

# 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 CONTACT WASSERSTEIN GEODESICS FOR NON-CONSERVATIVE SCHRÖDINGER BRIDGES

Anonymous authors

Paper under double-blind review

## ABSTRACT

The Schrödinger Bridge provides a principled framework for modeling stochastic processes between distributions; however, existing methods are limited by energy-conservation assumptions, which constrains the bridge’s shape preventing it from model varying-energy phenomena. To overcome this, we introduce the *non-conservative generalized Schrödinger bridge* (NCGSB), a novel, energy-varying reformulation based on contact Hamiltonian mechanics. By allowing energy to change over time, the NCGSB provides a broader class of real-world stochastic processes, capturing richer and more faithful intermediate dynamics. By parameterizing the Wasserstein manifold, we lift the bridge problem to a tractable geodesic computation in a finite-dimensional space. Unlike computationally expensive iterative solutions, our *contact Wasserstein geodesic* (CWG) is naturally implemented via a ResNet architecture and relies on a non-iterative solver with near-linear complexity. Furthermore, CWG supports guided generation by modulating a task-specific distance metric. We validate our framework on tasks including manifold navigation, molecular dynamics predictions, and image generation, demonstrating its practical benefits and versatility.

Project website: <https://sites.google.com/view/c-w-g>

## 1 INTRODUCTION

Inferring the stochastic process that most likely generates a set of sparse observations is a fundamental challenge, e.g., in cellular dynamics (Yeo et al., 2021; Zhang et al., 2024; Moon et al., 2019), meteorology (Franzke et al., 2015), and economics (Kazakevičius et al., 2021; Huang et al., 2024). Here, the target is not merely the distributions of observed data, but rather the underlying dynamics of cell populations, weather patterns, or economic phenomena, enabling reconstruction of missing intermediate states and predicting the systems’ future evolution.

The *Schrödinger Bridge* (SB, Schrödinger (1931)) is a powerful mathematical framework to address this. SB seeks the most likely stochastic path between marginals (i.e., observations), while being close to a reference process, typically Brownian motion. This offers a general stochastic optimal control perspective that encompasses both *Optimal Transport* (OT, Vargas et al. (2021)) and generative approaches such as diffusion models (Ho et al., 2020; Chen et al., 2024), which can be interpreted as optimal bridges with a Gaussian initial marginal (Bortoli et al., 2021). Unfortunately, current SB solvers operate on an infinite probability space and rely on iterative forward–backward stochastic simulations or progressive refinement of the reference dynamics. This leads to complex and costly optimizations, limiting adoption.

Solutions provided by the SB preserve the distribution’s energy throughout the full stochastic path. Here, energy is understood as a combination of *kinetic energy*, which reflects how fast samples move across the probability manifold, and *potential energy* from the underlying landscape. This energy preservation constrains the shape of the bridge and excludes stochastic paths with varying energy profiles, such as dissipative behaviors commonly encountered in real-world physical systems, e.g., storms gradually losing intensity in weather forecasting.

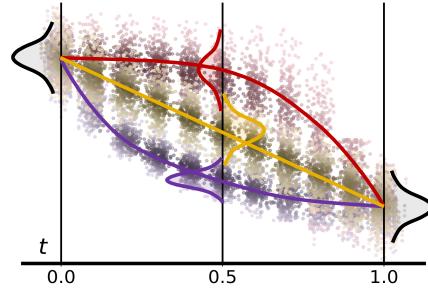


Figure 1: Probability paths obtained under energy-conserving (—), energy-decreasing (—), and energy-increasing (—) (details in App. F.2). Energy variation increases modeling flexibility in applications where distributions at intermediate time steps are of interest.

054 **This paper** overcomes the energy-preserving limitation and the iterative optimization schemes  
 055 of existing SB approaches by formulating a novel geometric generalization of the SB to model  
 056 non-conservative systems and by proposing a near-linear time solver. Specifically, we build on  
 057 the geometric perspective of the SB, which casts it as a flow governed by Hamiltonian dynamics  
 058 on the Wasserstein probability space (Sec. 3). Extending the Hamiltonian system to a contact  
 059 Hamiltonian (Zadra, 2023), we propose a more general energy-varying formulation of the SB prob-  
 060 lem: **The Non-Conservative Generalized SB** (NCGSB, Sec. 4). To make computations tractable,  
 061 we introduce the **Contact Wasserstein Geodesic** (CWG), which solves the NCGSB problem by  
 062 casting it as a geodesic computation, reshaping its cost functional into a Riemannian metric whose  
 063 induced distance is minimized. Discretizing the geodesic leads to geodesic segments that match  
 064 standard residual blocks. We show how this leads to an efficient solver that avoids outer iteration  
 065 loops and achieves near-linear complexity in both dimensionality and batch size. Additionally,  
 066 our approach allows for guided generation, steering the learned process via a task-dependent loss,  
 067 which in our geometric framework corresponds to modulating the Riemannian metric (Sec. 5). We  
 068 demonstrate our approach on benchmarks and tasks such as LiDAR manifold navigation, molecular  
 069 dynamics predictions, and image-based reconstruction of systems such as sea-surface temperature  
 and robotic pick-and-place (Sec. 6).

070 In summary, **we contribute**: (1) A novel non-conservative formulation of the Schrödinger Bridge  
 071 problem that models a wider range of real-world physical stochastic processes by building on  
 072 geometric flows governed by contact Hamiltonian dynamics; (2) The introduction of the Contact  
 073 Wasserstein Geodesic (CWG) framework, a general geometric solver compatible with all  
 074 Schrödinger Bridge variants, enabling fast and scalable computation; (3) a guided generation  
 075 method for steering the bridge learned by CWG according to new task specifications.

## 2 RELATED WORK

079 **Schrödinger bridges.** The SB problem imposes no constraints on the probabilistic path be-  
 080 yond matching the endpoint marginals, limiting its applicability when intermediate observations  
 081 exist. It also does not permit including known physical laws governing the system’s dynamics.  
 082 The *multi-marginal Schrödinger Bridge* (mmSB) treats intermediate observations (Theodoropoulos  
 083 et al., 2025) as constraints, which enables reconstruction of continuous dynamics without piecewise  
 084 approximations. In contrast, the *Generalized Schrödinger Bridge* (GSB) adds a state cost, allowing  
 085 potential energy functionals to be minimized along the probability path. This let us to model mean-  
 086 field interactions (Gaitonde et al., 2021; Ruthotto et al., 2020), conservative forces (Philippidis et al.,  
 087 1979; Noé et al., 2020), or geometric priors (Chen & Lipman, 2024; Liu et al., 2018).

088 **Non-conservative Schrödinger bridge formulations.** The SB problem assumes constant en-  
 089 ergy, preventing it from modeling non-conservative systems. The *momentum SB* augments the state  
 090 space with velocity (Theodoropoulos et al., 2025; Chen et al., 2023), allowing damping to be in-  
 091 corporated (Blessing et al., 2025; Sterling et al., 2025). This, however, doubles the state space  
 092 and increases computational cost. Other extensions replace the Brownian reference process with  
 093 the *Ornstein–Uhlenbeck (OU)* process (Orland, 2025; Zhang & Stumpf, 2025), introducing a non-  
 094 conservative prior for the target dynamics. The OU process defines a mean distribution with a  
 095 curl-free component that drives convergence and a divergence-free component that induces rotation.  
 096 Yet, this formulation lacks a mechanism for energy dissipation in the rotational dynamics. **We in-**  
 097 **introduce** a more general energy-varying framework in which dissipation naturally emerges across all  
 098 components of the dynamics, while only requiring a scalar augmentation of the state space.

099 **Schrödinger bridge solvers.** Matching-based iterative approaches (Bortoli et al., 2024; Shi et al.,  
 100 2023; Gushchin et al., 2024) have gained popularity by improving the scalability and robustness of  
 101 traditional *Iterative Proportional Fitting (IPF)* algorithms (Kullback, 1968; Léonard, 2013) through  
 102 Markovian projections, thereby avoiding the need for full trajectory storage. However, for GSB, re-  
 103 strictive assumptions like Gaussian probability paths limit expressivity (Liu et al., 2024; Tang et al.,  
 104 2025), while for mmSB, global trajectory consistency is compromised due to the piecewise nature  
 105 of the approach and its sensitivity to the initial choice of the reference process (Shen et al., 2025). To  
 106 overcome these limitations, the stochastic dynamics can be learned indirectly by leveraging the ana-  
 107 lytical optimality conditions of the SB problem (Vargas et al., 2021; Chen et al., 2022). This has been  
 108 shown to scale more effectively to both the GSB (Liu et al., 2022; Buzun et al., 2025) and the mmSB  
 109 problem (Theodoropoulos et al., 2025; Hong et al., 2025; Chen et al., 2023), due to the symmetries

108 inherent in these optimality conditions (App. A). However, these methods are limited by a classical  
 109 SB solver bottleneck: the computational overhead of their iterative nature, which alternates between  
 110 forward–backward passes or repeated dynamics integration and reference updates. **We propose** a  
 111 cheaper non-iterative solver that scales nearly linearly with both dimensionality and sample size.

112 **Schrödinger bridge guided generation.** The SB framework naturally extends to conditional  
 113 generation, where the marginals and the transition path depend on additional parameters or  
 114 objectives (Shi et al., 2022). Model guidance techniques (Song et al., 2023; Guo et al., 2024)  
 115 introduce a guidance term into the stochastic process, typically derived from the gradient of a  
 116 loss function. This approach steers the flow locally and resembles a gradient-based form of  
 117 optimal control. Alternatively, Raja et al. (2025) employs a global optimal control perspective.  
 118 Their approach generates full trajectories and chooses the one minimizing a task-specific action  
 119 functional. However, their method is deterministic and yields a single optimal path rather than a  
 120 posterior distribution over paths. Unlike previous guidance methods, **we propose** a hybrid approach  
 121 to guided generation within the NCGSB framework. By embedding a task-specific loss into the  
 122 potential, we reshape the Riemannian metric so that the resulting geodesic reflects the guidance  
 123 objective, allowing the learned dynamics to align with the desired outcome.

### 124 3 PRELIMINARIES

125 We introduce the necessary notation throughout the following sections, and summarize it in App. B.

126 **The Wasserstein manifold.** Before formally introducing the SB problem, we define its domain.  
 127 Let  $\mathcal{P}^+(\mathcal{M})$  denote the space of smooth, positive density functions supported on a manifold  $\mathcal{M}$ .  
 128 Each element  $\rho \in \mathcal{P}^+(\mathcal{M})$  is a function  $\rho(x) : \mathcal{M} \rightarrow \mathbb{R}^+$  satisfying  $\int_{\mathcal{M}} \rho(x) dx = 1$ . The  
 129 infinitesimal variation of the density at time  $t$  is the time derivative  $\partial_t \rho^t(x)$ , which lies in the tangent space  
 130  $\mathcal{T}_{\rho^t} \mathcal{P}^+(\mathcal{M})$ . The collection of all such tangent spaces forms the tangent bundle  $\mathcal{T}\mathcal{P}^+(\mathcal{M})$ . When  
 131 equipped with the Wasserstein metric,  $\mathcal{P}^+(\mathcal{M})$  becomes a Riemannian manifold (Ambrosio et al.,  
 132 2005). The corresponding metric tensor is defined as,

$$133 g^{\mathcal{W}_2}(\partial_t \rho^t, \partial_t \rho^t) = \int_{\mathcal{M}} \partial_t \rho^t(x) (-\Delta_{\rho^t})^\dagger \partial_t \rho^t(x) \rho^t(x) dx, \quad (1)$$

134 where  $(-\Delta_{\rho^t})^\dagger$  is the inverse of the weighted Laplacian operator  $\Delta_{\rho^t} = -\nabla_x \cdot (\rho^t \nabla_x)$  (Chow et al.,  
 135 2020), inducing an inner product on  $\mathcal{T}\mathcal{P}^+(\mathcal{M})$  and a distance  $d_{\mathcal{W}_2}(\rho_a, \rho_b)$  for  $\rho_a, \rho_b \in \mathcal{P}^+(\mathcal{M})$ .  
 136 The minimum-length curve  $\rho^t$  connecting the two distributions  $\rho_a, \rho_b$  is called a *geodesic*.

137 **Multi-marginal generalized Schrödinger bridge (mmGSB).** Given two endpoint densities  
 138  $\rho_a, \rho_b \in \mathcal{P}^+(\mathcal{M})$ , the SB problem (Schrödinger, 1931; Schrödinger, 1932) seeks the most probable  
 139 interpolating density path  $\rho^t$ . This minimizes the Kullback–Leibler divergence w.r.t. a reference  
 140 process  $\rho_{\text{ref}}^t$ , typically Brownian motion. The SB problem is equivalent to a stochastic optimal control  
 141 setting (Dai Pra, 1991), which minimizes the cost required to transport a set of diffusing particles  
 142 from an initial distribution  $\rho_a$  to a target distribution  $\rho_b$ . This dynamic reformulation of the OT  
 143 problem (Benamou & Brenier, 2000; Chen et al., 2014) has solutions corresponding to geodesics on  
 144 the Wasserstein manifold  $\mathcal{P}^+(\mathcal{M})$ . These trajectories are straight, since the classical SB problem  
 145 assumes that particles dynamics are unaffected by external potential functions  $U$ . This assumption,  
 146 however, limits our ability to model complex real-world physical systems.

147 Intermediate observations represented by marginal distributions  $\{\rho_m\}_{m=1}^M$  at specific time steps  
 148  $\{t_m\}_{m=1}^M$ , can also be incorporated as additional constraints (Chen et al., 2023; Tang et al., 2025).  
 149 This leads to the mmGSB problem,

$$150 \min_{v^t} J(v^t, \rho^t) = \int_0^1 \int_{\mathcal{M}} \left( \frac{1}{2} \|v^t(x)\|^2 + U(x) \right) \rho^t(x) dx dt; \quad (2a)$$

$$151 \text{s.t. } \partial_t \rho^t(x) + \nabla_x \cdot (\rho^t(x) v^t(x)) = \varepsilon \Delta_x \rho^t(x); \quad (2b)$$

$$152 \rho^0 = \rho_a, \rho^1 = \rho_b, \rho^{t_m} = \rho_m, \forall m \in 1, \dots, M. \quad (2c)$$

153 Here, the density evolution  $\rho^t$  is governed by the Fokker–Planck equation (2b), which generalizes  
 154 Brownian motion by incorporating a deterministic drift term  $v^t$  alongside a stochastic diffusion  
 155 term scaled by  $\varepsilon$ , to satisfy the boundary conditions in equation (2c). This drift  $v^t$  acts as the control

variable and ensures that the probability path interpolates between the given boundary marginals. The objective functional represents the kinetic energy associated with the drift and quantifies the deviation from the (uncontrolled) reference stochastic process.

**Wasserstein Hamiltonian flows and geodesics.** A convenient solution to the mmGSB problem (2) is to specify analytical optimality conditions (Sec. 2). These take the form of a Wasserstein Hamiltonian Flow (Chow et al., 2020), which describes a probability distribution evolving according to Hamiltonian dynamics. This evolution lies on planes tangent to  $\mathcal{P}^+(\mathcal{M})$ , specifically on the cotangent bundle  $\mathcal{T}^*\mathcal{P}^+(\mathcal{M})$ , the dual of  $\mathcal{T}\mathcal{P}^+(\mathcal{M})$ , and it is governed by the derivatives of a scalar Hamiltonian function  $H$ . However, their integration remains computationally expensive (Buzun et al., 2025; Wu et al., 2025), so we propose a geometric reformulation that results in a significant simplification. To this end, we introduce Proposition 1, a standard result from differential geometry (App. C.2), which is instrumental in lifting these equations to geodesics on  $\mathcal{P}^+(\mathcal{M})$ .

**Proposition 1.** *Let the optimality conditions of the mmGSB problem (2) be expressed in Hamiltonian form, yielding the optimal bridge  $\rho^t(x)$ . Then,  $\rho^t(x)$  can be viewed as a geodesic connecting the marginals in equation 2c w.r.t. the modified Riemannian metric  $g_J$ , known as the Jacobi metric (Abraham & Marsden, 2008).*

To access the Jacobi metric and determine the corresponding geodesic, we first derive the Hamiltonian optimality conditions of the mmGSB problem (2) using Lagrange multipliers (Cui et al., 2024). This introduces a potential function  $S^t(x)$ , whose gradient defines the drift via  $v^t(x) = \nabla_x S^t(x)$ . The potential enforces the dynamic constraint (2b) within the cost functional (2a), whose first variation yields the Hamiltonian optimality conditions,

$$\partial_t \rho^t(x) = \partial_S H(\cdot) = -\nabla_x \cdot (\rho^t(x) \nabla_x S^t(x)); \quad (3a)$$

$$\partial_t S^t(x) = -\partial_\rho H(\cdot) = -\frac{1}{2} \|\nabla_x S^t(x)\|^2 + \frac{1}{2} \varepsilon^2 \partial_\rho I(\rho^t(x)) + U(x), \quad (3b)$$

with the corresponding Hamiltonian,  $H(\rho^t, S^t) = \mathcal{K}(\rho^t, S^t) + \mathcal{F}(\rho^t)$ , defined as the sum of a kinetic energy  $\mathcal{K}$  and a potential energy  $\mathcal{F} = -U - I$ , dependent on the potential function  $U$  and the Fisher information  $I$ . A detailed derivation of these dynamics and the full Hamiltonian function is in App. D.1. To handle boundary conditions (2c), the Hamiltonian dynamics (3) are typically integrated backward in time, where the solution at each intermediate point  $(\rho^{t_m}, S^{t_m})$  serves as the initial condition for the next segment (Theodoropoulos et al., 2025). The potential function  $S^t(x)$  in equation (3a) is linked to the infinitesimal density variation  $\partial_t \rho^t \in \mathcal{T}\mathcal{P}^+(\mathcal{M})$  via the weighted Laplacian operator  $\Delta_{\rho^t}$ , introduced through the Wasserstein metric (1). This connection establishes a correspondence between the tangent bundle  $\mathcal{T}\mathcal{P}^+(\mathcal{M})$  and the cotangent bundle  $\mathcal{T}^*\mathcal{P}^+(\mathcal{M})$ , where  $S^t(x)$  naturally resides, and where the Hamiltonian dynamics of  $(\rho^t, S^t)$  unfold.

By Proposition 1, the Hamiltonian dynamics (3) corresponds to a geodesic flow on the underlying Wasserstein manifold  $\mathcal{P}^+(\mathcal{M})$ , which minimizes the Jacobi metric  $g_J = (H - \mathcal{F}) g^{\mathcal{W}_2}$ . The original metric  $g^{\mathcal{W}_2}$  (1) accounts only for the kinetic energy of the transport map  $\mathcal{K}$  by measuring distances between distributions. In contrast, the Jacobi metric  $g_J$  also includes the potential energy  $\mathcal{F}$ , which is maximized to attain values  $\mathcal{F} \approx H$ . Consequently, computing the geodesic between marginals under this metric is equivalent to solving the mmGSB problem (2).

## 4 THE NON-CONSERVATIVE GENERALIZED SCHRÖDINGER BRIDGE

**Non-conservative formulation.** The solution to the GSB problem (2) assumes a constant energy function  $H$ , and restricts the drift  $v^t$  to depend solely on the potential energy  $\mathcal{F}$ . This limits the model's flexibility in representing dynamics that cannot be described by a conservative potential, which reduces its ability to capture real-world processes involving energy dissipation and external interactions. To overcome this, we introduce the *non-conservative generalized Schrödinger bridge* (NCGSB), which allows for time-varying energy systems. To do so, we reformulate the cost functional  $J$  as the time integral of a new scalar state  $\mathcal{z}^t$ , representing the *Lagrangian action*, whose evolution depends recursively on itself. The NCGSB problem is formulated as follows,

$$\min_{v^t} J(v^t, \rho^t) = \int_0^1 \partial_t \mathcal{z}^t dt; \quad (4a)$$

$$\text{s.t. } \partial_t \mathbf{z}^t = \int_{\mathcal{M}} \left( \frac{1}{2} \|v^t(x)\|^2 + U(x) \right) \rho^t(x) dx - \gamma \mathbf{z}^t; \quad (4b)$$

$$\partial_t \rho^t(x) + \nabla_x \cdot (\rho^t(x) v^t(x)) = \varepsilon \Delta_x \rho^t(x); \quad (4c)$$

$$\rho^0 = \rho_a, \rho^1 = \rho_b, \rho^{t_m} = \rho_m, \forall m \in 1, \dots, M, \quad (4d)$$

where  $\gamma \in \mathbb{R}$  is a damping factor. The objective in equation 4a is no longer to minimize a static quantity, but rather a time-varying state  $z^t$ . Its dynamics (4b) depend explicitly on its current value. This recursive structure endows the system with a form of memory, as its evolution is influenced by the entire trajectory, implicitly encoded in  $z^t$ . Because non-conservative forces are path-dependent, augmenting the system's state space with the scalar  $z^t$  allows their effects to be modeled, enabling the system's energy to vary over time. The sign and magnitude of  $\gamma$  determine the direction and rate of this variation. By relaxing the implicit energy-conservation constraint of the GSB problem, our approach enhances the model's flexibility and improves the quality of the resulting optimal solution.

**Guided NCGSB.** NCGSB (4) can be extended to the guided generation setting by introducing a guiding function  $f$ , which steers the generative process toward desired conditions at any chosen time (Song et al., 2023; Guo et al., 2024). For a given time  $t_s$ , the guidance is expressed as  $y = f(x^{t_s})$ , with  $x^{t_s} \sim \rho^{t_s}$ . To enforce this form, the bridge  $\rho^t$  is steered according to  $\rho^t(x|y) = \frac{1}{Z} \rho^t(x) e^{-\|y - f(x^{t_s})\|^2}$ , where  $Z$  is a normalization constant and  $x^{t_s}$  denotes a sample from the predicted guided distribution  $\rho^{t_s}(x|y)$ . By Bayes' rule, the dynamics of the guided bridge  $\rho^t(x|y)$  acquire an additional guidance term via the drift  $v^t$ , determined by  $\|y - f(x^{t_s})\|^2$ . To perform a guided generation that enforces the constraint  $y = f(x^{t_s})$  while preserving the underlying data manifold, we incorporate  $y$  into the Lagrangian action constraint (4b) as (see App. D.3 for details),

$$\partial_t z^t = \int_{\mathcal{M}} \left( \frac{1}{2} \|v^t(x)\|^2 + U(x) + \|y - f(x^{t_s})\|^2 \right) \rho^t(x) dx - \gamma z^t. \quad (5)$$

**Wasserstein contact Hamiltonian flows and geodesics.** Analogously to mmGSB (2), understanding the dynamics of the optimality conditions in NCGSB (4) is essential for reformulating it as a geodesic computation. As detailed in App. D.2, we propose to leverage the contact Hamiltonian formalism (Kholodenko, 2013), an extension of classical Hamiltonian mechanics to non-conservative systems (App. C.1), to model the dynamics of the NCGSB optimality conditions as Wasserstein contact Hamiltonian flows. This generalizes Prop. 1, since the contact Hamiltonian dynamics defines a geodesic but on the extended space  $\mathcal{P}^+(\mathcal{M}) \times \mathbb{R}$  (Udriște, 2000; Testa et al., 2025). The contact Hamiltonian optimality conditions are,

$$\partial_t \rho^t(x) = \partial_S H(\cdot) = -\nabla_x \cdot (\rho^t(x) \nabla_x S^t(x)), \quad (6a)$$

$$\partial_t S^t(x) = -\partial_\rho H(\cdot) - S^t(x) \partial_z H(\cdot) = -\frac{1}{2} \|\nabla_x S^t(x)\|^2 + \frac{1}{2} \varepsilon^2 \partial_\rho I(\rho^t(x)) + U(x) - \gamma S^t(x) - \varepsilon \gamma \log \rho^t(x), \quad (6b)$$

$$\partial_t \mathcal{Z}^t = S^t(x) \partial_S H(\cdot) - H(\cdot) = \int_{\mathcal{M}} \left( \frac{1}{2} \|\nabla_x S^t(x)\|^2 + U(x) \right) \rho^t(x) dx + \frac{1}{2} \varepsilon^2 I(\rho^t) - \int_{\mathcal{M}} \varepsilon \gamma (\log \rho^t(x) - 1) \rho^t(x) dx - \gamma \mathcal{Z}^t, \quad (6c)$$

The corresponding contact Hamiltonian function is defined as,  $H(\rho^t, S^t, z^t) = \mathcal{K}(\rho^t, S^t) + \mathcal{F}(\rho^t) + \mathcal{B}(\rho^t) + \gamma z^t$ . This differs from its conservative counterpart in two ways. First, its explicit dependence on  $z^t$  allows the total energy to vary over time. Second, the potential energy is augmented by an entropy term,  $\mathcal{B}(\rho^t) = \int_{\mathcal{M}} \varepsilon(\log \rho^t(x) - 1) \rho^t(x) dx$ , producing an additional diffusion in the dynamics. As previously mentioned, for guided generation, an additional potential energy term  $\|y - f(x^t)\|^2$  can here be introduced to steer the flow. Geometrically, the dynamics of  $(\rho^t, S^t, z^t)$  can be interpreted as a flow on the cotangent bundle of the Wasserstein manifold, augmented by the scalar state  $z^t$ . That is, the dynamics unfold on the space  $\mathcal{T}^* \mathcal{P}^+(\mathcal{M}) \times \mathbb{R}$ .

