

Departures: Distributional Transport for Single-Cell Perturbation Prediction with Neural Schrödinger Bridges

Changxi Chi^{1,2}, Yufei Huang^{1,2}, Jun Xia³, Jiangbin Zheng^{1,2}, Yunfan Liu^{1,2}, Zelin Zang⁴ and Stan Z. Li²

¹Zhejiang University

²Westlake University

³The Hong Kong University of Science and Technology (Guangzhou)

⁴Centre for Artificial Intelligence and Robotics Hong Kong Institute of Science & Innovation, Chinese Academy of Sciences
{12563057, huangyufei}@zju.edu.cn

Abstract

Predicting single-cell perturbation outcomes directly advances gene function analysis and facilitates drug candidate selection, making it a key driver of both basic and translational biomedical research. However, a major bottleneck in this task is the unpaired nature of single-cell data, as the same cell cannot be observed both before and after perturbation due to the destructive nature of sequencing. Although some neural generative transport models attempt to tackle unpaired single-cell perturbation data, they either lack explicit conditioning or depend on prior spaces for indirect distribution alignment, limiting precise perturbation modeling. In this work, we approximate Schrödinger Bridge (SB), which defines stochastic dynamic mappings recovering the entropy-regularized optimal transport (OT), to directly align the distributions of control and perturbed single-cell populations across different perturbation conditions. Unlike prior SB approximations that rely on bidirectional modeling to infer optimal source-target sample coupling, we leverage Minibatch-OT based pairing to avoid such bidirectional inference and the associated ill-posedness of defining the reverse process. This pairing directly guides bridge learning, yielding a scalable approximation to the SB. We approximate two SB models, one modeling discrete gene activation states and the other continuous expression distributions. Joint training enables accurate perturbation modeling and captures single-cell heterogeneity. Experiments on public genetic and drug perturbation datasets show that our model effectively captures heterogeneous single-cell responses and achieves state-of-the-art performance.

Code — <https://github.com/ChangxiChi/Departures>

Introduction

Genetic perturbation and drug perturbation constitute the two primary avenues for studying and modeling single-cell responses to perturbations. Genetic perturbations, predominantly based on CRISPR technologies (Barrangou and Doudna 2016; Lino et al. 2018), target specific genes to reveal their functional roles, enabling systematic dissection of gene regulatory mechanisms at single-cell resolution. Drug perturbations, by contrast, modulate molecular pathways

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: Single-cell perturbation data are unpaired as RNA-seq is destructive.

through small-molecule compounds (Peidli et al. 2024), offering complementary insights into cellular responses and pharmacological effects. Together, these two perturbation strategies form the foundation of single-cell perturbation studies, providing mechanistic and translational perspectives on cell transitions under diverse experimental conditions. However, due to the high cost and limited scalability of single-cell experiments, it is impractical to measure all combinations of perturbations and cell types. This limitation highlights the need for computational models that can predict cellular responses under unseen perturbation conditions.

A major challenge in this task is that single-cell perturbation data are inherently unpaired (Bunne et al. 2023; Chi et al. 2025a). This arises from the destructive nature of RNA sequencing, which requires cell lysis to release RNA content (Mortazavi et al. 2008), as illustrated in Fig. 1. As a result, it is impossible to observe the same cell under both control and perturbed conditions. Some studies (Bunne et al. 2023; Cao et al. 2024) attempt to address this unpaired nature, but their use of unconditional models limits generalization to unseen perturbation conditions. In contrast, (Chi et al. 2025a) incorporates gene regulatory network (GRN) priors to model the effects of perturbations, providing a biologically grounded representation of cellular responses. It then aligns control and perturbed cell populations using dual conditional diffusion models with a shared prior space, enabling accurate prediction under unseen perturbation conditions. However, by aligning distributions indirectly through the prior space, it bypasses the underlying physical energy landscape, leading to transition paths that deviate from energy-optimal trajectories and fail to capture ideal optimal transport (OT).

In this paper, we propose **Departures (Distributional Transport for Single-Cell Perturbation Prediction with Neural Schrödinger Bridges)**, a generative framework designed to predict cellular gene expression under diverse genetic and molecular perturbation conditions. **By directly aligning the distributions of control and perturbed cells, the framework captures population-level transitions without relying on potentially noisy explicit cell pairing or restrictive latent space alignment. This method reduces information loss, improves robustness to biological variability and noise, and provides a principled, scalable solution to the challenges posed by unpaired single-cell perturbation data.**

While classical Schrödinger Bridge methods typically rely on alternating forward and backward updates (Rüschendorf 1995; Shi et al. 2023) to iteratively refine the source-target node pairing between marginal distributions, the backward process is often ill-defined in conditional settings (Fig. 3) (Tamir, Trapp, and Solin 2023; Tang et al. 2024). Moreover, training two coupled models significantly increases computational cost and reduces scalability. To address this, we build upon Minibatch OT (Fratras et al. 2021; Tong et al. 2023) to directly enhance source-target sample coupling during training, without requiring alternating bidirectional model updates for pairing consistency. Built upon this alignment, we adopt bridge matching (Liu et al. 2022; Peluchetti 2023) to learn the *Markovian projection* (Gyöngy 1986), yielding a tractable approximation to the SB between control and perturbed distributions. To faithfully capture single-cell perturbation outcomes, we learn two SB models, the first models discrete gene activation states, and the second learns continuous expression dynamics. Joint training of these bridges enhances both biological fidelity and generative robustness.

The main contributions of our work are as follows:

- We introduce **Departures**, a Schrödinger Bridge-based generative model that directly aligns the distributions of control and perturbed cells. To better approximate the Schrödinger Bridge, we use minibatch OT to compute source-target sample pairings that guide bridge matching training, eliminating the need for bidirectional iterative updates.
- To model single-cell responses to perturbations, we design two Schrödinger bridge models, one for discrete gene activation and one for continuous expression, which are trained jointly to improve fidelity and robustness.
- We demonstrate the superiority of **Departures** over existing methods on publicly available genetic and molecular perturbation datasets.

