CHARM-FLOW: CELLULAR HIERARCHY-AWARE RESPONSE MODELING FOR SINGLE-CELL PERTURBATIONS

Anonymous authors
Paper under double-blind review

ABSTRACT

Predicting single-cell perturbation responses from unpaired observations requires capturing both population-level expression changes and variation among responding cells. We introduce CHARM-Flow, a cellular hierarchy-aware generative framework that models a population-level response together with heterogeneous cellular variation around that response. A shared control-derived population–state–cell hierarchy coordinates response prediction, conditional source composition, and cell-level residual generation. Condition-dependent weights organize empirical control cells into a state-structured source, while each sampled source cell remains available as a persistent reference during generation. A control-derived gene partition parameterizes the source and evolving residual representations, and conditional flow matching provides the residual update mechanism. Across Replogle, PBMC, and Tahoe, CHARM-Flow captures perturbation-level population responses, achieving PDCorr values of 0.477, 0.870, and 0.959, respectively. Controlled analyses on Replogle show that state structured empirical sources and persistent source access improve agreement with observed cellular distributions within perturbations. The shared hierarchy connects population-level response prediction, empirical source construction, and heterogeneous cell-level generation.

1 INTRODUCTION

Single-cell perturbation experiments reveal transcriptional responses that vary across cells exposed to the same intervention (Dixit et al., 2016; Srivatsan et al., 2020). Genome-scale Perturb-seq links genetic perturbations to diverse expression phenotypes across large cell populations (Replogle et al., 2022). Even within a shared genetic background, pre-existing expression differences can be associated with distinct treatment outcomes, as illustrated by rare transcriptional states linked to drug resistance (Shaffer et al., 2017). This heterogeneity motivates modeling responses at multiple levels of cellular organization: the population provides a shared context, its constituent states describe internal structure, and individual cells vary within each state.

Single-cell transcriptomic assays are typically destructive, yielding unpaired control and perturbed populations. The same cell is generally not measured before and after an intervention, so individual transitions are unobserved, and prediction involves learning a conditional transformation between distributions (Bunne et al., 2023; Klein et al., 2025). Control measurements nevertheless provide observable structure for this task. Population composition captures shared context, groups of cells characterize transcriptional states, and observed expression profiles reflect within-state variation. We use this population–state–cell hierarchy to organize information that conditions the generator.

Existing methods address response generalization through latent perturbation representations, compositional condition modeling, and gene-relation priors (Lotfollahi et al., 2019; 2023; Roohani et al., 2024). Distributional approaches learn transformations between control and perturbed populations (Bunne et al., 2023; Klein et al., 2025), while set-based predictors incorporate population context (Adduri et al., 2026). Recent work explores condition-dependent mixture sources (Rubbi et al., 2026), hierarchical population–cell encoders (Chen et al., 2026), control-referenced expression architectures (Yu et al., 2026), and gene-module representations for conditional flow matching (Ruan

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.

et al., 2026). These developments motivate using a shared control-derived hierarchy to coordinate population-level response modeling, source composition, and cellular variation.

Our approach uses this hierarchy to model perturbation responses at population and cellular levels. Control-derived states define a condition-dependent empirical source whose observed member cells retain the heterogeneity present within each state. Cells sampled from this source remain available as references throughout generation, together with their local state context.

We introduce CHARM-FLOW, a cellular hierarchy-aware framework for modeling single-cell perturbation responses. A shared control-derived population–state–cell hierarchy coordinates response prediction, source composition, and cell-level generation. Condition-dependent state weights organize observed control cells into an empirical source that retains within-state variation. Sampled source cells remain available throughout generation as references, linking state context to individual cellular responses. Gene-level structure is modeled within this hierarchy.

We evaluate the framework on Replogle, PBMC, and Tahoe, covering genetic perturbations, cytokine stimulation, and drug treatments across distinct cellular contexts.

Our contributions are threefold:

- **Shared hierarchy for perturbation-response modeling.** A common control-derived population–state–cell hierarchy coordinates population response prediction, source composition, state context, and cell-level residual generation.
- **State-structured empirical source modeling.** Condition-dependent state weights govern source composition, and empirical control cells retain within-state variation. The source distribution follows the same hierarchy that conditions the response model.
- **Persistent source references for residual generation.** Starting from empirical control cells, the centered residual generator retains state context and each sampled source cell throughout generation, linking state-level organization to individual expression variation around the predicted response mean.

2 RELATED WORK

2.1 POPULATION CONTEXT AND CELL-STATE HETEROGENEITY

Single-cell perturbation models address generalization across cellular backgrounds and interventions using different response representations. scGen transfers perturbation effects through latent shifts (Lotfollahi et al., 2019), CPA represents basal state, perturbations, and covariates as composable components (Lotfollahi et al., 2023), and GEARS uses gene-relation priors to predict unseen genetic interventions and their combinations (Roohani et al., 2024). Population context also informs cell representations incorporate population context. State connects cell representations through set-level attention (Adduri et al., 2026), while SCALE combines each cell representation with a shared population summary before conditional latent transport (Chen et al., 2026).

Responses can also depend on a cell's basal state. CellCap uses attention and a sparse dictionary to associate basal cell states with transcriptional response programs (Xu et al., 2025). DeMixPert distinguishes basal-state-dependent systematic responses, perturbation-specific responses, and population variation (Liu et al., 2026). Population-context models emphasize the connection between shared context and local representations, while these heterogeneity models emphasize response dependence on basal state. CHARM-FLOW uses a shared population-state-cell hierarchy to coordinate response prediction, source composition, and cellular generation.

2.2 CONDITIONAL SOURCES FOR POPULATION TRANSPORT

Unpaired control and perturbed observations motivate distributional approaches to perturbation prediction. CellOT learns mappings between populations through neural optimal transport (Bunne et al., 2023), and CellFlow applies conditional flow matching to single-cell phenotype generation under complex perturbation conditions (Klein et al., 2025). CMonge learns optimal transport maps conditioned on experimental contexts (Driessen et al., 2026). OCOO-T uses a Transformer operating directly on continuous expression profiles, with expression patching and a Gaussian initial

distribution (Jiang et al., 2026). These approaches reflect different choices of representation space, source distribution, and transport parameterization.

The source distribution can itself adapt to the condition. CSFM studies condition-dependent sources and regularization for source learning (Kim et al., 2026), while SP-FM parameterizes both a learnable mixture base distribution and the vector field using experimental conditions (Rubbi et al., 2026). In CHARM-FLOW, mixture components correspond to observed control states, and condition-dependent weights determine their contribution to an empirical source. Observed control cells retain within-state variation, while source composition follows the same hierarchy used for response prediction.

2.3 CONTROL REFERENCES AND GENE-LEVEL DYNAMICS

Modeling individual expression changes involves both the use of control references and the organization of dependencies among genes. scDFM combines distributional alignment, gene-graph-masked encoding, and cross-differential attention to control representations (Yu et al., 2026). GeneGeoFlow combines a control-anchored residual field with gene geometry derived from biological networks (Li et al., 2026). Gene programs provide another way to organize expression, with scBIG learning program representations and inter-program interactions for conditional flow matching (Ruan et al., 2026).

In CHARM-FLOW, the sampled source cell remains available throughout generation, while population and state representations provide broader context. Gene-level dependencies are modeled within this source-referenced hierarchy using a control-derived partition.

3 METHOD

3.1 PERTURBATION-CONDITIONED POPULATION GENERATION

A perturbation reshapes the expression distribution of a cell population, producing both a population-level response and heterogeneous changes across cells. Single-cell assays do not provide paired measurements of the same cells before and after perturbation. We therefore formulate response prediction at the population level, rather than in terms of unobserved cellwise transitions. Let $\mathcal{X}$ denote the observed control cells and ξ the experimental condition, including the perturbation, cellular context, and associated covariates. We model the resulting population expression distribution under condition ξ, using $\mathcal{X}$ as the observed basal population.

CHARM-FLOW models this transformation using a population–state–cell hierarchy derived from the observed controls. The hierarchy determines how control states contribute to the empirical source and provides population and state context for both response prediction and population generation. Each sampled control cell also provides an individual source reference that remains available throughout generation. We denote the resulting hierarchical memory by $\mathcal{H} = \mathcal{E}_\theta(\mathcal{X}, \xi)$. For a requested population size N, the conditional distribution of predicted populations is defined by

$$P_{\theta,N}(\cdot \mid \mathcal{X}, \xi) = \left(\mathcal{F}_\theta^{\mathcal{X},\xi}\right)_{\#} Q_{\theta,N}(\cdot \mid \mathcal{X}, \xi). \tag{1}$$

The source law $Q_{\theta,N}$ governs the joint sampling of N control profiles and their state identities, denoted $(S, \boldsymbol{k})$. The generative map $\mathcal{F}_\theta^{\mathcal{X},\xi}$ transforms this sampled source population into a predicted perturbed population $\widehat{Y} \in \mathbb{R}_{\geq 0}^{N \times G}$, where G is the number of modeled expression coordinates. Equation (1) defines $P_{\theta,N}$ as the pushforward of the source law $Q_{\theta,N}$ under $\mathcal{F}_\theta^{\mathcal{X},\xi}$. Figure 1 summarizes the framework.

3.2 HIERARCHICAL ORGANIZATION OF CELLULAR HETEROGENEITY

The observed control population can contain distinct transcriptional states and substantial variation among cells within each state. To retain this structure, we represent the observed controls through a population–state–cell hierarchy, where the population provides global context, transcriptional states describe its internal structure, and individual cells retain variation within each state.

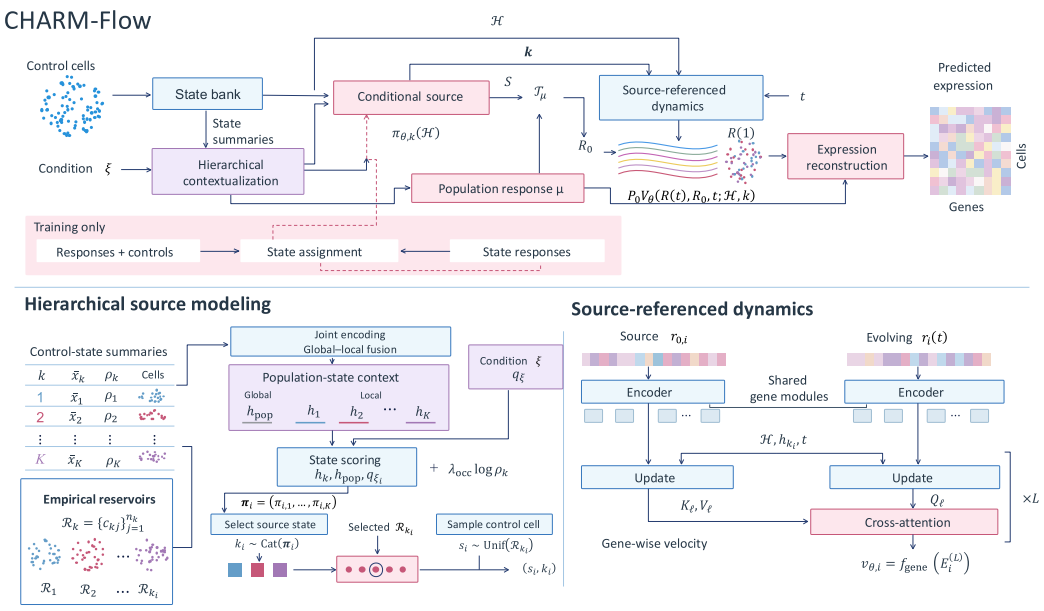

Figure 1: Overview of CHARM-Flow. The upper panel presents the overall CHARM-Flow pipeline, where a shared population–state representation informs population response prediction and empirical source construction, while sampled source cells remain available as individual references throughout residual generation to produce the predicted perturbed population. The lower-left panel details the global–local contextualization of control-derived states and shows how the resulting population–state representation guides state weighting and empirical cell sampling. The lower-right panel details the source-referenced gene-module updates and shows how the evolving representation incorporates information from the fixed source reference throughout residual generation. The shaded lower band denotes training-only supervision.

Within each control context, we derive transcriptional states by clustering control expression profiles. The context is defined by the cellular background and relevant experimental covariates, including cell-type annotations where applicable. Let z_i denote the fixed representation of control cell x_i used for state clustering. For a candidate group $\mathcal{B} \subseteq \mathcal{X}$, we define its clustering distance as

$$d(x_i, \mathcal{B}) = \|z_i - \bar{z}_{\mathcal{B}}\|_2^2, \tag{2}$$

where $\bar{z}_{\mathcal{B}}$ is the mean representation of the cells in $\mathcal{B}$.

Constrained bisecting k-means recursively splits the control population according to this distance:

$$\{\mathcal{X}_k\}_{k=1}^{K} = \mathrm{BisectingKMeans}(\mathcal{X}; d). \tag{3}$$

The resulting K groups form a disjoint partition of $\mathcal{X}$, with each group corresponding to one transcriptional state in the hierarchy.

For each state, we retain its member cells and summarize the state by its mean expression and relative abundance:

$$\bar{x}_k = \frac{1}{|\mathcal{X}_k|} \sum_{x_i \in \mathcal{X}_k} x_i, \qquad \rho_k = \frac{|\mathcal{X}_k|}{|\mathcal{X}|}. \tag{4}$$

Here $\bar{x}_k$ describes the basal expression profile of state k in the modeled expression space, while ρ_k gives its relative abundance within the control population. The hierarchical encoder uses these summaries, while the retained member cells provide observed within-state variation for source sampling.