The contact Hamiltonian flow evolving on the extended phase space  $\mathcal{T}^*\mathcal{P}^+(\mathcal{M}) \times \mathbb{R}$ , and interpolating between the marginal densities, induces a geodesic on the augmented manifold  $\mathcal{P}^+(\mathcal{M}) \times \mathbb{R}$ . This geodesic minimizes a Jacobi metric  $\tilde{g}_1 = (H - \mathcal{F} - \mathcal{B})^{q\mathcal{W}_2}$ , which generalizes the classical

270 Wasserstein metric by incorporating the potential energy of the contact Hamiltonian function. Computing  
 271 the geodesic under the Jacobi metric  $\tilde{g}_J$  corresponds to NCGSB (4). Unlike the conservative  
 272 case, the contact Hamiltonian  $H$  is no longer constant along the flow, allowing the total energy to  
 273 vary over time. This introduces an additional degree of freedom that can be leveraged to shape the  
 274 system's energy along the path over  $\mathcal{P}^+(\mathcal{M})$ . This is the reason that the geodesic  $(\rho^t, H^t)$  is defined  
 275 on the extended space  $\mathcal{P}^+(\mathcal{M}) \times \mathbb{R}$ .

## 277 5 CONTACT WASSERSTEIN GEODESICS (CWG)

279 **ResNet resembles a discrete geodesic.** Our objective is to compute a geodesic  $\rho^t$  on  $\mathcal{P}^+(\mathcal{M})$ ,  
 280 induced by the contact Hamiltonian dynamics, that is constrained to pass through a set of observed  
 281 marginals  $\{\rho_a, \rho_m, \rho_b\}$  (i.e., discretized distributions along the probability path). These constraints  
 282 naturally lead to a discretized parameterization of  $\rho^t$ , where the overall density transformation is  
 283 modeled as a composition of maps, each connecting a pair of consecutive observations. A ResNet  
 284 is ideally suited for this problem, as its sequential block structure directly mirrors this piecewise,  
 285 compositional nature of the approximated geodesic. Let  $\lambda$  be a fixed reference measure on  $\mathcal{P}^+(\mathcal{M})$   
 (e.g., a standard Gaussian or uniform distribution). We define a  $(K + 1)$ -block ResNet as follows,

$$286 \quad T_{\{\theta^k\}_{k=0}^K} = T_{\theta^K} \circ \cdots \circ T_{\theta^1} \circ T_{\theta^0}, \quad (7)$$

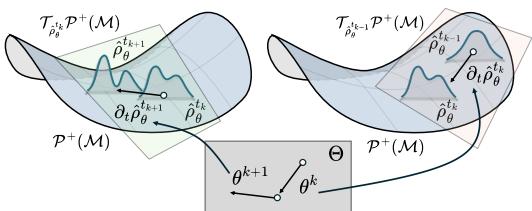
288 with parameters  $\{\theta^k\}_{k=0}^K \in \Theta$ . The process begins by sampling an initial batch of points  
 289  $x^s \sim \lambda$ , that is pushed forward through the first block to obtain  $x^{t_0} = T_{\theta^0}(x^s)$ . Then, the  
 290 parameters  $\theta^0$  are optimized such that the resulting pushforward reference measure approximates  
 291 the initial marginal  $\rho_\theta^{t_0} \approx \rho_a$ . Thereafter, each subsequent block  $k$  pushes forward the sample via  
 292  $x^{t_{k+1}} = T_{\theta^{k+1}}(x^{t_k})$ ,  $x^{t_k} \sim \rho_\theta^{t_k}$ . The full pushforward map induces,  
 293

$$294 \quad \rho_\theta^{t_{k+1}} = (T_{\theta^{k+1}})_\# \rho_\theta^{t_k} = \rho_\theta^{t_k} (T_{\theta^{k+1}}^{-1}(x^{k+1})) \det[\nabla_x T_{\theta^{k+1}}^{-1}(x^{k+1})]. \quad (8)$$

295 Starting from the reference measure  $\lambda$ , the ResNet parameters  $\{\theta^k\}_{k=0}^K$  define a sequence of  
 296 discrete probability transitions  $\{\partial_t \rho_\theta^{t_k}\}_{k=0}^K$ , which in turn specify the discrete family of densities  
 297  $\{\rho_\theta^{t_k}\}_{k=0}^K$ . Geometrically, the discretizations  $\{\partial_t \rho_\theta^{t_k}, \rho_\theta^{t_k}\}_{k=0}^K$ , provided by the ResNet, approximate  
 298  $(\rho^t, \partial_t \rho^t) \in \mathcal{TP}^+(\mathcal{M})$ , which can be seen as inducing a mapping from  $\mathcal{TP}^+(\mathcal{M})$  onto the  
 299 parameter space  $\Theta$  (Fig. 5). As stated in Proposition 2, the existence of such a map endows  
 300 the finite-dimensional space  $\Sigma$ , where the parameterized densities  $\rho_\theta^{t_k}$  reside, with the geometric  
 301 properties of  $\mathcal{P}^+(\mathcal{M})$ . This lifting enables faster and tractable computations for SB problems.  
 302

303 **Proposition 2.** *Approximate the evolution of the density  $\rho^t \in \mathcal{P}^+(\mathcal{M})$  by a series of  $K$  smooth  
 304 parametrized pushforwards  $T_{\theta^k}$ , with  $\theta^k$  belonging to a finite-dimensional space  $\Theta$ . If each push-  
 305 forward  $T_{\theta^k}$  is an immersion  $T_{\theta^k} : \Theta \rightarrow \mathcal{TP}^+(\mathcal{M})$ , then the parameter space  $\Theta$  can be endowed  
 306 with a Riemannian structure via the pullback of the Wasserstein metric  $g^{\mathcal{W}_2}$ . Consequently, the  
 307 contact Hamiltonian dynamics on  $\mathcal{TP}^+(\mathcal{M}) \times \mathbb{R}$  can be represented in the reduced phase space  
 308  $\Theta^* \times \mathbb{R}$ , with the associated geodesic on  $\mathcal{P}^+(\mathcal{M}) \times \mathbb{R}$  projected onto  $\Sigma \times \mathbb{R}$  (see App. E.1).*

309 Proposition 2 allows us to transform the geodesic computation from the infinite-dimensional  
 310  $\mathcal{P}^+(\mathcal{M})$  to a geodesic on a finite-dimensional parameterized space  $\Sigma$ , such that the resulting  
 311 geodesic flow on  $\Sigma \times \mathbb{R}$  evolves under the pullback of the Jacobi metric  $T_\theta \tilde{g}_J = \Phi^{t_k} T_\theta^* g^{\mathcal{W}_2}$ ,  
 312 where the scalar factor  $\Phi^{t_k} = H(\rho^{t_k}, S^{t_k}, z^{t_k}) - \mathcal{F}(\rho^{t_k}) - \mathcal{B}(\rho^{t_k})$ , encodes the potential energy.  
 313 Specifying the time evolution of  $H^{t_k}$  determines a unique parameterized bridge on  $\Sigma$ . This  
 314 formulation enables a tractable computation of geodesic flows to solve the NCGSB problem.  
 315 Although different parameterizations  $\{\theta^k\}_{k=0}^K$  may define distinct coordinate systems on  $\Sigma$ , the  
 316 geodesics solutions remain equivalent and share the same energy (Syrota et al., 2025).  
 317



318 **Figure 2:** Visualization of the ResNet trans-  
 319 formation. Two successive pushforwards  
 320  $\rho_\theta^{t_{k-1}} \rightarrow \rho_\theta^{t_k} \rightarrow \rho_\theta^{t_{k+1}}$  on  $\mathcal{P}^+(\mathcal{M})$  are shown as  
 321 local updates  $\partial_t \rho_\theta^{t_k}, \partial_t \rho_\theta^{t_{k+1}}$  on tangent spaces.  
 322 Each update is parameterized by  $\theta^k, \theta^{k+1} \in \Theta$ ,  
 323 defining local coordinates on  $\mathcal{TP}^+(\mathcal{M})$ . This  
 324 coordinate system is not unique.

324 **Geodesic computation.** The *contact Wasserstein geodesic* (CWG) corresponds to the discrete  
 325 path  $\{\rho_\theta^{t_k}\}_{k=0}^K$ , that approximates the NCGSB solution. This is trained to reconstruct the avail-  
 326 able marginals while minimizing the geodesic energy under the pullback Jacobi metric  $T_\theta^* \tilde{g}_J =$   
 327  $\Phi^{t_k} T_\theta^* g^{\mathcal{W}_2}$ . The initial and final marginals,  $\rho_a$  and  $\rho_b$ , are enforced at the path endpoints, corre-  
 328 sponding to the ResNet outputs at times  $t_0$  and  $t_K$ . When available, intermediate marginals  $\rho_m$  must  
 329 appear at time points matching the ResNet discretization for the condition to be enforced.

330 CWG training happens in two stages (Alg. 1 in App. E.2): (1) we optimize the first ResNet block to  
 331 match the initial marginal  $\rho_a$ , and (2) we find the optimal path by minimizing the loss,

$$\ell = \underbrace{d_{\mathcal{W}_2}^2(\rho_\theta^{t_K}, \rho_b)}_{\text{Terminal marginal}} + \underbrace{\sum_{m=1}^M d_{\mathcal{W}_2}^2(\rho_\theta^{t_{k_m}}, \rho_m)}_{\text{Intermediate marginals}} + \underbrace{\sum_{k=1}^K \Phi^{t_k} d_{\mathcal{W}_2}^2(\rho_\theta^{t_k}, \rho_\theta^{t_{k-1}})}_{\text{Energy minimization}}. \quad (9)$$

337 Here  $d_{\mathcal{W}_2}$  denotes the Wasserstein-2 distance between probability distributions. In practice, this  
 338 distance is approximated using empirical estimators based on samples drawn from the distributions.  
 339 Convergence of the algorithm is guaranteed in App. E.3. Its complexity is  $\mathcal{O}(NK(T_{\text{sh}} + D(LW +$   
 340  $\log N)))$ , scaling linearly in dimension  $D$  and nearly linearly in batch size  $N$  (see App. E.5), rather  
 341 than exponentially or quadratically (Hong et al., 2025). Unlike Chen et al. (2023); Shen et al. (2025),  
 342 our CWG avoids costly iteration loops and is only weakly affected by the number of marginals.

343 **Guided contact Wasserstein geodesics.** In the conditional setting, the Lagrangian action dy-  
 344 namics (4b) in NCGSB (4) is augmented as in equation 5. Here, the scaling factor  $\Phi^{t_k} =$   
 345  $H(\rho^{t_k}, S^{t_k}, z^{t_k}) - \mathcal{F}(\rho^{t_k}) - \mathcal{B}(\rho^{t_k})$  of the pullback Jacobi metric is augmented with the guid-  
 346 ance term  $\|y - f(x^{t_s})\|^2$ , to enforce the constraint  $y = f(x^{t_s})$  at time  $t_s$  of the generative process.  
 347 Under the ResNet parameterization, the desired distribution is approximated by  $x^{t_{k_s}} \approx x^{t_s}$  at time  
 348 step  $t_{k_s}$ , and the Jacobi metric is modified as  $\tilde{g}_J = (\Phi^{t_k} + \|y - f(x^{t_s})\|^2) g^{\mathcal{W}_2}$ , with  $x^{t_{k_s}} \sim \rho_\theta^{t_{k_s}}$ .  
 349 This penalizes geodesics crossing undesired regions at  $t_{k_s}$ . The loss for the guided optimization is

$$\ell = d_{\mathcal{W}_2}^2(\rho_\theta^{t_K}, \rho_b) + \sum_{m=1}^M d_{\mathcal{W}_2}^2(\rho_\theta^{t_{k_m}}, \rho_m) + \sum_{k=1}^K (\Phi^{t_k} + \|y - f(x^{t_s})\|^2) d_{\mathcal{W}_2}^2(\rho_\theta^{t_k}, \rho_\theta^{t_{k-1}}) + d_{\mathcal{W}_2}^{t_s}(\rho_\theta^{t_{k_s}}, \rho_s) \quad (10)$$

354 where the modified distance  $d_{\mathcal{W}_2}^{t_s}$  measures deviations between the generated distribution  $\rho_\theta^{t_{k_s}}$   
 355 and the intermediate marginal  $\rho_s$  at  $t_s$ , while incorporating the penalty for samples  $x_s$  that violate the  
 356 guidance constraint  $y = f(x_s)$ , c.f. App. F.1. In practice, this loss is optimized through a fine-tuning  
 357 procedure applied to a model initially trained without any guidance.

358 **Proof of concept.** We demonstrate our framework on a 2D distribution-matching task and guided  
 359 generation setting using the Two-Moons and Checkerboard benchmarks (Holderrieth & Erives,  
 360 2025). These lack intermediate marginals  $\{\rho_m\}_{m=1}^M$ , and the initial distribution  $\rho_a$  coincides with  
 361 the reference distribution  $\lambda$ . Hence, only the second step of Alg. 1 is needed. Figure 3 shows that our  
 362 method successfully generates the target distributions, and steers the generation to samples confined  
 363 to a subset of the target space (here, the upper half). This guided behavior is achieved via the term  
 364  $\|y - f(x^{t_s})\|^2$ , with  $t_s = 1$ ,  $f$  measuring 2D sample positions, and  $y$  defining the admissible region.

## 366 6 RESULTS

367 We benchmark our approach against four established baselines summarized in Table 1. Further  
 368 details of the experimental setups are provided in App. G.1.

370 Method	371 GSB	372 mmSB	373 Energy variation	374 Image Gen.	375 Guided Gen.
371 DSBM (Shi et al., 2023)	✗	✗	✗	✓	✗
372 SB-Flow (Bortoli et al., 2024)	✗	✗	✗	✓	✗
373 GSBM (Liu et al., 2024)	✓	✗	✗	✓	✗
374 SBIR (Shen et al., 2025)	✗	✓	✗	✗	✗
375 DM-SB (Chen et al., 2023)	✗	✓	✗	✗	✗
376 CWG (ours)	✓	✓	✓	✓	✓

377 Table 1: Comparison of our CWG with baselines designed to address various SB variants and types of problems.

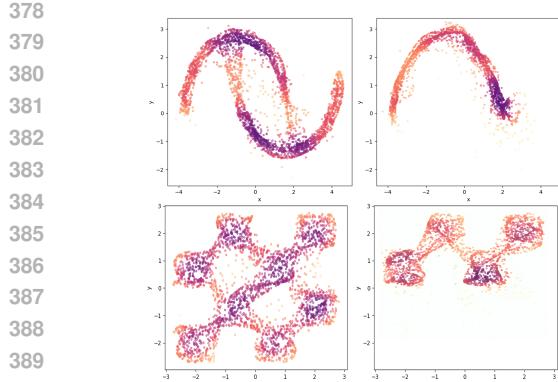


Figure 3: Two-Moons (top) and Checkerboard (bottom) benchmarks with guided variants (right).

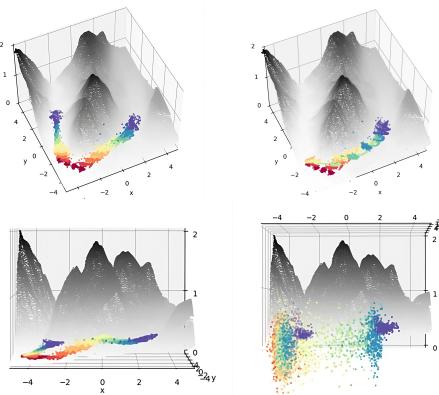


Figure 4: LiDAR Manifold Navigation: CWG before and after guidance (top), CWG vs GSBM (bottom).

**LiDAR manifold navigation.** First, we tackle a standard GSB task: computing a bridge evolving on a geometric manifold. We use the LiDAR scan of Mount Rainier (OpenTopography, 2025) as the reference surface, and we aim to connect two marginals while remaining on the manifold and favoring low-altitude regions. These conditions are encoded into the potential function  $U$  (see App. G.2). In this experiment, we do not model a physical system but instead compute an optimal transport map between the marginals under a conservative setting. Unlike our approach, DSBM and GSBM iteratively fit a deterministic path between the marginals, falling short on representing a posterior distribution. As a result, there is no guarantee that a Gaussian path remains on the manifold (Fig. 4). This leads to substantially higher-energy paths, as reflected by the Optimality metric in Table 2, and to inaccurate modeling of the target distribution, as indicated by the Feasibility metric (both metrics as defined in the baseline Liu et al. (2024)). In contrast, *our approach finds lower-energy solutions that respect the marginal constraints and converges significantly faster*. Furthermore, our method uniquely supports guided generation, illustrated here by steering the probabilistic path to the right side of the mount (Fig. 4; see App. G.2 for quantitative results).

Table 2: Optimality (67a) (↓), Feasibility (67b) (↓), and training time (tt) (↓) in LiDAR Navigation.

Metric	CWG (ours)	GSBM	DSBM
Optimality	<b>1.40<math>\pm</math>0.02</b>	2.18 $\pm$ 0.02	4.16 $\pm$ 0.01
Feasibility	<b>0.06<math>\pm</math>0.01</b>	0.83 $\pm$ 0.01	0.97 $\pm$ 0.01
tt (s)	<b>280<math>\pm</math>20</b>	1570 $\pm$ 50	1340 $\pm$ 50

Table 3: Wasserstein error at validation (↓) and training time (tt) (↓) in Single Cell Sequencing.

Metric	CWG (ours)	DM-SB	SBIR
$d_{\mathcal{W}_2}(x^{t_1})$	<b>1.11<math>\pm</math>0.06</b>	2.25 $\pm$ 0.01	1.92 $\pm$ 0.02
$d_{\mathcal{W}_2}(x^{t_3})$	<b>0.33<math>\pm</math>0.02</b>	1.64 $\pm$ 0.03	1.86 $\pm$ 0.02
tt (s)	<b>710<math>\pm</math>30</b>	38120 $\pm$ 1100	1740 $\pm$ 40

**Single cell sequencing.** Next, we reconstruct stem cell differentiation dynamics from a series of isolated cellular snapshots. We use the Embryoid Body (EB) dataset from Moon et al. (2019), which tracks cell state progression across five developmental stages  $[t_0, t_1, t_2, t_3, t_4]$ . Cell differentiation is fundamentally a non-conservative biological process (Zeevaert et al., 2020; Kinney et al., 2014) and the ability to model energy-varying bridges is essential. To evaluate generalization in regions with no available data, we split the dataset into a training set  $[t_0, t_2, t_4]$  and a validation set  $[t_1, t_3]$ . Accordingly, the former contains the distributions  $\{\rho_a, \rho_{m_2}, \rho_b\}$ , while the latter contains  $\{\rho_{m_1}, \rho_{m_3}\}$ . The geometry of the training distributions is encoded in the potential function  $U$ , which penalizes paths that stray from the observed data manifold. Minimizing  $U$  ensures the learned bridge remains close to the data manifold, enabling effective generalization. The combination of the data manifold guidance and an energy-varying bridge allows our approach to outperform other mmSB baselines in both reconstruction accuracy and computation time. Quantitative results are reported in Table 3, with additional details and an ablation study on the importance of energy variation provided in App. G.3.

**Image generation.** We also demonstrate our framework’s applicability to image generation tasks. Given sequences of images capturing the time evolution of physical phenomena, the model aims, from a single input image, to predict the most likely terminal state of the system, along with realistic intermediate frames at unseen time steps. Specifically, we use the NOAA OISST v2 High Resolution Dataset (Huang et al., 2021), which provides daily sea surface temperature averages over multiple years, and the BridgeData V2 (Walke et al., 2023) dataset, containing image snapshots of robotic manipulation tasks. Our NCGSB problem can also be applied to unpaired image-matching tasks using general image datasets that do not describe dynamical systems, such as MNIST, EMNIST, and

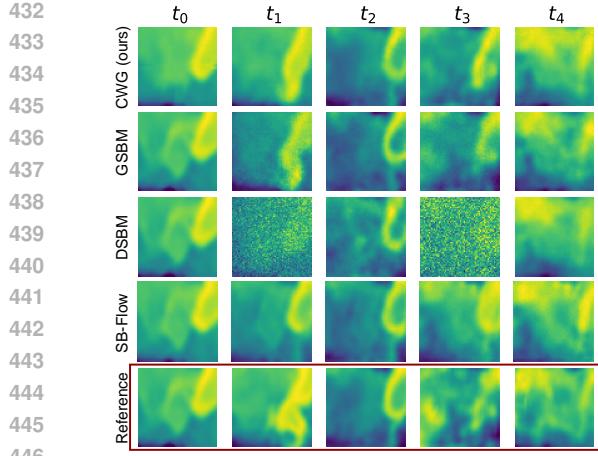


Figure 5: Predictions generated by CWG (ours, top), GSBM (second), DSBM (third), and SB-Flow (fourth). The red row displays the corresponding reference samples.

Table 4: FID scores at validation steps ( $\downarrow$ ) and training time (tt) ( $\downarrow$ ) in Sea Prediction (2020-2024).

Metric	CWG (ours)	GSBM	DSBM	SB-Flow
$\text{FID}(x^{t_1})$	$121 \pm 6$	$161 \pm 5$	$242 \pm 10$	$177 \pm 4$
$\text{FID}(x^{t_3})$	$160 \pm 7$	$186 \pm 7$	$236 \pm 10$	$190 \pm 7$
tt (60s)	$17 \pm 1$	$1227 \pm 54$	$318 \pm 15$	$83 \pm 5$

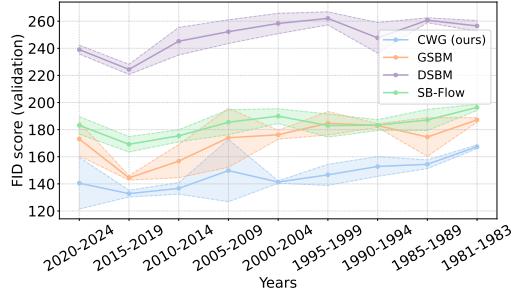


Figure 6: FID scores at validation time steps for four methods, evaluated over 9 tests (1981–2024). Our CWG scores are significantly lower than baselines.

