

000
001
002
003
004
005
006
007
008
009
010
011
012
013
014
015
016
017
018
019
020
021
022
023
024
025
026
027
028
029
030
031
032
033
034
035
036
037
038
039
040
041
042
043
044
045
046
047
048
049
050
051
052
053

BRANCHED SCHRÖDINGER BRIDGE MATCHING

Anonymous authors

Paper under double-blind review

ABSTRACT

Predicting the intermediate trajectories between an initial and target distribution is a central problem in generative modeling. Existing approaches, such as flow matching and Schrödinger Bridge Matching, effectively learn mappings between two distributions by modeling a single stochastic path. However, these methods are inherently limited to unimodal transitions and cannot capture *branched* or *divergent* evolution from a common origin to multiple distinct outcomes. To address this, we introduce **Branched Schrödinger Bridge Matching (BranchSBM)**, a novel framework that learns branched Schrödinger bridges. BranchSBM parameterizes multiple time-dependent velocity fields and growth processes, enabling the representation of population-level divergence into multiple terminal distributions. We show that BranchSBM is not only more expressive but also essential for tasks involving multi-path surface navigation, modeling cell fate bifurcations from homogeneous progenitor states, and simulating diverging cellular responses to perturbations.

1 INTRODUCTION

Tasks like crowd navigation and modeling cell-state transitions under perturbation involve learning a transport map between two empirically observed endpoint distributions, rather than a noisy prior required for denoising diffusion (Austin et al., 2021) and flow matching (Lipman et al., 2023). The Schrödinger Bridge (SB) (Schrödinger, 1931) problem seeks to identify an optimal stochastic map between a pair of endpoint distributions that minimizes the Kullback-Leibler (KL) divergence to an underlying reference process.

Schrödinger Bridge Matching (SBM) solves the SB problem by parameterizing a drift field that matches a mixture of conditional stochastic bridges between endpoint pairs that each minimize the KL divergence from a known reference process. Typically, SBM assumes conservation of mass from the initial to the target distribution, which fails to capture dynamical population behaviors such as growth and destruction of mass, commonly seen in single-cell population data. Furthermore, prior works focus on transporting samples from a pair of unimodal initial and target distributions via a single, continuous trajectory, without accounting for *branching* dynamics (Liu et al., 2023a; Tong et al., 2024b; Theodoropoulos et al., 2024; Liu et al., 2022; De Bortoli et al., 2021b), where a uniform population follows a branched trajectory that diverges toward multiple distinct target distributions.

The notion of branching is central to many real-world systems. For example, when a homogeneous cell population undergoes a perturbation such as gene knockouts or drug treatments, it frequently induces fate bifurcation as the cell population splits into multiple phenotypically distinct outcomes or commits to divergent cell fates (Shalem et al., 2014; Zhang et al., 2025a). These trajectories are observable in single-cell RNA sequencing (scRNA-seq) data, where each subpopulation independently evolves and undergoes growth or contraction along its trajectory toward a distinct terminal state (Dixit et al., 2016).

In this work, we introduce **Branched Schrödinger Bridge Matching (BranchSBM)**, a novel framework for learning stochastic transport maps from an unimodal initial distribution to multiple target distributions via branched trajectories.

Our main contributions can be summarized as follows:

1. We define the Branched Generalized Schrödinger Bridge problem and introduce **BranchSBM**, a novel matching framework that learns optimal branched trajectories from an initial distribution π_0 to multiple target distributions $\{\pi_{t,k}\}$.

054

055 2. We derive the Branched Conditional Stochastic Optimal Control (CondSOC) problem as the

056 sum of Unbalanced CondSOC objectives and leverage a multi-stage training algorithm to learn

057 the optimal branching drift and growth fields that transport mass along a branched trajectory.

058 3. We demonstrate the unique capability of BranchSBM to model dynamic branching trajectories

059 while matching multiple target distributions across various problems, including 3D navigation

060 over LiDAR manifolds (Section 5.1), modeling differentiating single-cell population dynamics

061 (Section 5.2), and predicting heterogeneous cell states after perturbation (Section 5.3).

062 **2 PRELIMINARIES**

064 **Schrödinger Bridge** Given a reference probability path measure \mathbb{Q} , the Schrödinger Bridge (SB)

065 problem aims to find an optimal path measure \mathbb{P}^{SB} that minimizes the Kullback-Leibler (KL) diver-

066 gence with \mathbb{Q} while satisfying the boundary distributions $\mathbb{P}_0 = \pi_0$ and $\mathbb{P}_1 = \pi_1$.

068
$$\mathbb{P}^{\text{SB}} = \min_{\mathbb{P}} \{\text{KL}(\mathbb{P} \parallel \mathbb{Q}) : \mathbb{P}_0 = \pi_0, \mathbb{P}_1 = \pi_1\} \quad (1)$$

070 where \mathbb{Q} is commonly defined as standard Brownian motion. For an extended background and formal

071 definition of Schrödinger Bridges, refer to Definition 3 and Appendix A.1.

072 **Generalized Schrödinger Bridge Problem** The solution to the standard SB problem minimizes

073 the kinetic energy of the conditional drift term $u_t(X_t)$ that preserves the endpoints drawn from the

074 coupling $(\mathbf{x}_0, \mathbf{x}_1) \sim \pi_{0,1}$ defined as

076
$$\min_{u_t} \int_0^1 \mathbb{E}_{p_t} \|u_t(X_t)\|^2 dt \quad \text{s.t.} \quad \begin{cases} dX_t = u_t(X_t) dt + \sigma dB_t \\ X_0 \sim \pi_0, \quad X_1 \sim \pi_1 \end{cases} \quad (2)$$

079 where dB_t is standard d -dimensional Brownian motion. To define more complex systems where the

080 *optimal* dynamics cannot be accurately captured by minimizing the standard squared-Euclidean cost

081 in entropic OT (Vargas et al., 2021), the Generalized Schrödinger Bridge (GSB) problem introduces

082 an additional non-linear state-cost $V_t(X_t)$ (Chen et al., 2021b; Chen & Georgiou, 2016; Liu et al.,

083 2022). The minimization objective becomes

084
$$\min_{u_t} \int_0^1 \mathbb{E}_{p_t} \left[\frac{1}{2} \|u_t(X_t)\|^2 + V_t(X_t) \right] dt \quad (3)$$

085 The state cost can also be interpreted as the *potential energy* of the system at state X_t .

088 **3 BRANCHED SCHRÖDINGER BRIDGE MATCHING**

091 **A key challenge** in trajectory matching is simulating trajectories that reconstruct multi-modal

092 marginals. Existing Schrödinger bridge and flow matching techniques have approached this problem

093 by reconstructing the weighted multi-modal distribution by simulating many independent particle

094 trajectories. To ensure that the intermediate trajectories follow a path distribution that approximates

095 that observed in the data, **multimarginal Schrödinger bridge matching** interpolates between

096 intermediate distributions along the path at specified time steps (Shen et al., 2024; Chen et al., 2023;

097 Theodoropoulos et al., 2025).

098 However, these methods remain limited in their ability to dynamically model *branched trajectories*,

099 where the marginal distributions have distinct modes that diverge in the state space following some

100 underlying energy landscape. When there is limited ground truth snapshot data of the intermediate

101 marginals, existing methods fail to accurately infer intermediate dynamics that follow some potential

102 energy manifold. Furthermore, simulating a branched path distribution containing multi-modal

103 marginals with independent particle simulations can suffer from mode collapse, where particles are

104 biased towards a high-density mode in the distribution or only traverse low-energy intermediate paths.

105 To address this challenge, we introduce **Branched Schrödinger Bridge Matching (BranchSBM)**, a

106 framework that infers intermediate branched dynamics by learning a *set* of diverging velocity networks

107 that accurately reconstruct a multi-modal target distribution following a potential energy landscape

V_t and growth networks that optimally distribute mass across the set of branches while recovering

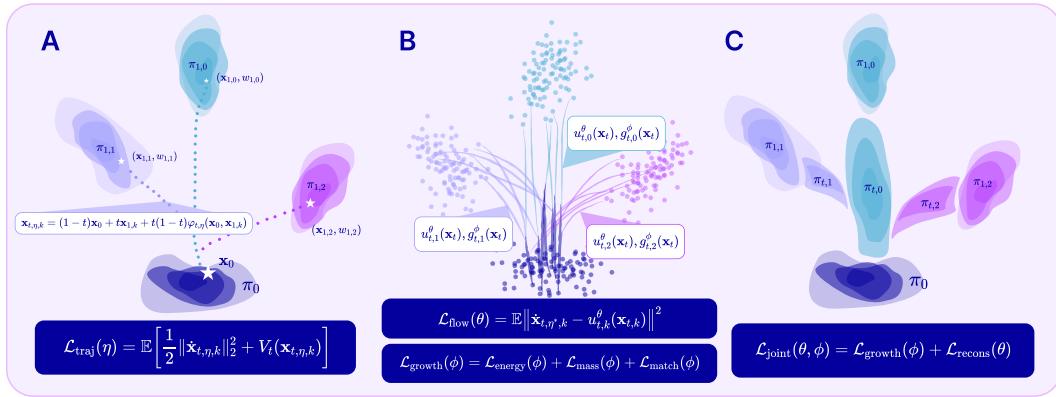


Figure 1: **Branched Schrödinger Bridge Matching** (A) Stage 1 trains a correction term that learns the optimal interpolant conditioned on endpoints (B) Stage 2 and 3 trains a separate flow and growth network for each branch independently (C) Stage 4 jointly optimizes the flow and growth networks to minimize the energy, mass, and matching loss.

the target density. Crucially, our approach does not require explicit samples from intermediate time points and can model branched, multi-modal dynamics from simulating a single sample from the initial distribution.

3.1 UNBALANCED CONDITIONAL STOCHASTIC OPTIMAL CONTROL

Unbalanced Generalized Schrödinger Bridge Problem Extending the definition of the Generalized Schrödinger Bridge (GSB) problem in Equation 3, we define the Unbalanced GSB problem by scaling the minimization objective by a time-dependent weight $w_t(X_t) = w_0 + \int_0^t g_s(X_s)ds$ that evolves according to a time-varying growth rate $g_t(X_t) : \mathbb{R}^d \times [0, 1] \rightarrow \mathbb{R}$.

$$\min_{u_t, g_t} \int_0^1 \mathbb{E}_{p_t} \left[\frac{1}{2} \|u_t(X_t)\|^2 + V_t(X_t) \right] w_t(X_t) dt \quad \text{s.t.} \quad \begin{cases} dX_t = u_t(X_t)dt + \sigma dB_t \\ X_0 \sim \pi_0, \quad X_1 \sim \pi_1 \\ w_0(X_0) = w_0^*, \quad w_1(X_1) = w_1^* \end{cases} \quad (4)$$

Unbalanced Conditional Stochastic Optimal Control (CondSOC) Now, we show that we can solve the Unbalanced GSB problem as an Unbalanced CondSOC problem where the optimal drift u_t and growth g_t minimize the expectation of the objective in (4) conditioned on pairs of endpoints.

Proposition 1 (Unbalanced Conditional Stochastic Optimal Control). *Suppose the marginal density can be decomposed as $p_t(X_t) = \int_{\pi_{0,1}} p_t(X_t | \mathbf{x}_0, \mathbf{x}_1) p_{0,1}(\mathbf{x}_0, \mathbf{x}_1) d\pi_{0,1}$, where $\pi_{0,1}$ is a fixed joint coupling of the data. Then, we can identify the optimal drift u_t^* and growth g_t^* that solves the Unbalanced GSB problem in (4) by minimizing the Unbalanced Conditional Stochastic Optimal Control objective given by*

$$\min_{u_t, g_t} \mathbb{E}_{(\mathbf{x}_0, \mathbf{x}_1) \sim \pi_{0,1}} \left[\int_0^1 \mathbb{E}_{p_t|0,1} \left[\frac{1}{2} \|u_t(X_t | \mathbf{x}_0, \mathbf{x}_1)\|^2 + V_t(X_t) \right] w_t(X_t) dt \right] \quad (5)$$

$$\text{s.t. } dX_t = u_t(X_t | \mathbf{x}_0, \mathbf{x}_1)dt + \sigma dB_t, \quad X_0 = \mathbf{x}_0, \quad X_1 = \mathbf{x}_1 \quad w_0(X_0) = w_0^*, \quad w_1(X_1) = w_1^* \quad (6)$$

where $w_t = w_0 + \int_0^t g_s(X_s)ds$ is the time-dependent weight initialized at w_0^* , u_t is the drift, g_t is the growth rate, and $\pi_{0,1}$ is the weighted coupling of paired endpoints $(\mathbf{x}_0, w_0^*, \mathbf{x}_1, w_1^*) \sim \pi_{0,1}$.

The proof is provided in Appendix C.1. This defines the objective for us to tractably solve the Unbalanced GSB problem by conditioning on a finite set of endpoint pairs in the dataset.

3.2 BRANCHSBM: SUM OF UNBALANCED CONDSOC PROBLEMS

Branched Generalized Schrödinger Bridge Problem Given the Unbalanced GSB problem, we define the Branched GSB problem as minimizing the sum of Unbalanced GSB problems across all branches. All mass begins along a primary path indexed $k = 0$ with initial weight 1. Over $t \in [0, 1]$,

mass is transferred across K secondary branches with initial weight 0 and target weight $w_{1,k}$ such that it minimizes the objective defined as

$$\begin{aligned} & \min_{\{u_{t,k}, g_{t,k}\}_{k=0}^K} \int_0^1 \left\{ \mathbb{E}_{p_{t,0}} \left[\frac{1}{2} \|u_{t,0}(X_{t,0})\|^2 + V_t(X_{t,0}) \right] w_{t,0} + \sum_{k=1}^K \mathbb{E}_{p_{t,k}} \left[\frac{1}{2} \|u_{t,k}(X_{t,k})\|^2 + V_t(X_{t,k}) \right] w_{t,k} \right\} dt \\ & \text{s.t. } dX_{t,k} = u_{t,k}(X_{t,k})dt + \sigma dB_t, \quad X_0 \sim \pi_0, \quad X_{1,k} \sim \pi_{1,k}, \quad w_{0,k} = \delta_{k=0}, \quad w_{1,k} = w_{1,k} \end{aligned} \quad (7)$$

When total mass across branches is conserved, we enforce $\sum_{k=0}^K w_{t,k} = 1$ for all $t \in [0, 1]$, which constrains the growth rates such that $g_{t,0}(X_{t,0}) + \sum_{k=1}^K g_{t,k}(X_{t,k}) = 0$. This ensures that mass lost from the primary branch (when $g_{t,0} < 0$) is redistributed among the secondary branches (where $g_{t,k} > 0$). The primary branch evolves from an initial weight of 1 according to $w_{t,0} = 1 + \int_0^t g_s(X_{s,0})ds$ and the K secondary branches grow from the primary branch from weight 0 according to $w_{t,k} = \int_0^t g_s(X_{s,k})ds$.

Branched Conditional Stochastic Optimal Control Following a similar procedure as shown for the Unbalanced GSB problem, we can reformulate the Branched GSB problem as solving the Branched CondSOC problem where we optimize a set of parameterized drift $\{u_{t,k}\}_{k=0}^K$ and growth $\{g_{t,k}\}_{k=0}^K$ networks by minimizing the energy of the conditional trajectories between paired samples $(\mathbf{x}_0, \{\mathbf{x}_{1,k}\}_{k=0}^K) \sim \{p_{0,1,k}\}_{k=0}^K$.

Proposition 2 (Branched Conditional Stochastic Optimal Control). *For each branch, let $p_{t,k}(X_{t,k}) = \mathbb{E}_{p_{0,1,k}}[p_{t,k}(X_{t,k} | \mathbf{x}_0, \mathbf{x}_{1,k})]$, where $\pi_{0,1,k}$ is the joint coupling distribution of samples $\mathbf{x}_0 \sim \pi_0$ from the initial distribution and $\mathbf{x}_{1,k} \sim \pi_{1,k}$ from the k th target distribution. Then, we can identify the set of optimal drift and growth functions $\{u_{t,k}^*, g_{t,k}^*\}_{k=0}^K$ that solve the Branched GSB problem in (3.2) by minimizing the sum of Unbalanced CondSOC objectives given by*

$$\begin{aligned} & \min_{\{u_{t,k}, g_{t,k}\}_{k=0}^K} \mathbb{E}_{(\mathbf{x}_0, \mathbf{x}_{1,0}) \sim \pi_{0,1,0}} \int_0^1 \left\{ \mathbb{E}_{p_{t|0,1,0}} \left[\frac{1}{2} \|u_{t,0}(X_{t,0})\|^2 + V_t(X_{t,0}) \right] w_{t,0} \right. \\ & \quad \left. + \sum_{k=1}^K \mathbb{E}_{(\mathbf{x}_0, \mathbf{x}_{1,k}) \sim \pi_{0,1,k}} \int_0^1 \mathbb{E}_{p_{t|0,1,k}} \left[\frac{1}{2} \|u_{t,k}(X_{t,k})\|^2 + V_t(X_{t,k}) \right] w_{t,k} \right\} dt \end{aligned} \quad (8)$$

$$\text{s.t. } dX_{t,k} = u_{t,k}(X_{t,k})dt + \sigma dB_t, \quad X_0 = \mathbf{x}_0, \quad X_{1,k} = \mathbf{x}_{1,k}, \quad w_{0,k} = \delta_{k=0}, \quad w_{1,k} = w_{1,k} \quad (9)$$

where $w_{t,0} = 1 + \int_0^t g_{s,1}(\mathbf{x}_{s,1})ds$ is the weight of the primary paths initialized at 1 and $w_{t,k} = \int_0^t g_{s,k}(\mathbf{x}_{s,k})ds$ are the weights of the K secondary branches initialized at 0.

The proof is given in Appendix C.2. This defines the objective for us to tractably solve the Branched GSB problem in Section 4 by conditioning on a discrete set of branched endpoint pairs in the dataset.

Remark 1. When $g_{t,0}(X_{t,0}) = 0$ and $g_{t,k}(X_{t,k}) = 0$ for all $t \in [0, 1]$ and $k \in \{1, \dots, K\}$, then the Branched CondSOC problem is the solution to the single path GSB problem.

4 LEARNING BRANCHSBM USING NEURAL NETWORKS

Given an initial data distribution π_0 and $K+1$ target distributions $\{\pi_{1,k}\}_{k=0}^K$, we aim to parameterize the optimal drift and growth fields that solve the Branched CondSOC problem in Proposition 2.

4.1 BRANCHED NEURAL INTERPOLANT OPTIMIZATION

Since the optimal trajectory under the state cost $V_t(X_t)$ follows a non-linear cost manifold, given a pair of endpoints $(\mathbf{x}_0, \mathbf{x}_{1,k})$, we train a neural path interpolant $\varphi_{t,\eta}(\mathbf{x}_0, \mathbf{x}_{1,k}) : \mathbb{R}^d \times \mathbb{R}^d \times [0, 1] \rightarrow \mathbb{R}^d$ that defines the intermediate state $\mathbf{x}_{t,\eta,k}$ and velocity $\dot{\mathbf{x}}_{t,\eta,k} = \partial_t \mathbf{x}_{t,\eta,k}$ at time t , which minimizes (2). We note that the neural interpolant has been introduced previously for single-path interpolants in Kapuśnian et al. (2024); Neklyudov et al. (2023). We define $\mathbf{x}_{t,\eta,k}$ to be bounded at the endpoints as given by

$$\mathbf{x}_{t,\eta,k} = (1-t)\mathbf{x}_0 + t\mathbf{x}_{1,k} + t(1-t)\varphi_{t,\eta}(\mathbf{x}_0, \mathbf{x}_{1,k}) \quad (10)$$

$$\dot{\mathbf{x}}_{t,\eta,k} = \mathbf{x}_{1,k} - \mathbf{x}_0 + t(1-t)\dot{\varphi}_{t,\eta}(\mathbf{x}_0, \mathbf{x}_{1,k}) + (1-2t)\varphi_{t,\eta}(\mathbf{x}_0, \mathbf{x}_{1,k}) \quad (11)$$

To optimize $\varphi_{t,\eta}(\mathbf{x}_0, \mathbf{x}_{1,k})$ such that it predicts the energy-minimizing trajectory, we minimize the trajectory loss $\mathcal{L}_{\text{traj}}$ defined as

$$\mathcal{L}_{\text{traj}}(\eta) = \sum_{k=0}^K \int_0^1 \mathbb{E}_{(\mathbf{x}_0, \mathbf{x}_{1,k}) \sim \pi_{0,1,k}} \left[\frac{1}{2} \|\dot{\mathbf{x}}_{t,\eta,k}\|_2^2 + V_t(\mathbf{x}_{t,\eta,k}) \right] dt \quad (12)$$

After convergence, Stage 1 returns the network $\varphi_{t,\eta}^*(\mathbf{x}_0, \mathbf{x}_{1,k})$ that generates the optimal conditional velocity $\dot{\mathbf{x}}_{t,\eta,k}^*$ which defines the matching objective in Stage 2. In Stage 2, we parameterize a set of neural drift fields $u_{t,k}^\theta(\mathbf{x}_{t,k}) : \mathbb{R}^d \times [0, 1] \rightarrow \mathbb{R}^d$ that generates the *mixture of bridges* defined in Stage 1 by minimizing the conditional flow matching loss (Lipman et al., 2023; Tong et al., 2024a).

$$\mathcal{L}_{\text{flow}}(\theta) = \sum_{k=0}^K \int_0^1 \mathbb{E}_{(\mathbf{x}_0, \mathbf{x}_{1,k}) \sim \pi_{0,1,k}} \left\| \dot{\mathbf{x}}_{t,\eta,k}^* - u_{t,k}^\theta(\mathbf{x}_{t,k}) \right\|_2^2 dt \quad (13)$$

Proposition 3 (Solving the GSB Problem with Stage 1 and 2 Training). *Stage 1 and Stage 2 training yield the optimal drift $u_t^*(X_t)$ that generates the optimal marginal probability distribution $p_t^*(X_t)$ that solves the GSB problem in (3).*

The proof is provided in Appendix C.3. Since the drift for each branch $u_{t,k}^\theta(X_t)$ are trained independently in Stage 2, we can extend this result across all $K+1$ branches and conclude that the sequential Stage 1 and Stage 2 training procedures yields the optimal set of drifts $\{u_{t,k}^*\}_{k=0}^K$ that generate the optimal probability paths $\{p_{t,k}^*\}_{k=0}^K$ that solves the GSB problem for each branch.

4.2 LEARNING THE ENERGY-MINIMIZING BRANCHING DYNAMICS

Branched Energy Loss To solve the Branched CondSOC problem defined in Proposition 2, we minimize a branched energy loss $\mathcal{L}_{\text{energy}}$ defined as

$$\begin{aligned} \mathcal{L}_{\text{energy}}(\theta, \phi) &= \int_0^1 \mathbb{E}_{\{p_{t,k}\}_{k=0}^K} \left\{ \underbrace{\left[\frac{1}{2} \|u_{t,0}^\theta(\mathbf{x}_{t,0})\|^2 + V_t(\mathbf{x}_{t,0}) \right] w_{t,0}^\phi}_{\text{primary trajectory}} + \underbrace{\sum_{k=1}^K \left[\frac{1}{2} \|u_{t,k}^\theta(\mathbf{x}_{t,k})\|^2 + V_t(\mathbf{x}_{t,k}) \right] w_{t,k}^\phi}_{K \text{ branches}} \right\} dt \\ &\text{s.t. } w_{t,0}^\phi = 1 + \int_0^t g_{s,1}^\phi(\mathbf{x}_{s,1}) ds, \quad w_{t,k}^\phi = \int_0^t g_{s,k}^\phi(\mathbf{x}_{s,k}) ds \end{aligned} \quad (14)$$

where $(\mathbf{x}_0, \mathbf{x}_{1,0})$ are the endpoints of the *primary path*. At time $t=0$, the primary path has weight 1 and the K branches have weights 0. Over $t \in [0, 1]$, the weight of the primary path changes according to $g_{t,0}^\phi(\mathbf{x}_{t,0})$ and supplies mass to the K branches, which grow at rates $g_{t,k}^\phi(\mathbf{x}_{t,k}) \geq 0$ (Lemma 2). Intuitively, the branched energy loss optimizes the branching growth rates such that they are non-zero when branching is favored over the primary path.

Weight Matching Loss We define a weight matching loss $\mathcal{L}_{\text{match}}$ that aims to minimize the difference between the predicted weights of each branch at $t=1$, obtained by integrating the growth function $g_{t,k}^\phi(X_t)$ over $t \in [0, 1]$, and the true weights of each terminal distribution $\{w_{1,k}^*\}_{k=0}^K$.

$$\mathcal{L}_{\text{match}}(\phi) = \sum_{k=0}^K \mathbb{E}_{p_{1,k}} \left(w_{1,k}^\phi - w_{1,k}^* \right)^2, \quad \text{s.t. } w_{1,k}^\phi(\mathbf{x}_{1,k}) = w_{0,k} + \int_0^1 g_{t,k}^\phi(\mathbf{x}_{t,k}) dt \quad (15)$$

where $w_{1,k}^* = N_k / N_{\text{total}}$ is the fraction of the population in the k th target distribution.

Mass Conservation Loss To ensure that the growth rate satisfies conservation of total mass at all times $t \in [0, 1]$, we define a mass loss $\mathcal{L}_{\text{mass}}$ that enforces the sum of the weights of all $K+1$ branches matches the true total weight at time t denoted as w_t^{total} .

$$\mathcal{L}_{\text{mass}}(\phi) = \int_0^1 \mathbb{E}_{\{p_{t,k}\}_{k=0}^K} \left[\left(\sum_{k=0}^K w_{t,k}^\phi(\mathbf{x}_{t,k}) - w_t^{\text{total}} \right)^2 + \sum_{k=0}^K \max(0, -w_{t,k}^\phi(\mathbf{x}_{t,k})) \right] dt \quad (16)$$

where $\max(0, -w_{t,k}^\phi)$ assigns an additional linear penalty for negative weight predictions. For the balanced branched SBM problem where the total mass is conserved, we have $w_t^{\text{total}} = 1$.

270 **Training the Growth Networks** In Stage 3, we train the growth networks $\{g_{t,k}^\phi\}_{k=0}^K$ by fixing the
 271 weights of the flow networks $\{u_{t,k}^\theta\}_{k=0}^K$ and minimizing the weighted combined loss $\mathcal{L}_{\text{growth}}$ with an
 272 additional growth penalty term $\|g_{t,k}^\phi\|_2^2$ to ensure coercivity of $\mathcal{L}_{\text{growth}}$
 273

$$\mathcal{L}_{\text{growth}}(\phi) = \lambda_{\text{energy}} \mathcal{L}_{\text{energy}}(\theta, \phi) + \lambda_{\text{match}} \mathcal{L}_{\text{match}}(\phi) + \lambda_{\text{mass}} \mathcal{L}_{\text{mass}}(\phi) + \lambda_{\text{growth}} \sum_{k=0}^K \|g_{t,k}^\phi\|_2^2 \quad (17)$$

278 We show in Lemma 2 that the optimal growth rates across the K secondary branches are non-
 279 decreasing; however, mass destruction can still be modeled by defining an additional branch with
 280 target weight equal to the ratio of mass lost over $t \in [0, 1]$. To ensure that the set of optimal growth
 281 functions g^* exists, we establish Proposition 4 (see proof in Appendix C.4).

282 **Proposition 4** (Existence of Optimal Growth Functions). *Assume the state space $\mathcal{X} \subseteq \mathbb{R}^d$ is a
 283 bounded domain within \mathbb{R}^d . Let the optimal probability density of branch k be a known non-negative
 284 function bounded in $[0, 1]$, denoted as $p_{t,k}^* : \mathcal{X} \times [0, 1] \rightarrow [0, 1] \in L^\infty(\mathcal{X} \times [0, 1])$. By Lemma 2, we
 285 can define the set of feasible growth functions in the set of square-integrable functions L^2 as*

$$\mathcal{G} := \{g = (g_{t,0}, \dots, g_{t,K}) \in L^2(\mathcal{X} \times [0, 1]) \mid g_{t,k}(\mathbf{x}) : \mathcal{X} \times [0, 1] \rightarrow \mathbb{R}, g_{t,k}(\mathbf{x}) \geq 0\} \quad (18)$$

288 Let the growth loss be the functional $\mathcal{L}(g) : L^2(\mathcal{X} \times [0, 1]) \rightarrow \mathbb{R}$. Then, there exists an optimal
 289 function $g^* = (g_{t,0}^*, \dots, g_{t,K}^*) \in L^2$ where $g_{t,k}^* \in \mathcal{G}$ such that $\mathcal{L}(g^*) = \inf_{g \in \mathcal{G}} \mathcal{L}(g)$ which can be
 290 obtained by minimizing $\mathcal{L}(g)$ over \mathcal{G} .
 291

292 **Final Joint Training** In the final Stage 4, we train the weights for both the flow and growth
 293 networks $\{u_{t,k}^\theta, g_{t,k}^\phi\}_{k=0}^K$ by minimizing $\mathcal{L}_{\text{growth}}$ from Stage 3 in addition to a reconstruction loss
 294 $\mathcal{L}_{\text{recons}}$ that ensures the endpoint distribution at time $t = 1$ is maintained.
 295

$$\mathcal{L}_{\text{recons}}(\theta) = \sum_{k=0}^K \mathbb{E}_{p_{1,k}} \sum_{\mathbf{x}_1 \in \mathcal{N}_n(\mathbf{x}_{1,k})} \max \left(0, \|\tilde{\mathbf{x}}_{1,k} - \mathbf{x}_{1,k}\|_2 - \epsilon \right) \quad (19)$$

300 where $\mathcal{N}_n(\mathbf{x}_{1,k})$ is the set of n -nearest neighbors to the reconstructed state $\tilde{\mathbf{x}}_{1,k} \sim p_{1,k}$ at time $t = 1$
 301 from the data points $\mathbf{x}_{1,k} \sim \pi_{1,k}$ at time $t = 1$.