3.3 Global–Local Contextualization

To combine state expression profiles with population context, we jointly encode the state summaries to obtain population and state representations.

For state k, we embed its mean expression and relative abundance $(\bar{x}_k, \rho_k)$ from Eq. (4) into a state representation e_k, and denote the collection of state representations by $\mathcal{S}_e = \{e_k\}_{k=1}^K$. A learned population token e_{pop} is processed together with $\mathcal{S}_e$ by a set encoder to obtain the population representation and intermediate state representations:

$$\big(h_{\mathrm{pop}}, \widetilde{H}_{\mathrm{loc}}\big) = \psi_{\mathrm{set}}\big(e_{\mathrm{pop}}, \mathcal{S}_e\big), \qquad h_k = f_{\mathrm{loc}}\big([\widetilde{h}_k, h_{\mathrm{pop}}]\big). \tag{5}$$

Here, h_{pop} summarizes the full control population, and $\widetilde{H}_{\mathrm{loc}} = [\widetilde{h}_1, \ldots, \widetilde{h}_K]$ contains the intermediate state representations. Each $\widetilde{h}_k$ is then combined with h_{pop} to obtain h_k, and we denote the resulting state bank by $H_{\mathrm{loc}} = [h_1, \ldots, h_K]$.

The relevance of individual transcriptional states can vary with the requested experimental condition. We therefore use the condition representation to retrieve state information relevant to the perturbation and cellular context. With q_ξ denoting the condition query, we define

$$h_\xi = f_{\mathrm{cond}}\big([h_{\mathrm{pop}}, q_\xi, \mathrm{Attn}(q_\xi, H_{\mathrm{loc}}, H_{\mathrm{loc}})]\big), \tag{6}$$

where f_{cond} combines the retrieved state information with the population representation and condition query to obtain the representation h_ξ for the requested condition.

Together, these representations form the hierarchical memory $\mathcal{H} = \mathcal{E}_\theta(\mathcal{X}, \xi)$, which carries population, state, and condition information into response prediction, source composition, and residual generation. Architectural details are provided in Appendix B.1.

3.4 State-Structured Empirical Source

To construct the source population, CHARM-Flow assigns weights to the transcriptional states according to the requested condition and samples single-cell expression profiles from the empirical distribution within each state. For state k, let $\pi_{\theta,k}(\mathcal{H})$ denote its normalized source weight and $\widehat{q}_k$ the empirical distribution over the retained single-cell expression profiles in that state. The resulting joint law over the sampled state k and single-cell profile s, together with its marginal over s, is

$$q_\theta(k, s \mid \mathcal{X}, \xi) = \pi_{\theta,k}(\mathcal{H})\widehat{q}_k(s), \qquad q_{\mathrm{src}}(s \mid \mathcal{X}, \xi) = \sum_{k=1}^{K} q_\theta(k, s \mid \mathcal{X}, \xi). \tag{7}$$

Sampling N cells together with their state identities from this law defines $Q_{\theta,N}$ in Eq. (1), yielding the sampled single-cell profiles S and corresponding state identities $\boldsymbol{k}$. The sampled single-cell profiles preserve the observed variation within each state, while their state identities provide the corresponding state context during generation.

The sampled source provides the cellular variation used for generation, while the response predictor specifies the population-level expression target $\mu \in \mathbb{R}^G_{\geq 0}$. To express the sampled control-cell variation in coordinates centered around the predicted population response, we map the sampled source profiles to

$$R_0 = \mathcal{T}_\mu(S; \mathcal{X}), \qquad \mathbf{1}^\top R_0 = 0. \tag{8}$$

The map $\mathcal{T}_\mu$ constructs a zero-mean initialization from the sampled source profiles while aligning their expression coordinates with the predicted population response μ.

Starting from R_0, we model the cell-level residual transformation with a vector field conditioned on the hierarchical memory, sampled state identities, and the fixed source reference:

$$\frac{\mathrm{d}R(t)}{\mathrm{d}t} = P_0 V_\theta\left(R(t), R_0, t; \mathcal{H}, \boldsymbol{k}\right), \qquad R(0) = R_0. \tag{9}$$

Here, $P_0 V = V - \mathbf{1}(N^{-1}\mathbf{1}^\top V)$ removes the velocity component shared across cells, preserving the zero-mean residual constraint during integration.

After integration, the perturbed population is reconstructed from the terminal residuals as

$$\widehat{Y} = \mathcal{F}_\theta^{\mathcal{X},\xi}(S, \boldsymbol{k}) = \mathcal{D}_\mu(R(1)), \tag{10}$$

where $\mathcal{D}_\mu$ reconstructs expression profiles from the terminal residuals while preserving the predicted population mean μ.

Together, the predicted population response, residual dynamics, and reconstruction define the generative map $\mathcal{F}_\theta^{\mathcal{X},\xi}$ in Eq. (1). Details of the coordinate construction and numerical integration are provided in Appendix B.

3.5 SOURCE-REFERENCED GENE-MODULE VECTOR FIELD

At the cellular level, perturbation responses can reflect both the basal expression profile of an individual cell and structured relationships among genes. Our vector field retains an individual source reference for each cell while representing gene expression through interactions among gene modules.

For cell i, let $r_i(t)$ denote its evolving residual and $r_{0,i}$ the fixed source residual. We encode the two residuals separately over a fixed gene partition $\{G_m\}_{m=1}^{M}$ derived from control expression covariation, yielding

$$E_i^{(0)} = \psi_{\mathrm{evo}}(r_i(t)), \qquad E_{i,\mathrm{src}}^{(0)} = \psi_{\mathrm{src}}(r_{0,i}). \tag{11}$$

To retain information from the sampled basal profile beyond initialization, we update the evolving representation using cross-attention to the source representation throughout generation:

$$E_i \leftarrow E_i + \mathrm{Attn}\left(E_i, E_{i,\mathrm{src}}, E_{i,\mathrm{src}}\right). \tag{12}$$

Cells in different basal transcriptional states can exhibit distinct responses to the same perturbation. We therefore condition both the evolving and source representations on the sampled state identity and hierarchical memory, situating the individual source reference within its state and population context. The source residual $r_{0,i}$ remains fixed along the generative trajectory, maintaining a stable reference to the sampled basal profile as the evolving representation changes.

After L decoding blocks, the evolving representation produces gene-level velocities through a gene-specific output mapping:

$$v_{\theta,i} = f_{\mathrm{gene}}\left(E_i^{(L)}\right), \qquad V_\theta = [v_{\theta,1}, \ldots, v_{\theta,N}]^\top. \tag{13}$$

The resulting gene-level velocity therefore reflects module-level expression interactions together with the individual basal profile and its state and population context. Architectural details are provided in Appendix B.1.

3.6 LEARNING

Because control and perturbed cells are observed independently, we assign observed perturbed cells to control states during training. We therefore form a training coupling between observed response cells and control states. For each observed response cell, we evaluate its compatibility with the response predicted from each control state under condition ξ. The resulting compatibility scores define soft assignments ω_{ik}, which quantify the training-time association between response cell i and state k.

These assignments guide the selection of source states during training and provide supervision for the predicted source composition. We match the group-averaged assignments to the corresponding source weights through

$$\mathcal{L}_{\mathrm{KL}} = \mathbb{E}_{\mathrm{groups}}\, \mathrm{KL}\left(\mathrm{sg}(\overline{\omega}) \| \overline{\pi}_\theta\right), \tag{14}$$

where $\overline{\omega}$ and $\overline{\pi}_\theta$ denote the soft assignments and predicted source weights averaged within each response group, respectively, and sg denotes stop-gradient. The soft assignments are used only during training; at generation, the source weights are predicted from the control population and requested condition without access to perturbed cells.

To supervise the residual transformation, we express the observed perturbed population in the same centered response coordinates as the sampled source:

$$R_1 = \Pi_\mu(Y) - \mathbf{1}\mu^\top, \qquad \mathbf{1}^\top R_1 = 0. \tag{15}$$

Together, R_0 and R_1 define the source and target endpoints used for flow matching. For each cell i, we independently sample $t_i \in (0, 1)$ and linearly interpolate between $r_{0,i}$ and $r_{1,i}$, yielding $R_{\boldsymbol{t}}$. We then optimize

$$\mathcal{L}_{\mathrm{FM}} = \mathbb{E}\left[\frac{1}{NG}\left\|P_0 V_\theta\left(R_{\boldsymbol{t}}, R_0, \boldsymbol{t}; \mathcal{H}, \boldsymbol{k}\right) - \left(R_1 - R_0\right)\right\|_F^2\right]. \tag{16}$$

Because both endpoints are centered and P_0 removes the component of the predicted velocity shared across the population, $\mathcal{L}_{\mathrm{FM}}$ focuses the regression on residual changes around the predicted population response. Although cellwise times are sampled independently during training, generation integrates all cells at a shared time; the relation between the two settings is discussed in Appendix B.6.

Population response prediction is supervised separately through $\mathcal{L}_{\mathrm{resp}}$, while $\mathcal{L}_{\mathrm{aux}}$ provides additional supervision for state responses and residual population structure. The complete training objective is

$$\mathcal{L} = \mathcal{L}_{\mathrm{FM}} + \lambda_{\mathrm{KL}}\mathcal{L}_{\mathrm{KL}} + \lambda_{\mathrm{resp}}\mathcal{L}_{\mathrm{resp}} + \mathcal{L}_{\mathrm{aux}}. \tag{17}$$

State assignments and residual flow computation are detailed in Appendix B; training objectives and optimization are described in Appendix C.

4 EXPERIMENTS

We evaluate CHARM-FLOW across Replogle, PBMC, and Tahoe, spanning genetic, cytokine, and drug perturbations, and use targeted Replogle analyses to examine state-structured source modeling and persistent individual source references.

4.1 EXPERIMENTAL SETUP

Data and evaluation. We evaluate separate CHARM-FLOW models on Replogle, PBMC, and Tahoe, covering 380 genetic perturbations in HepG2, 62 cytokines across 18 cell types and four evaluation donors, and 607 drug–dose labels across five evaluation cell lines, respectively. Across datasets, evaluation focuses on combinations of interventions and backgrounds held out from training, while the constituent factors may appear elsewhere in training and controls from the target background remain available.

Following the STATE protocol (Adduri et al., 2026), also used by PerturbDiff (Yuan et al., 2026), we report Pearson correlation between predicted and observed perturbation response vectors (PDCorr), differential expression overlap (DEOver), differential expression precision (DEPrec), direction agreement (DirAgr), perturbation discrimination scores under L_1, L_2, and cosine distances (PDS-L1, PDS-L2, and PDS-cos), and mean squared error (MSE) and mean absolute error (MAE) between predicted and observed pseudobulks with Cell-Eval 0.6.6. Targeted Replogle analyses additionally use normalized centered energy distance (NCED), centered multiscale maximum mean discrepancy (CMMD), and normalized centered sliced Wasserstein-1 distance (NCSW1). Baselines include CPA, STATE, CellFlow, Squidiff (He et al., 2026), PerturbDiff, and OCOO-T. Detailed evaluation populations, metric definitions, aggregation procedures, and sources of baseline results are provided in Appendix D.

Training setup. All CHARM-FLOW models and comparison variants use 14,000 optimizer updates with an effective batch size of 512 cells. The CHARM-FLOW results reported in the evaluation tables use training seed 42. Architecture, optimization, and generation details are provided in Appendix C.

4.2 PERTURBATION RESPONSE PREDICTION

Tables 1–3 summarize perturbation response prediction across the three datasets. On Replogle, CHARM-FLOW achieves the highest tabulated PDCorr, DEOver, DEPrec, and DirAgr, together with the lowest MSE and MAE. On PBMC and Tahoe, CHARM-FLOW reaches PDCorr values of 0.8703 and 0.9595, respectively, with high perturbation discrimination. The reported OCOO-T results are higher in DEOver and DEPrec on both datasets, reflecting differences among the evaluation metrics.

On Replogle, the transferred response prior combined with the target-context control mean, it reaches a PDCorr of 0.3956 and a DEPrec of 0.0874, compared with 0.4773 and 0.2430 for CHARM-FLOW (Appendix E.2).

To separate population-mean prediction from cellular variation, we replace each generated population with its corresponding predicted mean. PDCorr and MSE remain unchanged at the reported

Table 1: Replogle perturbation response prediction. Values use three decimal places, except MSE, which uses three significant figures. Bold and underlined entries denote the best and second-best values, respectively.

Method	PDCorr ↑	DEOver ↑	DEPrec ↑	DirAgr ↑	PDS-L1 ↑	PDS-L2 ↑	PDS-cos ↑	MSE ↓	MAE ↓
Mean	0.048	0.127	0.094	0.532	0.582	0.599	0.608	9.90e-2	0.206
Linear	0.058	0.068	0.074	0.535	0.519	0.517	0.667	1.30e-2	0.074
CPA	0.418	0.173	0.087	0.746	0.582	0.581	0.521	7.50e-2	0.054
STATE	0.437	0.196	0.193	0.778	0.788	0.802	0.676	6.40e-3	0.055
CellFlow	-0.003	0.110	0.093	0.472	0.507	0.507	0.495	9.49e-2	0.205
Squidiff	0.089	0.039	0.091	0.432	0.500	0.501	0.501	4.64e+0	2.077
PerturbDiff (Scratch)	0.340	0.190	0.174	0.702	0.690	0.700	0.575	1.47e-2	0.081
PerturbDiff (Finetuned)	0.376	0.214	0.192	0.723	0.758	0.762	0.639	1.46e-2	0.079
OCOO-T	0.437	0.232	0.240	0.766	**0.820**	**0.853**	**0.833**	7.90e-3	0.061
CHARM-FLOW	**0.477**	**0.254**	**0.243**	**0.796**	0.815	0.840	0.781	**5.83e-3**	**0.052**

Table 2: PBMC perturbation response prediction. Bold and underlined entries denote the best and second-best reported values, respectively.