**Flickr-Faces-HQ (FFHQ)** (Karras et al., 2019). For the high-resolution unpaired image-matching task on the FFHQ dataset, the SB is computed in a latent space after encoding the images with the pre-trained ALAE encoder (Pidhorskyi et al., 2020). For all the other datasets, we compute the SB directly in the ambient image space. To ensure that the generated frames remain faithful to the underlying data distribution, we introduce a potential function  $U$ , that penalizes deviations from the learned data manifold. Following the baseline (Liu et al., 2024), for samples  $x^t \sim \rho^t$ ,  $U(x^t)$  is defined as the reconstruction error, obtained via a VAE (Song & Itti, 2025). Details on the energy behavior and extended results are provided in Apps. G.4, G.5, G.6, and G.7.

For the sea temperature prediction task, we cluster data from 1981–2024 into five-year intervals. Using heatmaps from January, May, and September (i.e.,  $[t_0, t_2, t_4]$ ), our method predicts the temperature profiles of March and July (i.e.,  $[t_1, t_3]$ ). Our CWG produces cleaner, more accurate predictions than the baselines (Fig. 5). Since our framework operates efficiently in probability space and is not constrained by energy conservation, it achieves these results with an order of magnitude less computation (Table 4 for 2020-2024; Fig. 6 for all years). In the robotic task reconstruction, our model

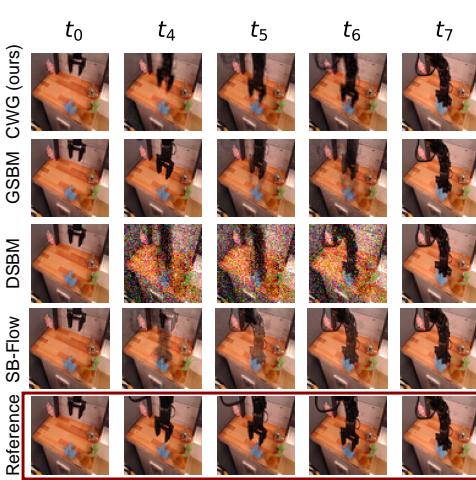


Figure 7: Reconstructions from CWG (top), GSBM (middle), and DSBM (bottom). Red row shows the reference.

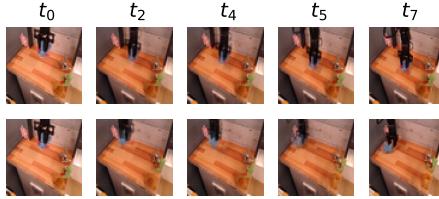


Figure 8: CWG outputs before (top) vs. after guidance (place the item left).

Metric	CWG(ours)	GSBM	DSBM	SB-Flow
$\text{FID}$	$19 \pm 1$	$40 \pm 2$	$150 \pm 1$	$73 \pm 1$
tt (h)	$0.5 \pm 0$	$25.3 \pm 2.2$	$7.6 \pm 0.7$	$1.4 \pm 0.1$

Table 5: FID score ( $\downarrow$ ) and training time (tt) ( $\downarrow$ ) in the robotic task reconstruction.

Metric	Standard	Guidance
Centroid	$35.8 \pm 11.1$	$22.3 \pm 2.9$
FID	$19.52 \pm 0.78$	$23.77 \pm 1.94$

Table 6: Item centroid position (px) and FID before vs. after guidance.

generates realistic intermediate frames connecting the initial and final states of a robot’s reaching motion (Fig. 7), and demonstrates consistently robust performance, outperforming baselines in image quality (Table 8). Also, Fig. 8 shows guided generation, where our model successfully steers the

486 placing motion task toward a target location on the left side of the table. This is achieved with only  
 487 a minimal drop in image quality, maintaining a clear advantage over competing methods (Table 6).  
 488

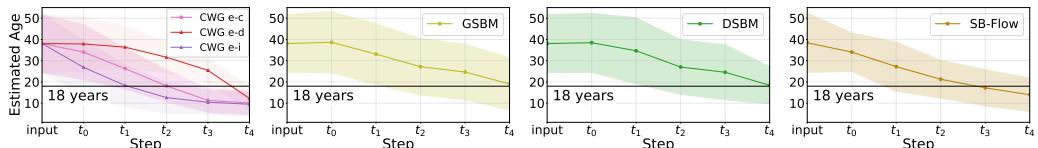
489 Regarding the unpaired image transfer experiments, we first evaluate the MNIST-to-EMNIST  
 490 transfer (Table 38, App. G.7) and validate the guided-bridge capabilities through a Gaussian  
 491 deblurring task. We then present the FFHQ transfer experiment, in which the SB maps images from  
 492 a starting *adult* distribution  $\rho_a$  to its closest analogue in a target *children* distribution  $\rho_b$ . As shown  
 493 in Table 7, while GSBM and DSBM produce images visually similar to the input (Optimality), they  
 494 fail to satisfy the boundary conditions of (2) and achieve poor Feasibility. In contrast, our CWG  
 495 model consistently generates images within the children distribution  $\rho_b$ , with average predicted  
 496 ages below 18, as illustrated in Fig. 10.

497 Table 7: Metrics for the FFHQ transfer experiment: Training time (tt) (↓), Optimality (72a) (↓), measuring  
 498 the geodesic distance to assess the transport cost between the two marginals ( $\rho_a, \rho_b$ ), and Feasibility (72b) (↓),  
 499 indicating how well the marginals are preserved, i.e., how closely the bridge endpoint aligns with  $\rho_b$ . Metric  
 < 18 (↓) indicates confidence that the final predicted images satisfy the boundary constraints.

Metric	CWG	GSBM	DSBM	SB-Flow
Optimality	$218.1 \pm 2.6$	$206.2 \pm 2.9$	<b><math>198.0 \pm 2.7</math></b>	$237.4 \pm 2.2$
Feasibility	<b><math>4.332 \pm 0.526</math></b>	$6.839 \pm 0.753$	$7.780 \pm 0.801$	$21.748 \pm 0.498$
< 18 (p-value)	$2.1 \times 10^{-9}$	$6.6 \times 10^{-1}$	$6.1 \times 10^{-1}$	$1.1 \times 10^{-2}$
tt (s)	<b><math>930 \pm 30</math></b>	$2650 \pm 30$	$2530 \pm 30$	$1490 \pm 30$



505  
 506 Input  
 507  
 508 CWG(ours)  
 509  
 510 GSBM  
 511  
 512 DSBM  
 513  
 514 SB - Flow  
 515  
 516  
 517  
 518  
 519  
 520  
 521 Figure 9: Adult → Child image generation on the FFHQ transfer experiment.



522  
 523  
 524  
 525  
 526 Figure 10: Average age predictions at each generation time step for the images shown in Figure 9, comparing  
 527 the baseline models with three CWG variants: energy-conserving, energy-increasing, and energy-decreasing.  
 528

## 529 7 CONCLUSION

530 Our work is motivated by the need to model intermediate time steps of Schrödinger bridges (SBs),  
 531 arising from the underlying dynamics of the observed physical system. As standard SBs conserve  
 532 energy across time, they cannot meaningfully encode such dynamics. To counter this, we introduced  
 533 the *non-conservative generalized Schrödinger bridge* (NCGSB), which extends the usual Hamiltonian  
 534 to its non-conservative counterpart, the *contact Hamiltonian*, allowing energy to vary. We  
 535 show that NCGSB is equivalent to geodesics on *contact Wasserstein manifolds*. This link leads to a  
 536 non-iterative and near-linear time algorithm for computing the non-conservative bridge, which can  
 537 practically be realized by a ResNet-like construction, easing its implementation. We show that these  
 538 theoretical contributions lead to a SB framework that is not only more expressive but also signifi-  
 539 cantly faster than existing approaches, as validated by the significant improvements achieved across  
 a range of diverse tasks.

540 REFERENCES  
541

542 Ralph Abraham and Jerrold E Marsden. *Foundations of mechanics*. Number 364. American Math-  
543 ematical Soc., 2008.

544 Luigi Ambrosio, Nicola Gigli, and Giuseppe Savaré. *Gradient flows: in metric spaces and in the*  
545 *space of probability measures*. Springer Science & Business Media, 2005. URL <https://doi.org/10.1007/b137080>.

546

547 Georgios Arvanitidis, Lars Kai Hansen, and Søren Hauberg. Latent space oddity: on the curvature  
548 of deep generative models. In *International Conference on Learning Representations*, 2018. URL  
549 <https://openreview.net/forum?id=SJzRZ-WCZ>.

550

551 Jean-David Benamou and Yann Brenier. A computational fluid mechanics solution to the Monge-  
552 Kantorovich mass transfer problem. *Numerische Mathematik*, 84(3):375–393, 2000. URL  
553 <https://doi.org/10.1007/s002110050002>.

554

555 Quentin Bertrand, Joey Bose, Alexandre Duplessis, Marco Jiralerspong, and Gauthier Gidel. On  
556 the stability of iterative retraining of generative models on their own data. In *International*  
557 *Conference on Learning Representations (ICLR)*, 2024. URL <https://openreview.net/forum?id=JORAfH2xFd>.

558

559 Denis Blessing, Julius Berner, Lorenz Richter, and Gerhard Neumann. Underdamped diffusion  
560 bridges with applications to sampling. In *International Conference on Learning Representations*  
561 (*ICLR*), 2025. URL <https://openreview.net/forum?id=Q1QTxFm0Is>.

562

563 Christian Blohmann. Lagrangian field theory. 2024. Lecture Notes.

564

565 Roberto Borelli, Agostino Dovier, and Federico Fogolari. Data structures and algorithms for  
566 k-th nearest neighbours conformational entropy estimation. *Biophysica*, 2(4):340–352, 2022.  
567 ISSN 2673-4125. doi: 10.3390/biophysica2040031. URL <https://www.mdpi.com/2673-4125/2/4/31>.

568

569 Valentin De Bortoli, James Thornton, Jeremy Heng, and Arnaud Doucet. Diffusion Schrödinger  
570 bridge with applications to score-based generative modeling. In *Advances in Neural Information*  
571 *Processing Systems (NeurIPS)*, 2021. URL <https://openreview.net/forum?id=9BnCwiXB0ty>.

572

573 Valentin De Bortoli, Iryna Korshunova, Andriy Mnih, and Arnaud Doucet. Schrodinger bridge flow  
574 for unpaired data translation. In *The Thirty-eighth Annual Conference on Neural Information*  
575 *Processing Systems*, 2024. URL <https://openreview.net/forum?id=1F32iCJFFa>.

576

577 Alessandro Bravetti, Hans Cruz, and Diego Tapias. Contact Hamiltonian mechanics. *Annals of*  
578 *Physics*, 376:17–39, 2017.

579

580 Nazar Buzun, Daniil Shlenskii, Maxim Bobrin, and Dmitry V. Dylov. Hota: Hamiltonian framework  
581 for optimal transport advection, 2025. URL <https://arxiv.org/abs/2507.17513>.

582

583 Minshuo Chen, Song Mei, Jianqing Fan, and Mengdi Wang. An overview of diffusion models: Ap-  
584 plications, guided generation, statistical rates and optimization. *arXiv preprint arXiv:2404.07771*,  
585 2024. URL <https://arxiv.org/abs/2404.07771>.

586

587 Ricky T. Q. Chen and Yaron Lipman. Flow matching on general geometries. In *International*  
588 *Conference on Learning Representations (ICLR)*, 2024. URL <https://openreview.net/forum?id=g7ohD1T1L>.

589

590 Tianrong Chen, Guan-Horng Liu, and Evangelos Theodorou. Likelihood training of Schrödinger  
591 bridge using forward-backward SDEs theory. In *International Conference on Learning Represen-*  
592 *tations (ICLR)*, 2022. URL <https://openreview.net/forum?id=nioAdKCEdXB>.

593

594 Tianrong Chen, Guan-Horng Liu, Molei Tao, and Evangelos Theodorou. Deep momentum  
595 multi-marginal Schrödinger bridge. In *Conference on Neural Information Processing Systems*  
596 (*NeurIPS*), 2023. URL <https://openreview.net/forum?id=ykvvv0gc4R>.

594 Xiwen Chen, Wenhui Zhu, Peijie Qiu, Hao Wang, Huayu Li, ZIHAN LI, Yalin Wang, Aristeidis  
 595 Sotiras, and Abolfazl Razi. FIC-TSC: Learning time series classification with Fisher information  
 596 constraint. In *International Conference on Machine Learning (ICML)*, 2025. URL <https://openreview.net/forum?id=Dr8msCnFYw>.

598 Yongxin Chen, Tryphon Georgiou, and Michele Pavon. On the relation between optimal transport  
 599 and Schrödinger bridges: A stochastic control viewpoint, 2014. URL <https://arxiv.org/abs/1412.4430>.

600 601

602 Shui-Nee Chow, Wuchen Li, and Haomin Zhou. Wasserstein Hamiltonian flows. *Journal of Dif-  
 603 ferential Equations*, 268(3):1205–1219, 2020. URL <https://doi.org/10.1016/j.jde.2019.08.046>.

604 605

606 Gregory Cohen, Saeed Afshar, Jonathan Tapson, and Andre Van Schaik. Emnist: Extending mnist  
 607 to handwritten letters. In *2017 international joint conference on neural networks (IJCNN)*, pp.  
 608 2921–2926. IEEE, 2017.

609 610

611 Jianbo Cui, Shu Liu, and Haomin Zhou. Stochastic Wasserstein Hamiltonian flows. *Journal of  
 Dynamics and Differential Equations*, 36(4):3885–3921, 2024. URL <https://doi.org/10.1007/s10884-023-10264-4>.

612 613

614 Paolo Dai Pra. A stochastic control approach to reciprocal diffusion processes. *Applied math-  
 ematics and Optimization*, 23(1):313–329, 1991. URL <https://doi.org/10.1007/BF01442404>.

615 616

617 M. do Carmo. *Riemannian Geometry*. Birkhäuser Basel, 1992.

618 619

620 Ruiqi Feng, Chenglei Yu, Wenhao Deng, Peiyan Hu, and Tailin Wu. On the guidance of flow  
 matching. In *Forty-second International Conference on Machine Learning*, 2025. URL <https://openreview.net/forum?id=pKaNgFzJBy>.

621 622

623 Jean Feydy. Geometric data analysis, beyond convolutions. *Applied Mathematics*, 3, 2020.

624 625

626 Christian L. E. Franzke, Terence J. O’Kane, Judith Berner, Paul D. Williams, and Valerio Lucarini.  
 627 Stochastic climate theory and modeling. *Wiley Interdisciplinary Reviews: Climate Change*, 6(1):  
 628 63–78, 2015. URL <https://doi.org/10.1002/wcc.318>.

629 630

631 Jason Gaitonde, Jon Kleinberg, and Éva Tardos. Polarization in geometric opinion dynamics. In  
 632 *ACM Conference on Economics and Computation (EC)*, pp. 499–519, 2021. URL <https://doi.org/10.1145/3465456.3467633>.

633 634

635 Hansjörg Geiges. A brief history of contact geometry and topology. *Expositiones Mathe-  
 maticae*, 19(1):25–53, 2001. ISSN 0723-0869. URL [https://doi.org/10.1016/S0723-0869\(01\)80014-1](https://doi.org/10.1016/S0723-0869(01)80014-1).

636 637

638 Hansjörg Geiges. An introduction to contact topology. *Cambridge Studies in Advanced Mathemat-  
 ics*, 109, 2008.

639 640

641 Yingqing Guo, Hui Yuan, Yukang Yang, Minshuo Chen, and Mengdi Wang. Gradient guidance  
 for diffusion models: An optimization perspective. In *Neural Information Processing Systems  
 (NeurIPS)*, 2024. URL <https://openreview.net/forum?id=X1QeUYBXke>.

642 643

644 Nikita Gushchin, Sergei Kholkin, Evgeny Burnaev, and Alexander Korotin. Light and op-  
 645 timal Schrödinger bridge matching. In *International Conference on Machine Learning  
 (ICML)*, pp. 17100–17122. PMLR, 2024. URL <https://openreview.net/forum?id=EWJn6hfZ4J>.

646 647

648 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic mod-  
 649 els. *Advances in neural information processing systems (NeurIPS)*, 33:6840–6851,  
 650 2020. URL [https://proceedings.neurips.cc/paper\\_files/paper/2020/file/4c5bcfec8584af0d967f1ab10179ca4b-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2020/file/4c5bcfec8584af0d967f1ab10179ca4b-Paper.pdf).

651 652

653 Peter Holderrieth and Ezra Erives. Introduction to flow matching and diffusion models, 2025. URL  
 654 <https://diffusion.csail.mit.edu/>.

648 Wanli Hong, Yuliang Shi, and Jonathan Niles-Weed. Trajectory inference with smooth Schrödinger  
 649 bridges. In *International Conference on Machine Learning (ICML)*, 2025. URL <https://openreview.net/forum?id=GHyvvWu1XC>.  
 650

651 Boyin Huang, Chunying Liu, Viva Banzon, Eric Freeman, Garrett Graham, Bill Hankins, Tom  
 652 Smith, and Huai-Min Zhang. Improvements of the daily optimum interpolation sea surface tem-  
 653 perature (doisst) version 2.1. *Journal of Climate*, 34(8):2923–2939, 2021.  
 654

655 Hongbin Huang, Minghua Chen, and Xiao Qiao. Generative learning for financial time series with  
 656 irregular and scale-invariant patterns. In *International Conference on Learning Representations*  
 657 (*ICLR*), 2024. URL <https://openreview.net/forum?id=CdjnzWsQax>.  
 658

659 Oliver Johns. *Analytical Mechanics for Relativity and Quantum Mechanics*. Oxford University  
 660 Press, 07 2005. ISBN 9780198567264. doi: 10.1093/acprof:oso/9780198567264.001.0001. URL  
 661 <https://doi.org/10.1093/acprof:oso/9780198567264.001.0001>.  
 662

663 Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative  
 664 adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern*  
 665 *recognition*, pp. 4401–4410, 2019.  
 666

667 Rytis Kazakevičius, Aleksejus Kononovicius, Bronislovas Kaulakys, and Vygintas Gontis. Under-  
 668 standing the nature of the long-range memory phenomenon in socioeconomic systems. *Entropy*,  
 669 23(9):1125, 2021. URL <https://doi.org/10.3390/e23091125>.  
 670

671 Arkady L Kholodenko. *Applications of contact geometry and topology in physics*. World Scientific,  
 672 2013.  
 673

674 Melissa A Kinney, Rabbia Saeed, and Todd C McDevitt. Mesenchymal morphogenesis of embryonic  
 675 stem cells dynamically modulates the biophysical microtissue niche. *Scientific reports*, 4(1):4290,  
 676 2014. URL <https://doi.org/10.1038/srep04290>.  
 677

678 Solomon Kullback. Probability densities with given marginals. *The Annals of Mathematical Statis-  
 679 tics*, 39(4):1236–1243, 1968. URL <https://www.jstor.org/stable/2239692>.  
 680

681 Flavien Léger and Wuchen Li. Hopf–cole transformation via generalized Schrödinger bridge prob-  
 682 lem. *Journal of Differential Equations*, 274:788–827, 2021. URL <https://doi.org/10.1016/j.jde.2020.10.029>.  
 683

684 David Levin. The approximation power of moving least-squares. *Mathematics of computation*, 67  
 685 (224):1517–1531, 1998.  
 686

687 Guan-Horng Liu, Tianrong Chen, Oswin So, and Evangelos Theodorou. Deep generalized  
 688 Schrödinger bridge. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.  
 689 URL <https://openreview.net/forum?id=fp33Nsh005>.  
 690

691 Guan-Horng Liu, Yaron Lipman, Maximilian Nickel, Brian Karrer, Evangelos Theodorou, and  
 692 Ricky T. Q. Chen. Generalized Schrödinger bridge matching. In *International Conference on*  
 693 *Learning Representations (ICLR)*, 2024. URL <https://openreview.net/forum?id=SoismgeX7z>.  
 694

695 Zhiyu Liu, Bo Wu, and Hai Lin. A mean field game approach to swarming robots control. In  
 696 *American Control Conference (ACC)*, pp. 4293–4298, 2018. URL <https://doi.org/10.23919/ACC.2018.8431807>.  
 697

698 Christian Léonard. A survey of the Schrödinger problem and some of its connections with optimal  
 699 transport, 2013. URL <https://arxiv.org/abs/1308.0215>.  
 700

701 Kevin R Moon, David Van Dijk, Zheng Wang, Scott Gigante, Daniel B Burkhardt, William S Chen,  
 702 Kristina Yim, Antonia van den Elzen, Matthew J Hirn, Ronald R Coifman, et al. Visualizing  
 703 structure and transitions in high-dimensional biological data. *Nature biotechnology*, 37(12):1482–  
 704 1492, 2019. URL <https://doi.org/10.1038/s41587-019-0336-3>.  
 705

702 Frank Noé, Alexandre Tkatchenko, Klaus-Robert Müller, and Cecilia Clementi. Machine learning  
 703 for molecular simulation. *Annual review of physical chemistry*, 71(1):361–390, 2020. URL  
 704 <https://doi.org/10.1146/annurev-physchem-042018-052331>.  
 705

706 OpenTopography. Southwest flank of mt. rainier, wa, 2025. URL <https://doi.org/10.5069/G9PZ56R1>. Distributed by OpenTopography. Accessed 2025-08-11.  
 707

708 Henri Orland. Schrödinger bridges for systems of interacting particles. *Physical Review E*, 112(1):  
 709 014115, 2025.  
 710

711 Stefano Peluchetti. Diffusion bridge mixture transports, Schrödinger bridge problems and generative  
 712 modeling. *Journal of Machine Learning Research (JMLR)*, 24(374):1–51, 2023. URL <https://www.jmlr.org/papers/v24/23-0527.html>.  
 713

714 Chris Philippidis, Chris Dewdney, and Basil J Hiley. Quantum interference and the quantum potential.  
 715 *Il Nuovo Cimento B (1971-1996)*, 52(1):15–28, 1979.  
 716

717 Stanislav Pidhorskyi, Donald A Adjeroh, and Gianfranco Doretto. Adversarial latent autoencoders.  
 718 In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp.  
 719 14104–14113, 2020.  
 720

721 Sanjeev Raja, Martin Šípka, Michael Psenka, Tobias Kreiman, Michal Pavelka, and Aditi S Krishnapriyan. Action-minimization meets generative modeling: Efficient transition path sampling  
 722 with the Onsager-Machlup functional. In *International Conference on Machine Learning (ICML)*,  
 723 2025. URL <https://openreview.net/forum?id=QwoGfQzuMa>.  
 724

725 Lars Ruthotto, Stanley J Osher, Wuchen Li, Levon Nurbekyan, and Samy Wu Fung. A machine  
 726 learning framework for solving high-dimensional mean field game and mean field control problems.  
 727 *Proceedings of the National Academy of Sciences*, 117(17):9183–9193, 2020. URL  
 728 <https://doi.org/10.1073/pnas.1922204117>.  
 729

730 Erwin Schrödinger. Sur la théorie relativiste de l'électron et l'interprétation de la mécanique  
 731 quantique. In *Annales de l'institut Henri Poincaré*, volume 2, pp. 269–310, 1932. URL  
 732 <http://eudml.org/doc/78968>.  
 733

734 Erwin Schrödinger. Über die umkehrung der naturgesetze. *Sitzungsberichte der Preussischen  
 735 Akademie der Wissenschaften, Physikalisch-Mathematische Klasse*, pp. 144–153, 1931. URL  
 736 <https://doi.org/10.1002/ange.19310443014>.  
 737

738 Yunyi Shen, Renato Berlinghieri, and Tamara Broderick. Multi-marginal Schrödinger bridges with  
 739 iterative reference refinement. In *International Conference on Artificial Intelligence and Statistics  
 (AISTATS)*, 2025. URL <https://openreview.net/forum?id=VcwZ3gtYFY>.  
 740

741 Yuyang Shi, Valentin De Bortoli, George Deligiannidis, and Arnaud Doucet. Conditional simulation  
 742 using diffusion Schrödinger bridges. In *Uncertainty in Artificial Intelligence (UAI)*, pp. 1792–  
 743 1802. PMLR, 2022. URL <https://proceedings.mlr.press/v180/shi22a.html>.  
 744

745 Yuyang Shi, Valentin De Bortoli, Andrew Campbell, and Arnaud Doucet. Diffusion Schrödinger  
 746 bridge matching. In *Advances on Neural Information Processing Systems (NeurIPS)*, 2023. URL  
 747 <https://openreview.net/forum?id=qy070HsJT5>.  
 748

749 Ana Cannas da Silva. Lectures on symplectic geometry. *Lecture Notes in Mathematics*, 1764, 2001.  
 750

751 Hae Jin Song and Laurent Itti. Riemannian-geometric fingerprints of generative models. In *2nd  
 752 Beyond Euclidean Workshop: Hyperbolic and Hyperspherical Learning for Computer Vision*,  
 753 2025. URL <https://openreview.net/forum?id=00LViVLqkN>.  
 754

755 Jiaming Song, Qinsheng Zhang, Hongxu Yin, Morteza Mardani, Ming-Yu Liu, Jan Kautz, Yongxin  
 756 Chen, and Arash Vahdat. Loss-guided diffusion models for plug-and-play controllable  
 757 generation. In *International Conference on Machine Learning (ICML)*, 2023. URL <https://proceedings.mlr.press/v202/song23k.html>.  
 758

756 Benjamin Sterling, Chad Gueli, and Mónica F. Bugallo. Critically-damped higher-order Langevin  
 757 dynamics. *arXiv preprint arXiv:2506.21741*, 2025. URL <https://arxiv.org/abs/2506.21741>.

