PERTURBSB: CONTROL-INFORMED SCHRÖDINGER BRIDGES FOR SINGLE-CELL PERTURBATION RESPONSE PREDICTION

Anonymous authors
Paper under double-blind review

ABSTRACT

Single-cell perturbation data are central to understanding cellular mechanisms and applications such as drug discovery. Yet perturbation effects vary across cellular states, making it difficult to predict responses in states without treated observations. We propose PerturbSB, a Schrödinger-bridge-inspired framework for transferring perturbation responses across cellular contexts. Its stochastic dynamics separate a learned intrinsic drift, shared by control and treated cells, from a state-dependent perturbation drift. Target control cells provide the initial distribution, while training matches predicted treated endpoints to observed treated populations in source contexts. Across three single-cell benchmarks, PerturbSB delivers competitive response predictions and outperforms established baselines on several evaluation measures. These results suggest that separating intrinsic and perturbation dynamics can help predict responses in unseen cellular states.

1 INTRODUCTION

The transcriptional effect of a perturbation depends on the cellular context. Single-cell measurements reveal this diversity, but they usually sample untreated and treated cells from separate populations. A predictor must therefore learn a response from unpaired data and apply it to a context for which only control cells are available.

Existing predictors represent this transfer in different ways. scGen applies latent vector arithmetic to a variational autoencoder representation (Lotfollahi et al., 2019), whereas CellOT learns a nonlinear map between control and treated distributions (Bunne et al., 2023). A stochastic transition offers a different parameterization: the perturbation-dependent drift is repeatedly evaluated at the evolving state. Its intermediate states are model-implied and need not correspond to measured biological time.

Schrödinger bridges (SBs) regularize stochastic transitions relative to a reference process. They have already been used to learn dynamics from unpaired snapshots and to predict perturbation responses (Tong et al., 2024; Chi et al., 2025; Luo et al., 2025). Our focus is the transfer setting: *can a stochastic response model learned from observed contexts predict a known perturbation in a context without treated observations?* The target context contributes control cells as initial states, while its treated cells remain unavailable during fitting.

We introduce PerturbSB, which shares a reference drift and diffusion between control and perturbed branches and adds a perturbation correction that is exactly zero for controls. Control cells supply the initial distribution; available treated populations supervise the predicted endpoint. A diffusion-weighted energy regularizes the correction. The reference is learned jointly with the response model, so this is a soft, SB-inspired objective rather than a classical bridge with a fixed reference. We assess its predictive value on Kang, Haber, and Hagai without interpreting simulated paths as observed cell trajectories.

Our contributions are:

- A control-informed stochastic transition with a shared reference and a state- and time-dependent perturbation correction.