Method	PDCorr ↑	DEOver ↑	DEPrec ↑	DirAgr ↑	PDS-L1 ↑	PDS-L2 ↑	PDS-cos ↑	MSE ↓	MAE ↓
Mean	0.642	0.564	0.544	0.740	0.730	0.777	0.787	8.49e-4	0.013
Linear	0.646	0.549	0.581	0.666	0.658	0.670	0.903	4.44e-4	0.009
CPA	0.181	0.488	0.515	0.539	0.576	0.613	0.597	1.10e-2	0.052
STATE	0.796	0.512	0.547	0.789	0.954	0.943	0.985	**1.41e-4**	0.005
CellFlow	0.628	0.270	0.350	0.652	0.639	0.689	0.725	3.58e-4	0.009
Squidiff	0.033	0.359	0.547	0.379	0.508	0.508	0.508	4.39e0	2.062
PerturbDiff (Scratch)	0.816	0.564	0.581	0.751	0.972	**0.975**	0.986	1.91e-4	0.006
PerturbDiff (Finetuned)	0.705	0.533	0.588	0.701	0.941	0.959	0.972	3.15e-4	0.008
OCOO-T	**0.924**	**0.760**	**0.763**	**0.937**	0.970	0.968	0.972	1.81e-4	0.006
CHARM-FLOW	0.8703	0.5475	0.5686	0.8125	**0.9724**	0.9732	**0.9867**	1.62e-4	**0.004877**

Table 3: Tahoe perturbation response prediction. Bold and underlined entries denote the best and second-best reported values, respectively.

Method	PDCorr ↑	DEOver ↑	DEPrec ↑	DirAgr ↑	PDS-L1 ↑	PDS-L2 ↑	PDS-cos ↑	MSE ↓	MAE ↓
Mean	0.205	0.430	0.504	0.595	0.515	0.508	0.508	3.09e-3	0.026
Linear	0.723	0.505	0.576	0.760	0.887	0.881	0.912	4.15e-4	0.009
CPA	0.425	0.502	0.504	0.710	0.654	0.593	0.585	8.89e-4	0.011
STATE	0.535	0.499	0.522	0.665	0.789	0.846	0.854	2.11e-3	0.021
CellFlow	0.269	0.183	0.313	0.638	0.608	0.564	0.550	3.10e-3	0.027
Squidiff	0.011	0.420	0.581	0.501	0.505	0.506	0.505	5.63e0	2.244
PerturbDiff (Scratch)	0.686	0.522	0.572	0.734	0.970	0.978	0.975	6.30e-4	0.012
PerturbDiff (Finetuned)	0.648	0.580	**0.598**	0.641	0.784	0.883	0.903	1.04e-3	0.017
OCOO-T	0.952	**0.634**	0.545	**0.959**	0.991	**0.993**	**0.994**	**1.11e-4**	**7.27e-3**
CHARM-FLOW	**0.9595**	0.4878	0.3261	0.9513	**0.9916**	0.9925	0.9916	7.12e-4	0.011424

precision, while DEOver decreases from 0.2542 to 0.1988 and DEPrec from 0.2430 to 0.1004 (Table 15). Thus, Cellular variation therefore contributes to differential expression prediction.

4.3 State Structure and Cellular Variation

Transcriptional state informs both empirical source composition and residual generation. We compare CHARM-FLOW with a marginal empirical source that omits intermediate state organization while retaining variation among observed control cells and persistent source references.

Compared with the marginal empirical source, state organization reduces NCED, CMMD, and NCSW1 by 30.49%, 24.12%, and 23.55%, respectively, indicating closer agreement between the generated and observed centered distributions.

Removing direct state context during residual generation leaves PDCorr and MSE nearly unchanged, while DEOver decreases from 0.2542 to 0.2372 and DEPrec from 0.2430 to 0.2356 (Table 12). Together, these results show that transcriptional state contributes in two complementary ways: state-

Table 4: Centered distribution comparison on Replogle, averaged over 380 perturbations. NCED, CMMD, and NCSW1 are computed from the control-derived representation; lower is better.

Source organization	NCED ↓	CMMD ↓	NCSW1 ↓
Marginal empirical source	0.0323	0.1428	0.2175
CHARM-FLOW	0.0224	0.1083	0.1663

structured source composition improves within-perturbation distribution fidelity, while state context during residual generation supports differential expression prediction.

4.4 PERSISTENT SOURCE REFERENCES DURING RESIDUAL GENERATION

In CHARM-FLOW, each sampled control cell remains available throughout residual generation as an individual reference to its basal expression profile. To distinguish this persistent reference from its role in initialization, we compare CHARM-FLOW with an initialization only variant that uses the same sampled cell to define the initial residual but does not retain its source representation as a separate reference during residual generation.

Table 5: Replogle response prediction with persistent source references or source use limited to initialization.

Source access	PDCorr ↑	DEOver ↑	DEPrec ↑	MSE ↓
Initialization only	0.4651	0.1565	0.1954	1.3434×10^{-2}
Persistent (CHARM-FLOW)	0.4773	0.2542	0.2430	5.8307×10^{-3}

Persistent source references yield higher PDCorr, DEOver, and DEPrec and lower MSE (Table 5). They also reduce NCED, NCGW1, and NStdL1 by 23.60%, 28.69%, and 45.62%, respectively (Table 17). Together, these comparisons show that retaining the sampled basal profile benefits both perturbation response prediction and the generated population distribution.

To test whether the trained generator directly uses this retained reference, we perform a paired inference intervention. The sampled basal profile still initializes the residual but is no longer supplied as a separate reference during subsequent generation, and the model parameters, sampled source cells, population mean computation, and terminal reconstruction remain unchanged. PDCorr and MSE remain unchanged, whereas NCED increases from 0.02239 to 0.30163 and DEPrec decreases from 0.24286 to 0.18322 (Table 18). The generator therefore uses the retained source representation during residual generation.

Additional distribution metrics and comparisons across individual perturbations are provided in Appendix E.

5 CONCLUSION

CHARM-FLOW models single-cell perturbation prediction responses through conditional population generation using the observed control population. A shared population–state–cell hierarchy coordinates population response prediction, state-structured empirical source construction, and residual generation that retains each sampled control cell as an individual reference. Across genetic, cytokine, and drug perturbations, CHARM-FLOW predicts population responses, while targeted Replogle comparisons show lower centered distribution errors when state organization and persistent source references are retained.

AI USE STATEMENT

AI-assisted tools were used to retrieve and summarize related work, survey the single-cell perturbation literature, and compare methodological practices reported by other studies. Through explicit instructions, they were also used to implement model modules, debug the code, and inspect experimental procedures. For the manuscript, AI-assisted tools were used for translation, language polishing, and checks for logical gaps and inconsistencies in the paper's argument.

Code availability. Code for reproducing the experiments is available at `https://github.com/paper-artifacts0925/CHARM-Flow`. Implementation details, training settings, and evaluation procedures are provided in the appendices.

REFERENCES

Abhinav K. Adduri, Dhruv Gautam, Beatrice Bevilacqua, Mohsen Naghipourfar, Alishba Imran, Rohan Shah, Noam Teyssier, Rishi Verma, Christopher Carpenter, Basak Eraslan, Francis Chalissery, Rajesh Ilango, Vishvak Subramanyam, Chiara Ricci-Tam, Sanjay Nagaraj, Aidan Winters, Mingze Dong, Stefanie Fellinger, Adam Krejci, Tilmann Burckstummer, Sravya Tirukkovular, Jeremy Sullivan, Brian S. Plosky, Nicholas D. Youngblut, Jure Leskovec, Luke A. Gilbert, Silvana Konermann, Patrick D. Hsu, Alexander Dobin, Dave P. Burke, Hani Goodarzi, and Yusuf H. Roohani. Predicting cellular responses to perturbation across diverse contexts with State. *Cell*, 2026. doi: 10.1016/j.cell.2026.07.052. URL https://doi.org/10.1016/j.cell.2026.07.052. Online publication; in press.

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.

Shuizhou Chen, Lang Yu, Xueqin Lin, Xinjie Mao, Songming Zhang, Xinyu Gu, Hao Wu, Sheng Xu, Kedu Jin, Lei Bai, Quan Qian, Qin Chen, Qiang Gao, Siqi Sun, and Zhangyang Gao. SCALE: Scalable Conditional Atlas-Level Endpoint Transport for Virtual Cell Perturbation Prediction, 2026. URL https://arxiv.org/abs/2603.17380.

Atray Dixit, Oren Parnas, Biyu Li, Jenny Chen, Charles P. Fulco, Livnat Jerby-Arnon, Nemanja Marjanovic, Danielle Dionne, Tyler Burks, Raktima Raychowdhury, Britt Adamson, Thomas M. Norman, Eric S. Lander, Jonathan S. Weissman, Nir Friedman, and Aviv Regev. Perturb-Seq: Dissecting Molecular Circuits with Scalable Single-Cell RNA Profiling of Pooled Genetic Screens. *Cell*, 167(7):1853–1866.e17, 2016. doi: 10.1016/j.cell.2016.11.038. URL https://doi.org/10.1016/j.cell.2016.11.038.

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

Siyu He, Yuefei Zhu, Daniel Naveed Tavakol, Haotian Ye, Yeh-Hsing Lao, Zixian Zhu, Cong Xu, Shradha Chauhan, Guy Garty, Raju Tomer, Gordana Vunjak-Novakovic, James Zou, Elham Azizi, and Kam W. Leong. Squidiff: Predicting Cellular Development and Responses to Perturbations Using a Diffusion Model. *Nature Methods*, 23(1):65–77, 2026. doi: 10.1038/s41592-025-02877-y. URL https://doi.org/10.1038/s41592-025-02877-y.

Danning Jiang, Zhiwen Yan, Qirun Wang, Zheming An, Yalong Zhao, and Lipeng Lai. OCOO-T : A Simple and Scalable Virtual Cell Model for Transcriptional Perturbation Response Prediction, 2026. URL https://arxiv.org/abs/2606.12838.

Junwan Kim, Jiho Park, Seonghu Jeon, and Seungryong Kim. Better Source, Better Flow: Learning Condition-Dependent Source Distribution for Flow Matching, 2026. URL https://arxiv.org/abs/2602.05951.

Dominik Klein, Jonas Simon Fleck, Daniil Bobrovskiy, Lea Zimmermann, Sören Becker, Alessandro Palma, Leander Dony, Alejandro Tejada-Lapuerta, Guillaume Huguet, Hsiu-Chuan Lin, Nadezhda Azbukina, Fátima Sanchís-Calleja, Theo Uscidda, Artur Szalata, Manuel Gander, Aviv Regev, Barbara Treutlein, J. Gray Camp, and Fabian J. Theis. CellFlow enables generative single-cell phenotype modeling with flow matching. *bioRxiv*, 2025. doi: 10.1101/2025.04.11.648220. URL https://doi.org/10.1101/2025.04.11.648220.

Quanquan Li, Yihe Chi, Liuyang Song, Hongbo Zhang, Jingyu Li, Xidong Xi, Conghua Wei, Yijie Sun, Yu Chen, Xin Liu, Qi Hu, Jing Ke, and Guitao Cao. Control-Anchored Residual Flow Matching Conditioned on Gene Geometry for Virtual Cell Perturbation Modeling, 2026. URL https://arxiv.org/abs/2608.06824.

Jiawen Liu, Xuechenxiao Cao, Yutong Li, Bing Liu, Jiaming Liang, Tinghe Zhang, Xiaoqi Sheng, and Hongmin Cai. DeMixPert: Decomposed Response Modeling with Gaussian Mixtures for OOD Single-Cell Perturbation Prediction, 2026. URL https://arxiv.org/abs/2608.23114.

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):e11517, 2023. doi: 10.15252/msb.202211517. URL https://doi.org/10.15252/msb.202211517.

Joseph M. Replogle, Reuben A. Saunders, Angela N. Pogson, Jeffrey A. Hussmann, Alexander Lenail, Alina Guna, Lauren Mascibroda, Eric J. Wagner, Karen Adelman, Gila Lithwick-Yanai, Nika Iremadze, Florian Oberstrass, Doron Lipson, Jessica L. Bonnar, Marco Jost, Thomas M. Norman, and Jonathan S. Weissman. Mapping information-rich genotype-phenotype landscapes with genome-scale Perturb-seq. *Cell*, 185(14):2559–2575.e28, 2022. doi: 10.1016/j.cell.2022.05.013. URL https://doi.org/10.1016/j.cell.2022.05.013.

Yusuf Roohani, Kexin Huang, and Jure Leskovec. Predicting transcriptional outcomes of novel multigene perturbations with GEARS. *Nature Biotechnology*, 42(6):927–935, 2024. doi: 10.1038/s41587-023-01905-6. URL https://doi.org/10.1038/s41587-023-01905-6.

