resolution (Gambardella et al., 2022). Most large-scale pharmacogenomic resources are built on bulk transcriptomic measurements and provide abundant labeled screening data across diverse drugs and cellular contexts (Chen et al., 2022; Van de Sande et al., 2023). However, bulk measurements average signals across heterogeneous cell populations and can therefore obscure cell-state-specific variation (Van de Sande et al., 2023). Advances in single-cell RNA sequencing (scRNA-seq) have made it possible to profile such heterogeneity at cellular resolution, but datasets with matched treatment labels remain limited. This imbalance between abundant supervision in bulk data and cellular resolution in single-cell data has prompted increasing efforts to transfer predictive knowledge from bulk measurements to single-cell data.

Recent studies have increasingly explored deep transfer learning for bulk-to-single-cell drug response prediction, including scDEAL (Chen et al., 2022), SCAD (Zheng et al., 2023), and DAGFormer (Yan et al., 2025); see Appendix A.1 for related work. The methodological scope has since expanded to multi-source domain adaptation (Liu et al., 2025) and knowledge-distillation-based transfer (Zhou et al., 2025), aiming to improve generalization across single-cell datasets. Despite these advances, adaptation remains largely representation-centric, seeking to make bulk and single-cell features comparable in a shared latent space. Yet the cross-modal shift between population-level

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.

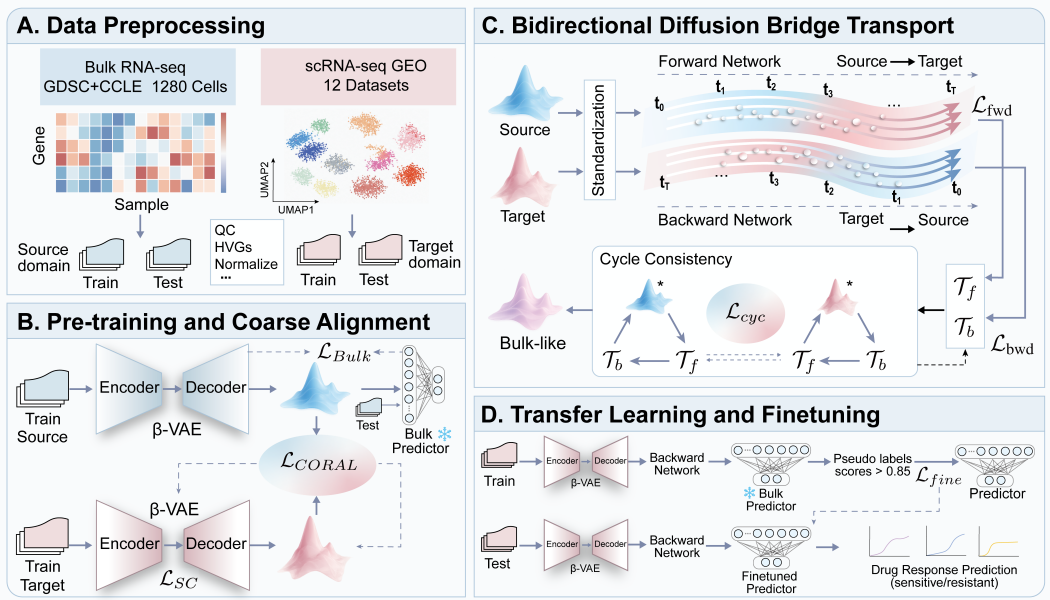

Figure 1: Overview of the scBridge framework. (A) Preprocessing and train–test partitioning of bulk and single-cell transcriptomic data. (B) β-VAE latent representation learning, bulk predictor pre-training, and CORAL-based coarse alignment. (C) Bidirectional diffusion bridge transport with alternating forward–backward optimization and cycle consistency. (D) Backward transport to the bulk-like space, confidence-aware pseudo-label finetuning, and single-cell drug response prediction.

bulk profiles and heterogeneous single-cell states can be highly nonlinear, making static representation alignment insufficient to explicitly characterize how one latent distribution is transformed into the other.

Beyond representation-level adaptation, distributional transport has also been used to relate unpaired single-cell populations. Waddington-OT reconstructs developmental trajectories from temporally sampled cells using optimal transport (Schiebinger et al., 2019), whereas CellOT learns neural transport maps between unpaired control and perturbed cells for perturbation-response prediction (Bunne et al., 2023). Subsequent work has moved toward stochastic transport, allowing transitions between cellular populations to be modeled probabilistically (Luo et al., 2025; Chi et al., 2026). These studies, however, are primarily concerned with transitions within the single-cell setting. Bulk-to-single-cell drug-response transfer differs in that it connects population-averaged measurements with heterogeneous cell-resolved profiles while relying largely on drug-response supervision from bulk pharmacogenomic data.

Based on the above motivation, we propose **scBridge**, as shown in Figure 1, a diffusion-based framework for single-cell drug response prediction. Specifically, we first learn a unified latent representation using a β-variational autoencoder (β-VAE) and perform coarse second-order statistical alignment through CORrelation ALignment (CORAL), providing a stable initialization. We then formulate cross-modality adaptation as an unpaired stochastic distribution transport problem. Departing from standard rigid alignment, we incorporate the concept of a Diffusion Schrödinger Bridge (DSB) to learn a continuous stochastic transport path between the bulk and single-cell latent distributions, allowing complex nonlinear cross-modal discrepancies to be resolved progressively. The bridge is parameterized by a pair of coupled forward and backward dynamics, governing bulk-to-single-cell and single-cell-to-bulk transport, respectively. We optimize these two dynamics through Iterative Proportional Fitting (IPF), together with cycle-consistency regularization to promote stable transport under unpaired supervision. Finally, a confidence-aware pseudo-label refinement strategy further adapts the transported representations for downstream prediction.

Our contributions are summarized as follows:

- We formulate bulk-to-single-cell drug response prediction as an unpaired stochastic distribution transport problem, shifting the focus from static representation alignment to modeling the transition between bulk and single-cell latent distributions.
- We develop **scBridge**, a bidirectional diffusion-bridge framework that learns cross-modal transport between unpaired bulk and single-cell representations while preserving structural consistency, and transfers bulk-trained drug-response knowledge to single cells without requiring paired cross-modal samples.
- Across 12 scRNA-seq datasets, scBridge consistently improves drug-response prediction and substantially reduces cross-domain discrepancy. Further analyses show robust performance under limited target-domain data and reveal biologically meaningful signals associated with drug response.

2 METHODS

Data Integration and Preprocessing. To mitigate batch effects and distributional discrepancies between bulk RNA-seq and single-cell (SC) datasets, we design a standardized preprocessing pipeline. For bulk data, we first perform feature selection based on gene dispersion to retain highly variable genes (HVGs), followed by Min-Max normalization to constrain expression values within $[0, 1]$, yielding a normalized feature matrix $\mathbf{x}^B \in \mathbb{R}^{N_B \times D_B}$, where N_B denotes the number of bulk samples and D_B the number of retained HVGs.

For the SC dataset, we perform cell-level and gene-level quality control, library size normalization, and log transformation, then select HVGs independently (Luecken & Theis, 2019), resulting in an expression matrix $\mathbf{x}^S \in \mathbb{R}^{N_S \times D_S}$, where N_S is the number of cells and D_S the number of SC HVGs. To obtain biologically meaningful clusters, we apply PCA and Leiden clustering (Traag et al., 2019). Finally, both modalities are subjected to unified Min-Max scaling to facilitate subsequent latent space learning.

Pre-training and Coarse Alignment. To unify the heterogeneous distributions between bulk and single-cell modalities, we employ a β-VAE encoder that maps input features $\mathbf{x}$ into a low-dimensional latent space $\mathcal{Z} \subset \mathbb{R}^L$ (Kingma & Welling, 2013; Higgins et al., 2017). The encoder $f_\phi : \mathcal{X} \to \mathcal{Z}$ parameterizes a posterior $q_\phi(\mathbf{z}|\mathbf{x}) = \mathcal{N}(\mathbf{z}; \boldsymbol{\mu}_\phi(\mathbf{x}), \mathrm{diag}(\boldsymbol{\sigma}^2_\phi(\mathbf{x})))$. The training objective is defined as:

$$\mathcal{L}_{\beta\text{-VAE}} = \mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x})}\big[-\log p_\theta(\mathbf{x}|\mathbf{z})\big] + \beta\, D_{KL}\big(q_\phi(\mathbf{z}|\mathbf{x}) \,\|\, p(\mathbf{z})\big), \tag{1}$$

where the expectation $\mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x})}[\cdot]$ is taken over the approximate posterior, $p_\theta(\mathbf{x}|\mathbf{z})$ denotes the decoder likelihood, $D_{KL}(\cdot\|\cdot)$ is the Kullback-Leibler divergence, and the prior $p(\mathbf{z}) = \mathcal{N}(0, \mathbf{I})$ is a standard isotropic Gaussian. The reconstruction term is implemented as mean squared error, and $\beta = 1.5$ in our experiments.

For the labeled bulk data $\{\mathbf{x}^B_i, y_i\}_{i=1}^{N_B}$, we first pretrain a bulk β-VAE encoder f_{ϕ_B} using $\mathcal{L}_{\beta\text{-VAE}}$ on the bulk training set. We then construct a separate predictor model g_{ψ_B} that shares the encoder weights from f_{ϕ_B} and appends a classification head P_ω. The predictor is fine-tuned with class-weighted cross-entropy loss:

$$\mathcal{L}_{\text{pred}}(\psi_B, \omega) = -\sum_i w_{y_i} \log P_\omega\big(g_{\psi_B}(\mathbf{x}^B_i)\big)_{y_i}, \tag{2}$$

where w balances the class frequencies. Importantly, g_{ψ_B} and the standalone encoder f_{ϕ_B} are two separate instances. The latter remains frozen during predictor training and is used solely for latent extraction in CORAL and DSB.

For single-cell data, the encoder f_{ϕ_S} is pretrained using the same β-VAE objective, but without any label information. Although SC response labels are available for dataset splitting and final evaluation, they are strictly hidden during β-VAE, CORAL, and DSB training.

To reduce the covariate shift between $\mathbf{z}^B = f_{\phi_B}(\mathbf{x}^B)$ and $\mathbf{z}^S = f_{\phi_S}(\mathbf{x}^S)$, we apply CORAL as a coarse alignment step (Sun & Saenko, 2016; Sun et al., 2016). The CORAL loss minimizes the Frobenius norm between their covariance matrices:

$$\mathcal{L}_{\text{CORAL}} = \frac{1}{4L^2} \|\mathbf{C}_{\mathbf{z}^B} - \mathbf{C}_{\mathbf{z}^S}\|_F^2, \tag{3}$$

where $\mathbf{C}_{\mathbf{z}^B}$ and $\mathbf{C}_{\mathbf{z}^S}$ are empirical covariance matrices of the bulk and single-cell latent features, respectively. This step fine-tunes f_{ϕ_S} while keeping f_{ϕ_B} fixed.

Bidirectional Diffusion Bridge Transport. To address the complex, nonlinear domain shifts between the bulk latent measure $\mathbb{P}_B$ and the single-cell latent measure $\mathbb{P}_S$, we draw inspiration from the concept of DSB (De Bortoli et al., 2021). Unlike standard linear alignments, we formulate the alignment as an entropy-regularized optimal transport problem (Léonard, 2013), learning a stochastic, topology-preserving mapping between unpaired heterogeneous distributions.