302 Our multi-stage training scheme decomposes the Branched CondSOC problem into two parts. We
 303 first independently learn an optimal drift field for each branch, which is a vector field over the state
 304 space that propagates mass flow in the direction of each target distribution. Then, we fix the drift fields
 305 and learn the growth dynamics that determine the optimal distribution of mass over the branches.
 306

307 5 EXPERIMENTS

310 We evaluate BranchSBM on a variety of branched matching tasks with different state costs $V_t(X_t)$,
 311 including multi-path LiDAR navigation (Section 5.1), modeling differentiating single-cell population
 312 dynamics (Section 5.2), and predicting heterogeneous cell-states after perturbation (Section 5.3).
 313

314 5.1 BRANCHED LiDAR SURFACE NAVIGATION

316 First, we evaluate BranchSBM for navigating branched paths along the surface of a 3-dimensional
 317 LiDAR manifold, from an initial distribution to two distinct target distributions (Figure 3).
 318

319 **Setup** We define a single initial Gaussian mixture π_0 and two target Gaussian mixtures $\pi_{1,0}, \pi_{1,1}$
 320 on either side of the mountain (Figure 3). We sample 5000 points i.i.d. from each of the Gaussian
 321 mixtures and assign all endpoints a target weight of $w_{1,0} = w_{1,1} = 0.5$. To ensure trajectories
 322 follow the LiDAR manifold, we define the state cost $V_t^{\text{LAND}}(X_t)$ as the data-dependent LAND
 323 metric (Kapuśnian et al., 2024; Arvanitidis et al., 2016), which assigns lower costs in regions near
 324 coordinates in the LiDAR dataset. Further experimental details are provided in Appendix E.3.

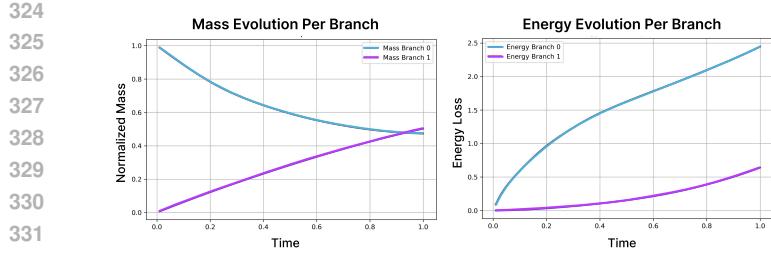


Figure 2: Plot of weight (left) and energy (right) calculated with (14) of each branch over time $t \in [0, 1]$. Mass is transferred from the primary branch to branch 1, and both converge to the target weight of 0.5 at $t = 1$. Both plots represent the average over trajectories from samples in the validation set.

Table 1: **Benchmark of BranchSBM against single-branch SBM on multi-path surface navigation.** Wasserstein distances (\mathcal{W}_1 and \mathcal{W}_2) between the reconstructed and ground-truth distributions with $N_{\text{steps}} = 100$ Euler steps at time $t = 1$ from validation samples in the initial distribution. Results are averaged over 5 independent runs.

Model	$\mathcal{W}_1 (\downarrow)$	$\mathcal{W}_2 (\downarrow)$
Single Branch SBM	0.975 ± 0.009	1.285 ± 0.007
BranchSBM	0.239 ± 0.001	0.309 ± 0.003

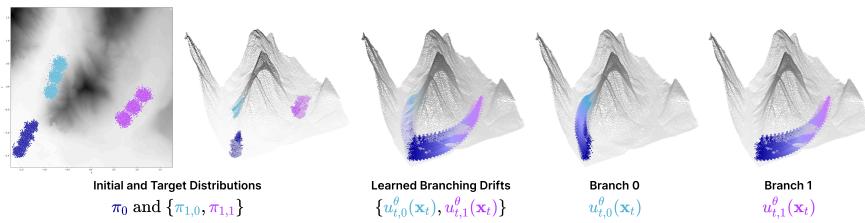


Figure 3: **Application of BranchSBM on Learning Branched Paths on a LiDAR Manifold.** Plots of the initial and target distributions, learned interpolants, and learned branched trajectories on the LiDAR manifold.

Results We show that BranchSBM can learn distinct, non-linear branched paths that curve along the 3-dimensional LiDAR manifold while minimizing the kinetic energy and state-cost. From the mass and energy curves in Figure 2, we see that mass begins in the primary branch (branch 0) and is gradually transferred to the secondary branch (branch 1) over $t \in [0, 1]$, with both curves converging to the target weight of 0.1 at $t = 1$. As mass is transferred, the slope of the cumulative energy curve decreases in branch 0 and increases in branch 1, reflecting the true energy dynamics. In Figure 3, we observe that the branching occurs at the edge of the inclined mountain, indicating that the model can determine the optimal branching time based on the paths of lowest potential energy. As shown in Table 1, BranchSBM reconstructs the endpoint distributions with significantly higher accuracy in comparison to single-branch SBM. In total, we demonstrate the capability of BranchSBM to learn branched trajectories on complex 3D manifolds.

5.2 DIFFERENTIATING SINGLE-CELL POPULATION DYNAMICS

BranchSBM is uniquely positioned to model single-cell population dynamics where a homogeneous cell population (e.g., progenitor cells) differentiates into several distinct subpopulation branches, each of which independently undergoes growth dynamics. Here, we demonstrate this capability on mouse hematopoiesis data (Sha et al., 2023; Weinreb et al., 2020) and [pancreatic \$\beta\$ -cell differentiation data](#) (Veres et al., 2019).

Mouse-Hematopoiesis Results

We use a dataset consisting of mouse hematopoiesis scRNA-seq data analyzed by a lineage tracing technique from (Sha et al., 2023; Weinreb et al., 2020). This data contains three time points t_i for $i \in \{0, 1, 2\}$ that are projected to two-dimensional representations $\mathbf{x} \in \mathbb{R}^2$ referred to as force-directed layouts or SPRING plots. From the plotted data, we can

Table 2: **Results for Modeling Single-Cell Differentiation.** Wasserstein distances (\mathcal{W}_1 and \mathcal{W}_2) between simulated and ground-truth cell distributions at time t_1 and t_2 on the validation dataset. BranchSBM reconstructs both intermediate and terminal states significantly better than single-branch SBM. Results are averaged over 5 independent runs.

Model	Single Branch SBM		BranchSBM		
	Time \downarrow	$\mathcal{W}_1 (\downarrow)$	$\mathcal{W}_2 (\downarrow)$	$\mathcal{W}_1 (\downarrow)$	$\mathcal{W}_2 (\downarrow)$
t_1		0.582 ± 0.020	0.703 ± 0.008	0.366 ± 0.034	0.479 ± 0.044
t_2		0.940 ± 0.075	1.037 ± 0.074	0.210 ± 0.042	0.265 ± 0.046

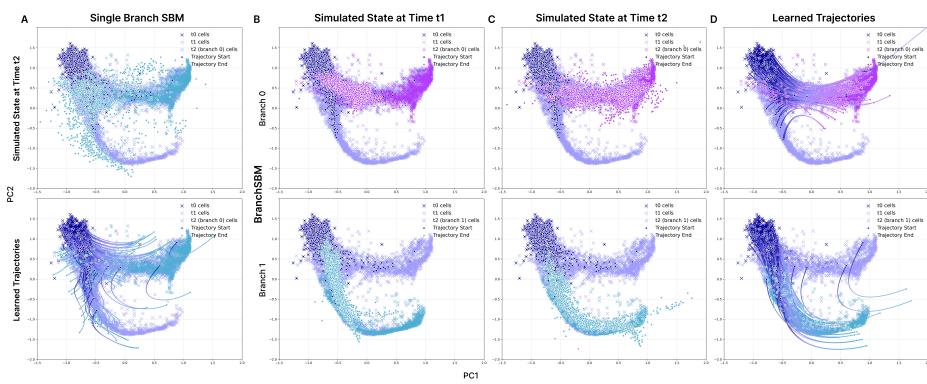


Figure 4: **Application of BranchSBM on Modeling Differentiating Single-Cell Population Dynamics.** Mouse hematopoiesis scRNA-seq data is provided for three time points t_0, t_1, t_2 . **(A)** Simulated states (top) and trajectories (bottom) at time t_1 using single-branch SBM. **(B)** Simulated states with BranchSBM at t_1 ($t = 0.5$) and **(C)** t_2 ($t = 1$). **(D)** Learned trajectories over the interval $t \in [t_0, t_2]$ on validation samples.

Table 3: **Results for pancreatic β -cell differentiation experiment.** 1-Wasserstein distance (\mathcal{W}_1) of intermediate distributions generated by DeepRUOT (Zhang et al., 2025c), CytoBridge (Zhang et al., 2025d), and BranchSBM (Ours) at eight different time points. † denotes values taken from (Zhang et al., 2025d).

Model	$t = 1$ \mathcal{W}_1	$t = 2$ \mathcal{W}_1	$t = 3$ \mathcal{W}_1	$t = 4$ \mathcal{W}_1	$t = 5$ \mathcal{W}_1	$t = 6$ \mathcal{W}_1	$t = 7$ \mathcal{W}_1
DeepRUOT †	8.0447 \pm 0.0005	8.0773 \pm 0.0021	7.6301 \pm 0.0032	8.0064 \pm 0.0042	7.9018 \pm 0.0117	8.3977 \pm 0.0102	7.8346 \pm 0.0109
CytoBridge †	8.0448 \pm 0.0005	8.0771 \pm 0.0021	7.6299 \pm 0.0032	8.0066 \pm 0.0043	7.9018 \pm 0.0117	8.3974 \pm 0.0102	7.8343 \pm 0.0109
BranchSBM (Ours)	11.9774 \pm 0.0000	7.4643 \pm 2.4031	11.5204 \pm 0.0000	11.2593 \pm 0.0000	10.2888 \pm 0.0000	8.7301 \pm 0.0000	6.8702 \pm 0.0000

observe two clear branches that indicate the differentiation of progenitor cells into two distinct cell fates (Figure 4). We use k -means clustering to define two distinct target distributions $\pi_{1,0}$ and $\pi_{1,1}$ of samples at time t_2 and set their target weights equal to $w_{1,0} = w_{1,1} = 0.5$ due to the equal ratio of cells (Figure 9). We used samples across all time steps t_i for $i \in \{0, 1, 2\}$ to define the data manifold via the LAND metric $V_{t,\eta}^{\text{LAND}}$. BranchSBM was trained on pairs sampled only from t_0 and t_2 , and samples from t_1 were held out for evaluation. For comparison, we trained a single-branch SBM model with both clusters at t_2 as the target distribution.

After evaluating the reconstructed distributions at the intermediate held-out time point t_1 and final time point t_2 ($t = 1$) from simulating validation samples from the initial distribution $\mathbf{x}_0 \sim \pi_0$. In Figure 4A, we observe that single-branch SBM trained with a single target distribution p_1 containing both terminal fates fails to learn distinct branched trajectories, and the simulated cell states at time t_2 do not reach either of the terminal distributions. In contrast, we show that BranchSBM simulates branched states at intermediate time steps not included in the training data while accurately reconstructing both target distributions with significantly lower 1-Wasserstein and 2-Wasserstein distances compared to the single-branch SBM model (Figure 4B-D; Table 2).

Pancreatic β -Cell Differentiation Results We use a dataset consisting of pancreatic β -cell differentiation data (Veres et al., 2019) containing 51,274 cells collected over eight time points as they evolve from human pluripotent stem cells to pancreatic β -like cells. This data contains eight time points t_i for $i \in \{0, \dots, 7\}$ that are projected to 30-dimensional PC representations $\mathbf{x} \in \mathbb{R}^{30}$. We use Leiden clustering to define $K = 11$ distinct target distributions of samples at time t_7 and set their target weights relative to the total weight at t_0 (App E.4). We used samples across all time steps to define the data manifold via the RBF metric $V_{t,\eta}^{\text{RBF}}$. BranchSBM was trained on pairs sampled only from t_0 and t_7 , and the distributions at all intermediate time points were inferred from learning to minimize the distance from the data manifold over time. For baselines, we compared two SOTA methods for single-cell trajectory modeling: DeepRUOT (Zhang et al., 2025c) and CytoBridge (Zhang et al., 2025d), which are trained explicitly on intermediate snapshots to reconstruct the stochastic dynamics of cells that evolve independently.

432 Notably, BranchSBM not only reconstructs the multi-modal terminal distribution at t_7 with superior
 433 accuracy against all baselines, but also produces intermediate trajectories that are competitive with
 434 models trained directly on intermediate snapshots using explicit reconstruction losses (Table 3).
 435 Leveraging the RBF state cost $V_{t,\eta}^{\text{RBF}}$, which encourages trajectories to remain on the underlying data
 436 manifold, BranchSBM effectively captures the true differentiation dynamics through the combined
 437 influence of the neural interpolant and the path-energy objective. These results demonstrate that
 438 BranchSBM scales reliably to a large number of branches and is particularly advantageous in settings
 439 where intermediate timepoints are sparse or unavailable, allowing the model to infer biologically
 440 meaningful trajectories even with limited temporal supervision.

441 5.3 CELL-STATE PERTURBATION MODELING

442 Predicting the effects of perturbation on cell state dynamics is a crucial problem for therapeutic
 443 design. In this experiment, we leverage BranchSBM to model the trajectories of a single cell line
 444 from a single homogeneous state to multiple heterogeneous states after a drug-induced perturbation.
 445 We demonstrate that BranchSBM is capable of modeling high-dimensional gene expression data and
 446 learning branched trajectories that accurately reconstruct diverging perturbed cell populations.

447 **Setup** For this experiment, we extract the data for a single cell line (A-549) under perturba-
 448 tion with Clonidine and Trametinib at $5 \mu\text{L}$, selected based on cell abundance and response
 449 diversity from the Tahoe-100M dataset (Zhang et al., 2025a). Since both drugs had over
 450 60K genes, we selected the top 2000 highly variable genes (HVGs) based on normalized ex-
 451 pression and performed principal component analysis (PCA) to find the top PCs that capture
 452 the variance in the data. We set the initial distribution at $t = 0$ to be a control DMSO-
 453 treated cell population and the target distributions at $t = 1$ to be distinct clusters in the drug-
 454 treated cell population. After clustering, we identified two divergent clusters in the Clonidine-
 455 perturbed population and three in the Trametinib-perturbed population (Appendix Figure 10).
 456 To determine the weights of each
 457 branch, we take the ratio of each cluster
 458 with respect to the total perturbed
 459 cell population (Appendix Table 9).
 460 For both experiments, we simulated
 461 the top 50 PCs, which capture approxi-
 462 mately 38% of the variance in the
 463 dataset. To further evaluate the scal-
 464 ability of BranchSBM on simulating
 465 trajectories in high-dimensional state
 466 spaces, we simulated the top 100 and
 467 150 PCs for Clonidine and compared
 468 the performance across dimensions.

469 Given that the intermediate trajectory
 470 between the control and perturbed
 471 state is unknown, we assume that the
 472 optimal trajectory both minimizes the
 473 kinetic energy of the drift field while
 474 minimizing the distance from the space of *feasible* cell states. We define the state cost $V_t(X_t)$ with
 475 the RBF metric (Kapuśnian et al., 2024; Arvanitidis et al., 2016), which pushes the intermediate
 476 trajectory to lie near states represented in the dataset. Further details are provided in Appendix E.5.

477 **Clonidine Perturbation Results** After multi-stage training of BranchSBM with $d \in \{50, 100, 150\}$
 478 PCs and two branched endpoints (Appendix Figure 6A), we simulated the final perturbed state of each
 479 branch at time $t = 1$ from the samples in the initial validation data distribution $\mathbf{x}_0 \sim \pi_0$ corresponding
 480 to the control DMSO condition. In Appendix Figure 6C, we demonstrate that BranchSBM accurately
 481 reconstructs the ground-truth distributions of endpoint 0 (top row) and endpoint 1 (bottom row)
 482 across increasing PC dimensions, capturing the location and spread of the dataset. To prove the
 483 necessity of our branched framework, we simulate the target distribution with only a single endpoint
 484 distribution p_1 containing both clusters with single-branch SBM and show that it only reconstructs
 485 the population of cells in endpoint 0, which represent cells closest to the control cells along PC2, and

Table 4: **Results for Clonidine Perturbation Modeling for Increasing Principal Component Dimensions.** Maximum-mean discrepancy (MMD) across all PCs and Wasserstein distances (\mathcal{W}_1 and \mathcal{W}_2) of top 2 PCs between ground truth and reconstructed distributions at $t = 1$ simulated from the validation data at $t = 0$. Results for single-branch SBM (50 PCs) and BranchSBM (2 branches) were averaged over 5 independent runs.

Model	RBF-MMD (\downarrow)	\mathcal{W}_1 (\downarrow)	\mathcal{W}_2 (\downarrow)
Single Branch SBM (50 PCs)	0.279 ± 0.024	5.124 ± 0.509	6.149 ± 0.463
BranchSBM			
50 PCs	0.065 ± 0.001	1.076 ± 0.085	1.224 ± 0.097
100 PCs	0.053 ± 0.002	1.832 ± 0.174	2.037 ± 0.174
150 PCs	0.083 ± 0.001	1.722 ± 0.064	1.931 ± 0.035

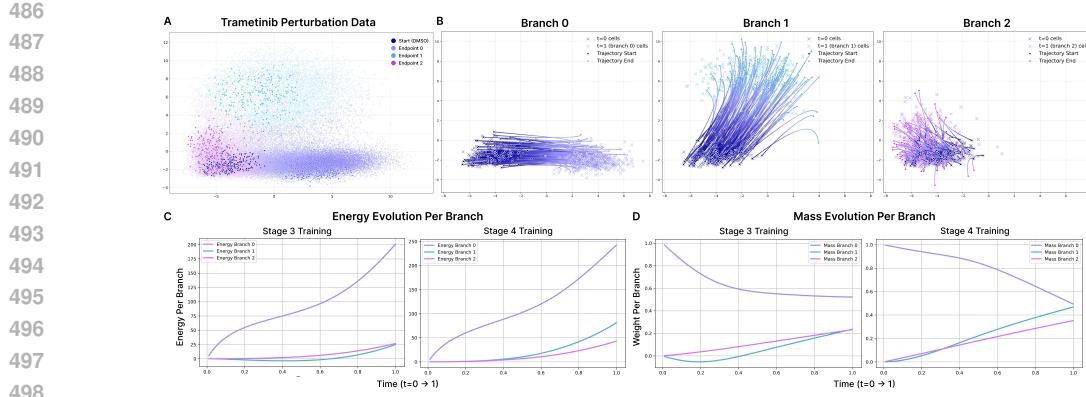


Figure 5: **Results for Trametinib Perturbation Modeling with BranchSBM.** (A) Gene expression data of DMSO control ($t = 0$) and cells after treatment with $5\mu M$ Trametinib ($t = 0$) with three distinct endpoints (purple, turquoise, and pink). (B) The simulated endpoints of the top 50 PCs at $t = 1$ on the validation data for each branch. (C) The evolution of cumulative energy across $t \in [0, 1]$ calculated as (14) along each branched trajectory after Stage 3 (growth with fixed drift) and Stage 4 (joint) training. (D) The evolution of mass across $t \in [0, 1]$ along each branched trajectory with target weights of $w_{1,0} = 0.603$, $w_{1,1} = 0.255$ and $w_{1,2} = 0.142$.

fails to differentiate cells in cluster 1 that differ from cluster 0 in higher-dimensional PCs (Appendix Figure 6B). Concretely, BranchSBM used across all PC dimensions outperforms single-branch SBM on only 50 PCs (Table 4), indicating that BranchSBM is required to model complex perturbation effects in high-dimensional gene expression spaces.

Trametinib Perturbation Results We further show that BranchSBM can scale beyond two branches, by modeling the perturbed cell population of Trametinib-treated cells, which diverge into three distinct clusters (Figure 5A). We trained BranchSBM with three endpoints and single-branch SBM with one endpoint containing all three clusters on the top 50 PCs. After simulating the trajectories over time $t \in [0, 1]$ on the validation cells in the control population, we show that BranchSBM generates clear trajectories to all three branched endpoints (Figure 5B) and reconstructs the overall target distribution with lower error compared to single-branch SBM (Table 5). In Figure 5C and D, we plot the evolution of cumulative energy calculated in $\mathcal{L}_{\text{energy}}(\theta, \phi)$ (14) and weight of each branch over $t \in [0, 1]$, demonstrating that BranchSBM’s multi-stage training scheme effectively learns the optimal trade-off between minimizing the energy across trajectories and matching the target weights of each branch.

6 CONCLUSION

In this work, we introduce **Branched Schrödinger Bridge Matching (BranchSBM)**, a novel matching framework that solves the Generalized Schrödinger Bridge (GSB) problem from an initial distribution to multiple weighted target distributions through the division of mass across learned branched trajectories. BranchSBM solves the branched SBM problem as a sum of Unbalanced Conditional Stochastic Optimal Control tasks, parameterizing branch-specific velocity and growth rates with neural networks to predict system trajectories with a single inference simulation. Through applications to nonlinear 3D navigation, cell differentiation, and perturbation-induced gene expression, we demonstrate that BranchSBM provides a unified and flexible framework for modeling complex branched dynamics across biological and physical systems.

540 REPRODUCIBILITY STATEMENT
541542 We have taken multiple steps to ensure reproducibility of BranchSBM. All theoretical derivations
543 are fully detailed in the main text and Appendix, with proofs of Propositions 1-4 provided in
544 Sections C.1-C.4. Our multi-stage training procedure is described explicitly, and pseudocode is
545 given in Algorithm 1. Experimental setups for each task are described in detail, including synthetic
546 LiDAR navigation (Appendix E.3), differentiating single-cell dynamics (Appendix E.4), and cell-state
547 perturbation modeling (Appendix E.5). We specify all hyperparameters and architectures in Table 11,
548 evaluation metrics such as Wasserstein distances and MMD are defined in Equations 95, 96, and 97,
549 and additional loss functions are given in Equations 12-17. Figures 1-5 and Tables 1-5 provide full
550 benchmarking comparisons to baselines. We will release all code, training data splits, and pretrained
551 BranchSBM models upon publication to allow complete reproducibility.
552553 ETHICS STATEMENT
554555 This work focuses on the development of a theoretical and computational framework for learning
556 branched Schrödinger bridges in generative modeling. All experiments are conducted on publicly
557 available or synthetic datasets, including 3D LiDAR manifolds, mouse hematopoiesis scRNA-seq data,
558 and the Tahoe-100M perturbation dataset, and do not involve any personally identifiable, sensitive
559 human subject data, or animal experiments. BranchSBM provides a general-purpose algorithm for
560 modeling stochastic processes with branching dynamics, with potential applications in biological
561 modeling, navigation, and physics. Possible risks, such as misuse in simulating sensitive biomedical
562 data, are mitigated by the fact that BranchSBM is a general generative framework that requires curated
563 training datasets and does not directly generate actionable biological interventions. We believe the
564 potential benefits of this work, in advancing generative modeling theory and enabling more accurate
565 modeling of branching systems in biology and beyond, outweigh possible risks. Code and models
566 will be released under a research-only license to encourage responsible use.
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581
582
583
584
585
586
587
588
589
590
591
592
593

594 REFERENCES
595

596 Sayali Anil Alatkar and Daifeng Wang. Artemis integrates autoencoders and schrodinger bridges to
597 predict continuous dynamics of gene expression, cell population and perturbation from time-series
598 single-cell data. *bioRxiv*, pp. 2025–01, 2025.

599 Georgios Arvanitidis, Lars Kai Hansen, and Søren Hauberg. A locally adaptive normal distribution.
600 *Advances in Neural Information Processing Systems*, 2016.

602 Lazar Atanackovic, Xi Zhang, Brandon Amos, Mathieu Blanchette, Leo J Lee, Yoshua Bengio,
603 Alexander Tong, and Kirill Neklyudov. Meta flow matching: Integrating vector fields on the
604 wasserstein manifold. *arXiv preprint arXiv:2408.14608*, 2024.

605 Jacob Austin, Daniel D. Johnson, Jonathan Ho, Daniel Tarlow, and Rianne van den Berg. Structured
606 denoising diffusion models in discrete state-spaces. *Advances in Neural Information Processing
607 Systems*, 2021.

609 Aymeric Baradat and Hugo Lavenant. Regularized unbalanced optimal transport as entropy mini-
610 mization with respect to branching brownian motion. *arXiv preprint arXiv:2111.01666*, 2021.

611 Bernard Beauzamy. *Introduction to Banach spaces and their geometry*, volume 68. Elsevier, 2011.

613 Barbora Benešová and Martin Kružík. Weak lower semicontinuity of integral functionals and
614 applications. *arXiv*, 2016.

616 Charlotte Bunne, Ya-Ping Hsieh, Marco Cuturi, and Andreas Krause. The schrödinger bridge between
617 gaussian measures has a closed form. *International Conference on Artificial Intelligence and
618 Statistics*, 2022a.

619 Charlotte Bunne, Laetitia Papaxanthos, Andreas Krause, and Marco Cuturi. Proximal optimal
620 transport modeling of population dynamics. In *International Conference on Artificial Intelligence
621 and Statistics*, pp. 6511–6528. PMLR, 2022b.

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

626 Tianrong Chen, Guan-Horng Liu, and Evangelos A. Theodorou. Likelihood training of schrödinger
627 bridge using forward-backward sdes theory. *International Conference on Learning Representations*,
628 2021a.

629 Tianrong Chen, Guan-Horng Liu, Molei Tao, and Evangelos A Theodorou. Deep momentum
630 multi-marginal schrödinger bridge. *arXiv preprint arXiv:2303.01751*, 2023.

632 Yongxin Chen and Tryphon Georgiou. Stochastic bridges of linear systems. *IEEE Transactions on
633 Automatic Control*, 61(2):526–531, 2016. doi: 10.1109/TAC.2015.2440567.

634 Yongxin Chen, Tryphon Georgiou, and Michele Pavon. On the relation between optimal transport
635 and schrödinger bridges: A stochastic control viewpoint. *Journal of Optimization Theory and
636 Applications*, 2016.

638 Yongxin Chen, Tryphon T. Georgiou, and Michele Pavon. The most likely evolution of diffus-
639 ing and vanishing particles: Schrodinger bridges with unbalanced marginals. *arXiv preprint
640 arXiv:2108.02879*, 2021b.

641 Lénaïc Chizat, Gabriel Peyré, Bernhard Schmitzer, and François-Xavier Vialard. Unbalanced optimal
642 transport: Dynamic and kantorovich formulations. *Journal of Functional Analysis*, 274(11):
643 3090–3123, 2018. ISSN 0022-1236.

644 Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transportation distances.
645 *Advances in Neural Information Processing Systems*, 2013.

647 Bernard Dacorogna. *Direct Methods in the Calculus of Variations*. Springer New York, 1989. ISBN
9780387357799.

648 Valentin De Bortoli, James Thornton, Jeremy Heng, and Arnaud Doucet. Diffusion schrödinger
 649 bridge with applications to score-based generative modeling. *Advances in Neural Information*
 650 *Processing Systems*, 2021a.

651 Valentin De Bortoli, James Thornton, Jeremy Heng, and Arnaud Doucet. Diffusion schrödinger
 652 bridge with applications to score-based generative modeling. *Advances in Neural Information*
 653 *Processing Systems*, 2021b.

654 Valentin De Bortoli, Iryna Korshunova, Andriy Mnih, and Arnaud Doucet. Schrodinger bridge
 655 flow for unpaired data translation. *Advances in Neural Information Processing Systems*, 37:
 656 103384–103441, 2024.

657 Atray Dixit, Oren Parnas, Biyu Li, Jenny Chen, Charles P Fulco, Livnat Jerby-Arnon, Nemanja D
 658 Marjanovic, Danielle Dionne, Tyler Burks, Raktima Raychowdhury, et al. Perturb-seq: dissecting
 659 molecular circuits with scalable single-cell rna profiling of pooled genetic screens. *cell*, 167(7):
 660 1853–1866, 2016.

