PERTURBRX: LEARNING TREATMENT-CONDITIONED LATENT TRANSITIONS FROM SINGLE-CELL PERTURBATIONS FOR PATIENT DRUG RESPONSE

Anonymous authors
Paper under double-blind review

ABSTRACT

Scarce data and tumor heterogeneity limit patient-level prediction of cancer treatment response. Many existing approaches predict response from patient profiles and drug representations without explicitly modeling treatment-induced molecular changes. We propose PerturbRx, a treatment-conditioned representation learning framework that learns drug- and dose-conditioned latent transitions from context-matched control and treated single-cell populations. The learned transition operator generalizes to compounds excluded from source-stage training, predicting held-out treatment effects more accurately than baselines (mean cosine similarity 0.229 vs. 0.198). We freeze and transfer the operator to pretreatment patient profiles, using the learned treatment-induced transformation as an additional representation of how each patient state may respond to therapy, without requiring matched post-treatment measurements. The transferred representation achieves the strongest pooled out-of-fold performance among evaluated methods on TCGA-485 (AUROC/AUPRC: 0.709/0.794) and an independent PDX cohort (0.685/0.612), with promising improvements beyond strong patient+drug controls in both domains.

1 INTRODUCTION

Scarce patient cohorts with molecular profiles linked to treatment-specific clinical outcomes and tumor heterogeneity limit patient-level cancer treatment-response prediction (Partin et al., 2023). Many methods therefore leverage large-scale cancer cell-line screens (Yang et al., 2013), despite substantial differences between cell lines and patient tumors (Gillet et al., 2013; Ben-David et al., 2017). Many approaches emphasize transferable representations, domain alignment, and response adaptation (He et al., 2022; Zhang et al., 2026), but do not explicitly model how a patient's molecular state changes under treatment. Such changes may provide complementary information, yet matched pre- and post-treatment patient profiles are rarely available at scale. Single-cell perturbation atlases provide context-matched control and treated populations across many interventions (Srivatsan et al., 2020; Zhang et al., 2025), enabling learning of treatment-induced state changes.

We introduce PerturbRx, a treatment-aware representation model that learns a drug- and dose-conditioned transition operator in latent space. Given a cellular state and intervention, PerturbRx predicts a latent transition rather than reconstructing post-treatment expression, following a joint-embedding predictive (world-model-style) formulation (LeCun, 2022). Because destructive single-cell assays do not provide paired before-and-after cells, we learn the operator through distributional supervision. We then freeze and transfer it to target profiles for response prediction.

Our transfer hypothesis is that intervention-conditioned structure, once learned in a shared transcriptomic representation space, may remain informative across measurement scales and biological domains, even as the source and target distributions shift. We test three questions:

1) Can transitions be learned from perturbation data and generalized to held-out compounds?
2) Do transferred representations add predictive signal beyond patient and drug representations?
3) Does the transferred representation remain informative across distinct target domains?

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.

We find that (1) transitions can be learned and generalized to held-out compounds; (2) the transferred representation achieves the strongest pooled out-of-fold performance among the evaluated methods on both TCGA-485 and PDX, with incremental gains beyond strong patient+drug controls; (3) the representation transfers across distinct target domains.

We instantiate PerturbRx using the Tahoe-100M single-cell perturbation dataset (Zhang et al., 2025), scFoundation embeddings (Hao et al., 2024), and ChemBERTa drug representations (Chithrananda et al., 2020). We evaluate response prediction in The Cancer Genome Atlas (TCGA) (The Cancer Genome Atlas Research Network et al., 2013) and a patient-derived xenograft (PDX) cohort (Gao et al., 2015). Our contributions are as follows:

- We introduce a distributionally supervised framework for learning drug- and dose-conditioned latent transitions from context-matched control and treated single-cell populations without cell-level pairing, and transfer the frozen operator to patient and PDX profiles.
- We demonstrate cross-scale transfer on TCGA and an independent PDX benchmark, where the transferred representation improves response prediction beyond strong patient+drug controls and remains informative across distinct target domains.
- We isolate the learned transition operator's contribution through objective- and data-matched source-architecture controls and show that its predictive value extends to treatments without exact source-stage support.

2 Related Work

Patient Drug Response Prediction and Preclinical-to-Patient Transfer. Many preclinical-to-patient drug-response methods transfer cell-line supervision or response-oriented representations through domain adaptation (CODE-AE (He et al., 2022); TransDRP (Liu & Li, 2025)), weak supervision (WISER (Shubham et al., 2024)), generative augmentation (GANDALF (Jayagopal et al., 2025)), or representation transfer (DeepSADR (Zhang et al., 2026)).

A complementary line of work incorporates treatment-induced molecular information. Dr.VAE (Rampášek et al., 2019) jointly models pre- and post-treatment expression with drug sensitivity, while CSG^{2}A (Bang et al., 2024) learns condition-specific gene interactions from LINCS L1000 for perturbation and response prediction. THERAPI (Sung et al., 2026) and PREDIKTOR (Bang et al., 2026) reuse pretrained CSG^{2}A for patient prediction.

In contrast, PerturbRx makes the treatment-induced state transition itself the object of transfer. It learns a drug- and dose-conditioned latent transition operator from control and treated single-cell populations and transfers it across scales to patient and PDX baseline tumor profiles for drug-response prediction. Target-domain prediction requires only pretreatment profiles and drug identity/dose, not post-treatment measurements.

Single-Cell Perturbation Modeling and Cross-Scale Transfer. Because single-cell measurements are destructive, perturbation models infer treatment effects from populations rather than paired cells. scGen (Lotfollahi et al., 2019) models perturbation effects as latent-space changes, while CPA (Lotfollahi et al., 2023) conditions responses on perturbations, doses, time points, and biological contexts. CellOT (Bunne et al., 2023) learns an optimal transport map between control and perturbed populations. CMonge (Driessen et al., 2026) learns conditional transport maps across drugs, doses, and cell types, enabling generalization across perturbations. Recent generative approaches include PerturbDiff (Yuan et al., 2026), which models responses in distribution space, and scPPDM (Liang et al., 2025), which uses latent diffusion conditioned on baseline state, drug structure, and dose to predict post-perturbation expression on Tahoe-100M.

In contrast, **PerturbRx uses single-cell perturbation data not as the final prediction domain, but to supervise learning a transferable treatment-transition operator**. The operator is then applied to pretreatment patient and PDX tumor states to predict therapeutic response, creating a direct bridge from cellular perturbation modeling to patient-level drug-response prediction.

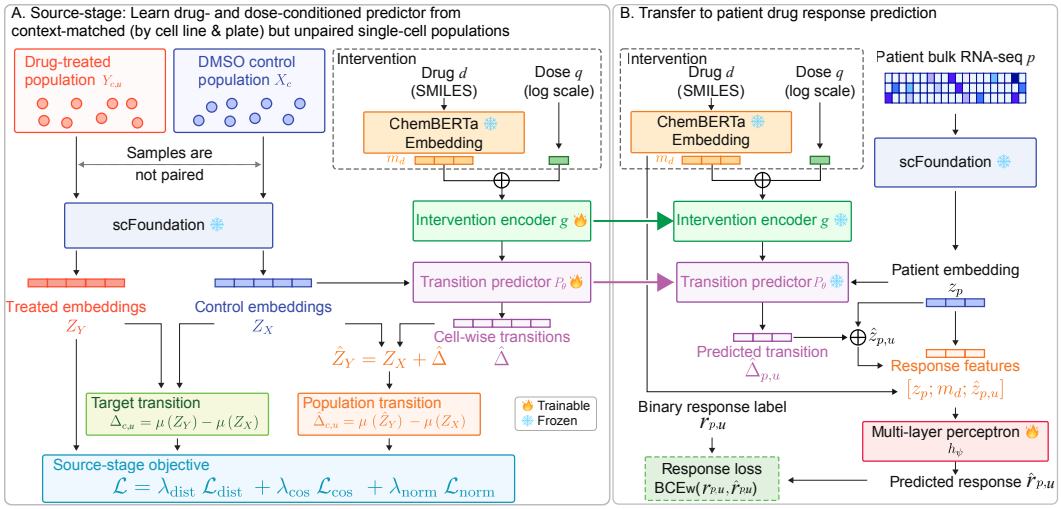


Figure 1: **Overview of PerturbRx. (A)** Source-stage learning from DMSO control and drug-treated single-cell populations matched by cell line and plate, but without cell-level pairing. A drug- and dose-conditioned predictor maps control embeddings and intervention features to cell-wise transitions, which are added to the controls to form post-intervention embeddings. We train the predictor using distributional, transition-direction, and transition-magnitude objectives. **(B)** Transfer of the transition predictor to patient response prediction. The predictor generates a treatment-conditioned representation, which is combined with patient and drug features for response prediction.

3 Method

Figure 1 summarizes the two-stage PerturbRx framework: learning an intervention-conditioned latent transition operator from perturbation populations, followed by frozen transfer to target-domain response prediction. We represent molecular states by gene-expression profiles, measured at single-cell resolution (scRNA-seq) in the source domain and as bulk aggregates (bulk RNA-seq) in the target domain. Source perturbation experiments provide separately observed control and treated cell populations X_c and $Y_{c,u}$ (defined below), rather than before-and-after measurements of the same cells. From these populations, we learn a treatment-conditioned transition operator P_θ. The frozen operator is then applied to target profiles x_p to predict binary treatment response $r_{p,u} \in \{0, 1\}$. Thus, source-stage learning uses perturbation supervision rather than drug-response labels; response supervision enters only when training the target-domain response head.

3.1 Source and Target Formulation

For source-stage training, we define an experimental context by cell line ℓ and plate b, and an intervention by drug d and dose q:

$$c = (\ell, b), \qquad u = (d, q). \tag{1}$$

DMSO serves as the vehicle control. For each condition (ℓ, b, d, q), we construct a context-matched control-treated population pair:

$$X_c = \{x_i : \text{cell line}(x_i) = \ell,\ \text{plate}(x_i) = b,\ \text{drug}(x_i) = \text{DMSO}\}, \tag{2}$$

$$Y_{c,u} = \{y_j : \text{cell line}(y_j) = \ell,\ \text{plate}(y_j) = b,\ \text{drug}(y_j) = d,\ \text{dose}(y_j) = q\}. \tag{3}$$

Here, x_i and y_j denote single-cell gene-expression profiles. The set $X_c = \{x_i\}_{i=1}^{n_c}$ denotes the DMSO control population for context c, and $Y_{c,u} = \{y_j\}_{j=1}^{m_{c,u}}$ denotes the population treated with intervention u in the same context. Their population sizes are $n_c = |X_c|$ and $m_{c,u} = |Y_{c,u}|$, respectively. We retain only conditions with sufficient control and treated cells.

In the target domain, we observe patient response records

$$\mathcal{D}_{\text{tgt}} = \{(x_p, u, r_{p,u})\}, \tag{4}$$

where x_p is the patient profile, $u = (d, q)$ is the intervention specification, and $r_{p,u}$ is the binary response label. Here, d is the observed treatment, whereas the target-domain value of q is treated as a representation hyperparameter rather than as an observed patient-specific exposure.

3.2 Intervention-Conditioned Latent Transition Predictor

Let E be a frozen pretrained encoder that maps cells into a d-dimensional latent space. scFoundation produces a concatenated cell embedding with $d = 3072$. We embed the two populations as

$$Z_X = E(X_c) = \{z_i^X\}_{i=1}^{n_c}, \qquad Z_Y = E(Y_{c,u}) = \{z_j^Y\}_{j=1}^{m_{c,u}}. \tag{5}$$

The intervention is encoded as

$$\alpha_u = g([m_d; \log q]), \tag{6}$$

where m_d is a frozen ChemBERTa embedding computed from the SMILES representation of drug d, q is the perturbation dose, log-transformed to represent multiplicative dose differences, and g is a trainable intervention encoder. The predictor $P_\theta : \mathbb{R}^d \times \mathbb{R}^r \to \mathbb{R}^d$, where r denotes the intervention-embedding dimension, maps a latent cell state and intervention embedding to a latent transition:

$$\hat{\Delta}_i = P_\theta(z_i^X, \alpha_u), \qquad \hat{z}_i^Y = z_i^X + \hat{\Delta}_i. \tag{7}$$

Importantly, $\hat{z}_i^Y$ is a treatment-conditioned representation of control cell i, not an estimate of any particular observed treated cell. The predicted treated population is $\hat{Z}_Y = \{\hat{z}_i^Y\}_{i=1}^{n_c}$.

$$\Delta_{c,u} = \mu(Z_Y) - \mu(Z_X), \qquad \hat{\Delta}_{c,u} = \frac{1}{n_c}\sum_{i=1}^{n_c}\hat{\Delta}_i = \mu\left(\hat{Z}_Y\right) - \mu\left(Z_X\right). \tag{8}$$

where $\mu(Z) = \frac{1}{|Z|}\sum_{z\in Z} z$. Architecture and optimization details are provided in Table 16.

3.3 Population-Level Training Objective

Maximum Mean Discrepancy (MMD) aligns the predicted and observed distributions but does not explicitly constrain the treatment-induced shift. We therefore complement MMD with cosine-based direction and L_2-norm magnitude constraints. The resulting source-stage objective is:

$$\mathcal{L}_{\text{src}} = \underbrace{\lambda_{\text{dist}}\text{MMD}^2(\hat{Z}_Y, Z_Y)}_{\text{distribution alignment}} + \lambda_{\cos}\underbrace{\left[1 - \cos\left(\hat{\Delta}_{c,u}, \Delta_{c,u}\right)\right]}_{\text{transition direction}} + \lambda_{\text{norm}}\underbrace{\left(\left\|\hat{\Delta}_{c,u}\right\|_2 - \|\Delta_{c,u}\|_2\right)^2}_{\text{transition magnitude}}. \tag{9}$$

We use $\lambda_{\text{dist}} = 1$, $\lambda_{\cos} = 0.1$ and $\lambda_{\text{norm}} = 10^{-3}$, selected using source-stage random-holdout validation only. We also evaluated a vector-displacement mean-squared error (MSE) term in ablations; adding it did not improve random-validation performance, so we excluded it from the final source objective. The same pattern was observed on the separate drug-held-out evaluation (Section 4.2). Details of the MMD estimator and objective selection are provided in Appendices B.2 and D, respectively.

3.4 Frozen Predictor Transfer to Patients

After source training, we freeze the intervention encoder and transition predictor and transfer them to the target. For a sample p with profile x_p and intervention $u = (d, q)$, the predictor produces

$$\hat{\Delta}_{p,u} = P_\theta(z_p, \alpha_u), \qquad \hat{z}_{p,u} = z_p + \hat{\Delta}_{p,u}, \tag{10}$$

where $z_p = E(x_p)$ is the scFoundation embedding of the target profile and α_u is the intervention embedding. In the target domain, q is treated as a representation hyperparameter selected using non-held-out data rather than as an estimate of patient-specific exposure (Appendix E).