759

760 Jürgen Struckmeier and Andreas Redelbach. Covariant hamiltonian field theory. *International Journal*  
 761 *of Modern Physics E*, 17(03):435–491, 2008.

762

763 Stas Syrota, Yevgen Zainchkovskyy, Johnny Xi, Benjamin Bloem-Reddy, and Søren Hauberg. Identifying  
 764 metric structures of deep latent variable models. In *International Conference on Machine*  
 765 *Learning (ICML)*, 2025. URL <https://openreview.net/forum?id=FHzkr8FRie>.

766

767 Sophia Tang, Yinuo Zhang, Alexander Tong, and Pranam Chatterjee. Branched Schrödinger bridge  
 768 matching. *arXiv preprint arXiv:2506.09007*, 2025. URL <https://arxiv.org/abs/2506.09007>.

769

770 Andrea Testa, Søren Hauberg, Tamim Asfour, and Leonel Rozo. Geometric contact flows: Contacto-  
 771 morphisms for dynamics and control. In *International Conference on Machine Learning (ICML)*,  
 772 2025. URL <https://openreview.net/forum?id=v2nQ1e78Rc>.

773

774 Panagiotis Theodoropoulos, Augustinos D Saravacos, Evangelos A Theodorou, and Guan-Horng  
 775 Liu. Momentum multi-marginal Schrödinger bridge matching. *arXiv preprint arXiv:2506.10168*,  
 2025. URL <https://arxiv.org/abs/2506.10168>.

776

777 Sahand Tokasi and Peter Pickl. Symplectic reduction of classical mechanics on shape space. *Foundations of Physics*, 52(5):106, 2022.

778

779 Alexander Tong, Jessie Huang, Guy Wolf, David Van Dijk, and Smita Krishnaswamy. Trajectorynet:  
 780 A dynamic optimal transport network for modeling cellular dynamics. In *International Conference on Machine Learning (ICML)*, pp. 9526–9536. PMLR, 2020. URL <https://proceedings.mlr.press/v119/tong20a.html>.

781

782

783 Constantin Udriște. Geometric dynamics. *Southeast Asian Bulletin of Mathematics*, 24(2):313–322,  
 2000.

784

785 Francisco Vargas, Pierre Thodoroff, Austen Lamacraft, and Neil Lawrence. Solving Schrödinger  
 786 bridges via maximum likelihood. *Entropy*, 23(9):1134, 2021. URL <https://doi.org/10.3390/e23091134>.

787

788

789 Homer Rich Walke, Kevin Black, Tony Z. Zhao, Quan Vuong, Chongyi Zheng, Philippe Hansen-  
 790 Estruch, Andre Wang He, Vivek Myers, Moo Jin Kim, Max Du, Abraham Lee, Kuan Fang,  
 791 Chelsea Finn, and Sergey Levine. Bridgedata v2: A dataset for robot learning at scale. In *Confer-  
 792 ence on Robot Learning (CoRL)*, 2023. URL <https://openreview.net/forum?id=f55M1AT1Lu>.

793

794

795 Zifan Wang, Alice Harting, Matthieu Barreau, Michael M Zavlanos, and Karl H Johansson. Source-  
 796 guided flow matching. *arXiv preprint arXiv:2508.14807*, 2025.

797

798 Hao Wu, Shu Liu, Xiaojing Ye, and Haomin Zhou. Parameterized Wasserstein Hamiltonian flow.  
 799 *SIAM Journal on Numerical Analysis*, 63(1):360–395, 2025. URL <https://doi.org/10.1137/23M159281X>.

800

801 Grace Hui Ting Yeo, Sachit D Saksena, and David K Gifford. Generative modeling of  
 802 single-cell time series with PRESCIENT enables prediction of cell trajectories with interven-  
 803 tions. *Nature communications*, 12(1):3222, 2021. URL <https://doi.org/10.1038/s41467-021-23518-w>.

804

805 Federico Zadra. *Topics in contact Hamiltonian systems: analytical and numerical perspectives*.  
 806 PhD thesis, University of Groningen, 2023.

807

808 Alireza Zaeemzadeh, Nazanin Rahnavard, and Mubarak Shah. Norm-preservation: Why residual  
 809 networks can become extremely deep? *IEEE transactions on pattern analysis and machine*  
*intelligence*, 43(11):3980–3990, 2020.

810 Kira Zeevaert, Mohamed H Elsafi Mabrouk, Wolfgang Wagner, and Roman Goetzke. Cell me-  
 811 chanics in embryoid bodies. *Cells*, 9(10):2270, 2020. URL <https://doi.org/10.3390/cells9102270>.

812

813 Jiaqi Zhang, Erica Larschan, Jeremy Bigness, and Ritambhara Singh. scNODE: generative model  
 814 for temporal single cell transcriptomic data prediction. *Bioinformatics*, 40(Supplement\_2):ii146-  
 815 ii154, 2024. URL <https://doi.org/10.1093/bioinformatics/btae393>.

816

817 Stephen Y. Zhang and Michael PH. Stumpf. Learning non-equilibrium diffusions with Schrödinger  
 818 bridges: from exactly solvable to simulation-free. *arXiv preprint arXiv:2505.16644*, 2025. URL  
 819 <https://arxiv.org/abs/2505.16644>.

820

821

## A EXTENDED STATE-OF-THE-ART ON SCHRÖDINGER BRIDGE SOLVERS

823 Existing methodologies for addressing the Schrödinger Bridge problem can be broadly divided into  
 824 two main categories, depending on their solution strategy: those that directly fit the stochastic dy-  
 825 namics on the probabilistic manifold, and those that leverage the analytic optimality conditions of  
 826 the problem to solve it. We review them in the sequel.

827 **Schrödinger Bridge Solvers via Dynamics Parametrization.** Traditional SB solvers often use  
 828 Iterative Proportional Fitting (IPF) (Kullback, 1968; Léonard, 2013), which alternates forward  
 829 and backward updates to successively match the initial and terminal marginals. However, IPF is  
 830 computationally expensive, as it stores full trajectories, and it suffers from error accumulation, nu-  
 831 matical instability, and reliance on strong priors (Vargas et al., 2021; Gushchin et al., 2024). Recent  
 832 matching-based approaches (Bortoli et al., 2024; Shi et al., 2023; Gushchin et al., 2024; Peluchetti,  
 833 2023) improve scalability and robustness by learning time-reversed drifts via Markovian projections,  
 834 circumventing the need for full trajectory storage and mitigating the IPF discretization errors. While  
 835 Liu et al. (2024) extended this idea to the GSB setting, their assumption of Gaussian probability  
 836 paths limits the model’s expressivity. To overcome this, Tang et al. (2025) proposed branched  
 837 dynamics using Gaussian mixtures. This enables more flexible path structures but at the expense  
 838 of higher computational cost. For mmSB, iterative reference refinement with piecewise SB inter-  
 839 polation (Shen et al., 2025) suffers from inconsistencies in global trajectory construction due to its  
 840 piecewise nature and shows high sensitivity to the choice of initial reference process. Alternatively,  
 841 Tong et al. (2020) proposed a continuous normalizing flow for deterministic interpolation, removing  
 842 noise from the reference process but preventing the construction of a true probabilistic bridge.

843 **Schrödinger Bridge Solvers via Optimality Conditions.** The optimality conditions of the  
 844 SB problem take the form of dynamical equations on a Hamiltonian phase space, driven by dual  
 845 potential functions (Chow et al., 2020). These conditions allow the exact dynamics to be recovered  
 846 via integration and provide a flexible framework for generalization through modifications of the  
 847 Hamiltonian. Furthermore, the state-dependent nature of the Hamiltonian framework offers a natural  
 848 way to obtain Markovian approximations of a stochastic process. Vargas et al. (2021) and Chen  
 849 et al. (2022) leveraged this view and solved the SB problem using control- and likelihood-based  
 850 approaches, both employing iterative forward-backward updates on the Hamiltonian dynamics. Liu  
 851 et al. (2022) extended this idea to the GSB setting, although without convergence guarantees. How-  
 852 ever, a critical bottleneck of these and derived methods is that a tractable integration of optimality  
 853 conditions relies on iterative updates of a reference process. For example, Buzun et al. (2025)  
 854 improved stability by directly modeling the dual potential and minimizing residuals of the Hamil-  
 855 tonian conditions, yet iterative updates incur significant computational overhead and may destabilize  
 856 training due to the dependence on self-generated samples (Bertrand et al., 2024). This issue persists  
 857 in mmSB settings (Theodoropoulos et al., 2025; Hong et al., 2025). Even when belief propagation  
 858 is used to reduce time complexity (Chen et al., 2023), scaling to high dimensions remains poor.  
 859 Therefore, while leveraging the optimality conditions offers clear advantages, it remains essential  
 860 to develop computationally efficient, non-iterative algorithms with favorable scaling properties.

861

862

863

864

## B NOTATION

865

866

867

868

Table 8: Summary of mathematical notation

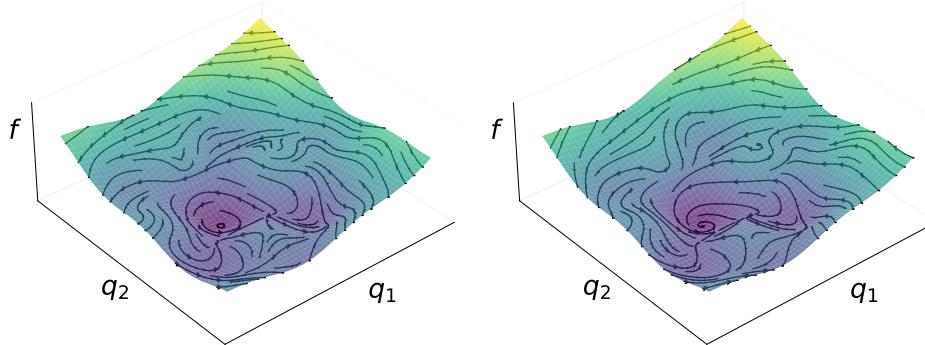
Notation	Description	Context
<b>Probability Distributions and Manifolds</b>		
$\mathcal{M}$	Underlying data manifold.	Sec. 3
$\mathcal{P}^+(\mathcal{M})$	The space of smooth, positive probability density functions supported on $\mathcal{M}$ (the Wasserstein manifold).	Sec. 3
$\rho(x), \rho^t(x)$	A probability density function and its time-dependent path.	Sec. 3
$\rho_a, \rho_b$	The initial and final endpoint marginal distributions.	Sec. 3
$\{\rho_m\}_{m=1}^M$	A set of intermediate marginal constraints at times $\{t_m\}_{m=1}^M$ .	Eq. (2)
$\lambda$	A fixed reference measure (e.g., standard Gaussian).	Sec. 5
<b>Transport Dynamics</b>		
$v^t(x)$	The drift term (control variable) in the Fokker-Planck equation.	Eq. (2)
$U(x)$	The external potential function (state cost).	Eq. (2)
$\epsilon$	Scaling factor for the stochastic diffusion term in the Fokker-Planck equation (related to entropy regularization).	Eq. (2)
$J$	The objective functional (cost) to be minimized both in the mmGSB problem and in the NCGSB problem.	Eq. (2)
$S^t(x)$	The potential function/Lagrange multiplier that acts as a canonical momentum, with $\nabla_x S^t(x) = v^t(x)$ .	Eq. (3)
$z^t$	The augmented scalar state representing the Lagrangian action.	Eq. (5)
$\gamma$	The damping factor in the Non-Conservative GSB formulation.	Eq. (5)
<b>Hamiltonian variables</b>		
$H(\rho^t, S^t)$	The Hamiltonian function, defined as the sum of kinetic ( $\mathcal{K}$ ) and potential ( $\mathcal{F}$ ) energy: $H(\rho^t, S^t) = \mathcal{K}(\rho^t, S^t) + \mathcal{F}(\rho^t)$ .	Eq. (3)
$H(\rho^t, S^t, z^t)$	The contact Hamiltonian function, defined as $H(\rho^t, S^t, z^t) = \mathcal{K}(\rho^t, S^t) + \mathcal{F}(\rho^t) + \mathcal{B}(\rho^t) + \gamma z^t$ .	Eq. (6)
$\mathcal{K}$	Kinetic energy of the transport map.	Sec. 3
$\mathcal{F}$	Potential energy, $\mathcal{F} = -U - I$ .	Sec. 3
$I$	Fisher information term arising from entropy regularization, defined as $I = \int_{\mathcal{M}} \ \nabla_x \log \rho^t(x)\ ^2 \rho^t(x) dx$ .	Sec. 3
$\mathcal{B}$	Entropy term, defined as $\mathcal{B} = \int_{\mathcal{M}} \epsilon(\log \rho^t(x) - 1) \rho^t(x) dx$ , producing an additional diffusion in the contact Hamiltonian dynamics.	Sec. 4
<b>Guidance Terms</b>		
$\ y - f(x^{t_s})\ ^2$	Guidance loss terms, penalizing the difference between the detected features and the desired values $y$	Eq. 5
$f(x^{t_s})$	Feature function applied to the sample $x^{t_s}$ at time $t_s$ . It can indicate a class, describe a property of the sample, or quantify the sample's value depending on the task.	Eq. 5
<b>Geometric Terms</b>		
$g^{\mathcal{W}_2}(\cdot, \cdot)$	The Wasserstein metric tensor on $\mathcal{P}^+(\mathcal{M})$ .	Eq. (1)
$d_{\mathcal{W}_2}(\rho_a, \rho_b)$	The Wasserstein distance between $\rho_a$ and $\rho_b$ .	Sec. 3
$\Delta_{\rho^t}$	The weighted Laplacian operator, $\Delta_{\rho^t} = -\nabla_x \cdot (\rho^t \nabla_x)$ .	Eq. (1)
$g_J$	The Jacobi metric, $g_J = (H - \mathcal{F})g^{\mathcal{W}_2}$ , which is minimized by the optimal bridge geodesic for the mmGSB problem.	Sec. 3
$\tilde{g}_J$	The Jacobi metric, $\tilde{g}_J = (H - \mathcal{F} - \mathcal{B})g^{\mathcal{W}_2}$ , which is minimized by the optimal bridge geodesic for the NCGSB problem.	Sec. 5
<b>ResNet parameterization</b>		
$T_{\{\theta^k\}_{k=0}^K}$	ResNet model.	Eq. (7)
$T_{\theta^k}$	A parametrized pushforward map (one block of the ResNet).	Eq. (7)
$\theta^k$	The parameters associated with the pushforward map $T_{\theta^k}$ .	Eq. (7)
$\Phi^k$	Scaling term for the Wasserstein distance, defined as $\Phi^{t_k} = H(\rho^{t_k}, S^{t_k}, z^{t_k}) - \mathcal{F}(\rho^{t_k}) - \mathcal{B}(\rho^{t_k})$ .	Sec. 5

918 C EXTENDED PRELIMINARIES ON DIFFERENTIAL GEOMETRY  
919

## 920 C.1 HAMILTONIAN AND CONTACT HAMILTONIAN DYNAMICS

921 Hamiltonian and contact Hamiltonian dynamics are governed by specific energy constraints that  
922 can be analyzed via differential geometry as flows on specialized manifolds. Hamiltonian dynamics  
923 is energy-conserving and evolves on a *symplectic manifold*. Contact Hamiltonian dynamics is more  
924 general, allowing for variable energy levels, and takes place on a *contact manifold*. Their formal  
925 definitions and key differences are discussed next.

926 **Symplectic and Contact Structures.** Let  $\mathcal{M}$  be a smooth compact manifold, and let  $\mathcal{T}_x\mathcal{M}$  denote  
927 the tangent space at  $x \in \mathcal{M}$ . The collection of all the tangent spaces identifies the tangent bundle  
928  $\mathcal{T}\mathcal{M} = \bigcup_{x \in \mathcal{M}} \mathcal{T}_x\mathcal{M}$ . A vector field  $X : \mathcal{M} \rightarrow \mathcal{T}\mathcal{M}$  assigns a tangent vector  $v$  to each point  $x \in \mathcal{M}$ .  
929 The set of all the vector fields over  $\mathcal{T}\mathcal{M}$  is denoted as  $\Gamma(\mathcal{T}\mathcal{M})$ . A differential 1-form  $\alpha : \mathcal{T}\mathcal{M} \rightarrow \mathbb{R}$   
930 is a smooth map field acting on vectors of the tangent bundle. For a smooth function  $f : \mathcal{M} \rightarrow \mathbb{R}$ ,  
931 the 1-form  $\alpha = df$  generalizes the gradient from Euclidean spaces. Specifically,  $df$  measures the  
932 variation of  $f$  under an infinitesimal displacement on  $\mathcal{M}$ . This displacement is locally described by a  
933 starting point  $x$  and a direction  $v$ , such that  $(x, v) \in \mathcal{T}\mathcal{M}$ . Alternatively, it can be globally expressed  
934 by a vector field  $X$ . The variation of  $f$  along the vector field  $X$  is given by  $df(X)$ . This variation  
935 is independent of the choice of reference frame. To preserve this invariance,  $df$  must transform  
936 covariantly with  $X$ . Consequently, the 1-form  $\alpha = df$  resides in the cotangent bundle  $\mathcal{T}^*\mathcal{M}$ ,  
937 the dual space to  $\mathcal{T}\mathcal{M}$ . The symplectic and contact structures provide two distinct mechanisms  
938 for associating a 1-form to a vector field, thereby establishing connections between the tangent and  
939 cotangent bundles. By considering the dynamics governed by the vector field and the scalar function  
940 defining the 1-form, a relationship between these elements emerges, as illustrated in Fig. 11.



953 Figure 11: The same scalar function  $f$ , associated with the 1-form  $\alpha = df$ , gives rise to two distinct vector  
954 fields under the symplectic (left) and contact (right) geometric structures. The streamlines of these vector fields  
955 are illustrated on a representation of the state manifold. In symplectic geometry, the streamlines are tangent  
956 to the level curves of  $f$ , representing isoenergetic trajectories where  $f$  remains constant, thus describing the  
957 dynamics of conservative systems. In contrast, in contact geometry, a single flow line can traverse different  
958 energy levels.

959 **Symplectic Geometry.** A differential 2-form  $\omega : \mathcal{T}\mathcal{M} \times \mathcal{T}\mathcal{M} \rightarrow \mathbb{R}$  is a skew-symmetric, bilinear,  
960 and smooth field of maps acting on pairs of tangent vectors. A 2-form is called *symplectic* if it is  
961 both closed ( $d\omega = 0$ ) and non-degenerate. The symplectic form lacks the properties required to  
962 define an inner product. However, it still establishes a fundamental relation between differential  
963 1-forms and vector fields: Given a 1-form  $df$ , the symplectic form  $\omega$  uniquely determines a vector  
964 field  $X_f$  that is tangent to the level sets of  $f$ . This relation is defined by,  
965

$$966 \quad df(X) = \omega(X_f, X), \quad \forall X \in \Gamma(\mathcal{T}\mathcal{M}). \quad (11)$$

967 By definition,  $f$  remains constant along the flow of  $X_f$ , which in turn preserves the symplectic  
968 form  $\omega$ , i.e.,  $\mathcal{L}_{X_f}\omega = 0$  where  $\mathcal{L}_{X_f}$  denotes the Lie derivative (Silva, 2001). In this framework, the  
969 function  $f$  is interpreted as a conserved energy, or equivalently, as a Hamiltonian  $H$ . The symplectic  
970 structure thereby endows  $\mathcal{M}$  with a natural geometric framework for formulating Hamiltonian  
971 dynamics (Tokasi & Pickl, 2022). The pair  $(\mathcal{M}, \omega)$  is referred to as a symplectic manifold. Notably,  
972 the non-degeneracy of  $\omega$  implies that  $\mathcal{M}$  must be even-dimensional.

**Contact Geometry.** While symplectic manifolds provide a geometric framework for modeling the dynamics of conservative systems in classical mechanics, a more general approach is required to describe non-conservative systems. This is addressed by contact manifolds, the odd-dimensional counterparts of symplectic manifolds (Geiges, 2001; Bravetti et al., 2017). A *contact manifold* is defined as  $(\mathcal{M}, \eta)$ , where  $\mathcal{M}$  is an odd-dimensional smooth manifold, and  $\eta$  is a non-degenerate 1-form known as the contact form (Geiges, 2008). The contact form satisfies the *maximal non-integrability* condition, meaning that the top-degree differential form  $\eta \wedge (d\eta)^d \neq 0$  is nowhere vanishing on  $\mathcal{M}$ . This form is constructed by taking the exterior product of  $\eta$  with the  $d$ -fold wedge product of its exterior derivative  $d\eta$ , i.e.,

$$(d\eta)^d = \underbrace{d\eta \wedge \cdots \wedge d\eta}_{d \text{ times}}. \quad (12)$$

The  $(2d + 1)$ -form defines a volume form on  $\mathcal{M}$ , ensuring that the hyperplanes  $\ker(\eta) \subset \mathcal{T}\mathcal{M}$ , constraining the dynamics on the contact manifold, do not form a foliation, i.e., they do not partition the manifold into lower-dimensional submanifolds (Geiges, 2001; 2008). Geometrically, this means that the contact distribution imposes *non-holonomic constraints*: it restricts the admissible directions of motion at each point without confining the dynamics to a fixed submanifold or energy level. This property is crucial for modeling systems where energy can change over time, enabling constraints on energy behavior without enforcing conservation.