Jiafa Ruan, Ruijie Quan, Liyang Xu, Zongxin Yang, and Yi Yang. Beyond Independent Genes: Learning Module-Inductive Representations for Single-Cell Gene Perturbation Prediction. In *Proceedings of the 43rd International Conference on Machine Learning*, volume 306 of *Proceedings of Machine Learning Research*, 2026.

Andrea Rubbi, Amir Akbarnejad, Mohammad Vali Sanian, Aryan Yazdan Parast, Hesam Asadollahzadeh, Arian Amani, Naveed Akhtar, Sarah Cooper, Andrew Bassett, Pietro Liò, Lassi Paavolainen, Sattar Vakili, and Mo Lotfollahi. Shortest-Path Flow Matching with Mixture-Conditioned Bases for OOD Generalization to Unseen Conditions, 2026. URL https://arxiv.org/abs/2601.11827.

Sydney M. Shaffer, Margaret C. Dunagin, Stefan R. Torborg, Eduardo A. Torre, Benjamin Emert, Clemens Krepler, Marilda Beqiri, Katrin Sproesser, Patricia A. Brafford, Min Xiao, Elliott Eggan, Ioannis N. Anastopoulos, Cesar A. Vargas-Garcia, Abhyudai Singh, Katherine L. Nathanson, Meenhard Herlyn, and Arjun Raj. Rare cell variability and drug-induced reprogramming as a mode of cancer drug resistance. *Nature*, 546(7658):431–435, 2017. doi: 10.1038/nature22794. URL https://doi.org/10.1038/nature22794.

Sanjay R. Srivatsan, José L. McFaline-Figueroa, Vijay Ramani, Laura Saunders, Jian Cao, Jonathan Packer, Hannah A. Pliner, David L. Jackson, Riza M. Daza, Lisa Christiansen, Robert Stewart, and Jay Shendure. Massively multiplex chemical transcriptomics at single-cell resolution. *Science*, 367(6473):45–51, 2020. doi: 10.1126/science.aax6234. URL https://doi.org/10.1126/science.aax6234.

Yang Xu, Stephen Fleming, Matthew Tegtmeyer, Steven A. McCarroll, and Mehrtash Babadi. Explainable modeling of single-cell perturbation data using attention and sparse dictionary learning. *Cell Systems*, 16(4):101245, 2025. doi: 10.1016/j.cels.2025.101245. URL https://doi.org/10.1016/j.cels.2025.101245.

Chenglei Yu, Chuanrui Wang, Bangyan Liao, and Tailin Wu. scDFM: Distributional Flow Matching Model for Robust Single-Cell Perturbation Prediction. In *International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=QSGanMEcUV.

Xinyu Yuan, Xixian Liu, Yashi Zhang, Zuobai Zhang, Hongyu Guo, and Jian Tang. PerturbDiff: Functional Diffusion for Single-Cell Perturbation Modeling, 2026. URL https://arxiv.org/abs/2602.19685.

A Data Preparation and Precomputed Model Inputs

A.1 Expression Coordinates and Dataset Splits

Each dataset uses its supplied prepared expression matrix and a fixed ordered set of $G = 2{,}000$ genes (`X_hvg`). During training, prepared expression values are rescaled as $x = x^{\text{prepared}}/10$, and predictions are multiplied by ten before evaluation. Gene identities and ordering are preserved within each dataset.

The evaluation splits hold out specified intervention–background combinations. In Replogle, validation and test cells correspond to the designated gene perturbations in HepG2. PBMC applies the designated cytokine lists within Donor1, Donor4, Donor9, and Donor12 while retaining cell type as a modeled context. Tahoe applies the designated drug–dose lists within five evaluation cell lines: `CVCL_1097`, `CVCL_1285`, `CVCL_1098`, `CVCL_0334`, and `CVCL_0480`. Thus, evaluation targets held-out intervention–background combinations while allowing the constituent factors to appear elsewhere in training. The complete intervention lists are retained in the dataset configurations.

Available controls, including those from the target backgrounds, define the control-derived representations used by the model. Response priors and response supervision use only treated cells from the training split. Table 6 summarizes the prepared datasets and precomputed model inputs.

Table 6: Prepared-data split counts and precomputed model inputs for the three CHARM-Flow models. Evaluation coverage and reference control populations are reported in Appendix D.

Data and precomputed inputs	Replogle	PBMC	Tahoe
Training treated rows	572,545	6,657,336	89,495,239
Validation treated rows	4,825	150,484	553,076
Test treated rows	26,878	2,260,453	8,270,319
Available control rows	39,165	629,701	2,330,156
Stored members per state (max.)	64	8	64
Gene modules	128	128	128

A.2 Dataset-Specific Conditions and Encoding

The experimental condition is written as $\xi = (p, c, b)$, comprising the intervention, biological context, and batch or donor covariate. These correspond to perturbed gene, cell line, and gem group for Replogle; cytokine, cell type, and donor for PBMC; and joint drug–dose label, cell line, and plate for Tahoe.

The condition encoder uses three 128-dimensional branches. For Replogle, the intervention branch projects the fixed 3,072-dimensional GenePT feature, whereas PBMC cytokines and Tahoe drug–dose labels use learned categorical embeddings. The context branch projects a fixed 3,072-dimensional embedding, and the batch branch uses a learned categorical embedding. The three branch representations are concatenated and projected to the shared condition representation

$$u_\xi = W_\xi \big[e_p(p) \parallel e_c(c) \parallel e_b(b) \big] + b_\xi \in \mathbb{R}^{2000}. \tag{18}$$

The hierarchy encoder maps u_ξ to the retrieval query q_ξ in Eq. (6), while the vector field applies separate projections to the same condition representation.

A.3 Control Contexts and State Reservoirs

The cellular hierarchy is constructed before response model training from non-targeting controls for Replogle, PBS-treated controls for PBMC, and DMSO controls with zero drug dose for Tahoe. Each control context defines one population $\mathcal{X}$. Replogle and Tahoe use one control bank per cell line, pooling across gem groups and plates, respectively, whereas PBMC uses one bank for each populated donor–cell-type pair, yielding 215 banks. In PBMC, donor therefore determines the control bank and also enters condition encoding; gem group and plate remain condition covariates in Replogle and Tahoe, respectively, but do not affect control bank selection.

Control states are constructed independently within each control bank. Control profiles are centered by their batch mean and shifted to the context-wide control mean for clustering, while reservoir members retain their prepared expression values. A randomized PCA with at most 64 components produces z_i, followed by constrained bisecting k-means with a minimum terminal-state size of 20. Leaves containing at least 40 cells are split in decreasing order of

$$\text{priority}(\mathcal{L}) = \left(\sum_{i \in \mathcal{L}} \|z_i - \bar{z}_{\mathcal{L}}\|_2^2 \right) \left[1 + \log \left(1 + \frac{|\mathcal{L}|}{20} \right) \right]. \tag{19}$$

Each binary split uses five initializations and at most 100 iterations, enforces 20-member children, and handles invalid or degenerate proposals deterministically. Recursion stops when no legal split remains or the capacity of 128 states is reached. PBMC adjusts this capacity using a target occupancy of 64 cells, and contexts with fewer than 20 controls remain single states. Construction uses seed 42 with deterministic context- and split-specific offsets.

For each terminal state $\mathcal{X}_k$, the full membership defines the centroid $\bar{x}_k$ and occupancy ρ_k used by the hierarchy. A separate empirical reservoir $\mathcal{R}_k = \{c_{kj}\}_{j=1}^{n_k}$ retains $n_k = \min(|\mathcal{X}_k|, R_{\max})$ observed controls, with $R_{\max} = 64$ for Replogle and Tahoe and $R_{\max} = 8$ for PBMC. Oversized states are subsampled uniformly without replacement using deterministic state-specific seeds. Generation draws states and reservoir members with replacement from this fixed support.

A.4 Control-Derived Gene Partition for Vector Field Computation

The gene partition is constructed from patterns of covariation among genes in control cells. Genes with similar control expression variation are grouped, and the resulting partition organizes both the source and evolving residual representations while preserving gene-specific velocity outputs. One partition is constructed for each dataset before response model training and remains fixed throughout training and generation.

The partition is constructed from 8,192 available control cells sampled uniformly without replacement using seed 42. Each sampled expression profile is centered by the control mean of its cell line, or by the corresponding donor–cell-type control mean for PBMC, so that the resulting variation reflects differences within a control context rather than differences between context means. For each gene, the centered values are divided by their standard deviation across the sampled cells; genes with standard deviation at most 10^{-8} are assigned zero standardized values. The standardized values are clipped to $[-10, 10]$ and collected in the control-by-gene matrix Z.

A randomized singular value decomposition of rank 64 with five power iterations gives $Z \approx U\Sigma V^\top$. Each gene is represented by the corresponding row of $V\Sigma$, normalized to unit length. The resulting gene features are partitioned into $M = 128$ groups using k-means with 20 initializations, at most 500 iterations, and seed 42. Group identifiers follow the order of their first member in the fixed gene list, and group sizes may differ.

Let $\{\mathcal{G}_m\}_{m=1}^M$ denote the resulting gene groups. At each vector field evaluation, the evolving residual and fixed source residual are aggregated over the same groups to form corresponding gene group tokens. The two tokenizers use separate learned parameters but share the same partition, so corresponding source and evolving tokens represent the same subsets of genes. This alignment supports cross attention from the source representation to the evolving representation. After group level decoding, gene-specific queries and output parameters map the updated evolving tokens to the full gene-level velocity $v_\theta \in \mathbb{R}^G$, as detailed in Appendix B.1.

A.5 Training Response Priors

The population response predictor combines the control mean of the target context, a response prior estimated from other training contexts, and a learned correction in Eq. (29). For intervention p, the prior uses responses observed in other training contexts to provide an estimate for the target context c, using only treated cells from the training split. Here c denotes a cell line for Replogle and Tahoe or a donor–cell-type pair for PBMC. On the model scale, the response to intervention p in training context c is

$$\delta^{\mathrm{tr}}_{c,p} = \frac{1}{|\mathcal{Y}^{\mathrm{tr}}_{c,p}|} \sum_{y \in \mathcal{Y}^{\mathrm{tr}}_{c,p}} y - \mu_c, \qquad \mu_c = \frac{1}{|\mathcal{X}_c|} \sum_{x \in \mathcal{X}_c} x. \tag{20}$$

For a target context c, we average the responses to intervention p over the other training contexts in which that intervention is observed:

$$\widetilde{\delta}_{c,p} = \begin{cases} |\mathcal{S}_{c,p}|^{-1} \displaystyle\sum_{c' \in \mathcal{S}_{c,p}} \delta^{\mathrm{tr}}_{c',p}, & |\mathcal{S}_{c,p}| > 0, \\ 0, & |\mathcal{S}_{c,p}| = 0. \end{cases} \tag{21}$$

Here, $\mathcal{S}_{c,p}$ denotes the set of supporting training contexts, and all such contexts receive equal weight. For PBMC, this set contains the other donors of the same cell type for which intervention p is available after applying the donor–cytokine split.

We calibrate the averaged response estimate to the target context independently for each gene. In Eq. (22), both sums run over $p \in \mathcal{P}^{\mathrm{tr}}_c$, while $[\cdot]_{[0,2]}$ and $\vee$ denote coordinatewise clipping and maximum, respectively:

$$\gamma_c = \left[\frac{\sum_p \widetilde{\delta}_{c,p} \odot \delta^{\mathrm{tr}}_{c,p}}{\left(\sum_p \widetilde{\delta}^{\odot 2}_{c,p} + \lambda_c \right) \vee 10^{-12} \mathbf{1}} \right]_{[0,2]}, \qquad \Delta_{\mathrm{prior}}(c,p) = \gamma_c \odot \widetilde{\delta}_{c,p}. \tag{22}$$

For each gene g, the penalty $\lambda_{c,g}$ is selected from $\{10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}, 1\}$ by up to five-fold cross-validation over the training interventions using MSE. The calibration is then refitted using all eligible interventions. If only one intervention is available, $\lambda = 0.01$ is used. If no calibration support is available for context c, we set $\gamma_c = \mathbf{1}$; if intervention p has no support from other contexts, we set $\Delta_{\mathrm{prior}}(c,p) = 0$.

B Model Construction and Computation

All computations below use the expression coordinates defined in Appendix A. A response group contains N cells that share the same intervention and biological context, while batch covariates may vary within the group.

B.1 Hierarchical Conditioning and Gene Module Residual Updates

The vector field combines population and state context with separate gene module representations for the evolving and source residuals. For state k, the shared expression encoder maps its centroid $\bar{x}_k$, computed from all assigned control cells, to $\phi(\bar{x}_k)$. This encoding and its numerical scale feature are jointly projected and combined with a projection of $\log \rho_k$ to form the state token e_k. The active state tokens and a learned population token enter the set encoder in Eq. (5). The condition query then retrieves h_{ξ_i} from the resulting state representations through Eq. (6).

Each vector field block accesses the hierarchical context at two levels. A compact memory contains the population representation, the sampled state representation, the retrieved condition representation, and the condition embedding. A full memory contains all active state representations together with the population representation, condition query, retrieved condition representation, and condition embedding. The compact memory provides local context for the sampled state, while the full memory retains the surrounding population context.