The frozen scFoundation pipeline maps target bulk RNA-seq profiles into the same $d = 3072$-dimensional representation space used by the source-stage predictor. This shared representation does

not imply source–target alignment or guarantee that source-domain transition geometry is preserved; we therefore treat cross-domain applicability as an empirical hypothesis evaluated by downstream transfer.

We refer to $\hat{\Delta}_{p,u}$ as a patient-conditioned transition because it depends on the patient and treatment. Without post-treatment profiles, we treat $\hat{\Delta}_{p,u}$ and $\hat{z}_{p,u}$ as transferred latent features rather than validated changes or post-treatment states.

During target-domain training, the response head predicts $\hat{r}_{p,u} = h_\psi([z_p; m_d; \hat{z}_{p,u}])$. We optimize the class-weighted binary cross-entropy loss

$$\mathcal{L}_{\text{resp}} = -\frac{1}{N} \sum_{(p,u)} \left[w_+ r_{p,u} \log \hat{r}_{p,u} + (1 - r_{p,u}) \log(1 - \hat{r}_{p,u}) \right], \tag{11}$$

where $w_+ = N_-/N_+$, with N_+ and N_- denoting the numbers of positive and negative training examples, respectively.

4 Experiments

4.1 Experimental Setup

We train the source-stage PerturbRx model on Tahoe-100M using frozen scFoundation embeddings and ChemBERTa drug representations, then transfer the intervention encoder and transition predictor to TCGA-485 and PDX datasets. TCGA-485 comprises 485 episodes from 441 patients across five drugs; PDX comprises 791 episodes from 177 models across 10 drugs. Appendix A details cohort construction.

We compare against CODE-AE, TransDRP, DeepSADR, and WISER, whose released implementations we can evaluate under our patient-level benchmark protocol. We discuss perturbation-aware methods not directly compatible with our benchmarks, including CSG^{2}A, THERAPI, and PREDIKTOR, in Appendix B.6. We use disjoint outer partitions: patient-grouped folds for TCGA-485 and model-disjoint folds for PDX. All model selection uses only non-held-out data, and all methods share the same held-out partitions while retaining their native training procedures (Appendix B).

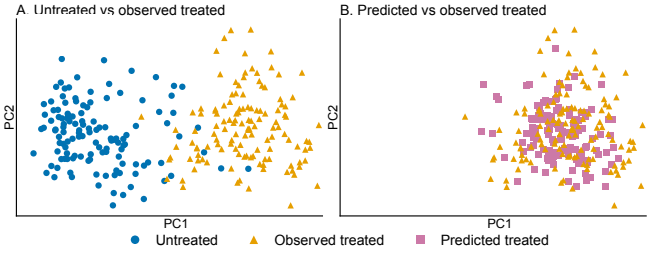
4.2 Learning Treatment-Induced Latent Transitions

We evaluate whether PerturbRx predicts treatment-induced latent transitions for held-out conditions and generalizes to compounds excluded from source-stage training. Figure 2 evaluates randomly held-out conditions and drug-held-out compounds. These splits probe distinct generalization regimes: random condition holdout primarily evaluates interpolation within the observed perturbation support, whereas drug-level holdout tests inductive transfer to compounds with no matched perturbation populations used during source-stage training. To isolate architectural effects, Linear, MLP, and conditional autoencoder (cAE) controls were trained in the same frozen scFoundation latent space using the same MMD + cosine + norm objective and matched source conditions (Appendix B.5).

On randomly held-out conditions, PerturbRx achieved the best aggregate performance across all three metrics among the evaluated predictors (Fig. 2C). Under the more stringent drug-held-out setting, PerturbRx achieved the highest mean transition cosine among the same-objective architectural controls: 0.229 versus 0.198 for cAE, 0.167 for MLP, and 0.137 for Linear (Fig. 2D). Across the 42 held-out drugs, the paired mean improvement over cAE was 0.032 (95% bootstrap CI: -0.020 to 0.086), with PerturbRx outperforming cAE for 23 of 42 compounds. The positive mean difference was preserved under leave-one-drug-out analysis (range: 0.020–0.040), indicating that the aggregate improvement was not attributable to a single compound, although performance varied substantially across drugs. Because we excluded all conditions associated with these compounds from source-stage training, this evaluation tests transfer through the molecular intervention representation rather than reuse of matched source perturbation populations.

We selected source objectives and checkpoints exclusively using the random source-validation split; we reserved the drug-held-out split for post-selection generalization assessment and did not use it for model or hyperparameter selection. This procedure selected the MMD + cosine + norm objective.

When evaluated subsequently on held-out drugs, the selected objective also outperformed MMD alone, whereas adding a vector-displacement MSE term reduced transition cosine and worsened both MMD and Δ MSE (Appendix Table 20).



C. Random held-out source conditions

Method	MMD ↓	Trans. Cosine ↑	Trans. MSE ↓
Identity	0.108	–	0.021
Global shift	0.105	0.219	0.020
Linear	0.108	0.418	0.021
MLP	0.094†	0.460†	0.016†
cAE	0.071	0.670	0.010
PerturbRx	**0.067**	**0.688**	**0.009**

Population-level metrics are averaged over held-out conditions. Bold, underline, and † indicate the top three results. cAE: conditional autoencoder.

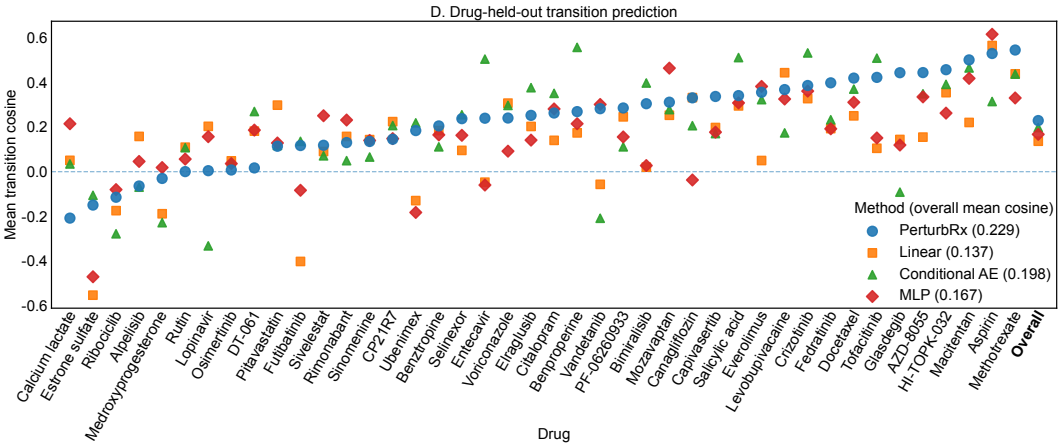


Figure 2: **PerturbRx predicts treatment-induced latent transitions for held-out source conditions and compounds. A–B,** PCA visualization of a held-out Bortezomib treatment condition (CVCL_C466, 0.05 μM) before and after prediction. **C,** Performance on randomly held-out conditions, where the condition is excluded from training but its drug may remain represented. **D,** Drug-wise transition cosine for 42 held-out drugs, with all drug conditions excluded from source-stage training. Relative to the same-objective cAE control, PerturbRx improved mean transition cosine by 0.032 (95% paired-bootstrap CI: -0.020 to 0.086) and achieved higher cosine for 23/42 drugs. Learned controls in **C–D** use the same MMD + cosine + norm objective as PerturbRx.

4.3 Patient Response Prediction with Transferred Treatment Transitions

We use TCGA-485 as the primary clinical benchmark, comprising 485 treatment episodes from 441 patients across five drugs with binary responder (R) and non-responder (NR) labels. Patient-grouped outer folds yield out-of-fold (OOF) predictions without patient overlap. We train a shared response head across treatments, with treatment represented by its ChemBERTa drug embedding. For the evaluated baselines, all five TCGA drugs were represented in their respective preclinical source datasets. In contrast, PerturbRx did not use cisplatin in Tahoe-100M, and it explicitly excluded docetaxel from its source-stage training.

PerturbRx achieved a pooled OOF AUROC of 0.709 and AUPRC of 0.794 (Table 1). Paired cluster-level comparisons were significant against WISER on both metrics and DeepSADR on AUPRC after Holm correction; the remaining comparisons were not significant.

Drug-stratified performance remained heterogeneous across treatments. Notably, docetaxel was held out from source-stage training, and cisplatin was absent from the source dataset. We therefore use representation ablations and transition-shuffling controls to assess whether transferred transitions provide signal beyond patient and treatment information.

To separate the contributions of baseline state, treatment identity, and transition information, we use a hierarchy of representation controls. Patient-only and Patient+Drug controls evaluate static

Table 1: Patient-response prediction on TCGA-485 under patient-grouped evaluation.

	Overall ($n = 485$)		Cisplatin ($n = 150$)		Docetaxel ($n = 51$)		Fluorouracil‡ ($n = 90$)		Gemcitabine‡ ($n = 112$)		Paclitaxel‡ ($n = 82$)	
Method	AUROC	AUPRC	AUROC	AUPRC	AUROC	AUPRC	AUROC	AUPRC	AUROC	AUPRC	AUROC	AUPRC
CODE-AE	0.648†	0.763†	0.584†	0.831†	0.421	0.579†	0.543	**0.691**	0.517†	0.415	0.531†	0.788†
Std. dev.	±0.026	±0.043	±0.193	±0.094	±0.113	±0.095	±0.093	±0.060	±0.109	±0.136	±0.088	±0.051
DeepSADR	0.678	0.740*	0.463	0.780	0.458	0.514	0.470	0.606	**0.531**	**0.435**	0.470	0.719
Std. dev.	±0.016	±0.024	±0.100	±0.040	±0.202	±0.193	±0.109	±0.078	±0.129	±0.120	±0.214	±0.086
TransDRP	0.630	0.768	**0.671**	**0.891**	0.592	0.657	**0.549**	0.652	0.392	0.335	0.529	0.792
Std. dev.	±0.028	±0.009	±0.134	±0.049	±0.230	±0.128	±0.164	±0.149	±0.056	±0.036	±0.198	±0.102
WISER	0.567*	0.677*	0.491	0.769	0.477†	0.525	0.464	0.599	0.521	0.396†	0.536	0.749
Std. dev.	±0.020	±0.041	±0.217	±0.102	±0.160	±0.082	±0.135	±0.082	±0.140	±0.138	±0.214	±0.125
PerturbRx	**0.709**	**0.794**	0.658	0.863	**0.644**	**0.699**	0.496†	0.611†	0.490	0.384	**0.617**	**0.814**
Std. dev.	±0.027	±0.030	±0.121	±0.034	±0.215	±0.108	±0.106	±0.089	±0.217	±0.148	±0.137	±0.094

Overall reports pooled exactly-once OOF performance across all 485 treatment episodes; drug-specific columns report within-drug results. The response predictor is shared across treatments and receives the drug representation as an input. Second-line values show standard deviation across five outer folds. ‡ denotes drugs represented in source-stage training. * denotes Holm-adjusted $p < 0.05$ versus PerturbRx. Bold, underline, and † denote the top three results.

patient and treatment representations without predicted transitions. The random-transition control tests whether adding a transition-like feature is sufficient, whereas the shuffled controls test whether correspondence between the predicted transition and the patient–treatment pair matters. Global shuffling disrupts both treatment and patient correspondence, while within-drug shuffling preserves treatment identity but breaks patient-specific transition correspondence.

Table 2: Ablation results on TCGA-485.

Representation		**AUROC↑**	**AUPRC↑**
Patient		0.614*	0.729*
	Std. dev.	±0.036	±0.018
Patient+Drug		0.705	0.777
	Std. dev.	±0.028	±0.045
Random transition		0.703†	0.770
	Std. dev.	±0.022	±0.036
Shuffled transition		0.665	0.753
	Std. dev.	±0.031	±0.017
Within-drug shuffle		0.690	0.771†
	Std. dev.	±0.033	±0.044
PerturbRx		**0.709**	**0.794**
	Std. dev.	±0.027	±0.030

Pooled OOF results are reported over five patient-grouped folds; Std. dev. denotes fold variation. * denotes Holm-adjusted $p < 0.05$ versus PerturbRx. Bold, underline, and † denote the top three results.

Table 2 compares static and treatment-conditioned representation controls on TCGA-485. PerturbRx achieved the highest AUROC and AUPRC (0.709/0.794) and significantly outperformed the patient-only representation after Holm correction. Notably, Patient+Drug performed comparably to or better than the evaluated baselines, indicating that the underlying patient and drug representations provide a strong predictive baseline. PerturbRx provided an additional gain over Patient+Drug and transition controls, although these differences were not significant after multiplicity correction. On pooled metrics, PerturbRx exceeded Patient+Drug by 0.004 AUROC (95% CI $[-0.022, 0.031]$) and 0.017 AUPRC (95% CI $[-0.024, 0.053]$). Thus, the ablations are consistent with matched transition information providing incremental predictive signal beyond static and shuffled controls, although the pairwise differences were not significant after multiplicity correction.

4.4 Cross-Domain Transfer to PDX Models

We next test whether the predictive and control patterns observed in TCGA-485 transfer to an independent target domain. The PDX benchmark provides a complementary cross-scale setting that differs from TCGA in both biological system and upstream expression quantification, while using the same frozen scFoundation representation pipeline.

Table 3: Results on the PDX benchmark.

Method	**AUROC↑**	**AUPRC↑**
Existing methods		
CODE-AE	0.467* ±0.010	0.407* ±0.013
DeepSADR	0.626±0.024	0.545±0.018
TransDRP	0.516* ±0.030	0.449* ±0.051
WISER	0.484* ±0.031	0.431* ±0.036
PerturbRx	**0.685**±0.050	**0.612**±0.043
Representation ablations		
PDX	0.615* ±0.062	0.528* ±0.052
PDX+Drug	0.673† ±0.044	0.589† ±0.062
Random transition	0.673† ±0.063	0.569±0.071
Shuffled transition	0.677±0.028	0.601±0.018
Within-drug shuffle	0.666±0.029	0.571±0.029

Pooled OOF results over five model-disjoint folds; ± denotes fold SD. * denotes Holm-adjusted $p < 0.05$ versus PerturbRx; bold, underline, and † denote first, second, and third.

PerturbRx achieved the highest pooled OOF performance across both the evaluated baseline methods and representation controls, with an AUROC of 0.685 and an AUPRC of 0.612 (Table 3). Among the existing methods, DeepSADR achieved the highest pooled performance (0.626/0.545), while the shuffled-transition condition provided the strongest representation control (0.677/0.601). PDX+Drug also provided a strong static control (0.673/0.589) without using the transferred transition feature. After Holm correction, PerturbRx significantly outperformed CODE-AE, TransDRP, and WISER on both AUROC and AUPRC, as well as the PDX-only representation. Differences from DeepSADR and the remaining representation controls were not significant. Overall, these results support the predictive value of the

transferred representation in PDX, with significant gains over the PDX-only representation and three of four evaluated external baselines.

Table 4: Drug-stratified AUROC on the shared 10-drug PDX benchmark.