Like symplectic geometry, contact geometry connects scalar functions to vector fields, enabling the description of dynamical systems (Zadra, 2023). Given an energy function  $H : \mathcal{M} \rightarrow \mathbb{R}$ , the dynamics on a contact manifold are defined by a contact Hamiltonian vector field  $X_H$ , as follows,

$$dH(X) = d\eta(X_H, X) - \mathcal{L}_{X_H} \eta(X), \quad \forall X \in \Gamma(\mathcal{T}\mathcal{M}). \quad (13)$$

Unlike symplectic geometry, where dynamics are confined to energy-preserving flows along the level sets of the Hamiltonian, contact geometry allows for an additional component of motion. Specifically, the dynamics on a contact manifold are not restricted to the term  $d\eta(X_H, X)$ , which lies tangent to the level sets of  $H$ , but also include a transverse component  $\mathcal{L}_{X_H} \eta(X)$ , arising from the non-degeneracy of the contact form. Consequently, while in symplectic geometry the symplectic form  $\omega$  is strictly preserved, contact geometry allows the contact form  $\eta$  to be preserved only up to a scaling factor  $a \in \mathbb{R}$  (Bravetti et al., 2017).

## C.2 RIEMANNIAN AND JACOBI METRICS

The Jacobi metric  $g_J$  is a rescaled version of a Riemannian metric  $g$  that allows Hamiltonian dynamics on the cotangent bundle  $\mathcal{T}^*\mathcal{M}$  to be represented as geodesics on the Riemannian manifold  $(\mathcal{M}, g)$ . The construction is detailed below.

**The Riemannian Metric.** Let  $\mathcal{M}$  be a smooth compact manifold. A Riemannian metric  $g : \mathcal{T}\mathcal{M} \times \mathcal{T}\mathcal{M} \rightarrow \mathbb{R}$  is a smooth, symmetric, and positive-definite bilinear field of maps defined on pairs of vectors in the tangent bundle. This enables the introduction of an inner product on the tangent spaces of the manifold, allowing us to measure distances and curve lengths. For a smooth curve  $x(t) : [t_0, t_1] \rightarrow \mathcal{M}$ , the length  $l$  w.r.t. the metric  $g$  is  $l = \int_{t_0}^{t_1} \sqrt{g(\dot{x}(t), \dot{x}(t))} dt$ , where  $\dot{x}(t) \in \mathcal{T}_{x(t)}\mathcal{M}$  is the vector tangent to the curve at  $x(t)$ . The curve minimizing this length between two points  $x(t_0)$  and  $x(t_1)$  on  $\mathcal{M}$  is called a *geodesic*. Geodesics generalize straight lines in Euclidean space to curved spaces, representing the shortest paths in the geometry induced by  $g$ .

**The Jacobi Metric.** The geodesic flow  $x(t)$  on a Riemannian manifold  $(\mathcal{M}, g)$  lifts to the joint evolution of coordinates  $(x(t), \alpha(x(t), \dot{x}(t)))$  on the cotangent bundle  $\mathcal{T}^*\mathcal{M}$  (Abraham & Marsden, 2008). This extended dynamics is governed by an energy function  $H(x, \alpha) : \mathcal{T}^*\mathcal{M} \rightarrow \mathbb{R} = g^{-1}(\alpha, \alpha)$ , which remains constant along the flow. A reparameterization  $ds = \sqrt{H}dt$  links the trajectory of the integrated dynamical system at time  $t$  on  $\mathcal{T}^*\mathcal{M}$  with the length of the corresponding geodesic on  $\mathcal{M}$ . This framework reveals a fundamental connection between geodesic flows and Hamiltonian dynamics in the special case where the Hamiltonian consists solely of a kinetic energy term. The cotangent bundle  $\mathcal{T}^*\mathcal{M}$  is naturally equipped with a symplectic structure, making it a symplectic manifold  $(\mathcal{T}^*\mathcal{M}, \omega)$ .

This formulation can be further generalized by introducing a potential energy function into the Hamiltonian, given by  $H(x, \alpha) = g^{-1}(\alpha, \alpha) + \mathcal{F}(x)$ . In this setting, the geodesic structure un-

1026 derlying the Hamiltonian flow is determined by the Jacobi metric,  
 1027

$$g_J = (H - \mathcal{F}(x)) g, \quad (14)$$

1029 which rescales the original metric  $g$  by a position-dependent conformal factor (Abraham & Marsden,  
 1030 2008). The corresponding time reparameterization takes the form  $ds = \sqrt{H - \mathcal{F}(x)} dt$ , restoring  
 1031 the interpretation of the trajectory as a geodesic with respect to the metric  $g_J$  (Udriște, 2000).  
 1032

## 1034 D INSIGHTS ON THE SCHRÖDINGER BRIDGE

### 1035 D.1 HAMILTONIAN STRUCTURE OF THE GENERALIZED SCHRÖDINGER BRIDGE

1038 This part presents the derivation of the Hamiltonian structure of the mmGSB problem (2), intro-  
 1039 duced in section 3, obtained via the method of Lagrange multipliers. To transform the constrained  
 1040 optimization problem into an unconstrained one, we incorporate the Fokker–Planck equation, scaled  
 1041 by the Lagrange multiplier  $S^t$ , into the original running cost  $\mathcal{L}$  (i.e., the Lagrangian), as follows,  
 1042

$$J(v^t, \rho^t, S^t) = \int_0^1 \mathcal{L}(v^t, \rho^t, S^t) dt; \quad (15)$$

$$\begin{aligned} \mathcal{L}(v^t, \rho^t, S^t) &= \int_{\mathcal{M}} \left( \frac{1}{2} \|v^t(x)\|^2 + U(x) \right) \rho^t(x) dx \\ &\quad + \int_{\mathcal{M}} S^t(x) \underbrace{(\partial_t \rho^t(x) + \nabla_x \cdot (\rho^t(x) v^t(x)) - \varepsilon \Delta \rho^t(x))}_{\text{Fokker-Planck equation}} dx. \end{aligned} \quad (16)$$

1050 The optimality conditions resulting from the extremization of the cost functional  $J$  in equa-  
 1051 tion (15) follow from the Euler–Lagrange equations, generalized to the setting of classical field  
 1052 theory (Blohm, 2024). In this framework, the arguments of the Lagrangian  $\mathcal{L}$ , in equation (16),  
 1053 are viewed as smooth fields defined over space and time. By setting to zero the variations of  $\mathcal{L}$  with  
 1054 respect to these fields, we obtain the stationarity conditions for  $J$ . For a generic field  $\psi^t(x)$ , the  
 1055 corresponding Euler–Lagrange equation takes the form,  
 1056

$$d_\psi \mathcal{L} = \partial_\psi \mathcal{L} + \partial_t(\partial_{\partial_t \psi} \mathcal{L}) + \nabla_x \cdot (\partial_{\nabla_x \psi} \mathcal{L}) + \Delta_x(\partial_{\Delta_x \psi} \mathcal{L}) = 0. \quad (17)$$

1058 Applying equation (17) to equation (16) for the fields  $v^t$ ,  $\rho^t$ , and  $S^t$ , we obtain the following system  
 1059 of optimality conditions,  
 1060

$$d_v \mathcal{L} = v^t(x) \rho^t(x) - \rho^t(x) \nabla_x S^t(x) = 0 \implies v^t(x) = \nabla_x S^t(x), \quad (18a)$$

$$d_\rho \mathcal{L} = \frac{1}{2} \|v^t(x)\|^2 - \partial_t S^t(x) - \nabla_x S^t(x) \cdot v^t(x) - \varepsilon \Delta_x S^t(x) + U(x) = 0, \quad (18b)$$

$$d_S \mathcal{L} = \partial_t \rho^t(x) + \nabla_x \cdot (\rho^t(x) v^t(x)) - \varepsilon \Delta_x \rho^t(x) = 0. \quad (18c)$$

1065 Substituting the expression for the optimal velocity from equation (18a) into equations (18b) and  
 1066 (18c), we obtain the following Hamiltonian system,  
 1067

$$\partial_t \rho^t(x) = \partial_S H(\cdot) = -\nabla_x \cdot (\rho^t(x) \nabla_x S^t(x)) + \varepsilon \Delta \rho^t(x), \quad (19a)$$

$$\partial_t S^t(x) = -\partial_\rho H(\cdot) = -\frac{1}{2} \|\nabla_x S^t(x)\|^2 - \varepsilon \Delta_x S^t(x) + U(x), \quad (19b)$$

1071 with the corresponding Hamiltonian function,  
 1072

$$H(\rho^t, S^t) = \frac{1}{2} \int_{\mathcal{M}} \|\nabla_x S^t(x)\|^2 \rho^t(x) dx - \int_{\mathcal{M}} U(x) \rho^t(x) dx + \varepsilon \int_{\mathcal{M}} S^t(x) \Delta_x \rho^t(x) dx. \quad (20)$$

1075 The Hamiltonian system (19) can be reformulated in a linear and decoupled form via the coordinate  
 1076 transformation from the original variables  $(\rho^t, S^t)$  to an alternative canonical set  $(\hat{\rho}^t, \hat{S}^t)$ . This  
 1077 transformation, known as Hopf–Cole coordinate transformation (Léger & Li, 2021; Chow et al.,  
 1078 2020), is derived from the generating function  $F$ ,  
 1079

$$F(\rho, \hat{S}) = \rho(x) \hat{S} + \varepsilon \rho(x) (\log \rho(x) - 1), \quad (21)$$

1080 and preserve the symplectic form, i.e.,  $dF = Sd\rho - \hat{S}d\hat{\rho}$ . Following the standard generating-  
 1081 function rules (Chapter 18.3 [Johns \(2005\)](#)),  
 1082

$$1083 \quad \hat{\rho}^t(x) = \partial_{\hat{S}}F(\cdot), \quad S^t(x) = \partial_{\rho}F(\cdot), \quad (22)$$

1084 the new coordinates are explicitly given by,  
 1085

$$1086 \quad \hat{\rho}^t(x) = \rho^t(x), \quad (23a)$$

$$1087 \quad \hat{S}^t(x) = S^t(x) - \varepsilon \log \rho^t(x). \quad (23b)$$

1088 The Hopf–Cole transformation is well established in the literature, not only for simplifying the  
 1089 mathematical form of the equations, but also for enabling the design of efficient numerical inte-  
 1090 gration schemes ([Léger & Li, 2021](#)). In our context, it is particularly advantageous for obtaining a  
 1091 Hamiltonian with separable kinetic and potential energy components. The Hamiltonian system (19)  
 1092 then becomes,  
 1093

$$1094 \quad \partial_t \hat{\rho}^t(x) + \nabla_x \cdot (\hat{\rho}^t(x) \nabla_x \hat{S}^t(x)) = 0, \quad (24a)$$

$$1095 \quad \partial_t \hat{S}^t(x) + \varepsilon \frac{1}{\hat{\rho}^t(x)} \partial_t \hat{\rho}^t(x) = -\frac{1}{2} \left\| \nabla_x \hat{S}^t(x) + \varepsilon \nabla_x \log \hat{\rho}^t(x) \right\|^2 \\ 1096 \quad - \varepsilon \Delta_x \hat{S}^t(x) - \varepsilon^2 \Delta_x \log \hat{\rho}^t(x) + U(x). \quad (24b)$$

1097 Expanding the squared norm in equation (24b) and substituting  $\partial_t \rho^t(x)$  from equation (24a) yields,  
 1098

$$1100 \quad \partial_t \hat{S}^t(x) - \varepsilon \frac{1}{\hat{\rho}^t(x)} \nabla_x \cdot (\hat{\rho}^t(x) \nabla_x \hat{S}^t(x)) = -\frac{1}{2} \left\| \nabla_x \hat{S}^t(x) \right\|^2 - \frac{1}{2} \varepsilon^2 \left\| \nabla_x \log \hat{\rho}^t(x) \right\|^2 \\ 1101 \quad - \varepsilon \nabla_x \hat{S}^t(x) \nabla_x \log \hat{\rho}^t(x) - \varepsilon \Delta_x \hat{S}^t(x) \\ 1102 \quad - \varepsilon^2 \Delta_x \log \hat{\rho}^t(x) + U(x). \quad (25)$$

1103 Using the identity,  
 1104

$$1105 \quad \varepsilon \frac{1}{\hat{\rho}^t(x)} \nabla_x \cdot (\hat{\rho}^t(x) \nabla_x \hat{S}^t(x)) = \varepsilon \Delta_x \hat{S}^t(x) + \varepsilon \nabla_x \hat{S}^t(x) \nabla_x \log \hat{\rho}^t(x), \quad (26)$$

1106 equation (25) simplifies to,  
 1107

$$1111 \quad \partial_t \hat{S}^t(x) = -\frac{1}{2} \left\| \nabla_x \hat{S}^t(x) \right\|^2 - \varepsilon^2 \frac{1}{2} \left\| \nabla_x \log \hat{\rho}^t(x) \right\|^2 - \varepsilon^2 \Delta_x \log \hat{\rho}^t(x) + U(x). \quad (27)$$

1112 This form reveals the emergence of the Fisher information, defined as,  
 1113

$$1114 \quad I(\hat{\rho}^t(x)) = \int_{\mathcal{M}} \left\| \nabla_x \log \hat{\rho}^t(x) \right\|^2 \hat{\rho}^t(x) dx, \quad (28a)$$

$$1115 \quad \partial_{\hat{\rho}} I(\hat{\rho}^t(x)) = -2 \Delta_x \log \hat{\rho}^t(x) - \left\| \nabla_x \log \hat{\rho}^t(x) \right\|^2. \quad (28b)$$

1116 Thus, equations (24a) and (27) admit the linear decoupled Hamiltonian formulation described in  
 1117 equation (3),  
 1118

$$1119 \quad \partial_t \hat{\rho}^t(x) = \partial_{\hat{S}} H(\cdot) = -\nabla_x \cdot (\hat{\rho}^t(x) \nabla_x \hat{S}^t(x)); \\ 1120 \quad \partial_t \hat{S}^t(x) = -\partial_{\hat{\rho}} H(\cdot) = -\frac{1}{2} \left\| \nabla_x \hat{S}^t(x) \right\|^2 + \frac{1}{2} \varepsilon^2 \partial_{\hat{\rho}} I(\hat{\rho}^t(x)) + U(x),$$

1121 governed by the Hamiltonian function,  
 1122

$$1123 \quad H(\hat{\rho}^t, \hat{S}^t) = \underbrace{\frac{1}{2} \int_{\mathcal{M}} \left\| \nabla_x \hat{S}^t(x) \right\|^2 \hat{\rho}^t(x) dx}_{\text{Kinetic energy } \mathcal{K}} - \underbrace{\int_{\mathcal{M}} U(x) \hat{\rho}^t(x) dx - \frac{1}{2} \varepsilon^2 I(\hat{\rho}^t(x))}_{\text{Potential energy } \mathcal{F}} \quad (29)$$

1124 In this formulation, the Fisher information contributes to the potential energy and encodes the effect  
 1125 of stochastic diffusion. Minimizing the Fisher information term promotes smoothness in the density  
 1126 and steers the Hamiltonian flow toward the target distribution  $\rho^1$ , providing greater robustness  
 1127

1134 due to the regularizing effect of diffusion. This mechanism has been studied in the literature and  
 1135 employed in control applications for its regularization properties (Chen et al., 2025).  
 1136  
 1137

## 1138 D.2 CONTACT HAMILTONIAN STRUCTURE OF THE NON-CONSERVATIVE GSB

1140 Here we derive the contact Hamiltonian formulation of the NCGSB problem, introduced in equa-  
 1141 tion (4) and discussed in section 4. The structure of the derivation closely mirrors that of the GSB in  
 1142 Appendix D.1, with one key distinction: the Lagrangian  $\mathcal{L}$  now depends explicitly on the accumu-  
 1143 lated action  $z^t$ . Specifically, the augmented cost functional and Lagrangian are given by,  
 1144

$$1145 J(v^t, \rho^t, S^t, z^t) = \int_0^1 \mathcal{L}(v^t, \rho^t, S^t, z^t) dt; \quad (30)$$

$$1147 \mathcal{L}(v^t, \rho^t, S^t, z^t) = \int_{\mathcal{M}} \left( \frac{1}{2} \|v^t(x)\|^2 + U(x) \right) \rho^t(x) dx - \gamma z^t \\ 1149 + \int_{\mathcal{M}} S^t(x) (\partial_t \rho^t(x) + \nabla_x \cdot (\rho^t(x) v^t(x)) - \varepsilon \Delta \rho^t(x)) dx, \quad (31)$$

1151 where the scalar  $\gamma$  acts as a damping factor. Since  $\mathcal{L}$  depends explicitly on the evolving action  $z^t$ , the  
 1152 problem lies outside the scope of classical variational calculus. Instead, it fits within the framework  
 1153 of non-conservative variational principles, where the cost functional  $J$  evolves dynamically with the  
 1154 system. This is naturally addressed by the Herglotz variational principle, which extends the Euler-  
 1155 Lagrange equations to systems with dissipative effects. The optimality conditions obtained from the  
 1156 variations of  $\mathcal{L}$ , namely, the Herglotz-type Euler–Lagrange equations, for a generic field argument  
 1157  $\psi^t(x)$  in this case take the form,  
 1158

$$1159 d_{\psi} \mathcal{L} = \partial_{\psi} \mathcal{L} + \partial_t (\partial_{\partial_t \psi} \mathcal{L}) + \nabla_x \cdot (\partial_{\nabla_x \psi} \mathcal{L}) + \Delta_x (\partial_{\Delta_x} \mathcal{L}) - \partial_z \mathcal{L} \partial_{\partial_t \psi} \mathcal{L} = 0. \quad (32)$$

1160 Applying equation (32) for the fields  $v^t$ ,  $\rho^t$ , and  $S^t$ , and recovering the dynamics of  $z^t$  from the  
 1161 optimization problem (4), we obtain the following system of optimality conditions,  
 1162

$$1163 d_v \mathcal{L} = v^t(x) \rho^t(x) - \rho^t(x) \nabla_x S^t(x) = 0 \implies v^t(x) = \nabla_x S^t(x), \quad (33a)$$

$$1164 d_{\rho} \mathcal{L} = \frac{1}{2} \|v^t(x)\|^2 - \partial_t S^t(x) - \nabla_x S^t(x) \cdot v^t(x) - \varepsilon \Delta_x S^t(x) + U(x) - \gamma S^t(x) = 0, \quad (33b)$$

$$1166 d_S \mathcal{L} = \partial_t \rho^t(x) + \nabla_x \cdot (\rho^t(x) v^t(x)) - \varepsilon \Delta_x \rho^t(x) = 0, \quad (33c)$$

$$1168 d_t z^t - \int_{\mathcal{M}} \left( \frac{1}{2} \|v^t(x)\|^2 + U(x) \right) \rho^t(x) dx - \gamma z^t = 0. \quad (33d)$$

1170 Compared to the optimality conditions for the GSB problem presented in equations (18), the set (33)  
 1171 includes an additional term,  $-\gamma S^t(x)$ , in equation (33b), which accounts for the dissipation term.  
 1172 By substituting the expression for the optimal velocity from equation (33a) into equations (33c),  
 1173 (33b), and equation (33d), we obtain the system of contact Hamiltonian dynamics,  
 1174

$$1174 \partial_t \rho^t(x) = \partial_S H(\cdot) = -\nabla_x \cdot (\rho^t(x) \nabla_x S^t(x)) + \varepsilon \Delta_x \rho^t(x), \quad (34a)$$

$$1175 \partial_t S^t(x) = -\partial_{\rho} H(\cdot) - S^t(x) \partial_z H(\cdot) = -\frac{1}{2} \|\nabla_x S^t(x)\|^2 - \varepsilon \Delta_x S^t(x) + U(x) - \gamma S^t(x), \quad (34b)$$

$$1178 \partial_t z^t = S^t(x) \partial_S H(\cdot) - H(\cdot) = \int_{\mathcal{M}} \left( \frac{1}{2} \|\nabla_x S^t(x)\|^2 + U(x) \right) \rho^t(x) dx + \gamma z^t, \quad (34c)$$

1180 with the associated contact Hamiltonian function given by,  
 1181

$$1182 H(\rho^t, S^t, z^t) = \frac{1}{2} \int_{\mathcal{M}} \|\nabla_x S^t(x)\|^2 \rho^t(x) dx - \int_{\mathcal{M}} U(x) \rho^t(x) dx + \varepsilon \int_{\mathcal{M}} S^t(x) \Delta_x \rho^t(x) dx - z^t. \quad (35)$$

1184 In this case as well, it is beneficial to derive a decoupled and linearized representation of the  
 1185 dynamics. To this end, we perform a coordinate transformation from the original variables  
 1186  $(\rho^t, S^t, z^t)$  to an alternative canonical set  $(\hat{\rho}^t, \hat{S}^t, \hat{z}^t)$ , while preserving the contact structure,  
 1187

$$1188 dz - S d\rho = d\hat{z} - \hat{S} d\hat{\rho}. \quad (36)$$

1188 This is achieved via a contact transformation generated by a generating function  $F$ , defined as,  
 1189

$$1190 \quad F(\rho, \hat{S}, \hat{z}^t) = \rho(x)\hat{S} + \varepsilon\rho(x)(\log \rho(x) - 1) + \hat{z}^t, \quad (37)$$

1191 which induces the following coordinate change according to the contact generating rules in  
 1192 [Struckmeier & Redelbach \(2008\)](#),  
 1193

$$1194 \quad \hat{\rho}^t(x) = \partial_{\hat{S}}F(\cdot), \quad (38a)$$

$$1195 \quad S^t(x) = \partial_{\rho}F(\cdot), \quad (38b)$$

$$1196 \quad \hat{z}^t = F(\cdot) - \hat{S}\hat{\rho}. \quad (38c)$$

1198 As a result, the transformed variables are given by  
 1199

$$1200 \quad \hat{\rho}^t(x) = \rho^t(x), \quad (39a)$$

$$1201 \quad \hat{S}^t(x) = S^t(x) - \varepsilon \log \rho^t(x), \quad (39b)$$

$$1202 \quad \hat{z}^t = z^t + \varepsilon \rho^t(x) (\log \rho^t(x) - 1). \quad (39c)$$

