

FEEDBACK SCHRÖDINGER BRIDGE MATCHING

Anonymous authors

Paper under double-blind review

ABSTRACT

Recent advancements in diffusion bridges for distribution transport problems have heavily relied on matching frameworks, yet existing methods often face a trade-off between scalability and access to optimal pairings during training. Fully unsupervised methods make minimal assumptions but incur high computational costs, limiting their practicality. On the other hand, imposing full supervision of the matching process with optimal pairings improves scalability, however, it can be infeasible in many applications. To strike a balance between scalability and minimal supervision, we introduce **Feedback Schrödinger Bridge Matching (FSBM)**, a novel *semi-supervised* matching framework that incorporates a small portion (less than 8% of the entire dataset) of pre-aligned pairs as *state feedback* to guide the transport map of non-coupled samples, thereby significantly improving efficiency. This is achieved by formulating a static Entropic Optimal Transport (EOT) problem with an additional term capturing the semi-supervised guidance. The generalized EOT objective is then recast into a dynamic formulation to leverage the scalability of matching frameworks. Extensive experiments demonstrate that FSBM accelerates training and enhances generalization by leveraging coupled pairs’ guidance, opening new avenues for training matching frameworks with partially aligned datasets.

1 INTRODUCTION

Transporting samples between distributions is a ubiquitous problem in machine learning. Given the rise of generative modeling, significant progress has been made using static (Goodfellow et al., 2014), deterministic (Chen et al., 2018; Biloš et al., 2021), or stochastic mappings (Ho et al., 2020; Song et al., 2020) that transport samples from noise to a more complex distribution. A prominent approach lies in diffusion models that simulate a Stochastic Differential Equation (SDE) to diffuse the data to noise, and learn the score function (Hyvärinen & Dayan, 2005; Vincent, 2011) to reverse the process (Anderson, 1982). However, these models present several limitations. For instance, the requirement to converge during the noising (forward) process to a Gaussian noise suggests that these models must run for a sufficient number of time steps to ensure the final distribution approximates Gaussian noise (Chen et al., 2021). Additionally, this suggests that these models begin their generative (backward) processes without any structural information about the data distribution, which implies randomness being introduced in the couplings that emerge from the diffusion models (Liu et al., 2023). Lastly, there is no guarantee that the optimal path interpolating the boundary distribution minimizes the kinetic energy. (Shi et al., 2023)

In an attempt to overcome these shortcomings, principled approaches that stem from Optimal Transport (Villani et al., 2009) have emerged. The most prominent alternative has been the Schrödinger Bridge (SB; Schrödinger (1931)), which has been shown to be equivalent to entropy-regularized Optimal Transport (EOT; Cuturi (2013); Léonard (2013); Pavon et al. (2021)), and can also be framed as a Stochastic Optimal Control (SOC) Problem (Chen et al., 2016; 2021). In particular, SB gained significant popularity in the realm of generative modeling following advancements proposing a training scheme based on the Iterative Proportional Fitting (IPF), a continuous state space extension of the Sinkhorn algorithm to solve the dynamic SB problem (De Bortoli et al., 2021; Vargas et al., 2021). Notably, SB generalizes standard diffusion models transporting data between arbitrary distributions π_0, π_1 with fully nonlinear stochastic processes, seeking the unique path measure that minimizes the kinetic energy. More recently, building on advancements in Bridge Matching methods (Peluchetti, 2023; Liu et al., 2022c), Shi et al. (2023) introduced Diffusion Schrödinger Bridge Matching (DSBM),

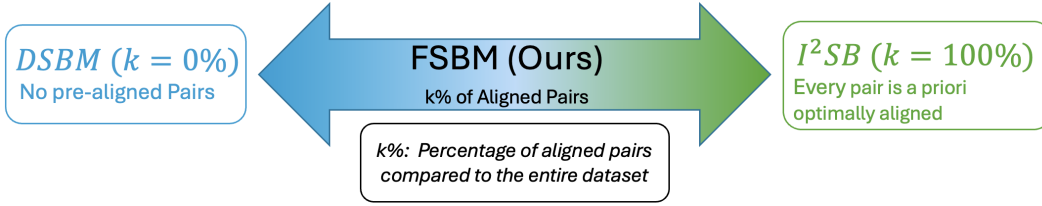


Figure 1: FSBM connecting existing Bridge Matching frameworks at the extremes, where the dataset is comprised of either fully aligned or fully non-aligned pairs

a framework that solves the Schrödinger Bridge problem while being significantly more efficient. Unlike previous methods, DSBM avoids the need to cache the full trajectories of forward and backward SDEs, making it more scalable and mitigating the time-discretization and "forgetting" issues encountered in earlier DSB techniques.

For training these matching frameworks, we start from two datasets that lack pre-aligned pairs—that is, we sample only from the marginal π_0, π_1 —and the framework is trained to determine the optimal coupling in the sense of least-energy transportation. This process represents an unsupervised learning approach, as no prior information about the coupling is provided. However, there are other applications—such as image restoration (Liu et al., 2023) and protein docking (Somnath et al., 2023)—in which we have prior knowledge of the coupling $\pi_{0,1}$, between pairs $(x_0, x_1) \sim \pi_{0,1}$ drawn presumably from the SB solutions. Consequently, using an unsupervised approach for such problems is not ideal, as it discards valuable information about the relationships between the samples of the two boundary distributions. An alternative approach suggests reframing these tasks as inverse problems. In this vein, Liu et al. (2023) proposed recently I^2SB , building diffusion bridge matching frameworks between coupled data to perform image restoration.

In practice, in most tasks, the available datasets might possess a few limited coupled pairs, however, the high cost of manually labeling large datasets renders fully supervised approaches infeasible. Unfortunately, existing matching frameworks can not effectively leverage the information in partially pre-aligned datasets. An alternative approach from Optimal Transport (OT) literature involves using guided transportation maps, where only a few pairs that belong to a Key-Point (KP) set are annotated. These pairs are utilized to guide the transportation mapping of the unpaired samples (Gu et al., 2023a), substantially reducing the need for extensive human labeling or expert guidance by leveraging a small number of source-target aligned sample pairs for training (Mustafa & Mantiuk, 2020). More specifically, recent approaches employ a “relation-preserving” scheme, which maintains the data’s relationship to the given pairs from the KP set (Mémoli, 2011; Sato et al., 2020), or a pairwise distance-preserving constraint (Gu et al., 2023b). However, the feasibility of adapting semi-supervised guidance in a dynamic Schrödinger Bridge or Bridge Matching setting remains an open question. In this vein, drawing inspiration from the guided OT schemes, we advocate a semi-supervised guided Schrödinger Bridge Matching framework.

In this work, we introduce a novel semi-supervised matching algorithm, **Feedback Schrödinger Bridge Matching (FSBM)**, designed to integrate information from partially aligned datasets. Drawing inspiration from optimal transport (OT) literature (Gu et al., 2022), our analysis begins from a static, semi-supervised OT problem, from which we derive a dynamic objective. Following recent advancements in matching frameworks, we adopt an alternating scheme, where the intermediate path and the coupling are optimized in two separate steps, resulting in a novel matching framework that leverages partially aligned data to guide the transport mapping of non-aligned samples. A key aspect of our approach is that the information from aligned samples is encoded as **state feedback** within the dynamic objective, effectively steering the transport of non-aligned data. This renders our FSBM a bridge between two extremes: unsupervised matching frameworks, which lack pre-aligned couplings (Shi et al., 2023), and fully supervised frameworks, where data is entirely pre-aligned (Liu et al., 2023; Somnath et al., 2023) (see Figure 1). Empirical results show that our algorithm generalizes better, is more robust to perturbations in the initial conditions, and exhibits reduced training time. Our contributions are summarized as follows:

- We introduce FSBM, the first matching framework leveraging information from partially aligned datasets.

- We begin our analysis from a static semi-supervised OT problem, and derive a modified dynamic formulation. Notably, our analysis can be extended to any selection of regularization functions.
- We introduce an Entropic Lagrangian extension of the variational gap objective used in BM frameworks to match the parameterized drift given a prescribed probability path
- Through extensive experimentation, we verify the remarkable capability of FSBM to generalize under a variety of unseen initial conditions, while simultaneously achieving faster training times. These results were consistent across a wide range of matching tasks from low-dimension crowd navigation, to high-dimensional opinion depolarization and image translation.

2 PRELIMINARIES

2.1 SCHRÖDINGER BRIDGES

The Schrödinger Bridge (Schrödinger, 1931) in the path measure sense is concerned with finding the optimal measure \mathbb{P}^* that minimizes the following optimization problem

$$\min_{\mathbb{P}} KL(\mathbb{P}|\mathbb{Q}), \quad \mathbb{P}_0 = \pi_0, \mathbb{P}_1 = \pi_1 \quad (1)$$

where \mathbb{Q} is a Markovian reference measure. Hence the solution of the dynamic SB \mathbb{P}^* is considered to be the closest path measure to \mathbb{Q} . A very convenient, and heavily used property of \mathbb{P}^* is its decomposition as $P^* = \int_{\mathbb{R}^d \times \mathbb{R}^d} \mathbb{Q}_{|0,1} d\pi^*(x_0, x_1)$, where π^* is the solution of the static SB $\min_{\pi \in \Pi(\pi_0, \pi_1)} KL(\pi|R)$, and $\mathbb{Q}_{|0,1}$ is the Bridge of the reference measure for pinned points (x_0, x_1) (Léonard, 2013). Another formulation of the dynamic SB crucially emerges by applying the Girsanov theorem in Eq. (1)

$$\min_{u_t, p_t} \int_0^1 \mathbb{E}_{p_t} [\|u_t\|^2] dt \quad \text{s.t.} \quad \frac{\partial p_t}{\partial t} = -\nabla \cdot (u_t p_t) + \frac{\sigma^2}{2} \Delta p_t, \quad \text{and} \quad p_0 = \pi_0, \quad p_1 = \pi_1 \quad (2)$$

Finally, note that the static SB is equivalent to the entropy regularized OT formulation (Pavon et al., 2021; Nutz, 2021).

$$\min_{\pi \in \Pi(\pi_0, \pi_1)} \int_{\mathbb{R}^d \times \mathbb{R}^d} \|X_0 - X_1\|^2 d\pi(X_0, X_1) + \epsilon KL(\pi|\pi_0 \otimes \pi_1) \quad (3)$$

This regularization term enabled efficient solution through the Sinkhorn algorithm and has presented numerous benefits, such as smoothness, and other statistical properties (Ghosal et al., 2022; Léger, 2021; Peyré et al., 2019).

2.2 SCHRÖDINGER BRIDGE MATCHING

As discussed above, SB Matching algorithms (Shi et al., 2023; Liu et al., 2024) have emerged recently, following advancements in matching frameworks (Liu et al., 2022c; Peluchetti, 2023; Lipman et al., 2022; Liu et al., 2022b), rendering the solution of the SB problem significantly more tractable. Specifically, this family of algorithms solves Eq. (2) by separating the training process into two alternating steps. The first step entails relaxing the boundary distributional constraints in Eq. (2) into just two endpoints, by fixing the coupling $\pi_{0,1}$, and drawing pairs of samples (x_0, x_1) .

$$\begin{aligned} & \min_{u_t|_{0,1}} \int_0^1 \mathbb{E}_{p_t|_{0,1}} [\|u_t|_{0,1}\|^2] dt \\ & \text{s.t.} \quad \frac{\partial p_t|_{0,1}}{\partial t} = -\nabla \cdot (p_t|_{0,1} u_t) + \frac{1}{2} \sigma^2 \Delta p_t|_{0,1}, \quad X_0 = x_0, \quad X_1 = x_1 \end{aligned} \quad (4)$$

This optimization problem returns the optimized intermediate bridges between the drawn pairs. Subsequently, the parameterized drift u_t^θ is matched given the prescribed marginal path from the previous step, by minimizing the following variational gap

$$\min_{\theta} \int_0^1 \mathbb{E}_{p_t} [\|u_t^* - u_t^\theta\|^2] dt \quad (5)$$

where u_t^* is the drift associated with the fixed intermediate path. Propagation of the SDE according to the matched parameterized drift $u_t^\theta : \mathbb{R}^d \times [0, 1] \rightarrow \mathbb{R}$

$$dX_t = u_t^\theta dt + \sigma dW_t, \quad X_0 \sim \pi_0, \quad X_1 \sim \pi_1 \quad (6)$$

induces the parameterized coupling $\pi_{0,1}^\theta$. In the next iteration of the BM algorithm, this updated coupling is used to generate improved pairs of $(x_0, x_1) \sim \pi_{0,1}^\theta$ employed to refine the marginal path.

3 METHODOLOGY

3.1 SEMI-SUPERVISED GUIDANCE

In the context of semi-supervised optimal transport, certain aligned couples of data points from the source and target distributions are employed to guide the matching of unpaired samples. These pairs are assumed by construction to be given by the optimal solution of the static SB. We denote with $\mathcal{K} = \{x_0^n, x_1^n\}_{n=1}^N$ the set with the aligned pairs, the Key-Point (KP) set, with N being the number of pairs in the set. Usually, $N \ll M$, where M is the total number of samples.

Assumption 3.1. All pairs from the KP set are considered to be solutions of the static Schrödinger Bridge.