Let $\{\mathcal{G}_m\}_{m=1}^{M}$ denote the fixed gene groups from Appendix A.4. For a residual vector r, both the evolving and source tokenizers use the same gene groups. For either tokenizer, group m is represented as

$$\psi_m(r) = \mathrm{LN}\left(e_m + \frac{1}{\sqrt{|\mathcal{G}_m|}} \sum_{g \in \mathcal{G}_m} r_g e_g\right) + f_{\mathrm{stat}}([\mathrm{mean}_{g\in\mathcal{G}_m}\, r_g, \mathrm{RMS}_{g\in\mathcal{G}_m}\, r_g])\,. \tag{23}$$

Here,

$$\mathrm{RMS}_{g\in\mathcal{G}_m}\, r_g = \sqrt{|\mathcal{G}_m|^{-1} \sum_{g\in\mathcal{G}_m} r_g^2 + 10^{-8}}.$$

The two tokenizers have separate learned gene, group, and projection parameters, while aggregating over the same gene groups.

Twelve decoding blocks update the evolving and source representations through self attention, cross attention from the source representation to the evolving representation, attention to the hierarchical context, and feed forward transformations. Time and condition representations modulate the adaptive normalization layers, with time represented as $1000t$. If $m(g)$ denotes the group containing gene g, the final evolving token produces

$$v_{\theta,ig} = \frac{1}{\sqrt{d}} w_g^\top \left[\mathrm{LN}\left(E^{(L)}_{i,m(g)} + q_g\right) + E^{(L)}_{i,m(g)}\right] + b_g, \qquad d = 384. \tag{24}$$

Gene-specific queries and output parameters distinguish individual genes within the same group. The evolving representation produces the velocity used for integration, while the source residual $r_{0,i}$ remains fixed along the trajectory and is retokenized at each vector field evaluation.

B.2 State-Structured Empirical Source Sampling

The empirical source is sampled in two stages: a condition-dependent state is selected, followed by a uniform draw from that state's retained controls. For active state k with reservoir $\mathcal{R}_k = \{c_{kj}\}_{j=1}^{n_k}$ and pre-truncation occupancy ρ_k, the scoring network produces $\ell_{ik} = f_\pi([h_k, h_{\mathrm{pop}}, q_{\xi_i}])$ and

$$\pi_{ik} = \frac{\mathbb{I}[k \in \mathcal{K}_{\mathrm{valid}}] \exp\big(\ell_{ik}/\tau + \lambda_{\mathrm{occ}} \log \max(\rho_k, 10^{-8})\big)}{\sum_{k' \in \mathcal{K}_{\mathrm{valid}}} \exp\big(\ell_{ik'}/\tau + \lambda_{\mathrm{occ}} \log \max(\rho_{k'}, 10^{-8})\big)}, \tag{25}$$

where $\tau = 1$ and $\lambda_{\mathrm{occ}} = 0.1$; inactive states are masked. Conditional on the selected state, member sampling gives

$$\Pr(k_i = k, j_i = j \mid \mathcal{X}, \xi_i) = \frac{\pi_{ik}}{n_k}, \qquad j \in \{1, \ldots, n_k\}, \qquad s_i = c_{k_i j_i}. \tag{26}$$

Draws are independent with replacement across generated cells. The sampled index k_i conditions the vector field, while s_i supplies the empirical expression used to construct the source coordinates.

B.3 Response Guided Assignments for Source Learning

The predicted weights π_i determine state selection during generation. During training, each observed response cell is used to construct a detached soft assignment ω_i over the control states. Let $\bar{\phi}_k = \phi(\bar{x}_k)$ denote the encoded centroid of state k, $d_k(\xi_i)$ its predicted latent response under condition ξ_i, and $o_{ik} = \phi(y_i) - \bar{\phi}_k$ the observed displacement relative to that state. Their compatibility reflects the agreement between the predicted response and observed displacement in overall deviation, direction, and magnitude:

$$\begin{aligned} C_{ik} = \frac{1}{d_\phi}\|o_{ik} - d_k(\xi_i)\|_2^2 + 0.25\,[1 - \cos(o_{ik}, d_k(\xi_i))] \\ + 0.05\;\left|\log \mathrm{RMS}_\delta(o_{ik}) - \log \mathrm{RMS}_\delta(d_k(\xi_i))\right|, \end{aligned} \tag{27}$$

where $\mathrm{RMS}_\delta(v) = \sqrt{d_\phi^{-1}\|v\|_2^2 + \delta}$ and $\delta = 10^{-6}$. After robust normalization of the compatibility costs within each response group, we obtain the soft assignment by solving

$$\begin{aligned} P^* &= \underset{\substack{P\geq 0\,;\; P\mathbf{1}=\mathbf{1}/N \\ P_{ik}=0 \text{ for inactive } k}}{\arg\min} \left[\langle P, \widetilde{C}\rangle + \epsilon \sum_{ik} P_{ik}(\log P_{ik} - 1) + \lambda_{\mathrm{mass}}\,\mathrm{KL}(P^\top \mathbf{1}\|\rho)\right], \\ \omega_{ik} &= N P^*_{ik}. \end{aligned} \tag{28}$$

We use $\epsilon = 0.1$ and set $\lambda_{\mathrm{mass}} = 0$ for Replogle and $\lambda_{\mathrm{mass}} = 1$ for PBMC and Tahoe. The mass term encourages the aggregate assignment across cells to remain close to the control state abundances ρ. The resulting assignment ω is detached from backpropagation. During training, source states are sampled from $\eta(s)\omega_i + [1 - \eta(s)]\pi_i$. The coefficient η remains equal to one through update 2,500 and then decreases to zero with a cosine schedule by update 6,000. A source KL loss trains the predicted weights π_i to match the guided assignments at the response group level. This schedule initially uses response guided state assignments for source selection and gradually transitions to the predicted source distribution used during generation, which does not require observed treated responses.

B.4 Response Coordinates and Mean Constrained Reconstruction

The population response branch combines the control mean μ_c, the response prior from Appendix A.5, and a bounded learned correction:

$$\begin{aligned} \widehat{\mu}_i^{(0)} &= \mu_c + \Delta_{\mathrm{prior}}(c, p) \\ &\quad + \alpha \tanh f_\mu([h_{\mathrm{pop}}, q_{\xi_i}, h_{\xi_i}, f_{\mathrm{prior}}(\Delta_{\mathrm{prior}}(c, p))]), \qquad \alpha = 0.05. \end{aligned} \tag{29}$$

Appendix C specifies the additional response adjustment used for Replogle and Tahoe and the donor conditioned correction used for PBMC. Let $\widehat{\mu}_i$ denote the resulting response prediction and $\varepsilon_i^{\mathrm{ctrl}}$ the sampled control cell's residual from its stored control population mean. We construct

$$s_i^{\mathrm{shift}} = \mathrm{sg}(\widehat{\mu}_i) + \varepsilon_i^{\mathrm{ctrl}}, \qquad \mu = \left[\frac{1}{N}\sum_i \widehat{\mu}_i\right]_+. \tag{30}$$

The sampled control residual is added without additional rescaling, preserving the observed cell-level variation carried by the source cell. When constructing the residual coordinates, the predicted group mean is detached from this path. Its learning is therefore governed separately by the population response losses.

To keep the generated population centered around the predicted response mean while preserving nonnegative expression, we constrain the reconstructed population to have group mean μ. Define

$$\mathcal{A}_\mu = \left\{U \in \mathbb{R}_{\geq 0}^{N\times G} : N^{-1}\mathbf{1}^\top U = \mu^\top\right\}, \qquad \Pi_\mu(Z) = \underset{U\in\mathcal{A}_\mu}{\arg\min}\,\|U - Z\|_F^2. \tag{31}$$

The projection Π_μ is then used to express both the sampled source and the observed response as residuals around the same population mean:

$$\begin{aligned} \mathcal{T}_\mu(S;\mathcal{X}) &= \Pi_\mu(S^{\text{shift}}) - \mathbf{1}\mu^\top, \\ R_0 &= \mathcal{T}_\mu(S;\mathcal{X}), \\ R_1 &= \Pi_\mu(Y) - \mathbf{1}\mu^\top. \end{aligned} \tag{32}$$

The same projection is used at reconstruction:

$$\mathcal{D}_\mu(R) = \Pi_\mu(\mathbf{1}\mu^\top + R). \tag{33}$$

By construction, both R_0 and R_1 have zero group mean. The residual dynamics therefore model cellular variation around the prescribed population response, while reconstruction restores the shared mean through $\mathcal{D}_\mu$. Up to numerical precision, the reconstruction satisfies

$$N^{-1}\mathbf{1}^\top\mathcal{D}_\mu(R) = \mu^\top, \qquad \mathcal{D}_\mu(R) \geq 0. \tag{34}$$

The projection is applied independently to each gene. The additional support constraints used during generation are described in Appendix C.4. For a singleton group, centering leaves no cellular variation, so $R_0 = R_1 = 0$ and reconstruction returns μ. For groupwise centering, we use $\mathsf{C}_0 = P_0$ when $N \geq 2$ and the identity map when $N = 1$.

B.5 Flow Training and Numerical Integration

Training uses the grouped residual coordinates defined above. For each response cell, we sample $t_i = \sigma(a_i)$ with $a_i \sim \mathcal{N}(-0.8, 0.8^2)$ and interpolate $R_{\boldsymbol{t},i} = (1 - t_i)R_{0,i} + t_i R_{1,i}$. The vector field is trained to predict $R_1 - R_0$ after centering its output within each response group. Endpoint construction, source learning, and the accompanying supervision follow Appendices B.3–B.4 and C. Group operations use the valid cells in each device-local microbatch; distributed training and gradient accumulation only aggregate their gradients.

Algorithm 1 Training CHARM-Flow

Require: Training responses, available control representations, fixed assets, parameters θ.
for each optimizer update s **do**
  Form local microbatches with response-group identifiers and valid-row masks.
  for each microbatch **do**
    Encode the required control banks and conditions into hierarchical memory $\mathcal{H}$.
    Predict state weights π, state responses, and population responses $\widehat{\mu}$.
    Compute detached assignments ω using Eq. (28).
    Draw source states from $\eta(s)\omega + [1 - \eta(s)]\pi$ and uniform members of their reservoirs.
    Construct shifted sources, detached group means, and R_0, R_1 using Eq. (32).
  Apply observed-support target projection when enabled.
    Sample independent logit-normal times t_i; form $r_{\boldsymbol{t},i} = (1 - t_i)r_{0,i} + t_i r_{1,i}$.
    Predict $V_\theta(R_{\boldsymbol{t}}, R_0, \boldsymbol{t}; \mathcal{H}, \boldsymbol{k})$ and apply C_0 to velocities by group.
    Evaluate flow, source-composition, population-response, and enabled auxiliary losses.
    Accumulate gradients with the reductions and coefficients in Appendix C.
  end for
  Clip gradients, update parameters, and update the exponential moving average.
end for

Starting from R_0, generation integrates the same centered field on a uniform grid. Each interval uses a Heun predictor–corrector update, except the final interval, which takes the Euler proposal directly. The final residual is reconstructed with $\mathcal{D}_\mu$ and, when enabled, the terminal support constraint. Generation groups and random seeds are listed in Appendix C.

B.6 Cellwise Time Sampling and Population Integration

The velocity network applies gene-module attention independently to individual cells. Its shared hierarchical memory is fixed during each velocity evaluation. Conditional on this memory, the sampled source, and the state identities, each raw velocity depends on the current coordinates and time of its own cell. Groupwise projection then couples the raw velocities through their mean.

Fix a response group with $N \geq 2$ valid cells, model parameters, and centered endpoints R_0, R_1. Write $\Delta_i = r_{1,i} - r_{0,i}$ and collect the time-independent conditioning of cell i in c_i. Along its interpolation path, define

$$\begin{aligned} r_i(t) &= (1-t)r_{0,i} + tr_{1,i}, \\ f_i(t) &= v_\theta(r_i(t), r_{0,i}, t; c_i), \qquad e_i(t) = f_i(t) - \Delta_i. \end{aligned} \tag{35}$$

Independent times query these cellwise paths at separate locations. At a common time, the interpolated population satisfies $\mathbf{1}^\top R(t) = 0$. Generation evaluates the same cellwise field at the common integration time and applies P_0, preserving this constraint throughout the numerical updates.

Projected regression bound. Let ν be the logit-normal time distribution used in training, and assume the errors in Eq. (35) have finite second moments. Define the independent-time and common-time regression risks on the same endpoint group by

$$\begin{aligned} L_{\mathrm{ind}} &= \frac{1}{NG}\mathbb{E}_{\boldsymbol{t}\sim\nu^{\otimes N}} \left\| P_0[e_i(t_i)]_{i=1}^N \right\|_F^2, \\ L_{\mathrm{sync}} &= \frac{1}{NG}\mathbb{E}_{t\sim\nu} \left\| P_0[e_i(t)]_{i=1}^N \right\|_F^2. \end{aligned} \tag{36}$$

Since $P_0(R_1 - R_0) = R_1 - R_0$, these expressions equal the corresponding projected velocity losses. They satisfy

$$L_{\mathrm{sync}} \leq \frac{N}{N-1} L_{\mathrm{ind}}. \tag{37}$$

Derivation. Let $m_i = \mathbb{E}_{t\sim\nu} e_i(t)$, $\bar{m} = N^{-1}\sum_i m_i$, and $\varepsilon_i(t) = e_i(t) - m_i$. Set

$$\begin{aligned} A &= \frac{1}{NG}\sum_i \|m_i - \bar{m}\|_2^2, \\ B &= \frac{1}{NG}\sum_i \mathbb{E}_{t\sim\nu}\|\varepsilon_i(t)\|_2^2, \\ C &= \frac{1}{G}\mathbb{E}_{t\sim\nu}\left\| \frac{1}{N}\sum_i \varepsilon_i(t) \right\|_2^2. \end{aligned} \tag{38}$$

The identity $\|P_0 E\|_F^2 = \sum_i \|E_i\|_2^2 - N\|N^{-1}\sum_i E_i\|_2^2$ and independence of the sampled times give

$$L_{\mathrm{ind}} = A + \left(1 - \frac{1}{N}\right) B, \qquad L_{\mathrm{sync}} = A + B - C. \tag{39}$$

All three terms are nonnegative, so $L_{\mathrm{sync}} \leq A + B \leq NL_{\mathrm{ind}}/(N-1)$. The inequality holds conditional on the endpoints and memory and therefore also after averaging over their distribution. For groups with sizes at least $n_{\min}$, the common upper-bound factor is $n_{\min}/(n_{\min}-1)$. Replogle training uses $n_{\min} = 32$, giving a factor of $32/31$.