661 Alice Driessen, Benedek Harsanyi, Marianna Rapsomaniki, and Jannis Born. Towards generalizable
 662 single-cell perturbation modeling via the conditional monge gap. *arXiv preprint arXiv:2504.08328*,
 663 2025.

664 Ivar Ekeland and Roger Temam. *Convex analysis and variational problems*. SIAM, 1999.

665 Jhanvi Garg, Xianyang Zhang, and Quan Zhou. Soft-constrained schrödinger bridge: a stochastic
 666 control approach. In *International Conference on Artificial Intelligence and Statistics*, pp. 4429–
 667 4437. PMLR, 2024.

668 George I Gavriilidis, Vasileios Vasileiou, Aspasia Orfanou, Naveed Ishaque, and Fotis Psomopoulos.
 669 A mini-review on perturbation modelling across single-cell omic modalities. *Computational and*
 670 *Structural Biotechnology Journal*, 23:1886, 2024.

671 Nikita Gushchin, Sergei Kholkin, Evgeny Burnaev, and Alexander Korotin. Light and optimal
 672 schrödinger bridge matching. In *Forty-first International Conference on Machine Learning*, 2024.

673 Guillaume Huguet, Daniel Sumner Magruder, Alexander Tong, Oluwadamilola Fasina, Manik
 674 Kuchroo, Guy Wolf, and Smita Krishnaswamy. Manifold interpolating optimal-transport flows for
 675 trajectory inference. *Advances in neural information processing systems*, 35:29705–29718, 2022.

676 Kacper Kapuśnian, Peter Potapchik, Teodora Reu, Leo Zhang, Alexander Tong, Michael Bronstein,
 677 Avishek Joey Bose, and Francesco Di Giovanni. Metric flow matching for smooth interpolations
 678 on the data manifold. *Advances in Neural Information Processing Systems*, 2024.

679 Jun Hyeong Kim, Seonghwan Kim, Seokhyun Moon, Hyeongwoo Kim, Jeheon Woo, and Woo Youn
 680 Kim. Discrete diffusion schrödinger bridge matching for graph transformation. *International*
 681 *Conference on Learning Representations*, 2024.

682 Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *International*
 683 *Conference for Learning Representations*, 2014.

684 Dominik Klein, Théo Uscidda, Fabian Theis, and Marco Cuturi. Genot: Entropic (gromov) wasser-
 685 stein flow matching with applications to single-cell genomics. *Advances in Neural Information*
 686 *Processing Systems*, 37:103897–103944, 2024.

687 Mutsumi Kobayashi, Misato Kobayashi, Junko Odajima, Keiko Shiota, Young Sun Hwang, Kotaro
 688 Sasaki, Pranam Chatterjee, Christian Kramme, Richie E Kohman, George M Church, et al.
 689 Expanding homogeneous culture of human primordial germ cell-like cells maintaining germline
 690 features without serum or feeder layers. *Stem Cell Reports*, 17(3):507–521, 2022.

691 Alexander Korotin, Nikita Gushchin, and Evgeny Burnaev. Light schrödinger bridge. *arXiv preprint*
 692 *arXiv:2310.01174*, 2023.

693 Christian Kramme, Alexandru M. Plesa, Helen H. Wang, Bennett Wolf, Merrick Pierson Smela,
 694 Xiaoge Guo, Richie E. Kohman, Pranam Chatterjee, and George M. Church. An integrated pipeline
 695 for mammalian genetic screening. *Cell Reports Methods*, 1(6):100082, October 2021.

702 Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching
 703 for generative modeling. *International Conference on Learning Representations*, 2023.

704

705 Peter Lippmann, Enrique Fita Sanmartín, and Fred A Hamprecht. Theory and approximate solvers
 706 for branched optimal transport with multiple sources. *Advances in Neural Information Processing
 707 Systems*, 35:267–279, 2022.

708

709 Guan-Horng Liu, Tianrong Chen, Oswin So, and Evangelos A. Theodorou. Deep generalized
 710 schrödinger bridge. *Advances in Neural Information Processing Systems*, 2022.

711

712 Guan-Horng Liu, Yaron Lipman, Maximilian Nickel, Brian Karrer, Evangelos A. Theodorou, and
 713 Ricky T. Q. Chen. Generalized schrödinger bridge matching. *International Conference on Learning
 714 Representations*, 2023a.

715

716 Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and
 717 transfer data with rectified flow. *International Conference on Learning Representations*, 2023b.

718

719 Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *International Conference
 720 for Learning Representations*, 2017.

721

722 Frederike Lübeck, Charlotte Bunne, Gabriele Gut, Jacobo Sarabia del Castillo, Lucas Pelkmans, and
 723 David Alvarez-Melis. Neural unbalanced optimal transport via cycle-consistent semi-couplings.
 724 *arXiv preprint arXiv:2209.15621*, 2022.

725

726 Christian Léonard. A survey of the schrödinger problem and some of its connections with optimal
 727 transport. *Discrete & Continuous Dynamical Systems-A*, 2014.

728

729 Stathis Megas, Daniel G Chen, Krzysztof Polanski, Moshe Eliasof, Carola-Bibiane Schönlieb, and
 730 Sarah A Teichmann. Estimation of single-cell and tissue perturbation effect in spatial transcriptomics
 731 via spatial causal disentanglement. In *The Thirteenth International Conference on Learning
 732 Representations*, 2025.

733

734 Kirill Neklyudov, Rob Brekelsmans, Alexander Tong, Lazar Atanackovic, Qiang Liu, and Alireza
 735 Makhzani. A computational framework for solving wasserstein lagrangian flows. *arXiv preprint
 736 arXiv:2310.10649*, 2023.

737

738 Maxence Noble, Valentin De Bortoli, Arnaud Doucet, and Alain Durmus. Tree-based diffusion
 739 schrödinger bridge with applications to wasserstein barycenters. *Advances in Neural Information
 740 Processing Systems*, 36:55193–55236, 2023.

741

742 Matteo Pariset, Ya-Ping Hsieh, Charlotte Bunne, Andreas Krause, and Valentin De Bortoli. Unbal-
 743 anced diffusion schrödinger bridge. *arXiv preprint arXiv:2306.09099*, 2023.

744

745 Michele Pavon, Esteban G Tabak, and Giulio Trigila. The data-driven schroedinger bridge. *arXiv*,
 746 2018.

747

748 Michele Pavon, Giulio Trigila, and Esteban G Tabak. The data-driven schrödinger bridge. *Commu-
 749 nica-
 750 nica-
 751 nica-
 752 nica-
 753 nica-
 754 nica-
 755 nica-
 756 nica-
 757 nica-
 758 nica-
 759 nica-
 760 nica-
 761 nica-
 762 nica-
 763 nica-
 764 nica-
 765 nica-
 766 nica-
 767 nica-
 768 nica-
 769 nica-
 770 nica-
 771 nica-
 772 nica-
 773 nica-
 774 nica-
 775 nica-
 776 nica-
 777 nica-
 778 nica-
 779 nica-
 780 nica-
 781 nica-
 782 nica-
 783 nica-
 784 nica-
 785 nica-
 786 nica-
 787 nica-
 788 nica-
 789 nica-
 790 nica-
 791 nica-
 792 nica-
 793 nica-
 794 nica-
 795 nica-
 796 nica-
 797 nica-
 798 nica-
 799 nica-
 800 nica-
 801 nica-
 802 nica-
 803 nica-
 804 nica-
 805 nica-
 806 nica-
 807 nica-
 808 nica-
 809 nica-
 810 nica-
 811 nica-
 812 nica-
 813 nica-
 814 nica-
 815 nica-
 816 nica-
 817 nica-
 818 nica-
 819 nica-
 820 nica-
 821 nica-
 822 nica-
 823 nica-
 824 nica-
 825 nica-
 826 nica-
 827 nica-
 828 nica-
 829 nica-
 830 nica-
 831 nica-
 832 nica-
 833 nica-
 834 nica-
 835 nica-
 836 nica-
 837 nica-
 838 nica-
 839 nica-
 840 nica-
 841 nica-
 842 nica-
 843 nica-
 844 nica-
 845 nica-
 846 nica-
 847 nica-
 848 nica-
 849 nica-
 850 nica-
 851 nica-
 852 nica-
 853 nica-
 854 nica-
 855 nica-
 856 nica-
 857 nica-
 858 nica-
 859 nica-
 860 nica-
 861 nica-
 862 nica-
 863 nica-
 864 nica-
 865 nica-
 866 nica-
 867 nica-
 868 nica-
 869 nica-
 870 nica-
 871 nica-
 872 nica-
 873 nica-
 874 nica-
 875 nica-
 876 nica-
 877 nica-
 878 nica-
 879 nica-
 880 nica-
 881 nica-
 882 nica-
 883 nica-
 884 nica-
 885 nica-
 886 nica-
 887 nica-
 888 nica-
 889 nica-
 890 nica-
 891 nica-
 892 nica-
 893 nica-
 894 nica-
 895 nica-
 896 nica-
 897 nica-
 898 nica-
 899 nica-
 900 nica-
 901 nica-
 902 nica-
 903 nica-
 904 nica-
 905 nica-
 906 nica-
 907 nica-
 908 nica-
 909 nica-
 910 nica-
 911 nica-
 912 nica-
 913 nica-
 914 nica-
 915 nica-
 916 nica-
 917 nica-
 918 nica-
 919 nica-
 920 nica-
 921 nica-
 922 nica-
 923 nica-
 924 nica-
 925 nica-
 926 nica-
 927 nica-
 928 nica-
 929 nica-
 930 nica-
 931 nica-
 932 nica-
 933 nica-
 934 nica-
 935 nica-
 936 nica-
 937 nica-
 938 nica-
 939 nica-
 940 nica-
 941 nica-
 942 nica-
 943 nica-
 944 nica-
 945 nica-
 946 nica-
 947 nica-
 948 nica-
 949 nica-
 950 nica-
 951 nica-
 952 nica-
 953 nica-
 954 nica-
 955 nica-
 956 nica-
 957 nica-
 958 nica-
 959 nica-
 960 nica-
 961 nica-
 962 nica-
 963 nica-
 964 nica-
 965 nica-
 966 nica-
 967 nica-
 968 nica-
 969 nica-
 970 nica-
 971 nica-
 972 nica-
 973 nica-
 974 nica-
 975 nica-
 976 nica-
 977 nica-
 978 nica-
 979 nica-
 980 nica-
 981 nica-
 982 nica-
 983 nica-
 984 nica-
 985 nica-
 986 nica-
 987 nica-
 988 nica-
 989 nica-
 990 nica-
 991 nica-
 992 nica-
 993 nica-
 994 nica-
 995 nica-
 996 nica-
 997 nica-
 998 nica-
 999 nica-
 1000 nica-
 1001 nica-
 1002 nica-
 1003 nica-
 1004 nica-
 1005 nica-
 1006 nica-
 1007 nica-
 1008 nica-
 1009 nica-
 1010 nica-
 1011 nica-
 1012 nica-
 1013 nica-
 1014 nica-
 1015 nica-
 1016 nica-
 1017 nica-
 1018 nica-
 1019 nica-
 1020 nica-
 1021 nica-
 1022 nica-
 1023 nica-
 1024 nica-
 1025 nica-
 1026 nica-
 1027 nica-
 1028 nica-
 1029 nica-
 1030 nica-
 1031 nica-
 1032 nica-
 1033 nica-
 1034 nica-
 1035 nica-
 1036 nica-
 1037 nica-
 1038 nica-
 1039 nica-
 1040 nica-
 1041 nica-
 1042 nica-
 1043 nica-
 1044 nica-
 1045 nica-
 1046 nica-
 1047 nica-
 1048 nica-
 1049 nica-
 1050 nica-
 1051 nica-
 1052 nica-
 1053 nica-
 1054 nica-
 1055 nica-
 1056 nica-
 1057 nica-
 1058 nica-
 1059 nica-
 1060 nica-
 1061 nica-
 1062 nica-
 1063 nica-
 1064 nica-
 1065 nica-
 1066 nica-
 1067 nica-
 1068 nica-
 1069 nica-
 1070 nica-
 1071 nica-
 1072 nica-
 1073 nica-
 1074 nica-
 1075 nica-
 1076 nica-
 1077 nica-
 1078 nica-
 1079 nica-
 1080 nica-
 1081 nica-
 1082 nica-
 1083 nica-
 1084 nica-
 1085 nica-
 1086 nica-
 1087 nica-
 1088 nica-
 1089 nica-
 1090 nica-
 1091 nica-
 1092 nica-
 1093 nica-
 1094 nica-
 1095 nica-
 1096 nica-
 1097 nica-
 1098 nica-
 1099 nica-
 1100 nica-
 1101 nica-
 1102 nica-
 1103 nica-
 1104 nica-
 1105 nica-
 1106 nica-
 1107 nica-
 1108 nica-
 1109 nica-
 1110 nica-
 1111 nica-
 1112 nica-
 1113 nica-
 1114 nica-
 1115 nica-
 1116 nica-
 1117 nica-
 1118 nica-
 1119 nica-
 1120 nica-
 1121 nica-
 1122 nica-
 1123 nica-
 1124 nica-
 1125 nica-
 1126 nica-
 1127 nica-
 1128 nica-
 1129 nica-
 1130 nica-
 1131 nica-
 1132 nica-
 1133 nica-
 1134 nica-
 1135 nica-
 1136 nica-
 1137 nica-
 1138 nica-
 1139 nica-
 1140 nica-
 1141 nica-
 1142 nica-
 1143 nica-
 1144 nica-
 1145 nica-
 1146 nica-
 1147 nica-
 1148 nica-
 1149 nica-
 1150 nica-
 1151 nica-
 1152 nica-
 1153 nica-
 1154 nica-
 1155 nica-
 1156 nica-
 1157 nica-
 1158 nica-
 1159 nica-
 1160 nica-
 1161 nica-
 1162 nica-
 1163 nica-
 1164 nica-
 1165 nica-
 1166 nica-
 1167 nica-
 1168 nica-
 1169 nica-
 1170 nica-
 1171 nica-
 1172 nica-
 1173 nica-
 1174 nica-
 1175 nica-
 1176 nica-
 1177 nica-
 1178 nica-
 1179 nica-
 1180 nica-
 1181 nica-
 1182 nica-
 1183 nica-
 1184 nica-
 1185 nica-
 1186 nica-
 1187 nica-
 1188 nica-
 1189 nica-
 1190 nica-
 1191 nica-
 1192 nica-
 1193 nica-
 1194 nica-
 1195 nica-
 1196 nica-
 1197 nica-
 1198 nica-
 1199 nica-
 1200 nica-
 1201 nica-
 1202 nica-
 1203 nica-
 1204 nica-
 1205 nica-
 1206 nica-
 1207 nica-
 1208 nica-
 1209 nica-
 1210 nica-
 1211 nica-
 1212 nica-
 1213 nica-
 1214 nica-
 1215 nica-
 1216 nica-
 1217 nica-
 1218 nica-
 1219 nica-
 1220 nica-
 1221 nica-
 1222 nica-
 1223 nica-
 1224 nica-
 1225 nica-
 1226 nica-
 1227 nica-
 1228 nica-
 1229 nica-
 1230 nica-
 1231 nica-
 1232 nica-
 1233 nica-
 1234 nica-
 1235 nica-
 1236 nica-
 1237 nica-
 1238 nica-
 1239 nica-
 1240 nica-
 1241 nica-
 1242 nica-
 1243 nica-
 1244 nica-
 1245 nica-
 1246 nica-
 1247 nica-
 1248 nica-
 1249 nica-
 1250 nica-
 1251 nica-
 1252 nica-
 1253 nica-
 1254 nica-
 1255 nica-
 1256 nica-
 1257 nica-
 1258 nica-
 1259 nica-
 1260 nica-
 1261 nica-
 1262 nica-
 1263 nica-
 1264 nica-
 1265 nica-
 1266 nica-
 1267 nica-
 1268 nica-
 1269 nica-
 1270 nica-
 1271 nica-
 1272 nica-
 1273 nica-
 1274 nica-
 1275 nica-
 1276 nica-
 1277 nica-
 1278 nica-
 1279 nica-
 1280 nica-
 1281 nica-
 1282 nica-
 1283 nica-
 1284 nica-
 1285 nica-
 1286 nica-
 1287 nica-
 1288 nica-
 1289 nica-
 1290 nica-
 1291 nica-
 1292 nica-
 1293 nica-
 1294 nica-
 1295 nica-
 1296 nica-
 1297 nica-
 1298 nica-
 1299 nica-
 1300 nica-
 1301 nica-
 1302 nica-
 1303 nica-
 1304 nica-
 1305 nica-
 1306 nica-
 1307 nica-
 1308 nica-
 1309 nica-
 1310 nica-
 1311 nica-
 1312 nica-
 1313 nica-
 1314 nica-
 1315 nica-
 1316 nica-
 1317 nica-
 1318 nica-
 1319 nica-
 1320 nica-
 1321 nica-
 1322 nica-
 1323 nica-
 1324 nica-
 1325 nica-
 1326 nica-
 1327 nica-
 1328 nica-
 1329 nica-
 1330 nica-
 1331 nica-
 1332 nica-
 1333 nica-
 1334 nica-
 1335 nica-
 1336 nica-
 1337 nica-
 1338 nica-
 1339 nica-
 1340 nica-
 1341 nica-
 1342 nica-
 1343 nica-
 1344 nica-
 1345 nica-
 1346 nica-
 1347 nica-
 1348 nica-
 1349 nica-
 1350 nica-
 1351 nica-
 1352 nica-
 1353 nica-
 1354 nica-
 1355 nica-
 1356 nica-
 1357 nica-
 1358 nica-
 1359 nica-
 1360 nica-
 1361 nica-
 1362 nica-
 1363 nica-
 1364 nica-
 1365 nica-
 1366 nica-
 1367 nica-
 1368 nica-
 1369 nica-
 1370 nica-
 1371 nica-
 1372 nica-
 1373 nica-
 1374 nica-
 1375 nica-
 1376 nica-
 1377 nica-
 1378 nica-
 1379 nica-
 1380 nica-
 1381 nica-
 1382 nica-
 1383 nica-
 1384 nica-
 1385 nica-
 1386 nica-
 1387 nica-
 1388 nica-
 1389 nica-
 1390 nica-
 1391 nica-
 1392 nica-
 1393 nica-
 1394 nica-
 1395 nica-
 1396 nica-
 1397 nica-
 1398 nica-
 1399 nica-
 1400 nica-
 1401 nica-
 1402 nica-
 1403 nica-
 1404 nica-
 1405 nica-
 1406 nica-
 1407 nica-
 1408 nica-
 1409 nica-
 1410 nica-
 1411 nica-
 1412 nica-
 1413 nica-
 1414 nica-
 1415 nica-
 1416 nica-
 1417 nica-
 1418 nica-
 1419 nica-
 1420 nica-
 1421 nica-
 1422 nica-
 1423 nica-
 1424 nica-
 1425 nica-
 1426 nica-
 1427 nica-
 1428 nica-
 1429 nica-
 1430 nica-
 1431 nica-
 1432 nica-
 1433 nica-
 1434 nica-
 1435 nica-
 1436 nica-
 1437 nica-
 1438 nica-
 1439 nica-
 1440 nica-
 1441 nica-
 1442 nica-
 1443 nica-
 1444 nica-
 1445 nica-
 1446 nica-
 1447 nica-
 1448 nica-
 1449 nica-
 1450 nica-
 1451 nica-
 1452 nica-
 1453 nica-
 1454 nica-
 1455 nica-
 1456 nica-
 1457 nica-
 1458 nica-
 1459 nica-
 1460 nica-
 1461 nica-
 1462 nica-
 1463 nica-
 1464 nica-
 1465 nica-
 1466 nica-
 1467 nica-
 1468 nica-
 1469 nica-
 1470 nica-
 1471 nica-
 1472 nica-
 1473 nica-
 1474 nica-
 1475 nica-
 1476 nica-
 1477 nica-
 1478 nica-
 1479 nica-
 1480 nica-
 1481 nica-
 1482 nica-
 1483 nica-
 1484 nica-
 1485 nica-
 1486 nica-
 1487 nica-
 1488 nica-
 1489 nica-
 1490 nica-
 1491 nica-
 1492 nica-
 1493 nica-
 1494 nica-
 1495 nica-
 1496 nica-
 1497 nica-
 1498 nica-
 1499 nica-
 1500 nica-
 1501 nica-
 1502 nica-
 1503 nica-
 1504 nica-
 1505 nica-
 1506 nica-
 1507 nica-
 1508 nica-
 1509 nica-
 1510 nica-
 1511 nica-
 1512 nica-
 1513 nica-
 1514 nica-
 1515 nica-
 1516 nica-
 1517 nica-
 1518 nica-
 1519 nica-
 1520 nica-
 1521 nica-
 1522 nica-
 1523 nica-
 1524 nica-
 1525 nica-
 1526 nica-
 1527 nica-
 1528 nica-
 1529 nica-
 1530 nica-
 1531 nica-
 1532 nica-
 1533 nica-
 1534 nica-
 1535 nica-
 1536 nica-
 1537 nica-
 1538 nica-
 1539 nica-
 1540 nica-
 1541 nica-
 1542 nica-
 1543 nica-
 1544 nica-
 1545 nica-
 1546 nica-
 1547 nica-
 1548 nica-
 1549 nica-
 1550 nica-
 1551 nica-
 1552 nica-
 1553 nica-
 1554 nica-
 1555 nica-
 1556 nica-
 1557 nica-
 1558 nica-
 1559 nica-
 1560 nica-
 1561 nica-
 1562 nica-
 1563 nica-
 1564 nica-
 1565 nica-
 1566 nica-
 1567 nica-
 1568 nica-
 1569 nica-
 1570 nica-
 1571 nica-
 1572 nica-
 1573 nica-
 1574 nica-
 1575 nica-
 1576 nica-
 1577 nica-
 1578 nica-
 1579 nica-
 1580 nica-
 1581 nica-
 1582 nica-
 1583 nica-
 1584 nica-
 1585 nica-
 1586 nica-
 1587 nica-
 1588 nica-
 1589 nica-
 1590 nica-
 1591 nica-
 1592 nica-
 1593 nica-
 1594 nica-
 1595 nica-
 1596 nica-
 1597 nica-
 1598 nica-
 1599 nica-
 1600 nica-
 1601 nica-
 1602 nica-
 1603 nica-
 1604 nica-
 1605 nica-
 1606 nica-
 1607 nica-
 1608 nica-
 1609 nica-
 1610 nica-
 1611 nica-
 1612 nica-
 1613 nica-
 1614 nica-
 16*

756 Yusuf Roohani, Kexin Huang, and Jure Leskovec. Predicting transcriptional outcomes of novel
 757 multigene perturbations with gears. *Nature Biotechnology*, 42(6):927–935, 2023.
 758

759 Jayoung Ryu, Charlotte Bunne, Luca Pinello, Aviv Regev, and Romain Lopez. Cross-modality
 760 matching and prediction of perturbation responses with labeled gromov-wasserstein optimal
 761 transport. *arXiv preprint arXiv:2405.00838*, 2024.

762 Geoffrey Schiebinger, Jian Shu, Marcin Tabaka, Brian Cleary, Vidya Subramanian, Aryeh Solomon,
 763 Joshua Gould, Siyan Liu, Stacie Lin, Peter Berube, et al. Optimal-transport analysis of single-cell
 764 gene expression identifies developmental trajectories in reprogramming. *Cell*, 176(4):928–943,
 765 2019.

766 Erwin Schrödinger. *Über die umkehrung der naturgesetze*. Verlag der Akademie der Wissenschaften
 767 in Kommission bei Walter De Gruyter u . . ., 1931.

768 Yutong Sha, Yuchi Qiu, Peijie Zhou, and Qing Nie. Reconstructing growth and dynamic trajectories
 769 from single-cell transcriptomics data. *Nature Machine Intelligence*, 6(1):25–39, 2023.

770 Ophir Shalem, Neville E Sanjana, Ella Hartenian, Xi Shi, David A Scott, Tarjei S Mikkelsen, Dirk
 771 Heckl, Benjamin L Ebert, David E Root, John G Doench, et al. Genome-scale crispr-cas9 knockout
 772 screening in human cells. *Science*, 343(6166):84–87, 2014.

773 Yunyi Shen, Renato Berlinghieri, and Tamara Broderick. Multi-marginal schr\" odinger bridges with
 774 iterative reference refinement. *arXiv preprint arXiv:2408.06277*, 2024.

775 Yuyang Shi, Valentin De Bortoli, Andrew Campbell, and Arnaud Doucet. Diffusion schrödinger
 776 bridge matching. *Advances in Neural Information Processing Systems*, 2023.

777 Merrick D Pierson Smela, Christian C Kramme, Patrick RJ Fortuna, Jessica L Adams, Rui Su,
 778 Edward Dong, Mutsumi Kobayashi, Garyk Brixi, Venkata Srikanth Kavirayuni, Emma Tysinger, et al.
 779 Directed differentiation of human ipscs to functional ovarian granulosa-like cells via transcription
 780 factor overexpression. *Elife*, 12:e83291, 2023.

781 Vignesh Ram Somnath, Matteo Pariset, Ya-Ping Hsieh, Maria Rodriguez Martinez, Andreas Krause,
 782 and Charlotte Bunne. Aligned diffusion schrödinger bridges. *Uncertainty in Artificial Intelligence*,
 783 2023.

784 J Michael Steele. *The Cauchy-Schwarz master class: an introduction to the art of mathematical
 785 inequalities*. Cambridge University Press, 2004.

786 Panagiotis Theodoropoulos, Nikolaos Komianos, Vincent Pacelli, Guan-Horng Liu, and Evangelos A.
 787 Theodorou. Feedback schrödinger bridge matching. *The Thirteenth International Conference on
 788 Learning Representations*, 2024.

789 Panagiotis Theodoropoulos, Augustinos D Saravacos, Evangelos A Theodorou, and Guan-Horng Liu.
 790 Momentum multi-marginal schrödinger bridge matching. *arXiv preprint arXiv:2506.10168*, 2025.

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

794 Alexander Tong, Kilian Fatras, Nikolay Malkin, Guillaume Huguet, Yanlei Zhang, Jarrid Rector-
 795 Brooks, Guy Wolf, and Yoshua Bengio. Improving and generalizing flow-based generative models
 796 with minibatch optimal transport. *Transactions on Machine Learning Research*, 2024a.

797 Alexander Y. Tong, Nikolay Malkin, Kilian Fatras, Lazar Atanackovic, Yanlei Zhang, Guillaume
 798 Huguet, Guy Wolf, and Yoshua Bengio. Simulation-free Schrödinger bridges via score and
 799 flow matching. *Proceedings of The 27th International Conference on Artificial Intelligence and
 800 Statistics*, 238:1279–1287, 2024b.

801 Francisco Vargas, Pierre Thodoroff, Neil D. Lawrence, and Austen Lamacraft. Solving schrödinger
 802 bridges via maximum likelihood. *arXiv preprint arXiv:2106.02081*, 2021.

810 Adrian Veres, Aubrey L Faust, Henry L Bushnell, Elise N Engquist, Jennifer Hyoje-Ryu Kenty,
 811 George Harb, Yeh-Chuin Poh, Elad Sintov, Mads Görtler, Felicia W Pagliuca, et al. Charting
 812 cellular identity during human in vitro β -cell differentiation. *Nature*, 569(7756):368–373, 2019.

813

814 Cédric Villani et al. *Optimal transport: old and new*, volume 338. Springer, 2008.

815 Dongyi Wang, Yuanwei Jiang, Zhenyi Zhang, Xiang Gu, Peijie Zhou, and Jian Sun. Joint velocity-
 816 growth flow matching for single-cell dynamics modeling. *arXiv preprint arXiv:2505.13413*,
 817 2025.

818

819 Gefei Wang, Yuling Jiao, Qian Xu, Yang Wang, and Can Yang. Deep generative learning via
 820 schrödinger bridge. *International Conference on Machine Learning*, 2021.

821 Caleb Weinreb, Alejo Rodriguez-Fraticelli, Fernando D. Camargo, and Allon M. Klein. Lineage
 822 tracing on transcriptional landscapes links state to fate during differentiation. *Science*, 367(6479),
 823 2020.

824

825 Toshiaki Yachimura, Hanbo Wang, Yusuke Imoto, Momoko Yoshida, Sohei Tasaki, Yoji Kojima,
 826 Yukihiro Yabuta, Mitinori Saitou, and Yasuaki Hiraoka. scegot: single-cell trajectory inference
 827 framework based on entropic gaussian mixture optimal transport. *BMC bioinformatics*, 25(1):388,
 828 2024.

829

830 Grace Hui Ting Yeo, Sachit D Saksena, and David K Gifford. Generative modeling of single-
 831 cell time series with prescient enables prediction of cell trajectories with interventions. *Nature
 832 communications*, 12(1):3222, 2021.