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

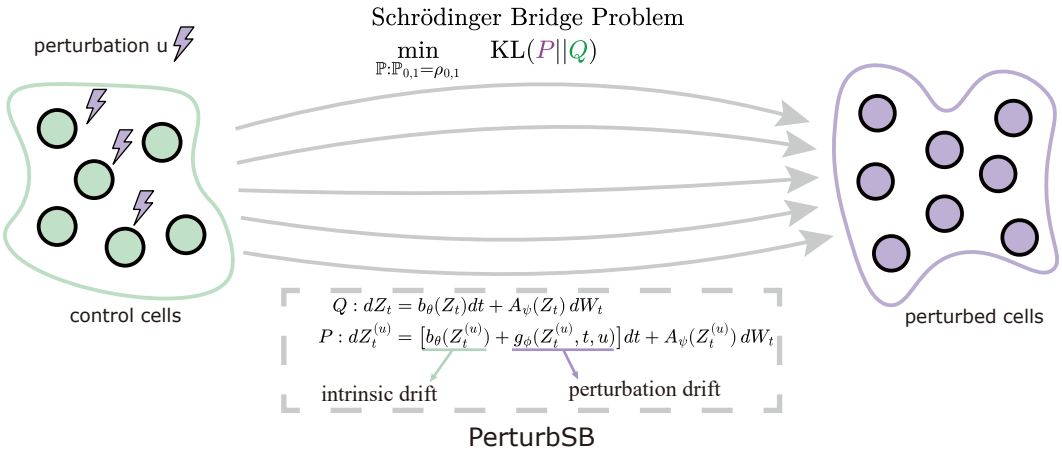

Figure 1: PerturbSB overview. Control cells initialize the transition. A shared reference drift and diffusion are combined with a gated perturbation correction, and available treated cells supervise the predicted endpoint. Intermediate states are model-implied.

- An endpoint distribution-matching objective with a diffusion-weighted relative path penalty.
- An evaluation across held-out cell types and species, including simple baselines and ablations that expose both gains and failures.

2 Related Work

Response prediction from unpaired populations. scGen learns an expression representation and transfers a latent perturbation displacement (Lotfollahi et al., 2019). CellOT uses input-convex neural networks to learn an optimal transport map (Bunne et al., 2023). We compare with their official implementations under shared data visibility and evaluation panels. A global mean shift and an identity predictor are also necessary controls: preserving the untreated distribution can perform well when the measured response is weak. Our use of a local stochastic correction changes the predictive parameterization, but does not establish greater expressiveness or accuracy than endpoint maps in general.

Stochastic transport and Schrödinger bridge. Simulation-free score and flow matching connects entropy-regularized transport to stochastic dynamics inferred from unpaired snapshots (Tong et al., 2024). Departures uses minibatch transport pairing to guide neural bridge learning for discrete gene activation and continuous expression (Chi et al., 2025). Direction-constrained diffusion SBs also address single-cell perturbation prediction (Luo et al., 2025). Thus, neither stochastic paths nor SBs for perturbation prediction are new in themselves. We focus on the shared reference drift and relative correction under held-out-context evaluation. These bridge methods are related work, not evaluated baselines; the experiments cannot establish superiority over them. Soft endpoint penalties have also been studied in SB control (Garg et al., 2024), while minibatch Wasserstein objectives have statistical properties distinct from population transport distances (Fatras et al., 2020).

3 Method

3.1 Problem and Representation

Let $X^0 \in \mathbb{R}^{N_0 \times G}$ and $X^1 \in \mathbb{R}^{N_1 \times G}$ contain normalized control and treated expression over G genes. The populations are unpaired. For each held-out context, treated cells are excluded from fitting and model selection, whereas an allowed subset of control cells remains available. The task is to predict the treated distribution from independent target-control inputs.

We reduce the G genes to an R-dimensional latent space via highly variable gene (HVG) selection and PCA. An encoder E maps expression to a latent state $z \in \mathbb{R}^R$, with decoder D. Write $Z^0 = E(X^0) \in \mathbb{R}^{N_0 \times R}$ and $Z^1 = E(X^1) \in \mathbb{R}^{N_1 \times R}$ for the encoded control and treated states; all dynamics are learned in this latent space and predictions are decoded back to gene space. The PCA representation is a modeling bottleneck, not a lossless description of the response.

3.2 A Schrödinger bridge model

A Schrödinger bridge (SB) learns the stochastic transition between two observed distributions that stays closest to a reference process. Given a reference diffusion Q and the encoded marginals ν^0, ν^1, it solves

$$\min_{P:\, P_0=\nu^0,\, P_1=\nu^1} \mathrm{KL}(P\|Q), \tag{1}$$

where ν^u is the encoded distribution for condition u. This classical problem assumes a fixed reference and exact marginals. Our finite-sample model uses soft distributional discrepancies and learns its reference jointly; its optimizer therefore need not solve Eq. (1).