Our analysis begins with the static entropic OT problem in Eq. (3), to which an additional regularization term is included to capture the interaction among unpaired samples and aligned samples, thereby guiding the transport map. We consider a distance-preserving scheme. We start by sampling a pair of non-coupled samples (X_0, X_1) . Let us assume that the i^{th} sample from the KP set $x_0^i \in \mathcal{K}$ is closest to X_0 in the source distribution. Then, the distance $d(X_0, x_0^i)$ is computed and we proceed to compute the distance $d(X_1, x_1^i)$. Therefore, for a fixed couple $(x_0^i, x_1^i) \in \mathcal{K}$, we define the Guidance function for the non-aligned pair (X_0, X_1) , and as

$$G(X_0, X_1) = \alpha \cdot (d(X_1, x_1^i) - d(X_0, x_0^i))^2 \quad (7)$$

where $\alpha \in \mathbb{R}$ is the guidance regularization hyperparameter. Intuitively, this regularization term creates a region of attraction, as illustrated in Figure 2. This attraction around each KP sample defines clusters that guide the neighboring samples from the source distribution to the target distribution, by preserving the distance to the respective KP sample. Adjusting the cost function of the entropic O.T. formulation in 3, to incorporate guidance from the coupled pairs yields

$$\min_{\pi \in \Pi(\pi_0, \pi_1)} \int (||X_0 - X_1||^2 + G(X_0, X_1)) d\pi(x_0, x_1) + \epsilon KL(\pi | \pi_0 \otimes \pi_1) \quad (8)$$

In this work, we start our analysis from the modified expression of Eq. (3) to exploit the flexibility of static OT formulation into adding the regularization term that is responsible for the guidance of the non-paired samples. However, our focal point is to derive an equivalent dynamic formulation, resembling Eq. (2), which provides a more thorough description of the transport process. From an algorithmic standpoint, our methodology is based on BM frameworks (see Section 2.2) in order to leverage their great efficiency and scalability in solving distribution matching problems.

3.2 FEEDBACK SCHRÖDINGER BRIDGE MATCHING

3.2.1 DYNAMIC OBJECTIVE

The theorem below provides a general framework to recast the entropy regularized semi-supervised O.T. problem in Eq. (8) in a dynamic setting, regardless of the selection for the guidance function G .

Theorem 3.2. Assume $X_0, X_1 \in \mathbb{R}^d$, with $X_0 \sim \pi_0$, and $X_1 \sim \pi_1$, and consider a stochastic random variable X_t connecting X_0 and X_1 , whose law is the continuous marginal probability path

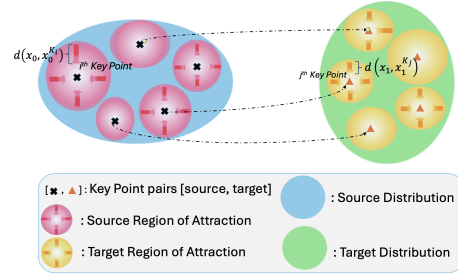


Figure 2: Source and Target Distributions with their KP set pairs

p_t joining π_0 , and π_1 . Starting from the static semi-supervised guided entropic OT problem in Eq. (8), we derive the following relaxed dynamic formulation.

$$\begin{aligned} \min_{p_t, u_t} \int_{\mathbb{R}^d \times [0,1]} & \left(\frac{1}{2} \|u_t\|^2 + \mathbb{E}_{p_{0|t}} [|u_t^\top \nabla G_{t,0}|] + \frac{\sigma^2}{2} \Delta G_{t,0} \right) p_t dx dt, \\ \text{s.t. } \frac{\partial p_t}{\partial t} &= -\nabla \cdot (u_t p_t) + \frac{\sigma^2}{2} \Delta p_t, \text{ and } p_0 = \pi_0, p_1 = \pi_1 \end{aligned} \quad (9)$$

Notice that our dynamic formulation is a modified and generalized version of Eq. (2) to capture the guidance from the KP samples. For degenerate guidance term ($G \equiv 0$), one readily obtains the standard entropic analogue of the fluid dynamic formulation (Benamou & Brenier, 2000; Gentil et al., 2017) in Eq. (2) which also corresponds to the objective of SB objective (De Bortoli et al., 2021; Shi et al., 2023).

From an algorithmic standpoint, we adopt recent advancements in Matching frameworks (Liu et al., 2024; Shi et al., 2023), and separate the training phase into two stages: first optimizing the intermediate path, and subsequently the parameterized drift.

3.2.2 INTERMEDIATE PATH OPTIMIZATION

As discussed above, the initial step of our methodology entails the optimization of the intermediate path. Let the marginal be expressed as a mixture of conditional probability paths $p_t = \int p_{t|0,1} \pi_{0,1} dx_0 dx_1$. More explicitly, we sample non-coupled pairs (x_0, x_1) and fix the coupling at the boundaries $\pi_{0,1}$, which is equivalent to fixing the drift u_t , that induces the coupling $\pi_{0,1}$, through Eq. (6). This results in relaxing the distributional constraints in Eq. (9) to just two endpoints. Consequently, Eq. (9) is modified as

Proposition 3.3. *Let the continuous marginal path p_t satisfy the following decomposition $p_t = \int p_{t|0,1} \pi_{0,1} dx_0 dx_1$. For optimizing with respect to the intermediate path, the coupling $\pi_{0,1}$ is frozen which results in Eq. (9) being recasted as*

$$\begin{aligned} \min_{p_{t|0,1}} \int_{\mathbb{R}^d \times [0,1]} & \left(\frac{1}{2} \|u_{t|0,1}\|^2 + |u_{t|0,1}^\top \nabla G(X_t, x_0)| + \frac{\sigma^2}{2} \Delta G(X_t, x_0) \right) p_{t|0,1} dx dt, \\ \text{s.t. } \frac{\partial p_{t|0,1}}{\partial t} &= -\nabla \cdot (p_{t|0,1} u_t) + \frac{1}{2} \sigma^2 \Delta p_{t|0,1}, \quad X_0 = x_0, \quad X_1 = x_1 \end{aligned} \quad (10)$$

Remark 3.4. Notice that in the extreme case in which the entire dataset is comprised of only pre-aligned couples, G would be trivial, and Eq. (10) would retrieve the diffusion bridges between optimal pairs, similarly to I^2SB (Liu et al., 2023). Conversely, in the other extreme case, where the KP set is the empty set, our algorithm retrieves DSBM.

Feedback The last two terms in Eq. (9) stem from the semi-supervised regularization. Interestingly, in stochastic optimal control literature ((Theodorou et al., 2010)) it has been reported that the inner product of the drift with the state cost acts as state feedback. Hence, the term $|\nabla G_{t|0}^\top u_{t|0,1}|$ is interpreted as state-feedback that steers the non-coupled samples, by projecting the conditioned drift on the gradient of $G_{t|0}$, as illustrated in Figure 3. Note that the trajectories of the KP samples along with their coupling are assumed to be given and fixed for each KP pair. Furthermore, the absolute value implies that the angle of projection is always acute, which greatly improves stability during training. The conditioning on the unpaired sample x_0 acts as an anchor, forcing unpaired samples to preserve their relative distance from the closest KP sample in the source distribution at each time-step. Finally, notice that having $u_{t|0,1}$, and G perpendicular does not suffice to minimize the guidance-related terms in Eq. (10), due to the term $\Delta G_{t|0} \geq 0$. Thus, the guidance-related terms are minimized only when $G_{t|0} \rightarrow 0$, indicating that the unpaired samples are aligned with the KP samples.

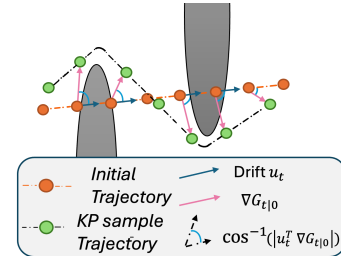


Figure 3: Initial Trajectory of non-aligned samples and KP sample trajectory guiding the transport

Recall that we have relaxed the distributional constraints in Eq. (9) to just two endpoints. We solve Eq. (10) for each pair $(x_0, x_1) \sim p_{0,1}^\theta$ and later marginalize to construct p_t . The conditioned probability path is approximated in a simulation-free manner as a Gaussian path $p_{t|0,1} \approx \mathcal{N}(I_{t|0,1}, \sigma_t \mathbf{I}_d)$, where $I_{t|0,1} \equiv I(t, x_0, x_1)$ is the interpolant function between the pinned endpoints $I_0 = x_0$, $I_1 = x_1$, and σ is the standard deviation with $\sigma_0 = \sigma_1 = 0$. This enables us to compute a closed-form expression for the conditional drift (Särkkä & Solin, 2019; Albergo et al., 2023).

$$u_{t|0,1} = \partial_t I_{t|0,1} + \left(\frac{\partial_t \sigma_t}{\sigma_t} - \frac{\nu^2}{2\sigma_t^2} \right) (X_t - I_t) \quad (11)$$

where ν is given by the definition of $\sigma_t = \nu \sqrt{t(1-t)}$. More details about the gaussian approximation of the interpolating path are left for Appendix B.1.

3.2.3 COUPLING OPTIMIZATION

Subsequently, a variational objective is derived to match the drift u_t^θ , given the prescribed path $p_{t|0,1}$ from the previous step. The inner product in Eq. (10) prompts us to express this variational gap through the Bregman divergence, for which the convex conjugate of Eq. (10) is needed.

Proposition 3.5. *The convex conjugate of the Lagrangian in Eq. (10) is defined as the Hamiltonian $H(x_t, a_t, t) = \sup_{u_{t|0,1}} \langle u_{t|0,1}, \nabla a_t \rangle - L(x_t, u_{t|0,1}, t)$. The optimization with respect to the drift yields $u_{t|0,1} = a_t - g(a_t^\top \nabla G_{t|0}) \nabla G_{t|0}$, where $\nabla G_{t|0} \equiv G(X_t, x_0)$, and $g(\cdot) : \mathbb{R} \rightarrow \mathbb{R}$ is given by*

$$g(a_t^\top \nabla G_{t|0}) = \begin{cases} \frac{a_t^\top \nabla G_{t|0}}{\|\nabla G_{t|0}\|^2} & \text{if } |a_t^\top \nabla G_{t|0}| \leq \|\nabla G_{t|0}\|^2 \\ \text{sgn}(a_t^\top \nabla G_{t|0}) & \text{else} \end{cases}$$

Finally, we obtain the Hamiltonian associated with the Lagrangian in Eq. (10)

$$H(x_t, a_t, t) = \frac{1}{2} \|a_t - g(a_t^\top \nabla G_{t|0}) \cdot \nabla G_{t|0}\|^2 - \frac{\sigma^2}{2} \Delta G(X_t, x_0) \quad (12)$$

Finally, the Hamiltonian in Eq. (36) facilitates us to introduce our Entropic Lagrangian Bridge Matching objective through the Bregman divergence $D_{L,H} = L(X_t, u_t, t) + H(X_t, a_t, t) - \langle a_t, u_t \rangle$, following recent advancements in bridging optimality gaps for general Lagrangian costs Neklyudov et al. (2023).

Proposition 3.6. *Given the optimal conditional drift $u_{t|0,1}^*$, we match the parameterized drift u_t^θ , by minimizing the optimality gap given by our Entropic Lagrangian Bridge Matching objective as*

$$\min_{\theta} \int_0^1 \mathbb{E}_{p_{0,1}} \mathbb{E}_{p_{t|0,1}} \|a_t^\theta - u_{t|0,1}^* - g(a_t^\top \nabla G_{t|0}) \cdot \nabla G_{t|0}\|^2 dt \quad (13)$$

From Proposition 3.5, the parameterized drift can be expressed as $u_t^\theta = a_t^\theta - g(a_t^\top \nabla G_{t|0}) \cdot \nabla G_{t|0}$. Therefore, the matching objective of Eq. (13) can be rewritten as the objective in Eq. (5)

$$\min_{\theta} \int_0^1 \mathbb{E}_{p_t} \|u_{t|0,1}^* - u_t^\theta\|^2 dt \quad (14)$$

3.3 TRAINING SCHEME

A summary of our training procedure is presented in Alg. 1. Our algorithm finds a solution (p_t^*, u_t^*) to Eq. (9) by iteratively optimizing Eq. (10), and 14. First, consider the coupling $\pi_{0,1}$, which we fix and draw pairs (x_0, x_1) . Let the i^{th} data point from the KPs set $x_0^i \in \mathcal{K}$ be closest to x_0 in the source distribution, the distance $d(x_0, x_0^i)$ between the two samples is evaluated. Next, for every sampled time-step t_k , we compute the distance $d(X_{t_k}, X_{t_k}^i)$ between X_{t_k} , and $X_{t_k}^i$, and evaluate $G(X_{t_k}, x_0)$ using these distances. Then, the alternating optimization of the intermediate path and coupling is performed. Keeping the coupling frozen, the optimization of Eq. (10) returns the optimized intermediate path $p_{t|0,1}^*$, which admits a closed form drift $u_{t|0,1}^*$ from Eq. (11). In the next step, the marginal path is fixed so that the coupling is optimized. More specifically, we match the parameterized drift u_t^θ to $u_{t|0,1}^*$ using Eq. (14), given the prescribed marginal p_t . The optimized parameterized drift u_t^θ induces improved coupling $\pi_{0,1}^\theta$ from Eq. (6), which will be used in the next iteration to forward propagate the dynamics.