Prior to DSB training, the latent representations $\mathbf{z}^B$ and $\mathbf{z}^S$ extracted from their respective encoders (f_{ϕ_B} and the CORAL-tuned f_{ϕ_S}) are standardized per dimension. We construct a diffusion process over $t \in [0, 1]$ that forms a continuous bridge between $\mathbf{z}_0 \sim \mathbb{P}_B$ and $\mathbf{z}_1 \sim \mathbb{P}_S$. The continuous stochastic differential equations are discretized into N steps (with step size $\Delta t = 1/N$), and the forward and backward drifts are parameterized by time-conditioned neural networks $\mathbf{v}_{\theta_f}$ and $\mathbf{v}_{\theta_b}$ (De Bortoli et al., 2021):

$$d\mathbf{z}_t = \mathbf{v}_{\theta_f}(\mathbf{z}_t, t)\, dt + \sqrt{2\sigma^2}\, d\mathbf{W}_t, \tag{4}$$

$$d\mathbf{z}_t = \mathbf{v}_{\theta_b}(\mathbf{z}_t, t)\, dt + \sqrt{2\sigma^2}\, d\mathbf{W}_t^{\leftarrow}, \tag{5}$$

where $\mathbf{W}_t$ and $\mathbf{W}_t^{\leftarrow}$ denote the standard and time-reversed Wiener processes, respectively, and $\sigma^2 = 0.05$ is the fixed diffusion variance.

The drift networks are optimized via an IPF procedure. Theoretically, IPF alternates between projecting the forward and backward path measures to minimize their Kullback-Leibler divergence. In our implementation, this path-level optimization translates to minimizing the mean squared error between the predicted drift and the empirical finite-difference velocity of reference trajectories. At iteration k, the forward network is updated using trajectories generated by rolling out target single-cell samples $\mathbf{z}^S$ through the previous backward network $\mathbf{v}_{\theta_b}^{(k-1)}$:

$$\mathcal{L}_{\text{fwd}}(\theta_f) = \mathbb{E}_{t,\mathbf{z}_t,\mathbf{z}_{t+\Delta t}}\left[\left\|\mathbf{v}_{\theta_f}(\mathbf{z}_t, t) - \frac{\mathbf{z}_{t+\Delta t} - \mathbf{z}_t}{\Delta t}\right\|_2^2\right]. \tag{6}$$

For the initial step ($k = 0$), we construct a deterministic, noise-free linear bridge $\mathbf{z}_t = (1 - t)\mathbf{z}^B + t\mathbf{z}^S$. For subsequent iterations ($k > 0$), reference trajectories are cached offline to drastically optimize computational efficiency. Symmetrically, the backward network is then updated using rollouts from bulk samples via the updated $\mathbf{v}_{\theta_f}^{(k)}$:

$$\mathcal{L}_{\text{bwd}}(\theta_b) = \mathbb{E}_{t,\mathbf{z}_t,\mathbf{z}_{t-\Delta t}}\left[\left\|\mathbf{v}_{\theta_b}(\mathbf{z}_t, t) - \frac{\mathbf{z}_{t-\Delta t} - \mathbf{z}_t}{\Delta t}\right\|_2^2\right]. \tag{7}$$

All rollouts are simulated with Euler-Maruyama updates, incorporating Gaussian noise scaled by $\sqrt{2\sigma^2\Delta t}$.

To further enforce approximate invertibility and stringently preserve the intrinsic biological structures of individual samples during cross-modal transport, we introduce a cycle-consistency regularizer (Zhu et al., 2017) exclusively in the $\mathbb{P}_B \to \mathbb{P}_S \to \mathbb{P}_B$ direction. This loss is computed using deterministic (noise-free) transport operators $\mathcal{T}_f$ and $\mathcal{T}_b$:

$$\mathcal{L}_{\text{cyc}}(\theta_f, \theta_b) = \mathbb{E}_{\mathbf{z}^B \sim \mathbb{P}_B}\left[\left\|\mathcal{T}_b\big(\mathcal{T}_f(\mathbf{z}^B)\big) - \mathbf{z}^B\right\|_2^2\right]. \tag{8}$$

This dual-constraint design promotes global alignment between the two domains while preserving cell-level identity and local topological relationships.

Transfer Learning and Finetuning. After DSB convergence, we transfer the single-cell latent representations into the bulk-like space for prediction. The training set SC embeddings $\mathbf{z}_{\text{train}}^S$ are standardized using the per-dimension mean μ_S and standard deviation σ_S of the SC domain: $\tilde{\mathbf{z}}^S = (\mathbf{z}_{\text{train}}^S - \mu_S)/\sigma_S$. The backward transport operator $\mathcal{T}_b$ (discretized over decreasing time indices $t = N-1, \ldots, 0$) projects these standardized embeddings to the standardized bulk domain. Finally, we denormalize using bulk statistics μ_B and σ_B:

$$\hat{\mathbf{z}}^{S_{\text{like}}} = \mathcal{T}_b(\tilde{\mathbf{z}}^S) \odot \sigma_B + \mu_B, \tag{9}$$

where $\odot$ denotes element-wise multiplication.

The bulk predictor P_ω, pre-trained on the labeled bulk data, is then used to generate pseudo-labels for the transported SC training samples (Lee et al., 2013). To reduce noise, we retain only samples with a maximum softmax probability (Sohn et al., 2020) exceeding a confidence threshold $\tau = 0.85$:

$$\mathcal{D}_{\text{conf}} = \left\{ (\hat{\mathbf{z}}_j^{S_{like}}, \hat{y}_j) \;\middle|\; \max P_\omega(\hat{\mathbf{z}}_j^{S_{like}}) > 0.85 \right\}, \tag{10}$$

$$\hat{y}_j = \arg\max P_\omega(\hat{\mathbf{z}}_j^{S_{like}}). \tag{11}$$

Only if the number of confident samples exceeds half the batch size do we fine-tune the predictor head P_ω (not the encoder) using the standard cross-entropy loss (without class weighting) on $\mathcal{D}_{\text{conf}}$:

$$\mathcal{L}_{\text{fine}}(\omega) = -\frac{1}{|\mathcal{D}_{\text{conf}}|} \sum_{(\hat{\mathbf{z}}^{S_{like}}, \hat{y}) \in \mathcal{D}_{conf}} \log P_\omega(\hat{\mathbf{z}}^{S_{like}})_{\hat{y}}. \tag{12}$$

For the held-out SC test set, the same standardization, backward transport, and denormalization steps are applied, and the fine-tuned predictor produces the final response probabilities. The test labels are used only for evaluation and are never involved in model parameter updates, thus ensuring a fair assessment of generalization.

The pseudocode and time complexity of the above procedure are provided in Appendix A.2.

3 RESULTS

Datasets. To systematically evaluate single-cell drug response prediction, we collect 12 publicly available scRNA-seq datasets with drug-treatment conditions from the National Center for Biotechnology Information (NCBI) Gene Expression Omnibus (GEO) (Clough et al., 2024). These datasets cover diverse cancer types, cell types, and drug treatment conditions, providing a heterogeneous benchmark for evaluating single-cell drug response prediction across different biological contexts. Detailed information for each dataset is summarized in Appendix Table A.1.

For bulk RNA-seq data, we use the same data resources as scDEAL (Chen et al., 2022), integrating cancer cell line gene expression profiles and drug response data from the Genomics of Drug Sensitivity in Cancer (GDSC) (Yang et al., 2012) and Cancer Cell Line Encyclopedia (CCLE) (Barretina et al., 2012). The integrated bulk dataset reported in scDEAL comprises 1,280 cancer cell lines, 1,557 drugs or chemical compounds, and gene expression profiles of 15,962 genes (Chen et al., 2022). This bulk dataset is used to learn gene-drug response associations and subsequently provide bulk-level drug response information for single-cell drug response prediction. Details of the data split strategy and evaluation protocol are provided in Appendix A.3.1.

Performance Evaluation of Single-Cell Drug Response Prediction. We compare scBridge with nine representative methods, with detailed descriptions of all baselines provided in Appendix A.3.2. As shown in Table 1, scBridge achieves the best average performance across the three evaluation metrics, with complete results on all 12 datasets reported in Appendix A.3.3. Compared with OT, which adopts a related distribution-level perspective, scBridge improves AUROC, AUPR, and F1 by 62.7%, 56.0%, and 69.6%, respectively. The gains over both feature-alignment and domain-adaptation baselines indicate that explicitly modeling the transition between bulk and single-cell latent distributions improves the transfer of drug-response knowledge across modalities.

To examine whether these average gains are maintained across individual datasets, Figure 2A compares scBridge with the strongest baseline, scDEAL, across the 12 scRNA-seq datasets. scBridge shows strong performance across datasets, with F1, AUROC, and AUPR exceeding 0.90 on several

Table 1: Average AUROC, AUPR, and F1 scores of scBridge and nine representative baseline methods across 12 datasets. (**Best**, <u>Second best</u>)

Metric	DeepSADR	DAGFormer	scATD	SCAD	scDEAL	DSN-DANN	OT	GFK	DAE	**scBridge**
AUROC ↑	0.779	0.601	0.578	0.551	<u>0.793</u>	0.548	0.533	0.497	0.446	**0.867**
AUPR ↑	0.767	0.633	0.625	0.578	<u>0.802</u>	0.550	0.563	0.497	0.433	**0.878**
F1 ↑	0.757	0.498	0.457	0.420	<u>0.758</u>	0.386	0.454	0.339	0.292	**0.770**

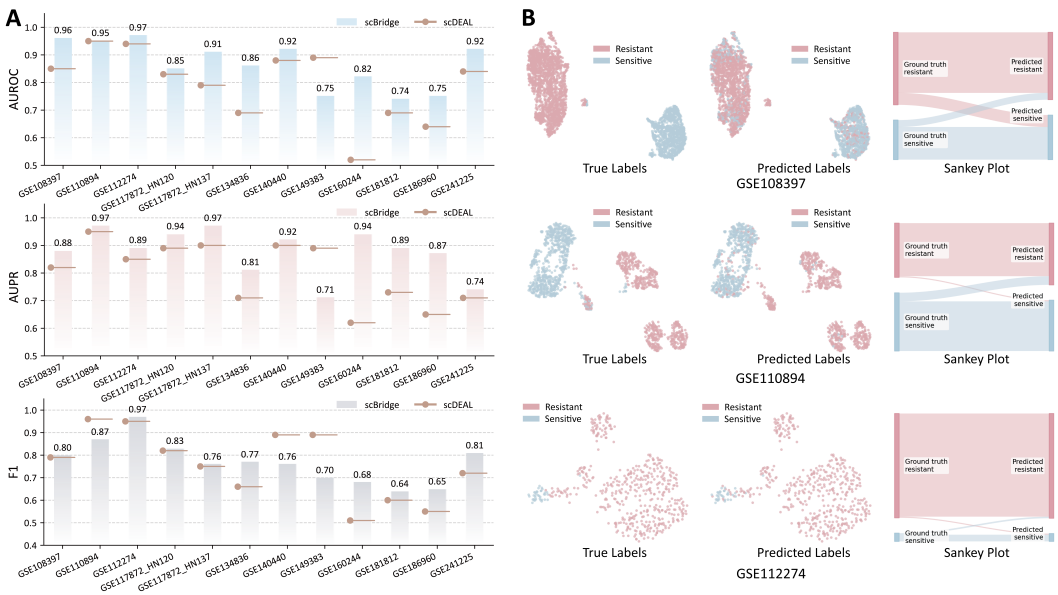

Figure 2: Performance evaluation of scBridge for single-cell drug response prediction. (A) Comparison of scBridge and scDEAL across 12 scRNA-seq datasets using AUROC, AUPR, and F1 scores. (B) UMAP visualizations of ground-truth and predicted drug-response labels, together with Sankey diagrams showing label correspondence, for three representative datasets.