Over a normalized transition interval $t \in [0, 1]$, we parameterize the transition as

$$\begin{aligned} dZ_t^{(u)} &= \big[b_\theta(Z_t^{(u)}) + g_\phi(Z_t^{(u)}, t, u)\big]\,dt + A_\psi(Z_t^{(u)})\,dW_t, \\ g_\phi(z,t,0) = 0, &\qquad A_\psi(z) = \mathrm{diag}\left(\sigma_{\psi,1}(z), \ldots, \sigma_{\psi,R}(z)\right), \end{aligned} \tag{2}$$

where b_θ is the reference drift and g_ϕ the perturbation correction. Both conditions share the same positive diagonal diffusion A_ψ, and the correction is exactly zero for controls ($u = 0$). The reference is a learned component of the predictor; it is not identified as the true unperturbed dynamics.

Let $Q_{\theta,\psi}$ be the law of the control process (drift b_θ) and $P_{\theta,\phi,\psi}$ that of the treated process (drift $b_\theta + g_\phi$), both starting from ν^0. For fixed parameters, common initial law and diffusion, and standard absolute-continuity and integrability conditions, Girsanov's identity gives a diffusion-weighted path KL:

$$\mathcal{E}(\theta,\phi,\psi) = \mathrm{KL}(P_{\theta,\phi,\psi}\|Q_{\theta,\psi}) = \frac{1}{2}\mathbb{E}^P \int_0^1 g_\phi(Z_t,t,1)^\top\, \Sigma_\psi(Z_t)^{-1}\, g_\phi(Z_t,t,1)\,dt, \tag{3}$$

where $\Sigma_\psi(z) = A_\psi(z)A_\psi(z)^\top = \mathrm{diag}\left(\sigma_{\psi,j}(z)^2\right)$ is the diffusion covariance. The energy $\mathcal{E}$ penalizes the correction in diffusion coordinates; Appendix B gives the derivation. Because the reference is learned jointly, this identity does not establish a unique solution or identify a biological trajectory.

3.3 Objective and prediction

The objective matches treated endpoints and regularizes the drift correction:

$$\mathcal{L} = \mathcal{L}_{\mathrm{W2}} + \lambda\mathcal{E}, \tag{4}$$

Here $\mathcal{L}_{\mathrm{W2}}$ compares the perturbed endpoint with treated cells from source contexts; control cells provide initial states, without a separate control endpoint loss. The main configuration uses dimension-normalized squared Wasserstein matching. For equally weighted batches a, b of size B, the matching cost is

$$\mathcal{D}_B(a,b) = \frac{1}{BR}\min_{\pi\in\mathfrak{S}_B}\sum_{i=1}^{B}\|a_i - b_{\pi(i)}\|_2^2, \tag{5}$$

with an exact assignment solver selecting the matching and gradients flowing through the selected costs. These assignments are optimization devices, not measured cell correspondences. The energy is integrated along perturbed rollouts under the Euler–Maruyama discretization

$$Z_{k+1}^{(u)} = Z_k^{(u)} + [b_\theta(Z_k^{(u)}) + g_\phi(Z_k^{(u)}, t_k, u)]\Delta t_k + A_\psi(Z_k^{(u)})\sqrt{\Delta t_k}\,\epsilon_k, \quad \epsilon_k \sim \mathcal{N}(0, I). \tag{6}$$

At prediction time, target controls are encoded, evolved under the perturbed branch, and reconstructed:

$$\widehat{\mu}^1 = D_\#\big(\nu^0 K^1\big), \qquad K^1(z,A) = \Pr\left(Z_1^{(1)} \in A \mid Z_0^{(1)} = z\right), \tag{7}$$

where ν^0 is the encoded target-control distribution. Noise generates a conditional response distribution, but its spread is not automatically calibrated uncertainty, and endpoint supervision does not establish individual-cell counterfactual accuracy or the correctness of intermediate paths.