Algorithm 1 Feedback Schrödinger Bridge Matching (FSBM)

```

1: Input: Boundary distributions  $\pi_0$ , and  $\pi_1$  and the pairs from the key-point (KP) set  $(x_0^i, x_1^i)$ 
2: Initialize  $u_t^\theta$ , and  $\pi_{0,1}$  as independent coupling  $\pi_0 \otimes \pi_1$ , sample trajectories from the KP pairs
3: repeat
4:   Sample  $x_0, x_1, X_{t_k}$  by forward propagating using  $u_t^\theta$ , on time steps  $0 < t_1 < \dots < t_K < 1$ 
5:   Find closest aligned sample  $x_0^i$  to  $x_0$ , and calculate  $d(x_0, x_0^i)$ 
6:   Using the same  $i^{\text{th}}$  data point from the KP set, calculate  $d(X_{t_k}, x_{t_k}^i)$ 
7:    $G(X_{t_k}, x_0)$  from Eq. (7)
8:   Fix the  $x_0, x_1$ , and obtain  $p_{t|0,1}^*$  using Eq. (10)
9:   Calculate  $u_{t|0,1}$  from Eq. (11)
10:  Update  $u_t^\theta$ , from Eq. (14)
11: until converges

```

4 EXPERIMENTS

We demonstrate the efficacy of our FSBM in a variety of distribution matching tasks, such as Crowd Navigation, Opinion Depolarization, and Unpaired Image Translation, compared against other state-of-the-art distribution matching frameworks such as GSBM, DSBM Light and Optimal Schrodinger Bridge Matching (LOSBM). All experiments were run on an NVidia RTX 4090.

4.1 CROWD NAVIGATION

We first showcase the efficacy of our FSBM in solving complex crowd navigation tasks and its superior capability to generalize under a variety of initial conditions, listed as: **i) Vanilla:** initial distribution is the same as the one the model was trained on, **ii) Perturbed Mean:** shift the mean of Vanilla, **iii) Perturbed STD:** shift the standard deviation (STD) of Vanilla, **iv) Uniform Distribution:** the initial distribution is defined as a uniform distribution. We compare our FSBM with the GSBM framework (Liu et al., 2024). Importantly, in the S-tunnel crowd navigation task, 50 keypoints were used to guide a population of 2560 particles, and in the V-neck task, only 20 KP samples were employed for the same number of particles. More details about the selection of the guidance function G , the interpolation of the key-point samples and additional plots are left for Appendix C.1.

Table 1 shows that FSBM generated more accurately the targeted distribution in both tasks. Additionally, Figure 5 depicts our FSBM consistently producing better trajectories, successfully avoiding the obstacles under all initial conditions, and generating final distributions that more closely matched the target distribution. In contrast, under GSBM, several particles were observed to diverge completely in the S-tunnel task under shifted initial conditions, causing the increase of the W^2 distance. Additionally, notice that the final distributions generated by GSBM (see navy blue particles in the upper row of Figure 5), even for the samples that did not diverge, were substantially more inaccurate when compared to our FSBM. Lastly, Fig. 4 illustrates that our FSBM required significantly less wall-clock time for training. This is attributed to the faster convergence of FSBM with regards to the number of

Table 1: W^2 distance between generated and ground truth distributions for GSBM and FSBM in (Lower is better)

	S-tunnel		V-neck	
	GSBM	FSBM	GSBM	FSBM
Vanilla	0.02	0.02	0.01	0.01
Perturbed Mean	8.87e+5	0.67	0.68	0.34
Perturbed STD	2.43e+10	0.34	3.03	0.11
Uniform Distribution	6.65e+9	0.49	0.10	0.10

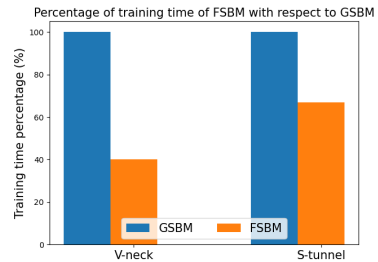


Figure 4: Training time percentage comparison between GSBM and FSBM in Crowd Navigation tasks

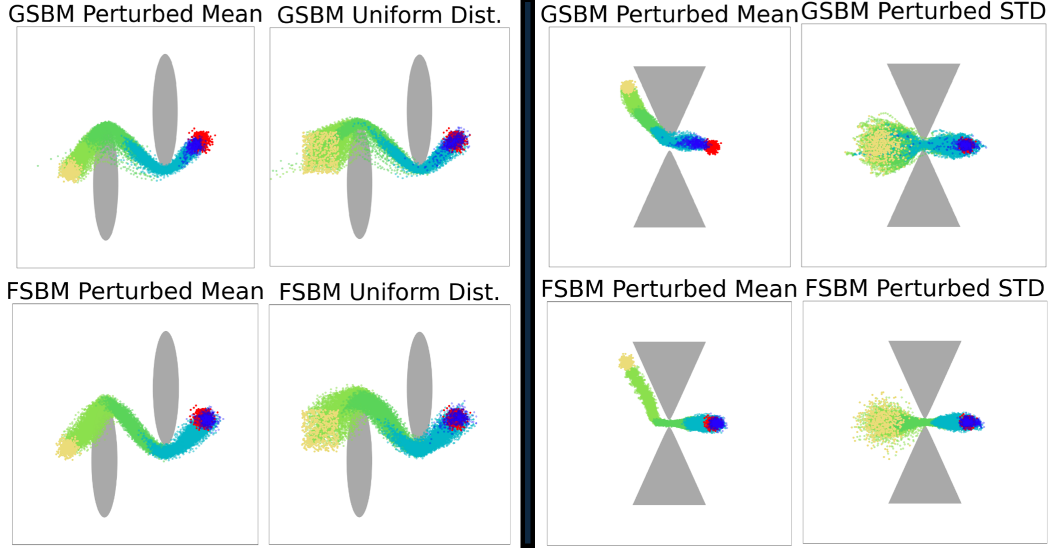


Figure 5: Crowd Navigation comparison between GSBM and FSBM in the S-tunnel, and V-neck under perturbed mean and uniform initial distribution, and in the V-neck under perturbed mean and perturbed STD. The color of particles means: i) **Yellow**: initial conditions, ii) **Green and Cyan**: intermediate trajectory, iii) **Red**: target distribution, iv) **Navy Blue**: generated distribution

epochs, but also due to GSBM requiring the computation of the entropy and congestion cost, which according to our experiments has been shown to increase training duration significantly.

4.2 OPINION DYNAMICS

Subsequently, we consider high-dimensional opinion depolarization, where each particle possesses a high-dimensional opinion $X_t \in \mathbb{R}^{100}$ that evolves through interaction with the other agents under polarizing dynamics

$$dX_t = f_{\text{polarize}}(X_t, t)dt + \sigma dW_t \quad (15)$$

which tend to segregate them into two groups of diametrically opposing opinions (see Figure 6: first row, second column). More information about the polarization base drift are left in Appendix C.2. To mitigate the segregation, we apply our FSBM method employing the keypoint samples to guide the rest of the population to a more uniform opinion, closer to the target distribution (see Figure 6: second row third column). We compare against DeepGSB and GSBM (Liu et al., 2022a; 2024), which employ state cost to mitigate polarizing dynamics. Lastly, it is worth noting that the size of the key-point set utilized for the opinion depolarization was approximately 2%, of the entire population.

Table 2: Comparison of W2 distance and KL divergence between target and generated distribution in High Dimensional ($X \in \mathbb{R}^{100}$) Opinion Dynamics among Deep GSB, GSBM and FSBM (ours) (Lower is better in all metrics)

Optimizer	Time (hours:min)	W2 Distance	KL divergence
DeepGSB	4:37	58.16	0.096
GSBM	5:29	43.41	0.034
FSBM	2:43	42.91	0.034

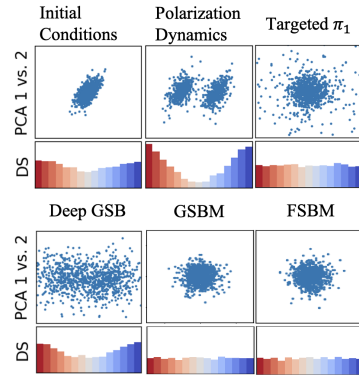


Figure 6: Opinion depolarization from our FSBM and GSBM, and DeepGSB

The results shown demonstrate that the utilization of certain opinion leaders, that guide the overall opinion of the population to depolarization renders our FSBM a highly efficient model to mitigate opinion segregation. The metrics used to evaluate the proximity of the generated and target distributions are the W_2^2 distance and the KL divergence. Additionally, observation of Table 2 and Figure 6 indicates that our FSBM retrieves a solution that is marginally closer to the targeted distribution compared to GSBM, albeit needing half of the training time. Both methods significantly outperform DeepGSB when it comes to their proximity to the target distribution.

4.3 IMAGE TRANSLATION

Finally, we present the results of our FSBM in two types of image translations: **i)** Gender translation, **ii)** Age translation. We follow the setup of (Korotin et al., 2023) with the pre-trained ALAE autoencoder (Pidhorskyi et al., 2020) on the 1024×1024 FFHQ dataset (Karras et al., 2019) to perform the translation in the latent space of dimensions 512×1 , enabling more efficient training and sampling. Notably, the number of aligned images was 4% of the total dataset for the gender translation and 8% for the age translation. Our method is compared with other state-of-the-art matching frameworks, such as DSBM (Shi et al., 2023) and Light and Optimal SBM (LOSBM)(Gushchin et al., 2024). Additional details are left for Appendix C.3.

Table 3: FID values in 4 translation tasks for DSBM, LOSBM, and our FSBM (Lower is better)

Optimizer	Man to Woman	Woman to Man	Old to Young	Young to Old
DSBM	15.90	17.02	16.83	17.76
L-SBM	18.89	17.24	17.89	19.32
FSBM	14.78	16.12	15.52	17.14

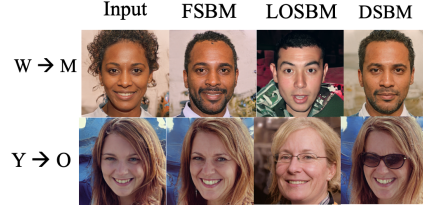


Figure 7: Translation comparison between our FSBM, LOSBM and DSBM

Table 3 shows that our FSBM achieves lower FID value, than the other methods, for 2048 generated images. Additionally, it is demonstrated that our FSBM recovers better couplings through the qualitative comparison in Figure 7. It is worth noting that the images obtained from DSBM and FSBM were very similar, which suggests that the two methods converge to close solutions, while the translations of LOSBM in the validation set were losing some of the features of the input image. Figure 8 presents additional examples of the couplings obtained from our FSBM. We can observe in all instances very accurate translations preserving many features identical, (i.e., smile, accessories, angle of the face, background color, etc) across all translation tasks. Notably, Table 4 demonstrates that FSBM required *only an additional 180 MB of VRAM* compared to DSBM, to manage the aligned latent vectors, which is tractable for most modern GPUS, while training DSBM and FSBM for the same number of epochs requires *virtually identical training duration*.

4.4 DISCUSSIONS

Impact of KP pairs In our methodology, it becomes clear that the number of KP pairs used to guide the rest of the samples affects the effectiveness of our algorithm. As depicted in Figure 9a, although an increase of the aligned pairs results in reduced training time, it is important to note that increasing the number of paired data points does not always lead to improved generalization. There appears to be a "sweet spot", which suggests that using too many KP pairs hinders the ability of our algorithm to generalize effectively akin to an overfitting-like effect. Lastly, it is important to highlight that the other hyperparameters were tuned for a smaller set of keypoints. As a result, increasing the number of aligned pairs may require re-tuning the hyperparameters to maintain optimal performance. Deeper understanding the effect of the aligned data points in the efficacy, and generalizability of the model is a topic for future work.