Table 2: Ablation study of the major components of scBridge across 12 datasets. (**Best**, Second best)

Variant	CORAL	DSB	Cycle	IPF	AUROC ↑	AUPR ↑	F1 ↑
w/o CORAL	×	✓	✓	✓	0.810	0.795	0.697
w/o DSB	✓	×	–	✓	0.698	0.675	0.623
w/o Cycle	✓	✓	×	✓	0.820	0.838	0.734
w/o IPF	✓	✓	✓	×	0.758	0.733	0.628
scBridge	✓	✓	✓	✓	0.867	0.878	0.770

datasets. For example, AUPR reaches 0.97 on GSE110894 and 0.92 on GSE140440. The performance margin is particularly large on several datasets: compared with scDEAL, scBridge improves AUROC by 57.7% on GSE160244 and F1 by 16.7% on GSE134836. These dataset-wise comparisons further show that the average gains are broadly reflected across individual single-cell datasets.

To further illustrate the effectiveness of our approach, we visualize the latent space embeddings using UMAP alongside Sankey diagrams to track prediction flows (Figure 2B). The visualizations confirm that our model effectively separates sensitive and resistant cell populations, maintaining high concordance with the ground-truth labels. The Sankey plots further corroborate this, showing minimal misclassification and robust label alignment across distinct cellular states. Beyond overall classification, we further analyze the targeting performance of the model on extreme response samples. Top-K precision analysis across all 12 datasets confirms that scBridge reliably prioritizes the most sensitive and resistant cells at high-confidence thresholds, with detailed curves provided in Appendix A.3.4.

Ablation Study. We further evaluate the contribution of each major component of scBridge. As shown in Table 2, removing any component consistently degrades performance. Removing DSB causes the most pronounced degradation, with relative decreases of 19.1%-23.1% across the three metrics, highlighting the importance of explicitly modeling cross-domain transport. IPF is similarly critical, with performance drops of at least 12.6%, confirming the benefit of iterative forward-backward optimization. Removing CORAL or cycle consistency also leads to consistent decreases of at least 6.6% and 4.6%, respectively, indicating their complementary roles in providing stable initialization and preserving transport structure. Overall, the consistent degradation observed af-

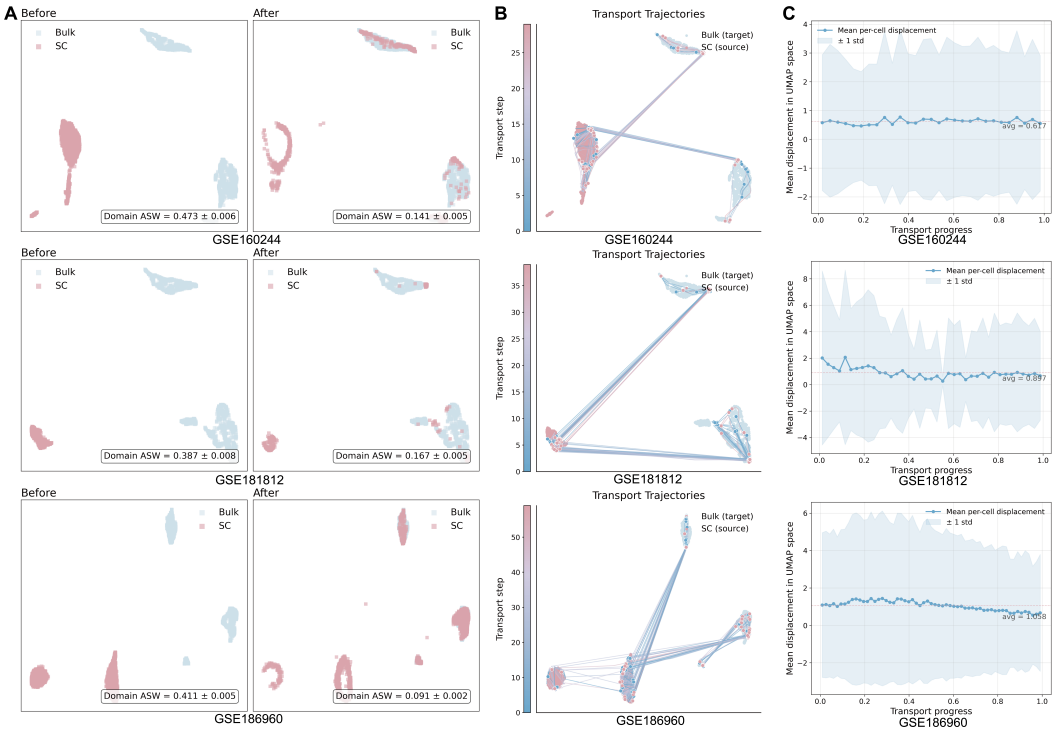

Figure 3: Cross-domain transport dynamics of scBridge across representative datasets. (A) Fixed-reference UMAP visualizations of bulk and single-cell representations before and after transport. (B) Deterministic SC-to-Bulk transport trajectories generated by the learned backward drift field. Each line traces the successive latent states of an individual single cell. (C) Stepwise displacement dynamics along the SC-to-Bulk transport process.

ter removing any component confirms that these modules are important for the effectiveness of scBridge. Detailed per-dataset ablation results are provided in Appendix Table A.5.

Analysis of the Learned Transport Bridge. To further characterize how scBridge reduces the discrepancy between single-cell and bulk representations, we analyze the deterministic SC-to-Bulk transport process generated by the learned backward drift. For each cell, intermediate latent states are recorded throughout transport and projected into a fixed UMAP reference space, while stepwise displacement is calculated to quantify how the magnitude of the correction evolves over transport time. As shown in Figure 3A, scBridge consistently reduces domain separability across the three representative datasets, with Domain ASW (Rousseeuw, 1987) decreasing by 56.8%–77.9%. The transported SC representations are distributed across multiple bulk-supported regions rather than concentrated into a single mode, indicating that the reduction in domain discrepancy is achieved through structured redistribution rather than trivial representation collapse. Additional visualization details and UMAP results for all 12 datasets are provided in Appendix A.3.6.

As shown in Figure 3B, individual cells follow markedly different transport directions and path lengths. Because each line traces the successive states of the same cell under the learned backward drift, these heterogeneous trajectories indicate that scBridge applies state-dependent and cell-specific corrections, rather than a uniform global transformation. Figure 3C shows that the magnitude of these corrections varies both over transport time and across cells. GSE160244 exhibits a relatively consistent correction magnitude throughout the transport process; GSE181812 shows stronger corrections at the early stage followed by reduced displacement, whereas GSE186960 maintains substantial transport during the early and intermediate stages before gradually decreasing toward the endpoint. The broad ± 1 SD ranges further indicate substantial cell-to-cell variation in transport magnitude, consistent with the heterogeneous path geometry observed in Figure 3B. Together, these

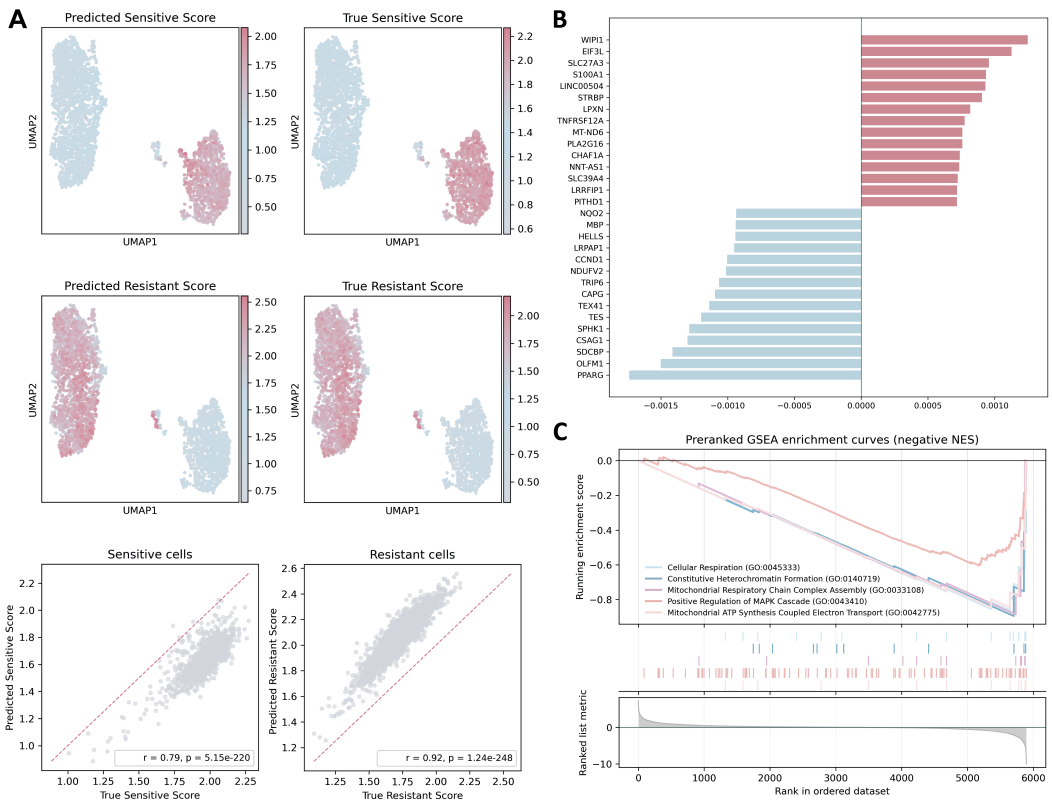

Figure 4: Model attribution and biological characterization of PLX4720 response in 451Lu melanoma cells. (A) DEG signature scores for model-predicted and experimentally defined sensitive and resistant states. (B) Integrated Gradients gene attributions for PLX4720 response prediction. (C) Preranked GSEA of genes ranked by IG attribution scores.

results show that scBridge accommodates dataset- and cell-specific domain discrepancies through adaptive SC-to-Bulk transport dynamics, while consistently achieving effective transport.

Model Attribution Reveals Mechanisms of PLX4720 Resistance. To determine whether the predicted drug-response states preserve biologically relevant transcriptional features, we compare model-predicted and experimentally defined PLX4720 response states in 451Lu melanoma cells from GSE108397 (Ho et al., 2018). Differentially expressed gene (DEG) signature scores derived from predicted labels recapitulate the spatial organization of experimentally defined sensitive and resistant transcriptional states on the UMAP manifold (Figure 4A). The DEG scores derived from predicted and experimental labels are strongly correlated in sensitive cells (Pearson $r = 0.79$, $P = 5.15 \times 10^{-220}$), with even stronger correspondence in resistant cells ($r = 0.92$, $P = 1.24 \times 10^{-248}$). These results indicate that the model not only distinguishes the two response states but also preserves the cell-level transcriptional structure associated with PLX4720 sensitivity and resistance.