4 EXPERIMENTS

4.1 PROTOCOL, BASELINES, AND METRICS

Tasks. The main evaluation considers three biological settings. Kang measures IFN-β stimulation in PBMCs (Kang et al., 2018); CD4 T cells, CD8 T cells, and CD14$^+$ monocytes are held out in turn. Haber measures intestinal epithelial responses to *H. polygyrus* infection (Haber et al., 2017); targets are Tuft, Enterocyte, and Stem cells, comparing controls with day-3 infection. Hagai measures LPS responses across species (Hagai et al., 2018); mouse and rat are held out in turn, using pig, rabbit, and the other rodent species as sources, with unstimulated 0 h versus LPS 6 h.

Visibility and preprocessing. Splits use seed 2026. Target treated cells are reserved for frozen evaluation; target controls are split between fitting and test input. Starting from raw counts, the reported pipeline uses library-size normalization to 10^4, $\log(1+x)$, 2,000 training-selected HVGs, PCA with $R = 30$, and standardization of PC coordinates. scGen uses its own VAE on the same gene panel; CellOT uses the shared PCA representation rather than its original autoencoder. Hagai uses 9,035 common one-to-one orthologs before fold-specific HVG selection. Source and target-control splits are at the cell level, with sample stratification where documented. These experiments do not test held-out donors or biological replicates.

Comparators and selection. We compare PerturbSB with identity, a source-derived global mean shift, official scGen 2.1.0, and official CellOT. Target treated data do not enter mean-shift estimation or model selection. The principal PerturbSB configuration uses W_2^2/R and $\lambda = 10^{-4}$. Validation determines checkpoints; the pilot selected the energy weight from $\{10^{-4}, 10^{-3}\}$, after which the small benchmark kept it fixed. Training budgets and documented implementation limits are given in Appendix C.

Evaluation and uncertainty. Response metrics compare predicted and observed treated-minus-control mean expression: Pearson-Δ measures direction, while ΔMAE and ΔRMSE measure magnitude error. Identity has zero predicted response, so its correlation is undefined. Distribution metrics are RBF MMD^2 with a training-derived median bandwidth and exact W_2^2/R in standardized PCA space. Centered metrics remove each population mean to assess shape. Metrics are interpreted within each dataset, not pooled across their different feature spaces. Evaluation uses repeated subsampling, with matched test subsets across methods. Where available, neural-method uncertainty is across three training runs after averaging five evaluation repeats; deterministic baselines instead vary across sampling repeats. The early Kang scGen wrapper did not apply its seed argument, making most nominal repeats identical. We therefore do not interpret its near-zero reported variation as stability or claim significance from these repeats.

4.2 HELD-OUT CELL TYPES IN KANG AND HABER

Table 1 presents all five evaluation metrics. In Kang, PerturbSB has the highest Pearson-Δ (0.9095), lowest ΔMAE (0.0421) and ΔRMSE (0.1230), and lowest W_2^2/R (0.6047). Its MMD^2 is close to, but higher than, CellOT's (0.0998 versus 0.0920). The predefined CD4 T-cell diagnostic is consistent with the aggregate response result: all ten selected interferon-response genes have the correct predicted induction direction. Figure 3 shows this agreement at the per-gene level on the CD4 target, where PerturbSB's predicted responses concentrate along the diagonal while the global shift and the neural baselines scatter more widely (identity, which predicts no response, is flat at zero). Predicted ISG15 and IFIT1 effects (2.09 and 1.33) are smaller than their reconstructed observed effects (2.41 and 1.55); Appendix E.1 gives the scope of this comparison.

In Haber, PerturbSB again has the highest macro response correlation (0.6667) and the lowest ΔMAE (0.0522) and ΔRMSE (0.1093). The correlation advantage over CellOT is small (0.0040); CellOT has slightly lower transport error (0.6021 versus 0.6029), and the global shift has the lowest MMD^2 (0.0509). Per-target correlations vary: CellOT leads on Tuft, while PerturbSB leads on Enterocyte (Appendix E). In the Tuft gene diagnostic, PerturbSB misses the observed downregulation of Dclk1 and Pou2f3. The favorable macro averages therefore do not imply uniformly accurate cell-type markers.