Method	Alpelisib‡	Buparlisib	Erlotinib‡	Fluorouracil‡	Gemcitabine‡	Paclitaxel‡	Ribociclib‡	Ruxolitinib	Tamoxifen	Trametinib‡
CODE-AE	0.464*	0.450	0.633	0.610	**0.620**	0.372	0.502	0.396	0.433	0.587†
Std. dev.	±0.083	±0.044	±0.274	±0.264	±0.202	±0.156	±0.182	±0.203	±0.400	±0.325
DeepSADR	0.587*	**0.595**	0.300	0.304	0.462†	0.511	0.497	0.762	0.719	0.430
Std. dev.	±0.079	±0.074	±0.217	±0.196	±0.349	±0.065	±0.082	±0.158	±0.206	±0.258
TransDRP	0.471*	0.462	0.517†	0.588†	0.373	0.542†	0.567†	0.450	0.181*	0.610
Std. dev.	±0.100	±0.114	±0.356	±0.163	±0.274	±0.239	±0.071	±0.237	±0.155	±0.114
WISER	0.512†*	0.486†	0.400	0.579	0.492	0.547	0.576	0.596†	0.624†	0.587†
Std. dev.	±0.049	±0.141	±0.418	±0.182	±0.232	±0.176	±0.108	±0.229	±0.340	±0.369
PerturbRx	**0.727**	0.569	**0.700**	**0.713**	0.462†	**0.623**	**0.588**	**0.803**	**0.752**	**0.830**
Std. dev.	±0.105	±0.097	±0.183	±0.205	±0.150	±0.211	±0.082	±0.109	±0.273	±0.172

Drug-specific values report mean AUROC across five model-disjoint outer folds; second-line values show the corresponding standard deviation. ‡ denotes drugs represented during PerturbRx source-stage training. Bold, underline, and (†) denote the best, second-best, and third-best mean results, respectively. * denotes Holm-adjusted $p < 0.05$ versus PerturbRx, with correction across the four external baselines separately within each drug and metric.

Drug-stratified evaluation showed substantial heterogeneity across treatments. PerturbRx achieved the highest AUROC for 8 of 10 drugs (Table 4) and the highest AUPRC for 7 of 10 drugs (Table 5). For erlotinib, PerturbRx achieved the highest AUROC and second-highest AUPRC, whereas DeepSADR ranked first for buparlisib and CODE-AE ranked first for gemcitabine. Drug-specific paired inference identified significant differences after Holm correction against all four external baselines for both metrics in alpelisib, against TransDRP for both metrics in tamoxifen, and against DeepSADR for AUPRC in trametinib.

Table 5: Drug-stratified AUPRC on the shared 10-drug PDX benchmark.

Method	Alpelisib‡	Buparlisib	Erlotinib‡	Fluorouracil‡	Gemcitabine‡	Paclitaxel‡	Ribociclib‡	Ruxolitinib	Tamoxifen	Trametinib‡
CODE-AE	0.489*	0.489	**0.690**	0.723	**0.814**	0.434	0.488	0.200	0.357	0.783
Std. dev.	±0.053	±0.046	±0.239	±0.181	±0.141	±0.115	±0.144	±0.110	±0.368	±0.155
DeepSADR	0.589*	**0.640**	0.456	0.467	0.721†	0.550	0.462	0.560	0.467	0.759*
Std. dev.	±0.072	±0.091	±0.108	±0.072	±0.208	±0.065	±0.106	±0.219	±0.315	±0.094
TransDRP	0.516*	0.504	0.624†	0.668	0.697	0.555†	0.526	0.207	0.160*	0.816†
Std. dev.	±0.100	±0.118	±0.269	±0.145	±0.198	±0.174	±0.141	±0.058	±0.028	±0.064
WISER	0.536†*	0.557†	0.601	0.676†	0.766	0.569	0.510†	0.450†	0.435†	0.832
Std. dev.	±0.067	±0.149	±0.260	±0.168	±0.117	±0.173	±0.120	±0.264	±0.344	±0.162
PerturbRx	**0.707**	0.602	0.678	**0.772**	0.715	**0.628**	**0.558**	**0.578**	**0.590**	**0.942**
Std. dev.	±0.106	±0.084	±0.182	±0.175	±0.085	±0.211	±0.070	±0.178	±0.391	±0.061

Drug-specific values report mean AUPRC across five model-disjoint outer folds; second-line values show the corresponding standard deviation. ‡ denotes drugs represented during PerturbRx source-stage training. Bold, underline, and (†) denote the best, second-best, and third-best mean results, respectively. * denotes Holm-adjusted $p < 0.05$ versus PerturbRx, with correction across the four external baselines separately within each drug and metric.

Notably, all 10 PDX drugs were represented in the source-domain response data used by the baseline methods, whereas buparlisib, ruxolitinib, and tamoxifen had no exact source-stage support for PerturbRx. Despite this, PerturbRx ranked first on both metrics for ruxolitinib and tamoxifen. Relative to the strongest baseline for each drug and metric, PerturbRx improved AUROC and AUPRC by 5.4% and 3.2%, respectively, for ruxolitinib (0.803 vs. 0.762; 0.578 vs. 0.560), and by 4.6% and 26.3% for tamoxifen (0.752 vs. 0.719; 0.590 vs. 0.467).

Table 6: Source-architecture controls.

	TCGA-485		PDX	
	AUROC	AUPRC	AUROC	AUPRC
Linear	0.665	0.755	0.670	0.586
Std. dev.	±0.025	±0.022	±0.017	±0.034
MLP	0.689	0.775	0.677	0.603
Std. dev.	±0.015	±0.020	±0.038	±0.038
cAE	0.699	0.792	0.676	0.571
Std. dev.	±0.020	±0.022	±0.034	±0.033
PerturbRx	**0.709**	**0.794**	**0.685**	**0.612**
Std. dev.	±0.027	±0.030	±0.050	±0.043

Pooled exactly-once OOF performance; ± denotes standard deviation across five outer folds. Only the source transition architecture is changed; the downstream response pipeline is fixed.

Architectural ablation To determine whether the gains arise from the learned source transition operator itself, we replaced it with strictly matched Linear, MLP, and conditional autoencoder (cAE) alternatives while keeping the downstream patient-response pipeline fixed. PerturbRx achieved the strongest pooled OOF AUROC and AUPRC on both TCGA-485 and PDX (Table 6), indicating that the learned transition operator contributes beyond the downstream prediction architecture alone.

Table 7: Response-feature decomposition.

	TCGA-485		PDX	
Representation	AUROC	AUPRC	AUROC	AUPRC
Predicted state only	0.630	0.735	0.658	0.561
$[\hat{z}_{p,u}]$	±0.047	±0.043	±0.040	±0.039
Predicted state+Drug	0.660	0.753	0.674	0.597
$[m_d; \hat{z}_{p,u}]$	±0.034	±0.027	±0.042	±0.026
PerturbRx	**0.709**	**0.794**	**0.685**	**0.612**
$[z_p; m_d; \hat{z}_{p,u}]$	±0.027	±0.030	±0.050	±0.043

Pooled exactly-once OOF performance; ± denotes standard deviation across five outer folds. $\hat{z}_{p,u}$ is the predicted treatment-conditioned state, m_d the drug representation, and z_p the baseline patient state.

Response feature decomposition We next examined how the response model uses the transferred signal (Table 7). The predicted treatment-conditioned state $\hat{z}_{p,u}$ remained predictive on its own; adding the drug representation m_d improved performance further, and the full input $[z_p; m_d; \hat{z}_{p,u}]$ performed best on both benchmarks. These results indicate that performance benefits from both the learned transition operator and combining its predicted treatment-conditioned state with explicit drug and baseline patient information.

5 Conclusion and Discussion

We introduced PerturbRx, a treatment-conditioned representation learning framework that learns drug- and dose-conditioned latent transitions from single-cell perturbation populations and transfers the learned operator to drug-response prediction. Our experiments address the three questions posed in the Introduction. First, the learned transition operator generalized to compounds excluded from source-stage training. Second, transferred transition representations provided incremental predictive signal beyond strong patient+drug representations, although the corresponding gains were not uniformly significant. Third, we observed the same predictive and ablation patterns across TCGA-485 and an independent PDX cohort, supporting transfer across distinct target domains.

The central contribution is therefore not a claim of uniform response improvement, but a transition operator that generalizes beyond source-stage compounds and transfers across biological scales without post-treatment measurements in the target domain.

Overall, the results support the hypothesis that intervention-conditioned transformations learned from large perturbation datasets can provide transferable structure across distinct target domains. The specific source of this signal remains unresolved: it may reflect patient-specific perturbation information, a more general nonlinear treatment-conditioned transformation of patient and drug representations, or both. Larger matched pre- and post-treatment cohorts will be needed to distinguish these mechanisms.

Several limitations remain. Clinical sample sizes are small, particularly for treatment-specific analyses, which limits statistical power and may partly explain why several pairwise improvements did not reach significance after multiplicity correction; treatment-specific estimates therefore remain uncertain. Source and target domains also differ substantially in biological system and measurement scale; inference dose is a representation setting rather than a patient-specific exposure estimate; and the learned transitions are not explicitly constrained by known pathways or mechanisms of action. More broadly, PerturbRx separates learning an intervention-conditioned state transformation from learning its association with an outcome. This principle may extend beyond cancer drug response to settings where rich perturbation data are available in a source domain but matched post-intervention observations are scarce in the target population. Extending the framework to multi-step transitions could further connect it to latent world models and counterfactual treatment-sequence modeling.

ACKNOWLEDGMENTS

Omitted for anonymous review; will be added at camera-ready.

AI USE STATEMENT

Generative AI tools assisted with research ideation, literature retrieval, code development and debugging, interpretation of results, drafting and language editing, and document formatting. The authors reviewed and verified all scientific claims, experimental design decisions, implementations, analyses, and reported results. The authors take full responsibility for the paper's content.

ETHICS STATEMENT

This study uses publicly available, de-identified molecular and treatment-response datasets, including TCGA and PDX. We did not recruit new human subjects or access identifiable patient information. We conducted the analyses for methodological research and do not intend them to support direct clinical decision-making. Predicted latent transitions and treatment-response scores should not be interpreted as validated clinical recommendations.

REPRODUCIBILITY STATEMENT

We document dataset access, cohort construction, preprocessing, model architectures, training objectives, hyperparameter search procedures, randomization procedures, and evaluation protocols in the appendices. The source and target datasets are publicly available, and the corresponding dataset sections provide access locations and preprocessing details. To support reproducibility during review, the code and supporting materials are available in an anonymized repository at `https://anonymous.4open.science/r/PerturbRx/README.md`. Upon publication, we will release the full code and processed metadata through a public GitHub repository and an archived Zenodo release, with the corresponding links included in the final version.

REFERENCES

Dongmin Bang, Bonil Koo, and Sun Kim. Transfer learning of condition-specific perturbation in gene interactions improves drug response prediction. *Bioinformatics*, 40(Supplement 1):i130–i139, 07 2024. ISSN 1367-4811. doi: 10.1093/bioinformatics/btae249. URL https://doi.org/10.1093/bioinformatics/btae249.

Dongmin Bang, Sugyun An, Inyoung Sung, Ilho Yun, Sun Kim, and Sangseon Lee. Predicting therapeutic outcome via aligning patient-specific knowledge graph and gene-level perturbation representations, 2026. URL https://arxiv.org/abs/2607.04557.

Uri Ben-David, Gavin Ha, Yuen-Yi Tseng, Noah F Greenwald, Coyin Oh, Juliann Shih, James M McFarland, Bang Wong, Jesse S Boehm, Rameen Beroukhim, and Todd R Golub. Patient-derived xenografts undergo mouse-specific tumor evolution. *Nature Genetics*, 49(11):1567–1575, 2017. doi: 10.1038/ng.3967. URL https://doi.org/10.1038/ng.3967.

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. doi: 10.1038/s41592-023-01969-x. URL https://doi.org/10.1038/s41592-023-01969-x.

Seyone Chithrananda, Gabriel Grand, and Bharath Ramsundar. Chemberta: Large-scale self-supervised pretraining for molecular property prediction, 2020. URL https://arxiv.org/abs/2010.09885.

Alice Driessen, Dhruva Abhijit Rajwade, Benedek Harsanyi, Marianna Rapsomaniki, and Jannis Born. Conditional monge gap enables generalizable single-cell perturbation modelling. *Nature Machine Intelligence*, 8(6):984–996, 2026. doi: 10.1038/s42256-026-01242-8. URL https://doi.org/10.1038/s42256-026-01242-8.

Hui Gao, Joshua M Korn, Stéphane Ferretti, John E Monahan, Youzhen Wang, Mallika Singh, Chao Zhang, Christian Schnell, Guizhi Yang, Yun Zhang, O Alejandro Balbin, Stéphanie Barbe, Hongbo Cai, Fergal Casey, Susmita Chatterjee, Derek Y Chiang, Shannon Chuai, Shawn M Cogan, Scott D Collins, Ernesta Dammassa, Nicolas Ebel, Millicent Embry, John Green, Audrey Kauffmann, Colleen Kowal, Rebecca J Leary, Joseph Lehar, Ying Liang, Alice Loo, Edward Lorenzana, E. Robert McDonald, Margaret E McLaughlin, Jason Merkin, Ronald Meyer, Tara L Naylor, Montesa Patawaran, Anupama Reddy, Claudia Röelli, David A Ruddy, Fernando Salangsang, Francesca Santacroce, Angad P Singh, Yan Tang, Walter Tinetto, Sonja Tobler, Roberto Velazquez, Kavitha Venkatesan, Fabian Von Arx, Hui Qin Wang, Zongyao Wang, Marion Wiesmann, Daniel Wyss, Fiona Xu, Hans Bitter, Peter Atadja, Emma Lees, Francesco Hofmann, En Li, Nicholas Keen, Robert Cozens, Michael Rugaard Jensen, Nancy K Pryer, Juliet A Williams, and William R Sellers. High-throughput screening using patient-derived tumor xenografts to predict clinical trial drug response. *Nature Medicine*, 21(11):1318–1325, 2015. doi: 10.1038/nm.3954. URL https://doi.org/10.1038/nm.3954.

Jean-Pierre Gillet, Sudhir Varma, and Michael M. Gottesman. The clinical relevance of cancer cell lines. *JNCI: Journal of the National Cancer Institute*, 105(7):452–458, 04 2013. ISSN 0027-8874. doi: 10.1093/jnci/djt007. URL https://doi.org/10.1093/jnci/djt007.

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(8):1481–1491, 2024. doi: 10.1038/s41592-024-02305-7. URL https://doi.org/10.1038/s41592-024-02305-7.

Di He, Qiao Liu, You Wu, and Lei Xie. A context-aware deconfounding autoencoder for robust prediction of personalized clinical drug response from cell-line compound screening. *Nature Machine Intelligence*, 4(10):879–892, 2022. doi: 10.1038/s42256-022-00541-0. URL https://doi.org/10.1038/s42256-022-00541-0.