We next apply Integrated Gradients (IG) (Sundararajan et al., 2017) to identify genes contributing to sensitive-versus-resistant predictions (Figure 4B), with positive attributions supporting PLX4720 sensitivity and negative attributions supporting resistance. WIPI1, EIF3L and SLC27A3 are among the strongest sensitivity-directed features, whereas PPARG, OLFM1, SDCBP, SPHK1, CCND1 and NDUFV2 show prominent resistance-directed contributions. Notably, increased CCND1 activity has been shown to reduce the sensitivity of BRAF-mutant melanoma cells to BRAF inhibition (Smalley et al., 2008), supporting the biological relevance of the attribution profile.

Preranked GSEA (Subramanian et al., 2005) of the IG-derived gene ranking further shows that the leading negative-NES gene sets are enriched for cellular respiration, mitochondrial respira-

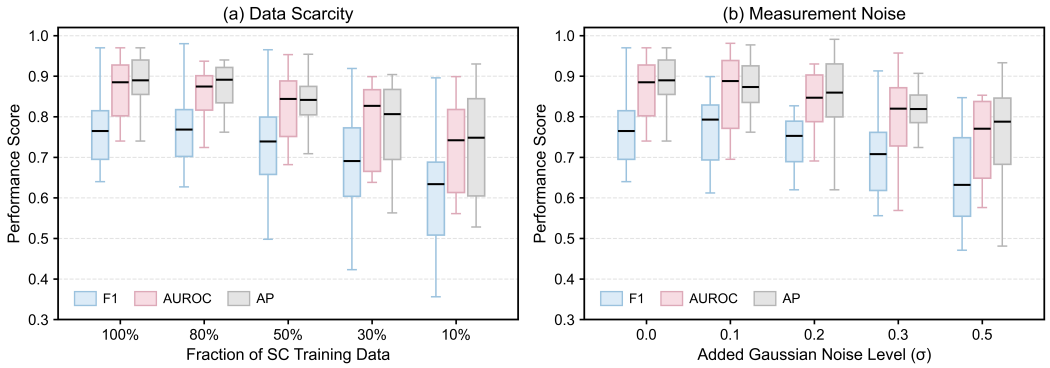

Figure 5: Robustness of scBridge to target-domain data scarcity and measurement noise. (a) Performance under progressively reduced fractions of SC training data, from 100% to 10%. (b) Performance under increasing Gaussian measurement noise, with σ ranging from 0 to 0.5.

tory chain complex assembly, mitochondrial ATP synthesis coupled to electron transport, positive regulation of the MAPK cascade and constitutive heterochromatin formation (Figure 4C). The MAPK-associated signal is consistent with sustained or reactivated ERK-MAPK signaling during acquisition of BRAF-inhibitor resistance in 451Lu cells (Villanueva et al., 2010). Meanwhile, enrichment of multiple mitochondrial respiration and electron-transport programs points to oxidative metabolic reprogramming in the resistance-associated state. BRAF inhibition has been shown to promote MITF-PGC1α-dependent mitochondrial biogenesis and oxidative phosphorylation (Haq et al., 2013), whereas melanoma cells with acquired BRAFi resistance exhibit increased mitochondrial respiration and greater dependence on mitochondrial metabolism (Baenke et al., 2016). Collectively, these resistance-associated features converge on two established axes of BRAF-inhibitor resistance, MAPK pathway reactivation and mitochondrial metabolic adaptation, providing biological support for the single-cell drug-response representation learned by the model.

Robustness Analysis. We further evaluate whether scBridge remains reliable when target single-cell (SC) information is limited or noisy. Specifically, we vary the available SC training data from 100% to 10% and add Gaussian noise with σ ranging from 0 to 0.5. Performance is evaluated across 12 datasets using F1, AUROC, and AUPR, with detailed results reported in Appendix Tables A.6 and A.7. As shown in Figure 5(a), performance decreases gradually with reduced SC training data, with substantial degradation appearing mainly at the extreme 10% setting. Even with only 50% of the target-domain data, the median F1, AUROC, and AUPR remain approximately 0.73, 0.82, and 0.83, respectively, indicating that scBridge can maintain effective knowledge transfer without densely sampled target data. Similarly, Figure 5(b) shows progressive degradation as measurement noise increases, while AUROC and AUPR remain relatively resilient even at $\sigma = 0.5$.

4 Conclusion

In this work, we introduce **scBridge**, a stochastic distribution transport framework for transferring drug-response knowledge from bulk transcriptomic data to single-cell settings. Rather than treating cross-domain adaptation as a static representation alignment problem, scBridge models the transition between bulk and single-cell latent distributions through a bidirectional diffusion bridge, combining coarse statistical alignment, iterative forward–backward transport optimization, cycle consistency, and confidence-aware prediction refinement. Across 12 scRNA-seq datasets, scBridge consistently achieves strong predictive performance and outperforms representative domain adaptation and distribution-matching baselines. Further analyses show that the learned transport reduces cross-domain discrepancy, captures heterogeneous cell- and dataset-specific transition dynamics, and remains robust under limited target-domain data and measurement noise. In addition, a melanoma case study using model attribution identifies biologically relevant signals associated with PLX4720 resistance. Together, these results demonstrate the potential of stochastic distribution transport as a flexible and effective framework for cross-domain single-cell drug response prediction.

Reproducibility Statement

Code and implementation details for reproducing the experiments are available at `https://anonymous.4open.science/r/scBridge-B55B/`. All datasets used in this study are included in the repository and are also publicly available from the original sources described in the manuscript.

AI Use Statement

Generative AI tools were used solely to aid in language polishing and improving the readability of the manuscript. They were not used in the development of the methods, experimental design, data processing, analysis, or interpretation of results. The authors reviewed all AI-assisted edits and take full responsibility for the final manuscript.

Ethics Statement

This work uses publicly available bulk transcriptomic and single-cell RNA sequencing datasets. No new human or animal experiments were conducted, and the method is intended for computational research rather than clinical decision making.

References

Alexandre F Aissa, Abul BMMK Islam, Majd M Ariss, Cammille C Go, Alexandra E Rader, Ryan D Conrardy, Alexa M Gajda, Carlota Rubio-Perez, Klara Valyi-Nagy, Mary Pasquinelli, et al. Single-cell transcriptional changes associated with drug tolerance and response to combination therapies in cancer. *Nature communications*, 12(1):1628, 2021.

Franziska Baenke, Barbara Chaneton, Matthew Smith, Niels Van Den Broek, Kate Hogan, Haoran Tang, Amaya Viros, Matthew Martin, Laura Galbraith, Maria R Girotti, et al. Resistance to braf inhibitors induces glutamine dependency in melanoma cells. *Molecular oncology*, 10(1):73–84, 2016.

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.

Charles C Bell, Katie A Fennell, Yih-Chih Chan, Florian Rambow, Miriam M Yeung, Dane Vassiliadis, Luis Lara, Paul Yeh, Luciano G Martelotto, Aljosja Rogiers, et al. Targeting enhancer switching overcomes non-genetic drug resistance in acute myeloid leukaemia. *Nature communications*, 10(1):2723, 2019.

Konstantinos Bousmalis, George Trigeorgis, Nathan Silberman, Dilip Krishnan, and Dumitru Erhan. Domain separation networks. *Advances in neural information processing systems*, 29, 2016.

Charlotte Bunne, Stefan G Stark, Gabriele Gut, Jacobo Sarabia Del Castillo, Mitch Levesque, Kjong-Van Lehmann, Lucas Pelkmans, Andreas Krause, and Gunnar Rätsch. Learning single-cell perturbation responses using neural optimal transport. *Nature methods*, 20(11):1759–1768, 2023.

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(1):6494, 2022.

Changxi Chi, Yufei Huang, Jun Xia, Jiangbin Zheng, Yunfan Liu, Zelin Zang, and Stan Z Li. Departures: Distributional transport for single-cell perturbation prediction with neural schrödinger bridges. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pp. 20544–20552, 2026.

Emily Clough, Tanya Barrett, Stephen E Wilhite, Pierre Ledoux, Carlos Evangelista, Irene F Kim, Maxim Tomashevsky, Kimberly A Marshall, Katherine H Phillippy, Patti M Sherman, et al. Ncbi geo: archive for gene expression and epigenomics data sets: 23-year update. *Nucleic acids research*, 52(D1):D138–D144, 2024.

Valentin De Bortoli, James Thornton, Jeremy Heng, and Arnaud Doucet. Diffusion schrödinger bridge with applications to score-based generative modeling. *Advances in neural information processing systems*, 34:17695–17709, 2021.

Pinar Demetci, Rebecca Santorella, Björn Sandstede, William Stafford Noble, and Ritambhara Singh. Scot: single-cell multi-omics alignment with optimal transport. *Journal of computational biology*, 29(1):3–18, 2022.

G Gambardella, G Viscido, B Tumaini, A Isacchi, R Bosotti, and D Di Bernardo. A single-cell analysis of breast cancer cell lines to study tumour heterogeneity and drug response. *Nature communications*, 13(1):1714, 2022.

Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario March, and Victor Lempitsky. Domain-adversarial training of neural networks. *Journal of machine learning research*, 17(59):1–35, 2016.

Caroline E Gleason, Mark A Dickson, Mary E Klein, Cristina R Antonescu, Rodrigo Gularte-Merida, Marimar Benitez, Juliana I Delgado, Raghu P Kataru, Mark Wei Yi Tan, Martina Bradic, et al. Therapy-induced senescence contributes to the efficacy of abemaciclib in patients with dedifferentiated liposarcoma. *Clinical Cancer Research*, 30(4):703–718, 2024.

Boqing Gong, Yuan Shi, Fei Sha, and Kristen Grauman. Geodesic flow kernel for unsupervised domain adaptation. In *2012 IEEE conference on computer vision and pattern recognition*, pp. 2066–2073. IEEE, 2012.

Rizwan Haq, Jonathan Shoag, Pedro Andreu-Perez, Satoru Yokoyama, Hannah Edelman, Glenn C Rowe, Dennie T Frederick, Aeron D Hurley, Abhinav Nellore, Andrew L Kung, et al. Oncogenic braf regulates oxidative metabolism via pgc1α and mitf. *Cancer cell*, 23(3):302–315, 2013.

Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner. beta-vae: Learning basic visual concepts with a constrained variational framework. In *International conference on learning representations*, 2017.

Yu-Jui Ho, Naishitha Anaparthy, David Molik, Grinu Mathew, Toby Aicher, Ami Patel, James Hicks, and Molly Gale Hammell. Single-cell rna-seq analysis identifies markers of resistance to targeted braf inhibitors in melanoma cell populations. *Genome research*, 28(9):1353–1363, 2018.

Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.

Dominik Klein, Giovanni Palla, Marius Lange, Michal Klein, Zoe Piran, Manuel Gander, Laetitia Meng-Papaxanthos, Michael Sterr, Lama Saber, Changying Jing, et al. Mapping cells through time and space with moscot. *Nature*, 638(8052):1065–1075, 2025.