Equation (37) relates the base projected velocity-regression risks along fixed endpoint interpolations under the same time distribution. It establishes control of common-time interpolation error by the independent-time objective. Numerical integration and the resulting endpoint distributions are evaluated separately using the generation procedure and population metrics described in Appendices C and D.

C Implementation and Training Details

C.1 Architecture and Optimization

Tables 7 and 8 summarize the principal architecture and optimization settings for the three CHARM-Flow models. Experiments use Python 3.9.23, PyTorch 2.7.1 with CUDA 12.8, and PyTorch Lightning 2.5.6. Precomputed model inputs and network computations are described in Appendices A and B, respectively.

Table 7: Principal architecture settings for the three CHARM-Flow models.

Setting	Replogle	PBMC	Tahoe
Context width / depth / heads	256 / 2 / 8	256 / 2 / 8	256 / 2 / 8
Decoder width / depth / heads	384 / 12 / 8	384 / 12 / 8	384 / 12 / 8
Feed forward expansion	4	4	4
Time embedding width	256	256	256

Donor conditioned response correction. For PBMC, the population response branch incorporates variation across donors using the difference between the PBS mean of the target donor and the PBS mean of the same cell type across donors, weighted by cell counts. Two encoders of width 128 combine this donor offset with the perturbation query through a gene projection of rank 64. The resulting correction is bounded coordinatewise by 0.05, capped at 0.75 times the RMS of the base response, and scaled by $n/(n+64)$ for n available control cells. The correction is set to zero when fewer than 20 controls are available.

All three models are trained from scratch for 14,000 optimizer updates. We use Adam with decoupled weight decay, $(\beta_1, \beta_2) = (0.9, 0.95)$, a peak learning rate of 10^{-4}, and weight decay of 0.01. The learning rate warms up for 500 updates, remains at its peak for the next 1,400 updates, and then follows a cosine schedule with a floor of 10^{-5}. The gradient norm is clipped at 1.0.

Training data are organized into 32-cell episodes that share the same perturbation and biological context; PBMC episodes additionally share the same donor. Matching rows are combined within each 64-cell local microbatch for groupwise projection and group level losses. Gradient accumulation and distributed averaging give an effective batch size of 512 cells. The corresponding batch covariate remains part of the condition encoding.

Table 8: Key optimization settings for the three CHARM-Flow models.

Setting	Replogle	PBMC	Tahoe
Effective cells per optimizer update	512	512	512
Optimizer updates	14,000	14,000	14,000
Training seed	42	42	42
EMA decay	0.9996	0.9996	0.9996

The CHARM-Flow results reported in Section 4 use the model state after 14,000 optimizer updates with training seed 42. Numerical integration follows Appendix B.5, while generation groups, random seeds, and terminal support constraints are described in Appendix C.4.

C.2 Shared Training Objectives

The shared training objective combines residual flow matching with supervision for population response prediction, conditional source composition, state response structure, and centered residual distributions. All loss terms use the internal expression scale defined in Appendix A.1, with prepared expression values divided by ten. Let $\mathrm{SL1}_b$ denote the elementwise smooth-ℓ_1 loss with transition b, averaged over the stated coordinates. The default transition is $b = 1$. Losses are averaged over cells, selected cell–state pairs, or eligible condition groups, depending on the objective. Group level objectives assign equal weight to each eligible condition group.

In Eq. (17), $\mathcal{L}_{\mathrm{FM}}$, $\mathcal{L}_{\mathrm{KL}}$, and $\mathcal{L}_{\mathrm{resp}}$ supervise the residual flow, conditional source composition, and population response, respectively. The auxiliary objective $\mathcal{L}_{\mathrm{aux}}$ contains the state response, latent matching, and centered residual distribution terms described below, together with the expression support and gene-level objectives in Section C.3. Table 9 lists the coefficients shared by the three CHARM-FLOW models.

Table 9: Coefficients of the shared training objective for the three CHARM-FLOW models.

Loss	Coefficient
Residual flow matching	1
Population response MSE	1
Population response direction	5×10^{-4}
Population response correlation	2×10^{-4}
Bounded correction squared norm	10^{-2}
Conditional source KL	10^{-4}
State response MSE	10^{-2}
State response direction	2×10^{-4}
State response log magnitude	5×10^{-4}
Latent matching	1.05×10^{-3}
Residual set energy	3×10^{-4}
Residual set log variance	7.5×10^{-5}
Residual set correlation	1.5×10^{-5}

Population response supervision. For each row with an available training response, let $\mu_i^\star = \mu_{c_i} + \Delta_i^\star$ denote the corresponding observed population mean, and let $\widehat{\Delta}_i = \widehat{\mu}_i - \mu_{c_i}$ denote the predicted response relative to the control mean of context c_i. Let b_i denote the bounded learned correction before the support adjustment in Section C.3. The population response objective $\mathcal{L}_{\mathrm{resp}}$ combines

$$\frac{1}{G}\|\widehat{\mu}_i - \mu_i^\star\|_2^2, \qquad 1 - \cos(\widehat{\Delta}_i, \Delta_i^\star), \qquad 1 - \mathrm{corr}_g(\widehat{\Delta}_i, \Delta_i^\star),$$

together with the mean squared magnitude of b_i. The first three terms are averaged over rows with an available training response, while the correction penalty is averaged over its coordinates. Cosine similarity uses $\epsilon = 10^{-8}$.

Source composition supervision. The conditional source objective uses the response guided assignments from Appendix B.3. The KL term compares the group averaged soft assignments with the corresponding group averaged predicted state probabilities, assigning equal weight to each response group. The assignment target is detached from backpropagation, so this loss trains the predicted source composition without propagating gradients through the training-time assignment procedure.

State response supervision. For each response cell, auxiliary state response supervision is applied to the states with the largest soft assignment weights ω_{ik}. Let $\mathcal{J}_i$ denote the selected states and K_i the number of active states. The number selected is $\min(r(s), K_i)$, where

$$r(s) = \begin{cases} 32, & 0 \le s < 500, \\ 16, & 500 \le s < 1200, \\ 8, & 1200 \le s < 2200, \\ 4, & 2200 \le s < 3500, \\ 2, & 3500 \le s < 5000, \\ 1, & s \ge 5000. \end{cases} \tag{40}$$

This subset is used only for auxiliary supervision; source sampling continues to use the full soft assignment distribution.

For a selected cell–state pair (i, k), let d_{ik}^{gene} denote the decoded response from state k and $o_{ik}^{\mathrm{gene}} = y_i - \bar{x}_k$ the observed expression displacement relative to that state centroid. Supervision compares these two responses through gene averaged squared error, directional disagreement, and log RMS

magnitude difference. With

$$\mathrm{RMS}(x) = \sqrt{G^{-1}\|x\|_2^2 + 10^{-6}},$$

the corresponding response quality is

$$\begin{aligned} Q(d,o) = G^{-1}\|d-o\|_2^2 \\ + 0.25\,\frac{\mathrm{RMS}(o)}{\mathrm{RMS}(o)+0.02}[1-\cos(d,o)] \\ + 0.05\,|\log \mathrm{RMS}(d) - \log \mathrm{RMS}(o)|. \end{aligned} \tag{41}$$

Pair supervision is modulated by a detached weight w_{ik} that reflects how the state-specific response prediction compares with the corresponding population response prediction. Let Q_i^{pop} denote the response quality of the population prediction and Q_{ik} that of state k. We define

$$g_{ik} = \mathrm{clip}\left(\frac{Q_i^{\mathrm{pop}} - Q_{ik}}{\max(|Q_i^{\mathrm{pop}}|, 10^{-3})}, -2, 2\right). \tag{42}$$

For epoch $e < 3$, we set $w_{ik} = 0.1$. Thereafter, with $f_e = \mathrm{clip}((e-3)/17, 0, 1)$,

$$\begin{aligned} w_{ik} &= \mathrm{sg}\left[u_e + (1-u_e)\sigma\left(\frac{g_{ik}-m_e}{T_e}\right)\right], \\ (u_e, m_e, T_e) &= (0.1 - 0.08 f_e,\ 0.30 - 0.28 f_e,\ 0.08 - 0.04 f_e). \end{aligned} \tag{43}$$

The weighted pair losses are averaged over the number of valid selected pairs, rather than normalized by the sum of w_{ik}. PBMC and Tahoe complete training before epoch three and therefore use $w_{ik} = 0.1$ throughout these state response terms, whereas Replogle enters the adaptive weighting schedule.

The latent matching objective provides corresponding supervision in the encoded representation space. It combines a weighted matching cost over the selected state pairs with an unweighted population matching cost. Both use the same deviation, direction, and log magnitude components as the assignment cost in Appendix B.3, with relative weights $(1, 0.25, 0.05)$. Observed latent displacements are detached. The training loss uses the raw matching costs, while the assignment procedure uses their median-normalized counterparts. The assignments and reliability weights are detached, while the predicted responses remain differentiable.

Residual distribution supervision. Following the centered response coordinates in Appendix B.4, this auxiliary objective compares predicted and observed cellular variation independently of the population response mean. A fixed 64-dimensional PCA is fitted to all available control cells on the internal expression scale, and each projected component is divided by the square root of its corresponding control variance. Within each training condition group, the reconstructed endpoint and observed response cells are centered separately and projected into this control-derived representation for the distribution losses below.

For projected sets $U = \{u_i\}_{i=1}^n$ and $V = \{v_j\}_{j=1}^m$, the energy term is

$$E(U,V) = \frac{1}{\sqrt{64}}\left[\frac{2}{nm}\sum_{i,j}\|u_i - v_j\|_2 - \frac{1}{n^2}\sum_{i,j}\|u_i - u_j\|_2 - \frac{1}{m^2}\sum_{i,j}\|v_i - v_j\|_2\right]_+. \tag{44}$$

The remaining terms compare componentwise log variances and correlation matrices using smooth-ℓ_1 losses, with population variances and $\epsilon = 10^{-6}$. Groups containing at least 32 cells contribute equally. To keep the energy computation tractable, we use at most 256 cells from each set, selected at approximately even intervals after ordering cells by their detached squared norm. The variance and correlation terms use all eligible cells.

The energy, log variance, and correlation terms use coefficients 3×10^{-4}, 7.5×10^{-5}, and 1.5×10^{-5}, respectively, as listed in Table 9. Because the reconstructed and observed populations are centered separately before projection, these losses compare within-group cellular variation rather than differences in population response means.

C.3 Expression Support and Gene-Level Statistics

In addition to the centered distribution objective above, all three models receive gene-level supervision for expression support and zero expression. The support head combines a 256-dimensional condition representation with the gene identity embeddings used by the evolving residual stream through a rank-64 bilinear interaction. Let $\widehat{\Delta}^{(0)}$ denote the population response before support modulation. The support logits combine learned gene and condition terms with the response magnitude score

$$\frac{\left|\log\left(1+[\mu_c+\widehat{\Delta}^{(0)}]_+\right)-\log(1+[\mu_c]_+)\right|-0.065}{0.015}.$$

A separate head predicts an adjustment to the control zero-rate logit, bounded in magnitude by four.

The supervision targets are constructed from training responses. Replogle uses training differential expression labels, whereas PBMC and Tahoe use the 256 genes with the largest absolute mean changes within each training condition. These targets are used only for training and are distinct from the statistical differential expression evaluation in Appendix D.3.

For Replogle and Tahoe, the predicted support probabilities also modulate the population response correction. Let p^{sup} denote the predicted support probabilities. The adjusted response is

$$\widehat{\Delta} = (1-\beta(s))\widehat{\Delta}^{(0)} + \beta(s)\,\mathrm{sg}(p^{\mathrm{sup}}) \odot \tanh(\widehat{\Delta}^{(0)}), \qquad \beta(s) = 0.1\min(1, s/500). \tag{45}$$

The support probabilities are detached along this path, so the response objective does not propagate gradients into the support prediction through this adjustment. PBMC instead trains the support and zero-rate heads using detached condition and response features and does not apply the adjustment in Eq. (45).

Table 10: Base coefficients for the additional gene-level training objectives. Warmup and gradient balancing are described below.