Allison P. Heath, Vincent Ferretti, Stuti Agrawal, Maksim An, James C. Angelakos, Renuka Arya, Rosita Bajari, Bilal Baqar, Justin H. B. Barnowski, Jeffrey Burt, Ann Catton, Brandon F. Chan, Fay Chu, Kim Cullion, Tanja Davidsen, Phuong-My Do, Christian Dompierre, Martin L. Ferguson, Michael S. Fitzsimons, Michael Ford, Miyuki Fukuma, Sharon Gaheen, Gajanan L. Ganji, Tzintzuni I. Garcia, Sameera S. George, Daniela S. Gerhard, Francois Gerthoffert, Fauzi Gomez, Kang Han, Kyle M. Hernandez, Biju Issac, Richard Jackson, Mark A. Jensen, Sid Joshi, Ajinkya Kadam, Aishmit Khurana, Kyle M. J. Kim, Victoria E. Kraft, Shenglai Li, Tara M. Lichtenberg, Janice Lodato, Laxmi Lolla, Plamen Martinov, Jeffrey A. Mazzone, Daniel P. Miller, Ian Miller, Joshua S. Miller, Koji Miyauchi, Mark W. Murphy, Thomas Nullet, Rowland O. Ogwara, Francisco M. Ortuño, Jesús Pedrosa, Phuong L. Pham, Maxim Y. Popov, James J. Porter, Raymond Powell, Karl Rademacher, Colin P. Reid, Samantha Rich, Bessie Rogel, Himanso Sahni, Jeremiah H. Savage, Kyle A. Schmitt, Trevar J. Simmons, Joseph Sislow, Jonathan Spring, Lincoln Stein, Sean Sullivan, Yajing Tang, Mathangi Thiagarajan, Heather D. Troyer, Chang Wang, Zhining Wang, Bedford L. West, Alex Wilmer, Shane Wilson, Kaman Wu, William P. Wysocki, Linda Xiang, Joseph T. Yamada, Liming Yang, Christine Yu, Christina K. Yung, Jean Claude Zenklusen, Junjun Zhang, Zhenyu Zhang, Yuanheng Zhao, Ariz Zubair, Louis M. Staudt, and Robert L. Grossman. The nci genomic data commons. *Nature Genetics*, 53(3):257–262, 2021. doi: 10.1038/s41588-021-00791-5. URL https://doi.org/10.1038/s41588-021-00791-5.

Aishwarya Jayagopal, Yanrong Zhang, Robert John Walsh, Tuan Zea Tan, Anand D Jeyasekharan, and Vaibhav Rajan. GANDALF: Generative attention based data augmentation and predictive modeling framework for personalized cancer treatment. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=WwmtcGr4lP.

Craig Knox, Mike Wilson, Christen M Klinger, Mark Franklin, Eponine Oler, Alex Wilson, Allison Pon, Jordan Cox, Na Eun (Lucy) Chin, Seth A Strawbridge, Marysol Garcia-Patino, Ray Kruger, Aadhavya Sivakumaran, Selena Sanford, Rahil Doshi, Nitya Khetarpal, Omolola Fatokun, Daphnee Doucet, Ashley Zubkowski, Dorsa Yahya Rayat, Hayley Jackson, Karxena Harford, Afia Anjum, Mahi Zakir, Fei Wang, Siyang Tian, Brian Lee, Jaanus Liigand, Harrison Peters, Ruo Qi (Rachel) Wang, Tue Nguyen, Denise So, Matthew Sharp, Rodolfo da Silva, Cyrella Gabriel, Joshua Scantlebury, Marissa Jasinski, David Ackerman, Timothy Jewison, Tanvir Sajed, Vasuk Gautam, and David S Wishart. Drugbank 6.0: the drugbank knowledgebase for 2024. *Nucleic Acids Research*, 52(D1):D1265–D1275, 01 2024. ISSN 0305-1048. doi: 10.1093/nar/gkad976. URL https://doi.org/10.1093/nar/gkad976.

Yann LeCun. A path towards autonomous machine intelligence, June 2022. URL https://openreview.net/forum?id=BZ5a1r-kVsf. Position paper, Version 0.9.2.

Zhaokang Liang, Shuyang Zhuang, Xiaoran Jiao, Weian Mao, Hao Chen, and Chunhua Shen. scppdm: A diffusion model for single-cell drug-response prediction, 2025. URL https://arxiv.org/abs/2510.11726.

Xuan Liu and Menglu Li. Knowledge-guided domain adaptation model for transferring drug response prediction from cell lines to patients. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 523–531, 2025. doi: 10.1609/aaai.v39i1.32032.

Mohammad Lotfollahi, F. Alexander Wolf, and Fabian J. Theis. scgen predicts single-cell perturbation responses. *Nature Methods*, 16(8):715–721, 2019. doi: 10.1038/s41592-019-0494-8. URL https://doi.org/10.1038/s41592-019-0494-8.

Mohammad Lotfollahi, Anna Klimovskaia Susmelj, Carlo De Donno, Leon Hetzel, Yuge Ji, Ignacio L. Ibarra, Sanjay R. Srivatsan, Mohsen Naghipourfar, Riza M. Daza, Beth Martin, Jay Shendure, Jose L. McFaline-Figueroa, Pierre Boyeau, F. Alexander Wolf, Nafissa Yakubova, Stephan Günnemann, Cole Trapnell, David Lopez-Paz, and Fabian J. Theis. Predicting cellular responses to complex perturbations in high-throughput screens. *Molecular Systems Biology*, 19(6):MSB202211517, 2023. doi: 10.15252/msb.202211517. URL https://doi.org/10.15252/msb.202211517.

Linh C. Nguyen, Stefan Naulaerts, Alejandra Bruna, Ghita Ghislat, and Pedro J. Ballester. Predicting cancer drug response in vivo by learning an optimal feature selection of tumour molecular profiles. *Biomedicines*, 9(10):1319, 2021. doi: 10.3390/biomedicines9101319.

Alexander Partin, Thomas S. Brettin, Yitan Zhu, Oleksandr Narykov, Austin Clyde, Jamie Overbeek, and Rick L. Stevens. Deep learning methods for drug response prediction in cancer: Predominant and emerging trends. *Frontiers in Medicine*, 10:1086097, 2023. doi: 10.3389/fmed.2023.1086097. URL https://www.frontiersin.org/journals/medicine/articles/10.3389/fmed.2023.1086097.

Ladislav Rampášek, Daniel Hidru, Petr Smirnov, Benjamin Haibe-Kains, and Anna Goldenberg. Dr.vae: improving drug response prediction via modeling of drug perturbation effects. *Bioinformatics*, 35(19):3743–3751, 10 2019. ISSN 1367-4803. doi: 10.1093/bioinformatics/btz158. URL https://doi.org/10.1093/bioinformatics/btz158.

Kumar Shubham, Aishwarya Jayagopal, Syed Mohammed Danish, Prathosh Ap, and Vaibhav Rajan. WISER: Weak supervision and supervised representation learning to improve drug response prediction in cancer. In Ruslan Salakhutdinov, Zico Kolter, Katherine Heller, Adrian Weller, Nuria Oliver, Jonathan Scarlett, and Felix Berkenkamp (eds.), *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 45228–45243. PMLR, 2024. URL https://proceedings.mlr.press/v235/shubham24a.html.

Sanjay R. Srivatsan, José L. McFaline-Figueroa, Vijay Ramani, Lauren Saunders, Junyue Cao, Jonathan Packer, Hannah A. Pliner, Dana L. Jackson, Riza M. Daza, Lena Christiansen, Fan Zhang, Frank Steemers, Jay Shendure, and Cole Trapnell. Massively multiplex chemical transcriptomics at single-cell resolution. *Science*, 367(6473):45–51, 2020. doi: 10.1126/science.aax6234. URL https://www.science.org/doi/abs/10.1126/science.aax6234.

Inyoung Sung, Dongmin Bang, Sun Kim, and Sangseon Lee. Deep learning for predicting patient drug response by transferring gene-level and cell-level knowledge to tumors. *npj Precision Oncology*, 10(1):310, 2026. doi: 10.1038/s41698-026-01491-1. URL https://doi.org/10.1038/s41698-026-01491-1.

The Cancer Genome Atlas Research Network, John N. Weinstein, Eric A. Collisson, Gordon B. Mills, Kenna R. Mills Shaw, Brad A. Ozenberger, Kyle Ellrott, Ilya Shmulevich, Chris Sander, and Joshua M. Stuart. The cancer genome atlas pan-cancer analysis project. *Nature Genetics*, 45(10):1113–1120, 2013. doi: 10.1038/ng.2764. URL https://doi.org/10.1038/ng.2764.

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, 01 2013. ISSN 0305-1048. doi: 10.1093/nar/gks1111. URL https://doi.org/10.1093/nar/gks1111.

Xinyu Yuan, Xixian Liu, Ya Shi Zhang, Zuobai Zhang, Hongyu Guo, and Jian Tang. Perturbdiff: Functional diffusion for single-cell perturbation modeling. In *Forty-third International Conference on Machine Learning*, 2026. URL https://openreview.net/forum?id=yBAqonxCdp.

Jesse Zhang, Airol A Ubas, Richard de Borja, Valentine Svensson, Nicole Thomas, Neha Thakar, Ian Lai, Aidan Winters, Umair Khan, Matthew G. Jones, John D. Thompson, Vuong Tran, Joseph Pangallo, Efthymia Papalexi, Ajay Sapre, Hoai Nguyen, Oliver Sanderson, Maria Nigos, Olivia Kaplan, Sarah Schroeder, Bryan Hariadi, Simone Marrujo, Crina Curca Alec Salvino, Guillermo Gallareta Olivares, Ryan Koehler, Gary Geiss, Alexander Rosenberg, Charles Roco, Daniele Merico, Nima Alidoust, Hani Goodarzi, and Johnny Yu. Tahoe-100m: A giga-scale single-cell perturbation atlas for context-dependent gene function and cellular modeling. *bioRxiv*, 2025. doi: 10.1101/2025.02.20.639398. URL https://www.biorxiv.org/content/early/2025/05/10/2025.02.20.639398.

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 *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=jrFJWpDZvq.

A Data and Cohort Construction

A.1 Tahoe-100M Source Dataset

We use Tahoe-100M as the source-stage single-cell perturbation dataset. We construct context-matched control and treated populations from DMSO and drug-treated cells within the same cell line and plate, with each condition indexed by drug and dose. Table 8 summarizes the resulting dataset.

Table 8: Statistics of the Tahoe-100M source-stage perturbation dataset. Observed treated conditions are observed cell-line, plate, drug, and dose combinations before the final source-modeling filters. Final modeling conditions additionally require context-matched DMSO controls with sufficient cell counts and an available molecular representation.

Statistic	Value
Post-QC cells	100.6M
Plates	14
Cell lines	50
Dose levels	0.05, 0.5, 5.0 μM
Observed treated conditions	65,093
Median cells per observed treated condition	1,278
Matched DMSO groups	700
Median DMSO cells per group	3,337
Final modeling drugs	377
Final modeling conditions	61,175
Source-training conditions	49,425
Random held-out conditions	5,492
Drug-held-out conditions	6,258
Drug-held-out compounds	42

A.2 TCGA-186 Patient Benchmark

We use a TCGA subset of the processed WISER-style patient drug response benchmark (Shubham et al., 2024) for target-stage evaluation. We retain patient-drug records with a patient identifier that is a TCGA barcode and a drug that is one of five anticancer agents: fluorouracil, temozolomide, gemcitabine, cisplatin, or sorafenib. Following the processed WISER-style benchmark, we define patient response labels from cancer relapse time after chemotherapy: records with relapse times greater than the benchmark-defined median are labeled as responders, whereas records with relapse times less than the median are labeled as non-responders. Thus, the target labels are binary relapse-time proxy labels derived from processed treatment-response and relapse-time annotations, rather than prospective RECIST clinical response endpoints.

We obtained TCGA bulk RNA-seq expression data from the National Cancer Institute Genomic Data Commons (GDC) (Heath et al., 2021). Specifically, we used the gene-level read-count files (`data_mrna_seq_read_counts.txt`) distributed in the TCGA treatment-response data repository for the TCGA-GDC study collections. We processed the expression files per study and combined them at the sample-by-gene level. TCGA sample barcodes were reduced to patient-level barcodes using the first 12 characters of each barcode (e.g., `TCGA-XX-YYYY-01A` $\rightarrow$ `TCGA-XX-YYYY`). For patients with multiple RNA-seq samples, we selected a single representative sample using a predefined sample-type priority that favors primary tumor samples over recurrent, metastatic, blood-derived normal, and solid-tissue-normal samples. The priority order was primary solid tumor (01), primary blood-derived cancer-peripheral blood (03), primary blood-derived cancer-bone marrow (09), recurrent solid tumor (02), metastatic (06), blood-derived normal (10), and solid tissue normal (11). This selection procedure follows the expression-matrix construction used for the GDC TCGA treatment-response cohorts.

After sample selection, we used the resulting gene-level count matrix as input to the target-stage expression preprocessing pipeline. We library-size-normalized counts, then $\log(1+x)$-transformed

them and encoded them using the frozen scFoundation model, yielding a 3072-dimensional patient representation. We used the same GDC-derived read-count source, patient-level sample-selection procedure, normalization, and scFoundation embedding pipeline for the TCGA-485 cohort, ensuring both TCGA benchmarks were represented in a common target-domain embedding space.

The TCGA treatment-response repository constructs its treatment cohorts from cBioPortal-style TCGA study directories containing treatment timelines and associated molecular data. The underlying treatment-response pipeline standardizes treatment outcomes to CR, PR, SD, and PD, normalizes therapeutic-agent names, consolidates treatment records into episodes, and identifies single-agent treatment episodes for downstream modeling.

After matching the response records to the GDC-derived expression profiles, the TCGA-186 benchmark contains 186 patient-drug records from 186 unique patients across the five eligible drugs. The benchmark includes 91 responders and 95 non-responders. Table 9 summarizes the drug-wise response composition of the benchmark, and Table 10 reports the corresponding distribution across TCGA cancer types and treatments.

Table 9: Drug-level composition of TCGA-186. Response labels are relapse-time proxies derived from the processed WISER-style benchmark.

Drug	Samples	Responders	Non-responders
Cisplatin	40	20	20
Fluorouracil	28	12	16
Gemcitabine	46	23	23
Sorafenib	26	13	13
Temozolomide	46	23	23
Total	186	91	95

A.3 TCGA-485 Treatment-Response Cohort

Treatment records and response labels. We constructed the TCGA-485 cohort from treatment timeline records in the TCGA Genomic Data Commons (GDC) (Heath et al., 2021). We retained pharmacologic treatment records with a non-empty therapeutic-agent annotation and a response that could be mapped to complete response (CR), partial response (PR), stable disease (SD), or progressive disease (PD). We mapped the source annotation "No Response" to PD. We then defined an objective-response-rate (ORR)-like binary endpoint as

$$y_{\mathrm{ORR}} = \begin{cases} 1, & \mathrm{CR\ or\ PR}, \\ 0, & \mathrm{SD\ or\ PD}. \end{cases}$$

Episode construction and drug eligibility. We normalized therapeutic agent names, mapping predefined salt forms and naming variants to a common active ingredient name. We consolidated treatment records into episodes using the TCGA dataset, patient identifiers, treatment dates, treatment types, and response annotations. When both treatment dates were missing, we retained the corresponding source records as separate episodes. We restricted the analysis to single-agent episodes.