833

834 Jesse Zhang, Airol A Ubas, Richard de Borja, Valentine Svensson, Nicole Thomas, Neha Thakar,
 835 Ian Lai, Aidan Winters, Umair Khan, and Matthew G. et al. Jones. Tahoe-100m: A giga-scale
 836 single-cell perturbation atlas for context-dependent gene function and cellular modeling. *bioRxiv*,
 837 2025a.

838

839 Stephen Zhang, Anton Afanassiev, Laura Greenstreet, Tetsuya Matsumoto, and Geoffrey Schiebinger.
 840 Optimal transport analysis reveals trajectories in steady-state systems. *PLoS computational biology*,
 841 17(12):e1009466, 2021.

842

843 Yuhui Zhang, Yuchang Su, Chenyu Wang, Tianhong Li, Zoe Wefers, Jeffrey Nirschl, James Burgess,
 844 Daisy Ding, Alejandro Lozano, Emma Lundberg, et al. Cellflow: Simulating cellular morphology
 845 changes via flow matching. *arXiv preprint arXiv:2502.09775*, 2025b.

846

847 Zhenyi Zhang, Tiejun Li, and Peijie Zhou. Learning stochastic dynamics from snapshots through
 848 regularized unbalanced optimal transport. *International Conference on Learning Representations*,
 849 2025c.

850

851

852

853

854

855

856

857

858

859

860

861

862

863

864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917
OUTLINE OF APPENDIX

In Appendix A, we provide an extended background on the relevant theory for learning optimal stochastic bridges (A.1) and simulating trajectories on the data manifold (A.2). In Appendix B, we discuss the relationship between our proposed formulation for BranchSBM and previous related works. Appendix C provides the theoretical basis for Sections 3 and 4, including formal proofs for Proposition 1 (C.1), Proposition 2 (C.2), Proposition 3 (C.3), and Proposition 4 (C.4). In Appendix E, we describe further details on experiments and hyperparameters used including specific details for each experiment, including multi-path LiDAR navigation (E.3), modeling differentiating single-cells (E.4), and modeling cell-state perturbations (E.4). Finally, we provide the pseudo code for the multi-stage training algorithm in Appendix F.

Notation We denote the state space as $\mathcal{X} \subseteq \mathbb{R}^d$ and time interval as $t \in [0, 1]$. The branches are indexed with $k \in \{0, \dots, K\}$. We denote initial data distribution at time $t = 0$ as π_0 and the terminal data distributions at $t = 1$ as $\{\pi_{1,k}\}_{k=0}^K$. The joint data distribution is denoted $\pi_{0,1,k}$ and a pair of samples is given by $(\mathbf{x}_0, \mathbf{x}_{1,k}) \sim \pi_{0,1,k}$. For simplicity, we denote $d(\mathbf{x}_0, \mathbf{x}_{1,k}) = d\pi_{0,1,k}$. Let $u_{t,k}(X_t)$ denote the marginal velocity field, $g_{t,k}(X_t)$ denote the growth rate, and $p_{t,k}(X_t)$ denote the marginal probability density, where we sometimes drop the input X_t for simplicity. In addition, we denote the conditional velocity field and probability density as $u_{t|0,1,k} \equiv u_{t|0,1,k}(X_t | \mathbf{x}_0, \mathbf{x}_{1,k})$ and $p_{t|0,1,k} \equiv p_{t|0,1,k}(X_t | \mathbf{x}_0, \mathbf{x}_{1,k})$ respectively. The optimal values for any quantity are superscripted with a \star symbol. We denote the parameterized flow neural networks as $u_{t,k}^\theta$ with parameters θ and the growth neural networks with $g_{t,k}^\phi$ with parameters ϕ . In the context of unbalanced endpoint distributions, we denote the true initial weight of a sample as w_0^* and the final weight of a sample from the k th target distribution as $w_{1,k}^*$. The predicted weights generated from the growth dynamics are given by $w_t(X_t)$, and we seek to match the predicted weight at time $t = 1$ given by $w_{1,k}(X_{1,k})$ to the true weight $w_{1,k}^*$. L^2 denotes the space of square integrable functions and $\|\cdot\|_{L^2}$ be the L^2 -norm in function space. L^∞ denotes the space of essentially bounded functions such that $\|f\|_\infty < \infty$.

A EXTENDED THEORETICAL BACKGROUND

A.1 LEARNING OPTIMAL STOCHASTIC BRIDGES

Pinned-Down Stochastic Bridges Let $\mathbb{Q} \in \mathcal{M}$ be a Markovian reference path measure that evolves over $t \in [0, 1]$ according to a drift field $f_t(X_t) : \mathbb{R}^d \rightarrow \mathbb{R}^d$ and stochastic d -dimensional Brownian motion $B_t \in \mathbb{R}^d$, via the SDE

$$dX_t = f_t(X_t)dt + \sigma_t dB_t, \quad X_0 \sim \pi_0 \quad (20)$$

Given \mathbb{Q} , consider a stochastic process $(X_t)_{t \in [0,1]}$ over the time interval $t \in [0, 1]$ pinned-down at the initial point $X_0 = \mathbf{x}_0$ and final point $X_1 = \mathbf{x}_1$ denoted as $\mathbb{Q}_{\cdot|0,1}(\cdot | \mathbf{x}_0, \mathbf{x}_1)$. Due to the endpoint conditions, $\mathbb{Q}_{\cdot|0,1}$ is not necessarily Markov, and evolves via the SDE

$$dX_t = \{f_t(X_t) + \sigma_t^2 \nabla_{\mathbf{x}} \log \mathbb{Q}_{1|t}(\mathbf{x}_1 | X_t)\}dt + \sigma_t dB_t, \quad X_0 = \mathbf{x}_0 \quad (21)$$

where $\nabla_{\mathbf{x}} \log \mathbb{Q}_{1|t}(\mathbf{x}_1 | X_t)$ is a non-Markovian score function that corrects the drift field $f_t(X_t)$ of the reference process such that it points toward the target endpoint \mathbf{x}_1 . Since $\log \mathbb{Q}_{1|t}(\mathbf{x}_1 | X_t)$ is the log-likelihood that the final state satisfies the condition $X_1 = \mathbf{x}_1$, the gradient defines how the log-likelihood changes with respect to the changes in the state \mathbf{x} at time t . The drift moves \mathbf{x} in the direction of the largest increase in log-likelihood given by the score function, which ensures that the process satisfies $X_1 = \mathbf{x}_1$ following the theory of Doob's h -transform (Rogers & Williams, 2000). Now, we can define the conditional probability distribution p_t as a mixture of pinned-down stochastic bridges over pairs of endpoints in the data coupling $\pi_{0,1} = \pi_0 \otimes \pi_1$ given by

$$p_t(\mathbf{x}) = \pi_{0,1} \mathbb{Q}_{t|0,1}(\mathbf{x} | \mathbf{x}_0, \mathbf{x}_1) = \int \mathbb{Q}_{t|0,1}(\mathbf{x} | \mathbf{x}_0, \mathbf{x}_1) d\pi_{0,1} \quad (22)$$

To simplify notation, we denote each conditional bridge as $\mathbb{Q}_{t|0,1}(X_t | \mathbf{x}_0, \mathbf{x}_1) = p_{t|0,1}(X_t | \mathbf{x}_0, \mathbf{x}_1)$ and the joint distribution $\pi_{0,1} = p_{0,1}(\mathbf{x}_0, \mathbf{x}_1)$. Now, we can rewrite the marginal p_t as

$$p_t(X_t) = \int p_{t|0,1}(X_t | \mathbf{x}_0, \mathbf{x}_1) p_{0,1} d(\mathbf{x}_0, \mathbf{x}_1) = \mathbb{E}_{p_{0,1}} [p_{t|0,1}(X_t | \mathbf{x}_0, \mathbf{x}_1)] \quad (23)$$

Furthermore, we denote $u_{t|0,1} \equiv u_{t|0,1}(X_t|\mathbf{x}_0, \mathbf{x}_1) = \{f_t(X_t) + \sigma_t^2 \nabla_{\mathbf{x}} \log \mathbb{Q}_{1|t}(\mathbf{x}_1|X_t)\}$ as the conditional drift that generates $p_{t|0,1} \equiv p_{t|0,1}(X_t|\mathbf{x}_0, \mathbf{x}_1)$, that satisfies the conditional Fokker-Planck equation

$$\frac{\partial}{\partial t} p_{t|0,1} = -\nabla \cdot (u_{t|0,1} p_{t|0,1}) + \frac{1}{2} \sigma^2 \Delta p_{t|0,1} \quad (24)$$

Definition 1 (Reciprocal Class). *Given our definition of the conditional bridge $\mathbb{Q}_{\cdot|0,1}$, we can define the reciprocal class, denoted $\mathcal{R}(\mathbb{Q})$, of the reference measure \mathbb{Q} as the class of path measures that share the same bridge as \mathbb{Q} , defined as $\mathcal{R}(\mathbb{Q}) = \{\Pi \mid \Pi(X_t|\mathbf{x}_0, \mathbf{x}_1) = \mathbb{Q}(X_t|\mathbf{x}_0, \mathbf{x}_1)\}$.*

Markovian Projections Given a mixture of conditional stochastic bridges $\Pi = \Pi_{0,1} \mathbb{Q}_{\cdot|0,1}$ under the reference measure \mathbb{Q} that require knowledge of the joint distribution $p_{0,1}$, we aim to project Π to the space of Markovian measures \mathcal{M} , where the drift dynamics $u_t(X_t)$ is only dependent on the current state X_t and require no knowledge on the endpoints. This allows us to parameterize a Markovian drift $u_t^\theta(X_t)$ that can transport samples from the initial distribution $\mathbf{x}_0 \sim \pi_0$ to samples from the target distribution $\mathbf{x}_1 \sim \pi_1$. To do this, we define the *Markovian projection* of Π (Shi et al., 2023; Liu et al., 2023b).

Definition 2 (Markovian Projection). *Given a conditional bridge $\mathbb{Q}_{\cdot|0,1}$ that evolves via the SDE in (21), we define a Markovian projection of the mixture of bridges $\Pi = \Pi_{0,1} \mathbb{Q}_{\cdot|0,1}$ as a Markov process $\mathbb{M}^* = \text{proj}_{\mathcal{M}}(\Pi) \in \mathcal{M}$ with the same marginals as Π such that $X_t \sim \Pi_t$ for all $t \in [0, 1]$, $X_1 \sim \pi_1$ and evolves via the SDE*

$$dX_t = \{f_t(X_t) + v_t^*(X_t)\} dt + \sigma_t dB_t \quad (25)$$

$$v_t^*(X_t) = \sigma^2 \mathbb{E}_{\Pi_{1|t}} [\nabla_{\mathbf{x}_t} \log \mathbb{Q}_{1|t}(X_1|X_t) | X_t = \mathbf{x}_t] \quad (26)$$

where $\Pi_{1|t}$ is the conditional distribution of X_1 under the mixture of bridges Π and $\nabla_{\mathbf{x}_t} \log \mathbb{Q}_{1|t}(X_1|X_t)$ points in the direction of greatest increase in the log-likelihood of the target endpoint $X_1 \sim \pi_1$ under the reference process \mathbb{Q} . In addition, the Markov projection \mathbb{M}^* minimizes the KL-divergence with the mixture of bridges $\mathbb{M}^* = \arg \min_{\mathbb{M} \in \mathcal{M}} \text{KL}(\Pi \parallel \mathbb{M})$ and can be obtained by parameterizing $v_t^\theta(X_t)$ and minimizing the dynamic formulation given by

$$\text{KL}(\Pi \parallel \mathbb{M}) = \mathbb{E}_{(\mathbf{x}_0, \mathbf{x}_1) \sim \Pi_{0,1}} \mathbb{E}_{\mathbf{x}_t \sim \Pi_{t|0,1}} \int_0^1 \frac{1}{2\sigma_t^2} \left[\|\sigma_t^2 \nabla_{\mathbf{x}_t} \log \mathbb{Q}_{1|t}(\mathbf{x}_1|\mathbf{x}_t) - v_t^\theta(\mathbf{x}_t)\|^2 \right] dt \quad (27)$$

In general, the Markovian projection of a reference measure \mathbb{Q} does not preserve the conditional bridge and is not in the reciprocal class $\mathcal{R}(\mathbb{Q})$. The unique path measure \mathbb{P} that is the Markovian projection of \mathbb{Q} , is in the reciprocal class $\mathcal{R}(\mathbb{Q})$, and preserves the endpoint distributions is called the *Schrödinger Bridge*.

Definition 3 (Schrödinger Bridge). *Given a reference measure \mathbb{Q} , a initial distribution π_0 , and final distribution π_1 . A path measure \mathbb{P} is the unique Schrödinger bridge if it satisfies*

1. \mathbb{P} is the Markovian projection of \mathbb{Q} , such that $\mathbb{P} = \text{proj}_{\mathcal{M}}(\mathbb{Q})$.
2. \mathbb{P} is in the reciprocal class of \mathbb{Q} , i.e. $\mathbb{P} \in \mathcal{R}(\mathbb{Q})$, such that it preserves the conditional bridge $\mathbb{P}(X_t|\mathbf{x}_0, \mathbf{x}_1) = \mathbb{Q}(X_t|\mathbf{x}_0, \mathbf{x}_1)$.
3. \mathbb{P} preserves the endpoint distributions $\mathbb{P}_0 = \pi_0$ and $\mathbb{P}_1 = \pi_1$.

The goal of Schrödinger Bridge Matching (SBM) is to estimate the Schrödinger Bridge that transports samples from an initial distribution π_0 to a final distribution π_1 given the optimal reference dynamics. We further discuss previous approaches to solving the SB problem in Appendix B.

A.2 SIMULATING TRAJECTORIES ON THE DATA MANIFOLD

Riemannian Manifolds and Metrics Since the interpolant $\mathbf{x}_{t,\eta}$ learned in Stage 1 is defined by minimizing a non-linear state cost $V_t(X_t)$, the resulting velocity field $u_t^\theta(X_t)$ lies on the tangent bundle $\mathcal{T}_{\mathbf{x}} \Omega$ of a smooth d -dimensional manifold $\Omega \in \mathbb{R}^d$ called a *Riemannian manifold*. Intuitively, a Riemannian manifold can be thought of as a smooth surface where the local curvature around a point $\mathbf{x} \in \Omega$ can be approximated by a tangent space $\mathcal{T}_{\mathbf{x}} \Omega$ that defines the set of directions in which

972 \mathbf{x} can move along the manifold. These *directions* are defined by a set of tangent vectors $u \in \mathcal{T}_{\mathbf{x}}\Omega$,
 973 which *pushes* the point \mathbf{x} along the manifold.
 974

975 In Stage 2, we seek to parameterize a vector field $u_t^\theta(X_t)$ that minimizes the *angle* from the tangent
 976 vector $\dot{\mathbf{x}}_{t,\eta} = \partial_t \mathbf{x}_{t,\eta}$ at each point \mathbf{x} on the manifold optimized in Stage 1. To do this, we must define
 977 the concepts of *length* and *angles* on Riemannian manifolds. First, we define a *location-dependent*
 978 *inner product* in Riemannian manifolds known as the *Riemannian metric* $g_{\mathbf{x}} : \mathcal{T}_{\mathbf{x}}\Omega \times \mathcal{T}_{\mathbf{x}}\Omega \rightarrow \mathbb{R}$ that
 979 maps two vectors $u, v \in \mathcal{T}_{\mathbf{x}}\Omega$ to a scalar that describes the relative direction and length of the two
 980 vectors. Formally, the Riemannian metric can be written as the *bilnear* and *positive definite* function
 981

$$g_{\mathbf{x}}(u, v) = u^\top \mathbf{G}(\mathbf{x})v = \langle u, \mathbf{G}v \rangle \text{ s.t. } \begin{cases} \forall u \neq 0 \quad g_{\mathbf{x}}(u, u) > 0 \\ \mathbf{G} \succ 0 \end{cases} \quad (28)$$

983 which defines the norm of a tangent vector as $\|u\|_{g_{\mathbf{x}}} = \sqrt{g_{\mathbf{x}}(u, u)}$. Now, we can decompose the
 984 Riemannian norm of the tangent vector $\|\dot{\mathbf{x}}_{t,\eta}\|_{g_{\mathbf{x}}}$ to get the loss defined in (12) as follows
 985

$$\mathcal{L}_{\text{traj}}(\eta) = \mathbb{E}_{t, (\mathbf{x}_0, \mathbf{x}_1) \sim p_{0,1}} [\|\dot{\mathbf{x}}_{t,\eta}\|_{g_{\mathbf{x}}}^2] \quad (29)$$

$$= \mathbb{E}_{t, (\mathbf{x}_0, \mathbf{x}_1) \sim p_{0,1}} \langle \dot{\mathbf{x}}_{t,\eta}, \mathbf{G}(\mathbf{x}_{t,\eta}) \dot{\mathbf{x}}_{t,\eta} \rangle \quad (30)$$

$$= \mathbb{E}_{t, (\mathbf{x}_0, \mathbf{x}_1) \sim p_{0,1}} [\|\dot{\mathbf{x}}_{t,\eta}\|_2^2 + \langle \dot{\mathbf{x}}_{t,\eta}, (\mathbf{G}(\mathbf{x}_{t,\eta}) - \mathbf{I}) \dot{\mathbf{x}}_{t,\eta} \rangle] \quad (31)$$

$$= \mathbb{E}_{t, (\mathbf{x}_0, \mathbf{x}_1) \sim p_{0,1}} [\|\dot{\mathbf{x}}_{t,\eta}\|_2^2 + V_t(\mathbf{x}_{t,\eta})] \quad (32)$$

991 where the state cost is defined as $V_t(\mathbf{x}_{t,\eta}) = \langle \dot{\mathbf{x}}_{t,\eta}, (\mathbf{G}(\mathbf{x}_{t,\eta}) - \mathbf{I}) \dot{\mathbf{x}}_{t,\eta} \rangle$.
 992

993 **Data-Dependent State Cost** Following Kapuśnian et al. (2024), we define the metric matrix $\mathbf{G}(\mathbf{x}_{t,\eta})$ described previously as the data-dependent LAND and RBF metrics of the form
 994 $\mathbf{G}_{\text{LAND}}(\mathbf{x}, \mathcal{D}) = \mathbf{G}_{\text{RBF}}(\mathbf{x}, \mathcal{D}) = (\text{diag}(\mathbf{h}(\mathbf{x})) - \varepsilon \mathbf{I})^{-1}$ which assigns higher cost (i.e. $\|\mathbf{G}(\mathbf{x})\|$
 995 is larger) when \mathbf{x} moves away from the support of the dataset \mathcal{D} . Specifically, given a dataset
 996 $\mathcal{D} = \{\mathbf{x}_i\}_{i=1}^N$, we define the elements $\mathbf{h}^{\text{LAND}}(\mathbf{x}) \in \mathbb{R}^d$ that scales down each dimension
 997 $j \in \{1, \dots, d\}$ in the LAND metric as
 998

$$h_j^{\text{LAND}}(\mathbf{x}) = \sum_{i=1}^N (x_i^j - x^j)^2 \exp\left(-\frac{\|\mathbf{x} - \mathbf{x}_i\|^2}{2\sigma^2}\right) \quad (33)$$

1001 where the \exp term is positive when there is a high concentration of data points around the point
 1002 \mathbf{x} (i.e. $\|\mathbf{x} - \mathbf{x}_i\|$ is small) and approaches 0 as the concentration of data around \mathbf{x} decreases (i.e.
 1003 $\|\mathbf{x} - \mathbf{x}_i\|$ is large). Writing $\langle \dot{\mathbf{x}}_{t,\eta}, \mathbf{G}(\mathbf{x}_{t,\eta}) \dot{\mathbf{x}}_{t,\eta} \rangle$ in terms of $h_j(\mathbf{x})$, we get
 1004

$$\mathcal{L}_{\text{traj}}(\eta) = \mathbb{E}_{t, (\mathbf{x}_0, \mathbf{x}_1) \sim p_{0,1}} \langle \dot{\mathbf{x}}_{t,\eta}, \mathbf{G}(\mathbf{x}_{t,\eta}) \dot{\mathbf{x}}_{t,\eta} \rangle \quad (34)$$

$$= \mathbb{E}_{t, (\mathbf{x}_0, \mathbf{x}_1) \sim p_{0,1}} \left[\sum_{j=1}^d \frac{(\dot{\mathbf{x}}_{t,\eta})_j^2}{h_j(\mathbf{x}_{t,\eta}) + \varepsilon} \right] \quad (35)$$

1009 When $h_j(\mathbf{x})$ is large, the loss is minimized, and when $h_j(\mathbf{x})$ is small, the loss is scaled up. While
 1010 the LAND metric effectively defines the data manifold in low dimensions, in high-dimensional state
 1011 spaces, setting a suitable variance σ in $h_j^{\text{LAND}}(\mathbf{x}_t)$ to ensure that the path does not deviate far from
 1012 the data manifold without overfitting is challenging. To overcome this limitation, the RBF metric
 1013 clusters the dataset into N_c clusters with centroids denoted as $\hat{\mathbf{x}}_n \in \mathbb{R}^d$ and trains a set of parameters
 1014 $\{\omega_{\alpha,n}\}_{n=1}^{N_c}$ to enforce $h_j(\mathbf{x}_i) \approx 1$ for all points in the dataset such that points \mathbf{x} within the data
 1015 manifold are also assigned $h_j(\mathbf{x}) \approx 1$. Specifically, h_j^{RBF} is defined as
 1016

$$h_j^{\text{RBF}}(\mathbf{x}) = \sum_{n=1}^{N_c} \omega_{n,j}(\mathbf{x}) \exp\left(-\frac{\lambda_{n,j}}{2} \|\mathbf{x} - \hat{\mathbf{x}}_n\|^2\right) \quad (36)$$

$$\lambda_n = \frac{1}{2} \left(\frac{\kappa}{|C_n|} \sum_{\mathbf{x} \in C_n} \|\mathbf{x} - \hat{\mathbf{x}}_n\|^2 \right)^{-2} \quad (37)$$

1022 where C_n denotes the n th cluster, κ is a tunable hyperparameter, and $\lambda_{n,j}$ is the bandwidth of cluster
 1023 n the j th dimension. To train the parameters, we minimize the following loss
 1024

$$\mathcal{L}_{\text{RBF}}(\{\omega_{\alpha,n}\}) = \sum_{\mathbf{x}_i \in \mathcal{D}} (1 - h_n^{\text{RBF}}(\mathbf{x}_i))^2 \quad (38)$$

1026 In our experiments, we use the LAND metric for the LiDAR and mouse hematopoiesis datasets
 1027 with dimensions $d = 2$ and $d = 3$ respectively, and the RBF metric for the perturbation modeling
 1028 experiment with gene expression data of dimensions $d \in \{50, 100, 150\}$.
 1029

1030 **Sampling on the Data Manifold** In Riemannian geometry, each Euclidean step $\Delta t \cdot u_t^\theta(\mathbf{x})$ following
 1031 the tangent vector $u_t^\theta(\mathbf{x})$ along the manifold requires mapping back to the manifold via an exponential
 1032 map $X_{t+\Delta t} = \exp_{\mathbf{x}}(\Delta t \cdot u_t^\theta(\mathbf{x}))$. In general, computing the exponential map $\exp_{\mathbf{x}}$ under the
 1033 Riemannian metric $\mathbf{G}(\mathbf{x})$ requires simulation of the geodesic flow (i.e., approximating the final state
 1034 at $t = 1$ under the initial conditions $\mathbf{x}_0 = \mathbf{x}$ and $\dot{\mathbf{x}}_0 = u_t^\theta(\mathbf{x})$).
 1035

1036 While the manifold defined by the data-dependent metric $\mathbf{G}(\mathbf{x}, \mathcal{D})$ induces a Riemannian geometry
 1037 in Euclidean space that follows an optimal cost structure, its underlying space is still Euclidean.
 1038 $\mathbf{G}(\mathbf{x}, \mathcal{D})$ just assigns varying costs of moving in the Euclidean space \mathbb{R}^d . Therefore, we can avoid
 1039 computing the exponential map and generate trajectories with simple Euclidean Euler integration
 1039 following

$$x_{t+\Delta t} = x_t + \Delta t \cdot u_t^\theta(x_t) \quad (39)$$

1040 where $\Delta t = 1/N_{\text{steps}}$ is the discretized step size.
 1041

1043 B COMPARISON TO EXISTING WORKS

1044 In this section, we discuss the relationship between our proposed formulation for Branched
 1045 Schrödinger Bridge Matching and related previous works. We establish reasons why BranchSBM is
 1046 the theoretically optimal formulation to solve the problem of modeling stochastic dynamical systems
 1047 with diverging trajectories over time by modeling branched Schrödinger bridges. We conclude by
 1048 introducing an alternative perspective of the Branched GSB problem defined in Section 3.2 as the
 1049 problem of modeling probabilistic trajectories of dynamic systems with nondeterministic states,
 1050 which BranchSBM is well-suited to solve.
 1051

1053 B.1 MODELING BRANCHED SCHRÖDINGER BRIDGES

1054 **Schrödinger Bridge Matching** Computational methods for solving the Schrödinger Bridge (SB)
 1055 problem for predicting trajectories between initial and target distributions have been extensively
 1056 studied in existing literature (De Bortoli et al., 2021a; Chen et al., 2021b; Korotin et al., 2023; Bunne
 1057 et al., 2022a; Chizat et al., 2018; Liu et al., 2023a; 2022; Shi et al., 2023; Kim et al., 2024; Wang
 1058 et al., 2021; Tong et al., 2024b; Peluchetti, 2024; Bunne et al., 2022a; Somnath et al., 2023; Gushchin
 1059 et al., 2024; Pavon et al., 2021; Garg et al., 2024; De Bortoli et al., 2024; Shen et al., 2024; Chen
 1060 et al., 2023; Noble et al., 2023). Previous work has framed the SB problem as an entropy-regularized
 1061 Optimal Control (EOT) problem (Cuturi, 2013; Léonard, 2014; Pavon et al., 2018) or a stochastic
 1062 optimal control (SOC) problem (Chen et al., 2016; 2021a; Liu et al., 2023a), which we build on in
 1063 this work.
 1064

1065 **Conditional Stochastic Optimal Control** Several works (Chen et al., 2016; 2021a; Liu et al.,
 1066 2023a) have reframed the SB problem as a Conditional Stochastic Optimal Control (CondSOC)
 1067 problem, which takes the canonical form

$$\min_{u_t} \int_0^1 \mathbb{E}_{p_t} \left[\frac{1}{2} \|u_t(X_t)\|^2 + V_t(X_t) \right] dt + \mathbb{E}_{p_1} [\phi(X_1)] \quad (40)$$

$$\text{s.t. } dX_t = u_t(X_t)dt + dB_t, \quad X_0 \sim \pi_0 \quad (41)$$

1068 where $\phi(X_1) : \mathcal{X} \rightarrow \mathbb{R}$ acts as a reconstruction loss that enforces that the distribution of $X_1 \sim p_1$
 1069 matches the true distribution π_1 . Due to the intractability of π_0, π_1 , GSBM (Liu et al., 2023a) uses
 1070 spline optimization to learn an optimal Gaussian probability path $p_{t|0,1}^* = \mathcal{N}(\mu_t, \gamma_t^2 \mathbf{I}_d)$ using only
 1071 samples $\mathbf{x}_0 \sim \pi_0$ and $\mathbf{x}_1 \sim \pi_1$ from the initial and terminal distributions. Although GSBM
 1072 does not require knowledge of the densities, it follows an iterative optimization scheme that alternates
 1073 between matching the drift u_t given a fixed marginal p_t , and updating the marginal given the drift.
 1074 This strategy can get stuck in suboptimal solutions and is sensitive to the initialization of u_t^θ . GSBM
 1075 is also limited to learning unimodal Gaussian paths between one source and one target distribution
 1076 with balanced mass, making it unsuitable for modeling tasks with multimodal terminal distributions
 1077 and splitting of mass over multiple distinct paths.
 1078

1080
 1081 **Regularized Unbalanced Optimal Transport** Several previous works have studied the unbalanced
 1082 optimal transport problem (Zhang et al., 2025c; Chen et al., 2021b; Lübeck et al., 2022; Pariset et al.,
 1083 2023); however, these approaches address a fundamentally different setting from the unbalanced
 1084 Generalized Schrödinger Bridge (GSB) problem considered in this work. Specifically, DeepRUOT
 1085 (Zhang et al., 2025c) solves the Regularized Unbalanced Optimal Transport (RUOT) problem by
 1086 parameterizing the canonical stochastic bridge SDE $dX_t = f_t(X_t)dt + \sigma_t dB_t$ with a probability
 1087 flow ODE given by

$$1088 \quad dX_t = \underbrace{\left\{ f_t(X_t) + \frac{1}{2} \sigma_t^2 \nabla_x \log p_t^\theta(X_t) \right\} dt}_{u_t^\theta(X_t)} \quad (42)$$

$$1089$$

$$1090$$

1091 where $u_t^\theta(X_t)$ is the drift of the probability flow ODE that is learned along with the probability density
 1092 $p_t^\theta(X_t)$ to derive the drift of the SDE $f_t(X_t)$. Unlike GSBM (Liu et al., 2023a), which enforces
 1093 a hard constraint on the endpoints, DeepRUOT models a probability density flow by minimizing
 1094 a reconstruction loss that encourages alignment with both intermediate and terminal distributions.
 1095 While effective when intermediate distributions are observed, the method fails to learn meaningful
 1096 trajectories in settings where these intermediate snapshots are unavailable, limiting its applicability in
 1097 many real-world scenarios.