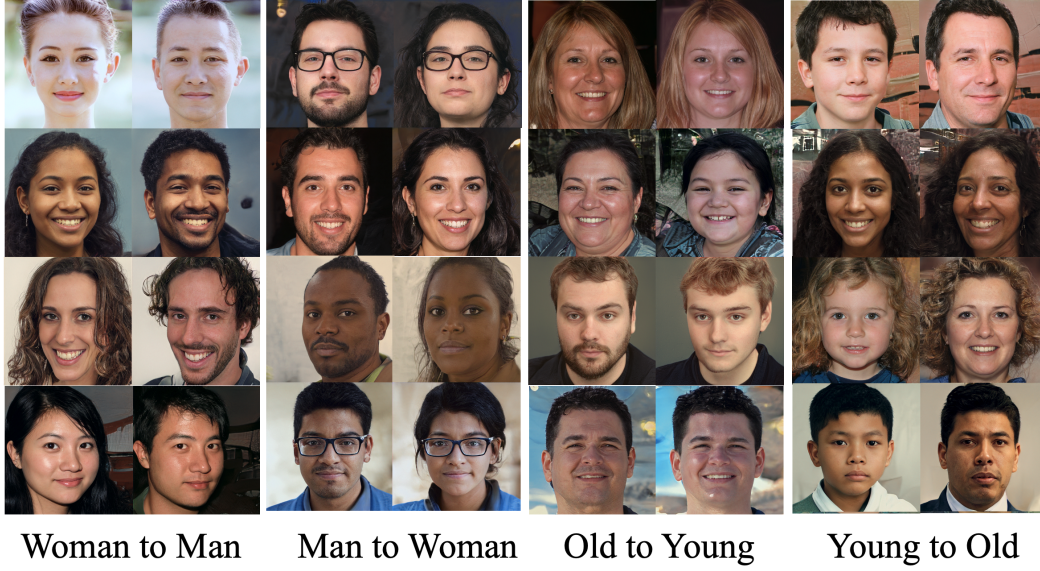
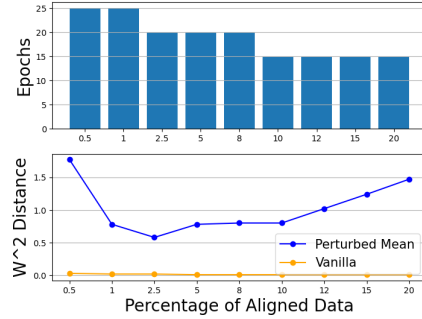


Figure 8: Image Translations using our FSBM

Table 4: Duration of each training epoch and memory usage (GB) in each of the different stages of Alg. 1, measured on the age translation task

	Epoch Time (hours:min)	Memory (GB)
<i>Guidance (lines 4-8)</i>	<i>00:01</i>	<i>0.18</i>
Path and Coupl. Opt. (line 9-11)	02:56	1.51

Figure 9: UP: Training Epochs, DOWN: The W^2 Distance under Vanilla and Perturbed Mean in the S-tunnel for varying the number of KP samples

Different selections for guidance function Defining the guidance function G is a crucial aspect of our FSBM framework. Although our analysis extends to any selection of function G , the specific type of guidance regularization term directly affects the nature of the attraction region and thus the guidance imposed by the KP samples. For instance, recall that in our analysis we used a relative distance-preserving guidance, while (Gu et al., 2022) employed a relation-preserving scheme, utilizing the Jensen-Shannon divergence. A more extensive comparison of different selections for the Guidance function and their effect on the efficacy of the model is a topic for future work.

5 CONCLUSION

In this work, we developed FSBM, a novel matching algorithm capable of using the information from pre-aligned data pairs, to guide the transport mapping of samples from the source to the target distribution. We demonstrated our algorithm’s enhanced generalizability and reduced training over prior methods through extensive experimentation. In future work, we wish to delve deeper into the effect of the KP samples on the matching procedure, along with using more intricate and task-specific guidance functions for further improvements. In summary, our work paves new ways for training diffusion models with partially aligned data.

REFERENCES

- Michael S. Albergo, Nicholas M. Boffi, and Eric Vanden-Eijnden. Stochastic interpolants: A unifying framework for flows and diffusions, 2023.
- Brian DO Anderson. Reverse-time diffusion equation models. *Stochastic Processes and their Applications*, 12(3):313–326, 1982.
- Jean-David Benamou and Yann Brenier. A computational fluid mechanics solution to the monge-kantorovich mass transfer problem. *Numerische Mathematik*, 84(3):375–393, 2000.
- Marin Biloš, Johanna Sommer, Syama Sundar Rangapuram, Tim Januschowski, and Stephan Günnemann. Neural flows: Efficient alternative to neural odes. *Advances in neural information processing systems*, 34:21325–21337, 2021.
- Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. *Advances in neural information processing systems*, 31, 2018.
- Tianrong Chen, Guan-Hong Liu, and Evangelos A Theodorou. Likelihood training of schrödinger bridge using forward-backward sdes theory. *arXiv preprint arXiv:2110.11291*, 2021.
- Yongxin Chen, Tryphon T Georgiou, and Michele Pavon. On the relation between optimal transport and schrödinger bridges: A stochastic control viewpoint. *Journal of Optimization Theory and Applications*, 169:671–691, 2016.
- Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in neural information processing systems*, 26, 2013.
- Valentin De Bortoli, James Thornton, Jeremy Heng, and Arnaud Doucet. Diffusion schrödinger bridge with applications to score-based generative modeling. *Advances in Neural Information Processing Systems*, 34:17695–17709, 2021.
- Rémi Flamary, Nicolas Courty, Alexandre Gramfort, Mokhtar Z. Alaya, Aurélie Boissunon, Stanislas Chambon, Laetitia Chapel, Adrien Corenflos, Kilian Fatras, Nemo Fournier, Léo Gautheron, Nathalie T.H. Gayraud, Hicham Janati, Alain Rakotomamonjy, Ievgen Redko, Antoine Rolet, Antony Schutz, Vivien Seguy, Danica J. Sutherland, Romain Tavenard, Alexander Tong, and Titouan Vayer. Pot: Python optimal transport. *Journal of Machine Learning Research*, 22(78):1–8, 2021. URL <http://jmlr.org/papers/v22/20-451.html>.
- Jason Gaitonde, Jon Kleinberg, and Éva Tardos. Polarization in geometric opinion dynamics. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pp. 499–519, 2021.
- Ivan Gentil, Christian Léonard, and Luigia Ripani. About the analogy between optimal transport and minimal entropy. In *Annales de la Faculté des sciences de Toulouse: Mathématiques*, volume 26, pp. 569–600, 2017.
- Promit Ghosal, Marcel Nutz, and Espen Bernton. Stability of entropic optimal transport and schrödinger bridges. *Journal of Functional Analysis*, 283(9):109622, 2022.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014.
- Xiang Gu, Yucheng Yang, Wei Zeng, Jian Sun, and Zongben Xu. Keypoint-guided optimal transport with applications in heterogeneous domain adaptation. *Advances in Neural Information Processing Systems*, 35:14972–14985, 2022.
- Xiang Gu, Liwei Yang, Jian Sun, and Zongben Xu. Optimal transport-guided conditional score-based diffusion model. *Advances in Neural Information Processing Systems*, 36:36540–36552, 2023a.
- Xiang Gu, Yucheng Yang, Wei Zeng, Jian Sun, and Zongben Xu. Keypoint-guided optimal transport, 2023b.
- Nikita Gushchin, Sergei Kholkin, Evgeny Burnaev, and Alexander Korotin. Light and optimal schrödinger bridge matching. In *Forty-first International Conference on Machine Learning*, 2024.

- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- Aapo Hyvärinen and Peter Dayan. Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research*, 6(4), 2005.
- Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 4401–4410, 2019.
- Alexander Korotin, Nikita Gushchin, and Evgeny Burnaev. Light schrödinger bridge. *arXiv preprint arXiv:2310.01174*, 2023.
- Flavien Léger. A gradient descent perspective on sinkhorn. *Applied Mathematics & Optimization*, 84(2):1843–1855, 2021.
- Christian Léonard. A survey of the schrödinger problem and some of its connections with optimal transport. *arXiv preprint arXiv:1308.0215*, 2013.
- Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022.
- Guan-Hong Liu, Tianrong Chen, Oswin So, and Evangelos Theodorou. Deep generalized schrödinger bridge. *Advances in Neural Information Processing Systems*, 35:9374–9388, 2022a.
- Guan-Hong Liu, Arash Vahdat, De-An Huang, Evangelos A Theodorou, Weili Nie, and Anima Anandkumar. Image-to-image schrödinger bridge. *arXiv preprint arXiv:2302.05872*, 2023.
- Guan-Hong Liu, Yaron Lipman, Maximilian Nickel, Brian Karrer, Evangelos A Theodorou, and Ricky TQ Chen. Generalized Schrödinger bridge matching. In *International Conference on Learning Representations*, 2024.
- Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. *arXiv preprint arXiv:2209.03003*, 2022b.
- Xingchao Liu, Lemeng Wu, Mao Ye, and Qiang Liu. Let us build bridges: Understanding and extending diffusion generative models, 2022c.
- I Loshchilov. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.
- Facundo Mémoli. Gromov–wasserstein distances and the metric approach to object matching. *Foundations of computational mathematics*, 11:417–487, 2011.
- Aamir Mustafa and Rafał K Mantiuk. Transformation consistency regularization—a semi-supervised paradigm for image-to-image translation. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVIII 16*, pp. 599–615. Springer, 2020.
- Kirill Neklyudov, Rob Brekelmans, Daniel Severo, and Alireza Makhzani. Action matching: Learning stochastic dynamics from samples. In *International conference on machine learning*, pp. 25858–25889. PMLR, 2023.
- Marcel Nutz. Introduction to entropic optimal transport. *Lecture notes, Columbia University*, 2021.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. In *Advances in neural information processing systems*, pp. 8026–8037, 2019.
- Michele Pavon, Giulio Trigila, and Esteban G Tabak. The data-driven schrödinger bridge. *Communications on Pure and Applied Mathematics*, 74(7):1545–1573, 2021.
- Stefano Peluchetti. Diffusion bridge mixture transports, schrödinger bridge problems and generative modeling. *Journal of Machine Learning Research*, 24(374):1–51, 2023.

- Gabriel Peyré, Marco Cuturi, et al. Computational optimal transport: With applications to data science. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, 2019.
- Stanislav Pidhorskyi, Donald A Adjeroh, and Gianfranco Doretto. Adversarial latent autoencoders. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14104–14113, 2020.
- Simo Särkkä and Arno Solin. *Applied stochastic differential equations*, volume 10. Cambridge University Press, 2019.
- Ryoma Sato, Marco Cuturi, Makoto Yamada, and Hisashi Kashima. Fast and robust comparison of probability measures in heterogeneous spaces. *arXiv preprint arXiv:2002.01615*, 2020.
- Erwin Schrödinger. *Über die umkehrung der naturgesetze*. Verlag der Akademie der Wissenschaften in Kommission bei Walter De Gruyter u . . . , 1931.
- Yuyang Shi, Valentin De Bortoli, Andrew Campbell, and Arnaud Doucet. Diffusion schrödinger bridge matching, 2023.
- Vignesh Ram Somnath, Matteo Pariset, Ya-Ping Hsieh, Maria Rodriguez Martinez, Andreas Krause, and Charlotte Bunne. Aligned diffusion schrödinger bridges. In *Uncertainty in Artificial Intelligence*, pp. 1985–1995. PMLR, 2023.
- Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020.
- Evangelos Theodorou, Jonas Buchli, and Stefan Schaal. A generalized path integral control approach to reinforcement learning. *The Journal of Machine Learning Research*, 11:3137–3181, 2010.
- Francisco Vargas, Pierre Thodoroff, Austen Lamacraft, and Neil Lawrence. Solving schrödinger bridges via maximum likelihood. *Entropy*, 23(9):1134, 2021.
- Cédric Villani et al. *Optimal transport: old and new*, volume 338. Springer, 2009.
- Pascal Vincent. A connection between score matching and denoising autoencoders. *Neural computation*, 23(7):1661–1674, 2011.

A APPENDIX

A.1 PROOF OF THEOREM 3.2

Assume $X_0, X_1 \in \mathbb{R}^d$, with $X_0 \sim \pi_0$, and $X_1 \sim \pi_1$, and consider a stochastic random variable X_t connecting X_0 and X_1 , whose law is the continuous marginal probability path p_t joining π_0 , and π_1 . Starting from the static semi-supervised guided entropic OT problem in Eq. (8), we modify it appropriately and obtain the following relaxed dynamic formulation.

$$\begin{aligned} \min_{p_t, u_t} \int_{\mathbb{R}^d \times [0,1]} & \left(\frac{1}{2} \|u_t\|^2 + \mathbb{E}_{p_{0|t}} [|u_t^\top G_{t,0}| + \frac{\sigma^2}{2} \Delta G_{t,0}] \right) p_t dx dt, \\ \text{s.t. } \frac{\partial p_t}{\partial t} &= -\nabla \cdot (u_t p_t) + \frac{\sigma^2}{2} \Delta p_t, \text{ and } p_0 = \pi_0, p_1 = \pi_1 \end{aligned} \quad (16)$$

Proof. First, let $c(x_0, x_1)$ be the cost from the static optimal transport expression, and let π_0, π_1 be the two given boundary laws. Recall the static formulation: $\min_{\pi} \int \frac{1}{2} \|x_0 - x_1\|^2 + G(x_0, x_1) - \epsilon H(\pi_{0,1})$, where $H(\pi_{0,1}) = -\int \log \pi_{0,1} \pi_{0,1} dx_0 dx_1$. Our objective is to derive a dynamic formulation for the static EOT problem with the cost function $c(x_0, x_1) = \|x_0 - x_1\|^2 + G(x_0, x_1)$. If we arrange the cost function

$$\min_{\pi} \int \left(\frac{1}{2} \|x_0 - x_1\|^2 + \log \pi_{0,1} \right) \pi_{0,1} dx_0 dx_1 + \int G(x_0, x_1) \pi_{0,1} dx_0 dx_1 \quad (17)$$

Assuming a Gaussian prior, notice that the first term is the relative entropy between the joint density optimized plan $\pi_{0,1}$, and the prior π^Q , $KL(\pi_{0,1} | \pi_{0,1}^Q)$, for which it holds that $\log \pi_{0,1}^Q \propto \frac{1}{2} \|x_0 - x_1\|^2$, hence $\pi^Q \propto e^{\frac{1}{2} \|x_0 - x_1\|^2}$. We rewrite the objective in Eq. (17), as $KL(\pi_{0,1} | \pi_{0,1}^Q) + \int G(x_0, x_1) \pi_{0,1} dx_0 dx_1$. In this work, our focal point is the derivation of a dynamic expression for the semi-supervised regularized cost function in 8. Through the prism of entropic interpolation [Gentil et al. \(2017\)](#), the solution of dynamical OT formulations correspond to solving ‘local’ versions of optimal transport problems between infinitesimally close distributions [Villani et al. \(2009\)](#).