Related Works and Preliminaries

Existing Perturbation Prediction Model

There are numerous methods designed to predict single-cell perturbation responses, encompassing both genetic and drug-induced perturbations. These methods typically employ either generative models (Lotfollahi, Wolf, and Theis 2019; Hetzel et al. 2022; Wu et al. 2022; Bereket and

Karaletsos 2024; Cui et al. 2024) or regression-based approaches (Roohani, Huang, and Leskovec 2024; He et al. 2024; Chi et al. 2025b). However, many of these models overlook the inherently unpaired nature of single-cell perturbation data and fail to explicitly model the relationship between control and perturbed samples. While a few studies have addressed the unpaired nature of the data (Bunne et al. 2023; Cao et al. 2024), their reliance on unconditional models limits their ability to generalize to unseen perturbation settings. (Chi et al. 2025a) addresses this by incorporating gene regulatory network (GRN) priors to model perturbation effects, and aligning control and perturbed cell populations via dual conditional diffusion models that share a common prior space. Additionally, a regression-based mask model is introduced to predict gene expression status, thereby improving generation quality. However, this implicit alignment bypasses the true physical energy landscape, resulting in transitions that deviate from energy-optimal trajectories and fail to recover the ideal OT. Moreover, the mask model is trained separately, leading to suboptimal efficiency.

Schrödinger Bridges

The path measure space $\mathcal{P}(\mathcal{C})$ refers to the space of probability measures over continuous trajectories $x : [0, T] \rightarrow \mathbb{R}^d$, that is, $\mathcal{P}(\mathcal{C}) = \mathcal{P}(C([0, T], \mathbb{R}^d))$. The Schrödinger Bridge (SB) problem (Schrödinger 1932) seeks a stochastic process $\mathbb{P}^* \in \mathcal{P}(\mathcal{C})$ that evolves between two given marginal distributions π_0 and π_T , while staying close to a reference process $\mathbb{Q} \in \mathcal{P}(\mathcal{C})$. It is formulated as:

$$\mathbb{P}^* = \operatorname{argmin}_{\mathbb{P} \in \mathcal{P}(\mathcal{C})} \{ \text{KL}(\mathbb{P} | \mathbb{Q}) : \mathbb{P}_0 = \pi_0, \mathbb{P}_T = \pi_T \} \quad (1)$$

The SB is an entropy-regularized dynamic Optimal Transport problem. Its path distribution implicitly defines an optimal pairing between nodes at the endpoints.

A common approach to solving this problem is Iterative Proportional Fitting (IPF, (Rüschendorf 1995)), which alternates KL projections onto path measures with fixed initial or terminal marginals. However, it enforces marginal consistency without explicitly modeling the bridge dynamics.

In contrast, following the formulation in (Shi et al. 2023), the reference process \mathbb{Q} is defined by a diffusion $d\mathbf{X}_t = f_t(\mathbf{X}_t)dt + \sigma_t d\mathbf{B}_t$, with $\mathbf{X}_0 \in \mathbb{Q}_0 = \pi_0$ and $(\mathbf{B}_t)_{t \in [0, T]}$ is Brownian motion. To enforce \mathbb{Q}_T also matches π_T , we construct a diffusion bridge $\mathbb{Q}_{|0, T}(\cdot | x_0, x_T)$ using Doob’s h -transform (Rogers and Williams 2000), such that it is conditioned to start at x_0 and end at x_T :

$$d\mathbf{X}_t^{0, T} = \{ f_t(\mathbf{X}_t^{0, T}) + \sigma_t^2 \nabla \log \mathbb{Q}_{T|t}(x_T | \mathbf{X}_t^{0, T}) \} dt + \sigma_t d\mathbf{B}_t \quad (2)$$

where $\mathbf{X}_0^{0, T} = x_0$. Given $\Pi_{0, T} = \pi_0 \otimes \pi_T$ and $\Pi = \Pi_{0, T} \mathbb{Q}_{|0, T}$, our goal is to find a Markov process $d\mathbf{Y}_t = \{ f_t(\mathbf{Y}_t) + v_t(\mathbf{Y}_t) \} dt + \sigma_t d\mathbf{B}_t$. The implicit analytical solution of this problem is given by:

$$v_t^*(x_t) = \sigma_t^2 \mathbb{E}_{\Pi_{T|t}} [\nabla \log \mathbb{Q}_{T|t}(\mathbf{X}_T | \mathbf{X}_t) | \mathbf{X}_t = x_t] \quad (3)$$

which corresponds to the *Markovian projection* (Gyöngy 1986). When $f_t = 0$ and $\sigma_t = \sigma$, $\mathbb{Q}_{|0, T}$ is a Brownian Bridge, and we have:

$$\mathbf{X}_t^{0, T} = \frac{t}{T} x_T + (1 - \frac{t}{T}) x_0 + \sigma (\mathbf{B}_t - \frac{t}{T} \mathbf{B}_T) \quad (4)$$