1098 **Learning Diverging Trajectories with Single Target SBM** To model diverging trajectories with
 1099 single-branch SBM where the target distribution is multi-modal, we can follow a set of N_{samples}
 1100 samples from the initial distribution π_0 to determine the distribution of samples that end at each of
 1101 the modes of the target distribution π_1 . While this can estimate the splitting of mass across different
 1102 trajectories, it does not explicitly learn the optimal distribution of mass in the latent space over time.
 1103

1104 Furthermore, if the path toward a specific cluster in the target distribution has lower potential than
 1105 that of the other clusters, mode collapse could occur, where all samples follow the same trajectory
 1106 without reaching the other clusters. BranchSBM learns to generate the correct mass distribution
 1107 over each of the target states by optimizing the growth term with respect to the matching loss $\mathcal{L}_{\text{match}}$
 1108 defined in Equation 15.

1109 With the standard SBM formulation, it is also challenging to determine the time at which branching
 1110 occurs and the population diverges toward different targets, as all samples follow stochastic trajectories.
 1111 With BranchSBM, we model population-wide branching dynamics with growth networks that
 1112 are trained to predict the origin of a branch from having zero mass ($w_{t,k} = 0$) and the rate at which it
 1113 grows/shrinks over time $\partial_t w_{t,k} = g_{t,k}(X_{t,k})$. The mass of a branch at any given time step can be
 1114 simulated with $w_{t,k}(X_{t,k}) = w_{0,k} + \int_0^t g_{s,k}^\phi(X_{s,k})ds$.

1115 **Branching Dynamics** While branched dynamics have been explored in the context of optimal
 1116 transport (Lippmann et al., 2022) and Brownian motion (Baradat & Lavenant, 2021), no previous
 1117 work has explicitly formulated or solved the branched Schrödinger bridge problem that seeks to
 1118 match an initial distribution to a multiple terminal distributions via stochastic bridges. For instance,
 1119 the concept of Branching Brownian Motion (BBM) introduced in Baradat & Lavenant (2021) models
 1120 a population of particles that each independently follow stochastic trajectories according to standard
 1121 Brownian motion with positive diffusivity $\nu > 0$. To model branching dynamics, each particle has
 1122 an independent branching rate $\lambda > 0$ that determines the probability that the particle undergoes
 1123 a *branching event*, defined as the particle dying and generating k new particles that then evolve
 1124 independently. The number of generated particles is sampled from a probability distribution over
 1125 non-negative integers $k \sim p = (p_k)_{k \in \mathbb{N}}$, where p_k is the probability of generating k particles at the
 1126 branching event and $\sum_{k \in \mathbb{N}} p_k = 1$. Given p_k , $q_k = \lambda p_k$ defines the rate of branching events that
 1127 generate k new particles and define a new probability measure $q = (q_k)_{k \in \mathbb{N}}$ called the *branching
 1128 mechanism*.

1129 While BBM defines branching as the generation of additional particles from a single particle following
 1130 an independent, temporal probability measure q , this model fails to model the *division* of mass across
 1131 multiple trajectories, where total mass remains constant but the mass of each branch changes. In
 1132 addition, BBM assumes that each branched particle undergoes independent Brownian motion, without
 1133 explicitly defining a terminal state or branch-specific drift. For these reasons, the BBM model is
 1134 unsuitable for modeling branching in the majority of real-world contexts, such as cell state transitions,

1134 where undifferentiated cells split probabilistically into distinct fates rather than proliferating in
 1135 number. In such systems, branching reflects a redistribution of probability mass over developmental
 1136 trajectories, governed by underlying regulatory patterns rather than purely stochastic reproduction.
 1137

1138 Where existing frameworks fall short is in modeling **meaningful energy-aware, conditional stochastic**
 1139 **trajectories with unbalanced and multi-modal dynamics**, which we address in this work.
 1140 Specifically, we formulate the Unbalanced CondSOC problem followed by the Branched CondSOC
 1141 problem that defines a *set* of optimal drifts and growth fields that define a set of branched trajectories
 1142 following optimal energy-minimizing trajectories defined by the state cost $V_t(X_t)$. Instead of spline
 1143 optimization, we leverage a parameterized network $\varphi_{t,\eta}^*(\mathbf{x}_0, \mathbf{x}_1)$ that learns to predict an optimal
 1144 interpolating path given a pair of endpoints.

1145 To model dynamic growth of mass along branched trajectories, we initialize a normalized population
 1146 weight of 1 at a *primary branch* ($k = 0$), that can split across K branches and generate weights
 1147 $\{w_{t,k}\}_{k=0}^K$ that evolve co-currently via learned growth rates $\{g_{t,k}^\phi\}_{k=0}^K$. This approach relies on
 1148 learning the growth rate over an entire population of samples across all branches rather than learning
 1149 independent growth rates of individual samples, which enforces stronger constraints during training
 1150 to ensure that the model captures true population growth dynamics.

1151 B.2 MODELING PERTURBATION RESPONSES

1152 In single-cell transcriptomics, perturbations such as gene knockouts, transcription factor induction,
 1153 or drug treatments frequently induce cell state transitions that diverge into distinct fates, reflecting
 1154 differentiation, resistance, and off-target effects (Shalem et al., 2014; Kramme et al., 2021; Dixit
 1155 et al., 2016; Gavriilidis et al., 2024; Zhang et al., 2025a; Kobayashi et al., 2022; Smela et al., 2023;
 1156 Pierson Smela et al., 2025; Yeo et al., 2021). Trajectory inference methods for modeling cell-state
 1157 dynamics with single-cell RNA sequencing data, including flow matching (Zhang et al., 2025b;
 1158 Rohbeck et al., 2025; Atanackovic et al., 2024; Tong et al., 2024b; Wang et al., 2025), Schrödinger
 1159 Bridge Matching (Alatkar & Wang, 2025; Tong et al., 2024b), and optimal transport (Zhang et al.,
 1160 2025c; Tong et al., 2020; Bunne et al., 2023; Driessens et al., 2025; Huguet et al., 2022; Klein et al.,
 1161 2024; Yachimura et al., 2024; Schiebinger et al., 2019; Zhang et al., 2021; Bunne et al., 2022b),
 1162 have been widely explored. While many algorithms have been developed to predict perturbation
 1163 responses (Bunne et al., 2023; Megas et al., 2025; Rohbeck et al., 2025; Roohani et al., 2023; Ryu
 1164 et al., 2024), they either predict only the terminal perturbed state without modeling the intermediate
 1165 cell-state transitions or model single unimodal perturbation trajectories. Schrödinger Bridge Matching
 1166 frameworks like $[\text{SF}]^2\text{M}$ (Tong et al., 2024b) and DeepGSB (Liu et al., 2022) have been shown to
 1167 effectively model stochastic transitions in biological systems, offering better scalability and sampling
 1168 efficiency than classical OT.

1169 However, these models are still limited to modeling trajectories between a single pair of boundary
 1170 distributions, limiting their ability to represent divergent trajectories arising from the same perturbed
 1171 initial state. This is a key limitation when modeling processes like fate bifurcation post-perturbation,
 1172 where a cell population exposed to the same stimulus may split into multiple phenotypically distinct
 1173 outcomes. BranchSBM extends the SBM framework to support multiple terminal marginals, enabling
 1174 modeling of stochastic bifurcations in a mathematically principled way. By learning a mixture of
 1175 conditional stochastic processes from a common source to multiple target distributions, BranchSBM
 1176 can capture the heterogeneity and uncertainty of cell fate decisions under perturbation. Moreover,
 1177 it retains the empirical tractability of previous SB-based models, requiring only samples from
 1178 distributions, and ensures that intermediate trajectories lie on the manifold of feasible cell states.

1179 B.3 MODELING PROBABILISTIC TRAJECTORIES WITH BRANCHSBM

1180 We conclude the discussion with an alternative interpretation of the Branched Schrödinger Bridge
 1181 problem that deviates from the branching population dynamics problem. We instead consider
 1182 Branched SB as a probabilistic trajectory matching problem, where each branch is one of multiple
 1183 possible trajectories that a sample X_t could follow. Since single-path SBM learns only a single
 1184 deterministic drift field $u_t(X_t)$ that determines the direction and flow of the SDE $dX_t = u_t^\theta(X_t)dt +$
 1185 σdB_t , the probabilistic aspect of the trajectory remains restricted to Brownian motion via σdB_t . This
 1186 fails to capture probabilistic dynamics where probability densities begin concentrated at a single state
 1187 but evolve into multi-modal probability densities $\{p_{t,k}(X_t)\}_{k=0}^K$ over multiple different states, each

of which evolve according to an SDE $dX_t = u_{t,k}(X_t)dt + \sigma_k dB_t$ with an independent drift term $u_{t,k}(X_t)$ and noise scaling σ_k .

Where other SBM frameworks fall short, BranchSBM is capable of modeling multiple probabilistic trajectories, where a system begins at a single deterministic state $X_0 = \mathbf{x}_0$ with probability $w_{1,0} = 1$ and evolves via multiple probabilistic trajectories that diverge in the state space governed by the SDEs $dX_t = u_{t,k}(X_t)dt + \sigma_k dB_t$. At time t , the system exists in a non-deterministic superposition of states $\{X_{t,k}, w_{t,k}\}_{k=0}^K$, each with a probability $w_{t,k}$ such that $\sum_{k=0}^K w_{t,k}^\phi = 1$. In addition, BranchSBM can model the evolution of the probability weights $w_{t,k}^\phi$ by parameterizing the *probabilistic growth rates* $\{g_{t,k}^\phi\}_{k=0}^K$ that preserve conservation of probability mass $\sum_{k=0}^K w_{t,k}^\phi = 1$ at all times $t \in [0, 1]$ by minimizing the mass loss $\mathcal{L}_{\text{mass}}(\phi)$ (16). This problem is prevalent in many biological and physical systems, where a system does not exist in a single deterministic state, but rather a *superposition* over a distribution of states.

Table 6: **BranchSBM enables robust modeling of branched Schrödinger bridges compared to existing frameworks.** BranchSBM can model both branching and unbalanced trajectories, follows intermediate trajectories governed by a task-specific state cost, requires only endpoint samples for training, and samples trajectories from only a single sample from the starting distribution.

Method	Models Branching	Models Unbalanced	Intermediate Dynamics	Requirements for Training	Requirements for Sampling
Generalized SBM Liu et al. (2023a)	No	No	Entropic OT with learned drift u_t^θ	Samples from π_0, π_1	Endpoints $(\mathbf{x}_0, \mathbf{x}_1) \sim p_{0,1}^\theta$
DeepRUOT (Zhang et al., 2025c)	Requires simulating multiple samples	Growth rate $g_t^\phi(X_t)$	Regularized OT with learned drift u_t^θ and density p_t^θ	Samples from π_0, π_1	Sample $\mathbf{x}_0 \sim \pi_0$
BranchSBM (Ours)	Simulates divergent trajectories and terminal states from single sample \mathbf{x}_0	Branched growth rates $\{g_{t,k}^\phi(X_t)\}_{k=0}^K$	Branched drifts $\{u_{t,k}^\theta\}_{k=0}^K$ that minimize state-cost V_t	Samples from π_0 and clusters $\{\pi_{1,k}\}_{k=0}^K$	Sample $\mathbf{x}_0 \sim \pi_0$

C THEORETICAL PROOFS

C.1 PROOF OF PROPOSITION 1

Proposition 1 (Unbalanced Conditional Stochastic Optimal Control) *Suppose the marginal density can be decomposed as $p_t(X_t) = \int_{\pi_{0,1}} p_t(X_t|\mathbf{x}_0, \mathbf{x}_1) p_{0,1}(\mathbf{x}_0, \mathbf{x}_1) d\pi_{0,1}$, where $\pi_{0,1}$ is a fixed joint coupling of the data. Then, we can identify the optimal drift u_t^* and growth g_t^* that solves the Unbalanced GSB problem in (4) by minimizing the Unbalanced Conditional Stochastic Optimal Control objective given by*

$$\min_{u_t, g_t} \mathbb{E}_{(\mathbf{x}_0, \mathbf{x}_1) \sim \pi_{0,1}} \left[\int_0^1 \mathbb{E}_{p_{t|0,1}} \left[\frac{1}{2} \|u_t(X_t|\mathbf{x}_0, \mathbf{x}_1)\|^2 + V_t(X_t) \right] w_t(X_t) dt \right] \quad (43)$$

$$\text{s.t. } dX_t = u_t(X_t|\mathbf{x}_0, \mathbf{x}_1)dt + \sigma dB_t, \quad X_0 = \mathbf{x}_0, \quad X_1 = \mathbf{x}_1 \quad w_0(X_0) = w_0^*, \quad w_1(X_1) = w_1^* \quad (44)$$

where $w_t = w_0 + \int_0^t g_s(X_s)ds$ is the time-dependent weight initialized at w_0^* , u_t is the drift, g_t is the growth rate, and $\pi_{0,1}$ is the weighted coupling of paired endpoints $(\mathbf{x}_0, w_0, \mathbf{x}_1, w_1) \sim \pi_{0,1}$.

Proof. We define the Unbalanced Generalized Schrödinger Bridge problem as the solution (u_t, p_t, g_t) to the energy minimization problem such that the unbalanced Fokker-Planck equation is satisfied.

$$\min_{u_t, g_t} \int_0^1 \mathbb{E}_{p_t} \left[\frac{1}{2} \|u_t(X_t)\|^2 + V_t(X_t) \right] \left(w_0 + \int_0^t g_s(X_s)ds \right) dt \quad (45)$$

$$\text{s.t. } \begin{cases} \frac{\partial}{\partial t} p_t(X_t) = -\nabla \cdot (u_t(X_t)p_t(X_t)) + \frac{\sigma^2}{2} \Delta p_t(X_t) + g_t(X_t)p_t(X_t) \\ p_0 = \pi_0, \quad p_1 = \pi_1 \end{cases} \quad (46)$$

Under the assumption that the joint probability $p_{0,1}(\mathbf{x}_0, \mathbf{x}_1)$ is fixed over all times $t \in [0, 1]$ and the marginal probability can be factorized as $p_t(X_t) = \mathbb{E}_{p_{0,1}}[p_t(X_t | \mathbf{x}_0, \mathbf{x}_1)]$, we can decompose the minimization objective into

$$\begin{aligned} & \int_0^1 \mathbb{E}_{p_t} \left[\frac{1}{2} \|u_t(X_t)\|^2 + V_t(X_t) \right] \left(w_0 + \int_0^t g_s(X_s) ds \right) dt \\ &= \int_0^1 \mathbb{E}_{p_{0,1}} \mathbb{E}_{p_{t|0,1}} \left[\frac{1}{2} \|u_t(X_t)\|^2 + V_t(X_t) \right] \left(w_0 + \int_0^t g_s(X_s) ds \right) dt \quad (\text{law of total expectation}) \\ &= \mathbb{E}_{p_{0,1}} \int_0^1 \mathbb{E}_{p_{t|0,1}} \left[\frac{1}{2} \|u_t(X_t)\|^2 + V_t(X_t) \right] \left(w_0 + \int_0^t g_s(X_s) ds \right) dt \quad (\text{Fubini's theorem}) \end{aligned} \quad (47)$$

which can be solved by identifying the conditional drift u_t that minimizes the expected objective value over all endpoint samples $(\mathbf{x}_0, \mathbf{x}_1) \sim p_{0,1}$. Under similar assumptions, we can decompose all terms in the Fokker-Planck equation. For the left-hand side, we have:

$$\frac{\partial}{\partial t} p_t(X_t) = \frac{\partial}{\partial t} \int_{\pi_{0,1}} p_t(X_t | \mathbf{x}_0, \mathbf{x}_1) p_{0,1}(\mathbf{x}_0, \mathbf{x}_1) d\pi_{0,1} = \mathbb{E}_{(\mathbf{x}_0, \mathbf{x}_1) \sim \pi_{0,1}} \left[\frac{\partial}{\partial t} p_t(X_t | \mathbf{x}_0, \mathbf{x}_1) \right] \quad (48)$$

For the divergence term, we have:

$$\begin{aligned} \nabla \cdot (u_t(X_t) p_t(X_t)) &= \nabla \cdot \left(u_t(X_t) \int_{\pi_{0,1}} p_t(X_t | \mathbf{x}_0, \mathbf{x}_1) p_{0,1}(\mathbf{x}_0, \mathbf{x}_1) d\pi_{0,1} \right) \\ &= \nabla \cdot \left(\int_{\pi_{0,1}} u_t(X_t) p_t(X_t | \mathbf{x}_0, \mathbf{x}_1) p_{0,1}(\mathbf{x}_0, \mathbf{x}_1) d\pi_{0,1} \right) \end{aligned} \quad (49)$$

By Fubini's Theorem and the linearity of the divergence operator, we can switch the order of integration to get

$$\begin{aligned} \nabla \cdot (u_t(X_t) p_t(X_t)) &= \int_{\pi_{0,1}} \nabla \cdot \left(u_t(X_t) p_t(X_t | \mathbf{x}_0, \mathbf{x}_1) \right) p_{0,1}(\mathbf{x}_0, \mathbf{x}_1) d\pi_{0,1} \\ &= \mathbb{E}_{(\mathbf{x}_0, \mathbf{x}_1) \sim \pi_{0,1}} \left[\nabla \cdot (u_t(X_t) p_t(X_t | \mathbf{x}_0, \mathbf{x}_1)) \right] \end{aligned} \quad (50)$$

For the Laplacian term, we have:

$$\begin{aligned} \frac{\sigma^2}{2} \Delta p_t(X_t) &= \frac{\sigma^2}{2} \nabla \cdot (\nabla p_t(X_t)) \\ &= \frac{\sigma^2}{2} \nabla \cdot \left(\nabla \int_{\pi_{0,1}} p_t(X_t | \mathbf{x}_0, \mathbf{x}_1) p_{0,1}(\mathbf{x}_0, \mathbf{x}_1) d\pi_{0,1} \right) \\ &= \frac{\sigma^2}{2} \int_{\pi_{0,1}} \left(\nabla \cdot \nabla p_t(X_t | \mathbf{x}_0, \mathbf{x}_1) \right) p_{0,1}(\mathbf{x}_0, \mathbf{x}_1) d\pi_{0,1} \\ &= \frac{\sigma^2}{2} \mathbb{E}_{(\mathbf{x}_0, \mathbf{x}_1) \sim \pi_{0,1}} [\Delta p_{t|0,1}(X_t | \mathbf{x}_0, \mathbf{x}_1)] \end{aligned} \quad (51)$$

For the growth term, we have:

$$\begin{aligned} g_t(X_t) p_t(X_t) &= g_t(X_t) \int_{\pi_{0,1}} p_t(X_t | \mathbf{x}_0, \mathbf{x}_1) p_{0,1}(\mathbf{x}_0, \mathbf{x}_1) d\pi_{0,1} \\ &= \int_{\pi_{0,1}} g_t(X_t) p_t(X_t | \mathbf{x}_0, \mathbf{x}_1) p_{0,1}(\mathbf{x}_0, \mathbf{x}_1) d\pi_{0,1} \\ &= \mathbb{E}_{(\mathbf{x}_0, \mathbf{x}_1) \sim \pi_{0,1}} [g_t(X_t) p_t(X_t | \mathbf{x}_0, \mathbf{x}_1)] \end{aligned} \quad (52)$$

Combining all the terms of the Fokker-Planck equation, we have shown that (46) can be rewritten as

$$\frac{\partial}{\partial t} p_t(X_t | \mathbf{x}_0, \mathbf{x}_1) = -\nabla \cdot (u_t p_t(X_t | \mathbf{x}_0, \mathbf{x}_1)) + \Delta p_{t|0,1}(X_t | \mathbf{x}_0, \mathbf{x}_1) + g_t p_t(X_t | \mathbf{x}_0, \mathbf{x}_1) \quad (53)$$

Therefore, the Unbalanced GSB problem is equivalent to solving the Unbalanced CondSOC problem, and we conclude our proof of Proposition 1.

1296 C.2 PROOF OF PROPOSITION 2
1297

1298 **Proposition 2 (Branched Conditional Stochastic Optimal Control)** *For each branch, let*
1299 *$p_{t,k}(X_{t,k}) = \mathbb{E}_{p_{0,1k}}[p_{t,k}(X_{t,k} | \mathbf{x}_0, \mathbf{x}_{1,k})]$, where $\pi_{0,1,k}$ is the joint coupling distribution of sam-*
1300 *ples $\mathbf{x}_0 \sim \pi_0$ from the initial distribution and $\mathbf{x}_{1,k} \sim \pi_{1,k}$ from the k th target distribution. Then,*
1301 *we can identify the set of optimal drift and growth functions $\{u_{t,k}^*, g_{t,k}^*\}_{k=0}^K$ that solve the Branched*
1302 *GSB problem in (3.2) by minimizing the sum of Unbalanced CondSOC objectives given by*

$$\begin{aligned} \min_{\{u_{t,k}, g_{t,k}\}_{k=0}^K} & \mathbb{E}_{(\mathbf{x}_0, \mathbf{x}_{1,0}) \sim \pi_{0,1,0}} \int_0^1 \left\{ \mathbb{E}_{p_{t|0,1,0}} \left[\frac{1}{2} \|u_{t,0}(X_{t,0})\|^2 + V_t(X_{t,0}) \right] w_{t,0} \right. \\ & \left. + \sum_{k=1}^K \mathbb{E}_{(\mathbf{x}_0, \mathbf{x}_{1,k}) \sim \pi_{0,1,k}} \int_0^1 \mathbb{E}_{p_{t|0,1,k}} \left[\frac{1}{2} \|u_{t,k}(X_{t,k})\|^2 + V_t(X_{t,k}) \right] w_{t,k} \right\} dt \end{aligned} \quad (54)$$

$$\text{s.t. } dX_{t,k} = u_{t,k}(X_{t,k})dt + \sigma dB_t, \quad X_0 = \mathbf{x}_0, \quad X_{1,k} = \mathbf{x}_{1,k}, \quad w_{0,k} = \delta_{k=0}, \quad w_{1,k} = w_{1,k} \quad (55)$$

1308 where $w_{t,0} = 1 + \int_0^t g_{s,0}(\mathbf{x}_{s,0})ds$ is the weight of the primary paths initialized at 1 and $w_{t,k} =$
1309 $\int_0^t g_{s,k}(\mathbf{x}_{s,k})ds$ are the weights of the K secondary branches initialized at 0.

1312 *Proof.* We extend the proof of Proposition 1 to the branching case by defining each branch k as an
1313 independent Unbalanced Generalized Schrödinger Bridge problem in (4) given by

$$\min_{u_{t,k}, g_{t,k}} \int_0^1 \mathbb{E}_{p_{t,k}} \left[\frac{1}{2} \|u_{t,k}(X_{t,k})\|^2 + V_t(X_{t,k}) \right] \left(w_{0,k} + \int_0^t g_{s,k}(X_{s,k})ds \right) dt \quad (56)$$

$$\text{s.t. } \begin{cases} \frac{\partial}{\partial t} p_{t,k}(X_{t,k}) = -\nabla \cdot (u_{t,k}(X_t) p_{t,k}(X_{t,k})) + \frac{\sigma^2}{2} \Delta p_{t,k}(X_{t,k}) + g_{t,k}(X_{t,k}) p_{t,k}(X_{t,k}) \\ p_0 = \pi_0, \quad p_{t,k} = \pi_{t,k} \end{cases} \quad (57)$$

1319 such that each branch independently solves the Fokker-Planck equation defined as $\frac{\partial}{\partial t} p_{t,k} = -\nabla \cdot$
1320 $(u_{t,k} p_{t,k}) + \frac{\sigma^2}{2} \Delta p_{t,k}$. Now, we show that the sum of unbalanced CondSOC problems still satisfies
1321 the global Fokker-Planck equation

$$\frac{\partial}{\partial t} p_t = -\nabla \cdot (u_t p_t) + \frac{\sigma^2}{2} \Delta p_t \quad (58)$$

1325 where we define p_t as the weighted sum of the branched distributions given by

$$p_t(X_t) = \sum_{k=0}^K w_{t,k} p_{t,k}(X_t) \quad (59)$$

1329 To obtain an expression for the global Fokker-Planck equation, we differentiate both sides and
1330 substitute the branched FPE as follows

$$\begin{aligned} \frac{\partial}{\partial t} p_t &= \frac{\partial}{\partial t} \left[\sum_{k=0}^K w_{t,k} p_{t,k} \right] \\ \frac{\partial}{\partial t} p_t &= \sum_{k=0}^K \left[w_{t,k} \left(\frac{\partial}{\partial t} p_{t,k} \right) + \left(\frac{\partial}{\partial t} w_{t,k} \right) p_{t,k} \right] \\ \frac{\partial}{\partial t} p_t &= \sum_{k=0}^K \left[w_{t,k} \left(-\nabla \cdot (u_{t,k} p_{t,k}) + \frac{\sigma^2}{2} \Delta p_{t,k} \right) + g_{t,k} p_{t,k} \right] \quad (\text{substitute branched FPE and } \frac{\partial}{\partial t} w_{t,k} = g_{t,k}) \\ \frac{\partial}{\partial t} p_t &= \sum_{k=0}^K \left[-w_{t,k} \nabla \cdot (u_{t,k} p_{t,k}) + w_{t,k} \frac{\sigma^2}{2} \Delta p_{t,k} + g_{t,k} p_{t,k} \right] \\ \frac{\partial}{\partial t} p_t &= \sum_{k=0}^K (-w_{t,k} \nabla \cdot (u_{t,k} p_{t,k})) + \sum_{k=0}^K w_{t,k} \frac{\sigma^2}{2} \Delta p_{t,k} + \sum_{k=0}^K g_{t,k} p_{t,k} \end{aligned} \quad (60)$$

1345 By linearity of the Laplacian, the diffusion term can be rewritten as

$$\sum_{k=0}^K w_{t,k} \frac{\sigma^2}{2} \Delta p_{t,k} = \frac{\sigma^2}{2} \Delta \underbrace{\left(\sum_{k=0}^K w_{t,k} p_{t,k} \right)}_{p_t} = \frac{\sigma^2}{2} \Delta p_t \quad (61)$$

1350 Since the global growth term $\sum_{k=0}^K g_{t,k} p_{t,k}$ is the sum of the growth dynamics across all branches
 1351 and is separate from the drift and diffusion dynamics, it doesn't alter the direction or motion of
 1352 particles along the branched fields. Thus, both the diffusion and growth terms satisfy the global FPE.
 1353

1354 Now, we set the divergence term in (60) equal to the global divergence $\nabla \cdot (u_t p_t)$ and derive
 1355 the expression for the total drift field $u_t(X_t)$ that satisfies the global FPE. By the linearity of the
 1356 divergence operator, we get

$$\begin{aligned} 1357 \quad -\nabla \cdot (u_t p_t) &= \sum_{k=0}^K (-w_{t,k} \nabla \cdot (u_{t,k} p_{t,k})) \\ 1358 \\ 1359 \quad -\nabla \cdot (u_t p_t) &= -\nabla \cdot \left(\sum_{k=0}^K w_{t,k} u_{t,k} p_{t,k} \right) \\ 1360 \\ 1361 \quad -\nabla \cdot (u_t p_t) &= -\nabla \cdot \underbrace{\left(\frac{1}{p_t} \sum_{k=0}^K w_{t,k} u_{t,k} p_{t,k} \right)}_{u_t} p_t \\ 1362 \\ 1363 \\ 1364 \\ 1365 \\ 1366 \end{aligned} \quad (62)$$

1367 Under the global FPE constraint, the drift u_t is defined as the mass-weighted average of the drift
 1368 fields for each branch, given by $u_t(X_t) = \frac{1}{p_t(X_t)} \sum_{k=0}^K w_{t,k}(X_t) u_{t,k}(X_t) p_{t,k}(X_t)$. Intuitively, this
 1369 means that in the global context, for any $X_t = \mathbf{x}_t$, the drift of state X_t along the dynamics of branch
 1370 k is scaled by the weight of \mathbf{x}_t at time t along branch k and the ratio of probability density of \mathbf{x}_t
 1371 under branch k over the total probability density of \mathbf{x}_t across all branches. Therefore, our definition
 1372 of the weighted drift decoupled over individual branches satisfies the global FPE equation in (58),
 1373 and this concludes the proof of Proposition 2.