Say Li Kong, Huipeng Li, Joyce A Tai, Elise T Courtois, Huay Mei Poh, Dawn Pingxi Lau, Yu Xuan Haw, Narayanan Gopalakrishna Iyer, Daniel Shao Weng Tan, Shyam Prabhakar, et al. Concurrent single-cell rna and targeted dna sequencing on an automated platform for comeasurement of genomic and transcriptomic signatures. *Clinical chemistry*, 65(2):272–281, 2019.

Dong-Hyun Lee et al. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on challenges in representation learning, ICML*, volume 3, pp. 896. Atlanta, 2013.

Christian Léonard. A survey of the schr\" odinger problem and some of its connections with optimal transport. *arXiv preprint arXiv:1308.0215*, 2013.

Guan-Horng Liu, Arash Vahdat, De-An Huang, Evangelos A Theodorou, Weili Nie, and Anima Anandkumar. I²sb: Image-to-image schrödinger bridge. *arXiv preprint arXiv:2302.05872*, 2023.

Guan-Horng Liu, Yaron Lipman, Maximilian Nickel, Brian Karrer, Evangelos Theodorou, and Ricky TQ Chen. Generalized schrödinger bridge matching. In *International Conference on Learning Representations*, volume 2024, pp. 14527–14552, 2024.

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.

Malte D Luecken and Fabian J Theis. Current best practices in single-cell rna-seq analysis: a tutorial. *Molecular systems biology*, 15(6):MSB188746, 2019.

Yiqing Luo, Lin Liu, Yaxin Fu, Yi Deng, and Lin Tang. Prediction of single-cell perturbation response based on direction-constrained diffusion schrödinger bridge. *Journal of Biomedical Informatics*, pp. 104915, 2025.

Danielle Maeser, Weijie Zhang, Yingbo Huang, and R Stephanie Huang. A review of computational methods for predicting cancer drug response at the single-cell level through integration with bulk rnaseq data. *Current opinion in structural biology*, 84:102745, 2024.

Daniel R Principe, Alexandre F Aissa, Sandeep Kumar, Thao ND Pham, Patrick W Underwood, Rakesh Nair, Rong Ke, Basabi Rana, Jose G Trevino, Hidayatullah G Munshi, et al. Calcium channel blockers potentiate gemcitabine chemotherapy in pancreatic cancer. *Proceedings of the National Academy of Sciences*, 119(18):e2200143119, 2022.

Diane Qiao, Richard C Wang, and Zhixiang Wang. Precision oncology: current landscape, emerging trends, challenges, and future perspectives. *Cells*, 14(22):1804, 2025.

Peter J Rousseeuw. Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *Journal of computational and applied mathematics*, 20:53–65, 1987.

Geoffrey Schiebinger, Jian Shu, Marcin Tabaka, Brian Cleary, Vidya Subramanian, Aryeh Solomon, Joshua Gould, Siyan Liu, Stacie Lin, Peter Berube, et al. Optimal-transport analysis of single-cell gene expression identifies developmental trajectories in reprogramming. *Cell*, 176(4):928–943, 2019.

Patricia M Schnepp, Greg Shelley, Jinlu Dai, Nicole Wakim, Hui Jiang, Atsushi Mizokami, and Evan T Keller. Single-cell transcriptomics analysis identifies nuclear protein 1 as a regulator of docetaxel resistance in prostate cancer cells. *Molecular Cancer Research*, 18(9):1290–1301, 2020.

Svetlana E Semina, Purab Pal, Nidhi S Kansara, Rosemary J Huggins, Elaine T Alarid, Geoffrey L Greene, and Jonna Frasor. Selective pressure of endocrine therapy activates the integrated stress response through nfκb signaling in a subpopulation of er positive breast cancer cells. *Breast Cancer Research*, 24(1):19, 2022.

Ankur Sharma, Elaine Yiqun Cao, Vibhor Kumar, Xiaoqian Zhang, Hui Sun Leong, Angeline Mei Lin Wong, Neeraja Ramakrishnan, Muhammad Hakimullah, Hui Min Vivian Teo, Fui Teen Chong, et al. Longitudinal single-cell rna sequencing of patient-derived primary cells reveals drug-induced infidelity in stem cell hierarchy. *Nature communications*, 9(1):4931, 2018.

Keiran SM Smalley, Mercedes Lioni, Maurizia Dalla Palma, Min Xiao, Brijal Desai, Suzanne Egyhazi, Johan Hansson, Hong Wu, Alastair J King, Patricia Van Belle, et al. Increased cyclin d1 expression can mediate braf inhibitor resistance in braf v600e–mutated melanomas. *Molecular cancer therapeutics*, 7(9):2876–2883, 2008.

Kihyuk Sohn, David Berthelot, Nicholas Carlini, Zizhao Zhang, Han Zhang, Colin Raffel, Ekin Dogus Cubuk, Alexey Kurakin, and Chun-Liang Li. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *Advances in neural information processing systems*, 33: 596–608, 2020.

Aravind Subramanian, Pablo Tamayo, Vamsi K Mootha, Sayan Mukherjee, Benjamin L Ebert, Michael A Gillette, Amanda Paulovich, Scott L Pomeroy, Todd R Golub, Eric S Lander, et al. Gene set enrichment analysis: a knowledge-based approach for interpreting genome-wide expression profiles. *Proceedings of the national academy of sciences*, 102(43):15545–15550, 2005.

Baochen Sun and Kate Saenko. Deep coral: Correlation alignment for deep domain adaptation. In *European conference on computer vision*, pp. 443–450. Springer, 2016.

Baochen Sun, Jiashi Feng, and Kate Saenko. Return of frustratingly easy domain adaptation. In *Proceedings of the AAAI conference on artificial intelligence*, volume 30, 2016.

Mukund Sundararajan, Ankur Taly, and Qiqi Yan. Axiomatic attribution for deep networks. In *International conference on machine learning*, pp. 3319–3328. PMLR, 2017.

Chayaporn Suphavilai, Shumei Chia, Ankur Sharma, Lorna Tu, Rafael Peres Da Silva, Aanchal Mongia, Ramanuj DasGupta, and Niranjan Nagarajan. Predicting heterogeneity in clone-specific therapeutic vulnerabilities using single-cell transcriptomic signatures. *Genome Medicine*, 13(1): 189, 2021.

Alexander Tong, Jessie Huang, Guy Wolf, David Van Dijk, and Smita Krishnaswamy. Trajectorynet: A dynamic optimal transport network for modeling cellular dynamics. In *International conference on machine learning*, pp. 9526–9536. PMLR, 2020.

Vincent A Traag, Ludo Waltman, and Nees Jan Van Eck. From louvain to leiden: guaranteeing well-connected communities. *Scientific reports*, 9(1):5233, 2019.

Bram Van de Sande, Joon Sang Lee, Euphemia Mutasa-Gottgens, Bart Naughton, Wendi Bacon, Jonathan Manning, Yong Wang, Jack Pollard, Melissa Mendez, Jon Hill, et al. Applications of single-cell rna sequencing in drug discovery and development. *Nature reviews Drug discovery*, 22(6):496–520, 2023.

Jessie Villanueva, Adina Vultur, John T Lee, Rajasekharan Somasundaram, Mizuho Fukunaga-Kalabis, Angela K Cipolla, Bradley Wubbenhorst, Xiaowei Xu, Phyllis A Gimotty, Damien Kee, et al. Acquired resistance to braf inhibitors mediated by a raf kinase switch in melanoma can be overcome by cotargeting mek and igf-1r/pi3k. *Cancer cell*, 18(6):683–695, 2010.

Pascal Vincent, Hugo Larochelle, Yoshua Bengio, and Pierre-Antoine Manzagol. Extracting and composing robust features with denoising autoencoders. In *Proceedings of the 25th international conference on Machine learning*, pp. 1096–1103, 2008.

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.

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.

Yuanpeng Zhang, Zhijian Huang, Ziyu Fan, Siyuan Shen, Yahan Li, Shangqian Wu, Min Wu, and Lei Deng. Deepsadr: Deep transfer learning with subsequence interaction and adaptive readout for cancer drug response prediction. In *International Conference on Learning Representations*, volume 2026, pp. 116485–116513, 2026.

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.

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.

Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *2017 IEEE international conference on computer vision (ICCV)*, pp. 2242–2251. Ieee, 2017.

APPENDIX

A.1 RELATED WORK

Cross-modal Transfer Learning for Single-Cell Drug Response Prediction. The scarcity of large-scale drug-response annotations at single-cell resolution has motivated the transfer of pharmacogenomic knowledge from bulk cancer cell lines to scRNA-seq data (Maeser et al., 2024). Early approaches such as CaDRReS-Sc leverage drug-response profiles learned from large-scale cell-line screens to estimate clone-specific therapeutic vulnerabilities from single-cell transcriptomes (Suphavilai et al., 2021). More recent deep learning methods explicitly address the distribution gap between bulk and single-cell measurements. scDEAL learns modality-specific representations and minimizes maximum mean discrepancy between bulk and single-cell latent spaces while preserving cellular heterogeneity through clustering regularization (Chen et al., 2022), whereas SCAD adopts adversarial domain adaptation to learn domain-invariant representations for transferring bulk drug-response labels to individual cells (Zheng et al., 2023). scATD further combines pretrained transcriptomic foundation models with bidirectional adaptive instance normalization and knowledge distillation to improve cross-dataset generalization and computational efficiency (Zhou et al., 2025). DAGFormer extends domain adaptation to graph-structured representations, incorporating cell-cell relationships and disentangling shared and modality-specific features (Yan et al., 2025). Despite their effectiveness, these approaches primarily formulate cross-modal transfer as representation alignment, where discrepancies are reduced through latent space matching, adversarial objectives, normalization, or graph-level adaptation. They therefore optimize endpoint similarity without explicitly modeling the continuous distributional transition from the bulk domain to the single-cell domain.

Optimal Transport for Single-Cell Distribution Mapping. Optimal transport (OT) provides a distribution-level perspective for relating unpaired cellular populations and has been increasingly adopted in single-cell analysis. Waddington-OT uses entropy-regularized optimal transport to reconstruct developmental couplings between cell populations observed at different time points, demonstrating that population-level transport can recover latent cellular trajectories from destructive snapshot measurements (Schiebinger et al., 2019). TrajectoryNet subsequently combines dynamic optimal transport with continuous normalizing flows to learn nonlinear paths between temporally observed cellular distributions (Tong et al., 2020). Beyond temporal dynamics, SCOT employs Gromov-Wasserstein optimal transport to align unpaired single-cell multi-omics datasets while preserving their within-domain geometries (Demetci et al., 2022). CellOT directly formulates perturbation response prediction as neural optimal transport between unpaired control and perturbed cell populations, enabling cell-specific responses to be inferred without one-to-one correspondences (Bunne et al., 2023). More recently, moscot provides a scalable framework for OT-based mapping across temporal, spatial, and multimodal single-cell datasets (Klein et al., 2025). These studies establish OT as a natural framework for matching heterogeneous cellular distributions. However, most existing approaches focus on couplings or deterministic transport maps for trajectory reconstruction, modality integration, or perturbation prediction, rather than learning a bidirectional stochastic path between bulk and single-cell transcriptomic distributions for downstream drug-response transfer.