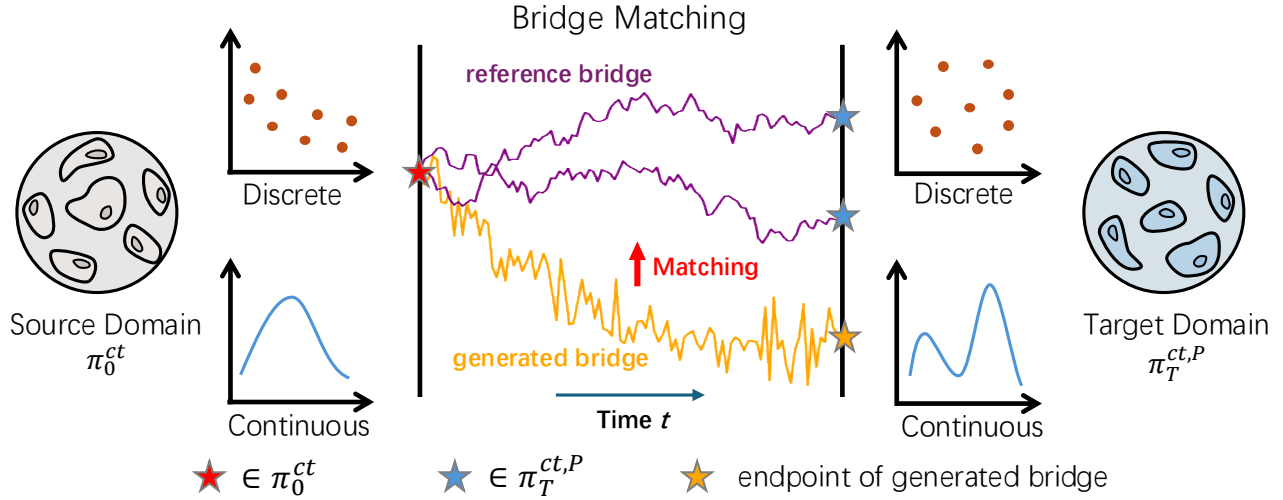


Figure 2: Overview of **Departures**. Conditioned on cell type ct , the model learns the distributional transition of single-cell gene expression profiles from the control population π_0^{ct} to perturbed population $\pi_T^{ct,P}$. The problem is formulated by decoupling into two components: (1) modeling the distributional shift of gene expression levels before and after perturbation (**Continuous**), and (2) modeling the distributional shift of gene activation status induced by perturbation (**Discrete**).

$$d\mathbf{X}_t^{0,T} = \{(x_T - \mathbf{X}_t^{0,T})/(T-t)\}dt + \sigma d\mathbf{B}_t \quad (5)$$

with $(\mathbf{B}_t - \frac{t}{T}\mathbf{B}_T) \sim \mathcal{N}(0, t(1 - \frac{t}{T})\text{Id})$, and \mathbf{B}_t is a standard Brownian motion.

The objective of the model v_θ is to predict v_t^* . Specifically, the bridge matching (Peluchetti 2023; Liu et al. 2022) operation learns the *Markovian projection* by optimizing the following objective:

$$\mathbb{E}_{\Pi_{t,T}}[\|(\mathbf{X}_T - \mathbf{X}_t)/(T-t) - v_\theta(t, \mathbf{X}_t)\|^2] \quad (6)$$

To generate sample paths of the continuous bridge process, we implement the dynamics using the Euler–Maruyama method (Heng et al. 2025), a standard numerical scheme for approximating solutions to stochastic differential equations (SDEs).

Based on the above, Iterative Markovian Fitting (IMF, (Shi et al. 2023)) is an algorithm designed to approximate Schrödinger bridges by alternating between *Markovian projection* and *Reciprocal projection* (Jamison 1970). It employs two time-symmetric models to match the marginals at both ends of the path, iteratively refining their source-target sample coupling to approximate SB. However, learning both forward and backward models is challenging due to the inherently one-directional nature of perturbations.

Methodology

We introduce the proposed model **Departures** in this section. The overview is shown in Fig. 2. Specifically, **Departures** learns the distributional transition mapping from control to perturbed samples, modeling both gene expression levels (**Continuous**) and gene activation states (**Discrete**) via bridge matching. To approximate the SB without requiring two time-reversal models, the node-level pairing is optimized using Minibatch OT (Fratras et al. 2021; Tong et al. 2023).

Problem Statement

In single-cell perturbation prediction, the objective is to infer the gene expression profile of a specific cell type ct under a given perturbation condition P . These conditions may stem from either genetic interventions or treatments with small-molecule compounds. For genetic perturbations, the condition is typically specified by the names of targeted genes, corresponding to gene knockout experiments. In contrast, small-molecule perturbations are characterized by the drug’s chemical structure along with its administered dosage. In our setting, instead of estimating a population-level summary (e.g., mean), the model generates a batch of samples under the given condition, aiming to represent the underlying distribution.

Distribution Transfer of Continuous Gene Expression

Suppose under cell type ct and perturbation condition P , we obtain a pairing $(x_0, x_T) \sim \gamma, \gamma \in \Pi(\pi_0^{ct}, \pi_T^{ct,P})$, where each $x_0, x_T \in \mathbb{R}^N$ denotes a gene expression vector of a single cell, with x_0 sampled from the control distribution π_0 , and x_T sampled from the perturbed distribution π_T^P . A detailed discussion of the joint distribution γ is provided later.

To learn the transition between distributions, it is necessary to build bridges between x_0 and x_t . Specifically, we construct a diffusion bridge using the Doob’s h -transform (Rogers and Williams 2000), following the assumption in (Liu et al. 2022; Shi et al. 2023) that $f_t = 0$ and $\sigma_t = \sigma$. Referring to Eq. 4, we have:

$$x_t^{0,T} = \frac{t}{T}x_T + (1 - \frac{t}{T})x_0 + \sigma(\mathbf{B}_t - \frac{t}{T}\mathbf{B}_T) \quad (7)$$

where $t \in [0, T]$, $\mathbf{B}_t \sim \mathcal{N}(0, \text{Id})$, and $(\mathbf{B}_t - \frac{t}{T}\mathbf{B}_T) \sim \mathcal{N}(0, t(1 - \frac{t}{T})\text{Id})$.

The continuous bridge dynamics between the initial state and the fixed terminal point x_T are described by the stochastic differential equation (SDE):

$$dx_t^{0:T} = v_t dt + \sigma d\mathbf{B}_t, \quad d\mathbf{B}_t \sim \mathcal{N}(0, \text{Id}) \quad (8)$$

where $v_t = (x_T - x_t^{0:T})/(T - t)$, σ denotes noise scale, and $z \sim \mathcal{N}(0, 1)$. Our objective is to train the model to learn the time-dependent drift term of this SDE, whose analytical solution is given by Eq. 3, as guided by the v_t prediction loss function in Eq. 6. In practice, sample paths of this SDE are generated by discretizing it using the Euler–Maruyama method (Heng et al. 2025).