Initially, consider the cost $c(X_0, X_1) = KL(\pi_{0,1} | \pi_{0,1}^Q) + \int G(x_0, x_1) \pi_{0,1} dx_0 dx_1$ that is associated with transporting samples from π_0, π_1 . Let a stochastic random variable path $(X_t)_{0 \leq t \leq 1}$ joining $X_0 \sim \pi_0$ to $X_1 \sim \pi_1$, namely X_t is an interpolation from X_0 to X_1 . This random variable is assumed to be the solution of the following SDE

$$dX_t = u_t dt + \sigma dW_t \quad (18)$$

Equivalent we can consider a marginal density induced by Eq. 18 $(p_t)_{0 \leq t \leq 1}$ which joins π_0 , and π_1 [Villani et al. \(2009\)](#) admits the Fokker Plank Equation (FPE):

$$\frac{\partial p_t}{\partial t} = -\nabla \cdot (u_t p_t) + \frac{1}{2} \sigma^2 \Delta p_t \quad (19)$$

Now, application of [[Villani et al. \(2009\)](#), Theorem 7.21] suggests that if we consider p_t be the solution to the infimum above, namely a minimizing curve. The endpoints of the random curve X_t , whose law admits an optimal transference plan, are an optimal coupling $(X_{t_i}, X_{t_{i+1}})$. Then from the definition of the optimal cost and the minimizing property of the X_t

$$\min_{p_t} \min_{X_t} \sum_{i=0}^{N-1} KL(p_{t_{i+1}} | q_{i \rightarrow i+1}) + \mathbb{E}_{p_t} \sum_{i=0}^{N-1} G(X_{t_i}, X_{t_{i+1}}) \quad (20)$$

for with $p_0 = \pi_0$, and $p_N = \pi_N$. It generally holds $\mathbb{E}_{p_t} c(x_i, x_{i+1}) \geq \mathbb{E}_{p_t} c(x_0, x_{i+1}) - \mathbb{E}_{p_t} c(x_0, x_i)$, with the equality being true for the minimizing path. Therefore, we can say that

$$\min_{p_t} \min_{X_t} \mathbb{E}_{p_t} c(X_i, X_{i+1}) = \min_{p_t} \min_{X_t} \mathbb{E}_{p_t} c(X_0, X_{i+1}) - \mathbb{E}_{p_t} c(X_0, X_i) \quad (21)$$

This yields the modified version of the dynamic analogue

$$\begin{aligned} \min_{p_t} \min_{X_t} \sum_{i=0}^{N-1} & KL(p_{t_{0,i+1}} | q_{0,i+1}) - KL(p_{t_{0,i}} | q_{0,i}) \\ & + \sum_{i=0}^{N-1} \mathbb{E}_{p_{t_{i+1},0}} [G(X_{t_{i+1}}, X_0)] - \mathbb{E}_{p_{t_{i,0}}} [G(X_{t_i}, X_0)] \end{aligned} \quad (22)$$

However, note that $KL(p_{t_0,i+1}|q_{0,i+1}) - KL(p_{t_0,i}|q_{0,i}) \leq KL(p_{t_i \rightarrow i+1}|q_{i \rightarrow i+1})$. Subsequently, note for any coupling between p_{t_i} , and $p_{t_{i+1}}$, it holds that

$$\mathbb{E}_{p_{t_{i+1},0}}[c(X_{t_{i+1}}, X_0)] - \mathbb{E}_{p_{t_i,0}}[c(X_{t_i}, X_0)] \leq \mathbb{E}_{p_{t_i \rightarrow i+1,0}}[c(X_{t_i}, X_0) - c(X_{t_{i+1}}, X_0)]$$

Finally, in order to ensure positivity of the difference $G(X_{t_i}, X_0) - G(X_{t_{i+1}}, X_0)$, we upper bound with its absolute value, and finally obtain the relaxed dynamic objective, which admits the following formulation

$$\min_{p_t, X_t} \sum_{i=0}^{N-1} \left| \mathbb{E}_{p_{t_{i+1},0}} G(X_{t_i}, X_0) - G(X_{t_{i+1}}, X_0) \right| + \sum_{i=0}^{N-1} KL(p_{t_i \rightarrow i+1}|q_{i \rightarrow i+1}) \quad (23)$$

For the first term, assuming the Markovian property of the base measure, we apply the Girsanov theorem in the infinitesimally small time-step $i \rightarrow i+1$:

$$KL(p_{i,i+1}|q_{i,i+1}) = \int_{t_i}^{t_{i+1}} \|u_{t_i}\|^2 dt = \|u_{t_i}\|^2 \Delta t_i \quad (24)$$

For the second term, we have

$$\left| \mathbb{E}_{p_{t_{i+1},0}} [G(X_{t_{i+1}}, X_0) - G(X_{t_i}, X_0)] \right|$$

For discretized random variable X_{t_i} obeying the following discretized SDE: $X_{t_{i+1}} = X_{t_i} + u_{t_i} \Delta t_i + \sigma_{t_i} \Delta W_{t_i}$, application of the discrete version of the Ito's Lemma yields

$$\sum_{i=0}^{N-1} \left| \mathbb{E}_{p_{t_{i+1},0}} \left[\left(\frac{\partial G_{t_i}}{\partial t} + u_t^T \nabla G_{t_i} + \frac{1}{2} \sigma_{t_i}^2 \text{Tr}(\nabla^2 G_{t_i}) \right) \Delta t_i + (\sigma_{t_i} \nabla G_{t_i}) \Delta W_{t_i} \right] \right| \quad (25)$$

It holds that $\frac{\partial G_{t_i}}{\partial t} = 0$, since G does not explicitly depend on t , and also assume that the Wiener process increment has zero expectation, hence the equation above is rewritten as

$$\sum_{i=0}^{N-1} \left| \mathbb{E} \left[(u_t^T \nabla G_{t_i} + \frac{1}{2} \sigma_{t_i}^2 \text{Tr}(\nabla^2 G_{t_i})) \Delta t_i \right] \right| \quad (26)$$

Now, we make use of the fact that $|\mathbb{E}[f(x)]| \leq \mathbb{E}[|f(x)|]$, to take the expectation out of the absolute value.

$$\begin{aligned} \sum_{i=0}^{N-1} \left| \mathbb{E} \left[u_t^T \nabla G + \frac{\sigma^2}{2} \text{Tr}(\nabla^2 G_{t_i}) \right) \Delta t \right] \right| &\leq \sum_{i=0}^{N-1} \mathbb{E} \left[|u_t^T \nabla G + \frac{\sigma^2}{2} \text{Tr}(\nabla^2 G_{t_i})| \Delta t \right] \\ &\leq \sum_{i=0}^{N-1} \mathbb{E} \left[(|u_t^T \nabla G| + \frac{\sigma^2}{2} \text{Tr}(\nabla^2 G_{t_i})) \Delta t \right] \end{aligned} \quad (27)$$

Note that G is convex, i.e., $\frac{\sigma^2}{2} \text{Tr}(\nabla^2 G) \geq 0$, thus we conclude that the expression in Eq. (34) is upper bound by

$$\mathbb{E}_{p_{t_{i+1},0}} \sum_{i=0}^{N-1} G(X_{t_i}, X_{t_{i+1}}) \leq \sum_{i=0}^{N-1} \mathbb{E}_{p_{t_{i+1},0}} \left[(|u_t^T \nabla G| + \frac{\sigma^2}{2} \text{Tr}(\nabla^2 G_{t_i})) \Delta t \right] \quad (28)$$

Finally, let us substitute back in the expression above the function G , and consider the continuous analogue of Eq. (24) and 28

$$\min_{p_t} \min_{u_t} \int_0^1 \mathbb{E}_{p_t} [\|u_t\|^2] + \mathbb{E}_{p_{t,0}} \left[(|u_t^T \nabla G_{t,0}| + \frac{\sigma^2}{2} \Delta G_{t,0}) \right] dt \quad (29)$$

where $G_{t,0} = G(X_t, X_0)$. Finally, we conclude the relaxed dynamic formulation is given by

$$\min_{p_t} \min_{u_t} \int_0^1 \mathbb{E}_{p_t} \left[\frac{1}{2} \|u_t\|^2 + \mathbb{E}_{p_{0|t}} [|u_t^T \nabla G_{t,0}| + \frac{\sigma^2}{2} \Delta G_{t,0}] \right] dt \quad (30)$$

subjected to the FPE: $\frac{\partial p_t}{\partial t} = -\nabla \cdot (u_t p_t) + \frac{1}{2} \sigma^2 \Delta p_t$, and under the boundary constraints: $p_0 = \pi_0$, and $p_1 = \pi_1$. \square

A.2 PROOF OF PROPOSITION 3.3

Let the continuous marginal path p_t satisfy the following decomposition $p_t = \int p_{t|0,1} \pi_{0,1} dx_0 dx_1$. For optimization with respect to the intermediate path, we freeze the joint distribution at the boundaries $\pi_{0,1}$, which results in Eq. (9) being recasted as

$$\begin{aligned} \min_{p_{t|0,1}} \int_0^1 \int_{\mathbb{R}^d} & \left(\frac{1}{2} \|u_{t|0,1}\|^2 + |u_{t|0,1}^\top \nabla G(X_t, x_0)| + \frac{\sigma^2}{2} \Delta G(X_t, x_0) p_{t|0,1} \right) dx dt, \\ \text{s.t. } \frac{\partial p_{t|0,1}}{\partial t} &= -\nabla \cdot (p_{t|0,1} u_t) + \frac{1}{2} \sigma^2 \Delta p_{t|0,1} \end{aligned} \quad (31)$$

Proof. Initially, let us deal with the dynamic objective in Eq. (30)

$$\min_{p_t} \min_{u_t} \int_0^1 \mathbb{E}_{p_t} \left[\frac{1}{2} \|u\|^2 + \mathbb{E}_{p_{0|t}} [|u_t^\top \nabla G_{t,0}| + \frac{\sigma^2}{2} \Delta G_{t,0}] \right] dt \quad (32)$$

The marginal is constructed as a mixture of conditional probability paths $p_t = \int p_{t|0,1} \pi_{0,1} dx_0 dx_1$. Now, we sample pairs of x_0, x_1 , and fix the couplings. This is equivalent to fixing the vector, which induces the couplings. Consequently, fixing the coupling results in recasting Eq. (30) as

$$\min_{p_{t|0,1}} \int \mathbb{E}_{p_t} \left[\frac{1}{2} \|u_{t|0,1}\|^2 + [|u_{t|0,1}^\top \nabla G_{t|0}| + \frac{\sigma^2}{2} \Delta G_{t|0}] \right] dt \quad (33)$$

where $u_{t|0,1}$ is used to emphasize the drift that corresponds to the conditional probability path, $G_{t|0} \equiv G(X_t, x_0)$. Additionally, notice that by fixing the coupling, the expectation $\mathbb{E}_{p_{0|t}}$ is dropped.

$$\min_{p_t} \min_{X_t} \mathbb{E}_{p_{t_{i+1}|0,1}} \sum_{i=0}^N G(X_{t_i}, X_{t_{i+1}}) \quad (34)$$

Now, we want to prove that the conditional path preserves the marginal. Our process to show that is to isolate from every term in the Fokker Plank the joint distribution $\pi_{0,1}$, and preserve the Fokker Plank formulation after freezing the coupling.

$$\begin{aligned} \bullet \frac{\partial p_t}{\partial t} &= \frac{\partial}{\partial t} \int p_{t|0,1} \pi_{0,1} dx_0 dx_1 = \int \frac{\partial p_{t|0,1}}{\partial t} \pi_{0,1} dx_0 dx_1 = \mathbb{E}_{\pi_{0,1}} \left[\frac{\partial p_{t|0,1}}{\partial t} \right] \\ \bullet \nabla \cdot (u_{t|0,1} p_t) &= \nabla \cdot (u_{t|0,1} \int p_{t|0,1} \pi_{0,1} dx_0 dx_1) = \int \nabla \cdot (u_{t|0,1} p_{t|0,1}) \pi_{0,1} dx_0 dx_1 = \\ &= \mathbb{E}_{\pi_{0,1}} [\nabla \cdot (u_{t|0,1} p_{t|0,1})] \\ \bullet \Delta p_t &= \nabla \cdot \nabla p_t = \nabla \cdot \nabla \int p_{t|0,1} \pi_{0,1} dx_0 dx_1 = \int \nabla \cdot \nabla p_{t|0,1} \pi_{0,1} dx_0 dx_1 = \mathbb{E}_{\pi_{0,1}} [\Delta p_{t|0,1}] \end{aligned}$$

where in every term we used Fubini's Theorem, to change the order between the integral from the definition of the marginal and the differentiations from the FPE. Therefore, we can conclude that by freezing the coupling the Eq. (30) is recasted as

$$\begin{aligned} \min_{p_{t|0,1}} \int \mathbb{E}_{p_t} & \left[\frac{1}{2} \|u_{t|0,1}\|^2 + [|u_{t|0,1}^\top \nabla G_{t|0}| + \frac{\sigma^2}{2} \Delta G_{t|0}] \right], \\ \text{s.t. } \frac{\partial p_{t|0,1}}{\partial t} &= -\nabla \cdot (u_{t|0,1} p_{t|0,1}) + \frac{\sigma^2}{2} \Delta p_{t|0,1} \end{aligned} \quad (35)$$