Candidate drugs were required to have at least 30 single-agent episodes, at least 10 ORR-positive episodes, at least 10 ORR-negative episodes, representation in at least three TCGA cancer types, and no single cancer type accounting for more than 80% of episodes. These criteria yielded five eligible drugs: cisplatin, docetaxel, fluorouracil, gemcitabine, and paclitaxel.

Expression matching and final benchmark. We matched treatment episodes to patient TCGA bulk RNA-seq profiles from the GDC. We used the GDC gene-level read-count profiles as the expression source and applied the same frozen scFoundation patient-embedding pipeline used for TCGA-186. We mapped sample-level TCGA barcodes to patient-level barcodes and selected a single representative patient tumor sample for each patient using the same sample-selection procedure. We then applied library-size normalization and a $\log(1+x)$ transformation before frozen scFoundation encoding, yielding 3072-dimensional patient embeddings. Thus, we processed TCGA-186 and

Table 10: Cancer type and drug composition of TCGA-186. Entries denote the number of unique patients associated with each cancer type and treatment. We assigned cancer types using TCGA project metadata from the Genomic Data Commons.

Cancer type	Cis.	Flu.	Gem.	Sor.	Tem.	Total
Adrenocortical	0	0	0	1	0	1
Bladder	0	0	4	0	0	4
Brain	0	0	0	0	42	42
Cervical	5	0	0	0	0	5
Colorectal	0	8	0	0	0	8
Esophagus/Stomach	1	20	0	0	0	21
Gallbladder	0	0	3	0	0	3
Head and Neck	22	0	0	0	0	22
Kidney	0	0	1	6	0	7
Liver	0	0	2	19	0	21
Lung	2	0	3	0	1	6
Melanoma	0	0	0	0	3	3
Pancreas	0	0	30	0	0	30
Prostate	0	0	1	0	0	1
Sarcoma	1	0	2	0	0	3
Uterine	9	0	0	0	0	9
Total	40	28	46	26	46	186

Table 11: Drug-level composition of the final shared TCGA-485 analytic cohort after input-availability filtering across evaluated methods. ORR-positive denotes complete or partial response, whereas ORR-negative denotes stable or progressive disease. An exact Tahoe match indicates agreement under harmonized drug-name or canonical-SMILES matching.

Drug	Episodes	ORR+	ORR-	Prevalence	Tahoe match
Cisplatin	150	119	31	0.793	No
Docetaxel	51	28	23	0.549	Yes
Fluorouracil	90	56	34	0.622	Yes
Gemcitabine	112	42	70	0.375	Yes
Paclitaxel	82	61	21	0.744	Yes
Total	485	306	179	0.631	–

TCGA-485 in the same target-domain expression space, using the same frozen patient-embedding pipeline.

Specimen timing sensitivity. TCGA eligibility criteria generally favored treatment-naive tumor specimens; however, episode-specific specimen procurement dates were unavailable for the selected RNA-seq samples, precluding uniform verification that each specimen preceded its corresponding treatment episode. GDC annotations identified documented neoadjuvant-treatment exceptions, including one primary-tumor specimen associated with a Paclitaxel episode. Excluding this episode from the fixed OOF evaluation had negligible effects on PerturbRx performance (AUROC, 0.709 to 0.708; AUPRC, 0.794 to 0.794), leaving the main conclusions unchanged.

Common-cohort filtering. To maintain a common evaluation cohort across PerturbRx and all baseline methods, we retained only treatment episodes with the required patient tumor expression profile, drug representation, response label, and shared preprocessing outputs. The resulting benchmark comprised 485 treatment episodes from 441 unique patients across five drugs and 24 cancer types. We retained multiple treatment episodes from the same patient; accordingly, we grouped all evaluation splits by patient identifier to prevent patient overlap between training and test sets.

The final TCGA-485 analytic cohort comprised 485 treatment episodes from 441 unique patients, spanning 24 cancer types and five drugs. Of these episodes, 306 were ORR-positive, and 179 were ORR-negative, corresponding to an overall positive-class prevalence of 0.631. Table 11 summarizes

Table 12: Cancer-type and drug composition of the canonical TCGA-485 cohort. Counts denote treatment episodes. The cohort contains 485 episodes from 441 unique patients across 24 cancer types and five drugs.

Cancer type	Cisplatin	Docetaxel	Fluorouracil	Gemcitabine	Paclitaxel	Total
BRCA	0	25	0	4	48	77
PAAD	1	0	10	66	0	77
CESC	62	0	4	1	2	69
STAD	5	2	51	0	1	59
HNSC	44	3	0	1	2	50
BLCA	10	3	0	17	1	31
LUAD	6	7	0	2	3	18
UCS	2	3	0	1	10	16
COAD	0	0	15	0	0	15
UCEC	4	1	0	1	9	15
LUSC	3	3	0	5	3	14
ESCA	0	0	5	1	1	7
PLMESO	4	0	0	3	0	7
READ	0	0	5	0	0	5
CHOL	1	0	0	4	0	5
HGSOC	1	1	0	1	1	4
PRAD	0	3	0	0	0	3
SKCM	1	0	0	1	1	3
PRCC	1	0	0	1	0	2
HCC	0	0	0	2	0	2
THYM	2	0	0	0	0	2
NSGCT	2	0	0	0	0	2
ACC	1	0	0	0	0	1
MNET	0	0	0	1	0	1
Total	**150**	**51**	**90**	**112**	**82**	**485**

the drug-level composition of the cohort, whereas Table 12 reports the corresponding distribution of treatment episodes across cancer types and drugs.

A.4 PDX Cohort Construction

We constructed an independent patient-derived xenograft (PDX) treatment-response benchmark from the NIBR-PDXE resource (Gao et al., 2015) by matching patient tumor bulk RNA-seq profiles to treatment-response records at the PDX-model level. Individual PDX models could contribute multiple treatment episodes corresponding to different drugs.

Response-label construction. We derived treatment response from longitudinal tumor-volume measurements and the associated categorical response annotations (Gao et al., 2015; Nguyen et al., 2021). To obtain a consistent endpoint, we recomputed the response category for each evaluable model-treatment episode from the reported *BestResponse* and *BestAvgResponse* measurements using the NIBR-PDXE response criteria. Complete response (CR) required $\mathrm{BestResponse} < -95$ and $\mathrm{BestAvgResponse} < -40$; partial response (PR) required $\mathrm{BestResponse} < -50$ and $\mathrm{BestAvgResponse} < -20$; stable disease (SD) required $\mathrm{BestResponse} < 35$ and $\mathrm{BestAvgResponse} < 30$; and all remaining evaluable episodes were classified as progressive disease (PD). We assigned categories hierarchically in the order CR, PR, SD, and PD. The recomputed CR/PR/SD/PD base categories agreed with the corresponding base categories in the provided response annotations for all evaluable records.

Following the original NIBR-PDXE response definition (Nguyen et al., 2021), we grouped CR, PR, and SD as responders and PD as non-response:

$$y_{\mathrm{PDX}} = \begin{cases} 1, & \text{CR, PR, or SD,} \\ 0, & \text{PD.} \end{cases}$$

This benchmark-specific definition differs from TCGA-485, where SD is classified as non-response. Records without sufficient information to assign one of these four response categories were excluded.

Expression data and molecular preprocessing. We obtained patient tumor bulk RNA-seq profiles from the published NIBR-PDXE expression resource accompanying the original study (Gao et al., 2015). Unlike the TCGA target cohorts, which were constructed from GDC gene-level read counts, the PDX source matrix provides precomputed gene-level RNA-seq expression quantified as FPKM. We transposed the source matrix from a gene-by-PDX model to a model-by-gene representation and collapsed duplicate gene symbols by averaging their expression values.

To match the frozen scFoundation encoder, we aligned PDX expression profiles to the 19,264-gene scFoundation vocabulary, assigning zero to genes absent from the PDX expression matrix. We then normalized each model-level expression profile to a total expression of $10{,}000$ and transformed it using $\log(1 + x)$. We encoded the resulting profiles using the same frozen scFoundation model used for the TCGA target cohorts, yielding a 3072-dimensional tumor representation for each PDX model. Thus, we shared the downstream normalization and frozen representation model across PDX and TCGA target evaluations, while the upstream expression quantification differed: the PDX cohort used published FPKM values, whereas the TCGA cohorts were constructed from GDC gene-level read counts.

We harmonized drug names to canonical treatment identities. We matched retained small molecules to valid chemical structures and represented them using the frozen 768-dimensional ChemBERTa encoder used throughout the target-stage evaluations. Episodes lacking a matched expression profile, valid response label, or valid molecular representation were excluded.

Shared 10-drug benchmark. For cross-method comparison, we restricted evaluation to the 10 drugs shared across all evaluated methods: fluorouracil, buparlisib, alpelisib, ruxolitinib, ribociclib, erlotinib, gemcitabine, paclitaxel, tamoxifen, and trametinib. The resulting benchmark comprised 791 treatment episodes from 177 PDX models, including 343 responders and 448 non-responders, corresponding to an overall positive-response prevalence of 0.434. Because individual PDX models could contribute multiple treatment episodes, we assigned all episodes from a given model to the same outer fold to prevent overlap between training and test sets. Details of the model-disjoint evaluation protocol and target-domain tuning procedure are provided in Appendix B.8.3 and Appendix E.

Tumor-type composition. The final 10-drug benchmark included breast cancer (BRCA), cutaneous melanoma (CM), colorectal cancer (CRC), non-small-cell lung cancer (NSCLC), and pancreatic ductal adenocarcinoma (PDAC). Drug coverage was not uniform across tumor types, and some tumor-drug combinations were sparsely represented or absent. We therefore report tumor-stratified response prevalence only as a descriptive cohort characteristic, not as an estimate of tumor-type-specific treatment efficacy. Table 13 summarizes the tumor-level cohort composition, while Table 14 shows the distribution of treatment episodes across tumor-drug combinations.

Cross-cohort expression preprocessing. TCGA and PDX use different source expression quantifications: TCGA is derived from gene-level read counts, whereas the available PDX profiles are provided as FPKM values. We therefore apply the cohort-specific preprocessing described above before scFoundation encoding. This does not make the upstream quantification schemes equivalent, and the resulting preprocessing difference adds another source of target-domain heterogeneity.

A.5 Cross-benchmark Drug Coverage in Tahoe-100M

Cross-benchmark molecular coverage. To characterize how directly the Tahoe-100M source chemical space supports the target drugs, we examined all unique drugs appearing across TCGA-186, TCGA-485, and the shared PDX benchmark. Table 15 reports exact Tahoe-100M support and

Table 13: Tumor-type composition of the shared 10-drug PDX benchmark. Parenthetical labels denote the corresponding TCGA project codes. Response prevalence denotes the proportion of episodes labeled CR, PR, or SD.

Tumor type	Models	Episodes	Drugs	Response prevalence
BRCA	38	227	6	0.401
CM (SKCM)	32	64	2	0.375
CRC (COAD/READ)	43	168	7	0.470
NSCLC (LUAD/LUSC)	25	124	5	0.387
PDAC (PAAD)	39	208	6	0.486
Total	177	791	10	0.434

Table 14: Tumor-drug composition of the shared 10-drug PDX benchmark. Entries denote treatment-episode counts for each tumor type-drug combination.

Tumor type	Alpelisib	Buparlisib	Erlotinib	Fluorouracil	Gemcitabine	Paclitaxel	Ribociclib	Ruxolitinib	Tamoxifen	Trametinib	Total
BRCA	38	38	0	0	0	38	38	37	38	0	227
PDAC	36	35	0	0	34	0	35	33	0	35	208
CRC	42	40	0	41	1	0	42	1	0	1	168
NSCLC	25	25	25	0	0	24	25	0	0	0	124
CM	0	32	0	0	0	0	32	0	0	0	64
Total	**141**	**170**	**25**	**41**	**35**	**62**	**172**	**71**	**38**	**36**	**791**

Table 15: Chemical coverage of target drugs in Tahoe-100M. Exact support indicates whether the target drug is present in Tahoe-100M by normalized drug name or canonicalized molecular structure. For each target drug, we report the nearest non-identical Tahoe compound by cosine similarity of ChemBERTa embeddings. Exact target matches were excluded from nearest-neighbor searches.

Drug	Primary target / mechanism	Exact Tahoe support	Nearest non-identical Tahoe compound	ChemBERTa cosine
Alpelisib	PI3Kα	Yes	Sonidegib	0.906
Buparlisib	Pan-PI3K	No	Bimiralisib	0.991
Cisplatin	DNA cross-linking	No	Sodium Salicylate	0.753
Docetaxel	Microtubule stabilization	Yes	Docetaxel (Trihydrate)	0.996
Erlotinib	EGFR	Yes	XRK3F2	0.823
Fluorouracil	TYMS / DNA-RNA incorporation	Yes	Methylthiouracil	0.882
Gemcitabine	DNA incorporation / RNR	Yes	Decitabine	0.922
Paclitaxel	Microtubule stabilization	Yes	Vincristine	0.947
Ribociclib	CDK4/6	Yes	Palbociclib	0.902
Ruxolitinib	JAK1/2	No	Pimitespib	0.824
Sorafenib	RAF / VEGFR / PDGFR	No	Regorafenib	0.988
Tamoxifen	Estrogen receptor (SERM)	No	Diphenhydramine	0.837
Temozolomide	DNA methylation	No	Pentoxifylline	0.839
Trametinib	MEK1/2	Yes	Relugolix	0.938

Note: Primary target/mechanism annotations were curated from DrugBank (Knox et al., 2024) and denote the principal molecular target or pharmacologic mechanism of each drug. Interpret similarity estimates for cisplatin cautiously, as conventional SMILES-based molecular representations are not optimized for metal complexes. Exact Tahoe support indicates presence in Tahoe-100M and does not imply inclusion in the final source-training partition.

the nearest non-identical Tahoe compound under the frozen ChemBERTa representation used by PerturbRx.

B Training, Implementation, and Evaluation Protocols

Unless otherwise stated, all target-domain experiments use frozen predictor transfer: the source-trained intervention encoder and transition predictor remain fixed, and only the target-domain response head is trained using response labels.

B.1 Source-Stage Training Procedure

The PerturbRx pipeline consists of two stages: source-stage perturbation learning and target-stage patient response transfer.