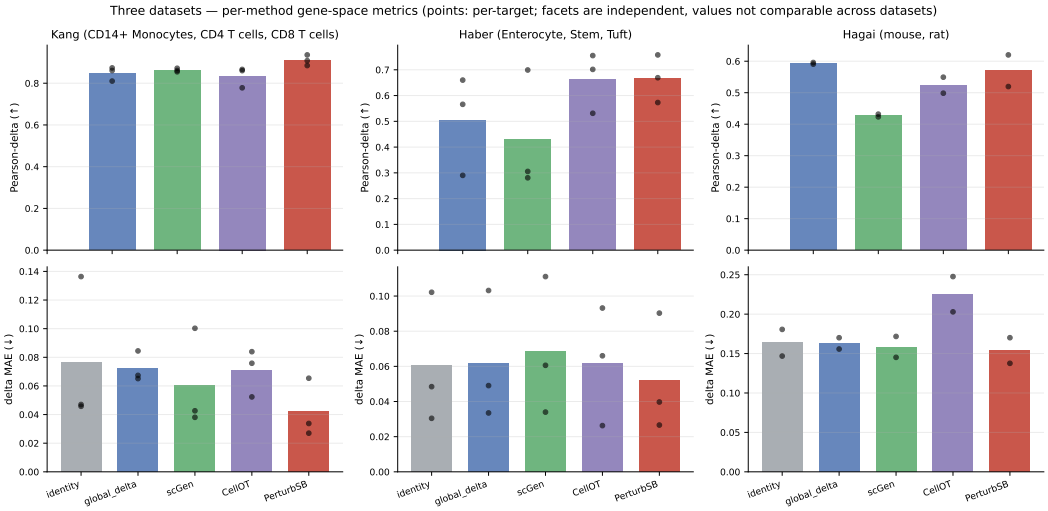

Figure 2: Response prediction across Kang, Haber, and Hagai. Bars summarize targets within each dataset and points show individual targets. Scales differ across datasets; points are not training seeds. Numerical comparisons are given in Table 1.

Table 1: Response prediction on held-out cell types (Kang, Haber) and species (Hagai). Kang and Haber are equal-weight macro means across three targets; Hagai macro averages mouse and rat, which are also shown separately. Lower is better except Pearson-Δ. Bold marks the best reported mean within a block, including ties, without implying statistical significance. Identity has undefined response correlation. Metrics should not be compared numerically across datasets.

Dataset / target	Method	MMD^2 ↓	Pearson-Δ ↑	ΔMAE↓	ΔRMSE↓	W_2^2/R ↓
Kang	Identity	0.2238	—	0.0764	0.2831	1.0303
	Global shift	0.1608	0.8482	0.0723	0.1987	0.8134
	scGen	0.1696	0.8624	0.0603	0.1709	0.7428
	CellOT	**0.0920**	0.8348	0.0707	0.1969	0.7270
	PerturbSB	0.0998	**0.9095**	**0.0421**	**0.1230**	**0.6047**
Haber	Identity	0.0744	—	0.0603	0.1387	0.7576
	Global shift	**0.0509**	0.5054	0.0619	0.1239	0.7168
	scGen	0.1924	0.4285	0.0685	0.1370	0.9651
	CellOT	0.0820	0.6627	0.0618	0.1242	**0.6021**
	PerturbSB	0.0617	**0.6667**	**0.0522**	**0.1093**	0.6029
Hagai / mouse	Identity	0.229	—	0.181	0.410	1.161
	Global shift	**0.180**	**0.591**	**0.170**	**0.329**	1.020
	scGen	0.245	0.432	0.172	0.376	**0.945**
	CellOT	0.206	0.549	0.203	0.403	1.107
	PerturbSB	0.250	0.520	**0.170**	0.375	0.963
Hagai / rat	Identity	**0.125**	—	0.147	0.315	0.908
	Global shift	0.132	0.596	0.156	0.266	0.930
	scGen	0.146	0.423	0.145	0.304	0.627
	CellOT	0.166	0.498	0.248	0.436	0.819
	PerturbSB	**0.125**	**0.620**	**0.138**	**0.256**	**0.587**
Hagai / macro	Identity	0.177	—	0.164	0.362	1.035
	Global shift	**0.156**	**0.5933**	0.163	**0.298**	0.975
	scGen	0.195	0.4276	0.158	0.340	0.786
	CellOT	0.186	0.5239	0.225	0.419	0.963
	PerturbSB	0.188	0.5698	**0.154**	0.316	**0.775**