□

A.3 PROOF OF PROPOSITION 3.5

The convex conjugate of the Lagrangian in Eq. (10) is defined as the Hamiltonian $H(x_t, a_t, t) = \sup_{u_{t|0,1}} \langle u_{t|0,1}, \nabla a_t \rangle - L(x_t, u_{t|0,1}, t)$. The optimization with respect to the drift $u_{t|0,1}$ yields $u_{t|0,1} = a_t - g(a_t^\top \nabla G_{t|0}) \nabla G_{t|0}$, where $\nabla G_{t|0} \equiv G(X_t, x_0)$, and

$$g(a_t^\top \nabla G_{t|0}) = \begin{cases} \frac{a_t^\top \nabla G_{t|0}}{\|\nabla G_{t|0}\|^2} & \text{if } |a_t^\top \nabla G_{t|0}| \leq \|\nabla G_{t|0}\|^2 \\ \text{sgn}(a_t^\top \nabla G_{t|0}) & \text{else} \end{cases}$$

Finally, we obtain the Hamiltonian associated with the Lagrangian in Eq. (10)

$$H(x_t, a_t, t) = \frac{1}{2} \|a_t - g(a_t^\top \nabla G_t|_0) \cdot \nabla G_t|_0\|^2 - \frac{\sigma^2}{2} \Delta G(X_t, x_0) \quad (36)$$

Proof. We begin with the definition of the Hamiltonian, dependent on the gradient field a_t .

$$H(x_t, a_t, t) = \sup_{u_{t|0,1}} \langle u_{t|0,1}, a_t \rangle - L(x_t, u_{t|0,1}, t) \quad (37)$$

A more detailed analysis on the gradient field is given in Neklyudov et al. (2023). Given the convexity of Eq. (10), we compute the Hamiltonian through the subgradient of the absolute value $\partial(|u_t^\top \nabla G_t|)$. Optimizing with respect to the conditioned drift $u_{t|0,1}$ yields the following

$$0 = a_t - (u_t + \partial(u_t^\top \nabla G_t)(\nabla G_t)) \Rightarrow u_{t|0,1}^* = a_t - \partial(u_t^\top \nabla G_t) \cdot (\nabla G_t) \quad (38)$$

where $\partial(\cdot)$ denotes the subgradient. Note that

$$\begin{cases} u_t^\top \nabla G_t > 0 : u_t = a_t - (\nabla G_t) \Rightarrow (a_t - (\nabla G_t))^\top \nabla G_t \geq 0 \Rightarrow a_t^\top \nabla G_t \geq \|\nabla G_t\|^2 \geq 0 \\ u_t^\top \nabla G_t < 0 : u_t = a_t + (\nabla G_t) \Rightarrow (a_t + (\nabla G_t))^\top \nabla G_t \leq 0 \Rightarrow a_t^\top \nabla G_t \leq -\|\nabla G_t\|^2 \leq 0 \\ u_t^\top \nabla G_t = 0 : \partial(u_t^\top \nabla G_t) = [-1, 1], \text{ from which we infer } a_t^\top \nabla G_t = \alpha \|\nabla G_t\|^2, \text{ for } \alpha \in [-1, 1]. \end{cases}$$

This essentially means that the inner products $u_t^\top \nabla G_t$, and $a_t^\top \nabla G_t$ share the same sign, namely $\text{sgn}(u_t^\top \nabla G_t) = \text{sgn}(a_t^\top \nabla G_t)$, when $u_t^\top \nabla G_t \neq 0$, otherwise $a_t^\top \nabla G_t = \alpha \|\nabla G_t\|^2$.

$$\begin{cases} \partial(u_t^\top \nabla G_t) = \text{sgn}(a_t^\top \nabla G_t), & \text{for } |a_t^\top \nabla G_t| \geq \|\nabla G_t\|^2 \\ \partial(u_t^\top \nabla G_t) = \frac{a_t^\top \nabla G_t}{\|\nabla G_t\|^2}, & \text{for } |a_t^\top \nabla G_t| \leq \|\nabla G_t\|^2 \end{cases}$$

We define function $g : \mathbb{R} \rightarrow \mathbb{R}$, to be equal to:

$$g(a_t^\top \nabla G_t) = \begin{cases} \text{sgn}(a_t^\top \nabla G_t), & \text{for } |a_t^\top \nabla G_t| \geq \|\nabla G_t\|^2 \\ \frac{a_t^\top \nabla G_t}{\|\nabla G_t\|^2}, & \text{for } |a_t^\top \nabla G_t| \leq \|\nabla G_t\|^2 \end{cases}$$

and substitute it in Eq. (37), which yields an expression for u_t as a function of the gradient field

$$u_{t|0,1} = a_t - g(a_t^\top \nabla G_t) \cdot \nabla G_t \quad (39)$$

Now, we will provide an expression for the Hamiltonian dependent only on the gradient field a_t and the feedback term $\nabla G_t(X_t, x_0)$.

$$H(x_t, a_t, t) = \langle a_t - g_t \nabla G_t, a_t \rangle - \left(\frac{1}{2} \|a_t - g_t \nabla G_t\|^2 - \langle a_t - g_t \nabla G_t, \nabla G_t \rangle \right) - \frac{\sigma^2}{2} \Delta G_t \quad (40)$$

Therefore, we conclude that the Hamiltonian can be expressed as:

$$H(x_t, a_t, t) = \frac{1}{2} \|a_t - g(a_t^\top \nabla G_t) \cdot \nabla G_t\|^2 - \frac{\sigma^2}{2} \Delta G_t \quad (41)$$

□

A.4 PROOF OF PROPOSITION 3.6

Given the optimal conditional drift $u_{t|0,1}^*$, we match the parameterized drift u_t^θ , by minimizing the optimality gap given by our Entropic Lagrangian Bridge Matching objective as

$$\min_{\theta} \int_0^1 \mathbb{E}_{p_{0,1}} \mathbb{E}_{p_{t|0,1}} \|a_t^\theta - u_{t|0,1}^* - g(a_t^\top \nabla G_t) \cdot \nabla G_t|_0\|^2 dt \quad (42)$$

From Proposition 3.5, we find that the parameterized drift can be expressed as $u_t^\theta = a_t^\theta - g(a_t^\top \nabla G_t) \cdot \nabla G_t|_0$. Therefore, the matching objective of Eq. (13) can be rewritten as follows

$$\min_{\theta} \int_0^1 \mathbb{E}_{p_t} \|u_{t|0,1}^* - u_t^\theta\|^2 dt \quad (43)$$

Proof. Given the prescribed path, we seek to learn a parameterized drift field that matches the optimal conditioned paths, acquired from the previous step, which minimized the Lagrangian \mathcal{L} . Following advancements in matching frameworks, we extend the notion of bridging the variational gap for an entropy regularized optimal transport problem with general langriangian cost. We will first prove the following proposition which shows that we can express the variational gap that we need to bridge is expressed through the Bregman divergence.

Proposition A.1. *For a Lagrangian function $L(X_t, u_t, t)$, the convex conjugate is given by the Hamiltonian: $H(x_t, a_t, t) = \sup_{u_t} \langle a_t, u_t \rangle - L(x_t, u_t, t)$, dependent on the gradient field a_t . The variational gap between the optimal drift and the parameterized gradient field is expressed as the following Bregman divergence.*

Proof. Starting from a general minimization problem for a Lagrangian strictly convex with u_t

$$\min_{u_{t|0,1}} \int_0^1 \int_{\mathbb{R}^d} L(x_t, u_{t|0,1}, t) p_{t|0,1} dx dt, \text{ subj.: } \frac{\partial p_{t|0,1}}{\partial t} = -\nabla \cdot (p_{t|0,1} u_{t|0,1}) + \frac{1}{2} \sigma^2 \Delta p_{t|0,1}, \text{ and } p_0 = x_0, p_1 = x_1 \quad (44)$$

and expressing it as a Lagrange optimization problem yields, with a_t is the Lagrange multiplier enforcing the Fokker Plank equation constraint

$$\begin{aligned} \mathcal{L}(\theta) &= \min_{u_{t|0,1}} \int_0^1 \int_{\mathbb{R}^d} L(x_t, u_{t|0,1}, t) p_{t|0,1} dx dt, \text{ subj.: } \frac{\partial p_{t|0,1}}{\partial t} = -\nabla \cdot (p_{t|0,1} u_{t|0,1}) + \frac{1}{2} \sigma^2 \Delta p_{t|0,1} \\ &= \sup_{a_t} \min_{u_{t|0,1}} \int_0^1 \int_{\mathbb{R}^d} L(x_t, u_{t|0,1}, t) p_{t|0,1} dx dt + \int_0^1 \int_{\mathbb{R}^d} a_t \left(\frac{\partial p_{t|0,1}}{\partial t} + \nabla \cdot (p_{t|0,1} u_{t|0,1}) - \frac{1}{2} \sigma^2 \Delta p_{t|0,1} \right) dx dt \\ &= \sup_{a_t} \min_{u_{t|0,1}} \underbrace{\int_0^1 \int_{\mathbb{R}^d} L(x_t, u_{t|0,1}, t) p_{t|0,1} dx dt}_{I_1} + \underbrace{\int_0^1 \int_{\mathbb{R}^d} a_t \frac{\partial p_{t|0,1}}{\partial t} dx dt}_{I_2} \\ &\quad + \underbrace{\int_0^1 \int_{\mathbb{R}^d} a_t \nabla \cdot (p_{t|0,1} u_{t|0,1}) dx dt}_{I_3} - \underbrace{\int_0^1 \int_{\mathbb{R}^d} \frac{a_t}{2} \sigma^2 \Delta p_{t|0,1} dx dt}_{I_4} \end{aligned} \quad (45)$$

We separate each one of the integrals and perform integration by parts

- $I_2 = \int_0^1 \int_{\mathbb{R}^d} a_t \frac{\partial p_{t|0,1}}{\partial t} dx dt = [a_t p_{t|0,1}]_0^1 - \int_0^1 \int_{\mathbb{R}^d} p_{t|0,1} \frac{\partial a_t}{\partial t} dx dt$
- $I_3 = \int_0^1 \int_{\mathbb{R}^d} a_t \nabla \cdot (p_{t|0,1} u_{t|0,1}) dx dt = - \int_0^1 \int_{\mathbb{R}^d} \langle \nabla a_t, u \rangle p_{t|0,1} dt dx + \frac{\sigma^2}{2} \int_0^1 \int_{\mathbb{R}^d} \langle \nabla p_{t|0,1}, d\mathbf{n} \rangle dt$
- $I_4 = \int_0^1 \int_{\mathbb{R}^d} \langle \nabla p_{t|0,1}, d\mathbf{n} \rangle dt - \int_0^1 \int_{\mathbb{R}^d} a_t \Delta p_{t|0,1} dt dx = \int_0^1 \int_{\mathbb{R}^d} \langle \nabla a_t, \nabla p_{t|0,1} \rangle dx dt - \int_0^1 \int_{\mathbb{R}^d} \Delta a_t p_{t|0,1} dx dt + \frac{\sigma^2}{2} \int_0^1 \int_{\mathbb{R}^d} \langle p_{t|0,1}, d\mathbf{n} \rangle dt$

Substituting them back to Eq. (45), we obtain

$$\begin{aligned} \mathcal{L} &= \sup_{a_t} \min_{u_{t|0,1}} \int_0^1 \int_{\mathbb{R}^d} \left(L(x_t, u_{t|0,1}, t) - \frac{\partial a_t}{\partial t} - \langle \nabla a_t, u \rangle - \Delta a_t \right) p_{t|0,1} dx dt - \mathbb{E}_{p_0}[x_0] + \mathbb{E}_{p_1}[x_1] \\ &\quad + \frac{\sigma^2}{2} \int_0^1 \int_{\mathbb{R}^d} \langle p_{t|0,1}, d\mathbf{n} \rangle \\ &= \sup_{a_t} \int_0^1 \int_{\mathbb{R}^d} \left(-H(x_t, a_t, t) - \frac{\partial a_t}{\partial t} - \Delta a_t \right) p_{t|0,1} dx dt - \mathbb{E}_{p_0}[x_0] + \mathbb{E}_{p_1}[x_1] + \frac{\sigma^2}{2} \int_0^1 \int_{\mathbb{R}^d} \langle p_{t|0,1}, d\mathbf{n} \rangle \\ &= \min_{a_t} \int_0^1 \int_{\mathbb{R}^d} \left(H(x_t, a_t, t) + \frac{\partial a_t}{\partial t} + \Delta a_t \right) p_{t|0,1} dx dt + \mathbb{E}_{p_0}[x_0] - \mathbb{E}_{p_1}[x_1] - \frac{\sigma^2}{2} \int_0^1 \int_{\mathbb{R}^d} \langle p_{t|0,1}, d\mathbf{n} \rangle \end{aligned} \quad (46)$$