1204 By substituting the new coordinates (39) into the dynamical system (34), we obtain the following  
 1205 transformed contact Hamiltonian dynamics,  
 1206

$$1206 \quad \partial_t \hat{\rho}^t(x) = \partial_{\hat{S}}H(\cdot) = -\nabla_x \cdot (\hat{\rho}^t(x) \nabla_x \hat{S}^t(x)),$$

$$1207 \quad \partial_t \hat{S}^t(x) = -\partial_{\hat{\rho}}H(\cdot) - \hat{S}^t(x)\partial_{\hat{z}}H(\cdot) = -\frac{1}{2}\|\nabla_x \hat{S}^t(x)\|^2 + \frac{1}{2}\varepsilon^2 \partial_{\hat{\rho}}I(\hat{\rho}^t(x)) + U(x)$$

$$1208 \quad \quad \quad - \gamma \hat{S}^t(x) - \varepsilon \gamma \log \hat{\rho}^t(x),$$

$$1209 \quad \partial_t \hat{z}^t = \hat{S}^t(x)\partial_{\hat{S}}H(\cdot) - H(\cdot) = \int_{\mathcal{M}} \left( \frac{1}{2}\|\nabla_x \hat{S}^t(x)\|^2 + U(x) \right) \hat{\rho}^t(x) dx + \frac{1}{2}\varepsilon^2 I(\hat{\rho}^t)$$

$$1210 \quad \quad \quad - \int_{\mathcal{M}} \varepsilon \gamma (\log \hat{\rho}^t(x) - 1) \hat{\rho}^t(x) dx - \gamma \hat{z}^t,$$

$$1211 \quad \quad \quad$$

1216 with the associated contact Hamiltonian function given by,  
 1217

$$1217 \quad H(\hat{\rho}^t, \hat{S}^t, \hat{z}^t) = \underbrace{\frac{1}{2} \int_{\mathcal{M}} \|\nabla_x \hat{S}^t(x)\|^2 \hat{\rho}^t(x) dx}_{\text{Kinetic energy } \mathcal{K}} + \underbrace{\gamma \hat{z}^t}_{\text{Non-conservative potential}}$$

$$1218 \quad \quad \quad - \underbrace{\int_{\mathcal{M}} U(x) \hat{\rho}^t(x) dx - \frac{1}{2}\varepsilon^2 I(\hat{\rho}^t)}_{\text{Potential energy } \mathcal{F}} + \underbrace{\int_{\mathcal{M}} \varepsilon \gamma (\log \hat{\rho}^t(x) - 1) \hat{\rho}^t(x) dx}_{\text{Entropy } \mathcal{B}}. \quad (40)$$

$$1219 \quad \quad \quad$$

$$1220 \quad \quad \quad$$

$$1221 \quad \quad \quad$$

$$1222 \quad \quad \quad$$

$$1223 \quad \quad \quad$$

$$1224 \quad \quad \quad$$

$$1225 \quad \quad \quad$$

$$1226 \quad \quad \quad$$

$$1227 \quad \quad \quad$$

$$1228 \quad \quad \quad$$

$$1229 \quad \quad \quad$$

$$1230 \quad \quad \quad$$

$$1231 \quad \quad \quad$$

$$1232 \quad \quad \quad$$

$$1233 \quad \quad \quad$$

$$1234 \quad \quad \quad$$

$$1235 \quad \quad \quad$$

$$1236 \quad \quad \quad$$

$$1237 \quad \quad \quad$$

$$1238 \quad \quad \quad$$

$$1239 \quad \quad \quad$$

$$1240 \quad \quad \quad$$

$$1241 \quad \quad \quad$$

1242 D.3 GUIDED SCHRÖDINGER BRIDGE  
1243

1244 We consider the bridge  $\rho^t$ , computed between the terminal marginals  $\rho_a$  and  $\rho_b$ , using any variant  
1245 of the SB problem (e.g., GSB, mmSB, NCGSB). This process can be modified to enforce desired  
1246 conditions  $y$  at any chosen time  $t_s$  defined as,  
1247

$$y = f(x^{t_s}), \quad x^{t_s} \sim \rho^{t_s}, \quad (41)$$

1248 while preserving the underlying data manifold. Conditioning in this way modifies the probability  
1249 flow  $\rho^t$  (Guo et al., 2024) as,  
1250

$$\rho^t(x | y) = \frac{1}{Z} \rho^t(x) e^{-\|y - f(x^{t_s})\|^2}, \quad (42)$$

1253 where  $Z$  is a normalization constant, and  $x^{t_s}$  denotes a sample from the predicted prescribed distri-  
1254 bution  $\rho^{t_s}(x - y)$ . This weight biases the generation toward samples that satisfy the desired property  
1255  $y$ . In a dynamical setting, we perform this conditioning by incorporating a control term  $G^t$  into the  
1256 Fokker–Planck dynamics,

$$\partial_t \rho^t(x) + \nabla_x \cdot [\rho^t(x) (v^t(x) + G^t(x))] = \varepsilon \Delta_x \rho^t(x). \quad (43)$$

1259 A naive choice such as  $G^t(x) \propto \nabla_{x^t} f(x^t)$  often drives the dynamics off the data manifold, produc-  
1260 ing unrealistic samples far from it and the target distribution  $\rho_b$ . Instead, Bayes’ rule provides the  
1261 correct structure of the guidance term. The gradient of the conditional log-likelihood decomposes  
1262 as,

$$\nabla_{x^t} \log \rho^t(x^t | y) = \nabla_{x^t} \log \rho^t(x^t) + \underbrace{\nabla_{x^t} \log \rho^t(y | x^t)}_{\text{estimated by } G^t}. \quad (44)$$

1266 Substituting the conditional form from equation (42) yields,

$$G^t(x) = \nabla_{x^t} \log e^{-\|y - f(x^{t_s})\|^2} = -\nabla_{x^t} \|y - f(x^{t_s})\|^2. \quad (45)$$

1269 Specifically, in the NCGSB framework (4), this guidance is incorporated directly into the drift  $v^t$  by  
1270 adding the penalty  $\|y - f(x^{t_s})\|^2$  to the Lagrangian in the action dynamics constraint (4b),  
1271

$$\begin{aligned} \min_{v^t} \quad & J(v^t, \rho^t) = \int_0^1 \partial_t z^t dt, \\ \text{s.t.} \quad & \partial_t z^t = \int_{\mathcal{M}} \left( \frac{1}{2} \|v^t(x)\|^2 + U(x) + \|y - f(x^{t_s})\|^2 \right) \rho^t(x) dx - z^t, \\ & \partial_t \rho^t(x) + \nabla_x \cdot (\rho^t(x) v^t(x)) = \varepsilon \Delta_x \rho^t(x), \\ & \rho^0 = \rho_a, \quad \rho^1 = \rho_b, \quad \rho^{t_m} = \rho_m, \quad \forall m \in \{1, \dots, M\}. \end{aligned} \quad (46)$$

1280 Here, the inclusion of  $\|y - f(x^{t_s})\|^2$  in the Lagrangian produces the desired  $-\nabla_{x^t} \|y - f(x^{t_s})\|^2$   
1281 correction in the drift, while preserving the Schrödinger bridge structure and constraints.  
1282

1283 E RESNET PARAMETERIZATION FOR DISCRETE GEODESICS  
1284

## 1285 E.1 GEOMETRIC INTERPRETATION

1286 As stated in the main paper, our objective is to compute a geodesic  $\rho^t$  on  $\mathcal{P}^+(\mathcal{M})$ , induced by  
1287 the contact Hamiltonian dynamics, that is constrained to pass through a set of observed marginals  
1288  $\{\rho_a, \rho_m, \rho_b\}$  (i.e., discretized distributions along the probability path). These constraints naturally  
1289 lead to a discretized parameterization of  $\rho^t$ , where the overall density transformation is modeled as a  
1290 composition of maps, each connecting a pair of consecutive observations. In this context, a ResNet  
1291 architecture is ideally suited for this problem, as its sequential block structure directly mirrors this  
1292 piecewise, compositional nature of the approximated geodesic.  
1293

1294 Geometrically, each parameterized pushforward defines a vector  $\partial_t \rho_\theta^{t_k}$  in the tangent space of  $\rho_\theta^{t_{k-1}}$ ,  
1295 representing its change rate. The pair  $(\rho_\theta^{t_{k-1}}, \partial_t \rho_\theta^{t_k})$  thus corresponds to a point on the tangent  
bundle  $\mathcal{TP}^+(\mathcal{M})$ , with parameters  $\theta^k \in \Theta$  representing one of the possible coordinate charts for

1296 this update. As such, the parameter space  $\Theta$  forms a finite-dimensional subspace of  $\mathcal{TP}^+(\mathcal{M})$  (see  
1297 Fig. 5). The block transformation defines a smooth immersion  $T_{\theta^k} : \Theta \rightarrow \mathcal{TP}^+(\mathcal{M})$  with full-rank  
1298 Jacobian  $\nabla_x T_{\theta^k}$ , ensuring the pullback of the Wasserstein metric  $g^{\mathcal{W}_2}$  to  $\Theta$ , denoted  $T_{\theta}^* g^{\mathcal{W}_2}$ , is  
1299 well-defined and induces a Riemannian structure. This Riemannian metric identifies  $\Theta$  with its dual  
1300  $\Theta^*$  via the standard tangent–cotangent isomorphism (do Carmo, 1992). Consequently, the contact  
1301 Hamiltonian dynamics on  $\mathcal{T}^*\mathcal{P}^+(\mathcal{M}) \times \mathbb{R}$  can be equivalently expressed in the reduced phase space  
1302  $\Theta^* \times \mathbb{R}$  (Wu et al., 2025). At the same time, the Wasserstein manifold is approximated by the  
1303 finite-dimensional submanifold  $\mathcal{P}_{\theta}^+(\mathcal{M})$ , whose tangent space is  $\mathcal{TP}_{\theta}^+(\mathcal{M}) = \Theta$ .

## 1304 E.2 TRAINING ALGORITHM

---

### 1307 Algorithm 1 Training the Contact Wasserstein Geodesic (CWG) Framework

---

1308 **Input:** Dataset: samples from marginals  $x_{a, \{i,j\}} \sim \rho_a$ ,  $x_{b, \{i,j\}} \sim \rho_b$ ,  $x_{m, \{i,j\}} \sim \{\rho_m\}_{m=1}^N$ .

1309 **Output:** A trained ResNet  $T_{\{\theta^k\}_{k=0}^K}$ .

#### 1311 Part I: Initialization to Match the Initial Marginal

```

1312 1: for  $i = 1$  to  $E$  do                                 $\triangleright$  Epoch loop
1313 2:   for  $j = 1$  to  $B$  do                       $\triangleright$  Batch loop
1314 3:      $x^s \sim \lambda$                           $\triangleright$  Sample from reference distribution
1315 4:      $x^{t_0} = T_{\theta^0}(x^s)$                    $\triangleright$  Apply initial block
1316 5:      $\min_{\theta^0} d_{\mathcal{W}_2}^2(x^{t_0}, x_{a, \{i,j\}})$   $\triangleright$  Match initial marginal
1317 6:   end for
1318 7: end for

```

#### 1319 Part II: Geodesic Optimization

```

1320 8: for  $i = 1$  to  $E$  do
1321 9:   for  $j = 1$  to  $B$  do
1322 10:     $x^s \sim \lambda$                           $\triangleright$  Sample from reference distribution
1323 11:     $\{x^{t_0}, x^{t_1}, \dots, x^{t_K}\} = T_{\{\theta^k\}_{k=0}^K}(x^s)$   $\triangleright$  Full ResNet transformation
1324 12:     $\min_{\theta \setminus \theta^0} \ell(\{x^{t_k}\}_{k=0}^K, x_{b, \{i,j\}}, x_{m, \{i,j\}})$   $\triangleright$  Minimize geodesic loss (equation (9))
1325 13:   end for
1326 14: end for

```

---

## 1327 E.3 STABILITY CONSIDERATIONS

1329 The proposed solver discretizes the geodesic flow on the Wasserstein manifold and can be interpreted  
1330 as a single-shooting direct method for solving the optimal control problem of determining  
1331 the optimal set of parameters  $\{\theta^{k*}\}_{k=0}^K$ . Specifically, the ResNet-based parametrization  $T_{\{\theta^k\}_{k=0}^K}$   
1332 produces a discrete sequence of intermediate batches  $\{x_{\theta}^{t_k}\}_{k=0}^K$ , converting the continuous optimal  
1333 control problem into a nonlinear programming problem. This problem is then solved using standard  
1334 static optimization methods, with gradients computed through backpropagation. The loss function  
1335 in equation 9,

$$1337 \ell = \underbrace{d_{\mathcal{W}_2}^2(\rho_{\theta}^{t_K}, \rho_b)}_{\text{Terminal marginal}} + \underbrace{\sum_{m=1}^M d_{\mathcal{W}_2}^2(\rho_{\theta}^{t_{km}}, \rho_m)}_{\text{Intermediate marginals}} + \underbrace{\sum_{k=1}^K \Phi^{t_k} d_{\mathcal{W}_2}^2(\rho_{\theta}^{t_k}, \rho_{\theta}^{t_{k-1}})}_{\text{Energy minimization}},$$

1341 involving the computation of the Wasserstein distance, which can be explicitly rewritten as,

$$1343 \ell = \underbrace{\sum_{i,l}^N \pi_{il} \|x_i^{t_K} - y_l^{t_K}\|^2}_{\text{Terminal marginal}} + \underbrace{\sum_{m=1}^M \sum_{i,l}^N \pi_{il} \|x_i^{t_{km}} - y_l^{t_m}\|^2}_{\text{Intermediate marginals}} + \underbrace{\sum_{k=0}^{K-1} \sum_{i,l}^N \Phi^k \pi_{il} \|x_i^{t_{k+1}} - x_i^{t_k}\|^2}_{\text{Energy minimization}}, \quad (47)$$

1348 where  $\pi_{il}$  denotes the optimal transport plan between samples  $i$  and  $l$  the data batches, optimized  
1349 in order to minimize the transport cost and preserve the constraint,  $\sum_{il}^N \pi_{il} = 1$ , ensuring full  
mass conservation between batches. Importantly, the optimization of the transport plan  $\pi_{il}$  is a

1350 subproblem carried out for a fixed set of parameters  $\{\theta^k\}_{k=0}^K$ . This subproblem is a linear program  
 1351 whose solution depends Lipschitz-continuously on  $\{\theta^k\}_{k=0}^K$ , i.e.,  
 1352

$$1353 \quad \left\| \pi \left( \{\theta^{k'}\}_{k=0}^K \right) - \pi \left( \{\theta^k\}_{k=0}^K \right) \right\| \leq L_\pi \left\| \{\theta^{k'}\}_{k=0}^K - \{\theta^k\}_{k=0}^K \right\|, \quad (48)$$

1355 where  $L_\pi$  is a scalar constant. The two differences quantify, respectively, the change in the transport  
 1356 plan induced by a perturbation in the parameter set, and the perturbation in the parameter set itself.  
 1357

1358 The computation of the gradient of the loss in equation (47), which is responsible for updating the  
 1359 set of parameters according to,  
 1360

$$1360 \quad \{\theta_{(j+1)}^k\}_{k=0}^K = \{\theta_{(j)}^k\}_{k=0}^K - \alpha \nabla_\theta \ell, \quad (49)$$

1361 leads to the appearance of three types of terms:  
 1362

- 1363 •  $\nabla_\theta \pi_{il}$ , which is Lipschitz in  $\{\theta^k\}_{k=0}^K$  because the transport plan  $\pi$  is the solution of a con-  
 1364     vex optimization problem whose dependence on the parameters is Lipschitz-continuous;  
 1365
- 1366 •  $\nabla_\theta \|x_i^{t_k} - y_l^{t_k}\|^2$ , where  $x_i^{t_k}$  depends nonlinearly on  $\{\theta^k\}_{k=0}^K$  due to the nonlinear ResNet  
 1367     parametrization  $T_{\{\theta^k\}_{k=0}^K}$  used for expressivity. This implies that multiple different param-  
 1368     eter sets may minimize this term. However, around a local optimal point the problem of  
 1369     minimizing the Euclidean distance is locally convex;  
 1370
- 1371 •  $\nabla_\theta \Phi^{t_k}$ , which is a nonlinear function of  $\{x_\theta^{t_k}\}_{k=0}^K$ . If  $\Phi^{t_k}$  is at least piece-wise smooth,  
 1372     it is possible to identify a local minimum for  $\{x_\theta^{t_k}\}_{k=0}^K$ , and therefore a corresponding  
 1373     local minimum for  $\{\theta^k\}_{k=0}^K$ , depending on the initialization of the parameters. Even in the  
 1374     most complex scenario considered in this paper, the energy-varying SB case, the required  
 1375     smoothness assumptions are satisfied. Specifically, for  $\Phi^{t_k} = H(\rho^{t_k}, S^{t_k}, z^{t_k}) + U(\rho^{t_k}) +$   
 1376      $I(\rho^{t_k}) - \mathcal{B}(\rho^{t_k})$ ,  $H$  is a smooth function depending on a hyperparameter,  $U$  is implemented  
 1377     either via at least piecewise smooth learning models (Apps. G.4, G.5, G.7) or via a LAND  
 1378     metric (Arvanitidis et al., 2018) with a smooth kernel (Apps. G.2, G.3), and both  $I$  and  $\mathcal{B}$   
 1379     are smooth functions.

1380 Therefore, each update  $\{\theta_{(j)}^k\}_{k=0}^K \rightarrow \{\theta_{(j+1)}^k\}_{k=0}^K$  does not destabilize the training and  
 1381 converges toward a local minimum, which geometrically corresponds to one of the possible coordinate  
 1382 charts that can describe the tangent bundle  $\mathcal{TP}^+(\mathcal{M})$  (see Appendix E.1 for the geometrical  
 1383 interpretation). The stability of the gradient for long time horizons is moreover guaranteed by the  
 1384 ResNet architecture itself, where the skip connections prevent detrimental effects such as exploding  
 1385 or vanishing gradients (Zaeemzadeh et al., 2020).  
 1386

#### 1388 E.4 GEODESIC PARAMETERIZATION: TIME AND ARC-LENGTH PERSPECTIVES

1389 The proposed ResNet architecture  $T_{\{\theta^k\}_{k=0}^K}$  discretizes the geodesic  $\rho^t$  on  $\mathcal{P}^+(\mathcal{M})$  into a piecewise-  
 1390 linear path segmented at nodes  $\{\rho_\theta^{t_k}\}_{k=0}^K$ , thereby defining a specific sampling of points along the  
 1391 geodesic. By default, the discretization obtained by minimizing the curve energy,  
 1392

$$1394 \quad d_{\mathcal{W}_2}^2(\rho_a, \rho_b) = \sum_{k=1}^K d_{\mathcal{W}_2}^2(\rho_\theta^{t_k}, \rho_\theta^{t_{k-1}}), \quad (50)$$

1393 corresponds to a uniform arc-length discretization  $s$ , where each segment has approximately the  
 1394 same length, i.e.,  $d_{\mathcal{W}_2}^2(\rho_\theta^{t_k}, \rho_\theta^{t_{k-1}}) \approx C$ . The relationship between the arc-length discretization  $s$   
 1395 and the time discretization  $t$ , which describes the progression along the geodesic with respect to  
 1396 the time variable, is given by  $ds = \sqrt{\mathcal{K}_{\text{phy}}} dt$ , where  $\mathcal{K}_{\text{phy}}$  denotes the kinetic energy of the system  
 1397 whose dynamics is interpreted as a geodesic. In the energy-conserving case,  $\mathcal{K}_{\text{phy}}$  remains constant,  
 1398 and the time and arc-length discretizations differ only by a constant rescaling of the geodesic length.  
 1399 Nevertheless, the discretization nodes  $\{\rho_\theta^{t_k}\}_{k=0}^K$  remain unchanged, uniformly distributed in space  
 1400 and in time. In contrast, when modeling an energy-varying system whose associated trajectory  
 1401

1404 projects onto a geodesic, the curve energy takes the form,  
 1405

$$1406 \quad 1407 \quad 1408 \quad d_{\mathcal{W}_2}^2(\rho_a, \rho_b) = \sum_{k=1}^K \Phi^{t_k} d_{\mathcal{W}_2}^2(\rho_\theta^{t_k}, \rho_\theta^{t_{k-1}}), \quad (51)$$

1409 where the scaling factor  $\Phi^{t_k} = H(\rho^{t_k}, S^{t_k}, z^{t_k}) - \mathcal{F}(\rho^{t_k}) - \mathcal{B}(\rho^{t_k})$  accounts for the use of the  
 1410 Jacobi metric. This scaling factor also determines the length of each segment, which is inversely  
 1411 proportional to it, i.e.,  $d_{\mathcal{W}_2}^2(\rho_\theta^{t_k}, \rho_\theta^{t_{k-1}}) \approx \frac{C}{\Phi^{t_k}}$ . Consequently, the arc-length discretization  
 1412  $d_{\mathcal{W}_2}(\rho_\theta^{t_k}, \rho_\theta^{t_{k-1}}) = \Delta s$  is non-uniform, while the time discretization  $\Delta t = \sqrt{\frac{C}{\kappa}}$  remains uniform,  
 1413 as the nodes are assumed to be sampled at constant time intervals. The origin of this discrepancy  
 1414 lies in the kinetic energy term relating the arc-length and time increments,  $\Delta s = \sqrt{\mathcal{K}_{\text{phy}} \Delta t}$ , which  
 1415 is no longer constant in the energy-varying case.  
 1416

1417 This reasoning provides insight into the relationship between the energy term  $\Phi^{t_k}$  used in the  
 1418 geodesic minimization and the physical kinetic energy of the modeled system. Specifically, since  
 1419 the geodesic minimization problem yields,

$$1420 \quad 1421 \quad d_{\mathcal{W}_2}^2(\rho_\theta^{t_k}, \rho_\theta^{t_{k-1}}) = \frac{C}{\Phi^{t_k}}, \quad (52)$$

1422 while, as established above,

$$1424 \quad 1425 \quad d_{\mathcal{W}_2}^2(\rho_\theta^{t_k}, \rho_\theta^{t_{k-1}}) = \Delta s^2 = \mathcal{K}_{\text{phy}} \Delta t^2 = \mathcal{K}_{\text{phy}} \frac{C}{\kappa}, \quad (53)$$

1426 we obtain  $\Phi^{t_k} = \frac{\kappa}{\mathcal{K}_{\text{phy}}}$ . This result shows that the energy employed in the geodesic computation  
 1427 is inversely proportional to the physical kinetic energy, scaled by the constant  $\kappa$ . This relationship  
 1428 arises from the coupling between the arc-length and time discretization of the system's trajectory.  
 1429

## 1431 E.5 TIME COMPLEXITY AND PRACTICAL CONSIDERATIONS

1432 At each iteration of the geodesic optimization of Alg. 1, we sample  $N$  points  $x^s \in \mathcal{M}$  of dimension  
 1433  $D$  from the reference distribution  $\lambda$  and pass them through the ResNet. Assuming its architecture  
 1434 consists of  $K$  blocks, each being an MLP of  $L$  layers and hidden dimension  $W$ , the computational  
 1435 cost of this forward pass is  $\mathcal{O}(N D K L W)$ . The loss function  $\ell$  in equation (9) requires  $M + K$   
 1436 evaluations of the Wasserstein distance between sample batches, and  $K$  evaluations of the factor  
 1437  $\Phi^{t_k}$ . For the Wasserstein distance, we employ the `geomloss` library (Feydy, 2020), which uses  
 1438 the Sinkhorn algorithm with time complexity  $\mathcal{O}(N(D + T_{\text{sh}}))$ , where  $T_{\text{sh}}$  is the number of Sinkhorn  
 1439 iterations until convergence (Feydy, 2020). To evaluate  $\Phi^{t_k}$ , we apply its complete definition,  
 1440

$$1442 \quad 1443 \quad \Phi^{t_k} = H^{t_k} + \int_{\mathcal{M}} \left[ U(x^{t_k}) - 2\varepsilon(\log(2\rho_\theta^{t_k}) - 1) \right] \rho_\theta^{t_k} dx + \frac{1}{2}\varepsilon^2 I(\rho_\theta^{t_k}), \quad (54)$$

1444 to a batch of samples. As discussed in Sec. 3, the Fisher information  $I(\rho_\theta^{t_k})$  in the potential energy  
 1445 originates from the entropy regularization in the SB formulation and does not require explicit  
 1446 computation. Its effect is implicitly captured by the entropy-regularized Wasserstein distance. Therefore,  
 1447 evaluating  $\Phi^{t_k}$  reduces to computing the scalar functions  $H$  and  $U$ , along with estimating the  
 1448 entropy term  $\log(2\rho_\theta^{t_k}) \rho_\theta^{t_k}$ , which is the computational bottleneck. This term can be estimated using  
 1449 a  $k$ -NN entropy estimator with time complexity  $\mathcal{O}(D N \log N)$  (Borelli et al., 2022). Considering  
 1450 all components, and given that  $K \geq M$ , the overall time complexity becomes,  
 1451

$$1452 \quad 1453 \quad \mathcal{O}(N d K L W) + \mathcal{O}(N(K+M)(D+T_{\text{sh}})) + \mathcal{O}(K d N \log N) \approx \mathcal{O}\left(N K (T_{\text{sh}} + D(LW + \log N))\right). \quad (55)$$

1454 Our method demonstrates highly favorable scaling properties, offering a significant advantage over  
 1455 existing approaches. Notably, its computational complexity scales linearly with data dimensionality  
 1456  $D$  and nearly linearly with the batch size  $N$ . Furthermore, our model's performance is only weakly  
 1457 influenced by the number of marginals and it circumvents the expensive outer iteration loops. This  
 1458 stands in stark contrast to existing methods, as elaborated next:

1458 • Hong et al. (2025) scales quadratically in  $N$  and exponentially in  $D$ , with time complexity,  
 1459

$$1460 \quad \mathcal{O}(T_{\text{iter}} DKN^2 \exp(\alpha(D + 1))) , \quad (56)$$

1461 where  $T_{\text{iter}}$  denotes the number of Belief Propagation iterations required for convergence,  
 1462 and  $\alpha$  is a parameter determined by the order of the Gaussian process and structural as-  
 1463 sumptions on their representation basis.