1374 **Remark 1 (Reduction to Single Path GSBM)** When $g_{t,0}(X_{t,0}) = 0$ and $g_{t,k}(X_{t,k}) = 0$ for all
 1375 $t \in [0, 1]$ and $k \in \{1, \dots, K\}$, then the Branched CondSOC problem is the solution to the single
 1376 path GSB problem given by

$$1377 \quad \min_{u_t} \int \mathbb{E}_{p_{t|0,1}} \left[\frac{1}{2} \|u_t(X_t)\|^2 + V_t(X_t) \right] dt \quad \text{s.t.} \quad \begin{cases} dX_t = u_t(X_t) dt + \sigma dB_t \\ X_0 \sim \pi_0, \quad X_1 \sim \pi_1 \end{cases} \quad (63)$$

1381 where the probability density $p_{t|0,1}(X_t | \mathbf{x}_0, \mathbf{x}_1)$ is conditioned explicitly on a pair of endpoints
 1382 $(\mathbf{x}_0, \mathbf{x}_1) \sim \pi_{0,1}$ drawn from the joint distribution.

1384 C.3 PROOF OF PROPOSITION 3

1386 **Proposition 3 (Solving the GSB Problem with Stage 1 and 2 Training)** *Stage 1 and Stage 2
 1387 training yield the optimal drift $u_t^*(X_t)$ that generates the optimal marginal probability distribution
 1388 $p_t^*(X_t)$ that solves the GSB problem in (3).*

1389 *Proof.* Let the marginal probability distribution $p_t^*(X_t)$ and corresponding drift $u_t^*(X_t)$ define the
 1390 optimal solution to the GSB problem. It suffices to show that

- 1392 Given $(\mathbf{x}_0, \mathbf{x}_1) \sim \pi_{0,1}$, Stage 1 training with the trajectory loss $\mathcal{L}_{\text{traj}}(\eta)$ (12) yields the inter-
 1393 polant $\mathbf{x}_{t,\eta}^*$ and time-derivative $\dot{\mathbf{x}}_{t,\eta}^*$ that define $p_t^*(X_t)$.
- 1395 Given $p_t^*(X_t)$, Stage 2 training with the explicit flow matching loss $\mathcal{L}_{\text{flow}}(\theta)$ (13) yields the
 1396 optimal drift $u_t^*(X_t)$.

1397 To prove Part 1, we establish the following Lemma.

1398 **Lemma 1.** *Given the Markovian reference process \mathbb{Q} with drift $v_t^*(X_t)$ that is the minimizer of the
 1399 unconstrained GSB problem in (3), Stage 1 training returns the velocity field $\dot{\mathbf{x}}_{t,\eta}^*$ that generates
 1400 reciprocal projection $\Pi^* = \text{proj}_{\mathcal{R}(\mathbb{Q})}(\mathbb{P})$ of the path measure $\mathbb{P} = p_{0,1}$.*

1401 *Proof of Lemma.* It suffices to prove that $\dot{\mathbf{x}}_{t,\eta}^*$ minimizes the KL-divergence with $v_t^*(X_t)$ under
 1402 endpoint constraints $X_0 = \mathbf{x}_0$ and $X_1 = \mathbf{x}_1$.

1404 Consider the unconstrained Markovian drift v_t^* that is the minimizer of the energy function given by
 1405

$$1406 \quad v_t^* = \arg \min_{v_t} \int_0^1 \mathbb{E}_{p_t} \left[\frac{1}{2} \|v_t(X_t)\|^2 + V_t(X_t) \right] dt \quad (64)$$

1407 The class of interpolants is given by parameters η is defined as
 1408

$$1409 \quad \mathbf{x}_{t,\eta} = (1-t)\mathbf{x}_0 + t\mathbf{x}_1 + t(1-t)\varphi_{t,\eta}(\mathbf{x}_0, \mathbf{x}_1) \quad (65)$$

$$1410 \quad \dot{\mathbf{x}}_{t,\eta} = \mathbf{x}_1 - \mathbf{x}_0 + t(1-t)\dot{\varphi}_{t,\eta}(\mathbf{x}_0, \mathbf{x}_1) + (1-2t)\varphi_{t,\eta}(\mathbf{x}_0, \mathbf{x}_1) \quad (66)$$

1411 Stage 1 training yields the optimal interpolant $\dot{\mathbf{x}}_{t,\eta}^*$ that minimizes the energy function across all time
 1412 points $t \in [0, 1]$, defined as
 1413

$$1414 \quad \dot{\mathbf{x}}_{t,\eta}^* = \arg \min_{\dot{\mathbf{x}}_{t,\eta}} \int_0^1 \mathbb{E}_{p_t} \left[\frac{1}{2} \|\dot{\mathbf{x}}_{t,\eta}^*(\mathbf{x}_0, \mathbf{x}_1)\|^2 + V_t(\mathbf{x}_{t,\eta}^*) \right] dt \quad (67)$$

1415 Therefore, we aim to prove that $\dot{\mathbf{x}}_{t,\eta}^*$ is the velocity field corresponding to $\Pi^* = p_{0,1}\mathbb{Q}_{\cdot|0,1}$, which is
 1416 the reciprocal projection of $p_{0,1}$ onto the class of path measures that share the same bridge marginals
 1417 conditioned on a pair of endpoints, called the reciprocal class $\mathcal{R}(\mathbb{Q})$ (Léonard, 2014).

1418 By definition, the reciprocal projection Π^* is defined as the minimizer of the KL-divergence with
 1419 $\mathbb{P} = p_{0,1}$ that lies within the reciprocal class $\mathcal{R}(\mathbb{Q})$ of measures that match the bridge conditions of
 1420 \mathbb{Q} .
 1421

$$1422 \quad \Pi^* = \arg \min_{\Pi \in \mathcal{R}(\mathbb{Q})} \text{KL}(\mathbb{P} \parallel \Pi) \quad (68)$$

1423 The reciprocal projection matches the endpoint constraints of \mathbb{P} while following the bridge conditions
 1424 $\mathbb{Q}_{\cdot|0,1} = \mathbb{Q}(\cdot|X_0 = \mathbf{x}_0, X_1 = \mathbf{x}_1)$ (Léonard, 2014). Therefore, we can write the generating
 1425 velocity field $\dot{\mathbf{x}}_{t,\eta}^*$ as
 1426

$$1427 \quad \dot{\mathbf{x}}_{t,\eta}^*(\mathbf{x}_0, \mathbf{x}_1) = \mathbb{E}_{\mathbf{x}_t \sim \mathbb{Q}} [\dot{\mathbf{x}}_t | X_0 = \mathbf{x}_0, X_1 = \mathbf{x}_1] \quad (69)$$

1428 Since both v_t^* generating \mathbb{Q} and $\dot{\mathbf{x}}_{t,\eta}^*$ minimize the same energy function, the objective in (68) reduces
 1429 to determining $\dot{\mathbf{x}}_{t,\eta}^*$ that is the minimizer of the dynamic formulation of KL divergence between the
 1430 reference process \mathbb{Q} and all path measures $\mathbb{P} = p_{0,1}$ that preserve endpoint constraints given by
 1431

$$1432 \quad \dot{\mathbf{x}}_{t,\eta}^* = \arg \min_{\dot{\mathbf{x}}_{t,\eta}} \text{KL}(\mathbb{P} \parallel \mathbb{Q}) \quad (70)$$

$$1433 \quad = \arg \min_{\dot{\mathbf{x}}_{t,\eta}} \mathbb{E}_{(\mathbf{x}_0, \mathbf{x}_1) \sim p_{0,1}} \left[\int_0^1 \frac{1}{2} \|\dot{\mathbf{x}}_{t,\eta}^* - v_t^*(\mathbf{x}_{t,\eta})\|^2 dt \right] \quad (71)$$

1434 Therefore, we conclude that $\dot{\mathbf{x}}_{t,\eta}^*$ generates the reciprocal projection $\Pi^* = p_{0,1}\mathbb{Q}_{\cdot|0,1}$.
 1435

1436 By Lemma 1, we know that v_t^* is the optimal drift energy function in the GSB problem without
 1437 endpoint constraints and $\dot{\mathbf{x}}_{t,\eta}^*$ is the reciprocal projection that follows the dynamics of v_t^* while
 1438 preserving the coupling $p_{0,1}$. Therefore, we can define $u_{t|0,1}^*(X_t | \mathbf{x}_0, \mathbf{x}_1) \equiv \dot{\mathbf{x}}_{t,\eta}^*(\mathbf{x}_0, \mathbf{x}_1)$ as the
 1439 conditional drift that generates the conditional probability distribution $p_{t|0,1}^*(X_t | \mathbf{x}_0, \mathbf{x}_1)$ that satisfies
 1440 the Fokker-Planck equation
 1441

$$1442 \quad \frac{\partial}{\partial t} p_{t|0,1} = -\nabla \cdot (u_{t|0,1} p_{t|0,1}) + \frac{1}{2} \sigma^2 \Delta p_{t|0,1} \quad (72)$$

1443 Given that $p_{t|0,1}^*(X_t | \mathbf{x}_0, \mathbf{x}_1)$ is the optimal bridge that solves the GSB problem for a pair of endpoints
 1444 $(\mathbf{x}_0, \mathbf{x}_1) \sim p_{0,1}$, we can define the marginal probability distribution p_t^* as the mixture of bridges
 1445

$$1446 \quad p_t^* = \mathbb{E}_{(\mathbf{x}_0, \mathbf{x}_1) \sim \pi_{0,1}} [p_{t|0,1}^*(X_t | \mathbf{x}_0, \mathbf{x}_1)] \quad (73)$$

1447 which concludes the proof of Part 1.
 1448

1449 For Part 2 of the proof, we aim to show that Stage 2 training yields the optimal Markovian drift
 1450 $u_t^*(X_t)$ that generates $p_t^*(X_t)$. To do this, we write the Fokker-Planck equation for $p_t^*(X_t)$ in terms
 1451

of the conditional bridge $p_{t|0,1}$ and drift field $u_{t|0,1}$ to extract an expression for $u_t^*(X_t)$ that satisfies it. Starting from the definition of p_t^* , we have

$$\begin{aligned}
 p_t^* &= \int p_{t|0,1}^* p_{0,1} d\mathbf{x}_{0,1} \\
 \frac{\partial}{\partial t} p_t^* &= \int \left(\frac{\partial}{\partial t} p_{t|0,1}^* \right) p_{0,1} d\mathbf{x}_{0,1} \\
 &= \int \left(-\nabla \cdot (u_{t|0,1} p_{t|0,1}^*) + \frac{1}{2} \sigma^2 \Delta p_{t|0,1}^* \right) p_{0,1} d\mathbf{x}_{0,1} \\
 &= \int \left(-\nabla \cdot (u_{t|0,1} p_{t|0,1}^*) \right) p_{0,1} d\mathbf{x}_{0,1} + \int \left(\frac{1}{2} \sigma^2 \Delta p_{t|0,1}^* \right) p_{0,1} d\mathbf{x}_{0,1} \\
 &= -\nabla \cdot \int (u_{t|0,1} p_{t|0,1}^*) p_{0,1} d\mathbf{x}_{0,1} + \frac{1}{2} \sigma^2 \Delta \int p_{t|0,1}^* p_{0,1} d\mathbf{x}_{0,1} \\
 &= -\nabla \cdot \underbrace{\int (u_{t|0,1} p_{t|0,1}^*) p_{0,1} d\mathbf{x}_{0,1}}_{(u_t^* p_t^*)} + \underbrace{\frac{1}{2} \sigma^2 \Delta \int p_{t|0,1}^* p_{0,1} d\mathbf{x}_{0,1}}_{p_t^*} \tag{74}
 \end{aligned}$$

For (74) to satisfy the Fokker-Planck equation, we set the first term equal to $(u_t^* p_t^*)$ and solve for u_t^* to get

$$u_t^* p_t^* = \int (u_{t|0,1}^* p_{t|0,1}^*) p_{0,1} d\pi_{0,1} \tag{75}$$

$$u_t^* = \frac{\mathbb{E}_{\pi_{0,1}} [u_{t|0,1}^* p_{t|0,1}^*]}{p_t^*} \tag{76}$$

Therefore, the optimal Markovian drift (or *Markovian projection*) is the *average* of the conditional drifts defined in part 1 as $u_{t|0,1}^* \equiv \dot{\mathbf{x}}_{t,\eta}^*$ over the joint distribution $p_{0,1}$. This means that the minimizer of the conditional flow matching loss $\mathcal{L}_{\text{flow}}(\theta)$ in (13) defined as the expected mean-squared error between a Markovian drift field $u_t(X_t)$ and $\dot{\mathbf{x}}_{t,\eta}^*$ over pairs $(\mathbf{x}_0, \mathbf{x}_1) \sim p_{0,1}$ in the dataset is the optimal drift $u_t^*(X_t)$ that solves the GSB problem.

$$u_t^*(X_t) = \arg \min_{\theta} \mathcal{L}_{\text{flow}}(\theta) = \arg \min_{u_t^{\theta}} \int_0^1 \mathbb{E}_{(\mathbf{x}_0, \mathbf{x}_1) \sim p_{0,1}} \|\dot{\mathbf{x}}_{t,\eta}^* - u_t^{\theta}(X_t)\|^2 dt \tag{77}$$

which concludes the proof of Proposition 3. Since the drift for each branch $u_{t,k}^{\theta}(X_t)$ are trained independently in Stage 2, we can extend this result across all $K+1$ branches and conclude that the sequential Stage 1 and Stage 2 training procedures yields the optimal set of drifts $\{u_{t,k}^*\}_{k=0}^K$ that generate the optimal probability paths $\{p_{t,k}^*\}_{k=0}^K$ that solves the GSB problem for each branch.

C.4 PROOF OF PROPOSITION 4

Lemma 2. Suppose the optimal drift field $u_{t,k}^* : \mathbb{R}^d \rightarrow \mathbb{R}^d$ and probability density $p_{t,k}^* : \mathbb{R}^d \rightarrow \mathbb{R}$ that minimizes the GSB problem in (3) is well-defined over the state space $\mathcal{X} \subseteq \mathbb{R}^d$ for each branch. Then, the optimal weight $w_{t,k}^*$ at any of the secondary branches $k \in \{1, \dots, K\}$ is non-decreasing over the interval $t \in [0, 1]$. Equivalently, the optimal growth rates $g_{t,k}^*(X_t) \geq 0$ for all $t \in [0, 1]$.

Proof. We will prove this lemma by contradiction. Suppose there exists a branch k that decreases in weight over the time interval $[t_1, t_2]$ for $0 \leq t_1 < t_2 \leq 1$, such that $w_{t_1,k} > w_{t_2,k}$. We know that the target weight at time $t = 1$ is non-negative $w_{1,k} \geq 0$ and the total weight across all branches is conserved (i.e. $\sum_{k=0}^K w_{t,k}(X_t) = w_t^{\text{total}}$) and non-decreasing (i.e. $\partial_t w_t^{\text{total}} > 0$ for all $t \in [0, 1]$). In this proof, we let $\mathcal{E}_k(t_1, t_2) = \int_{t_1}^{t_2} \left[\frac{1}{2} \|u_{t,k}^*(X_t)\|_2^2 + V_t(X_t) \right]$ denote the energy of following the dynamics of the k th branch over the time interval $[t_1, t_2]$.

Then, there can only be two possible reasons for the loss of mass along a branch k : (1) the mass is destroyed, or (2) the mass is transferred to a different branch $j \neq k$.

1512 **Case 1.** The destruction of mass directly violates the assumption that the total mass across all
 1513 branches is conserved and non-decreasing. So, we only need to consider the possibility of Case 2.
 1514

1515 **Case 2.** Suppose the mass is transferred to a different branch $j \neq k$ over the interval $[t_1, t_2]$. By
 1516 Proposition 3, Stage 1 and 2 training yields the optimal velocity fields $\{u_{t,k}^*(X_t)\}_{k=0}^K$ that generate
 1517 the optimal interpolating probability density $\{p_{t,k}^*(X_t)\}_{k=0}^K$ that independently minimize the GSB
 1518 problem in (3).

1519 If mass is transferred from branch k to branch j over the interval $[t_1, t_2]$, it must be compensated
 1520 for over time $[t_2, 1]$ to reach the target weight $w_{1,k} > 0$. Then, without loss of generality, we can
 1521 consider two sub-cases: (1.1) the mass is compensated from the primary branch, and (1.2) the mass
 1522 that diverges to the j th branch returns to the k th branch following a continuous trajectory.

1523 *Case 1.1.* Since all mass along the K secondary branches originates from the primary branch, this
 1524 implies that there exists a positive weight $\tilde{w} > 0$ that first follows the dynamics of branch k and
 1525 is transferred to branch j , contributing to the final weight of the j th endpoint, and the total weight
 1526 supplied from the primary branch to branch k is $w_{1,k} + \tilde{w}$. Given that each branch has no capacity
 1527 constraints, it follows that the dynamics along branch k over $[0, t_1]$ and branch j over $[t_1, 1]$ is optimal
 1528 for all mass reaching endpoint j . Formally, we express this in terms of energy as

$$\mathcal{E}_k(0, t_1) + \mathcal{E}_j(t_1, 1) < \mathcal{E}_j(0, 1) \quad (78)$$

1529 which contradicts the assumption that the dynamics of branch j given by $(u_{t,j}^*, p_{t,j}^*)$ are optimal over
 1530 the state space \mathcal{X} .
 1531

1533 *Case 1.2.* In this case, there exists a positive mass that follows the dynamics of branch k over the
 1534 interval $[0, t_1]$, the dynamics of branch j over $[t_1, t_2]$, and back to branch k over $[t_2, 1]$. Similarly to
 1535 Case 1.1, given that each branch has no capacity constraints, this implies that this concatenation of
 1536 dynamics is optimal for all mass reaching endpoint k . Expressing in terms of energy, we have

$$\mathcal{E}_k(0, t_1) + \mathcal{E}_j(t_1, t_2) + \mathcal{E}_k(t_2, 1) < \mathcal{E}_k(0, 1) \quad (79)$$

1538 which contradicts the assumption that the dynamics of branch k given by $(u_{t,k}^*, p_{t,k}^*)$ are optimal over
 1539 the state space \mathcal{X} .
 1540

1541 Given that both sub-cases lead to a contradiction of the optimality assumption, we conclude that mass
 1542 along each branch cannot be transferred to another branch and is non-decreasing over $t \in [0, 1]$.

1543 Note that we do not need to consider the case where the mass is compensated from another secondary
 1544 branch $\ell \neq j$, as this would imply that mass is transferred from branch ℓ to branch k , which is not
 1545 possible under the same argument.

1547 **Proposition 4 (Existence of Optimal Growth Functions)** *Assume the state space $\mathcal{X} \subseteq \mathbb{R}^d$ is a
 1548 bounded domain within \mathbb{R}^d . Let the optimal probability density of branch k be a known non-negative
 1549 function bounded in $[0, 1]$, denoted as $p_{t,k}^* : \mathcal{X} \times [0, 1] \rightarrow [0, 1] \in L^\infty(\mathcal{X} \times [0, 1])$. By Lemma 2, we
 1550 can define the set of feasible growth functions in the set of square-integrable functions L^2 as*

$$\mathcal{G} := \{g = (g_{t,0}, \dots, g_{t,K}) \in L^2(\mathcal{X} \times [0, 1]) \mid g_{t,k}(\mathbf{x}) : \mathcal{X} \times [0, 1] \rightarrow \mathbb{R}, g_{t,k}(\mathbf{x}) \geq 0\} \quad (80)$$

1552 Let the growth loss be the functional $\mathcal{L}(g) : L^2(\mathcal{X} \times [0, 1]) \rightarrow \mathbb{R}$. Then, there exists an optimal
 1553 function $g^* = (g_{t,0}^*, \dots, g_{t,K}^*) \in L^2$ where $g_{t,k}^* \in \mathcal{G}$ such that $\mathcal{L}(g^*) = \inf_{g \in \mathcal{G}} \mathcal{L}(g)$ which can be
 1554 obtained by minimizing $\mathcal{L}(g)$ over \mathcal{G} .
 1555

1556 *Proof.* This proof draws on several concepts from functional and convex analysis. For a comprehensive
 1557 background on these concepts, see Benešová & Kružík (2016). We follow the *direct method in the*
 1558 *calculus of variations* (Dacorogna, 1989) for proving there exists a minimizer for the functional
 1559 $\mathcal{L}_{\text{growth}}(g)$ in (17) with the following steps:

1. Show that the set \mathcal{G} of *feasible* growth functions is convex and closed under the weak topology
 of the set of square integrable functions L^2 (Lemma 3).
2. Show that the minimizing sequence $\{g^{(n)}\}$ has a weakly convergent subsequence (Lemma 4).
3. Show that the functional $\mathcal{L}_{\text{growth}}(g)$ is weakly lower semi-continuous (Lemmas 5, 6, 7, 8).

1560 We prove each with a sequence of Lemmas.
 1561
 1562
 1563
 1564
 1565

1566 **Lemma 3.** *The set of feasible growth functions $\mathcal{G} := \{g_{t,k}(\mathbf{x}) : \mathcal{X} \times [0, 1] \rightarrow \mathbb{R} \in L^2(\mathcal{X} \times [0, 1]) \mid g_{t,k}(\mathbf{x}) \geq 0\}$ is convex and closed over the weak topology of L^2 .*

1569 *Proof.* We first prove convexity and then closure under the weak topology of L^2 .

1570 *Proof of Convexity.* To prove that the set of functions \mathcal{G} is convex, we first define what it means for a
1571 set of functions to be convex.

1572 **Definition 4** (Convex Set). *A set of functions \mathcal{G} is said to be convex if any convex combination of two
1573 functions in the set $g_1, g_2 \in \mathcal{G}$ is also in the set. Formally, \mathcal{G} if*

$$1575 \quad \forall g, g' \in \mathcal{G}, \quad \forall \lambda \in [0, 1], \quad g^\lambda = \lambda g + (1 - \lambda)g' \in \mathcal{G} \quad (81)$$

1577 Recall our definition of \mathcal{G} as the set of functions that return strictly non-negative values:

$$1578 \quad \mathcal{G} := \{g_{t,k}(\mathbf{x}) : \mathcal{X} \times [0, 1] \rightarrow \mathbb{R} \in L^2(\mathcal{X} \times [0, 1]) \mid g_{t,k}(\mathbf{x}) \geq 0\} \quad (82)$$

1580 Given any $\lambda \in [0, 1]$, we have

$$\begin{aligned} 1581 \quad g_{t,k}^\lambda(\mathbf{x}) &= \lambda g_{t,k}(\mathbf{x}) + (1 - \lambda)g'_{t,k}(\mathbf{x}) \\ 1582 \quad g_{t,k}^\lambda(\mathbf{x}) &\geq \lambda \cdot 0 + (1 - \lambda) \cdot 0 \\ 1583 \quad g_{t,k}^\lambda(\mathbf{x}) &\geq 0 \end{aligned} \quad (83)$$

1585 which means $g_{t,k}^\lambda \in \mathcal{G}$ and \mathcal{G} is convex.

1587 *Proof of Closure.* First, we define what it means to be closed under the weak topology of L^2 .

1588 **Definition 5** (Closure Under Weak Topology of L^2). *A set of functions \mathcal{G} is said to be closed in the
1589 weak topology of L^2 if the statement is true: if a sequence of functions $\{g^{(n)} : g^{(n)} \in \mathcal{G}\}$ indexed by
1590 n , converges to some function $g^{(\infty)} \in L^2$ as $n \rightarrow \infty$, then $g^{(\infty)} \in \mathcal{G}$.*

1592 To show that \mathcal{G} is closed under weak topology of L^2 , we need to show that all sequences $\{g_{t,k}^{(n)} : g_{t,k}^{(n)} \in \mathcal{G}\}$ converge to $g \in \mathcal{G}$ as $n \rightarrow \infty$, such that $g \geq 0$. The proof follows directly from Fatou's
1593 Lemma (Ekeland & Temam, 1999) which states that given a sequence of non-negative, measurable
1594 functions $\{g_{t,k}^{(n)} : g_{t,k}^{(n)} \in \mathcal{G}\}$, the following is true

$$1597 \quad 0 \leq \int_{\mathcal{X}} \liminf_{n \rightarrow \infty} g_{t,k}^{(n)}(\mathbf{x}) d\mathbf{x} \leq \liminf_{n \rightarrow \infty} \int_{\mathcal{X}} g_{t,k}^{(n)}(\mathbf{x}) d\mathbf{x} \quad (\text{Fatou's Lemma})$$

1600 meaning that the limit of a converging sequence of non-negative functions $g_{t,k}^{(\infty)}$ is also non-negative
1601 over the state space \mathcal{X} , and thus is in the set of feasible growth functions $g_{t,k}^{(\infty)} \in \mathcal{G}$, concluding our
1602 proof.

1603 **Lemma 4.** *Given a minimizing sequence $\{g^{(n)} \in \mathcal{G}\}$ under the functional $\mathcal{L}(g) : L^2 \rightarrow \mathbb{R}$ such
1604 that $\mathcal{L}(g^{(n)}) \rightarrow \inf_{g \in \mathcal{G}} \mathcal{L}(g)$, there exists a subsequence $\{g^{(n_i)} \in \mathcal{G}\}$ that converges weakly in L^2 to
1605 some limit $g^* \in \mathcal{G}$.*

1607 *Proof.* It suffices to show the following:

- 1609 1. The functional $\{g^{(n)}\}$ is bounded in L^2 such that there exists a positive value M where
1610 $\|g^{(n)}\|_{L^2} \leq M$ for all n .
- 1611 2. The space of feasible functions \mathcal{G} is reflexive, such that all bounded sequences have a weakly
1612 convergent subsequence in L^2 .

1613 The proof of Part 1 follows from the growth penalty term in $\mathcal{L}(g)$ defined as the squared-norm $\|g\|_{L^2}^2$
1615 of the growth functional. This term ensures that $\mathcal{L}(g)$ is *coercive*, such that

$$1616 \quad \|g^{(n)}\|_{L^2} \rightarrow \infty \implies \mathcal{L}(g^{(n)}) \rightarrow \infty \quad (84)$$

1617 which ensures that the sequence does not diverge to infinity in norm without incurring a penalty from
1618 the loss functional. Given that $\mathcal{L}(g^{(n)}) \rightarrow \inf_{g \in \mathcal{G}} \mathcal{L}(g)$, where $\inf_{g \in \mathcal{G}} \mathcal{L}(g) < \infty$, it follows from
1619 coercivity that $\|g^{(n)}\|_{L^2}$ does not diverge to infinity and $\{g^{(n)}\}$ is bounded in L^2 .

1620 The proof of Part 2 follows from the well-established result that for $1 < p < \infty$, the space L^p is
 1621 reflexive (Beauzamy, 2011). Therefore, by reflexivity of L^2 , we have that every minimizing sequence
 1622 $\{g^{(n)}\}$ has a weakly convergent subsequence $\{g^{(n_i)}\}$, such that $g^{(n_i)} \rightarrow g^*$ in L^2 , concluding our
 1623 proof.

1624 Before proving that each component of the loss functional is weakly lower semi-continuous, we
 1625 establish the definitions for weakly continuous and weakly lower semi-continuous functionals.

1627 **Definition 6** (Weakly Continuous Functionals). *A functional $\mathcal{L} : L^2 \rightarrow \mathbb{R}$ is said to be weakly
 1628 continuous if it satisfies*

$$1629 \quad n \rightarrow \infty \implies g^{(n)} \rightarrow g \implies \mathcal{L}(g^{(n)}) \rightarrow \mathcal{L}(g) \quad (85)$$

1631 such that if a sequence $\{g^{(n)} : g^{(n)} \in \mathcal{G}\}$ converges $g^{(n)} \rightarrow g$ as $n \rightarrow \infty$, then the functional also
 1632 converges $\mathcal{L}(g^{(n)}) \rightarrow \mathcal{L}(g)$.

1633 **Definition 7** (Weakly Lower Semi-Continuous Functionals). *A functional $\mathcal{L} : L^2 \rightarrow \mathbb{R}$ is said to be
 1634 weakly lower semi-continuous (w.l.s.c.) if it satisfies*

$$1636 \quad n \rightarrow \infty \implies g^{(n)} \rightarrow g \implies \liminf_{n \rightarrow \infty} \mathcal{L}(g^{(n)}) \geq \mathcal{L}(g) \quad (86)$$

1638 such that if a sequence $\{g^{(n)} : g^{(n)} \in \mathcal{G}\}$ converges $g^{(n)} \rightarrow g$ as $n \rightarrow \infty$, then the functional is
 1639 lower bounded by $\mathcal{L}(g)$. By definition, weak continuity implies w.l.s.c.