Loss	Coefficient
Soft signed rank effect	2×10^{-4}
Soft detection rate logit	2×10^{-4}
Expression log variance	10^{-4}
Expression quantiles	10^{-4}
Effect score pairwise ranking	2×10^{-5}
Effect score support classification	10^{-4}
Support and zero rate head	10^{-3}
Endpoint zero rate consistency	5×10^{-4}
Support informed response signature	10^{-3}

Expression statistic losses. These objectives compare predicted and observed treated populations through four gene-level summaries: a signed rank effect, detection rate, variance, and expression quantiles. For gene g, the differentiable signed rank effect between a treated population X and its control population C is

$$u_g(X, C) = 2\,\mathrm{mean}_{i,j}\,\sigma\left(\frac{X_{ig} - C_{jg}}{T_g}\right) - 1, \qquad T_g = \mathrm{clip}(0.2\,\mathrm{std}(C_{:g}), 0.003, 0.03)\,.$$

Predicted and observed signed rank effects are compared using $\mathrm{SL1}_{0.1}$.

The soft detection rate for gene g is $\mathrm{mean}_i\,\sigma(X_{ig}/0.003)$. Rates are clipped to $[10^{-4}, 1-10^{-4}]$ before comparing their logits with smooth-ℓ_1. Log variance is computed from the population variance with an additive 10^{-6}. Quantile supervision uses levels $(0, 0.25, 0.5, 0.75, 0.9)$; predicted and observed quantiles are scaled by

$$\mathrm{clip}\left(\mathrm{std}(Y_{:g}) + 0.003,\ 0.003,\ 0.1\right)$$

before smooth-ℓ_1 comparison.

The signed rank effect also defines a gene-level score used to distinguish support genes from other genes:

$$z_g = |u_g| \sqrt{\frac{3 n_t n_c}{n_t + n_c + 1}},$$

where n_t and n_c are the numbers of training treated and control cells, respectively, each lower bounded by two. To encourage support genes to receive stronger effect scores than non-support genes, the pairwise ranking loss compares up to 256 support genes with the weakest scores against up to 256 non-support genes with the strongest scores:

$$\mathrm{mean}_{g^+, g^-} \, \mathrm{softplus}\left(\frac{z_{g^-} - z_{g^+}}{0.5} \right).$$

The ranking term is zero when either set is empty. A complementary support classification loss applies binary cross-entropy to the logit $(z_g - 4)/0.5$, with weight two on negative labels.

These effect-based objectives are used only as training surrogates; the held-out differential expression metrics are computed independently using the statistical procedure in Appendix D.3. Each statistic is first averaged over genes and then over eligible condition groups containing at least eight cells. Training targets and controls are detached.

Support head and endpoint losses. Support supervision is first aggregated within each training condition group. The support head is trained with binary cross-entropy on the gene-level support labels together with a smooth-ℓ_1 loss between the predicted and target support fractions. The latter term has relative weight 0.25. The zero-rate head is supervised by the Bernoulli KL divergence from the observed to the predicted gene-wise zero probabilities, with each gene weighted by $n_t/(n_t + 64)$ according to the number of available treated cells. These losses are averaged over genes and condition groups.

A separate endpoint consistency term aligns the reconstructed expression population with the predicted zero probabilities. For gene g, the endpoint zero rate is represented by

$$\mathrm{mean}_i \exp\Big(-[\widehat{X}_{ig}]_+ / 0.003 \Big)$$

and compared with the predicted condition-level zero probability using Bernoulli KL. The predicted probability is detached in this comparison, and both quantities are clipped to $[10^{-6}, 1 - 10^{-6}]$.

Support informed response signature. This objective regularizes the bounded population response correction relative to the training response prior. Let p denote the detached support probabilities, b the condition averaged bounded correction, Δ_0 the training response prior, and $\Delta^\star$ the observed training response.

The correction is encouraged to match the support-weighted residual between the observed training response and the response prior:

$$p \odot \mathrm{clip}\left(\Delta^\star - \Delta_0, -0.05, 0.05 \right).$$

The comparison averages mean absolute difference, RMS difference with 10^{-8} inside the square root, and cosine distance. A second term penalizes correction magnitude outside the predicted support through $\mathrm{mean}\,|(1 - p) \odot b|$.

Two additional terms constrain the corrected response $\Delta_0 + b$. A margin loss compares its distance to the corresponding training response with its distance to the closest alternative training response in the same context, using margin 0.005. The final term penalizes the positive excess of the corrected response distance to its target over the corresponding distance of the uncorrected prior. Distance comparisons use the same three measures with equal weight. Conditions without an alternative training response omit the margin term. The four components have unit relative weight, and the overall objective has coefficient 10^{-3}.

Warmup and gradient balancing. The additional gene-level objectives are introduced with a linear warmup over the first 500 optimizer updates. For objectives whose gradient contribution is

additionally controlled, let a denote the weighted auxiliary loss, b its reference objective, and z the representation at which their gradient magnitudes are compared. We use

$$\widetilde{a} = \mathrm{sg}\left[\min\left(1, \gamma \frac{\|\nabla_z b\|_2}{\max(\|\nabla_z a\|_2, 10^{-12})}\right)\right] a. \tag{46}$$

The six expression statistic objectives are scaled jointly with $\gamma = 0.3$, using the flow matching objective as the reference and the raw velocity as z. Endpoint zero rate consistency uses the same reference with $\gamma = 0.2$. The response signature objective uses the bounded correction as z and the combined population, source, and state supervision as its reference, also with $\gamma = 0.2$. The support and zero rate head objective uses the 500-update warmup without gradient balancing.

C.4 Generation and Reconstruction

Generation groups. Group centering and mean constrained reconstruction are applied within generation groups. Within a sampling microbatch, cells sharing the same intervention and biological context form one generation group; PBMC groups additionally share the same donor. A single evaluation population may therefore combine several generation groups when its cells span multiple sampling microbatches.

Sampling protocol. Replogle generation uses seed 42 for both the data stream and model sampling. PBMC generation is divided into seven fixed cytokine subsets, each assigned a deterministic seed derived from 42; predictions from all subsets are combined before evaluation. For Tahoe, the evaluation protocol considers doses of 0.05, 0.5, and $5.0\,\mu$M. It deterministically selects 245 eligible groups at each dose under the double OOD split, where a group is defined by drug, dose, cell line, and plate. Duplicate cell identities are removed within each group, after which 64 treated cells are sampled uniformly without replacement. Each group uses a seed computed from seed 42 and its identifier. The model random generator is reset to seed 42 before generation.

Expression support. Reconstruction maps the terminal residuals back to nonnegative expression with prescribed group mean μ. Gene-wise expression support is handled separately. For training endpoints, the observed response support is incorporated into the target projection. During generation, the support pattern is determined by the learned zero-rate prediction.

For a generated group of size N, let $z_g^{(0)}$ denote the number of cells whose current expression for gene g is zero. The predicted zero probability is

$$q_g = \sigma\left(\frac{1}{N}\sum_i \ell_{ig}^{\mathrm{zero}}\right),$$

and the requested zero count is

$$\begin{aligned} z_g &= \min\left(z_g^{\max}, \max\left\{z_g^{(0)}, \lfloor N q_g + \tfrac{1}{2} \rfloor\right\}\right), \\ z_g^{\max} &= \begin{cases} N-1, & \mu_g > 0, \\ N, & \mu_g = 0. \end{cases} \end{aligned} \tag{47}$$

Existing zeros are retained, and additional zeros are assigned preferentially to cells with the smallest expression values for gene g. The remaining positive entries are then projected onto the nonnegative simplex with total mass $N\mu_g$. Consequently, zero-mean genes remain zero throughout the group, while genes with positive mean retain nonzero expression in at least one cell.

D Evaluation Protocols and Metrics

D.1 Evaluation Units and Control Populations

Following the STATE evaluation protocol (Adduri et al., 2026), also used by PerturbDiff (Yuan et al., 2026), we compute the main metrics with Cell-Eval 0.6.6 from unpaired observed and generated populations using the fixed 2,000-gene panel on the prepared expression scale. Each prediction matches the treated-cell count of its observed evaluation population. The CHARM-Flow models and comparison variants reported in this paper are trained and scored locally. Published results for CPA, STATE, CellFlow, Squidiff, and PerturbDiff follow the unified Cell-Eval benchmark of Yuan et al. (2026); OCOO-T results follow Jiang et al. (2026). Values are retained at the precision reported by the corresponding studies. Table 11 summarizes the evaluation populations, reference controls, and macro-averaging units.

Table 11: Evaluation populations, reference controls, and macro-averaging units. Treated counts apply to both observed and generated populations.

Evaluation	Treated cells	Controls	Macro-averaging units
Replogle	26,878	4,976	380 genetic perturbations
PBMC	2,260,453	235,478	62 pooled cytokines
Tahoe	47,040	204,536	607 pooled drug–dose labels

Replogle. All selected test cells from 380 genetic perturbations in HepG2 are evaluated against a shared non-targeting control population. Metrics are computed per perturbation and macro-averaged.

PBMC. Cells from 18 cell types and four donors are pooled within each of 62 cytokines and compared with pooled PBS controls. Cytokines are macro-averaged, so cell types and donors contribute in proportion to their cell counts within each cytokine.

Tahoe. Sixty-four treated cells are sampled from each of 735 drug–dose–cell-line–plate groups across five evaluation cell lines. Groups are pooled by drug–dose into 607 evaluation units, compared with pooled DMSO controls, and macro-averaged. Each selected group contributes 64 observed and 64 generated cells to its drug–dose population.

D.2 Pseudobulk Response and Discrimination Metrics

Let u index an evaluation unit, and let $\bar{x}_u$, $\bar{\hat{x}}_u$, and $\bar{c}_u \in \mathbb{R}^G$ denote its observed, predicted, and reference control pseudobulks. We define response vectors $\delta_u = \bar{x}_u - \bar{c}_u$ and $\hat{\delta}_u = \bar{\hat{x}}_u - \bar{c}_u$.

PDCorr. PDCorr is the Pearson correlation across genes between the observed and predicted response vectors:

$$\mathrm{PDCorr}_u = \mathrm{corr}_g(\delta_{ug}, \hat{\delta}_{ug}). \tag{48}$$

Pseudobulk errors. MSE and MAE compare the observed and predicted pseudobulks directly:

$$\mathrm{MSE}_u = G^{-1}\|\bar{\hat{x}}_u - \bar{x}_u\|_2^2, \qquad \mathrm{MAE}_u = G^{-1}\|\bar{\hat{x}}_u - \bar{x}_u\|_1. \tag{49}$$

Perturbation discrimination. Each predicted response is compared with the observed response vectors of all candidate labels in the same dataset. If $r_u^{(d)} \in \{0, \ldots, |\mathcal{U}| - 1\}$ is the zero-based rank of the correct label under distance d, then

$$\mathrm{PDS}_u^{(d)} = 1 - \frac{r_u^{(d)}}{|\mathcal{U}|}, \qquad d \in \{L_1, L_2, \cos\}. \tag{50}$$

The three distances are L_1, L_2, and cosine distance between response vectors; controls are excluded from the candidate pool. For Replogle, the query gene is omitted from the distance calculation when it is present in the modeled panel, whereas PBMC and Tahoe use the full panel.

D.3 Differential-Expression Metrics

Testing and correction. Observed and predicted treated cells are tested separately against the specified controls using a two-sided Mann–Whitney U test for each gene. Benjamini–Hochberg correction is applied separately to the observed and predicted results over all evaluation-unit–gene tests within each dataset: $380 \times 2{,}000$ for Replogle, $62 \times 2{,}000$ for PBMC, and $607 \times 2{,}000$ for Tahoe. Genes with $q < 0.05$ are considered significant.

Significant genes are ranked by absolute log-base-two fold change. Cell-Eval computes this fold change from the transformed pseudobulk mean $\exp(\bar{x}) - 1$ on the stored log-expression scale. Let ℓ_{ug} and $\hat{\ell}_{ug}$ denote the observed and predicted signed values.

DEOver and DEPrec. Let A_u and B_u be the observed and predicted significant-gene lists, ordered by decreasing absolute fold change, and let $A_u[:k]$ denote the first $\min(k, |A_u|)$ entries. The adaptive-length metrics are

$$\begin{aligned} \mathrm{DEOver}_u &= \frac{|A_u \cap B_u[:|A_u|]|}{|A_u|}, \\ \mathrm{DEPrec}_u &= \frac{|A_u[:|B_u|] \cap B_u|}{|B_u|}. \end{aligned} \tag{51}$$

Direction agreement. Direction agreement is evaluated on the observed significant genes, whether or not they are significant in the prediction:

$$\mathrm{DirAgr}_u = \frac{1}{|A_u|} \sum_{g \in A_u} \mathbf{1}\{\mathrm{sign}(\ell_{ug}) = \mathrm{sign}(\hat{\ell}_{ug})\}. \tag{52}$$

D.4 Centered Within-Perturbation Distribution Metrics

The Replogle centered distribution analyses project observed and predicted cells through the same fixed, control-derived 64-dimensional PCA basis. Each projected component is divided by the square root of its corresponding control variance. The two populations are centered separately within each perturbation, removing their pseudobulk displacement. Perturbations with fewer than eight cells on either side are excluded, and larger populations are capped at 256 cells. Scores are macro-averaged over the same 380 held-out perturbations.

All three distances are lower-is-better. Normalized centered energy distance (NCED) divides empirical energy distance by the observed population's mean pairwise spread. Centered multiscale maximum mean discrepancy (CMMD) uses RBF bandwidths at 0.5, 1, and 2 times a common control-derived median distance. Normalized centered sliced Wasserstein-1 (NCSW1) averages Wasserstein-1 distance over 512 fixed Gaussian directions and normalizes it by the observed sliced pairwise spread.