Now, consider the optimality gap between a parameterized gradient field, and an optimal gradient field a_t^* , which corresponds to a loss function \mathcal{L}^* . We compute the optimality gap as

$$\begin{aligned}
\mathcal{L}^\theta - \mathcal{L}^* &= \int_0^1 \int_{\mathbb{R}^d} (H(x_t, a_t, t) + \frac{\partial a_t}{\partial t} + \Delta a_t) p_{t|0,1} dx dt + \mathbb{E}_{p_0}[x_0] - \mathbb{E}_{p_1}[x_1] - \frac{\sigma^2}{2} \int \oint a_t \langle p_{t|0,1}, d\mathbf{n} \rangle \\
&\quad - \int_0^1 \int_{\mathbb{R}^d} (H(x_t, a_t^*, t) + \frac{\partial a_t^*}{\partial t} + \Delta a_t^*) p_{t|0,1} dx dt + \mathbb{E}_{p_0}[x_0] - \mathbb{E}_{p_1}[x_1] - \frac{\sigma^2}{2} \int \oint a_t^* \langle p_{t|0,1}, d\mathbf{n} \rangle \\
&= \int_0^1 \int_{\mathbb{R}^d} (H(x_t, a_t, t) - H(x_t, a_t^*, t)) p_{t|0,1} dx dt + \underbrace{\int_0^1 \int_{\mathbb{R}^d} \frac{\partial a_t}{\partial t} p_{t|0,1} dx dt + \mathbb{E}_{p_0}[x_0] - \mathbb{E}_{p_1}[x_1]}_{-\int_0^1 \int_{\mathbb{R}^d} \frac{\partial p_{t|0,1}}{\partial t} a_t dx dt} \\
&\quad - \underbrace{\int_0^1 \int_{\mathbb{R}^d} \frac{\partial a_t^*}{\partial t} p_{t|0,1} dx dt + \mathbb{E}_{p_0}[x_0] - \mathbb{E}_{p_1}[x_1]}_{\int_0^1 \int_{\mathbb{R}^d} \frac{\partial p_{t|0,1}}{\partial t} a_t^* dx dt} + \underbrace{\int_0^1 \int_{\mathbb{R}^d} \Delta a_t p_{t|0,1} dx dt - \frac{\sigma^2}{2} \int \oint a_t \langle p_{t|0,1}, d\mathbf{n} \rangle}_{\int_0^1 \int_{\mathbb{R}^d} \Delta p_{t|0,1} a_t dx dt} \\
&\quad - \underbrace{\int_0^1 \int_{\mathbb{R}^d} \Delta a_t^* p_{t|0,1} dx dt - \frac{\sigma^2}{2} \int \oint a_t^* \langle p_{t|0,1}, d\mathbf{n} \rangle}_{\int_0^1 \int_{\mathbb{R}^d} \Delta p_{t|0,1} a_t^* dx dt} \\
&= \int_0^1 \int_{\mathbb{R}^d} (H(x_t, a_t, t) - H(x_t, a_t^*, t)) p_{t|0,1} dx dt + \int_0^1 \int_{\mathbb{R}^d} (a_t - a_t^*) \left(-\frac{\partial p_{t|0,1}}{\partial t} + \Delta p_{t|0,1} \right) dx dt \\
&= \int_0^1 \int_{\mathbb{R}^d} (H(x_t, a_t, t) - H(x_t, a_t^*, t)) p_{t|0,1} dx dt + \int_0^1 \int_{\mathbb{R}^d} (a_t - a_t^*) (\nabla \cdot (p_{t|0,1} u_{t|0,1}^*)) dx dt \\
&= \int_0^1 \int_{\mathbb{R}^d} (H(x_t, a_t, t) - H(x_t, a_t^*, t)) - \langle \nabla(a_t - a_t^*), u_t^* \rangle p_{t|0,1} dx dt \\
&= \int_0^1 \int_{\mathbb{R}^d} H(x_t, a_t, t) + L(x_t, u_{t|0,1}^*) - \langle a_t, u_{t|0,1}^* \rangle p_{t|0,1} dx dt \\
&= \int_0^1 \int_{\mathbb{R}^d} D_{L,H}[a_t : u_{t|0,1}^*] p_{t|0,1} dx dt
\end{aligned} \tag{47}$$

where the second to last equality stems from $L(x, u_{t|0,1}^*) = \sup_{a_t} \langle a_t, u_{t|0,1} \rangle - H(x, a_t, t)$, and the last equality is the definition of the Bregman divergence, finishing the proof of the proposition. \square

Therefore, to minimize the optimality gap, we sample sufficiently many uncoupled pairs (x_0, x_1) , and parameterize the gradient field $a_t \rightarrow a_t^\theta$, and try to match the parameterized gradient field on the optimized drift u_t^* . We express our matching objective as an expectation over the sampled pairs of optimality gap derived from the previous proposition

$$\min_{\theta} \mathbb{E}_{\pi_{0,1}}[\mathcal{L}^\theta - \mathcal{L}^*] = \min_{\theta} \mathbb{E}_{\pi_{0,1}} \left[\int_0^1 \int_{\mathbb{R}^d} H(x_t, a_t^\theta, t) + L(x_t, u_{t|0,1}^*) - \langle a_t^\theta, u_{t|0,1}^* \rangle p_{t|0,1} dx dt \right] \tag{48}$$

Substituting for the expressions we have for the functions L , and H , we obtain

$$\begin{aligned}
\mathbb{E}_{\pi_{0,1}}[\mathcal{L}^\theta - \mathcal{L}^*] &= \int_0^1 \int_{\mathbb{R}^d} \left(\frac{1}{2} \|a_t^\theta - g_t \cdot \nabla G_t\|^2 - \frac{\sigma^2}{2} \Delta G_t \right. \\
&\quad \left. + \frac{1}{2} \|u_{t|0,1}^*\|^2 + |u_{t|0,1}^*|^\top \nabla G_t + \frac{\sigma^2}{2} \Delta G_t - \langle a_t^\theta, u_{t|0,1}^* \rangle \right) p_{t|0,1} \pi_{0,1} dx dt
\end{aligned} \tag{49}$$

At this point, we take cases for the different values of g_t . **First:** $\text{sgn}(a_t^\top \nabla G_t) = \text{sgn}(u_{t|0,1}^{\star\top} \nabla G_t) > 0$, and $g(a_t^\top \nabla G_t) = 1$:

$$\begin{aligned}\mathbb{E}_{\pi_{0,1}}[\mathcal{L}^\theta - \mathcal{L}^*] &= \int_0^1 \int_{\mathbb{R}^d} \left(\frac{1}{2} \|a_t^\theta - \nabla G_t\|^2 + \frac{1}{2} \|u_{t|0,1}^*\|^2 + u_{t|0,1}^{\star\top} \nabla G_t - \langle a_t^\theta, u_{t|0,1}^* \rangle \right) p_t dx dt \\ &= \int_0^1 \int_{\mathbb{R}^d} \left(\frac{1}{2} \|a_t^\theta - \nabla G_t\|^2 + \frac{1}{2} \|u_{t|0,1}^*\|^2 - \langle u_{t|0,1}^*, a_t^\theta - \nabla G_t \rangle \right) p_t dx dt \\ &= \int_0^1 \int_{\mathbb{R}^d} \left(\frac{1}{2} \|a_t^\theta - \nabla G_t - u_{t|0,1}^*\|^2 \right) p_t dx dt\end{aligned}\tag{50}$$

Second: $\text{sgn}(a_t^\top \nabla G_t) = \text{sgn}(u_{t|0,1}^{\star\top} \nabla G_t) < 0$, and $g(a_t^\top \nabla G_t) = -1$:

$$\begin{aligned}\mathbb{E}_{\pi_{0,1}}[\mathcal{L}^\theta - \mathcal{L}^*] &= \int_0^1 \int_{\mathbb{R}^d} \left(\frac{1}{2} \|a_t^\theta + \nabla G_t\|^2 + \frac{1}{2} \|u_{t|0,1}^*\|^2 - u_{t|0,1}^{\star\top} \nabla G_t - \langle a_t^\theta, u_{t|0,1}^* \rangle \right) p_t dx dt \\ &= \int_0^1 \int_{\mathbb{R}^d} \left(\frac{1}{2} \|a_t^\theta + \nabla G_t\|^2 + \frac{1}{2} \|u_{t|0,1}^*\|^2 - \langle u_{t|0,1}^*, a_t^\theta + \nabla G_t \rangle \right) p_t dx dt \\ &= \int_0^1 \int_{\mathbb{R}^d} \left(\frac{1}{2} \|a_t^\theta + \nabla G_t - u_{t|0,1}^*\|^2 \right) p_t dx dt\end{aligned}\tag{51}$$

Third: $(u_{t|0,1}^{\star\top} \nabla G_t) = 0$, , and $g(a_t^\top \nabla G_t) = \lambda$, with $\lambda \in [-1, 1]$:

$$\begin{aligned}\mathbb{E}_{\pi_{0,1}}[\mathcal{L}^\theta - \mathcal{L}^*] &= \int_0^1 \int_{\mathbb{R}^d} \left(\frac{1}{2} \|a_t^\theta - \lambda \nabla G_t\|^2 + \frac{1}{2} \|u_{t|0,1}^*\|^2 - \langle a_t^\theta, u_{t|0,1}^* \rangle \right) p_t dx dt \\ &= \int_0^1 \int_{\mathbb{R}^d} \left(\frac{1}{2} \|a_t^\theta - \lambda \nabla G_t\|^2 + \frac{1}{2} \|u_{t|0,1}^*\|^2 + \lambda u_{t|0,1}^{\star\top} \nabla G_t - \langle a_t^\theta, u_{t|0,1}^* \rangle \right) p_t dx dt \\ &= \int_0^1 \int_{\mathbb{R}^d} \left(\frac{1}{2} \|a_t^\theta - \nabla G_t\|^2 + \frac{1}{2} \|u_{t|0,1}^*\|^2 - \langle u_{t|0,1}^*, a_t^\theta + \lambda \nabla G_t \rangle \right) p_t dx dt \\ &= \int_0^1 \int_{\mathbb{R}^d} \left(\frac{1}{2} \|a_t^\theta - \lambda \nabla G_t - u_{t|0,1}^*\|^2 \right) p_t dx dt\end{aligned}\tag{52}$$

where the second equality comes from the fact that we added the quantity $\lambda u_{t|0,1}^{\star\top} \nabla G_t = 0$. Finally, we can succinctly express the optimality gap combining Eq. (50), and 51:

$$\min_{\theta} \mathbb{E}_{\pi_{0,1}}[\mathcal{L}^\theta - \mathcal{L}^*] = \min_{\theta} \int_0^1 \int_{\mathbb{R}^d} \left(\frac{1}{2} \|a_t^\theta - g(a_t^\theta, \nabla G_t) \nabla G_t - u_{t|0,1}^*\|^2 \right) p_t dx dt\tag{53}$$

But notice that by the definition of u_t^θ in Eq. (39), we have $u_t^\theta = a_t^\theta - g(a_t^\theta, \nabla G_{t|0}) \nabla G_{t|0}$, hence the minimization problem of the variational gap in Eq. (53) is expressed as

$$\min_{\theta} \int_0^1 \int_{\mathbb{R}^d} \frac{1}{2} \|u_{t|0,1}^* - u_t^\theta\|^2 p_t dx dt\tag{54}$$

□

B DISCUSSIONS

B.1 GAUSSIAN PATH APPROXIMATION

Finally, following recent advancements in matching frameworks (Liu et al., 2024; Albergo et al., 2023), the conditioned probability path is approximated in a simulation-free manner, as a Gaussian path

$$p_{t|0,1} \approx \mathcal{N}(I_{t|0,1}, \sigma_t \mathbf{I}_d)\tag{55}$$

, where the process is given from

$$X_t = I_{t|0,1} + \sigma_t Z$$

Algorithm 2 Spline Optimization

```

1: Input:  $x_0, x_1, \{X_{t_k}\}$  where  $0 < t_1 < \dots < t_K < 1$ 
2: Initialize  $I_t, \sigma_t$ 
3: repeat
4:   Sample  $X_t \sim (I_t, \sigma_t^2 \mathbf{I}_d)$  with
5:   Evaluate the objective  $G(X_t, x_0)$ 
6:   Estimate the conditional drift  $u_t(X_t|x_0, x_1)$  Eq. (11)
7:   Take gradient step w.r.t control pts.  $\{X_{t_k}, \sigma_{t_k}\}$ 
8: until converges
9: Return  $p_{t|0,1}$  parametrized by optimized  $I_t, \sigma_t$ 

```

, where $I_{t|0,1} \equiv I(t, x_0, x_1)$ is the interpolant function between the pinned endpoints, and $Z \sim \mathcal{N}(0, \mathbf{I}_d)$ is the Brownian motion. evaluate I_t , and σ_t through the following equations for the vector field and the score function (Särkkä & Solin, 2019)

$$\partial_t X_t = \partial I_t + \frac{\partial_t \sigma_t}{\sigma_t} (X_t - I_t), \quad \nabla \log p_t(X_t) = -\frac{1}{\sigma_t^2} (X_t - I_t)$$

Lastly, using these expressions, we obtain the following expression for the conditional drift

$$u_t(X_t|x_0, x_1) = \partial_t I_t + \frac{\partial_t \sigma_t}{\sigma_t} (X_t I_t) - \frac{\nu^2}{2\sigma_t^2} (X_t I_t)$$

where ν is given by the definition of $\sigma_t = \nu \sqrt{t(1-t)}$. We can see that substitution of $I_t := (1-t)x_0 + tx_1$ and $\sigma_t := \nu \sqrt{t(1-t)}$, to the conditional drift expression indeed yields the desired drift of the Brownian bridge $\frac{x_1 - X_t}{1-t}$.