Rather than regressing the velocity field v_t directly, we instead predict the endpoint x_T , which can be equivalently transformed to the drift v_t . This alternative formulation leads to more stable optimization and reduced numerical error (Huang et al. 2024).

Moreover, due to the sparsity of gene expression data, directly computing the loss over all genes can lead to mode collapse (Chi et al. 2025a). To address this, we compute the loss only over genes with non-zero expression after perturbation. A separate module for predicting gene expression status will be introduced in a later section. As a result, the model x_θ focuses on expressed genes during training. The final objective function is defined as follows:

$$\mathcal{L}_{cont} = \mathbb{E}_{t, (x_0, x_T) \sim \gamma} \left[\frac{\|d_T \odot (x_T - x_\theta(t, x_t^{0:T}, ct, P))\|^2}{\sum_i d_{T,i}} \right] \quad (9)$$

where \odot denotes Hadamard Product, $d_T \in \{0, 1\}^N$ is a binary mask vector defined as:

$$d_{T,i} = \begin{cases} 1, & \text{if } x_{T,i} \neq 0 \\ 0, & \text{otherwise} \end{cases} \quad (10)$$

Distribution Transfer of Discrete Gene Status

As previously mentioned, gene expression data are high-dimensional and sparse. To prevent the model from treating all genes equally and thus causing mode collapse, we apply a mask to the loss in Eq. 9 based on the true gene expression status at the endpoint. Therefore, it is necessary to train a separate module capable of accurately predicting whether a gene is expressed or not.

Given a pairing $(x_0, x_T) \sim \gamma$, we convert it into discrete labels $(d_0, d_T) \sim \gamma_d$ using Eq. 10. However, in the discrete setting, directly applying the interpolation in Eq. 7 to the discrete labels is not meaningful, as the labels represent categorical states rather than continuous values. We define the intermediate states of the bridge process as stochastic mixtures of the two endpoints according to the following distribution:

$$d_{t,i} = \begin{cases} x_{T,i} & \text{with prob} = \kappa(t) \\ x_{0,i} & \text{with prob} = (1 - \kappa(t)) \end{cases} \quad (11)$$

where we set $\kappa(t) = \frac{t}{T}$ is monotonic increasing with $\kappa(0) = 0$ and $\kappa(T) = 1$. This process captures the discrete interpolation from the initial state to the terminal state by randomly selecting each coordinate from x_0 or x_T according to $\kappa(t)$.




Figure 3: Perturbations induce natural, directional transitions from control to treated states. In contrast, reconstructing control states from perturbed ones is ill-posed, and training a backward model adds significant computational overhead.

Unlike diffusion bridges in continuous spaces governed by SDE, the discrete bridge corresponds to a continuous-time Markov chains (CTMC, (Campbell et al. 2024; Gat et al. 2024)) on discrete space characterized by the transition probability:

$$p_{t+h|t}(d_{t+h}|d_t) = \delta(d_{t+h}, d_t) + hu_t(d_{t+h}, d_t) + o(h) \quad (12)$$

where $d_t, d_{t+h} \in \{0, 1\}^N$, δ is Kronecker delta function, and $u_t(d_{t+h}, d_t)$ represents transition rate from d_t to d_{t+h} . Our goal is to learn the transition rate function u_t , which fully characterizes the dynamics of the CTMC.

As shown in (Lipman et al. 2024), this can be achieved by learning the conditional distribution $p_{T|t}$ as:

$$u_t(d_{t+h}, d_t) = \mathbb{E}_{d_T} \left[\frac{\kappa(t)}{1 - \kappa(t)} [\delta(d_{t+h}, d_T) - \delta(d_{t+h}, d_t)] \times p_{T|t}(d_T|d_t) \right] \quad (13)$$

with $p_{T|t}(d_T|d_t) = \mathbb{E}_{(d_0, d_T) \sim \gamma_d} p_{0,1|t}(d_0, d_T|d_t)$. We aim to approximate this conditional posterior $p_{T|t}(\cdot|d_t)$ using a parameterized neural network.

Intuitively, learning the posterior transition probabilities $p_{T|t}(\cdot|d_t)$ is correspond to the conditional distribution of the bridge’s terminal state given the intermediate state. The objective function is as follows:

$$\mathcal{L}_{disc} = \mathbb{E}_{t, (d_0, d_T) \sim \gamma_d} [D(d_T, d_\theta(t, d_t, ct, P))] \quad (14)$$

where $D(\cdot, \cdot)$ denotes the cross-entropy loss between the true discrete state d_T and the prediction.

Initialization Pairing for Schrödinger Bridge

We now discuss the joint distribution γ . As introduced earlier, IMF(Shi et al. 2023) employs two time-symmetric neural networks to model the forward and backward processes, matching the marginal distributions at both ends of the trajectory through Markovian projection and reciprocal projection. A core design goal of IMF is to reduce error accumulation during node rectification, where the joint distribution γ is iteratively updated. By obtaining higher-quality pairing

between nodes from γ , it subsequently facilitates the learning of the diffusion bridge. Through iterative optimization, the process converges to the true SB.

However, a major obstacle in applying IMF to conditional generation tasks is the need to learn models in both forward and backward directions. The backward process of generating from the perturbed distribution back to the control distribution is ill-defined (Fig. 3), which makes bidirectional training impractical (Tamir, Trapp, and Solin 2023; Tang et al. 2024).