Schrödinger Bridge-inspired Stochastic Distribution Transport. Schrödinger Bridge (SB) provides a stochastic perspective for connecting two prescribed probability distributions by modeling the most likely evolution of a reference diffusion process subject to endpoint constraints (Léonard, 2013). DSB extends this idea to high-dimensional data by parameterizing the transport with coupled forward and backward stochastic dynamics and optimizing them iteratively through procedures related to IPF (De Bortoli et al., 2021). Subsequent studies have further developed Schrödinger Bridge-based learning for generative modeling, image-to-image translation, and high-dimensional distribution matching, demonstrating its ability to capture complex nonlinear transitions between distributions (Liu et al., 2023; 2024). Compared with conventional domain adaptation methods that primarily focus on matching source and target endpoints, these approaches offer a path-based view in which intermediate distributional evolution is explicitly modeled. Inspired by this perspective, we reformulate bulk-to-single-cell adaptation as a stochastic distribution transport problem and adopt the bidirectional transport and iterative optimization principles of DSB to learn continuous transitions between bulk and single-cell latent distributions.

A.2 ALGORITHM

The detailed procedure of the proposed transport method is summarized in Algorithm 1.

Algorithm 1 Iterative Proportional Fitting for Diffusion Schrödinger Bridge and Deterministic Transport

Require: Source Bulk latents $\mathbf{z}^B \sim \mathbb{P}_B$, Target SC latents $\mathbf{z}^S \sim \mathbb{P}_S$ (standardized), Diffusion steps N, Step size $\Delta t = 1/N$, Diffusion variance σ^2

Ensure: Optimized backward drift network $\mathbf{v}_{\theta_b}^{(K)}$, Transported SC features $\hat{\mathbf{z}}^{S_{like}}$

Step 1: IPF Optimization for Diffusion Bridge

$\mathbf{z}_t \leftarrow (1-t)\mathbf{z}^B + t\mathbf{z}^S, \quad \forall t \in \{0, \Delta t, \dots, 1\}$ ▷ Initialize deterministic linear bridge ($k = 0$)

for iteration $k = 1, 2, \dots, K$ **do**

 # Phase 1: Backward Rollout & Forward Network Update

$\mathbf{z}_1 \leftarrow \mathbf{z}^S$

for $n = N$ **downto** 1 **do** ▷ Time $t = n\Delta t$

$\epsilon \sim \mathcal{N}(0, \mathbf{I})$

$\mathbf{z}_{t-\Delta t} \leftarrow \mathbf{z}_t - \mathbf{v}_{\theta_b}^{(k-1)}(\mathbf{z}_t, t)\Delta t + \sqrt{2\sigma^2 \Delta t}\,\epsilon$ ▷ Euler-Maruyama backward simulation

end for

$\mathcal{L}_{\text{fwd}} \leftarrow \mathbb{E}_{t,\mathbf{z}_t}\left[\left\|\mathbf{v}_{\theta_f}(\mathbf{z}_t, t) - \frac{\mathbf{z}_{t+\Delta t} - \mathbf{z}_t}{\Delta t}\right\|_2^2\right]$ ▷ Empirical velocity matching

$\theta_f^{(k)} \leftarrow \theta_f^{(k-1)} - \eta \nabla_{\theta_f}\big(\mathcal{L}_{\text{fwd}} + \lambda \mathcal{L}_{\text{cyc}}(\theta_f, \theta_b)\big)$ ▷ Update forward drift with cycle constraint

 # Phase 2: Forward Rollout & Backward Network Update

$\mathbf{z}_0 \leftarrow \mathbf{z}^B$

for $n = 0$ **to** $N - 1$ **do** ▷ Time $t = n\Delta t$

$\epsilon \sim \mathcal{N}(0, \mathbf{I})$

$\mathbf{z}_{t+\Delta t} \leftarrow \mathbf{z}_t + \mathbf{v}_{\theta_f}^{(k)}(\mathbf{z}_t, t)\Delta t + \sqrt{2\sigma^2 \Delta t}\,\epsilon$ ▷ Euler-Maruyama forward simulation

end for

$\mathcal{L}_{\text{bwd}} \leftarrow \mathbb{E}_{t,\mathbf{z}_t}\left[\left\|\mathbf{v}_{\theta_b}(\mathbf{z}_t, t) - \frac{\mathbf{z}_{t-\Delta t} - \mathbf{z}_t}{\Delta t}\right\|_2^2\right]$ ▷ Empirical velocity matching

$\theta_b^{(k)} \leftarrow \theta_b^{(k-1)} - \eta \nabla_{\theta_b}\big(\mathcal{L}_{\text{bwd}} + \lambda \mathcal{L}_{\text{cyc}}(\theta_f, \theta_b)\big)$ ▷ Update backward drift with cycle constraint

end for

Step 2: Deterministic Transport and Denormalization

$\tilde{\mathbf{z}}^S \leftarrow (\mathbf{z}^S_{\text{train}} - \mu_S)/\sigma_S$ ▷ Standardize SC embeddings dimension-wise

$\hat{\mathbf{z}}_1 \leftarrow \tilde{\mathbf{z}}^S$

for $n = N$ **downto** 1 **do** ▷ Time $t = n\Delta t$

$\hat{\mathbf{z}}_{t-\Delta t} \leftarrow \hat{\mathbf{z}}_t - \mathbf{v}_{\theta_b}^{(K)}(\hat{\mathbf{z}}_t, t)\Delta t$ ▷ Noise-free backward operator $\mathcal{T}_b$ (var $\sigma^2 = 0$)

end for

$\hat{\mathbf{z}}^{S_{like}} \leftarrow \hat{\mathbf{z}}_0 \odot \sigma_B + \mu_B$ ▷ Domain statistics restoration via denormalization

return $\mathbf{v}_{\theta_b}^{(K)}, \hat{\mathbf{z}}^{S_{like}}$

A.3 EXPERIMENT DETAILS

A.3.1 DATASET DETAILS AND EVALUATION PROTOCOL

Dataset Details. We use 12 publicly available scRNA-seq drug-response datasets from the NCBI Gene Expression Omnibus (GEO). These datasets cover multiple cancer types and anticancer drugs and contain experimentally annotated sensitive and resistant cell populations. GSE117872 includes two distinct cohorts, HN120 and HN137, which we treat as separate datasets. Table A.1 summarizes the datasets used in our experiments.

Table A.1: Summary of the 12 scRNA-seq drug-response datasets used in this study, including dataset accession, drug treatment, numbers of sensitive and resistant cells, total number of cells, cancer type, and original reference.

Dataset	Drug	Sensitive	Resistant	Number	Cancer type	Reference
GSE108397	PLX.4720	1022	1880	2902	Melanoma	(Ho et al., 2018)
GSE110894	I.BET.762	734	685	1419	Acute myeloid leukemia	(Bell et al., 2019)
GSE112274	GEFITINIB	37	470	507	Non-small cell lung cancer	(Kong et al., 2019)
GSE117872_120	Cisplatin	361	187	548	Oral squamous cell carcinoma	(Sharma et al., 2018)
GSE117872_137	Cisplatin	403	165	568	Oral squamous cell carcinoma	(Sharma et al., 2018)
GSE134836	ERLOTINIB	399	381	780	Non-small cell lung carcinoma	(Aissa et al., 2021)
GSE140440	DOCETAXEL	162	162	324	Prostate cancer	(Schnepp et al., 2020)
GSE149383	ERLOTINIB	1201	1529	2730	Non-small cell lung cancer	(Aissa et al., 2021)
GSE160244	Osimertinib	1228	375	1603	Non-small cell lung carcinoma	(Aissa et al., 2021)
GSE181812	TAMOXIFEN	332	110	442	ER positive breast cancer	(Semina et al., 2022)
GSE186960	GEMCITABINE	2647	1208	3855	Pancreatic Ductal Adenocarcinoma	(Principe et al., 2022)
GSE241225	PALBOCICLIB	114	506	620	Liposarcoma	(Gleason et al., 2024)

Evaluation Protocol. For each single-cell dataset, we perform a stratified split based on the drug response labels, with 80% of the cells used for training and 20% reserved for testing. An additional 10% of the training set is used for validation during β-VAE pretraining, resulting in approximately 72%, 8%, and 20% of the samples for training, validation, and testing, respectively. Single-cell response labels are used only for data splitting and final evaluation, and are not involved in β-VAE, CORAL, or DSB/IPF optimization.

For the bulk data, 10% of the samples are reserved for testing, and 10% of the remaining samples are used for validation, resulting in approximately 81%, 9%, and 10% of the samples for training, validation, and testing, respectively. We conduct five independent runs using different random seeds and report the mean and standard deviation of AUROC, AUPR, and F1 score.

A.3.2 BASELINES

In the following, we provide a detailed overview of the comparative models:

scDEAL (Chen et al., 2022). A deep transfer learning framework for predicting drug responses at the single-cell level by transferring pharmacogenomic knowledge from large-scale bulk RNA-seq data. It employs separate denoising autoencoders to extract latent representations from bulk and single-cell expression profiles, and aligns the two domains through maximum mean discrepancy while preserving cellular heterogeneity with cell-cluster regularization. The learned bulk drug-response predictor is subsequently transferred to scRNA-seq data for single-cell drug sensitivity prediction.

SCAD (Zheng et al., 2023). An adversarial domain adaptation framework for transferring drug-response knowledge from bulk RNA-seq cell lines to single-cell transcriptomic data. It jointly trains a feature extractor and a domain discriminator to learn domain-invariant representations between bulk and single-cell samples, while a drug-response predictor is supervised by drug sensitivity labels from the bulk source domain. By reducing the domain shift between the two data modalities, SCAD enables drug sensitivity prediction for individual cells with limited single-cell pharmacogenomic annotations.

scATD (Zhou et al., 2025). A high-throughput transfer learning framework for single-cell drug response prediction that leverages pretrained transcriptomic foundation models. It extracts RNA-seq representations using models such as scFoundation and Geneformer, and employs bidirectional adaptive instance normalization to align bulk and single-cell feature distributions without requiring

additional parameter training for new patient datasets. It further incorporates knowledge distillation to reduce computational cost and feature-attribution methods to identify drug-resistance biomarkers at cohort, patient, and single-cell levels.

DAGFormer (Yan et al., 2025). A graph-based domain adaptation framework for single-cell cancer drug response prediction by integrating bulk RNA-seq and scRNA-seq data. It constructs cellular neighbor graphs to model intercellular relationships and employs graph domain adaptation with adversarial training to learn domain-invariant representations across bulk and single-cell domains. A dual-domain decoupling mechanism further separates shared and modality-specific features, enabling effective transfer of bulk drug-response knowledge while preserving biological characteristics specific to each domain.

DeepSADR (Zhang et al., 2026). A deep transfer learning framework for transferring drug-response knowledge from cancer cell lines to patients by explicitly modeling fine-grained drug-genomic interactions. It decomposes drugs and genomic profiles into biologically meaningful subsequences and constructs bipartite interaction graphs between drug substructures and gene functional subsequences, which are encoded using a supervised graph autoencoder. A Set Transformer-based adaptive readout then learns domain-invariant drug-response representations, facilitating effective knowledge transfer from large-scale cell-line data to scarce patient cohorts.

DSN-DANN (Bousmalis et al., 2016; Ganin et al., 2016). A domain adaptation framework that combines Domain Separation Networks with domain-adversarial learning to disentangle domain-invariant and domain-specific information. It employs a shared encoder to capture representations common to bulk and single-cell data together with domain-specific private encoders, while reconstruction and orthogonality constraints encourage the shared and private components to preserve complementary information. A domain discriminator trained adversarially on the shared representation further suppresses modality-specific signals, allowing the drug-response predictor supervised on labeled bulk samples to operate on features that are more transferable to the single-cell domain.