Algorithm 1 PerturbRx source-stage training

Require: Valid matched populations $\mathcal{P} = \{(X_{\ell,b}^{\mathrm{DMSO}}, Y_{\ell,b,d,q}, u)\}$
Require: Frozen encoder E, intervention encoder g, transition predictor P_θ
for each training step **do**
  Sample $(X_{\ell,b}^{\mathrm{DMSO}}, Y_{\ell,b,d,q}, u) \sim \mathcal{P}$
  Sample mini-populations $X \subset X_{\ell,b}^{\mathrm{DMSO}}$ and $Y \subset Y_{\ell,b,d,q}$
  Compute $Z_X = E(X)$ and $Z_Y = E(Y)$
  Encode intervention $\alpha_u = g([m_d; \log q])$
  Predict $\hat{\Delta}_i = P_\theta(z_i^X, \alpha_u)$ for each $z_i^X \in Z_X$
  Form $\hat{Z}_Y = \{z_i^X + \hat{\Delta}_i : z_i^X \in Z_X\}$
  Compute $\mathcal{L}_{\mathrm{src}}$ using Equation 9
  Update trainable parameters using $\mathcal{L}_{\mathrm{src}}$
end for

Stage 1: source-stage perturbation learning. We first construct valid matched population pairs from the single-cell perturbation atlas. For each treated condition (ℓ, b, d, q), treated cells are matched to DMSO control cells from the same cell line and plate context (ℓ, b). We retain conditions only if both the control and treated pools meet minimum cell-count requirements. During training, we sample mini-populations

$$X \subset X_{\ell,b}^{\mathrm{DMSO}}, \qquad Y \subset Y_{\ell,b,d,q},$$

embed cells with the fixed pretrained encoder E, and train the intervention encoder g and transition predictor P_θ using the distributional and delta-aware objective in Equation 9. This stage learns drug- and dose-conditioned predictors that transport control-cell latent distributions toward matched treated-cell latent distributions.

Stage 2: target-domain response transfer. The frozen expression pipeline embeds patient tumor profiles, and the source-trained intervention encoder and transition predictor generate treatment-conditioned transition features. Only the response head h_ψ is trained using target-domain response labels and a class-weighted binary cross-entropy loss.

This two-stage design decouples intervention learning from patient response supervision. Source-stage training uses large-scale perturbation data to learn treatment-associated transformations of latent molecular states, while patient-stage training evaluates whether these learned intervention predictors provide a useful signal when only patient tumor profiles and response labels are available.

B.2 Source-Stage Objective Details

The source-stage objective is

$$\mathcal{L}_{\mathrm{src}} = \mathcal{L}_{\mathrm{dist}} + \lambda_{\cos}\mathcal{L}_{\cos} + \lambda_{\mathrm{norm}}\mathcal{L}_{\mathrm{norm}}. \tag{12}$$

Distributional alignment. We instantiate the distributional term as the biased empirical multi-kernel maximum mean discrepancy:

$$\mathcal{L}_{\mathrm{dist}} = \mathrm{MMD}^2(\hat{Z}_Y, Z_Y). \tag{13}$$

For $A = \{a_i\}_{i=1}^n$ and $B = \{b_j\}_{j=1}^m$,

$$\mathrm{MMD}^2(A, B) = \frac{1}{n^2}\sum_{i,i'} k(a_i, a_{i'}) + \frac{1}{m^2}\sum_{j,j'} k(b_j, b_{j'}) - \frac{2}{nm}\sum_{i,j} k(a_i, b_j). \tag{14}$$

We use a multi-kernel Gaussian kernel,

$$k(x, x') = \sum_{\sigma \in \mathcal{S}} \exp\left(-\frac{\|x - x'\|_2^2}{2\sigma^2}\right), \tag{15}$$

with bandwidth set

$$\mathcal{S} = \{8, 16, 32, 64, 128\}.$$

Transition-direction alignment. The cosine term aligns the direction of the predicted and empirical population transitions:

$$\mathcal{L}_{\cos} = 1 - \frac{\hat{\Delta}_{c,u}^{\top} \Delta_{c,u}}{\left(\|\hat{\Delta}_{c,u}\|_2 + \epsilon\right)\left(\|\Delta_{c,u}\|_2 + \epsilon\right)}, \tag{16}$$

where $\epsilon > 0$ is a small numerical constant used to avoid division by zero.

Transition-norm alignment. We additionally constrain the magnitude of the predicted population transition using

$$\mathcal{L}_{\text{norm}} = \left(\left\|\hat{\Delta}_{c,u}\right\|_2 - \|\Delta_{c,u}\|_2\right)^2. \tag{17}$$

This term penalizes discrepancies in transition magnitude independently of transition direction. Together, the cosine and norm terms constrain the direction and magnitude of the predicted population transition.

Loss-weight selection. We use $\lambda_{\cos} = 0.1$, and $\lambda_{\text{norm}} = 10^{-3}$. We selected the weights using source-stage validation without access to target-domain response labels. The loss-weight search space and selection procedure are reported in Appendix D.

B.3 Source-Stage Validation and Generalization Protocol

We evaluate source-stage generalization using two complementary held-out partitions constructed from the final $61{,}175$ Tahoe-100M source conditions spanning 377 drugs. We obtained these conditions by matching treated populations to cell-line- and plate-specific DMSO controls, requiring at least 64 cells in both treated and control populations and restricting to compounds with available molecular embeddings. Each condition corresponds to a context-matched control-treated population pair indexed by cell line, plate, drug, and dose.

Drug-held-out evaluation. We first construct a compound-disjoint evaluation partition by withholding complete drugs together with all associated cell-line, plate, and dose conditions. Drugs are ordered using split seed 0 and removed as indivisible groups until approximately 10% of the final source conditions are withheld. Because we remove complete drugs as groups, the resulting fraction may differ slightly from the nominal 10%.

This procedure yields $6{,}258$ held-out conditions spanning 42 drugs, leaving $54{,}917$ conditions for subsequent training and random-validation partitioning. We excluded fluorouracil and gemcitabine from drug-held-out selection so that their Tahoe-100M perturbation observations remained available for the corresponding patient-transfer analyses.

The drug-held-out partition evaluates inductive generalization to compounds for which no matched perturbation populations are used during source-stage training. The frozen ChemBERTa representation of each held-out compound remains available at inference, so this setting evaluates generalization of the learned treatment-conditioned transition operator rather than absence of all prior molecular information about the compound.

Random held-out validation. From the $54{,}917$ conditions remaining after drug-level withholding, we sample approximately 10% uniformly at random using random-validation seed 1. This produces $5{,}492$ random held-out conditions, with the remaining $49{,}425$ conditions used for source-stage training. The training and random held-out partitions both contain drugs represented on both sides of the split and are therefore not drug-disjoint.

The random held-out partition evaluates generalization to unseen *context-intervention conditions* within the observed source support. Because conditions are sampled independently rather than grouped by cell line, plate, drug, or dose, a held-out condition may share one or more of these factors with the training set. Accordingly, this evaluation primarily measures interpolation across observed perturbation factors and should not be interpreted as strict generalization to unseen drugs, cell lines, plates, or doses.

Evaluation metrics. We evaluate source-stage performance using MMD between the predicted and observed treated latent populations, cosine similarity between the predicted and empirical population transitions, and mean-squared error between the corresponding population transition vectors (Δ MSE). Lower MMD indicates better treated-population alignment, while higher transition cosine and lower Δ MSE indicate better agreement with the empirical population transition.

B.4 Model Architecture and Hyperparameters

We use a frozen scFoundation encoder for expression profiles. Single-cell and target-domain expression profiles are represented in the same $d = 3072$-dimensional scFoundation output space. Drug features are obtained from the frozen ChemBERTa model ("seyonec/ChemBERTa-zinc-base-v1") and concatenated with a log-dose feature before the trainable intervention encoder g produces the intervention embedding α_u.

Table 16 summarizes the architecture and model-selection settings used in the main PerturbRx experiments. We select the source-stage model exclusively using Tahoe-100M validation data and keep it frozen throughout target-domain evaluation. We use the same response-head architecture across target benchmarks. Meanwhile, hyperparameter selection and cross-validation grouping follow benchmark-specific evaluation protocols because the cohorts differ in sample size, repeated-measures structure, and the natural unit of independence (patients for TCGA and PDX models for the xenograft benchmark). These procedures are described in Appendix E.

B.5 Source-stage predictor controls

To determine whether simpler cell-wise predictors could reproduce PerturbRx's source-stage performance, we compared it against an identity mapping, a global mean transition, a linear transition predictor, a nonlinear MLP, and a cAE. All learned controls operate in the same 3072-dimensional frozen scFoundation space and use the same raw intervention condition $[m_d; \log q] \in \mathbb{R}^{769}$, consisting of the frozen 768-dimensional ChemBERTa representation and log-transformed dose. We trained them on the same source-training conditions and evaluated them on the same randomly held-out source conditions using identical sampled populations.

Identity and global-mean controls. The identity control applies no transition,

$$\hat{z}_i^Y = z_i^X,$$

so that the predicted population transition is zero; the global-mean control applies a single transition vector to every condition. This vector is estimated exclusively from the source-training partition as the mean empirical population transition across training conditions and is then held fixed during evaluation.

Cell-wise linear predictor. The linear control independently predicts a residual transition for each control-cell embedding from the concatenation of the normalized cell state and normalized raw intervention condition:

$$\hat{\Delta}_i = W\big[\mathrm{LN}(z_i^X); \mathrm{LN}([m_d; \log q])\big] + b, \qquad \hat{z}_i^Y = z_i^X + \hat{\Delta}_i.$$

Thus, this control tests whether a single linear state- and intervention-conditioned mapping is sufficient to reproduce the observed population transition.

Cell-wise MLP. The nonlinear MLP uses the same normalized cell-state and intervention inputs as the linear control but replaces the linear mapping with a two-hidden-layer feed-forward network,

$$3841 \to 1024 \to 1024 \to 3072,$$

with GELU activations and dropout 0.1 after each hidden layer. The 3072-dimensional output is interpreted as a residual transition,

$$\hat{z}_i^Y = z_i^X + \hat{\Delta}_i.$$

Unlike PerturbRx, this control uses the raw ChemBERTa-plus-dose condition directly and includes no dedicated learned intervention encoder.

Table 16: Architecture and hyperparameter details for PerturbRx.

Component	Configuration
Source-stage transition learning	
Cell encoder	Frozen scFoundation, $d = 3072$
Drug encoder	Frozen ChemBERTa, $d = 768$
Intervention encoder g	Input $768 + 1$; hidden/output $512/512$
Transition predictor P_θ	Input $3072 + 512$; hidden 1024; output 3072
Dropout	0.1
Optimizer	AdamW; LR 10^{-4}; weight decay 10^{-4}
Batch construction	32 conditions; 128 cells per population
Training	100 epochs
Source objective	$\mathrm{MMD} + \lambda_{\cos}\mathcal{L}_{\cos} + \lambda_{\mathrm{norm}}\mathcal{L}_{\mathrm{norm}}$
Loss weights	$\lambda_{\cos} = 0.1,\ \lambda_{\mathrm{norm}} = 10^{-3}$
Target-domain response prediction	
Frozen components	scFoundation, ChemBERTa, g, and P_θ
Treatment-conditioned representation	$\hat{z}_{p,u} = z_p + \hat{\Delta}_{p,u}$
Response input	$[z_p; m_d; \hat{z}_{p,u}]$
Response-head architecture	MLP selected by inner cross-validation
Architecture candidates	1–2 hidden layers; GELU/SiLU/ReLU; optional LayerNorm
Objective	Class-weighted binary cross-entropy
Optimizer	AdamW
Hidden dimension	$\{64, 128, 256\}$
Dropout	$\{0, 0.2, 0.4, 0.6\}$
Learning rate	$\{3 \times 10^{-4}, 10^{-3}, 3 \times 10^{-3}\}$
Weight decay	$\{0, 0.1, 0.2\}$
Inference dose	$\{0.05, 0.5, 5.0\}\ \mu$M
Selection criterion	Inner-validation AUROC; AP and validation loss as tie-breakers
Outer evaluation	Exactly-once prediction on each held-out outer fold
Benchmark-specific evaluation	
TCGA-186	5 patient-disjoint outer folds; 3-fold inner CV
TCGA-485	5 patient-grouped outer folds; selection within outer-training data
PDX shared 10-drug	5 model-disjoint outer folds; selection within non-held-out models

Conditional autoencoder. The Conditional AE separately compresses the cell state and intervention condition before predicting the residual transition. The cell encoder is

$$3072 \to 1024 \to 512,$$

with LayerNorm, GELU, and dropout 0.1. The intervention encoder maps

$$769 \to 512 \to 512,$$

with LayerNorm and GELU. The resulting 512-dimensional cell and intervention representations are concatenated and decoded through

$$1024 \to 1024 \to 3072,$$

with GELU and dropout 0.1. The decoder output is again interpreted as $\hat{\Delta}_i$ and added residually to z_i^X.

Training and evaluation. The three learned predictor controls were initialized with seed 123 and trained with AdamW using a learning rate of 10^{-4}, weight decay 10^{-4}, 128 sampled cells per condition, and the same sampled source-training populations as PerturbRx. To isolate architectural effects from objective design, all learned controls were trained using the same final source objective as PerturbRx,

$$\mathcal{L}_{\mathrm{control}} = \mathcal{L}_{\mathrm{MMD}} + 0.1\mathcal{L}_{\cos} + 10^{-3}\mathcal{L}_{\mathrm{norm}},$$

with no vector-displacement MSE term. We trained each model for 100 epochs and selected, for each architecture, the checkpoint maximizing transition cosine on the random source-validation split. We then evaluated the same selected checkpoint with the deterministic full-split evaluator on both the random validation and drug-held-out evaluation splits. These comparisons therefore provide objective, matched architectural controls, although model capacities were not explicitly parameter-matched.

B.6 Baseline Implementations

We compare PerturbRx against four representative baselines for patient drug-response prediction and cell-line-to-patient transfer: CODE-AE (He et al., 2022), TransDRP (Liu & Li, 2025), DeepSADR (Zhang et al., 2026), and WISER (Shubham et al., 2024). These methods span domain adaptation, transfer learning, and weakly supervised target adaptation using pharmacogenomic cell-line data.

We selected baselines based on compatibility with the target evaluation setting, prioritizing methods that support patient-level drug-response prediction on the same evaluation units without altering their original prediction task or introducing new method-specific learning objectives.

For both TCGA-485 and PDX, all target drugs were represented in the source-domain response data used by each evaluated baseline method, whereas several target drugs were not represented in PerturbRx source-stage training.

These baselines were originally developed for transferring information from cell-line drug-response data and do not natively use large single-cell perturbation atlases such as Tahoe-100M. We therefore evaluate each method under its native source-data assumptions rather than introducing method-specific adaptations to incorporate Tahoe-100M, while enforcing the same held-out target-domain evaluation boundaries used for PerturbRx.

CODE-AE. (He et al., 2022) CODE-AE learns shared and domain-specific gene-expression representations and aligns source cell lines and target-domain samples in a shared latent space. We then predict drug responses using the learned representation. We use the repository implementation adapted to the benchmark-specific expression matrices, drug identifiers, and fixed evaluation splits while preserving the method-specific training procedure.

TransDRP. (Liu & Li, 2025) TransDRP transfers cell-line drug-response information to patient samples through drug-aware response modeling and domain-adaptation objectives. We retain the method-specific model and training procedure while harmonizing gene features, drug identifiers, evaluation partitions, and prediction outputs for the common benchmark.