4.3 Cross-species transfer in Hagai

The two target species lead to different rankings. On rat, PerturbSB has the best reported Pearson-Δ (0.620), ΔMAE (0.138), ΔRMSE (0.256), and W_2^2/R (0.587); its MMD^2 (0.125) ties identity at the reported precision. On mouse, it ties the global shift for the lowest ΔMAE (0.170) but trails it in correlation, MMD^2, and ΔRMSE. The global shift also has the strongest Hagai macro correlation

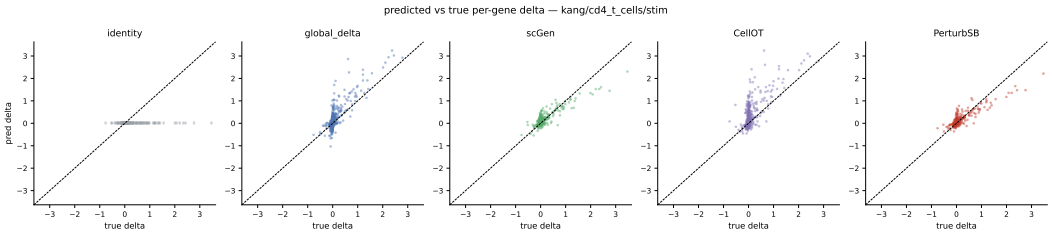

Figure 3: Predicted versus true per-gene response (Δ) for the Kang CD4 T-cell target under IFN-β stimulation. Each panel is one method and each point is a gene; the diagonal marks perfect agreement. PerturbSB's points track the diagonal most tightly, consistent with its highest Pearson-Δ (0.9095).

(0.5933), MMD^2 (0.156), and ΔRMSE (0.298), whereas PerturbSB leads in macro ΔMAE (0.154) and W_2^2/R (0.775). Three training runs and two held-out species do not establish a general cross-species advantage.

Better response means do not necessarily preserve population shape. On mouse, PerturbSB has a predicted-to-observed variance ratio of 0.239 and centered MMD^2 of 0.0211, compared with 1.283 and 0.0034 for CellOT. On rat the corresponding variance ratios are 0.643 and 1.004. Despite its lower centered transport error, PerturbSB contracts variation more strongly than CellOT (Appendix D).

4.4 Endpoint geometry and energy have distinct effects

The completed 2×2 ablation varies full-dimensional versus sliced Wasserstein matching and independently includes or removes energy with $\lambda = 10^{-4}$. It covers three held-out targets and three training seeds in each of Kang and Haber (72 configurations, including reused main runs). Sliced matching uses 64 directions. With no energy, changing from sliced to full matching raises Pearson-Δ from 0.884 to 0.912 on Kang and from 0.640 to 0.670 on Haber. It lowers Kang MMD^2 from 0.136 to 0.105 but raises Haber MMD^2 from 0.0498 to 0.0715.

With full matching, adding energy changes Pearson-Δ from 0.912 to 0.910 on Kang and from 0.670 to 0.667 on Haber, while MMD^2 falls from 0.105 to 0.0998 and from 0.0715 to 0.0617, respectively. Thus endpoint geometry changes the correlation–distribution tradeoff, and energy has a small, metric-dependent effect in these tasks. The ablation does not isolate the effect of the stochastic process or establish a uniquely SB-driven improvement.