1640 **Lemma 5.** *The functional $\mathcal{L}_{\text{match}}(g) : L^2 \rightarrow \mathbb{R}$ defined as*

$$\begin{aligned} 1642 \quad \mathcal{L}_{\text{match}}(g) &= \sum_{k=0}^K \mathbb{E}_{p_{t,k}^*} \left(w_{0,k} + \int_0^1 g_{t,k} dt - w_{1,k}^* \right)^2 \\ 1643 &= \sum_{k=0}^K \left(w_{0,k} + \int_0^1 \mathbb{E}_{p_{t,k}^*}[g_{t,k}] dt - w_{1,k}^* \right)^2 \quad (\text{linearity of expectation}) \\ 1644 &= \sum_{k=0}^K \left(\int_0^1 \int_{\mathcal{X}} p_{t,k}^* g_{t,k} d\mathbf{x} dt + c \right)^2 \end{aligned} \quad (87)$$

1650 where $c = w_{0,k} - w_{1,k}^*$ is a constant, is weakly lower semi-continuous in L^2 (w.l.s.c.).

1652 *Proof of Lemma.* First, we show that the map $\phi(g_{t,k}) = \int_{\mathcal{X}} p_{t,k}^* g_{t,k} d\mathbf{x}$ is a bounded linear functional
 1653 in L^2 . By linearity of integration, we have that for two functions $g, g' \in \mathcal{G}$, $\phi(cg + c'g') =$
 1654 $c\phi(g) + c'\phi(g')$, so $\phi(\cdot)$ is a linear map. To establish that $\phi(\cdot)$ is bounded in L^2 , we must show that
 1655 $|\phi(g)| \leq C\|g\|_{L^2}$. By the Cauchy-Schwarz Inequality (Steele, 2004), we have

$$1657 \quad |\phi(g_{t,k})| = |\langle p_{t,k}^*, g_{t,k} \rangle| = \left| \int_{\mathcal{X}} p_{t,k}^* g_{t,k} d\mathbf{x} \right| \leq \|p_{t,k}^*\|_{L^2} \|g_{t,k}\|_{L^2} \quad (88)$$

1660 which is valid as $p_{t,k}^*, g_{t,k} \in L^2$. Since $\|p_{t,k}^*\|_{L^2}$ is a fixed constant with respect to $g_{t,k}$, we have
 1661 shown that $|\phi(g_{t,k})| \leq C\|g_{t,k}\|_{L^2}$ and $g_{t,k}$ is bounded. By the definition of weak topology on L^2 ,
 1662 all bounded linear functionals are weakly continuous, such that $\lim_{n \rightarrow \infty} \phi(g^{(n)}) \rightarrow \phi(g)$. Given
 1663 that $\phi(g)$ is weakly continuous, as $g^{(n)} \rightarrow g^*$, we have $\phi(g^{(n)}) \rightarrow \phi(g^*)$. Since the square function
 1664 $\psi(\cdot) = (\cdot)^2$ is convex and continuous and bounded below by $\psi \geq 0$, the function does not contain
 1665 discontinuities and $\lim_{n \rightarrow \infty} (\phi(g^{(n)}) - c)^2 \geq (\phi(g^*) - c)^2$, which is the definition of a w.l.s.c.
 1666 functional. Since the sum of w.l.s.c. functionals is also w.l.s.c., we conclude our proof.

1667 **Lemma 6.** *The functional $\mathcal{L}_{\text{energy}}(g) : L^2 \rightarrow \mathbb{R}$ defined as*

$$\begin{aligned} 1669 \quad \mathcal{L}_{\text{energy}}(g) &= \int_0^1 \mathbb{E}_{p_{t,k}^*} \left[\frac{1}{2} \|u_{t,k}(X_t)\|^2 + V_t(X_t) \right] \left(w_{0,k} + \int_0^t g_{s,k}(X_s) ds \right) dt \\ 1670 &= \int_0^1 \alpha(t) \left(w_{0,k} + \int_0^t \int_{\mathcal{X}} p_{s,k}^* g_{s,k} d\mathbf{x} ds \right) dt \quad (\text{linearity of expectation}) \end{aligned}$$

1673 where $\alpha(t)$ is a constant not dependent on $g_{t,k}$, is weakly lower semi-continuous in L^2 (w.l.s.c.).

1674 *Proof of Lemma.* Following a similar proof as Lemma 5, we have that the map $\phi(g_{s,k}) =$
 1675 $\int_{\mathcal{X}} p_{s,k}^* g_{s,k} d\mathbf{x}$ is a bounded linear functional, and thus is weakly continuous in L^2 . Given that
 1676 $\phi(g)$ is bounded, $t \in [0, 1]$, and linearity of integration holds, the integral of $\phi(g)$ over the interval
 1677 $s \in [0, t]$ remains a bounded linear functional in L^2 , and thus is weakly continuous in L^2 . Following
 1678 similar logic, the outer integral $\int_0^1 \alpha(t) \left(w_{0,k} + \int_0^t \phi(g) ds \right) dt$ is also a bounded linear functional
 1679 given that both $\alpha(t)$ and $w_{0,k}$ are constants, and we have shown that $\mathcal{L}_{\text{energy}}(g)$ is weakly continuous.
 1680 Since weak continuity implies w.l.s.c. and the sum of w.l.s.c. functionals is also w.l.s.c., we conclude
 1681 our proof.
 1682

1683 **Lemma 7.** *The mass loss functional $\mathcal{L}_{\text{mass}}(g) : L^2 \rightarrow \mathbb{R}$ defined as*

$$\begin{aligned} 1685 \quad \mathcal{L}_{\text{mass}}(g) &= \int_0^1 \mathbb{E}_{\{p_{t,k}\}_{k=0}^K} \left[\sum_{k=0}^K w_{t,k}^\phi(\mathbf{x}_{t,k}) - w_t^{\text{total}} \right]^2 dt \\ 1686 &= \int_0^1 \left[\sum_{k=0}^K \left(w_{0,k} + \int_0^t \mathbb{E}_{p_{t,k}}[g_{t,k}] \right) - w_t^{\text{total}} \right]^2 dt \quad (\text{linearity of expectation}) \\ 1687 &= \int_0^1 \left[1 + \sum_{k=0}^K \int_0^t \mathbb{E}_{p_{t,k}}[g_{t,k}] - w_t^{\text{total}} \right]^2 dt \end{aligned} \quad (89)$$

1693 where w_t^{total} is a constant not dependent on $g_{t,k}$, is weakly lower semi-continuous in $(L^2)^{K+1}$
 1694 (w.l.s.c.). Note that we do not need to account for the negative penalty loss as we assume the growth
 1695 function is strictly positive.
 1696

1697 *Proof of Lemma.* Following a similar proof as Lemma 5, we have that the map $\phi(g) = \int_{\mathcal{X}} p_{s,k}^* g_{s,k} d\mathbf{x}$
 1698 is a bounded linear functional, and thus is weakly continuous in L^2 . Since the sum of bounded linear
 1699 functionals across $K+1$ branches is also weakly continuous in L^2 , following similar logic to Lemma
 1700 5, the composition of a weakly continuous functional with the convex and continuous square map is
 1701 w.l.s.c.. Since the sum of w.l.s.c. functionals is also w.l.s.c., $\mathcal{L}_{\text{mass}}(g)$ is w.l.s.c., which concludes our
 1702 proof.

1703 **Lemma 8.** *The combined Stage 3 training loss $\mathcal{L}_{\text{growth}}(g) = \lambda_{\text{energy}} \mathcal{L}_{\text{energy}}(g) + \lambda_{\text{match}} \mathcal{L}_{\text{match}}(g) +$*
 1704 $\lambda_{\text{mass}} \mathcal{L}_{\text{mass}}(g) + \lambda_{\text{growth}} \|g\|_{L^2}$ *is weak lower semi-continuous in L^2 (w.l.s.c.).*
 1705

1706 *Proof of Lemma.* Given the weighted sum of w.l.s.c. functionals $\sum_{i=1}^M \lambda_i \mathcal{L}_i(g)$ that each satisfy
 1707 $\lim_{n \rightarrow \infty} \mathcal{L}_i(g^{(n)}) \geq \mathcal{L}_i(g)$ for constants $\{\lambda_i\}_{i=1}^M$ for a converging sequence $g^{(n)} \rightarrow g$ as $n \rightarrow \infty$, it
 1708 follows easily that
 1709

$$\lim_{n \rightarrow \infty} \sum_{i=1}^M \mathcal{L}_i(g^{(n)}) \geq \sum_{i=1}^M \mathcal{L}_i(g) \quad (90)$$

1713 By definition, the norm in L^2 is lower semi-continuous in the strong topology and thus the weak
 1714 topology. Since $\mathcal{L}_{\text{energy}}$, $\mathcal{L}_{\text{match}}$, and $\mathcal{L}_{\text{mass}}$ are w.l.s.c. by Lemmas 5, 6, and 7, we conclude that
 1715 $\mathcal{L}_{\text{growth}}(g)$ is w.l.s.c., which concludes the proof.

1716 In total, we have shown:
 1717

1. The set \mathcal{G} of feasible growth functions is convex and closed under the weak topology of the set
 1719 of square integrable functions L^2 (Lemma 3).
2. The minimizing sequence $\{g^{(n)}\}$ has a weakly convergent subsequence (Lemma 4).
3. The functional $\mathcal{L}(g)$ is weakly lower semi-continuous (Lemmas 5, 6, 7, 8).

1724 Thus, by the *direct method in the calculus of variations*, we have shown

$$\exists g^* = (g_{t,0}^*, \dots, g_{t,1}^*) \in \mathcal{G} \text{ s.t. } \mathcal{L}(g^*) = \inf_{g \in \mathcal{G}} \mathcal{L}(g) \quad (91)$$

1725 which concludes the proof of Proposition 4.
 1726
 1727

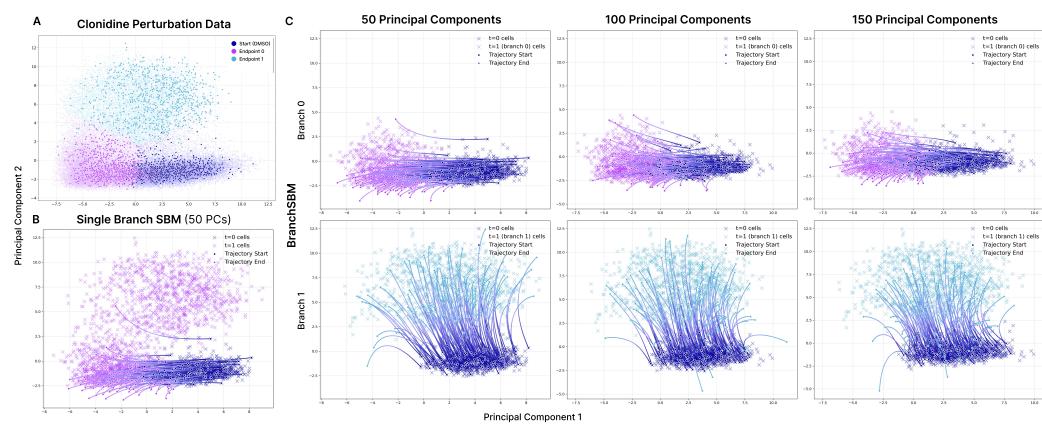


Figure 6: **Results for Clonidine Perturbation Modeling.** (A) Gene expression data of DMSO control (set to $t = 0$) and cell states (set to $t = 1$) after Clonidine perturbation with two distinct endpoints (pink and purple). (B) The simulated trajectories for single-branch SBM on the top 50 PCs with both clusters. All samples take the low-energy path without reaching the second cluster. (C) The simulated endpoints of the top 50, 100, and 150 PCs at $t = 1$ on the validation data for each branch.

D ADDITIONAL EXPERIMENTS AND DISCUSSIONS

D.1 COMPARISON TO SINGLE-BRANCH SCHRÖDINGER BRIDGE MATCHING

Setup For each experiment, we compared the performance of BranchSBM against single-branch SBM. Instead of clustering the data at $t = 1$ into distinct endpoint distributions, we left the unclustered data as a single target distribution p_1 and let the model learn the optimal Schrödinger Bridge from the initial distribution π_0 . For the single-branch task, we assume mass conservation and set the weights for all samples from π_0 and π_1 to 1.0, while keeping the model architecture, state-cost V_t , and hyperparameters equivalent to BranchSBM. Since single-branch SBM does not require modeling the growth of separate branches, we train only Stages 1 and 2 to optimize the drift field u_t^θ of the single branch. For evaluation, the trajectories of validation samples from the initial distribution $\mathbf{x}_0 \sim \pi_0$ were simulated over $t \in [0, 1]$ and compared with the ground truth distribution at $t = 1$. For BranchSBM, we take the overall distribution generated across all branches and compare it to the overall ground truth distribution.

Modeling Mouse Hematopoiesis Differentiation In Figure 7, we provide a side-by-side comparison of the reconstructed distributions at time t_1 and t_2 using BranchSBM (top) with two branches and single-branch SBM (bottom), as well as the learned trajectories over the time interval $t \in [t_1, t_2]$. While single-branch SBM produces samples that loosely capture the two target cell fates, the resulting distributions display high variance and fail to align with the true differentiation trajectories. In contrast, BranchSBM generates intermediate distributions that are sharply concentrated along the correct developmental paths, more faithfully reflecting the underlying branching structure of the data.

Modeling Clonidine and Trametinib Perturbation For both Clonidine and Trametinib, we performed the single-branch experiment on the top 50 principal components identified by PCA. After training, we simulated all the validation samples from the initial distribution π_0 by integrating the single velocity field u_t^θ to $t = 1$. We evaluated the performance of the single-branch parameterization by computing the RBF-MMD (97) of all PCs and \mathcal{W}_1 (95) and \mathcal{W}_2 (96) distances of the top-2 PCs of the simulated samples at time $t = 1$ with the ground truth data points.

In Table 4, we show that BranchSBM with two branches trained on gene expression vectors across all dimensions $d \in \{50, 100, 150\}$ outperforms single-branch SBM on dimension $d = 50$ in reconstructing the distribution of cells perturbed with Clonidine. In Table 5, we further show improved performance of BranchSBM with three branches compared to single-branch SBM.

In Figure 8A and 8B, we see that single-branch SBM only reconstructs cluster 0, while failing to generate samples from cluster 1 for Clonidine perturbation or clusters 1 and 2 for Trametinib

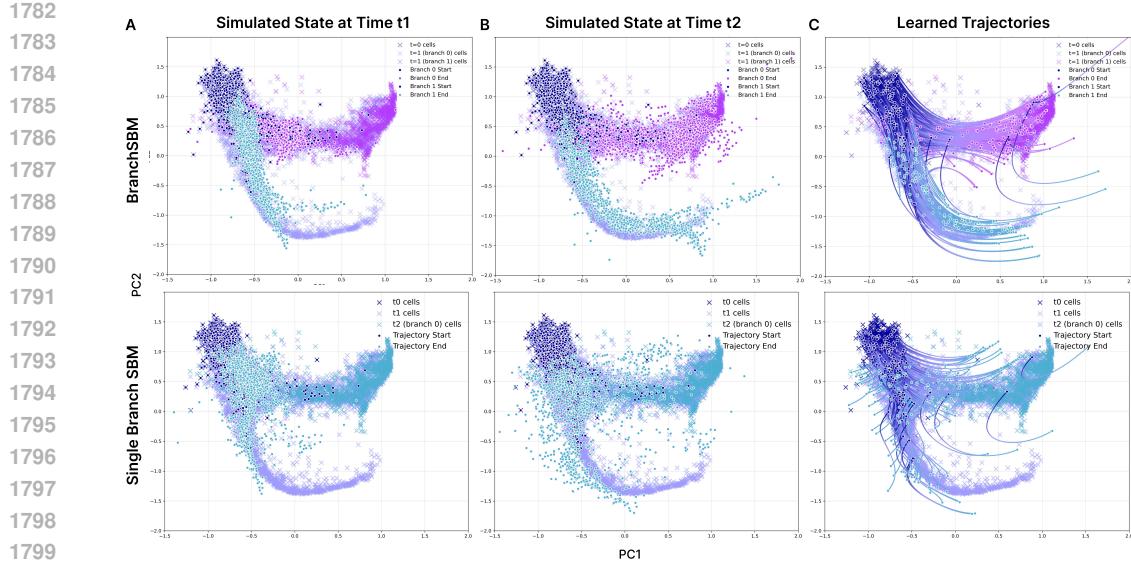


Figure 7: **Comparison of BranchSBM to Single-Branch SBM for Cell-Fate Differentiation.** Mouse hematopoiesis scRNA-seq data is provided for three time points t_0, t_1, t_2 . (A, B) Distribution of simulated cell states at time (A) t_1 and (B) t_2 across both branches for BranchSBM (top) and single-branch SBM (bottom). (C) Learned trajectories of BranchSBM and single-branch SBM over the $t \in [t_1, t_2]$ on validation samples.

perturbation. The endpoints for both perturbation experiments are clustered largely on variance in the first two or three principal components (PCs), where PC1 captures the divergence from the control state to the perturbed clusters and PC2 and higher captures the divergence between clusters in the perturbed population, where cluster 0 is closest to the control state along PC2 and cluster 1 and 2 are farther from the control. From Figure 8A and B, we can conclude that single-branch SBM is not expressive enough to capture the complexities of higher-dimensional PCs and follows the most obvious trajectory from the control to cluster 0, resulting in an inaccurate representation of the perturbed cell population.

In contrast, we demonstrate that BranchSBM is capable of stimulating trajectories to both clusters in the population perturbed with Clonidine (Figure 8B) and three clusters in the population perturbed with Trametinib (Figure 8D), generating branched distributions that accurately capture the location and spread of the perturbed cell distribution in the dataset.

D.2 EFFECT OF FINAL JOINT TRAINING ON LOSSES

In Table 7, we show the final losses after convergence, summed across each branch and averaged over the batch size, for Stage 3 training of only the growth networks and Stage 4 joint training of the flow and growth networks discussed in Section 4. All losses are calculated exactly as shown in Section 4. Crucially, we find that the final joint training stage refines the parameters of both the flow and growth networks simultaneously to minimize the energy loss $\mathcal{L}_{\text{energy}}(\theta, \phi)$ (14) while ensuring the growth parameters maintain minimal losses across $\mathcal{L}_{\text{match}}(\phi)$ (15) and $\mathcal{L}_{\text{mass}}(\phi)$ (16) for all experiments. This indicates that jointly optimizing both the drift and growth dynamics leads to further refinement towards modeling the optimal branching trajectories in the data.

E EXPERIMENT DETAILS

E.1 MULTI-STAGE TRAINING

To ensure stable training while incorporating all loss functions, we introduce a multi-stage training approach (Algorithm 1).

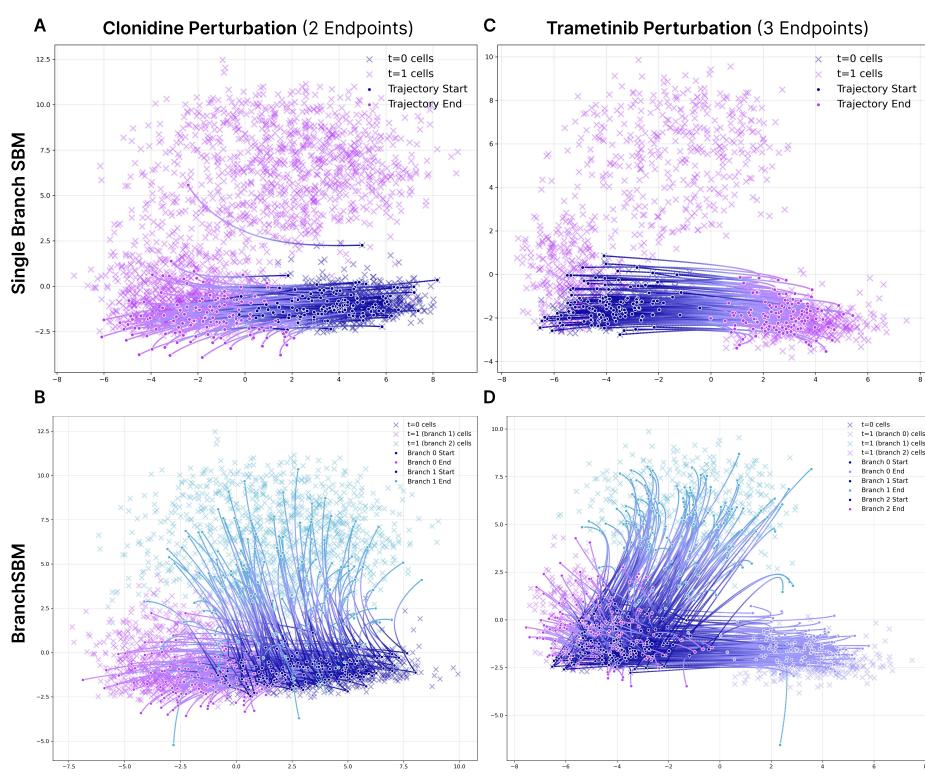


Figure 8: **Comparison of BranchSBM to Single-Branch SBM for Perturbation Modelling.** (A, B) Clonidine perturbation trajectories with two target clusters generated by (A) single-branch SBM and (B) BranchSBM from the validation data. (C, D) Trametinib perturbation trajectories with three target clusters generated by (C) single-branch SBM and (D) BranchSBM. In both experiments, single-branch SBM only generated states in cluster 0 and not cluster 1 or 2, whereas BranchSBM reconstructed all perturbed clusters via branched trajectories.

Table 7: **Validation Losses for Stage 3 and 4 Training Across Experiments.** Losses are summed across both branches and averaged over batch size. The final Stage 4 joint training stage that refines both the flow and growth networks simultaneously minimize the energy loss $\mathcal{L}_{\text{energy}}(\theta, \phi)$ (14) from Stage 3 while ensuring the growth parameters maintain minimal losses across $\mathcal{L}_{\text{match}}(\phi)$ (15) and $\mathcal{L}_{\text{mass}}(\phi)$ (16)

Experiment	Stage 3			Stage 4		
	$\mathcal{L}_{\text{energy}}(\theta, \phi)$	$\mathcal{L}_{\text{mass}}(\theta, \phi)$	$\mathcal{L}_{\text{match}}(\theta, \phi)$	$\mathcal{L}_{\text{energy}}(\theta, \phi)$	$\mathcal{L}_{\text{mass}}(\theta, \phi)$	$\mathcal{L}_{\text{match}}(\theta, \phi)$
LiDAR	1.276	3.0×10^{-5}	0.007	0.768	2.0×10^{-5}	0.102
Mouse Hematopoiesis	2.209	1.2×10^{-4}	0.054	1.918	5.0×10^{-5}	0.076
Chlonidine Perturbation	36.469	0.030	0.109	25.798	0.053	0.153
Trametinib Perturbation	35.834	0.023	0.078	32.843	0.017	0.056

Stage 1 First, we train a neural interpolant $\varphi_{t,\eta}(\mathbf{x}_0, \mathbf{x}_{t,k}) : \mathbb{R}^d \times \mathbb{R}^d \times [0, 1] \rightarrow \mathbb{R}^d$ that takes the endpoints of a branched coupling and defines the optimal interpolating state X_t at time t by minimizing the energy function $\mathcal{L}_{\text{traj}}(\eta)$ in (12). This is used to calculate the optimal conditional velocity $\dot{\mathbf{x}}_{t,\eta,k}$ that preserves the endpoints $X_0 = \mathbf{x}_0$ and $X_{1,k} = \mathbf{x}_{1,k}$ for the flow matching objective in Stage 2.

Stage 2 Next, we train a set of flow networks $\{u_{t,k}^\theta\}_{k=0}^K$ that generate the optimal interpolating trajectories for each branch with the conditional flow matching objective $\mathcal{L}_{\text{flow}}$ in (13).

1890 **Stage 3** We freeze the parameters of the flow networks and only train the growth networks $\{g_{t,k}^\phi\}_{k=0}^K$
 1891 by minimizing $\mathcal{L}_{\text{growth}}$ in (17).
 1892

1893 **Stage 4** Finally, we unfreeze the parameters of both the flow and growth networks and jointly
 1894 train $\{u_{t,k}^\theta, g_{t,k}^\phi\}_{k=0}^K$ by minimizing the growth loss $\mathcal{L}_{\text{growth}}$ in (17) from Stage 3 in addition to the
 1895 distribution reconstruction loss $\mathcal{L}_{\text{recons}}$ in (19).
 1896

1897 **E.2 GENERAL TRAINING DETAILS**
 1898

1900
 1901 **Leiden Clustering** To identify branch endpoints in the dataset, we apply an automated clustering
 1902 pipeline to the final time point ($t = 7$), where the differentiated cell-states are most clearly defined.
 1903 First, we construct a 20-nearest-neighbor (kNN) graph on the final-timepoint embeddings and
 1904 run Leiden community detection using the `ModularityVertexPartition` objective. Leiden
 1905 clustering is particularly suited and the state-of-the-art method for clustering single-cell data, given
 1906 that it is robust to noise, can handle heterogeneous cluster sizes and shapes, and guarantees well-
 1907 connected clusters, all while being fully unsupervised and automatic. This gives us stable and
 1908 biologically meaningful terminal clusters compared to alternative methods such as K-means clustering,
 1909 which assumes spherical clusters and struggles with sparse high-dimensional manifolds. After initial
 1910 clustering, we merge any cluster with fewer than `min_cells` cells into the nearest large cluster
 1911 based on Euclidean centroid distance.
 1912

1913 **Model Architecture** We parameterized the branched trajectory $\varphi_{t,\eta}(\mathbf{x}_0, \mathbf{x}_1)$ with a 3-layer MLP
 1914 with Scaled Exponential Linear Unit (SELU) activations. The endpoint pair $(\mathbf{x}_0, \mathbf{x}_1)$ and the time
 1915 step t are concatenated into a single $(2d + 1)$ -dimensional vector and used as input to the model.
 1916 Similarly, we parameterize each flow network $u_{t,k}^\theta(\mathbf{x}_t)$ and growth network $g_{t,k}^\theta(\mathbf{x}_t)$ with the same
 1917 3-layer MLP and SELU activations but takes the interpolating state \mathbf{x}_t and time t concatenated into a
 1918 $(d + 1)$ -dimensional vector.
 1919

1920 To ensure that the growth rates of all secondary branches are non-negative (i.e. for all $k \in \{1, \dots, K\}$),
 1921 $g_{t,k}(X_t) \geq 0$, we apply an additional softplus activation to the output of the 3-layer MLP in the
 1922 growth networks, defined as $\text{softplus}(\cdot) = \log(1 + \exp(\cdot))$, which is smooth function that transforms
 1923 negative values to be positive near 0. For the growth network of the primary branch ($k = 0$), we
 1924 allow for both positive and negative growth, as all mass starts at the primary branch and flows into
 1925 the secondary branches, but the primary branch itself can grow as well, depending on whether mass
 1926 is conserved.
 1927

1928 **State Cost V_t** Depending on the dimensionality of the data type, we set the state cost $V_t(X_t) : \mathbb{R}^d \rightarrow [0, +\infty]$ to be either the LAND or RBF metric discussed in Appendix A.2. For the experiments
 1929 on LiDAR ($d = 3$) and Mouse Hematopoiesis scRNA-seq ($d = 2$) data, we used the LAND metric
 1930 (33) with hyperparameters $\sigma = 0.125$ and $\varepsilon = 0.001$.
 1931

1932 To avoid the task of setting a suitable variance σ for the high-dimensional gene expression space
 1933 $d \in \{50, 100, 150\}$, we use the RBF metric (37) that learns parameters to ensure the regions within
 1934 the data manifold have low cost and regions far from the data manifold have high cost. Using the
 1935 training scheme in Kapuśniki et al. (2024), we identified N_c cluster centers with k -means clustering,
 1936 and trained the parameters $\{\omega_{j,n}\}_{n=1}^{N_c}$ by minimizing \mathcal{L}_{RBF} (38) on the training data. We found that
 1937 setting the number of cluster centers N_c too low resulted in non-decreasing \mathcal{L}_{RBF} and that increasing
 1938 N_c for higher dimensions enabled more effective training. Furthermore, we found that modeling
 1939 higher-dimensional PCs required setting a larger κ , which determines the *spread* of the RBF kernel
 1940 around each cluster center. The specific values for N_c and κ depending on the dimension of principal
 1941 components are provided in Table 11.
 1942

1943 **Optimal Transport Coupling** Since our experiments consist of unpaired initial and target distributions
 1944 and we seek to minimize the energy of the interpolating bridge, we define pairings $(\mathbf{x}_0, \mathbf{x}_{1,k})$
 1945 using the optimal transport plan $\pi_{0,1,k}^*$ that minimizes the distance between the initial distribution
 1946 π_0 and each target distribution $\pi_{1,k}$ in probability space. Specifically, we define $\pi_{0,1,k}^*$ as the
 1947

1944 2-Wasserstein transport plan (Villani et al., 2008) between π_0 and $\pi_{1,k}$ defined as
 1945

$$\pi_{0,1,k}^* = \arg \min_{\pi_{0,1,k} \in \Pi} \int_{\pi_0 \otimes \pi_{1,k}} \|\mathbf{x}_0 - \mathbf{x}_{1,k}\|_2^2 d\pi(\mathbf{x}_0, \mathbf{x}_{1,k}) \quad (92)$$