1464 • Chen et al. (2023) proposed a method with time complexity,  
 1465

$$1466 \quad \mathcal{O}(T_{\text{iter}} KN(2D)^2) , \quad (57)$$

1467 where  $T_{\text{iter}}$  denotes the number of Bregman iterations needed for convergence. The  
 1468 quadratic dependence on  $2D$  arises from the state space being doubled by incorporating  
 1469 velocity, which comes at the cost of computational efficiency.

1470 • Shen et al. (2025) employed a score-matching SB solver (Vargas et al., 2021) for the mmSB  
 1471 problem with computational complexity,  
 1472

$$1473 \quad \mathcal{O}(T_{\text{iter}} MN(2D)) , \quad (58)$$

1474 where  $T_{\text{iter}}$  again denotes the number of Bregman iterations needed for convergence, with  
 1475 reported values between 2000 and 4000. Although the method is nearly linear in both  $N$   
 1476 and  $D$ , it depends explicitly on the number of intermediate marginals  $M$  and suffers from  
 1477 additional outer-loop iterations.

1478 As highlighted in this analysis, our CWG method advances the state of the art as a remarkably  
 1479 efficient computational approach: it scales linearly in all relevant terms and does not require iterations.  
 1480 The baselines (Shi et al., 2023; Bortoli et al., 2024; Liu et al., 2024) do not explicitly report  
 1481 their computational efficiency, so the comparison with our method is addressed empirically in the  
 1482 experiments.

## F IMPLEMENTATION DETAILS

### F.1 EMPIRICAL APPROXIMATION OF THE WASSERSTEIN-2 DISTANCE

1483 Let  $\rho_c$  and  $\rho_d$  be two probability distributions from which we draw batches of samples  $\{x_{c,i}\}_{i=1}^N \sim$   
 1484  $\rho_c$  and  $\{x_{d,j}\}_{j=1}^M \sim \rho_d$ , respectively. The Wasserstein-2 distance, denoted by  $d_{\mathcal{W}_2}(\rho_c, \rho_d)$ , measures  
 1485 the minimal cost of transporting mass between these two distributions. To approximate this distance  
 1486 empirically, we first construct a cost matrix  $C \in \mathbb{R}^{N \times M}$ , where each entry,

$$1492 \quad C_{ij} = \|x_{c,i} - x_{d,j}\|^2 , \quad (59)$$

1493 represents the squared Euclidean distance between sample  $x_{c,i}$  and sample  $x_{d,j}$ . A transport plan is  
 1494 then defined as a matrix  $\pi \in \mathbb{R}_+^{N \times M}$  that assigns how much mass to move from each  $x_{c,i}$  to each  
 1495  $x_{d,j}$ , minimizing the total transport cost weighted by  $C$ . To avoid degenerate solutions where all  
 1496 mass is concentrated on a few points, entropy regularization is introduced, encouraging smoother  
 1497 and more distributed transport plans. For this computation, we employ the `SamplesLoss` function  
 1498 from the `geomloss` library (Feydy, 2020), with parameters resumed in Table 9.

1499 In the conditional generation setting, the probability flow  $\rho^t(x | y)$  is conditioned on the feature  $y =$   
 1500  $f(x^{t_s})$ , with  $x^{t_s} \sim \rho_s$ , to ensure that the generated samples satisfy  $y$  at time step  $t_s$ . Therefore, when  
 1501 comparing the prescribed distribution  $\rho_s$  with a marginal  $\rho_m$ , a modified Wasserstein-2 distance  
 1502  $d'_{\mathcal{W}_2}(\rho_s, \rho_m)$  incorporating the feature penalty is used. This distance is defined via the cost matrix

$$1504 \quad C_{ij} = \|x_i^{t_s} - x_{m,j}\|^2 + \|y - f(x_{m,j})\|^2 , \quad (60)$$

1505 which penalizes transport plans assigning mass to samples  $x_{m,j}$  inconsistent with the conditioning  
 1506 feature  $y$ .

### F.2 TEASER IMAGE

1507 The system illustrated in Figure 1 depicts the time parameterization of three different Schrödinger  
 1508 bridges connecting two 1D Gaussian distributions with identical variance but different means. The

1512 energy-conserving bridge (—) corresponds to the solution of the mmGSB problem (2) in the specific  
 1513 case of  $U = 0$  and without intermediate marginals  $\rho_m$ . The computed bridge can be interpreted as  
 1514 a piecewise geodesic on the Wasserstein manifold  $\Sigma$  endowed with the pullback Wasserstein metric  
 1515  $T^*g^W$ .

1516 In contrast, the two energy-varying curves (—, —) instead correspond to the solutions of  
 1517 the NCGSB problem (4) in the case of no intermediate marginals and  $U = 0$ . These  
 1518 trajectories are piecewise geodesics with respect to the pullback Jacobi metric  $T_\theta^*\tilde{g}_J =$   
 1519  $(H(\rho_\theta^{t_k}) - \mathcal{F}(\rho_\theta^{t_k}) - \mathcal{B}(\rho_\theta^{t_k}))T^*g^{\mathcal{W}_2}$ . Here, the potential energy term  $\mathcal{F} = -U - I$ , re-  
 1520 duces to the Fisher information contribution, yielding the equivalent expression  $T_\theta^*\tilde{g}_J =$   
 1521  $(H(\rho_\theta^{t_k}) + I(\rho_\theta^{t_k}) - \mathcal{B}(\rho_\theta^{t_k}))g^{\mathcal{W}_2}$ . Minimizing the geodesic energy with respect to this metric  
 1522 yields trajectories that tend to reduce the scaling factor  $(H(\rho_\theta^{t_k}) + I(\rho_\theta^{t_k}) - \mathcal{B}(\rho_\theta^{t_k}))$ . As a result,  
 1523 these paths preferentially traverse low-energy regions of  $\Sigma$  while concentrating the discretization  
 1524 nodes  $\{\rho_\theta^{t_k}\}_{k=0}^K$  in the higher-energy portions of the curve, since long segments between nodes in  
 1525 such regions are penalized by the large value of the scaling factor. Specifically, the geodesic com-  
 1526 putation exhibits the following behaviors:

- 1528 • The minimization of the Fisher information  $I(\rho_\theta^{t_k})$ , promoting smoother and more regular  
 1529 trajectories by penalizing rapid variations in the density. This ensures stable and coherent  
 1530 dynamics along the path.
- 1531 • The maximization of the entropy term  $\mathcal{B}(\rho_\theta^{t_k})$ , upper-bounded by the Hamiltonian  $H(\rho_\theta^{t_k})$ ,  
 1532 since the metric  $T_\theta^*\tilde{g}_J$  must remain positive. Consequently, higher Hamiltonian values  
 1533 correspond to higher-entropy paths.
- 1534 • The concentration of the discretization nodes in regions of higher energy (i.e., where  
 1535  $H(\rho_\theta^{t_k})$  is larger). As the evolution of the Hamiltonian can be prescribed, selecting a de-  
 1536 creasing or increasing law results in a stochastic trajectory that slows down or speeds up  
 1537 over time, respectively, as discussed in Appendix E.4.

1538 These behaviors can be observed in the two energy-varying probability paths illustrated in Figure 1.  
 1539 In the red curve (—), the Hamiltonian function, ranging from 1.5 to 0.5, is defined as,

$$1541 \quad H^{t_k} = 0.5 + \frac{\mu(x^{t_K}) - \mu(x^{t_k})}{\mu(x^{t_K}) - \mu(x^{t_0})}, \quad (61)$$

1543 where  $\mu(x^{t_k})$  denotes the mean of the samples  $x^{t_k} \sim \rho_\theta^{t_k}$ . The Hamiltonian is thus decreasing,  
 1544 producing a dynamics that evolves slowly at the beginning and accelerates toward the end of the  
 1545 time interval. The opposite behavior is observed in the purple curve (—), where the Hamiltonian  
 1546 increases from 0.01 to 1.01,

$$1548 \quad H^{t_k} = 1.01 - \frac{\mu(x^{t_K}) - \mu(x^{t_k})}{\mu(x^{t_K}) - \mu(x^{t_0})}. \quad (62)$$

1550 In this case, the dynamics gradually slows down over time. Moreover, the lower overall energy  
 1551 level results in reduced entropy, as evidenced by the narrower Gaussian distribution observed at the  
 1552 midpoint of the bridge in the purple curve compared to the red one.

## 1553 G EXTENDED RESULTS

### 1554 G.1 EXPERIMENTAL SETUP

1557 The LiDAR Manifold Navigation and Cell Sequencing experiments were conducted on a machine  
 1558 equipped with 13th Gen Intel® Core™ i7-13850HX CPUs. The Image Generation experiment was  
 1559 run on a system with an NVIDIA GeForce RTX5090 GPU (32GB VRAM, CUDA12.9, driver ver-  
 1560 sion 575.64.03). Results for the LiDAR Manifold Navigation (Table 2), Single Cell Sequencing  
 1561 (Table 3), Sea Temperature Prediction (Table 4 and Appendix G.4), Robot Task Reconstruction  
 1562 (Table 8), FFHQ Transfer (Table 7), and MNIST-to-EMNIST (Table 38) experiments are based on  
 1563 a total of ten evaluations obtained from training runs with different initial conditions. The tables  
 1564 report the mean and standard deviation of these distributions.

1565 The ResNet architecture varies depending on the task. For the LiDAR Manifold Navigation, the  
 1566 Cell Sequencing and the Unpaired Transfer experiments, each block is an MLP that processes the

1566 output of the previous block and generates an update, which is added to the input with a step size  
 1567  $\tau$ :  $x \leftarrow x + \tau \text{block}(x)$ . Details of this architecture are provided in Table 11. In contrast, for the  
 1568 Image Generation experiment, the input consists of images, and each block is implemented as a 2D  
 1569 U-Net. Details of this architecture are provided in Table 12. The total number of CWG parameters  
 1570 used across the different experiments, alongside the baselines, is summarized in Table 13.  
 1571

Parameter	Value
Entropy	0.05
Euclidean norm order	2
Scaling	0.7

Table 9: SampleLoss parameters.

Parameter	Lidar	Cell	Sea	Robot	FFHQ	Mnist
# samples $N$	1000	1000	200	100	1000	300
Weight $w_b$	10	10	100	10	1	1
Weight $w_m$	—	10	100	—	—	—
Weight $w_g$	1	1	1	1	1	1

Table 10: Loss function (9) parameters for the experiments.

Component	ResNet		
	Lidar	Cell	FFHQ
Number of blocks $K$	20	5	6
Layers per block	3	3	3
Layer hidden size	30	128	1024
Step size $\tau$	0.1	0.1	0.1
Input dimension $d$	3	5	512

Table 11: Configuration of the ResNet

Component	Image ResNet		
	Sea	Robot	Mnist
Number of blocks $K$	5	8	6
Input channels	1	3	1
Output channels	1	3	1
Layers per block	1	2	2
Downsampling blocks	2 (32 and 64 channels)		
Upsampling blocks	2 (32 and 64 channels)		
Step size $\tau$	1	1	1
Input dimension $d$	4096	12288	784

Table 12: Configuration of the Image ResNet.

1597 During training, the weights  $\{w_b, w_m, w_g\}$  balancing the loss terms in equation (9),  
 1598

$$1599 \ell = w_b d_{\mathcal{W}_2}^2(\rho_{\theta}^{t_K}, \rho_b) + w_m \sum_{m=1}^M d_{\mathcal{W}_2}^2(\rho_{\theta}^{t_{k_m}}, \rho_m) + w_g \sum_{k=1}^K \Phi^{t_k} d_{\mathcal{W}_2}^2(\rho_{\theta}^{t_k}, \rho_{\theta}^{t_{k-1}}),$$

1602 along with the number of samples  $N$  used for Wasserstein distance estimation, are experiment-  
 1603 specific and summarized in Table 10. The GPU memory consumption for the two image generation  
 1604 experiments, comparing our CWG method with the baseline approaches, is reported in Table 14.  
 1605 Due to the explicit handling of the full probability distribution (albeit in discretized form) within the  
 1606 ResNet architecture, CWG exhibits particularly high memory requirements. In contrast, methods  
 1607 such as GSBM and DSBM model only the drift component and subsequently integrate the dynamics.  
 1608 While this makes them significantly more demanding in terms of computation time, they are more  
 1609 memory-efficient than CWG.

Experiment	DSBM	SB-Flow	GSBM	SBIRR	DM-SB	CWG
LiDAR Navigation	0.14 M	—	0.33 M	—	—	0.017 M
Cell Sequencing	—	—	—	0.013 M	2.5 M	0.017 M
Sea Temperature	79.2 M	11 M	271 M	—	—	5.2 M
Robot Task	79.2 M	11 M	271 M	—	—	7.3 M
Unpaired Transfer	5.7 M	6.6 M	5.7 M	—	—	15.7 M
Mnist-to-Emnist	4.6 M	3.3 M	4.6 M	—	—	5.2 M

Table 13: Number of parameters (in millions, M) of the models reported in Table 1

Methodology	CWG	GSBM	DSBM
Sea Temperature (MB)	$25200 \pm 200$	$9600 \pm 300$	$9000 \pm 200$
Robot Task (MB)	$16100 \pm 200$	$12500 \pm 300$	$10200 \pm 200$

Table 14: Comparison of GPU memory consumption across the methods evaluated in the Image Generation experiment. CWG shows a decrease in memory usage for the Robotic Task Reconstruction experiments, whereas the other methods exhibit an increase due to the reduced batch size used in this training (Table 10).

## G.2 LiDAR MANIFOLD NAVIGATION

The LiDAR dataset (OpenTopography, 2025) consists of point clouds contained in the domain  $[-5, 5]^3 \subset \mathbb{R}^3$ . The objective of the experiment is to construct a bridge across the data manifold for connecting two distributions while avoiding regions of high elevation and remaining closely aligned with the manifold structure. The initial distribution  $\rho_a$  is composed by a mixture of 4 Gaussian distributions, the target distribution  $\rho_b$  is composed of 2 Gaussians on the two sides of the mountain. The manifold shape is incorporated in the problem through the potential function  $U$ , inherited from the baseline (Liu et al., 2024),

$$\int_{\mathcal{M}} U(x) \rho^t(x) dx = \int_{\mathcal{M}} (U_{\text{manifold}}(x) + U_{\text{height}}(x)) \rho^t(x) dx, \quad (63)$$

$$U_{\text{manifold}}(x) = w_{\text{manifold}} \|\psi(x) - x\|^2, \quad U_{\text{height}}(x) = w_{\text{height}} \|\psi^{(z)}(x)\|^2.$$

Here,  $\psi(x)$  denotes the projection of a point  $x$  onto an approximate tangent plane, estimated from its  $p$  nearest neighbors on the data manifold. The notation  $\psi^{(z)}(x)$  refers to the  $z$ -coordinate of  $\psi(x)$ , i.e., the height of the fitted plane. The weights  $w_{\text{manifold}}$  and  $w_{\text{height}}$  control the relative importance of the two terms in the potential function. We now detail the construction of  $\psi(x)$ . Let  $N_p(x) = \{x_1^l, \dots, x_p^l\}$  denote the set of  $p$  nearest neighbors of  $x \in \mathbb{R}^3$  in the dataset. To approximate the local tangent plane, we employ a moving least-squares (MLS) procedure (Levin, 1998). Specifically, the plane parameters  $(a, b, c)$  are obtained by solving,

$$\min_{a, b, c} \frac{1}{p} \sum_{i=1}^p w(x, x_i^l) \left( ax_i^{l(x)} + bx_i^{l(y)} + c - x_i^{l(z)} \right)^2, \quad (64)$$

where the superscripts indicate coordinates and the weights are defined as,

$$w(x, x_i^l) = \exp\left(-\frac{\|x - x_i^l\|}{\gamma}\right), \quad (65)$$

with  $\gamma$  being a scaling parameter. Given the fitted plane, the projection operator  $\psi(x)$  is defined as,

$$\psi(x) = x - \frac{x^\top n + c}{\|n\|^2} n, \quad n = [a \ b \ -1]^\top, \quad (66)$$

where  $n$  denotes the plane's normal vector. Differentiation through  $\psi$  naturally restricts gradients to this tangent plane, thereby ensuring that optimization of the state cost  $U$  evolves within the geometry of the data manifold. The values of the parameters used in the computation of the projection operator  $\psi(x)$  are provided in Table 15.

The quantitative results reported in Table 2 aim to characterize the two key aspects of the Schrödinger Bridge problem, as formulated in the stochastic optimization problems (2) and (4). The Optimality metric reports the length of the stochastic bridge, computed as the square root of the cost functional  $J$ . It quantifies the ability of the approach to find low transport-cost solutions connecting the marginals  $\rho_a$  and  $\rho_b$ . In contrast, the Feasibility metric is computed as the Wasserstein distance between the marginals and the nodes of the discretized bridge approximating them. This metric measures how well the boundary constraints in the optimization problems are satisfied. Formally, these metrics are defined as,

$$\text{Optimality} = \sqrt{J}; \quad (67a)$$

$$\text{Feasibility} = d_{\mathcal{W}_2}(\rho_\theta^{t_0}, \rho_a) + d_{\mathcal{W}_2}(\rho_\theta^{t_K}, \rho_b). \quad (67b)$$

1674  
1675

In the guided generation setting, the feature function  $f$ , defines as follows,

1676  
1677

$$f(x_b) = \text{ReLU}(w_x x_b^{(x)} - b_x) + \text{ReLU}(w_y x_b^{(y)} - b_y), \quad x_b \sim \rho_b \quad (68)$$

1678  
1679  
1680  
1681

is used to penalize samples from the terminal marginal distribution  $\rho_b$  on the left side of the mountain. The parameters used in this experiment are listed in Table 16, and the results of guidance fine-tuning are reported in Table 17. The training time (tt) metric shows that guiding the generation requires only 10.7% of the original training time.

1682  
1683

Table 15: Potential function  $U$  parameters for the LiDAR Manifold Navigation experiment.

Parameter	Value
Weight manifold $w_{\text{manifold}}$	5
Weight height $w_{\text{height}}$	1
Spatial scaling $\gamma$	0.1
# neighbor points $p$	20

1684  
1685  
1686  
1687  
1688  
1689  
1690

Table 17: Bridge energy  $J$  (↓) with penalty  $f$  in the LiDAR Manifold Navigation task, reported for our CWG method before and after guidance fine-tuning.

1691  
1692  
1693

Metric	before	after
$J$	$12.31 \pm 0.18$	$2.49 \pm 0.43$
tt (s)	$280 \pm 20$	$310 \pm 25$

1694

### G.3 SINGLE CELL SEQUENCING

1695  
1696  
1697  
1698  
1699  
1700  
1701  
1702  
1703  
1704  
1705  
1706  
1707  
1708  
1709  
1710  
1711  
1712

The Embryoid Body (EB) stem cell differentiation dataset (Moon et al., 2019) captures cell state progression across five developmental stages  $[t_0, t_1, t_2, t_3, t_4]$  over a 27-day period. Snapshots were collected at five discrete time intervals:  $t_0 \in [0, 3]$ ,  $t_1 \in [6, 9]$ ,  $t_2 \in [12, 15]$ ,  $t_3 \in [18, 21]$ , and  $t_4 \in [24, 27]$ . These stages involve significant structural changes, with cells moving and reorganizing within increasingly stiff tissue while consuming and releasing mechanical energy (Zeevaert et al., 2020; Kinney et al., 2014). Consequently, the resulting dynamics exhibit energy dissipation and are better described by the NCGSB framework than by energy-conserving models. In this experiment, we evaluate the framework’s ability to generalize to regions with no available data by dividing the dataset into a training set  $[t_0, t_2, t_4]$  and a validation set  $[t_1, t_3]$ . The geometry of the data manifold is incorporated into the NCGSB problem through a potential function  $U$ , defined as,

$$U(x^t) = \frac{1}{N_1} \sum_{i=1}^{N_1} \left[ \frac{1}{N_2} \sum_{j=1}^{N_2} e^{\frac{1}{\gamma} \|x_j^t - x_i\|^2} (x_j^t - x_i)^2 \right]^{-1}, \quad (69)$$

1713  
1714  
1715  
1716  
1717  
1718  
1719  
1720  
1721  
1722  
1723  
1724  
1725  
1726  
1727

where  $x_j^t \sim \rho^t$  denotes a sample from the posterior distribution, and  $x_i \sim \{\rho_a, \rho_m, \rho_b\}$  are samples from the marginal distributions.  $N_1$  and  $N_2$  indicate the number of samples taken from  $\rho^t$  and  $\{\rho_a, \rho_m, \rho_b\}$ , respectively, while  $\gamma$  represents a spatial scaling parameter. The exponential term acts as a kernel-like weight that reduces the influence in the inverse summation  $\sum_{i=1}^{N_1} [\cdot]^{-1}$  of the data points far from the bridge. Globally, the potential function  $U$  measures the distance of the bridge  $\rho^t$  from the available data  $\{\rho_a, \rho_m, \rho_b\}$  (the training set) and its minimization guides the construction of a bridge that stays close to the known manifold while generalizing effectively to regions without data (the validation set). While other approaches leveraging this dataset (Tong et al., 2020; Shen et al., 2025) first embed the data into a 100-dimensional feature space using principal component analysis (PCA) and then restrict the analysis to the first five dimensions, this procedure excessively linearizes and flattens the data manifold, making navigation trivial and eliminating the need for intermediate marginals (Shen et al., 2025). To better preserve the manifold’s geometry, we instead apply the PHATE algorithm (Moon et al., 2019) to the 100-dimensional representation, producing a 5-dimensional nonlinear embedding that more faithfully captures the original structure.