Gene-wise distribution and dispersion. Gene-wise metrics use all cells from each eligible perturbation and the same fixed 2,000-gene panel. Normalized centered gene-wise Wasserstein-1 (NCGW1) sums Wasserstein-1 distances between separately centered observed and predicted gene distributions and divides by the sum of observed gene-wise mean pairwise absolute differences. NStdL1 and NIQRL1 sum absolute errors in gene-wise sample standard deviations and interquartile ranges, respectively, and normalize by the corresponding sums of observed statistics. The predicted-to-observed dispersion ratios divide the sums of predicted statistics by the corresponding observed sums within each perturbation and are then macro-averaged. ZeroRateMAE averages absolute gene-wise zero-fraction errors on uncentered expression. DetectionRateMAE uses the fraction exceeding 10^{-6} on the model's expression scale.

D.5 Aggregation

Metrics are computed within each scoring unit and then macro-averaged. Reported standard deviations summarize variation across these units. Empty significant-gene lists contribute zero to the corresponding overlap, precision, or direction-agreement metric.

E Additional Model Analyses

The analyses below use the same 380 Replogle perturbations held out from training. Controlled model comparisons and distribution analyses use 14,000 optimizer updates and training seed 42. Table 13 additionally reports independent training runs with different seeds.

E.1 Controlled Model Comparisons

Table 12 compares three variants that remove distinct components of CHARM-FLOW. The marginal empirical source removes the intermediate state organization while retaining empirical control cell variation and persistent source references. The population conditioned vector field removes direct state context during residual generation while leaving the state structured source unchanged. The initialization only variant uses the same sampled control cell for residual initialization but removes subsequent access to its source representation. Each variant is trained independently end to end.

Table 12: Controlled comparisons on Replogle. The variants remove state organization, direct state context during residual generation, or persistent source access, respectively.

Configuration	PDCorr ↑	DEOver ↑	DEPrec ↑	MSE ↓
CHARM-FLOW	0.4773	0.2542	0.2430	5.8307×10^{-3}
Marginal empirical source	0.4755	0.2267	0.2311	5.9128×10^{-3}
Population-conditioned vector field	0.4778	0.2372	0.2356	5.8506×10^{-3}
Initialization only	0.4651	0.1565	0.1954	1.3434×10^{-2}

Training stability. Table 13 compares three independently trained CHARM-FLOW models under the same 380 perturbation evaluation protocol. Population response and centered distribution metrics are similar across the three training seeds.

Table 13: Training stability of CHARM-FLOW on Replogle across three independent runs. Mean ± SD is computed across runs under the same 380 perturbation evaluation protocol.

Training seed	PDCorr ↑	MSE ↓	DEOver ↑	DEPrec ↑	NCED ↓	CMMD ↓	NCSW1 ↓
42	0.4773	0.005831	0.2542	0.2430	0.0224	0.1083	0.1663
2026	0.4737	0.005969	0.2411	0.2339	0.0238	0.1137	0.1781
2027	0.4751	0.005957	0.2519	0.2416	0.0233	0.1118	0.1752
Mean ± SD	0.4754 ± 0.0018	0.005919 ± 0.000076	0.2491 ± 0.0070	0.2395 ± 0.0049	0.0232 ± 0.0007	0.1113 ± 0.0027	0.1732 ± 0.0062

Residual generation comparisons. We compare CHARM-FLOW with two alternatives that differ in how the centered residual is transformed. Source only applies no learned residual transformation and returns the prepared empirical source in centered response coordinates. Direct MLP replaces the time conditioned vector field and numerical integration with a three layer conditional MLP that predicts the centered endpoint displacement in a single step, while retaining the same auxiliary supervision and terminal reconstruction as CHARM-FLOW. This comparison isolates the form of the learned residual transformation under a common training and evaluation protocol.

Table 14: Residual-generation comparisons on 380 held-out Replogle perturbations. All learned configurations use 14,000 optimizer updates and training seed 42. Source-only returns the prepared empirical source. Direct-MLP retains the same auxiliary supervision and terminal reconstruction while replacing the residual field and integration with a one-step mapping.

Configuration	NCED ↓	NStdL1 ↓	DEOver ↑	DEPrec ↑
Source-only endpoint	0.0242	0.1684	0.2222	0.1941
Direct-MLP endpoint	0.0266	0.4292	0.1369	0.1685
CHARM-FLOW	0.0224	0.1658	0.2542	0.2430

Under the shared training and reconstruction protocol, the complete model yields lower NCED and NStdL1 and higher DEOver and DEPrec than Direct-MLP. Source-only provides the untransformed empirical source endpoint for the same hierarchical source and population response framework.

E.2 Population Mean and Training Response Prior Diagnostics

Table 15 compares CHARM-FLOW with two fixed predictors that separate population response prediction from cellular variation. Mean only repeats each reconstruction group's predicted population mean across all of its cells, leaving the group mean unchanged while eliminating cell to cell variation within the group. The training prior predictor combines the target context control mean with the calibrated response prior from Eq. (22). Both predictors require no additional training.

Table 15: Population mean and training response prior diagnostics on Replogle. Mean only repeats each predicted group mean across its cells, while Training prior combines the target context control mean with the calibrated response prior.

Metric	CHARM-FLOW	Mean only	Training prior
PDCorr	0.4773	0.4773	0.3956
DEOver	0.2542	0.1988	0.2050
DEPrec	0.2430	0.1004	0.0874
DirAgr	0.7957	0.7957	0.7259
PDS-L1	0.8153	0.8153	0.5996
PDS-L2	0.8402	0.8402	0.6296
PDS-cos	0.7814	0.7814	0.7453
MSE	5.8307×10^{-3}	5.8307×10^{-3}	7.5207×10^{-3}
MAE	0.052061	0.052061	0.054671

Because Mean only uses the same predicted group means as CHARM-FLOW, PDCorr and MSE remain unchanged at the reported precision. In contrast, DEOver decreases from 0.2542 to 0.1988 and DEPrec from 0.2430 to 0.1004 when cell to cell variation is removed. Thus, at fixed predicted group means, cellular variation contributes substantially to the differential expression metrics.

The training prior predictor reaches a PDCorr of 0.3956 and an MSE of 7.52×10^{-3}, compared with 0.4773 and 5.83×10^{-3} for CHARM-FLOW. The calibrated cross context prior therefore provides substantial population response signal before the learned response correction.

E.3 State Structured Source Analyses

Centered distribution metrics follow Appendix D.4. Table 16 reports both the relative error reduction from state structured source organization and the fraction of perturbations for which it yields lower centered distribution error than the marginal empirical source.

Table 16: State structured versus marginal empirical source on 380 Replogle perturbations. The table reports relative error reduction and the fraction of perturbations with lower centered distribution error for the state structured source.

Metric	Mean error reduction ↑	Lower error fraction ↑
NCED	30.49%	89.21%
CMMD	24.12%	89.74%
NCSW1	23.55%	92.63%

The reduction in centered distribution error is consistent across individual perturbations. The state structured source yields lower NCED for 89.21% of the 380 perturbations, lower CMMD for 89.74%, and lower NCSW1 for 92.63%. Together with the lower macro averaged errors in Table 4, these results show that the effect of state structured source organization extends across a broad set of perturbations.

E.4 Persistent Source Reference Analyses

Persistent source access and cellular variation. Table 17 characterizes generated cellular variation using centered distribution, gene wise dispersion, and expression support metrics. It compares persistent source access with initialization only access over the same 380 Replogle perturbations. In both models, the sampled control cell defines the initial residual. In CHARM-FLOW, its source representation remains available throughout residual generation, whereas the initialization only model uses the sampled cell only for initialization. The two models are trained independently under the same training and generation framework.

Table 17: Persistent and initialization only source access on 380 Replogle perturbations. Lower error fractions refer to persistent source access.

Metric	Initialization only	Persistent	Reduction	Lower error fraction
NCED	0.02936	0.02244	23.60%	85.26%
CMMD	0.13277	0.10831	18.42%	85.53%
NCSW1	0.20534	0.16629	19.02%	89.47%
NCGW1	0.33962	0.24217	28.69%	99.47%
NStdL1	0.30498	0.16585	45.62%	100.00%
NIQRL1	0.39829	0.31145	21.80%	95.00%
ZeroRateMAE	0.06459	0.05430	15.92%	95.53%
DetectionRateMAE	0.06707	0.05433	18.99%	94.74%

Persistent source access lowers all eight distribution, dispersion, and expression support errors in Table 17. NCED decreases from 0.02936 to 0.02244, NCGW1 from 0.33962 to 0.24217, and NStdL1 from 0.30498 to 0.16585. The same pattern holds across individual perturbations: persistent access has lower NCED for 324 of 380 perturbations, lower NCGW1 for 378, and lower NStdL1 for all 380.

Direct use of the retained source reference. Table 18 applies a paired inference intervention to the same trained CHARM-FLOW model. The sampled control cells still define the initial residuals, while source reference cross attention is omitted during subsequent residual generation. This separates persistent source access from the source cell's role in initialization.

Table 18: Paired inference intervention on the same trained CHARM-FLOW model. Both conditions use the sampled control cells for residual initialization; Initialization only omits source reference cross attention during subsequent residual generation.

Source access	PDCorr ↑	MSE ↓	NCED ↓	CMMD ↓	NCSW1 ↓	NCGW1 ↓	NStdL1 ↓	DEOver ↑	DEPrec ↑
Initialization only	0.47732	5.831×10^{-3}	0.30163	0.42091	0.98517	0.78313	0.89709	0.22839	0.18322
Persistent	0.47732	5.831×10^{-3}	0.02239	0.10811	0.16623	0.24225	0.16582	0.25396	0.24286

When the retained source representation is omitted after initialization, PDCorr and MSE remain unchanged, while NCED increases from 0.02239 to 0.30163 and DEPrec decreases from 0.24286 to 0.18322. The retained source representation therefore continues to influence residual generation beyond its role in initialization.

E.5 Effects Across Individual Perturbations

Figure 2 shows how the aggregate effects of state organization and persistent source access extend across individual Replogle perturbations. The state structured source has lower NCED than the marginal empirical source for 339 of the 380 perturbations. Persistent source access yields higher DEOver than initialization only access for 255 perturbations.

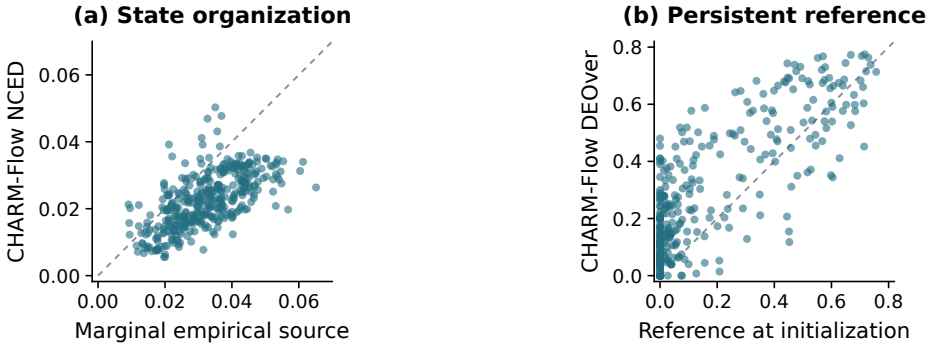

Figure 2: Comparison across the 380 Replogle perturbations held out from training. Each point represents one perturbation. Panel (a) compares NCED between the state structured and marginal empirical sources, and panel (b) compares DEOver between persistent and initialization only source access. The diagonal marks equal values.

E.6 Control Derived Representation Diagnostics

State partition consistency. The split sample analysis examines the consistency of the control derived state partition across disjoint subsets of its member cells. Each state reservoir is divided into two halves. One half is used to estimate the state centroid, and cells in the other half are ranked by distance to all state centroids within the same cell line. The analysis covers 512 states across 4 cell lines and 13,801 held out control cells. Table 19 reports top k identification and rank based statistics for these cells.

Table 19: State partition consistency on Replogle. State centroids are estimated from one half of each reservoir and evaluated using cells from the other half. The random baseline assumes uniform ranking among the 128 states within each cell line.

Statistic	Held out cells	Random baseline
Top 1 state identification	0.4087	0.0078
Top 5 state identification	0.8126	0.0391
State averaged Top 1	0.3922	0.0078
State averaged Top 5	0.7874	0.0391
Mean reciprocal rank	0.5810	0.0424
Median assigned state rank	2	64.5

For held out control cells, the assigned state ranks first for 40.87% of cells and within the top five for 81.26%, with a median assigned state rank of 2. These results indicate that the fixed control derived partition is consistent across disjoint subsets of its member cells.

Gene partition structure. The control derived gene partition groups genes according to covariation in control cells. Table 20 compares correlations among genes assigned to the same group with those from size matched random groups.

Table 20: Control correlations within gene groups on Replogle, compared with size matched random groups.

Statistic	Gene groups	Random baseline
Within group control Pearson (pair weighted)	0.0663	0.0105
Within group control Pearson (group averaged)	0.0452	0.0105

Genes assigned to the same control derived group show stronger pair weighted and group averaged Pearson correlation than genes in size matched random groups. This pattern indicates that the fixed partition retains structured gene covariation in the control population.