DeepSADR. (Zhang et al., 2026) We evaluate DeepSADR using its original model architecture and training code within a leak-free benchmark adapter. For each outer split, we restrict target-domain model selection and adaptation to the corresponding non-held-out samples. Where the original implementation uses an internally generated validation split, we replace it with the canonical benchmark-specific inner training and validation partitions so that no held-out evaluation labels are used for model selection.

WISER. (Shubham et al., 2024) We evaluate WISER using its original domain-adaptation and weak-supervision framework. The method generates target-domain pseudo-labels, selects reliable target examples, and combines them with labeled source-domain data to train the final response predictor. We adapt the released implementation to the benchmark-specific held-out splits and harmonized input data while preserving the core WISER training and weak-supervision procedure.

Drug-response modeling granularity. The evaluated methods retain their implementation-specific modeling granularity. We evaluate CODE-AE, DeepSADR, and the WISER baseline using drug-specific response models, whereas TransDRP jointly models multiple drugs. PerturbRx uses a single multi-drug response model in which the ChemBERTa embedding and the corresponding treatment-conditioned representation represent drug identity.

Baseline inputs and supervision. All methods were evaluated on the same benchmark-specific target records and held-out evaluation partitions, while retaining their method-specific input representations and source-domain supervision. CODE-AE and WISER use drug-specific response models based on their native gene-expression inputs. In contrast, TransDRP jointly models multiple drugs and additionally performs target-domain adaptation using only the corresponding outer-training partition. DeepSADR uses patient gene-expression features together with drug molecular representations and performs drug-specific target-domain training. In contrast, PerturbRx uses frozen scFoundation patient embeddings, ChemBERTa drug embeddings, and treatment-conditioned representations learned from Tahoe-100M. Thus, the comparison standardizes the target evaluation cohort and held-out boundaries rather than forcing all methods to use identical input representations or source data.

Perturbation-aware methods not included as direct baselines. We also considered the perturbation-aware methods CSG^{2}A, THERAPI, and PREDIKTOR discussed in Section 2. We did not include them as direct baselines because we could not evaluate them on our fixed TCGA and PDX partitions using their released implementations without introducing method-specific changes to their original tasks or pipelines. Evaluating CSG^{2}A in our setting would require adapting its LINCS-to-cell-line drug-sensitivity formulation to patient-level binary response prediction, including expression-space harmonization and a change from drug-sensitivity regression to the target response task. Reproducing THERAPI on our fixed TCGA and PDX partitions would require additional patient-specific preprocessing and representation-generation components not available in the released artifacts. PREDIKTOR did not provide a publicly released training and inference implementation at the time of our evaluation; its paper states that code and trained parameters will be made available upon request. We therefore restrict direct comparisons to methods with released implementations that we can evaluate on the same patient–drug prediction units and held-out partitions while preserving their original learning objectives.

Evaluation protocol. Baseline implementations preserve their method-specific architectures, objectives, and supervision regimes as far as possible while enforcing a common held-out evaluation boundary. Benchmark-specific modifications are limited to feature and identifier harmonization, predefined split handling, validation-set construction where required to prevent leakage, and prediction export.

Where applicable, we select models, apply early stopping, and perform target-domain adaptation for each baseline without access to the corresponding held-out evaluation partition. For PerturbRx, response-head and representation settings are selected using the benchmark-specific procedures described in Appendix E. No method uses response labels from the held-out outer-evaluation partition for model selection or adaptation.

We do not require identical hyperparameters or optimization procedures across methods. Instead, we standardize comparisons at the evaluation level: all methods use the same held-out target-domain partitions and respect the same test-data exclusion boundary. TCGA-186 results are evaluated across five patient-disjoint outer folds, TCGA-485 across five patient-grouped outer folds, and PDX across five PDX-model-disjoint outer folds.

B.7 Target-stage representation controls

We evaluated the following target-stage representation controls.

Patient. The response head receives only the patient tumor representation, $[z_p]$.

Patient+Drug. The response head receives the patient and drug representations, $[z_p; m_d]$, without a treatment-conditioned representation.

Random transition. For each drug, we estimated the element-wise mean and standard deviation of predicted transitions from the corresponding training partition and sampled synthetic transitions from the resulting diagonal Gaussian distribution. We generated validation and held-out test transitions using statistics estimated exclusively from the corresponding training partition.

Shuffled transitions. We randomly reassigned predicted transitions using donors from the full training transition pool, disrupting both drug–transition and patient–transition correspondence. Held-out examples received only transitions drawn from the corresponding training partition.

Within-drug shuffle. Predicted transitions were randomly reassigned using donors from training examples receiving the same drug, preserving drug identity while disrupting patient-specific transition correspondence. Validation and held-out test examples received only same-drug transitions drawn from the corresponding training partition.

B.8 Benchmark-Specific Target-Domain Evaluation Protocols

Across all target benchmarks, we use fixed five-fold outer partitions and restrict all target-domain model selection to the corresponding outer-training partition. For PerturbRx, we keep the pretrained intervention encoder and transition predictor frozen throughout evaluation. Response-head hyperparameters and, for transition-based representations, inference-dose settings are selected using only non-held-out data under the benchmark-specific inner-validation procedure described in Appendix E. After model selection, we refit the selected response model on the complete outer-training partition and evaluate it once on the held-out outer-test partition. We do not use outer-test response labels or performance metrics for hyperparameter selection, early stopping, representation selection, dose selection, or target-domain adaptation.

B.8.1 TCGA-186: Seen-Drug, Held-Out-Patient Evaluation

TCGA-186 contains 186 patient–drug response records from 186 unique patients; because each patient contributes exactly one record, patient-level and record-level grouping are equivalent. We use a single canonical five-fold outer partition stratified jointly by drug identity and response label, with mutually exclusive and exhaustive test folds and the same fixed assignment across evaluated methods.

Within each outer-training partition, PerturbRx uses three stratified inner folds based on joint drug–response strata. We rank candidate configurations by mean inner-validation AUROC, with mean inner-fold average precision and validation loss as tie-breakers. Concatenating predictions from the five held-out outer folds yields one exactly-once OOF prediction for every patient. We compute AUROC and AUPRC separately within each drug and report their unweighted macro-average as the primary aggregate metric, thereby giving each treatment equal weight despite differences in cohort size and response prevalence. Drug-specific metrics are computed from the same exactly-once OOF predictions.

B.8.2 TCGA-485: Patient-Grouped Evaluation

We evaluate TCGA-485 using five patient-grouped outer folds, with all treatment episodes from the same patient assigned to the same fold. The benchmark contains five drugs shared across the evaluated methods: cisplatin, docetaxel, fluorouracil, gemcitabine, and paclitaxel. We use the same fixed outer-fold assignment for all methods.

We restrict all response-model fitting, feature normalization, model selection, and method-specific target-domain adaptation to the corresponding outer-training partition. For PerturbRx, the response model uses the patient tumor embedding, ChemBERTa drug embedding, and predicted treatment-conditioned representation. Predictions from the five mutually exclusive outer-test folds are concatenated to obtain exactly-once OOF predictions. We report pooled OOF AUROC and AUPRC as the primary aggregate metrics and additionally report drug-specific performance from the corresponding OOF predictions.

B.8.3 PDX: Model-Disjoint Evaluation

We evaluate the PDX benchmark using five outer folds grouped by PDX model identifier, so all treatment episodes from the same model remain within a single fold. After excluding the single PDX model lacking a tumor-type annotation, the primary cross-method comparison uses the 10 drugs shared across the evaluated methods, fluorouracil, buparlisib, alpelisib, ruxolitinib, ribociclib,

erlotinib, gemcitabine, paclitaxel, tamoxifen, and trametinib, yielding 791 model–drug treatment episodes from 177 PDX models.

For each outer fold, we restrict all response-model fitting, method-specific target-domain adaptation, pseudo-labeling, feature normalization, and supervised PDX fine-tuning to the non-held-out models. For PerturbRx, the response predictor uses the PDX embedding, ChemBERTa drug embedding, and treatment-conditioned representation. For transition-based representations, the inference dose is selected from $q \in \{0.05, 0.5, 5.0\}\,\mu\mathrm{M}$ jointly with the response-model settings using only non-held-out models.

Predictions from the five model-disjoint outer-test folds are concatenated to obtain exactly-once OOF predictions. Aggregate AUROC and AUPRC are computed from the pooled OOF predictions, with fold-wise variability reported separately. Drug-specific AUROC and AUPRC are computed within each held-out outer fold and summarized as mean $\pm$ standard deviation across the five model-disjoint folds.

C Additional Results and Diagnostics

C.1 Drug-held-out source-stage generalization

We further evaluated whether the source-stage transition predictor generalizes to compounds excluded entirely from source-stage training. Before constructing the random held-out split, we withheld 6,258 source conditions spanning 42 drugs, with all cell-line, plate, and dose conditions for each held-out drug excluded from source-stage training. This drug-held-out evaluation differs from the random held-out evaluation in Section B.3: the latter evaluates unseen context-intervention conditions for drugs represented during training, whereas the drug-held-out setting evaluates inductive generalization to unseen compounds.

Figure 2D reports transition-direction agreement across the drug-held-out set. Under the matched MMD + cosine + norm source objective, PerturbRx achieved the highest overall transition cosine among the evaluated predictors (0.229), compared with Conditional AE (0.198), MLP (0.167), and Linear (0.137). Performance nevertheless varied substantially across individual compounds, indicating that generalization to unseen drugs remains treatment-dependent rather than uniform across the held-out chemical space.

We additionally evaluated distributional alignment and transition-displacement error on the same drug-held-out set (Figure 3). PerturbRx achieved the lowest overall MMD (0.111), compared with Conditional AE (0.115), MLP (0.121), and Linear (0.155), and the lowest overall Δ MSE (0.021), compared with Conditional AE (0.022), MLP (0.023), and Linear (0.034). Across these objective-matched comparisons, PerturbRx achieved the best aggregate performance among the evaluated predictors on all three drug-held-out metrics: transition direction, treated-state distribution matching, and transition-displacement error.

C.2 Secondary Prior-Work Benchmark: TCGA-186

We evaluate whether treatment-conditioned representations transferred from the source perturbation domain predict patient response in TCGA-186 under a common patient-disjoint, exactly-once out-of-fold (OOF) protocol. We compare PerturbRx with existing methods and use representation ablations and negative controls to assess the contribution of the treatment-conditioned representation beyond static patient and drug representations.

Table 17 reports treatment-stratified exactly-once OOF performance and drug-macro AUROC/AUPRC, which give each treatment equal weight. PerturbRx achieved the strongest drug-macro performance among the evaluated methods (0.604/0.607 AUROC/AUPRC), compared with 0.492/0.531 for the strongest baseline, CODE-AE.

Performance was heterogeneous across treatments: PerturbRx achieved the highest AUROC and AUPRC for cisplatin, gemcitabine, and sorafenib, but not for fluorouracil or temozolomide. Thus, the advantage was not uniform across drugs, motivating the representation controls to assess whether treatment-conditioned representations contribute beyond static patient and drug information.

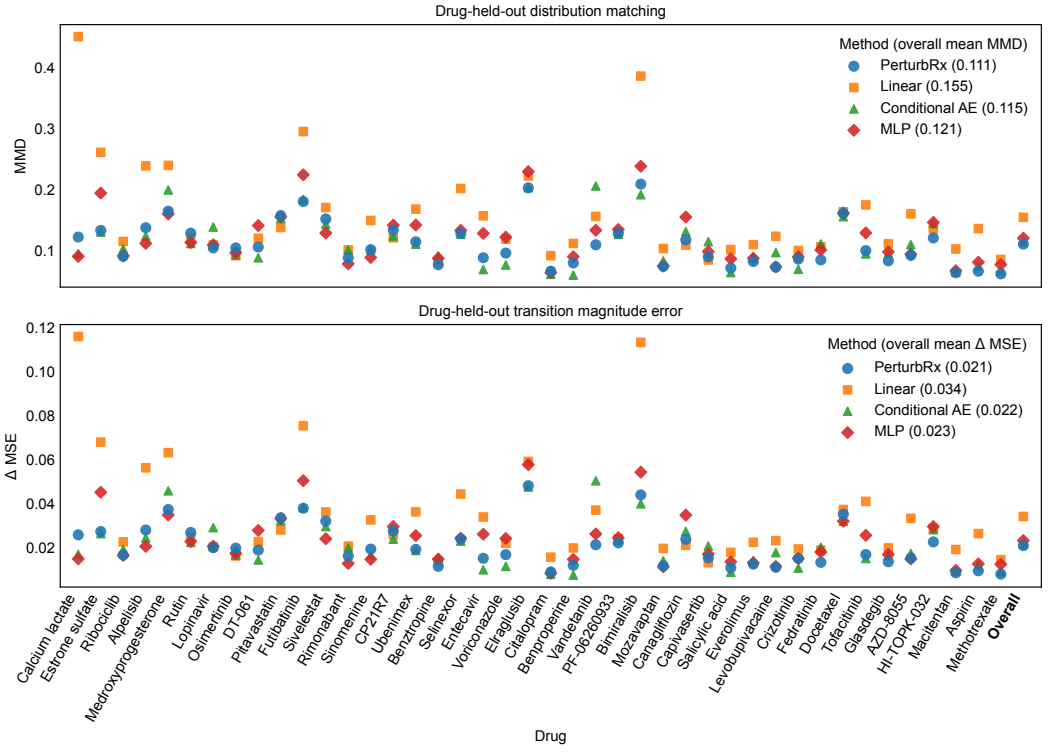


Figure 3: **Drug-held-out source-stage generalization under complementary transition-prediction metrics.** Drug-wise performance for compounds excluded entirely from source-stage training. **A.** MMD between predicted and observed treated latent populations; lower values indicate better distributional alignment. **B.** Mean-squared error between predicted and empirical population transitions (Δ MSE); lower values indicate better displacement fidelity. Legend values report aggregate performance across the drug-held-out evaluation.

Table 17: Patient-response prediction on the TCGA-186 benchmark.