Optimal Transport (OT) (Bunne et al., 2023). A distribution-matching framework for domain adaptation that aligns source and target domains by explicitly minimizing the cost of transporting probability mass between their empirical distributions. It estimates an optimal transport plan that establishes soft correspondences between samples across domains while favoring mappings between nearby representations under a predefined ground cost. The resulting coupling provides an aligned representation of bulk and single-cell transcriptomic distributions, enabling drug-response knowledge learned from labeled bulk samples to be transferred to the unlabeled single-cell domain.

GFK (Geodesic Flow Kernel) (Gong et al., 2012). An unsupervised domain adaptation method that models domain shift as a continuous transition between low-dimensional source and target subspaces. It represents the source and target subspaces as points on a Grassmann manifold and integrates an infinite sequence of intermediate subspaces along the geodesic connecting them, yielding a geodesic flow kernel that captures gradual changes in domain geometry. By expressing bulk and single-cell samples through this domain-adaptive kernel, GFK reduces distributional discrepancies while allowing drug-response information from the source domain to generalize to the target domain.

DAE (Denoising Autoencoder) (Vincent et al., 2008). A robust representation-learning baseline that learns low-dimensional features by reconstructing clean transcriptomic profiles from stochastically corrupted inputs. The denoising objective encourages the encoder to capture stable and informative structure in gene-expression data rather than noise or input-specific variations, producing representations that are more robust to perturbations and measurement variability. These learned latent features are subsequently used for downstream drug response prediction, providing a non-adversarial baseline for evaluating the contribution of explicit cross-domain alignment.

A.3.3 Detailed Performance Comparison with Baseline Methods

The detailed performance results of scBridge and all baseline methods across the 12 datasets are reported in Tables A.2–A.4, including AUROC, AUPR, and F1 scores, respectively.

Table A.2: AUROC scores of scBridge and baseline methods across 12 datasets.

Dataset	**scBridge**	DeepSADR	DAGFormer	scATD	SCAD	scDEAL	DSN-DANN	OT	GFK	DAE
GSE108397	0.960	0.960	0.895	0.893	0.230	0.850	0.636	0.516	0.584	0.518
GSE110894	0.950	0.944	0.721	0.946	0.529	0.950	0.501	0.490	0.447	0.398
GSE112274	0.970	0.862	0.922	0.744	0.728	0.940	0.735	0.690	0.677	0.622
GSE117872_120	0.850	0.828	0.356	0.253	0.215	0.830	0.566	0.510	0.512	0.466
GSE117872_137	0.910	0.813	0.314	0.325	0.491	0.790	0.436	0.510	0.386	0.341
GSE134836	0.860	0.700	0.559	0.335	0.535	0.690	0.597	0.490	0.535	0.487
GSE140440	0.920	0.692	0.723	0.692	0.953	0.880	0.681	0.480	0.634	0.571
GSE149383	0.750	0.807	0.579	0.725	0.491	0.890	0.507	0.513	0.461	0.422
GSE160244	0.820	0.843	0.408	0.534	0.663	0.520	0.462	0.670	0.401	0.359
GSE181812	0.740	0.570	0.614	0.497	0.553	0.690	0.382	0.523	0.338	0.286
GSE186960	0.750	0.667	0.458	0.481	0.534	0.640	0.417	0.500	0.379	0.329
GSE241225	0.920	0.658	0.665	0.512	0.692	0.840	0.653	0.505	0.612	0.548
AVG	0.867	0.779	0.601	0.578	0.551	0.793	0.548	0.533	0.497	0.446

Table A.3: AUPR scores of scBridge and baseline methods across 12 datasets.

Dataset	**scBridge**	DeepSADR	DAGFormer	scATD	SCAD	scDEAL	DSN-DANN	OT	GFK	DAE
GSE108397	0.880	0.933	0.808	0.828	0.251	0.820	0.584	0.360	0.539	0.472
GSE110894	0.970	0.940	0.755	0.971	0.588	0.950	0.552	0.514	0.412	0.361
GSE112274	0.890	0.346	0.340	0.249	0.258	0.850	0.671	0.680	0.628	0.583
GSE117872_120	0.940	0.887	0.571	0.509	0.501	0.890	0.629	0.660	0.571	0.442
GSE117872_137	0.970	0.909	0.590	0.634	0.757	0.900	0.502	0.710	0.451	0.395
GSE134836	0.810	0.703	0.673	0.421	0.522	0.710	0.548	0.510	0.493	0.451
GSE140440	0.920	0.686	0.820	0.808	0.942	0.900	0.623	0.480	0.578	0.529
GSE149383	0.710	0.769	0.497	0.633	0.418	0.890	0.451	0.446	0.424	0.319
GSE160244	0.940	0.941	0.719	0.805	0.883	0.620	0.529	0.765	0.476	0.431
GSE181812	0.890	0.797	0.858	0.803	0.809	0.730	0.447	0.759	0.395	0.334
GSE186960	0.870	0.822	0.686	0.663	0.689	0.650	0.483	0.686	0.441	0.382
GSE241225	0.740	0.471	0.282	0.176	0.319	0.710	0.579	0.180	0.558	0.497
AVG	0.878	0.767	0.633	0.625	0.578	0.802	0.550	0.563	0.497	0.433

Table A.4: F1 scores of scBridge and baseline methods across 12 datasets.

Dataset	**scBridge**	DeepSADR	DAGFormer	scATD	SCAD	scDEAL	DSN-DANN	OT	GFK	DAE
GSE108397	0.800	0.896	0.813	0.712	0.391	0.790	0.498	0.510	0.452	0.387
GSE110894	0.870	0.873	0.415	0.641	0.437	0.960	0.312	0.310	0.263	0.204
GSE112274	0.970	0.907	0.424	0.635	0.619	0.950	0.592	0.720	0.536	0.491
GSE117872_120	0.830	0.770	0.351	0.177	0.177	0.820	0.417	0.650	0.372	0.328
GSE117872_137	0.760	0.786	0.523	0.159	0.389	0.750	0.238	0.310	0.192	0.161
GSE134836	0.770	0.638	0.465	0.378	0.580	0.660	0.456	0.320	0.398	0.352
GSE140440	0.760	0.613	0.447	0.350	0.325	0.890	0.549	0.340	0.501	0.439
GSE149383	0.700	0.712	0.538	0.403	0.485	0.890	0.331	0.403	0.287	0.241
GSE160244	0.680	0.813	0.649	0.664	0.664	0.510	0.276	0.580	0.223	0.195
GSE181812	0.640	0.643	0.100	0.100	0.410	0.600	0.196	0.410	0.154	0.118
GSE186960	0.650	0.661	0.507	0.558	0.503	0.550	0.251	0.150	0.216	0.179
GSE241225	0.810	0.767	0.741	0.711	0.054	0.720	0.514	0.741	0.476	0.406
AVG	0.770	0.757	0.498	0.457	0.420	0.758	0.386	0.454	0.339	0.292

A.3.4 Top-K Precision Analysis

To further evaluate whether scBridge assigns high confidence to genuinely drug-sensitive cells, we perform a Top-K precision analysis across all 12 single-cell datasets, as shown in Figure A.1. Overall, the precision curves remain substantially above the prevalence-based random-selection baseline over the high-confidence portion of the ranking, indicating that sensitive cells are preferentially concentrated among samples receiving the highest predicted sensitivity probabilities. At the top 10% operating point, scBridge achieves a mean precision of 0.90 across datasets (range: 0.55-1.00), with five datasets reaching a precision of 1.00. This enrichment remains evident when the selection fraction is matched to the prevalence of sensitive cells in each dataset: the mean precision is 0.79, compared with a mean prevalence of 0.51. The effect is particularly pronounced in highly imbalanced datasets. For example, in GSE112274, where sensitive cells constitute only 6.9% of the population,

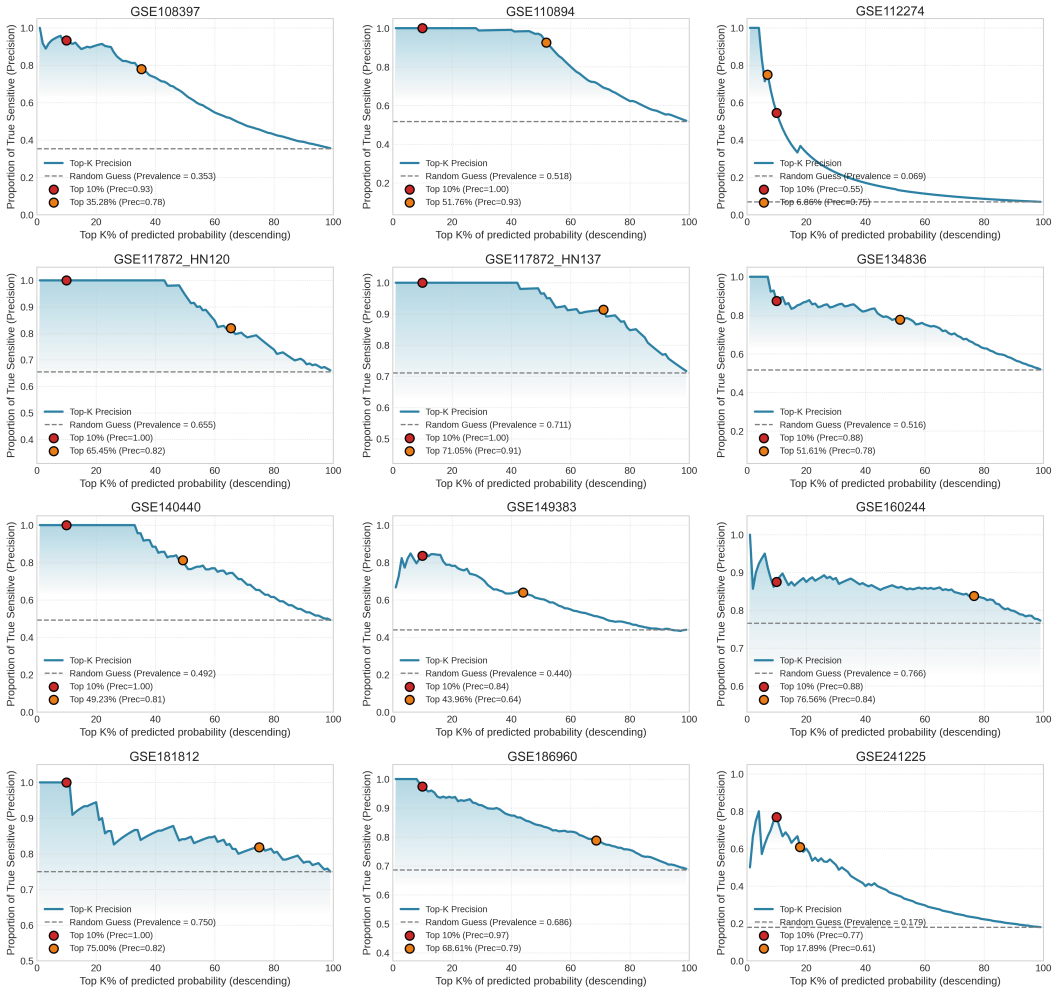

Figure A.1: Top-K precision analysis of scBridge across 12 scRNA-seq datasets.

the top 10% predictions attain a precision of 0.55, while selecting the top 6.9% of cells yields a precision of 0.75. Similarly, for GSE241225, the sensitive-cell prevalence is 17.9%, whereas the top 10% precision reaches 0.77 and the prevalence-matched precision reaches 0.61. As K increases, precision gradually approaches the dataset prevalence, as expected when progressively lower-confidence cells are included. Together, these results indicate that scBridge provides not only accurate global discrimination but also effective ranking of high-confidence drug-sensitive cells, which is particularly relevant for prioritizing candidate responsive subpopulations in class-imbalanced single-cell settings.