To address this problem, we draw inspiration from Mini-batch OT (Fratras et al. 2021; Tong et al. 2023) to construct local OT pairings by selecting subsets of control samples for all samples under each perturbation condition. Specifically, under perturbation condition P and cell type ct , we obtain a batch of perturbed cell samples $\{x_T^{(i)}\}_{i=1}^B$. Correspondingly, we sample an equal number of control samples $\{x_0^{(i)}\}_{i=1}^B$ from the same cell type ct to serve as the source for pairing. Then we compute an optimal transport (OT) plan π_{batch} between these two sets of samples using the Sinkhorn algorithm (Cuturi 2013), which efficiently solves the entropically regularized OT problem. The OT plan defines a joint distribution over matched pairs $(x_0^{(j)}, x_T^{(i)})$ that approximately minimizes the transport cost within each batch. Based on these pairs, we apply Eq. 10 to derive the corresponding discrete gene activation states $(d_0^{(j)}, d_T^{(i)})$. The continuous and discrete pairings define the endpoints of the bridge, guiding the model to learn *Markovian projections* via bridge matching (Eq. 9, Eq. 14), thereby approximating the Schrödinger Bridge.

Implementation and Generation

Before training, we perform log lp normalization on the gene expression matrix using Scanpy (Wolf, Angerer, and Theis 2018), followed by the selection of the top N highly variable genes (HVGs). In modeling various perturbation conditions, we adopt the approach of (Chi et al. 2025a), which incorporates the Gene Regulation Network (GRN) to model perturbation conditions. During training, we initialize the source-target sample couplings at the beginning of each epoch using the strategy described in the **for Schrödinger bridge Section**. We jointly train the model x_θ predicting changes in continuous gene expression values and the model d_θ predicting changes in discrete gene activation status introduced above, with the following loss function:

$$\mathcal{L} = \mathcal{L}_{cont} + \mathcal{L}_{disc} \quad (15)$$

In predicting the perturbation results, we perform sampling separately from the two models to obtain predictions of gene expression values \hat{x}_T and gene activation states \hat{d}_T . Specifically, when predicting perturbation outcomes of cells of type ct under perturbation condition P , we iteratively generate the results according to the following procedure, as introduced in Eq. 8:

$$dx_t = \frac{(x_\theta(t, x_t, ct, P) - x_t)}{(T - t)} dt + \sigma d\mathbf{B}_t \quad (16)$$

where at time $t = 0$, x_0 denotes a sample drawn from the control group of the same cell type ct , and at time $t = T$, $x_T = \hat{x}_T$ represents the final predicted outcome.

When predicting the discrete activation states of genes, we follow (Lipman et al. 2024) and define the conditional transition probability of moving from state d_t at time t to the next state after a small increment h as follow:

$$p_{t+h|t}(\cdot|d_t) = d_t + h \times \frac{d_\theta(t, d_t, ct, P)}{T - t} \quad (17)$$

where d_0 denotes the discrete initial state representing gene activation of a control sample of cell type ct . The next state d_{t+h} is then sampled from the distribution $p_{t+h|t}(\cdot|d_t)$, and by iteratively applying this sampling procedure over time steps, the final predicted discrete activation state $d_T = \hat{d}_T$ is obtained.

In practice, to generate predictions, we first sample a real control group example x_0 from cell ct and its corresponding discretized state d_0 (Eq. 10) as the respective starting points for the continuous and discrete models. We then perform iterative sampling using Eq. 16 and Eq. 17 to obtain \hat{x}_T and \hat{d}_T . The final predicted outcome \tilde{x}_T is given by:

$$\tilde{x}_T = \hat{x}_T \odot \hat{d}_T \quad (18)$$

where \odot denotes Hadamard Product.

Experiments and Results

Dataset

In the main experiments, we leverage diverse datasets covering both genetic and chemical perturbations to evaluate our method, including the Adamson CRISPR knockout dataset (Adamson et al. 2016) and the sci-Plex3 chemical perturbation dataset (Srivatsan et al. 2020). The Adamson dataset comprises gene expression profiles from 87 distinct single-gene perturbations in a single cell type, while the sci-Plex3 dataset spans 187 chemical perturbations across four dosage levels and three different cell types. Each perturbation condition in both datasets is represented by the average profile of hundreds of single cells, ensuring robust statistical coverage. We consider 5,000 genes in the Adamson dataset and 2,000 genes in sci-Plex3, providing a rich and representative foundation for model training and evaluation. By evaluating our model across different perturbation modalities and biological settings, we demonstrate its broad applicability and effectiveness.

Experiment Settings

We conduct experiments on two widely-used datasets in the single-cell perturbation modeling field, Adamson (Adamson et al. 2016) and SciPlex3 (Srivatsan et al. 2020). For the Adamson dataset, we randomly split the perturbed genes into 70% for training and 30% for testing across all gene perturbation types. All control group cells are included in the training set. For the SciPlex3 dataset, all control group cells are included in the training set. For each experimental condition (defined by a unique combination of drug, dosage, and cell type), we assign the entire group of corresponding cells to the test set with a 30% probability; otherwise, it is