Spline optimization. Efficient optimization is achieved by parameterizing I_t, σ_t respectively as d - and 1-D splines using several knots (i.e., control points) sampled sparsely and uniformly along the time steps $0 < t_1 < \dots < t_K < 1$: $X_{t_k} \in \mathbb{R}^d$ and $\sigma_{t_k} \in \mathbb{R}$. Among those knots, we fit the time-dependent interpolant function I_t , and the standard deviation σ_t

$$I_t := \text{Spline}(t; x_0, \{X_{t_k}\}, x_1) \quad \sigma_t := \text{Spline}(t; \sigma_0=0, \{\sigma_{t_k}\}, \sigma_1=0). \quad (56)$$

Notice that the parameterization in Eq. (56) satisfies the boundary in Eq. (55), hence remains as a feasible solution to Eq. (10). The number of control points K is much smaller than the discretization steps ($K \leq 30$ for all experiments). This significantly reduces the memory complexity, compared to prior methods caching entire discretized SDEs (e.g. (De Bortoli et al., 2021; Chen et al., 2021)).

We follow the algorithm proposed in (Liu et al., 2024), shown in Alg. 2, which, crucially, involves no simulation of an SDE. This is attributed to the utilization of independent samples from $p_{t|0,1}$, which are known in closed form. Finally, recall that since we solve Eq. (10) for each pair $(x_0, x_1) \sim p_{0,1}^\theta$ and later marginalize to construct p_t .

C ADDITIONAL DETAILS ON THE EXPERIMENTS

In this section, we provide further information about the experiments run to validate the efficacy of our FSBM.

Baselines We compare the efficacy of our FSBM in the crowd navigation tasks against the state-of-the-art matching framework GSBM (Liu et al., 2024), with the default hyperparameters. For the opinion depolarization task, we compare against the GSBM, and the DeepGSB (Liu et al., 2022a) also run with their default hyperparameters. For the DeepGSB, we adopt the ‘‘actor-critic’’ parameterization which yields better performance. Finally, in the image translation experiments, we compare against with DSBM (Shi et al., 2023), and Light and Optimal SBM (LOSBM) (Korotin et al., 2023). In our experiments, DSBM was implemented by considering trivial guidance returning the analytic solution of non-guided Brownian bridges. All other frameworks were run using their official implementation, and default hyperparameters. All methods, including our FSBM, are implemented in PyTorch (Paszke et al., 2019).

Guidance Function Recall that our analysis begins with the static entropic OT problem in Eq. (3), to which an additional regularization term is included to capture the interaction among unpaired samples and aligned samples, thereby guiding the transport map. Recall the definition for the Guidance Function

$$G(X_0, X_t) = \alpha \cdot (d(X_t, x_t^i) - d(X_0, x_0^i))^2 \quad (57)$$

where for our experiments, the selected distance between the non-coupled samples and the fixed KP sample was the L_1 norm $d(X_t, x_t^i) = \|X_t - x_t^i\|_{L_1}$. The i^{th} aligned data point was the closest keypoint sample to non-aligned particle X_t at time $t = 0$. To compute the guidance function, we need the trajectories of the KP samples. These are assumed to be given only through their samples and fixed for each KP pair. In practice in the next sections, we discuss how we acquired those trajectories for each of our experiments.

Network Architectures For all experiments, we adopted the same architectures from GSBM, which consists of 4 to 5 residual blocks with sinusoidal time embedding. All networks are trained from scratch, without utilizing any pretrained checkpoint, and optimized with AdamW (Loshchilov, 2017)

Forward and Backward Scheme Training of our FSBM entails a ‘forward-backward’ scheme originally proposed in DSBM (Shi et al., 2023). This necessitates the parameterization of two drifts, one for the forward SDE and another for the backward. During odd epochs, we simulate the coupling from the forward drift, solve the corresponding conditional lagrangian problem in Eq. (10), and then match the resulting p_t with the backward drift. On the other hand, during even epochs, we perform the reverse procedure, matching the forward drift with the marginal p_t derived from the backward drift. As discussed in Liu et al. (2024), this alternating scheme improves the performance of the matching algorithm, averting the matching framework from bias accumulation, as the forward drift always matches the terminal distribution π_1 , and the backward drift is ensured to match the source distribution π_0 .

C.1 CROWD NAVIGATION

We revisit the efficacy of our FSBM in solving complex crowd navigation tasks and its superior capability to generalize under a variety of initial conditions. More explicitly, the two compared frameworks were trained on transporting the samples $X_0 \sim \pi_0$ to $X_1 \sim \pi_1$. In the evaluation phase, the task is always to transport the samples to the same target distribution as in the training phase, however, the initial conditions change throughout the different tasks. For the first task (Vanilla), we draw new samples from the same distribution $\pi_0 = \mathcal{N}(\mu_0, \sigma_0)$, for the second task we shift the mean of the distribution, for the third task we increase the standard deviation, and in the fourth task we sample from a uniform distribution. The initial conditions for each task are the following:

- Vanilla: $\mathcal{N}([-11, -1], 0.5)$
- Perturbed Mean: $\mathcal{N}([-11, -4], 0.5)$
- Perturbed STD: $\mathcal{N}([-11, -1], 3)$
- Uniform Distribution with geometric mean $[-11, -1]$

where $\mathcal{N}([\mu_1, \mu_2], \sigma^2)$ represents a Gaussian with mean the point $[\mu_1, \mu_2]$, and standard deviation equal to σ . Note that for the S-tunnel crowd navigation, we used 100 KPs, which is approximately 4%, of the entire population, and in the V-neck task only 20 KP samples were employed, which is approximately 1% of the population. In our experiments, for the trajectories of the aligned data, we utilized the POT python package (Flamary et al., 2021), to acquire optimal pairings, and GSBM for one epoch to obtain the trajectories between the endpoints. Figures 10, and 11 illustrate the trajectories of the coupled samples, that guide the matching of the unpaired particles. Finally, Figure 13 presents additional plots of FSBM in the crowd navigation tasks examined under the variety of initial conditions mentioned above.

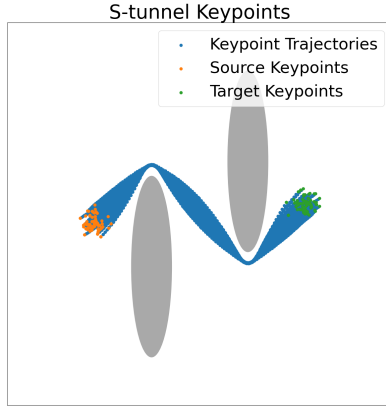


Figure 10: Interpolating path of the KP samples for the S-tunnel

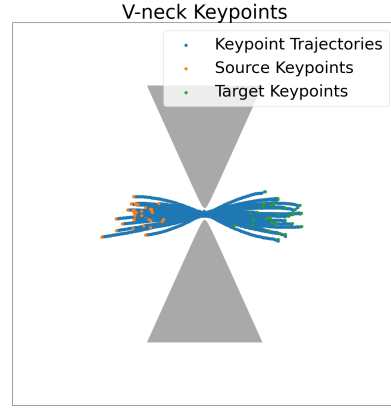


Figure 11: Interpolating path of the KP samples for the V-neck

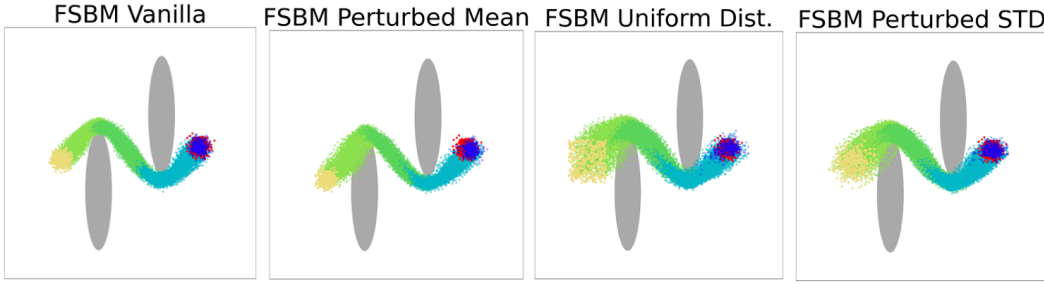


Figure 12: Additional Figures on S-tunnel using our FSBM. The color of particles means: **i) Yellow:** initial conditions, **ii) Green and Cyan:** intermediate trajectory, **iii) Red:** target distribution, **iv) Navy Blue:** generated distribution

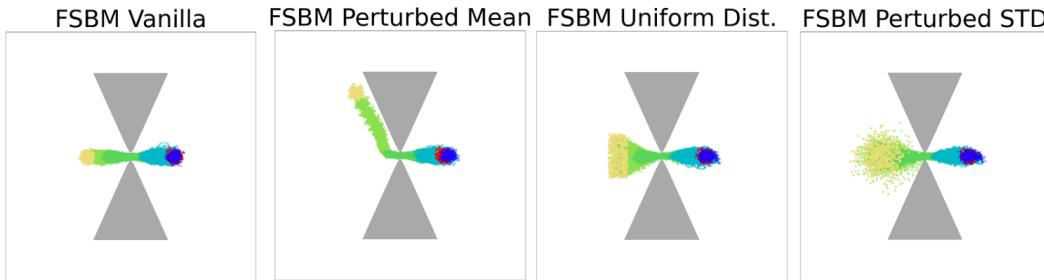


Figure 13: Additional Figures on V-neck using our FSBM. The color of particles means: **i) Yellow:** initial conditions, **ii) Green and Cyan:** intermediate trajectory, **iii) Red:** target distribution, **iv) Navy Blue:** generated distribution

C.2 OPINION DYNAMICS

Subsequently, we revisit the high-dimensional opinion depolarization, of a population evolves under polarizing dynamics which tend to segregate them into two groups of diametrically opposing opinions.

Polarization drift We use the same polarization drift from DeepGSB and GSBM Liu et al. (2022a; 2024), based on the party model (Gaitonde et al., 2021). At each time step t , all agents receive the same random information $\xi_t \in \mathbb{R}^d$ sampled independently of p_t , then react to this information according to

$$f(x; p_t, \xi_t) := \mathbb{E}_{y \sim p_t}[a(x, y, \xi_t)\bar{y}], \quad a(x, y, \xi_t) := \begin{cases} 1 & \text{if } \text{sign}(\langle x, \xi_t \rangle) = \text{sign}(\langle y, \xi_t \rangle) \\ -1 & \text{otherwise} \end{cases}, \quad (58)$$

where x, y are opinions and $\bar{y} := \frac{y}{\|y\|^{\frac{1}{2}}}$ and $a(x, y, \xi_t)$ is the agreement function indicating whether the two opinions agree on the information ξ_t . This suggests that the agents tend to reject opinions, with which they disagree, while they accept opinions closer to theirs. As the dynamics of this model evolve

Recall that the size of the key-point set utilized for the opinion depolarization was approximately 2%, of the entire population. In our opinion depolarization experiments, the optimal pairings were obtained through the POT python package, however, in contrast to the crowd navigation tasks, in this case, the trajectories of the opinion guide particles were obtained by linear interpolating between initial and target positions.



Woman to Man

Man to Woman

Old to Young

Young to Old

Figure 14: Additional Translation Examples

C.3 IMAGE TRANSLATIONS

Finally, we review the experiments in the two types of image translations: **i)** Gender translation, **ii)** Age translation. We follow the setup of (Gushchin et al. (2024)) with the pre-trained ALAE autoencoder (Pidhorskyi et al. (2020)) on 1024×1024 FFHQ dataset (Karras et al. (2019)) to perform the translation in the latent space of dimension 512×1 enable more efficient training and sampling.

Recall that the number of aligned images was 4% of the total number of images for the gender translation and 8% for the age translation. In our translation experiments, the aligned pairings, used to guide the rest of the matching algorithm, were obtained by employing a lightweight SB algorithm using Gaussian mixture parameterization (Korotin et al., 2023). Theoretically, this algorithm provides guarantees to solve the SB, and empirically the quality of the images generated in the training set of this algorithm was deemed satisfactory to constitute our KP set. Finally, Figure 14 presents additional translation instances using our FSBM.