	Drug-macro		**Cisplatin** ($n = 40$)		**Fluorouracil**‡ ($n = 28$)		**Gemcitabine**‡ ($n = 46$)		**Sorafenib** ($n = 26$)		**Temozolomide** ($n = 46$)	
Methods	**AUROC↑**	**AUPRC↑**	**AUROC↑**	**AUPRC↑**	**AUROC↑**	**AUPRC↑**	**AUROC↑**	**AUPRC↑**	**AUROC↑**	**AUPRC↑**	**AUROC↑**	**AUPRC↑**
CODE-AE	0.492	0.531	0.440	0.552†	**0.490**	**0.421**	0.348	0.419	0.598	0.667	**0.586**	**0.596**
Std. dev.	±0.033	±0.049	±0.248	±0.156	±0.209	±0.197	±0.278	±0.123	±0.091	±0.132	±0.201	±0.100
DeepSADR	0.455†	0.496†	0.490†	0.541	0.411	0.394†	0.456†	0.496	0.509†	0.548	0.408	0.504
Std. dev.	±0.033	±0.014	±0.204	±0.122	±0.286	±0.189	±0.214	±0.163	±0.160	±0.116	±0.195	±0.116
TransDRP	0.426	0.488	0.383	0.474	0.401	0.390	0.465	0.467	0.355	0.569†	0.527†	0.538†
Std. dev.	±0.063	±0.057	±0.234	±0.152	±0.211	±0.168	±0.292	±0.140	±0.194	±0.158	±0.117	±0.129
WISER	0.452	0.493	0.505	0.557	0.411†	0.385	0.399	0.476†	0.420	0.501	0.524	0.548
Std. dev.	±0.029	±0.047	±0.120	±0.120	±0.253	±0.221	±0.091	±0.080	±0.145	±0.082	±0.264	±0.194
PerturbRx	**0.604**	**0.607**	**0.725**	**0.774**	0.432	0.411	**0.552**	**0.577**	**0.740**	**0.750**	0.571	0.525
Std. dev.	±0.051	±0.064	±0.135	±0.101	±0.242	±0.226	±0.215	±0.153	±0.354	±0.246	±0.152	±0.086

Drug-specific metrics use exactly-once OOF predictions; Drug-macro is the unweighted mean across five drugs. Std. dev. denotes fold variation. ‡ denotes drugs represented in source-stage training. Bold, underline, and † denote the top three results.

We assessed whether the treatment-conditioned representation improves within-treatment discrimination by computing AUROC and AUPRC within each drug and macro-averaging across treatments.

Table 18: Ablation results on TCGA-186.

Representation	AUROC↑	AUPRC↑
Patient	0.595	**0.615**
	Std. dev. ±0.059	±0.044
Patient+Drug	0.592†	0.590
	Std. dev. ±0.070	±0.048
Random transition	0.567	0.558
	Std. dev. ±0.049	±0.071
Shuffled transition	0.579	0.606†
	Std. dev. ±0.079	±0.069
Within-drug shuffle	0.586	0.589
	Std. dev. ±0.052	±0.039
PerturbRx	**0.604**	0.607
	Std. dev. ±0.051	±0.064

Metrics are computed from exactly-once OOF predictions within each drug and reported as unweighted drug-macro averages.

PerturbRx achieved the highest AUROC (0.604), whereas the patient-only representation achieved the highest AUPRC (0.615 versus 0.607 for PerturbRx). Thus, the treatment-conditioned representation improved AUROC over the static controls but did not uniformly improve both metrics (Table 18).

The within-drug shuffle, which preserves treatment identity while disrupting patient-specific transition correspondence, reduced performance to 0.586/0.589 AUROC/AUPRC. This is consistent with patient-specific transition correspondence contributing signal beyond treatment identity. However, the patient-only baseline, using the frozen scFoundation tumor representation, highlights the limited resolution of this benchmark for isolating the incremental contribution of treatment-conditioned features. We therefore treat TCGA-186 primarily as a benchmark for comparison with prior work and examine this contribution in the larger TCGA-485 cohort and independent PDX benchmark.

C.3 Paired inference

For aggregate benchmark-level comparisons, we aligned exactly-once OOF predictions across methods and performed paired resampling at the independent cluster level (patients for TCGA and PDX models for PDX). We used 10,000 paired bootstrap replicates for 95% confidence intervals and 10,000 paired cluster-level randomization/permutation replicates for two-sided p-values. We applied Holm correction across comparators separately for each benchmark and metric.

Paired drug-specific inference used exactly-once out-of-fold predictions after canonical key alignment. We performed paired resampling at the independent-cluster level: patients for TCGA-186 and TCGA-485, and PDX models for PDX. We evaluated AUROC and average precision (AUPRC) separately within each drug. For each comparison, we used 10,000 paired bootstrap replicates to obtain 95% confidence intervals and 10,000 paired cluster-level randomization/permutation replicates to obtain two-sided p-values. We applied Holm correction within each benchmark $\times$ drug $\times$ metric family across comparators. As a secondary, more stringent analysis, we also applied Holm correction within each benchmark $\times$ metric family across all drugs and comparators. We treated single-class bootstrap replicates as invalid, excluded them from percentile calculations, and counted them separately.

D Source-Stage Model Selection and Hyperparameter Tuning

We performed all source-stage model selection using Tahoe-100M source validation data only; we did not use TCGA or PDX response labels to select the source objective, molecular representation, or loss weights.

Validation metrics. For each validation condition, we compare the predicted population-level transition with the empirical population-level transition:

$$\hat{\Delta} = \mu\Big(\hat{Z}_Y\Big) - \mu(Z_X)\,, \qquad \Delta = \mu(Z_Y) - \mu(Z_X)\,.$$

Our primary source-stage checkpoint-selection metric is the cosine similarity,

$$s_{\cos} = \cos\Big(\hat{\Delta}, \Delta\Big)\,.$$

Higher values indicate better agreement between the predicted and empirical transition directions. During training, the corresponding cosine loss is defined as

$$\mathcal{L}_{\cos} = 1 - s_{\cos},$$

and is therefore minimized. We report cosine similarity rather than cosine loss for validation results.

We use two source-stage held-out partitions. The random validation split contains randomly held-out source conditions from drugs represented during source training and is used for model selection. The drug-held-out evaluation split contains compounds excluded entirely from source-stage training and assesses generalization to unseen drugs. Within each training run, we selected the checkpoint that maximized transition cosine on the random validation split. We then evaluated the same selected checkpoint on both the random and drug-held-out splits using a common deterministic full-split evaluator. The corrected source benchmark contains 5,492 random-validation conditions and 6,258 drug-held-out conditions spanning 42 held-out drugs.

In addition to transition cosine, we report the maximum mean discrepancy (MMD) between predicted and observed treated-state populations and the mean-squared error between predicted and empirical population-level transition vectors,

$$\Delta\mathrm{MSE} = \frac{1}{D}\left\|\hat{\Delta} - \Delta\right\|_2^2.$$

Higher transition cosine and lower MMD and Δ MSE indicate better source-stage prediction.

Loss-weight sweep. We first evaluated the source-stage objective

$$\mathcal{L} = \lambda_{\mathrm{dist}}\mathcal{L}_{\mathrm{dist}} + \lambda_{\cos}\mathcal{L}_{\cos} + \lambda_{\mathrm{mse}}\mathcal{L}_{\mathrm{mse}} + \lambda_{\mathrm{norm}}\mathcal{L}_{\mathrm{norm}},$$

with $\lambda_{\mathrm{dist}} = 1$. For the initial MMD-based sweep, we evaluated

$$\lambda_{\cos} \in \{0, 0.1, 0.3, 1.0\}, \qquad \lambda_{\mathrm{mse}} \in \{0, 0.3, 1.0, 3.0\}, \qquad \lambda_{\mathrm{norm}} \in \{0, 10^{-3}\},$$

yielding $4 \times 4 \times 2 = 32$ configurations. We trained each configuration for 300 optimization steps under otherwise identical source-stage settings. This short-run sweep identified promising loss-weight regimes for subsequent longer-run comparisons.

Distributional objective and molecular representation. We compared MMD with entropically regularized Sinkhorn optimal transport as the population-level distributional loss. We additionally evaluated three pretrained molecular representations: ChemBERTa-zinc-base-v1, ChemBERTa-77M-MTR, and MoLFormer. We evaluated molecular representations on the same corrected 377-drug source subset, enabling matched source-stage comparisons. Based on these source-only comparisons, we retained ChemBERTa-zinc-base-v1 for the subsequent distributional-objective and objective-component analyses.

We next compared longer-run MMD and Sinkhorn candidate configurations using the same corrected source splits and common evaluation procedure. All runs were trained for up to 100 epochs, and we evaluated the checkpoint that maximized the random-validation transition cosine. Table 19 reports the resulting condition-weighted metrics.

Table 19: Source-stage comparison of MMD- and Sinkhorn-based distributional objectives. All models use ChemBERTa-zinc-base-v1 and $\lambda_{\mathrm{norm}} = 10^{-3}$. For each run, the checkpoint maximizing transition cosine on the random validation split was selected and evaluated with the same deterministic full-split evaluator on both random and drug-held-out source conditions. Higher transition cosine and lower MMD and Δ MSE are better. Bold, underline, and † denote the best, second-best, and third-best results within each split and metric, respectively.

			Random held-out			Drug held-out		
Distribution loss	$\boldsymbol{\lambda_{\cos}}$	$\boldsymbol{\lambda_{\mathrm{mse}}}$	**MMD ↓**	**Δ cosine ↑**	**Δ MSE ↓**	**MMD ↓**	**Δ cosine ↑**	**Δ MSE ↓**
MMD	0.1	1.0	**0.0682**	**0.6823**	**0.0096**	<u>0.1167</u>	0.1891	0.0226
MMD	0.1	3.0	<u>0.0694</u>	<u>0.6698</u>	<u>0.0099</u>	**0.1153**	0.1895	0.0223
Sinkhorn ($\epsilon = 5$)	0.3	1.0	0.1895	0.6145	0.0114	0.2379	0.2371†	0.0222†
Sinkhorn ($\epsilon = 10$)	0.3	0.3	0.1892	0.6079	0.0117	0.2328	<u>0.2399</u>	<u>0.0218</u>
Sinkhorn ($\epsilon = 20$)	1.0	0.3	0.1840†	0.6282†	0.0112†	0.2311†	**0.2706**	**0.0210**

The distributional objectives behaved differently across validation regimes. MMD-based candidates achieved substantially lower treated-state MMD and stronger transition metrics on randomly held-out conditions. On the drug-held-out split, the Sinkhorn candidates achieved higher transition cosine

and slightly lower Δ MSE, whereas the MMD candidates retained substantially lower treated-state MMD. These results motivated a focused analysis of the transition-aware terms within the MMD objective rather than selecting a final source objective from the distributional-loss comparison alone.

Objective-component ablation. We next evaluated the contribution of the transition-aware objective components using matched 100-epoch MMD-based source-stage runs. We compared four objectives: MMD alone; MMD with vector-displacement MSE and transition- norm alignment; MMD with cosine and transition-norm alignment; and the full MMD + cosine + vector-MSE + norm objective. For each objective, we selected the checkpoint that maximized transition cosine on the random validation split and evaluated that same checkpoint using the same deterministic full-split evaluation procedure on both held-out splits.

On the random validation split, MMD + cosine + norm performed best across all three metrics and was therefore selected as the final source objective. The same selected configuration also performed best across all three metrics on the separate drug-held-out split (Table 20).

Adding vector-displacement MSE did not improve performance on the random validation split. The same pattern was observed on the drug-held-out split: adding the MSE term reduced transition cosine from 0.2290 to 0.1891 and increased Δ MSE from 0.0209 to 0.0226.

Table 20: Source-stage objective-component ablation using matched 100-epoch MMD-based training runs. For each objective, we selected the checkpoint that maximized transition cosine on the random validation split and evaluated it with the same deterministic full-split evaluator under both random and drug-held-out source conditions. Higher transition cosine and lower MMD and Δ MSE are better. Bold, underline, and † denote the best, second-best, and third-best results within each column, respectively.

	Random held-out			Drug held-out		
Source objective	**MMD ↓**	**Δ cosine ↑**	**Δ MSE ↓**	**MMD ↓**	**Δ cosine ↑**	**Δ MSE ↓**
MMD	0.0690†	0.6369	0.0101	<u>0.1132</u>	0.1437	<u>0.0218</u>
MMD + 1.0 MSE + 10^{-3} norm	0.0692	0.6554†	0.0101†	0.1218	0.1735†	0.0245
MMD + 0.1 cos + 10^{-3} norm	**0.0671**	**0.6882**	**0.0093**	**0.1114**	**0.2290**	**0.0209**
MMD + 0.1 cos + 1.0 MSE + 10^{-3} norm	<u>0.0682</u>	<u>0.6823</u>	<u>0.0096</u>	0.1167†	<u>0.1891</u>	0.0226†

Final source objective. We selected the final source objective based only on the random source-validation split, without using the drug-held-out split or any target-domain response labels or performance metrics. Within each training run, we selected the checkpoint maximizing transition cosine on the random source-validation split. We used the separate drug-held-out split only to assess generalization to compounds absent from source-stage training.

The initial distributional-objective comparison showed that MMD-based models provided substantially better treated-state distribution matching on the random split. At the same time, Sinkhorn-based candidates achieved stronger transition direction agreement for held-out drugs. The subsequent matched MMD objective-component ablation resolved this trade-off on the random validation split: MMD + cosine + norm achieved the best MMD, transition cosine, and Δ MSE among the four MMD objective variants and was therefore selected as the final source-stage objective,

$$\mathcal{L}_{\text{source}} = \mathcal{L}_{\text{MMD}} + 0.1\mathcal{L}_{\cos} + 10^{-3}\mathcal{L}_{\text{norm}}.$$

The selected objective also achieved the best values for all three metrics on the separate drug-held-out evaluation split. ChemBERTa-zinc-base-v1 was retained as the molecular representation. We did not use TCGA or PDX response labels for these source-stage decisions.

E Target-Domain Hyperparameter and Representation-Setting Selection

We performed target-domain model selection separately from source-stage model selection. We fixed the pretrained intervention encoder and transition predictor before target-domain response

modeling; we selected only the response head and target-domain representation settings using response data from the corresponding outer-training partition.

Response-head architecture and hyperparameters. The response head takes the treatment-conditioned representation

$$[z_p; m_d; \hat{z}_{p,u}], \qquad \hat{z}_{p,u} = z_p + \hat{\Delta}_{p,u},$$

and outputs the predicted treatment-response probability. We used the same predefined architecture and optimization search space across representation conditions. Candidate response heads comprised one- or two-hidden-layer MLPs with GELU, SiLU, or ReLU activations and optional LayerNorm. We varied hidden dimension ($\{64, 128, 256\}$), dropout ($\{0, 0.2, 0.4, 0.6\}$), learning rate ($\{3 \times 10^{-4}, 10^{-3}, 3 \times 10^{-3}\}$), and weight decay ($\{0, 0.1, 0.2\}$).

Within each outer-training partition, we evaluated candidate configurations by three-fold inner cross-validation. We ranked them by mean inner-validation AUROC, using mean inner-fold average precision and validation loss as tie-breakers. We determined early-stopping epochs exclusively from inner-validation data.

Inference-dose setting. For transition-based representations, we included inference dose in the same inner-loop model-selection procedure, with candidate values:

$$q \in \{0.05, 0.5, 5.0\}\,\mu\mathrm{M}.$$

We treat this dose as a representation hyperparameter supplied to the frozen source-stage intervention encoder rather than as a calibrated clinical or in vivo exposure.

Separation from source-stage model selection. We selected source-stage objectives, molecular representations, and loss weights exclusively using Tahoe-100M validation data and fixed them before target-domain evaluation. We used no held-out outer-test labels for target-domain model selection. Benchmark-specific grouping, refitting, and exactly-once outer-fold evaluation procedures are described in Appendix B.8.