Dataset	Methods	All		DE20		DE40	
		E-distance(\downarrow)	EMD ² (\downarrow)	E-distance(\downarrow)	EMD ² (\downarrow)	E-distance(\downarrow)	EMD ² (\downarrow)
Adamson	GRAPE	0.7905 ± 0.0484	0.0444 ± 0.0024	0.7514 ± 0.0575	0.1528 ± 0.0234	0.7648 ± 0.0565	0.1503 ± 0.0182
	GEARS	0.8721 ± 0.1304	0.0531 ± 0.0027	0.7884 ± 0.1245	0.1298 ± 0.0324	0.7935 ± 0.0544	0.1221 ± 0.0231
	GraphVCI	1.3182 ± 0.9763	0.3026 ± 0.1953	2.4499 ± 0.2446	1.2457 ± 0.5183	2.6327 ± 0.4950	1.0801 ± 0.0866
	scGPT	2.1368 ± 0.0441	0.1724 ± 0.0355	1.2571 ± 0.3373	0.3895 ± 0.1032	1.4484 ± 0.3087	0.3781 ± 0.0866
	Departures (Ours)	0.5955 ± 0.1212	0.0356 ± 0.0053	0.5028 ± 0.1942	0.0953 ± 0.0487	0.5243 ± 0.2243	0.1013 ± 0.0442
sci-Plex3	chemCPA	0.7847 ± 0.1029	0.0838 ± 0.0081	0.4717 ± 0.1571	0.1836 ± 0.0358	0.5008 ± 0.1659	0.1784 ± 0.0261
	CPA	0.9894 ± 0.1336	0.1357 ± 0.0461	0.9737 ± 0.9768	0.3761 ± 0.0667	1.0794 ± 1.1890	0.3856 ± 0.0387
	GraphVCI	0.8393 ± 0.1823	0.0986 ± 0.0108	0.4958 ± 0.1275	0.2016 ± 0.0379	0.5174 ± 0.1347	0.1861 ± 0.0288
	Departures (Ours)	0.5478 ± 0.1461	0.0254 ± 0.0051	0.2764 ± 0.1246	0.0623 ± 0.0317	0.2827 ± 0.1113	0.0565 ± 0.0340

Table 1: Performance comparison on Adamson and sci-Plex3 datasets, evaluated using E-distance and EMD² on all genes, top 20, and top 40 differentially expressed (DE) genes.

assigned to the training set. We select the top 5,000 highly variable genes (HVGs) for the Adamson dataset and the top 2,000 HVGs for the SciPlex3 dataset. All baseline methods are evaluated using the same gene selection to ensure a fair comparison.

The model is trained using the AdamW (Loshchilov 2017) optimizer with a learning rate of 0.001 and a batch size of 64. During inference, we use 50 uniformly spaced time steps in Eq. 16 and Eq. 17 (i.e., a step size of 0.02). We set the noise scale σ in Eq. 16 to 0.2. All our method and its competitors are conducted using two Nvidia A100 GPU.

For evaluation, we follow previous work and adopt distribution-aware metrics to account for the strong heterogeneity observed in single-cell data. Specifically, we use Energy Distance (E-distance) to capture overall distributional alignment by considering both inter-group and intra-group distances, and Earth Mover’s squared Distance (EMD²) to quantify gene-level shifts by measuring the minimal cost to align predicted and true distributions. These metrics together provide a comprehensive and robust assessment of model performance at both the population and gene levels. Detailed computation procedures can be found in the Appendix.

Departures Outperform Existing Methods

To assess the effectiveness of our model in predicting gene expression under perturbations, we compare it against several baseline methods, including GraphVCI (Wu et al. 2022), scGPT (Cui et al. 2024), CPA (Lotfollahi et al. 2023), chemCPA (Hetzl et al. 2022), GEARS (Roohani, Huang, and Leskovec 2024), and GRAPE (Chi et al. 2025b). For regression-based methods (GEARS and GRAPE), we augment their outputs by adding the predicted gene expression changes to different control samples. This promotes diver-

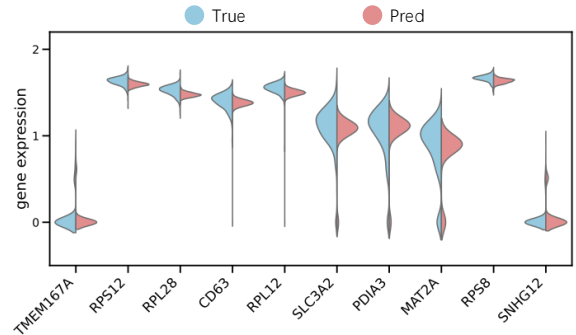


Figure 4: Violin plots comparing predicted and actual expression levels of the top differentially expressed (DE) genes under the TMEM167A knockout condition, which was unseen during training, from the Adamson dataset.

sity in the generated profiles and ensures a more equitable comparison.

Table 1 shows that **Departures** outperforms graphVCI, scGPT, GEARS, GRAPE, which rely on forced pairing of perturbed and unperturbed cells during training. This forced matching can limit the model’s ability to capture true biological variability and cell heterogeneity. Although GEARS and GRAPE inherit control group variability by adding a fixed predicted delta to control samples, this approach does not fully translate the control group’s diversity into the predicted results. Because they do not condition their predictions on the starting point information, their performance is inferior to that of Departures.

Methods such as (Hetzl et al. 2022) and (Lotfollahi et al.




Figure 5: UMAP visualization of predicted and actual gene activation states. The left panel shows results for the unseen HSD17B12 perturbation condition from the Adamson test set. The right panel presents predictions for a held-out perturbation (compound Sodium – dosage 0.001 – cell type MCF7) from the sciplex3 test set.

2023) focus solely on reconstructing perturbed cells without modeling the transitional dynamics from unperturbed states. By relying only on predictions of mean and variance, these approaches oversimplify the perturbation process, limiting their ability to capture the inherent complexity and heterogeneity of cellular responses. In contrast, our model starts from the initial state and guides the state transitions through the learned Schrödinger bridge dynamics. Moreover, as illustrated in Fig. 4, our model accurately captures the distribution of top differentially expressed (DE) genes after perturbation.

	Adamson	sciplex3
PCC_all	0.9910±0.0078	0.9652±0.0690
PCC_DE20	0.9515±0.0685	0.9437±0.0969
PCC_DE40	0.9493±0.0716	0.9520±0.0722

Table 2: Pearson correlation coefficients (PCC) between the predicted gene activation probabilities after perturbation and the true probabilities.