A.3.5 Ablation Study

Detailed results of the ablation experiments are reported in Table A.5.

A.3.6 Details of the Learned Transport Bridge Analysis

For the transport analysis in Figure 3, bulk and single-cell samples are first represented in the β-VAE latent space. For Figure 3A, we construct a fixed-reference UMAP space and project the single-cell representations before and after deterministic SC-to-Bulk transport into the same coordinate system, enabling direct comparison of cross-domain mixing. Domain separation is quantified using the average silhouette width (Domain ASW) based on Bulk/SC domain labels, where values closer

Table A.5: Ablation study of scBridge across 12 datasets.

Dataset	w/o DSB			w/o CORAL			w/o Cycle			w/o IPF		
	AUROC	AUPR	F1	AUROC	AUPR	F1	AUROC	AUPR	F1	AUROC	AUPR	F1
GSE108397	0.89	0.79	0.78	0.91	0.82	0.80	0.92	0.84	0.81	0.87	0.77	0.75
GSE110894	0.64	0.59	0.60	0.89	0.91	0.81	0.91	0.93	0.83	0.83	0.83	0.72
GSE112274	0.62	0.68	0.57	0.83	0.83	0.81	0.94	0.85	0.93	0.78	0.74	0.65
GSE117872	0.78	0.85	0.73	0.70	0.87	0.75	0.82	0.90	0.78	0.79	0.72	0.65
GSE117872	0.82	0.87	0.65	0.84	0.89	0.68	0.87	0.93	0.71	0.88	0.85	0.61
GSE134836	0.74	0.72	0.67	0.87	0.73	0.69	0.82	0.76	0.72	0.74	0.72	0.63
GSE140440	0.84	0.73	0.74	0.89	0.83	0.74	0.78	0.88	0.79	0.82	0.81	0.74
GSE149383	0.68	0.61	0.60	0.76	0.63	0.62	0.71	0.67	0.65	0.63	0.60	0.57
GSE160244	0.44	0.35	0.46	0.76	0.87	0.60	0.79	0.91	0.63	0.78	0.83	0.53
GSE181812	0.46	0.49	0.41	0.70	0.81	0.56	0.70	0.85	0.59	0.63	0.77	0.50
GSE186960	0.67	0.77	0.55	0.69	0.67	0.57	0.71	0.83	0.60	0.54	0.53	0.49
GSE241225	0.79	0.65	0.71	0.88	0.68	0.73	0.87	0.70	0.77	0.81	0.63	0.69

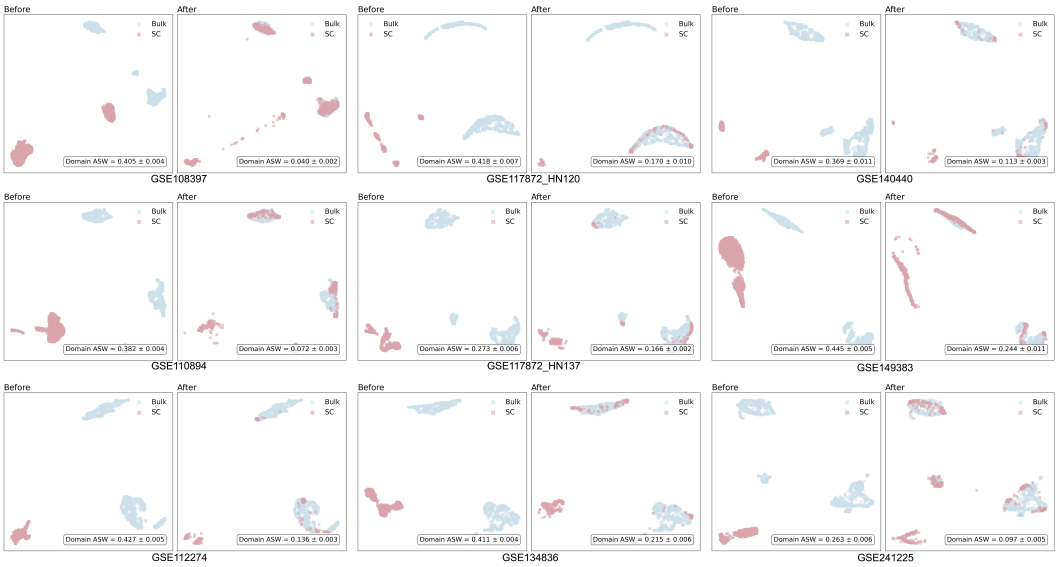

Figure A.2: Cross-domain transport visualization across the remaining nine scRNA-seq datasets.

to zero indicate stronger domain mixing. For Figure 3B, we explicitly trace the deterministic SC-to-Bulk transport induced by the learned backward drift field. Starting from each single-cell latent representation, the state is propagated through the learned backward dynamics over the discrete transport steps, with the stochastic noise term disabled to isolate the learned transport direction. The intermediate latent state at every step is retained, producing a complete trajectory from the initial single-cell representation toward the bulk domain. All intermediate states are subsequently projected into the same fixed-reference UMAP space, allowing the direction, path length, and cell-to-cell heterogeneity of the learned transport to be visualized directly. For clarity, up to 80 cells are randomly sampled for trajectory visualization, with trajectory color indicating transport progress from the single-cell starting state toward the bulk-domain endpoint. For Figure 3C, we quantify the Euclidean displacement between consecutive transport states for each cell and summarize the mean displacement and its variation across cells as a function of transport progress. This provides a quantitative view of how the magnitude of the learned correction evolves along the transport path.

To complement the representative datasets, Figure A.2 shows the before-and-after transport visualization for the remaining nine datasets. Consistent reductions in Domain ASW are observed across all datasets, accompanied by increased overlap between the transported single-cell and bulk representations. The magnitude of the residual domain discrepancy varies across datasets, further supporting the dataset-dependent transport behavior observed in the main analysis.

A.3.7 ROBUSTNESS ANALYSIS

We further supplement the robustness analysis with detailed per-dataset results in Tables A.6 and A.7, providing a comprehensive evaluation of the model's performance under varying levels of target-domain data scarcity and measurement noise.

Table A.6: Robustness of scBridge to reduced single-cell training data across 12 datasets.

Dataset	100%			80%			50%			30%			10%		
	F1	AUROC	AUPR	F1	AUROC	AUPR	F1	AUROC	AUPR	F1	AUROC	AUPR	F1	AUROC	AUPR
GSE108397	0.800	0.960	0.880	0.809	0.928	0.847	0.797	0.898	0.827	0.712	0.899	0.770	0.673	0.899	0.728
GSE110894	0.870	0.950	0.970	0.890	0.894	0.940	0.837	0.908	0.954	0.820	0.869	0.904	0.782	0.839	0.930
GSE112274	0.970	0.970	0.890	0.980	0.937	0.883	0.965	0.953	0.858	0.919	0.873	0.867	0.896	0.858	0.747
GSE117872	0.830	0.850	0.940	0.803	0.860	0.909	0.786	0.849	0.938	0.698	0.718	0.869	0.673	0.714	0.856
GSE117872	0.760	0.910	0.970	0.769	0.896	0.933	0.692	0.840	0.925	0.623	0.829	0.848	0.642	0.796	0.841
GSE134836	0.770	0.860	0.810	0.726	0.858	0.809	0.685	0.842	0.780	0.684	0.825	0.661	0.602	0.662	0.605
GSE140440	0.760	0.920	0.920	0.843	0.917	0.918	0.807	0.846	0.858	0.761	0.866	0.778	0.734	0.770	0.750
GSE149383	0.700	0.750	0.710	0.724	0.752	0.670	0.659	0.690	0.602	0.616	0.638	0.563	0.356	0.593	0.528
GSE160244	0.680	0.820	0.940	0.637	0.837	0.936	0.655	0.762	0.853	0.568	0.649	0.870	0.528	0.561	0.857
GSE181812	0.640	0.740	0.890	0.628	0.755	0.900	0.498	0.720	0.830	0.423	0.670	0.835	0.450	0.620	0.777
GSE186960	0.650	0.750	0.870	0.627	0.724	0.843	0.634	0.682	0.813	0.540	0.652	0.706	0.441	0.564	0.604
GSE241225	0.810	0.920	0.740	0.768	0.889	0.762	0.792	0.885	0.709	0.808	0.861	0.592	0.626	0.811	0.549

Table A.7: Robustness of scBridge to measurement noise in the single-cell domain across 12 datasets.

Dataset	0			0.1			0.2			0.3			0.5		
	F1	AUROC	AUPR	F1	AUROC	AUPR	F1	AUROC	AUPR	F1	AUROC	AUPR	F1	AUROC	AUPR
GSE108397	0.800	0.960	0.880	0.862	0.981	0.886	0.814	0.921	0.857	0.810	0.918	0.806	0.745	0.837	0.762
GSE110894	0.870	0.950	0.970	0.847	0.924	0.977	0.781	0.899	0.981	0.759	0.864	0.865	0.820	0.840	0.933
GSE112274	0.970	0.970	0.890	0.899	0.980	0.845	0.964	0.899	0.862	0.913	0.957	0.821	0.847	0.853	0.814
GSE117872	0.830	0.850	0.940	0.823	0.860	0.924	0.772	0.821	0.931	0.756	0.833	0.817	0.713	0.721	0.828
GSE117872	0.760	0.910	0.970	0.771	0.936	0.930	0.746	0.873	0.991	0.660	0.864	0.907	0.633	0.845	0.894
GSE134836	0.770	0.860	0.810	0.752	0.854	0.806	0.707	0.800	0.736	0.648	0.789	0.724	0.622	0.770	0.616
GSE140440	0.760	0.920	0.920	0.815	0.947	0.884	0.827	0.915	0.930	0.758	0.894	0.850	0.759	0.825	0.831
GSE149383	0.700	0.750	0.710	0.681	0.700	0.768	0.700	0.744	0.681	0.631	0.569	0.595	0.571	0.633	0.481
GSE160244	0.680	0.820	0.940	0.694	0.783	0.959	0.630	0.810	0.894	0.580	0.765	0.881	0.506	0.654	0.892
GSE181812	0.640	0.740	0.890	0.612	0.695	0.863	0.620	0.691	0.822	0.556	0.602	0.841	0.471	0.604	0.749
GSE186960	0.650	0.750	0.870	0.692	0.737	0.847	0.658	0.752	0.821	0.561	0.617	0.810	0.494	0.576	0.705
GSE241225	0.810	0.920	0.740	0.819	0.916	0.762	0.760	0.930	0.620	0.769	0.807	0.658	0.631	0.771	0.553