5 Discussion and limitations

The benchmarks separate two questions that are often conflated: whether a model predicts the average transcriptional response and whether it reproduces the treated population. On Kang and Haber, PerturbSB reduces mean-response errors and attains the highest response correlations, but other methods lead on at least one distribution metric. The cross-species results sharpen this distinction. PerturbSB performs well for the rat target, whereas a source-derived global shift is stronger on most mouse metrics and on several Hagai macro metrics. Thus, the evidence supports useful context-dependent prediction, not a uniform advantage of stochastic transport over a simple response baseline. The close Haber comparisons, in particular, are too small to establish a reliable ordering from the available runs.

The mouse results also identify a failure mode that aggregate response scores can hide. Although its mean-response error ties the best reported value, PerturbSB captures substantially less treated-cell variation than observed. A stochastic transition therefore does not, by itself, ensure faithful cellular heterogeneity. Since training matches the perturbed endpoint distribution rather than individual cells or observed intermediate states, its simulated paths should be interpreted as a predictive construction, not as identified biological trajectories. Evaluating centered distribution metrics alongside response metrics is necessary to expose this distinction.

The ablation constrains the methodological interpretation. Full-dimensional endpoint matching improves response correlation relative to sliced matching in both Kang and Haber, while its effect on MMD^2 changes sign across datasets. Adding energy to full matching modestly improves MMD^2 but does not improve correlation in either benchmark. These comparisons establish neither that the reference drift nor that stochasticity causes the gains. The path-KL identity characterizes the relative-energy term for a fixed reference; it does not make the learned, softly matched process an exact Schrödinger bridge. A direct comparison with an endpoint-only predictor and controlled removal of the reference drift and diffusion would test the value of the process model more specifically.

The scope of transfer remains limited by evaluation design. Cell-level splits do not test generalization to unseen donors or biological replicates, and three nominal training runs do not quantify biological uncertainty; the Kang scGen seed issue further limits run-to-run comparisons. Hagai is evaluated on only two held-out species, with incomplete ortholog identifier matching for rat and pig, while PCA may discard response components. Replicate- or donor-held-out evaluations, sensitivity to ortholog and feature selection, and explicit checks of distributional spread would provide stronger evidence of biological transfer. The final training code and resolved configurations are also needed to verify the implementation details distinguished in Appendix C.

6 Conclusion

We studied how to transfer a known perturbation to cellular contexts for which only untreated cells are available. PerturbSB addresses this problem with a stochastic transition that separates a shared reference drift from a state-dependent perturbation correction. It learns from unpaired populations by matching treated endpoint distributions and regularizing the correction relative to the reference process. Evaluations across held-out cell types and species show competitive response prediction, including improvements over established baselines in several settings. The ablations further show that endpoint matching and path regularization affect predictive behavior differently. Together, these results establish control-informed stochastic transitions as a useful framework for perturbation transfer and underscore the importance of evaluating both response effects and treated-population distributions.

Reproducibility Statement

Appendix A documents data sources, preprocessing, data visibility, and evaluation metrics; Appendix B states the assumptions behind the path-energy identity. Appendix C summarizes the experimental settings recorded in the reports and identifies implementation details that require checking against the supplied code. The results in this manuscript were transcribed from the experimental reports rather than independently rerun during manuscript preparation.

Ethics Statement

This work analyzes previously collected single-cell datasets and reports computational response predictions. These predictions are not validated treatment recommendations or individual-cell causal measurements. Cell-type, donor, and species coverage limits their applicability. Human-derived data should be used under the original access and consent conditions; the manuscript does not introduce new participant recruitment or release individual-level records.