Discrete Model Effectively Captures Gene Activation States

To evaluate the effectiveness of our discrete bridge model in predicting gene activation states after perturbation, we showcase the Pearson correlation coefficients (PCC) between predicted and true gene activation probabilities in Table 2. Specifically, under each perturbation condition, we first predict the activation states of k cells $\{\hat{d}_T^{(i)}\}^k \in \{0, 1\}^{k \times N}$ using Eq. 17, and then compute gene-wise activation probabilities. Similarly, we estimate the ground-truth activation probabilities from the real cell samples under the same condition, and compute the PCC between predicted and true probabilities. Additionally, we visualize the predicted gene activation states of individual cells in Fig. 5. The high PCCs and visualization results show strong alignment between predicted and actual gene activation states.




Figure 6: Ablation study results on Adamson.

Ablation Study

To further evaluate the effectiveness of **Departures**, we compare it with the following methods through an ablation study. 1)**w/o Discrete**: Only uses a single model to capture continuous gene expression, without discrete activation prediction model. 2)**w/o OT**: Node pairing is performed without Minibatch-OT based pairing, relying on random pairing instead. 3)**OT cosine**: Uses cosine similarity as the transport cost when computing the OT-based pairing between unperturbed and perturbed distributions. 4)**OT euclidean**: Uses Euclidean distance as the transport cost when computing the OT-based pairing between unperturbed and perturbed distributions. The results are shown in Fig.6.

Experimental results show that modeling gene activation with a discrete model plays a crucial role in accurately predicting gene expression. It helps the prediction focus on important genes rather than overfitting sparse regions, thereby preventing mode collapse. Besides, OT-based pairing, inspired by the Schrödinger Bridge, outperforms random pairing by learning a probabilistically consistent transport plan that better captures biological transitions.

Conclusion

In this work, we present **Departures**, a generative framework for distribution-level prediction of single-cell responses under various perturbation conditions, via an approximation of the Schrödinger Bridge. Considering that the reverse process of conditional generation under gene perturbation is ill-posed, we employ Minibatch-OT to obtain a relatively better source-target sample coupling. This approach avoids the need to learn bidirectional conditional models with iterative pairing updates as before. Based on this pairing, we learn two *Markovian projections* through bridge matching, one capturing discrete gene activation patterns and the other modeling continuous expression changes. Together, these enhance biological fidelity and generative robustness, enabling an efficient approximation of the Schrödinger Bridge. This approach learns distributional transitions between pre- and post-perturbation cells, effectively addressing the unpaired nature of single-cell perturbation data. By integrating gene regulatory priors and jointly modeling discrete and continuous dynamics, it improves biological fidelity and prediction accuracy, offering a new paradigm for single-cell perturbation modeling.

Acknowledgments

This work was supported by National Science and Technology Major Project (No. 2022ZD0115101), National Natural Science Foundation of China Project (No. 623B2086), National Natural Science Foundation of China Project (No. U21A20427), Project (No. WU2022A009) from the Center of Synthetic Biology and Integrated Bioengineering of Westlake University, and the Zhejiang Province Selected Funding for Postdoctoral Research Projects (ZJ2025113).