1946 where $\pi_0 \otimes \pi_{1,k}$ is the set of all possible couplings between the endpoint distributions. For each of
 1947 the branches, the dataset was paired such that $(\mathbf{x}_0, \mathbf{x}_{1,k}) \sim \pi_{0,1,k}^*$.
 1948

1949 **Training** We train each stage for a maximum of 100 epochs. For Stage 1, we used the Adam
 1950 optimizer (Kingma & Ba, 2014) with a learning rate of 1.0×10^{-4} to train $\varphi_{t,\eta}(\mathbf{x}_0, \mathbf{x}_1)$. For Stage 2,
 1951 3, and 4, we used the AdamW optimizer (Loshchilov & Hutter, 2017) with weight decay 1.0×10^{-5}
 1952 and learning rate 1.0×10^{-3} to train each flow network $u_{t,k}^\theta$ and growth network $g_{t,k}^\phi$. All experiments
 1953 were performed on one NVIDIA A100 GPU. We trained on the LiDAR and mouse hematopoiesis
 1954 data with a batch size of 128 and the Clonidine and Trametinib perturbation data with a batch size of
 1955 128, each divided with a 0.9/0.1 train/validation split. All hyperparameters across experiments are
 1956 provided in Table 11.
 1957

1958 **Computational Overhead** Although we train $K + 1$ flow and growth networks, the overall training
 1959 time remains comparable to that of single-branch SBM, since the networks for each branch is trained
 1960 only on the subset of data corresponding to its respective target distribution. While the method does
 1961 incur higher space complexity, we find that simple MLP architectures suffice for strong performance,
 1962 suggesting that scalability is not a major concern. BranchSBM also significantly reduces inference
 1963 time, as predicting branching population dynamics requires simulating only a single sample from the
 1964 initial distribution, unlike other models that require simulating large batches of samples.
 1965

1966 E.3 LiDAR EXPERIMENT DETAILS

1967 **LiDAR Data** We used the same LiDAR manifold from Liu et al. (2023a); Kapuśniatek et al. (2024).
 1968 The data is a collection of three-dimensional point clouds within 10 unit cubes $[-5, 5]^3 \subset \mathbb{R}^3$ that
 1969 span the surface of Mount Rainier. Given any point $\mathbf{x} \in \mathbb{R}^3$ in the three-dimensional space, we
 1970 project it onto the LiDAR manifold by identifying the k -nearest neighbors $\{\mathbf{x}_1, \dots, \mathbf{x}_k\}$ and fitting a
 1971 2D tangent plane to the set of neighbors
 1972

$$\arg \min_{a,b,c} \frac{1}{k} \sum_{i=1}^k \exp \left(\frac{-\|\mathbf{x} - \mathbf{x}_i\|}{\tau} \right) (a\mathbf{x}_i^{(x)} + b\mathbf{x}_i^{(y)} + c - \mathbf{x}_i^{(z)})^2 \quad (93)$$

1973 where $\tau = 0.001$ following Liu et al. (2023a). Then, we solve for the tangent plane $ax + by + c = z$
 1974 using the Moore-Penrose pseudoinverse from $k = 20$ neighbors. From the tangent plane, we can
 1975 project any point \mathbf{x} to the LiDAR manifold with the function $\pi(\mathbf{x})$ defined as
 1976

$$\pi(\mathbf{x}) = \mathbf{x} - \left(\frac{\mathbf{x}^\top \mathbf{v} + c}{\|\mathbf{v}\|^2} \right) \mathbf{v}, \quad \text{where } \mathbf{v} = \begin{bmatrix} a \\ b \\ -1 \end{bmatrix} \quad (94)$$

1977 **Synthetic Distributions** To reformulate the experiment in Liu et al. (2023a) as a branching problem,
 1978 we define a single initial distribution and two divergent target distributions. Specifically, we define a
 1979 single initial distribution $\pi_0 = \mathcal{N}(\mu_0, \sigma_0)$ as a mixture of four Gaussians and two target distributions
 1980 $\pi_{1,0}, \pi_{1,1}$ on either side of the mountain, both as mixtures of three Gaussians. The exact parameters
 1981 of each Gaussian are given in Table 8.
 1982

1983 We sample a total of 5000 points i.i.d. from each of the Gaussian mixtures $\{\mathbf{x}_0^i\}_{i=1}^{5000} \sim \pi_0$,
 1984 $\{\mathbf{x}_{1,0}^i\}_{i=1}^{5000} \sim \pi_{1,0}$, and $\{\mathbf{x}_{1,1}^i\}_{i=1}^{5000} \sim \pi_{1,1}$. The data points are projected to the LiDAR mani-
 1985 fold with the projection function $\pi(\mathbf{x})$ in (94).
 1986

1987 **Evaluation Metrics** To determine how closely the simulated trajectories match the ground truth
 1988 trajectories, we compute the Wasserstein-1 (\mathcal{W}_1) and Wasserstein-2 (\mathcal{W}_2) distances defined as
 1989

$$\mathcal{W}_1 = \left(\min_{\pi \in \Pi(p,q)} \int \|\mathbf{x} - \mathbf{y}\|_2 d\pi(\mathbf{x}, \mathbf{y}) \right) \quad (95)$$

$$\mathcal{W}_2 = \left(\min_{\pi \in \Pi(p,q)} \int \|\mathbf{x} - \mathbf{y}\|_2^2 d\pi(\mathbf{x}, \mathbf{y}) \right)^{1/2} \quad (96)$$

1998
1999
2000
2001
2002
2003
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023
2024
2025
2026
2027
2028
2029
2030
2031
2032
2033
2034
2035
2036
2037
2038
2039
2040
2041
2042
2043
2044
2045
2046
2047
2048
2049
2050
2051

Table 8: Synthetic Gaussian mixture distribution parameters for LiDAR experiment. 5000 datapoints are drawn i.i.d. from each of the Gaussian mixtures and paired randomly $(\mathbf{x}_0, \mathbf{x}_{1,0}, \mathbf{x}_{1,1})$ to define the training dataset. A visualization of the training data on the LiDAR manifold is provided in Figure 3.

	Distribution	μ	σ
π_0	$\mathcal{N}(\mu_0, \sigma_0)$	$(-4.5, -4.0, 0.5), (-4.2, -3.5, 0.5), (-4.0, -3.0, 0.5), (-3.75, -2.5, 0.5)$	0.02
$\pi_{1,0}$	$\mathcal{N}(\mu_{1,0}, \sigma_{1,0})$	$(-2.5, -0.25, 0.5), (-2.25, 0.675, 0.5), (-2, 1.5, 0.5)$	0.03
$\pi_{1,1}$	$\mathcal{N}(\mu_{1,1}, \sigma_{1,1})$	$(2, -2, 0.5), (2.6, -1.25, 0.5), (3.2, -0.5, 0.5)$	0.03

where p denotes the ground truth distribution and q denotes the predicted distribution. After training the velocity and growth networks on samples from the initial Gaussian mixture π_0 and target Gaussian mixtures $\pi_{1,0}$ and $\pi_{1,1}$, we evaluate \mathcal{W}_1 and \mathcal{W}_2 of the reconstructed distribution simulated from the validation points in the initial distribution π_0 against the true distribution at $t = 1$.

E.4 DIFFERENTIATING SINGLE-CELL EXPERIMENT DETAILS

Mouse Hematopoiesis scRNA-seq Data We used the mouse hematopoiesis dataset from [Zhang et al. \(2025c\)](#) consisting of three timesteps t_0, t_1, t_2 , with a total of 1429 cells at t_0 , 3781 cells from t_1 , and 5788 cells from t_2 . The data points at t_0 form a homogeneous cluster, while the data points at t_2 are clearly divided into two distinct cell fates. We performed k -means clustering with $k = 2$ clusters to create branching on the cells at t_2 , splitting the cells into two clusters: endpoint 0 with 2902 cells and endpoint 1 with 2886 cells. Since the two endpoints are near equal in ratio, we set the final weights of both endpoints as 0.5 (i.e. $w_{1,0} = w_{1,1} = 0.5$). To match the size of the t_0 samples, we randomly sampled 1429 samples from each of these two clusters and used them as the endpoints for branches 0 and 1, respectively. Training and validation follow a 0.9/0.1 split ratio.

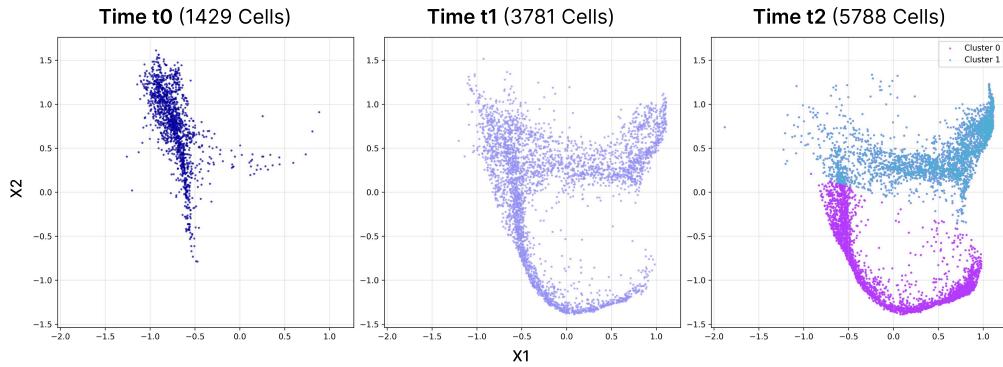


Figure 9: **Mouse Hematopoiesis Single-Cell RNA Sequencing Data Plotted by Time Point.** Real scRNA-seq data is projected to 2D force-directed SPRING plots ([Sha et al., 2023](#); [Weinreb et al., 2020](#); [Zhang et al., 2025c](#)). There is a clear divergence of cell fate between times t_0 (left), t_1 (middle), and t_2 (right) from the initial homogeneous progenitor cells into two distinct cell fates (shown in pink and purple in the t_2 plot). Cells at time t_2 are clustered into endpoint 0 (pink; 2902 cells) and endpoint 1 (turquoise; 2886 cells).

Pancreatic β -Cell Differentiation Data The pancreatic β -cell dataset from [Veres et al. \(2019\)](#) contains 51,274 cells over eight time points that evolve from human pluripotent stem cells to pancreatic β -like cells. Following ([Zhang et al., 2025d](#)), the gene expression representations are projected to 30-dimensional PCA space. After Leiden clustering with $k = 20$ and `min_cells = 100`, we obtain 11 clusters for $K = 11$ branches.

Evaluation Metrics To determine how closely the reconstructed distributions along the trajectory match the ground truth distributions, we compute the 1-Wasserstein (\mathcal{W}_1) ([95](#)) and 2-Wasserstein (\mathcal{W}_2) ([96](#)) distances similar to the LiDAR experiment. After training the velocity and growth networks

2052 on samples from the initial cell distribution π_{t_1} and differentiated target cell distributions $\pi_{t_2,0}$ and
 2053 $\pi_{t_2,1}$, we evaluate \mathcal{W}_1 and \mathcal{W}_2 between the reconstructed branched distributions at the intermediate
 2054 time t_1 ($p_{t_1,0}$ and $p_{t_1,1}$) and the target distributions at the final time t_2 ($p_{t_2,0}$ and $p_{t_2,1}$) simulated
 2055 from the validation points in the initial distribution π_0 and the true distributions π_{t_1} and π_{t_2} .
 2056

2057 **E.5 CELL-STATE PERTURBATION MODELING EXPERIMENT DETAILS**

2059 **Tahoe Single-Cell Perturbation Data** The Tahoe-100M dataset consists of 50 cell lines and over
 2060 1000 different drug-dose conditions (Zhang et al., 2025a). For this experiment, we extract the data for
 2061 a single cell line (A-549) under two drug perturbation conditions selected based on cell abundance
 2062 and response diversity.

2063 Clonidine at 5 μ L was selected first due to having the largest number of cells at this dosage, while
 2064 Trametinib was chosen as the second drug based on its second-highest cell count under the same con-
 2065 dition. For both drugs, we selected the top 2000 highly variable genes (HVGs) based on normalized
 2066 expression and projected the data into a 50-dimensional PCA space, which captured approximately
 2067 38% of the total variance in both cases.

2068 We applied K -nearest neighbor ($K = 50$) and conducted Leiden clustering separately for drugged
 2069 versus DMSO control conditions. The most abundant DMSO cluster was selected as the initial state
 2070 ($t = 0$). For Clonidine, we identified two clusters that were most distinct from the DMSO control
 2071 along PC1 and PC2, respectively. These were selected as two distinct endpoints for branches 1 and
 2072 2 at $t = 1$. We applied centroid-based sampling to obtain balanced training sets of 1033 cells per
 2073 cluster (Figure 10).

2074 For Trametinib, we extended the branching up to three endpoints. From its Leiden clustering
 2075 results, we identified three clusters that were the most divergent from the DMSO control clusters
 2076 along PC1, PC2, and PC3, respectively. All selected clusters contained at least 100 cells and were
 2077 subsampled to 381 cells each for branch training. The remaining cells were clustered with K -means
 2078 into three (Clonidine) or four (Trametinib) groups to construct the metrics dataset. The training and
 2079 validation dataset split followed a 0.9/0.1 ratio. The final visualization utilized the first two principal
 2080 components.

2082 **Table 9: Training cluster cell counts for perturbation experiments.**

	Clonidine		Trametinib		
	Cluster 0	Cluster 1	Cluster 0	Cluster 1	Cluster 2
Original Cell Count	1675	1033	1622	686	381
Initial Weight $w_{0,k}$	1.0	0	1.0	0	0
Target Weight $w_{1,k}$	0.619	0.381	0.603	0.255	0.142

2092 **Evaluation Metrics** To quantify the alignment of the reconstructed and ground-truth distri-
 2093 butions for the cell-state perturbation experiment on principal component (PCs) dimensions
 2094 $d \in \{50, 100, 150\}$, we calculate the Maximum Mean Discrepancy with the RBF kernel (RBF-
 2095 MMD) on all predicted PCs and the 1-Wasserstein (\mathcal{W}_1) (95) and 2-Wasserstein (\mathcal{W}_2) (96) distances
 2096 on the top-2 PCs.

2097 This choice is made because the Wasserstein distance between empirical measures (i.e., discrete point
 2098 clouds) is known to converge to the true Wasserstein distance between the underlying distributions at
 2099 a rate that scales as $\varepsilon = \mathcal{O}(n^{-1/d})$ for some constant C , where n is the number of samples and d is
 2100 the dimensionality of the state space. Therefore, an exponential number of samples is required to
 2101 achieve the same estimation accuracy in high dimensions.

2103 Computing the Wasserstein distance on the top-2 principal components provides a statistically stable
 2104 and computationally tractable approximation that preserves the major axes of variation in the data. In
 2105 addition, we evaluate the RBF-MMD metric on *all simulated PCs*, ensuring that higher-dimensional
 reconstruction accuracy is still quantitatively assessed.

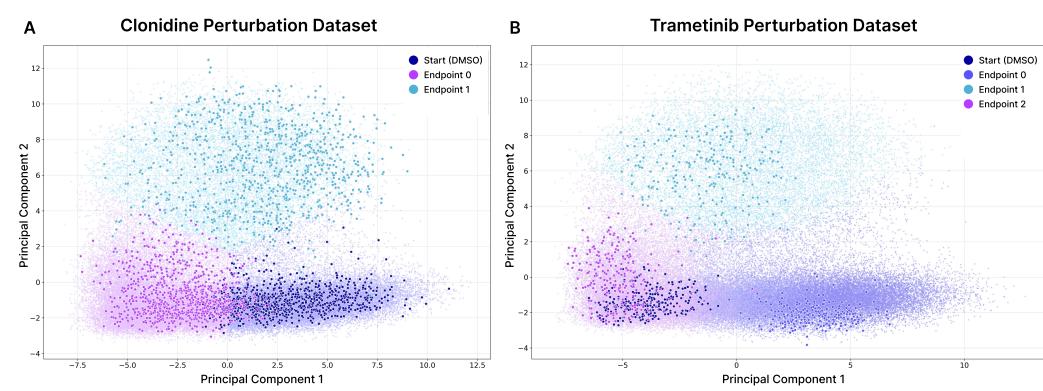


Figure 10: **Clustered Cell-State Perturbation Data from the Tahoe-100M Dataset.** PCA was conducted on all cells for the control DMSO-treated and two drug-treated populations, and clustered. Plots show divergence along the top 2 PCs. **(A)** Clonidine-treated cells ($5\mu M$) are plotted in pink (endpoint 0), and turquoise (endpoint 1), and the distribution of control cells is plotted in navy. **(B)** Trametinib-treated cells ($5\mu M$) are plotted in purple (endpoint 0), turquoise (endpoint 1), and pink (endpoint 2), and the distribution of control cells is plotted in navy.

Given the predicted distribution p and true distribution q and n samples from each distribution $\{\mathbf{x}_i \sim p\}_{i=1}^n$ and $\{\mathbf{y}_i \sim q\}_{i=1}^n$, the RBF-MMD between p and q is calculated as

$$\text{MMD}(p, q) = \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n k_{\text{mix}}(\mathbf{x}_i, \mathbf{x}_j) + \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n k_{\text{mix}}(\mathbf{y}_i, \mathbf{y}_j) - \frac{2}{n^2} \sum_{i=1}^n \sum_{j=1}^n k_{\text{mix}}(\mathbf{x}_i, \mathbf{y}_j) \quad (97)$$

where $k_{\text{mix}}(\cdot, \cdot)$ is a mixture of RBF kernel functions defined as

$$k_{\text{mix}}(\mathbf{x}, \mathbf{y}) = \frac{1}{|\Sigma|} \sum_{\sigma \in \Sigma} \exp\left(-\frac{\|\mathbf{x} - \mathbf{y}\|^2}{2\sigma^2}\right) \quad (98)$$

where $\Sigma = \{0.01, 0.1, 1, 10, 100\}$ is the set of values that determine how much the distances between pairs of points are scaled when computing the overall discrepancy. The equations for 1-Wasserstein (\mathcal{W}_1) and 2-Wasserstein (\mathcal{W}_2) distances are provided in (95) and (96) respectively.

E.6 COMPUTATIONAL COSTS

For each experiment, we report the GPU hours on a single NVIDIA A100 GPU, demonstrating that only a minor increase in compute time is required for increasing the number of branches.

Table 10: **Total training time for each experiment on a NVIDIA A100 GPU for different dimensions and number of branches K .**

Experiment (Branches)	Total Training Time (min)
Mouse Hematopoiesis ($K = 1$)	13m 4s
Mouse Hematopoiesis ($K = 2$)	14m 3s
Clonidine 50D ($K = 1$)	17m 22s
Clonidine 50D ($K = 2$)	22m 50s
Clonidine 100D ($K = 2$)	23m 3s
Clonidine 150D ($K = 2$)	18m 31s
Trametinib ($K = 1$)	4m 52s
Trametinib ($K = 3$)	6m 44s

2160 E.7 HYPERPARAMETER SELECTION AND DISCUSSION
2161

2162 In this section, we present the hyperparameters used in each experiment. While the model architecture
2163 remained largely the same across experiments, we increased the hidden dimension to 1024 for di-
2164 mensions $d \in \{50, 100, 150\}$. For low-dimensional data, we found that increasing model complexity
2165 underperforms in comparison to lower hidden dimensions, and we established that a hidden dimension
2166 of 64 achieves relatively optimal performance for $d \in \{2, 3\}$. While beyond the scope of this study,
2167 we believe that further exploration of diverse model architectures and hyperparameter tuning could
2168 improve the performance of BranchSBM. Exploration of diverse task-dependent state costs for novel
2169 applications is another exciting extension of our work.

2170 Table 11: **Hyperparameter settings for different datasets.** The Clonidine perturbation experiment is split into
2171 three columns for each of the three dimensions of principal components (PCs) used $d \in \{50, 100, 150\}$.
2172

2173 Parameter	2174 Dataset					
	2175 LiDAR	2176 Mouse Hematopoiesis scRNA	2177 Clonidine Perturbation	2178 50PCs	2179 100PCs	2180 150PCs
2181 branches	2	2	2	2	2	3
2182 data dimension	3	2	50	100	150	50
2183 batch size	128	128	32	32	32	32
2184 λ_{energy}	1.0	1.0	1.0	1.0	1.0	1.0
2185 λ_{mass}	100	100	100	100	100	100
2186 λ_{match}	1.0×10^3	1.0×10^3	1.0×10^3	1.0×10^3	1.0×10^3	1.0×10^3
2187 λ_{recons}	1.0	1.0	1.0	1.0	1.0	1.0
2188 λ_{growth}	0.01	0.01	0.01	0.01	0.01	0.01
2189 V_t	LAND	LAND	RBF	RBF	RBF	RBF
2190 RBF N_c	-	-	150	300	300	150
2191 RBF κ	-	-	1.5	2.0	3.0	1.5
2192 hidden dimension	64	64	1024	1024	1024	1024
2193 lr $\varphi_{t,\eta}$	1.0×10^{-4}	1.0×10^{-4}	1.0×10^{-4}	1.0×10^{-4}	1.0×10^{-4}	1.0×10^{-4}
2194 lr u_t^θ	1.0×10^{-3}	1.0×10^{-3}	1.0×10^{-3}	1.0×10^{-3}	1.0×10^{-3}	1.0×10^{-3}
2195 lr g_t^ϕ	1.0×10^{-3}	1.0×10^{-3}	1.0×10^{-3}	1.0×10^{-3}	1.0×10^{-3}	1.0×10^{-3}

2188
2189
2190
2191
2192
2193
2194
2195
2196
2197
2198
2199
2200
2201
2202
2203
2204
2205
2206
2207
2208
2209
2210
2211
2212
2213

2214 F TRAINING ALGORITHM
22152216 Here, we provide the pseudocode for BranchSBM’s multi-stage training algorithm for stable
2217 optimization of the velocity and growth networks over the K branched trajectories.
22182219 **Algorithm 1** Multi-Stage Training of **BranchSBM**
2220

```

2221 1: Stage 1: Learning the Branched Neural Interpolants
2222 2: while Training do
2223 3:    $\forall k, (\mathbf{x}_0, \mathbf{x}_{1,k}) \sim \pi_{0,1,k}^*, t \sim \mathcal{U}(0, 1)$ 
2224 4:   for  $k = 0$  to  $K$  do
2225 5:      $\mathbf{x}_{t,\eta,k} \leftarrow (1-t)\mathbf{x}_0 + t\mathbf{x}_{1,k} + t(1-t)\varphi_{t,\eta}(\mathbf{x}_0, \mathbf{x}_{1,k})$ 
2226 6:      $\dot{\mathbf{x}}_{t,\eta,k} \leftarrow \mathbf{x}_1 - \mathbf{x}_0 + t(1-t)\dot{\varphi}_{t,\eta}(\mathbf{x}_0, \mathbf{x}_{1,k}) + (1-2t)\varphi_{t,\eta}(\mathbf{x}_0, \mathbf{x}_{1,k})$ 
2227 7:     Compute  $V_t(\mathbf{x}_{t,\eta,k})$  given the task-specific definition
2228 8:      $\mathcal{L}_{\text{traj}}(\eta) \leftarrow \int_0^1 \left[ \frac{1}{2} \|\dot{\mathbf{x}}_{t,\eta,k}\|^2 + V_t(\mathbf{x}_{t,\eta,k}) \right] dt$ 
2229 9:     Update  $\varphi_{t,\eta}$  using gradient  $\nabla_\eta \mathcal{L}_{\text{traj}}(\eta)$ 
2230 10:   end for
2231 11: end while
2232 12: Stage 2: Initial Training of Velocity Networks
2233 13: while Training do
2234 14:   Initialize  $K + 1$  flow networks  $\{u_{t,k}^\theta(X_t)\}_{k=0}^K$ 
2235 15:   for  $k = 0$  to  $K$  do
2236 16:     Calculate  $\mathbf{x}_{t,\eta,k}$  and  $\dot{\mathbf{x}}_{t,\eta,k}$  with the trained network  $\varphi_{t,\eta}^*(\mathbf{x}_0, \mathbf{x}_{1,k})$ 
2237 17:      $\mathcal{L}_{\text{flow}}(\theta) \leftarrow \|\dot{\mathbf{x}}_{t,\eta,k} - u_{t,k}^\theta(\mathbf{x}_{t,\eta,k})\|_2^2$ 
2238 18:     Update  $u_{t,k}^\theta$  using gradient  $\nabla_\theta \mathcal{L}_{\text{flow}}(\theta)$ 
2239 19:   end for
2240 20: end while
2241 21: Stage 3: Initial Training of Growth Networks
2242 22: while Training do
2243 23:   Freeze parameters of flow networks and initialize  $K + 1$  growth networks  $\{g_{t,k}^\phi(X_t)\}_{k=0}^K$ 
2244 24:   for  $t = 0$  to  $1$  do
2245 25:     for  $k = 0$  to  $K$  do
2246 26:        $\mathbf{x}_{t,k} \leftarrow \int_0^t u_{s,k}^\theta(\mathbf{x}_{s,k}) ds$ 
2247 27:       if  $k = 0$  then
2248 28:          $w_{t,k}^\phi \leftarrow 1 + \int_0^t g_{t,k}^\phi(\mathbf{x}_{s,k}) ds$ 
2249 29:       else
2250 30:          $w_{t,k}^\phi \leftarrow \int_0^t g_{t,k}^\phi(\mathbf{x}_{s,k}) ds$ 
2251 31:       end if
2252 32:        $\mathcal{L}_{\text{energy}}(\phi) \leftarrow \mathcal{L}_{\text{energy}}(\phi) + \int_t^{t+\Delta t} \left[ \frac{1}{2} \|u_{t,k}^\theta\|^2 + V_t(\mathbf{x}_{t,k}) \right] w_{t,k}^\phi$ 
2253 33:     end for
2254 34:      $\mathcal{L}_{\text{mass}} \leftarrow \left( \sum_{k=0}^K w_{t,k}^\phi - w^{\text{total}} \right)^2$ 
2255 35:   end for
2256 36:    $\mathcal{L}_{\text{match}} \leftarrow \sum_{k=0}^K \left( w_{1,k}^\phi(\mathbf{x}_{1,k}) - w_{1,k}^* \right)^2$ 
2257 37:    $\mathcal{L}_{\text{recons}}(\theta) \leftarrow \sum_{k=0}^K \sum_{\mathbf{x}_{1,k} \in \mathcal{N}_n(\mathbf{x}_{1,k})} \max(0, \|\tilde{\mathbf{x}}_{1,k} - \mathbf{x}_{1,k}\|_2 - \epsilon)$ 
2258 38:    $\mathcal{L}_{\text{growth}}(\phi) \leftarrow \lambda_{\text{energy}} \mathcal{L}_{\text{energy}}(\theta, \phi) + \lambda_{\text{match}} \mathcal{L}_{\text{match}}(\phi) + \lambda_{\text{mass}} \mathcal{L}_{\text{mass}}(\phi) + \lambda_{\text{growth}} \sum_{k=0}^K \|g_{t,k}^\phi\|^2$ 
2259 39:   Update  $g_{t,k}^\phi$  using gradient  $\nabla_\phi \mathcal{L}_{\text{growth}}(\phi)$ 
2260 40: end while
2261 41: Stage 4: Final Joint Training
2262 42: while Training do
2263 43:   Unfreeze parameters of flow networks  $\{u_{t,k}^\theta(X_t)\}_{k=0}^K$ 
2264 44:   Repeat steps of Stage 3 and calculate  $\mathcal{L}_{\text{joint}}(\theta, \phi) \leftarrow \mathcal{L}_{\text{growth}}(\theta, \phi) + \mathcal{L}_{\text{recons}}(\theta)$ 
2265 45:   Jointly update  $u_{t,k}^\theta$  and  $g_{t,k}^\phi$  for all branches using gradients  $\nabla_\theta \mathcal{L}_{\text{joint}}(\theta, \phi)$  and  $\nabla_\phi \mathcal{L}_{\text{joint}}(\theta, \phi)$ 
2266 46: end while
2267

```

2268 **G USE OF LARGE LANGUAGE MODELS (LLMs)**
22692270 We acknowledge the use of large language models (LLMs) to assist in polishing and editing parts of
2271 this manuscript. LLMs were used to refine phrasing, improve clarity, and ensure consistency of style
2272 across sections. All technical content, experiments, analyses, and conclusions were developed by the
2273 authors, with LLM support limited to language refinement and editorial improvements.
2274

2275

2276

2277

2278

2279

2280

2281

2282

2283

2284

2285

2286

2287

2288

2289

2290

2291

2292

2293

2294

2295

2296

2297

2298

2299

2300

2301

2302

2303

2304

2305

2306

2307

2308

2309

2310

2311

2312

2313

2314

2315

2316

2317

2318

2319

2320

2321