AI Use Statement

Generative AI assisted with manuscript organization and drafting, translation of experimental reports, literature lookup, interpretation and consistency checking of reported results, presentation of the mathematical formulation, LaTeX preparation, and figure-formatting code. Quantitative results were transcribed from existing experimental reports; no new model-training result was generated during this writing process. The authors are responsible for verifying the final text, references, mathematical claims, and agreement with the experimental implementation. This statement describes assistance in preparing this manuscript and does not certify the provenance of earlier experimental code.

REFERENCES

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*, 20:1759–1768, 2023. doi: 10.1038/s41592-023-01969-x. URL https://doi.org/10.1038/s41592-023-01969-x.

Changxi Chi, Yufei Huang, Jun Xia, Jiangbin Zheng, Yunfan Liu, Zelin Zang, and Stan Z. Li. Departures: Distributional transport for single-cell perturbation prediction with neural Schrödinger bridges. *arXiv preprint arXiv:2511.13124*, 2025. URL https://arxiv.org/abs/2511.13124.

Kilian Fatras, Younes Zine, Rémi Flamary, Rémi Gribonval, and Nicolas Courty. Learning with minibatch Wasserstein: asymptotic and gradient properties. In *International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pp. 2131–2141, 2020. URL https://proceedings.mlr.press/v108/fatras20a.html.

Jhanvi Garg, Xianyang Zhang, and Quan Zhou. Soft-constrained Schrödinger bridge: a stochastic control approach. In *International Conference on Artificial Intelligence and Statistics*, volume 238 of *Proceedings of Machine Learning Research*, pp. 4429–4437, 2024. URL https://proceedings.mlr.press/v238/garg24a.html.

Adam L. Haber, Moshe Biton, Noga Rogel, Rebecca H. Herbst, Karthik Shekhar, Christopher Smillie, Grace Burgin, Toni M. Delorey, Michael R. Howitt, Yarden Katz, et al. A single-cell survey of the small intestinal epithelium. *Nature*, 551:333–339, 2017. doi: 10.1038/nature24489. URL https://doi.org/10.1038/nature24489.

Tzachi Hagai, Xi Chen, Ricardo J. Miragaia, Raghd Rostom, Tomás Gomes, Natalia Kunowska, Johan Henriksson, Jong-Eun Park, Valentina Proserpio, Giacomo Donati, et al. Gene expression variability across cells and species shapes innate immunity. *Nature*, 563:197–202, 2018. doi: 10.1038/s41586-018-0657-2. URL https://doi.org/10.1038/s41586-018-0657-2.

Hyun Min Kang, Meena Subramaniam, Sasha Targ, Michelle Nguyen, Lenka Maliskova, Elizabeth McCarthy, Eunice Wan, Simon Wong, Lauren Byrnes, Cristina M. Lanata, Rachel E. Gate, Sara Mostafavi, Alexander Marson, Noah Zaitlen, Lindsey A. Criswell, and Chun Jimmie Ye. Multiplexed droplet single-cell RNA-sequencing using natural genetic variation. *Nature Biotechnology*, 36:89–94, 2018. doi: 10.1038/nbt.4042. URL https://doi.org/10.1038/nbt.4042.

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

Yiqing Luo, Lin Liu, Yaxin Fu, Yi Deng, and Lin Tang. Prediction of single-cell perturbation response based on direction-constrained diffusion Schrödinger bridge. *Journal of Biomedical Informatics*, 170:104915, 2025. doi: 10.1016/j.jbi.2025.104915. URL https://doi.org/10.1016/j.jbi.2025.104915.

Alexander Y. Tong, Nikolay Malkin, Kilian Fatras, Lazar Atanackovic, Yanlei Zhang, Guillaume Huguet, Guy Wolf, and Yoshua Bengio. Simulation-free Schrödinger bridges via score and flow matching. In *International Conference on Artificial Intelligence and Statistics*, volume 238 of *Proceedings of Machine Learning Research*, pp. 1279–1287, 2024. URL https://proceedings.mlr.press/v238/tong24a.html.