References

- Adamson, B.; Norman, T. M.; Jost, M.; Cho, M. Y.; Nuñez, J. K.; Chen, Y.; Villalta, J. E.; Gilbert, L. A.; Horlbeck, M. A.; Hein, M. Y.; et al. 2016. A multiplexed single-cell CRISPR screening platform enables systematic dissection of the unfolded protein response. *Cell*, 167(7): 1867–1882.
- Barrangou, R.; and Doudna, J. A. 2016. Applications of CRISPR technologies in research and beyond. *Nature biotechnology*, 34(9): 933–941.
- Bereket, M.; and Karaletsos, T. 2024. Modelling cellular perturbations with the sparse additive mechanism shift variational autoencoder. *Advances in Neural Information Processing Systems*, 36.
- Bunne, C.; Stark, S. G.; Gut, G.; Del Castillo, J. S.; Levesque, M.; Lehmann, K.-V.; Pelkmans, L.; Krause, A.; and Rätsch, G. 2023. Learning single-cell perturbation responses using neural optimal transport. *Nature methods*, 20(11): 1759–1768.
- Campbell, A.; Yim, J.; Barzilay, R.; Rainforth, T.; and Jaakkola, T. 2024. Generative flows on discrete state-spaces: Enabling multimodal flows with applications to protein co-design. *arXiv preprint arXiv:2402.04997*.
- Cao, Y.; Zhao, X.; Tang, S.; Jiang, Q.; Li, S.; Li, S.; and Chen, S. 2024. scButterfly: a versatile single-cell cross-modality translation method via dual-aligned variational autoencoders. *Nature Communications*, 15(1): 2973.
- Chi, C.; Xia, J.; Huang, Y.; Zhou, J.; Li, S.; Liu, Y.; Yu, C.; and Li, S. Z. 2025a. Unlasting: Unpaired Single-Cell Multi-Perturbation Estimation by Dual Conditional Diffusion Implicit Bridges. *arXiv preprint arXiv:2506.21107*.
- Chi, C.; Xia, J.; Zhou, J.; Cheng, J.; Yu, C.; and Li, S. Z. 2025b. GRAPE: Heterogeneous Graph Representation Learning for Genetic Perturbation with Coding and Non-Coding Biotype. *arXiv preprint arXiv:2505.03853*.
- Cui, H.; Wang, C.; Maan, H.; Pang, K.; Luo, F.; Duan, N.; and Wang, B. 2024. scGPT: toward building a foundation model for single-cell multi-omics using generative AI. *Nature Methods*, 21(8): 1470–1480.
- Cuturi, M. 2013. Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in neural information processing systems*, 26.
- Fatras, K.; Zine, Y.; Majewski, S.; Flamary, R.; Gribonval, R.; and Courty, N. 2021. Minibatch optimal transport distances; analysis and applications. *arXiv preprint arXiv:2101.01792*.
- Gat, I.; Remez, T.; Shaul, N.; Kreuk, F.; Chen, R. T.; Synnaeve, G.; Adi, Y.; and Lipman, Y. 2024. Discrete flow matching. *Advances in Neural Information Processing Systems*, 37: 133345–133385.
- Gyöngy, I. 1986. Mimicking the one-dimensional marginal distributions of processes having an Itô differential. *Probability theory and related fields*, 71(4): 501–516.
- He, S.; Zhu, Y.; Tavakol, D. N.; Ye, H.; Lao, Y.-H.; Zhu, Z.; Xu, C.; Chauhan, S.; Garty, G.; Tomer, R.; et al. 2024. Squidiff: Predicting cellular development and responses to perturbations using a diffusion model. *bioRxiv*, 2024–11.
- Heng, J.; De Bortoli, V.; Doucet, A.; and Thornton, J. 2025. Simulating diffusion bridges with score matching. *Biometrika*, asaf048.
- Hetzel, L.; Boehm, S.; Kilbertus, N.; Günnemann, S.; Theis, F.; et al. 2022. Predicting cellular responses to novel drug perturbations at a single-cell resolution. *Advances in Neural Information Processing Systems*, 35: 26711–26722.
- Huang, Y.; Zhang, O.; Wu, L.; Tan, C.; Lin, H.; Gao, Z.; Li, S.; Li, S.; et al. 2024. Re-Dock: towards flexible and realistic molecular docking with diffusion bridge. *arXiv preprint arXiv:2402.11459*.
- Jamison, B. 1970. Reciprocal processes: The stationary Gaussian case. *The Annals of Mathematical Statistics*, 41(5): 1624–1630.
- Lino, C. A.; Harper, J. C.; Carney, J. P.; and Timlin, J. A. 2018. Delivering CRISPR: a review of the challenges and approaches. *Drug delivery*, 25(1): 1234–1257.
- Lipman, Y.; Havasi, M.; Holderrieth, P.; Shaul, N.; Le, M.; Karrer, B.; Chen, R. T.; Lopez-Paz, D.; Ben-Hamu, H.; and Gat, I. 2024. Flow matching guide and code. *arXiv preprint arXiv:2412.06264*.
- Liu, X.; Wu, L.; Ye, M.; and Liu, Q. 2022. Let us build bridges: Understanding and extending diffusion generative models. *arXiv preprint arXiv:2208.14699*.
- Loshchilov, I. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Lotfollahi, M.; Klimovskaia Susmelj, A.; De Donno, C.; Hetzel, L.; Ji, Y.; Ibarra, I. L.; Srivatsan, S. R.; Naghipourfar, M.; Daza, R. M.; Martin, B.; et al. 2023. Predicting cellular responses to complex perturbations in high-throughput screens. *Molecular systems biology*, 19(6): e11517.
- Lotfollahi, M.; Wolf, F. A.; and Theis, F. J. 2019. scGen predicts single-cell perturbation responses. *Nature methods*, 16(8): 715–721.
- Mortazavi, A.; Williams, B. A.; McCue, K.; Schaeffer, L.; and Wold, B. 2008. Mapping and quantifying mammalian transcriptomes by RNA-Seq. *Nature methods*, 5(7): 621–628.
- Peidli, S.; Green, T. D.; Shen, C.; Gross, T.; Min, J.; Garda, S.; Yuan, B.; Schumacher, L. J.; Taylor-King, J. P.; Marks, D. S.; et al. 2024. scPerturb: harmonized single-cell perturbation data. *Nature Methods*, 21(3): 531–540.
- Peluchetti, S. 2023. Non-denoising forward-time diffusions. *arXiv preprint arXiv:2312.14589*.

- Rogers, L. C. G.; and Williams, D. 2000. *Diffusions, Markov processes, and martingales*, volume 2. Cambridge university press.
- Roohani, Y.; Huang, K.; and Leskovec, J. 2024. Predicting transcriptional outcomes of novel multigene perturbations with GEARS. *Nature Biotechnology*, 42(6): 927–935.
- Rüschemdorf, L. 1995. Convergence of the iterative proportional fitting procedure. *The Annals of Statistics*, 1160–1174.
- Schrödinger, E. 1932. Sur la théorie relativiste de l'électron et l'interprétation de la mécanique quantique. In *Annales de l'institut Henri Poincaré*, volume 2, 269–310.
- Shi, Y.; De Bortoli, V.; Campbell, A.; and Doucet, A. 2023. Diffusion schrödinger bridge matching. *Advances in Neural Information Processing Systems*, 36: 62183–62223.
- Srivatsan, S. R.; McFaline-Figueroa, J. L.; Ramani, V.; Saunders, L.; Cao, J.; Packer, J.; Pliner, H. A.; Jackson, D. L.; Daza, R. M.; Christiansen, L.; et al. 2020. Massively multiplex chemical transcriptomics at single-cell resolution. *Science*, 367(6473): 45–51.
- Tamir, E.; Trapp, M.; and Solin, A. 2023. Transport with support: Data-conditional diffusion bridges. *arXiv preprint arXiv:2301.13636*.
- Tang, Z.; Hang, T.; Gu, S.; Chen, D.; and Guo, B. 2024. Simplified diffusion schrödinger bridge. *arXiv preprint arXiv:2403.14623*.
- Tong, A.; Fatras, K.; Malkin, N.; Huguet, G.; Zhang, Y.; Rector-Brooks, J.; Wolf, G.; and Bengio, Y. 2023. Improving and generalizing flow-based generative models with minibatch optimal transport. *arXiv preprint arXiv:2302.00482*.
- Wolf, F. A.; Angerer, P.; and Theis, F. J. 2018. SCANPY: large-scale single-cell gene expression data analysis. *Genome biology*, 19: 1–5.
- Wu, Y.; Barton, R. A.; Wang, Z.; Ioannidis, V. N.; De Donno, C.; Price, L. C.; Voloch, L. F.; and Karypis, G. 2022. Predicting cellular responses with variational causal inference and refined relational information. *arXiv preprint arXiv:2210.00116*.