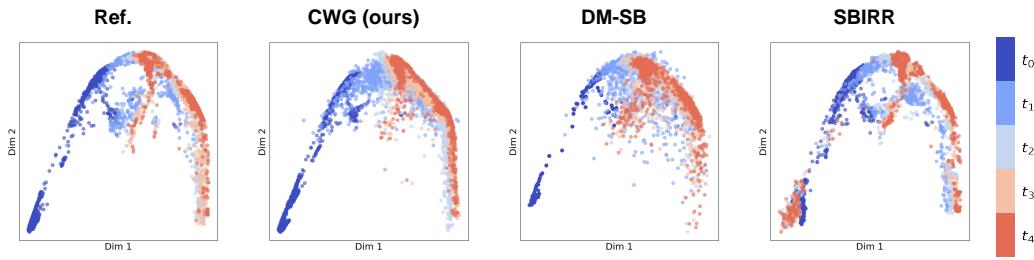
1728  
1729

In the results reported in Tables 18 and 19, corresponding to the cell dynamics reconstruction shown in Fig. 12, the scalar Hamiltonian function  $H^{t_k}$  in the pullback Jacobi metric,  $T_\theta^* \tilde{g}_J = (H^{t_k} -$

1728  $\mathcal{F}(\rho^t) - \mathcal{B}(\rho^t)$   $T_\theta^* g^{\mathcal{W}_2}$  used in the CWG approach, is linearly varied from an initial value  $H^{t_0} =$   
 1729 0.82 to a final value  $H^{t_K} = 1$ , along the trajectory defined by the first principal dimension  $D_1$ , i.e.,  
 1730

$$1731 \quad H^{t_k} = \frac{\mu(x_{D_1}^{t_K}) - \mu(x_{D_1}^{t_k})}{\mu(x_{D_1}^{t_K}) - \mu(x_{D_1}^{t_0})} (H^{t_0} - H^{t_K}) + H^{t_0}, \quad x^{t_k} \sim \rho^{t_k}. \quad (70)$$

1734 This law is treated as a hyperparameter of the methodology. All ten experiments reported are con-  
 1735 ducted under this Hamiltonian behavior. This increasing trend, associated with the reduction of  
 1736 physical energy (see Appendix E.4), aligns with the underlying biological process and has been  
 1737 shown to produce better results than the energy-conserving case, where  $H^{t_k} = 1$  is kept constant  
 1738 (see Table 21 and Figure 13a). As detailed in Appendix E.2, the behavior of  $H^{t_k}$  plays a crucial  
 1739 role in determining how quickly the nodes of the stochastic path depart from the initial reference  
 1740 distribution and approach the subsequent reference marginal. In Figure 13b, the dataset exhibits a  
 1741 decreasing Wasserstein distance  $d_{\mathcal{W}_2}$  between successive cell snapshots, reflecting that cell differen-  
 1742 tiation progresses more rapidly at early stages (Zeevaert et al., 2020). Training with the Hamiltonian  
 1743 behavior defined in equation (70) enables the model to capture this trend more accurately.  
 1744



1755 Figure 12: Reconstructions from the models in the Single Cell Sequencing experiment.

1762 Table 18: Wasserstein error ( $\downarrow$ ) for all the  
 1763 time points in Single Cell Sequencing.

Metric	CWG (ours)	DM-SB	SBIRR
$d_{\mathcal{W}_2}(x^{t_0})$	<b>0.10</b> $\pm$ 0.01	0.59 $\pm$ 0.01	0.86 $\pm$ 0.02
$d_{\mathcal{W}_2}(x^{t_1})$	<b>1.11</b> $\pm$ 0.06	2.25 $\pm$ 0.01	1.92 $\pm$ 0.02
$d_{\mathcal{W}_2}(x^{t_2})$	<b>0.16</b> $\pm$ 0.01	1.17 $\pm$ 0.01	1.05 $\pm$ 0.02
$d_{\mathcal{W}_2}(x^{t_3})$	<b>0.33</b> $\pm$ 0.02	1.64 $\pm$ 0.03	1.86 $\pm$ 0.02
$d_{\mathcal{W}_2}(x^{t_4})$	<b>0.11</b> $\pm$ 0.02	1.03 $\pm$ 0.01	1.38 $\pm$ 0.02

1764 Table 19: Maximum Mean Discrepancy ( $\times 10^{-2}$ )  
 1765 ( $\downarrow$ ) for the time points in Single Cell Sequencing.

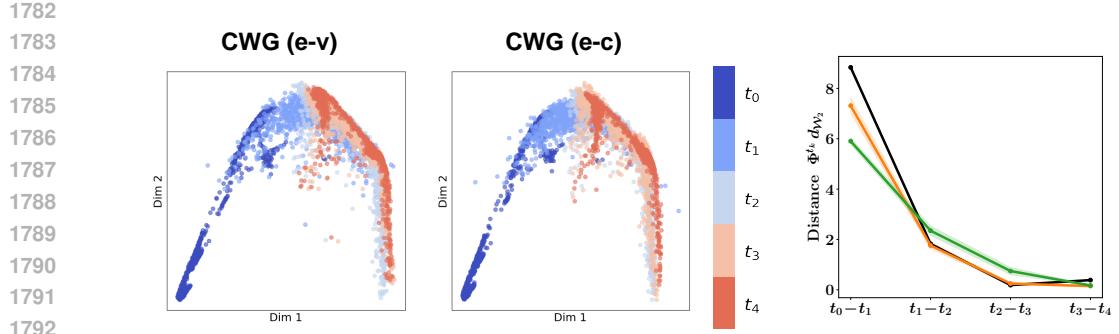
Metric	CWG (ours)	DM-SB	SBIRR
$\text{MMD}(x^{t_0})$	0.9 $\pm$ 0.08	<b>0.3</b> $\pm$ 0.02	1.0 $\pm$ 0.11
$\text{MMD}(x^{t_1})$	<b>7.0</b> $\pm$ 0.21	10.1 $\pm$ 0.33	9.9 $\pm$ 0.48
$\text{MMD}(x^{t_2})$	<b>1.2</b> $\pm$ 0.12	2.9 $\pm$ 0.18	2.5 $\pm$ 0.17
$\text{MMD}(x^{t_3})$	<b>3.5</b> $\pm$ 0.22	4.3 $\pm$ 0.20	6.7 $\pm$ 0.44
$\text{MMD}(x^{t_4})$	<b>0.2</b> $\pm$ 0.05	1.5 $\pm$ 0.09	5.9 $\pm$ 0.36

1775 Table 20: Single Cell Sequencing Parameters.

Parameter	Value
Spatial scaling $\gamma$	0.3
# bridge samples $N_1$	1000
# marginal samples $N_2$	3000

1776 Table 21: Wasserstein error at validation ( $\downarrow$ ) and training  
 1777 time (tt) ( $\downarrow$ ) in the ablation study comparing the energy-  
 1778 varying (e-v) and energy-conserving (e-c) versions of  
 1779 CWG on the Single Cell Sequencing task.

Metric	CWG (e-v)	CWG (e-c)
$d_{\mathcal{W}_2}(x^{t_1})$	<b>1.11</b> $\pm$ 0.06	1.31 $\pm$ 0.06
$d_{\mathcal{W}_2}(x^{t_3})$	<b>0.33</b> $\pm$ 0.02	0.49 $\pm$ 0.03
tt (s)	710 $\pm$ 30	710 $\pm$ 30



1782  
1783  
1784  
1785  
1786  
1787  
1788  
1789  
1790  
1791  
1792  
1793  
1794  
1795  
1796  
1797  
1798  
1799  
1800  
1801  
1802  
1803  
1804  
1805  
1806  
1807  
1808  
1809  
1810  
1811  
1812  
1813  
1814  
1815  
1816  
1817  
1818  
1819  
1820  
1821  
1822  
1823  
1824  
1825  
1826  
1827  
1828  
1829  
1830  
1831  
1832  
1833  
1834  
1835

(a) Reconstruction provided by the CWG method, in the energy-varying (e-v) and energy-conserving (e-c) cases.

(b) Geodesic distances.

Figure 13: Visual reconstruction (left) and geodesic distances of cell snapshots for the reference dataset (—), CWG (e-v) (—), and CWG (e-c) (—) (right). Snapshots are equally spaced in time, but cell differentiation progresses more rapidly at the beginning, as indicated by the larger early distances in the reference. The energy profile in CWG (e-v) more accurately captures this temporal dynamics.

#### G.4 SEA TEMPERATURE PREDICTION

The **NOAA OISST v2 High Resolution Dataset** (Huang et al., 2021) is a long-term Climate Data Record that integrates observations from multiple platforms (satellites, ships, buoys, and Argo floats) into a global gridded product. For this experiment, we use daily averages of sea surface temperature in the Gulf of Mexico between 1981 and 2024, represented as  $64 \times 64$  single-channel images. We cluster the measurements over five-year periods and select five representative months to define five time frames: January ( $t_0$ ), March ( $t_1$ ), May ( $t_2$ ), July ( $t_3$ ), and September ( $t_4$ ). Each month corresponds to a distribution of images, denoted as  $\{\rho_a$  (January),  $\rho_{m_1}$  (March),  $\rho_{m_2}$  (May),  $\rho_{m_3}$  (July),  $\rho_b$  (September)  $\}$ . A sample from one of these distributions is a heatmap of the Gulf’s temperature for a specific day in the corresponding month of the specified five-year period. The goal of this test is to evaluate our method’s ability to interpolate across missing time frames, generating realistic temperature maps for months without data. To this end, we partition the dataset into a training set  $\{t_0, t_2, t_4\}$  and a validation set  $\{t_1, t_3\}$ , and assess the quality of predictions on the held-out months. To encourage generalization beyond the training data, we introduce a potential function  $U$  that penalizes deviations from the learned data manifold. Building on the approach of [Song & Itti \(2025\)](#), where generative models are evaluated by measuring the distance between their outputs and a geometric manifold of real images learned by a VAE, we adopt a similar strategy. Specifically, we use a [state-of-the-art VAE architecture](#), with parameters listed in Table 22, to learn the manifold of the training images in our dataset. The potential function  $U(x^t)$ , for samples  $x^t \sim \rho^t$ , is then defined as the squared distance between a bridge sample  $x^t$  and its VAE-projected reconstruction  $\tilde{x}^t = \text{VAE}(x^t)$ :  $U(x^t) = \|x^t - \tilde{x}^t\|^2$ . Results for each five-year periods are presented below.

In these results, the scalar Hamiltonian function  $H^{t_k}$  in the pullback Jacobi metric,  $T_\theta^* \tilde{g}_J = (H^{t_k} - \mathcal{F}(\rho^t) - \mathcal{B}(\rho^t)) T_\theta^* g^{\mathcal{W}_2}$  is linearly varied from an initial value  $H^{t_0} = 1.36$  to a final value  $H^{t_K} = 1.0$ . This law is treated as a hyperparameter of the methodology, and all ten experiments reported in the following tables were conducted under this Hamiltonian behavior. In the geodesic computation, the solution path seeks to minimize  $T_\theta^* \tilde{g}_J$ . The functional  $\mathcal{F}(\rho^t)$  measures the distance with respect to the data manifold, i.e., it evaluates whether the samples from the bridge remain coherent with the images observed during training, and it is minimized along the trajectory. Since  $H^{t_k}$  is prescribed, the entropy term  $\mathcal{B}(\rho^t)$  is effectively maximized up to the limit set by  $H^{t_k}$ , as the metric cannot become negative. Hence, the assigned Hamiltonian energy  $H^{t_k}$  determines the admissible entropy level of the bridge. In this experiment, a decreasing trend of  $H^{t_K}$ , corresponding to an increase in physical energy (see Appendix E.4), was observed to be beneficial for modeling the warmer months, which appear to be more distant from the rest of the dataset in terms of Wasserstein distances compared to the colder months. A higher final energy  $H^{t_K}$ , and thus a higher final entropy  $\mathcal{B}$ , was observed to be beneficial for modeling the warmer months. This effect can be associated with the increased thermodynamic entropy of such cases, leading to more diverse samples and larger

variability across the data manifold. By prescribing a higher energy level, the model is encouraged to capture this diversification, exploring regions of the data manifold that in other contexts (Arvanitidis et al., 2018) are regarded as uncertain and are typically avoided.

1839

1840

Table 22: Architecture of the ConvVAE used in the Sea Temperature Prediction experiment. All Conv2D and ConvTranspose2D layers use kernel size 4, stride 2, and padding 1, followed by ReLU activations (except the last layer, which uses Sigmoid).

1841

1842

1843

Stage	Layer (channels)	Output size
Input	Single-channel image	$1 \times 64 \times 64$
Encoder	Conv2D ( $1 \rightarrow 32$ )	$32 \times 32 \times 32$
	Conv2D ( $32 \rightarrow 64$ )	$64 \times 16 \times 16$
	Conv2D ( $64 \rightarrow 128$ )	$128 \times 8 \times 8$
	Flatten	8192
Latent space	Linear $\rightarrow \mu$	5
	Linear $\rightarrow \log \sigma^2$	5
Decoder	Linear $\rightarrow$ reshape	$128 \times 8 \times 8$
	ConvT2D ( $128 \rightarrow 64$ )	$64 \times 16 \times 16$
	ConvT2D ( $64 \rightarrow 32$ )	$32 \times 32 \times 32$
	ConvT2D ( $32 \rightarrow 1$ ), Sigmoid	$1 \times 64 \times 64$

1844

1845

1846

1847

1848

1849

1850

1851

1852

1853

1854

1855

1856

1857

1858

1859

1860

1861

Table 23: FID scores at training and validation steps ( $\downarrow$ ), and training time (tt) ( $\downarrow$ ) in Sea (2020–2024).

Metric	CWG	GSBM	DSBM	SB-Flow
$\text{FID}(x^{t_0})$	$41.56 \pm 1.89$	—	—	—
$\text{FID}(x^{t_1})$	<b><math>121.47 \pm 5.61</math></b>	$160.68 \pm 4.54$	$242.26 \pm 9.94$	<b><math>176.76 \pm 4.41</math></b>
$\text{FID}(x^{t_2})$	$51.51 \pm 5.52$	$54.47 \pm 5.88$	$56.51 \pm 5.78$	<b><math>49.08 \pm 4.27</math></b>
$\text{FID}(x^{t_3})$	<b><math>159.53 \pm 7.38</math></b>	$185.54 \pm 7.11$	$235.83 \pm 10.44$	<b><math>189.53 \pm 7.42</math></b>
$\text{FID}(x^{t_4})$	$61.48 \pm 6.79$	$59.14 \pm 5.52$	$58.39 \pm 6.13$	<b><math>59.96 \pm 5.85</math></b>
tt (s)	<b><math>1030 \pm 50</math></b>	$73600 \pm 3200$	$19100 \pm 900$	<b><math>4950 \pm 320</math></b>

1862

1863

1864

1865

1866

1867

1868

1869

1870

1871

1872

1873

Table 24: FID scores at training and validation steps ( $\downarrow$ ) in Sea Temperature (2015–2019).

Metric	CWG	GSBM	DSBM	SB-Flow
$\text{FID}(x^{t_0})$	$42.96 \pm 2.07$	—	—	—
$\text{FID}(x^{t_1})$	<b><math>130.33 \pm 5.94</math></b>	$145.97 \pm 6.12$	$220.76 \pm 8.61$	<b><math>174.84 \pm 5.47</math></b>
$\text{FID}(x^{t_2})$	$58.72 \pm 5.46$	$62.13 \pm 5.69$	$65.47 \pm 6.16$	<b><math>60.88 \pm 5.11</math></b>
$\text{FID}(x^{t_3})$	$135.25 \pm 6.73$	$142.82 \pm 7.27$	$228.14 \pm 9.26$	<b><math>163.53 \pm 5.92</math></b>
$\text{FID}(x^{t_4})$	$63.02 \pm 6.14$	$59.73 \pm 5.86$	$61.29 \pm 6.05$	<b><math>59.04 \pm 5.76</math></b>

1880

1881

1882

1883

1884

Table 25: FID scores at training and validation steps ( $\downarrow$ ) in Sea Temperature (2010–2014).

Metric	CWG	GSBM	DSBM	SB-Flow
$\text{FID}(x^{t_0})$	$45.19 \pm 2.27$	—	—	—
$\text{FID}(x^{t_1})$	<b><math>132.47 \pm 6.58</math></b>	$168.93 \pm 6.24$	$255.36 \pm 10.19$	<b><math>171.03 \pm 6.38</math></b>
$\text{FID}(x^{t_2})$	$60.57 \pm 6.08$	$59.79 \pm 5.98$	$61.89 \pm 6.23$	<b><math>58.26 \pm 5.72</math></b>
$\text{FID}(x^{t_3})$	$140.84 \pm 6.01$	$144.60 \pm 6.83$	$235.08 \pm 9.07$	<b><math>179.81 \pm 7.74</math></b>
$\text{FID}(x^{t_4})$	$54.92 \pm 5.23$	$51.63 \pm 5.55$	$50.27 \pm 5.94$	<b><math>51.38 \pm 5.15</math></b>

Table 26: FID scores at training and validation steps (tt) ( $\downarrow$ ) Sea Temperature (2005–2009).

Metric	CWG	GSBM	DSBM	SB-Flow
$\text{FID}(x^{t_0})$	$45.19 \pm 2.27$	—	—	—
$\text{FID}(x^{t_1})$	<b><math>172.69 \pm 7.62</math></b>	$195.83 \pm 8.04$	$260.97 \pm 10.92$	<b><math>194.62 \pm 6.93</math></b>
$\text{FID}(x^{t_2})$	$56.08 \pm 5.52$	$59.57 \pm 5.87$	$62.03 \pm 6.29$	<b><math>60.87 \pm 5.44</math></b>
$\text{FID}(x^{t_3})$	<b><math>126.87 \pm 6.94</math></b>	$152.26 \pm 6.63$	$243.58 \pm 10.39$	<b><math>176.32 \pm 6.21</math></b>
$\text{FID}(x^{t_4})$	$66.43 \pm 6.58$	$63.17 \pm 6.34$	$64.86 \pm 6.63$	<b><math>64.37 \pm 5.59</math></b>

Table 27: FID scores at training and validation steps ( $\downarrow$ ) in Sea Temperature (2000–2004).

Metric	CWG	GSBM	DSBM	SB-Flow
$\text{FID}(x^{t_0})$	$45.19 \pm 2.27$	—	—	—
$\text{FID}(x^{t_1})$	<b><math>140.48 \pm 7.93</math></b>	$172.96 \pm 8.34$	$250.86 \pm 11.08$	<b><math>184.58 \pm 7.49</math></b>
$\text{FID}(x^{t_2})$	$56.73 \pm 5.36$	$56.97 \pm 5.79$	$59.64 \pm 6.10$	<b><math>59.53 \pm 5.77</math></b>
$\text{FID}(x^{t_3})$	<b><math>142.18 \pm 7.01</math></b>	$179.42 \pm 6.98$	$265.78 \pm 10.67$	<b><math>195.34 \pm 7.62</math></b>
$\text{FID}(x^{t_4})$	$68.29 \pm 6.97$	$65.91 \pm 6.68$	$64.73 \pm 6.91$	<b><math>64.23 \pm 5.42</math></b>

Table 28: FID scores at training and validation steps ( $\downarrow$ ) in Sea Temperature (1995–1999).

Metric	CWG	GSBM	DSBM	SB-Flow
$\text{FID}(x^{t_0})$	$45.19 \pm 2.27$	—	—	—
$\text{FID}(x^{t_1})$	<b><math>138.92 \pm 7.35</math></b>	$176.14 \pm 7.03$	$257.26 \pm 10.93$	<b><math>174.53 \pm 7.12</math></b>
$\text{FID}(x^{t_2})$	$60.19 \pm 6.01$	$63.48 \pm 6.12$	$66.07 \pm 6.57$	<b><math>62.95 \pm 6.06</math></b>
$\text{FID}(x^{t_3})$	<b><math>154.37 \pm 8.14</math></b>	$193.28 \pm 8.47$	$266.86 \pm 11.34$	<b><math>191.46 \pm 7.93</math></b>
$\text{FID}(x^{t_4})$	$69.57 \pm 7.02$	$66.83 \pm 6.69$	$66.18 \pm 6.94$	<b><math>65.62 \pm 7.50</math></b>

Table 29: FID scores at training and validation steps ( $\downarrow$ ) in Sea Temperature (1990–1994).

Metric	CWG	GSBM	DSBM	SB-Flow
$\text{FID}(x^{t_0})$	$45.19 \pm 2.27$	—	—	—
$\text{FID}(x^{t_1})$	<b><math>145.59 \pm 7.82</math></b>	$184.37 \pm 7.46$	$258.97 \pm 11.26$	<b><math>179.43 \pm 7.83</math></b>
$\text{FID}(x^{t_2})$	$72.38 \pm 6.42$	$78.69 \pm 6.71$	$82.13 \pm 7.04$	<b><math>75.75 \pm 6.27</math></b>
$\text{FID}(x^{t_3})$	<b><math>160.08 \pm 7.97</math></b>	$181.67 \pm 8.19$	$236.46 \pm 11.59$	<b><math>187.26 \pm 8.24</math></b>
$\text{FID}(x^{t_4})$	$73.42 \pm 7.19$	$70.68 \pm 6.87$	$69.83 \pm 7.09$	<b><math>69.47 \pm 6.98</math></b>

Table 30: FID scores at training and validation steps ( $\downarrow$ ) in Sea Temperature (1985–1989).

Metric	CWG	GSBM	DSBM	SB-Flow
$\text{FID}(x^{t_0})$	$45.19 \pm 2.27$	—	—	—
$\text{FID}(x^{t_1})$	<b><math>151.27 \pm 8.13</math></b>	$188.79 \pm 7.68$	$262.47 \pm 11.64$	<b><math>194.87 \pm 7.61</math></b>
$\text{FID}(x^{t_2})$	$66.89 \pm 6.31$	$69.19 \pm 6.47$	$71.23 \pm 7.08$	<b><math>67.83 \pm 6.20</math></b>
$\text{FID}(x^{t_3})$	$157.58 \pm 8.73$	$160.37 \pm 8.91$	$258.96 \pm 11.78$	<b><math>179.36 \pm 8.42</math></b>
$\text{FID}(x^{t_4})$	$66.02 \pm 7.41$	$62.87 \pm 7.16$	$61.59 \pm 7.34$	<b><math>63.37 \pm 7.19</math></b>

Table 31: FID scores at training and validation steps ( $\downarrow$ ) in Sea Temperature (1981–1983).

Metric	CWG	GSBM	DSBM	SB-Flow
$\text{FID}(x^{t_0})$	$45.19 \pm 2.27$	—	—	—
$\text{FID}(x^{t_1})$	<b><math>166.08 \pm 8.91</math></b>	$188.79 \pm 9.02$	$252.59 \pm 12.03$	<b><math>194.22 \pm 8.59</math></b>
$\text{FID}(x^{t_2})$	$69.29 \pm 6.65$	$71.68 \pm 6.74$	$73.87 \pm 7.26$	<b><math>70.75 \pm 7.36</math></b>
$\text{FID}(x^{t_3})$	<b><math>168.73 \pm 8.46</math></b>	$185.57 \pm 8.08$	$260.47 \pm 11.92$	<b><math>198.63 \pm 8.45</math></b>
$\text{FID}(x^{t_4})$	$77.83 \pm 7.68$	$74.97 \pm 7.34$	$74.29 \pm 7.57$	<b><math>74.87 \pm 7.09</math></b>

1944  
1945

## G.5 ROBOTIC TASK RECONSTRUCTION

1946 **BridgeData V2** (Walke et al., 2023) is a large and diverse dataset of robotic manipulation behaviors,  
 1947 designed to advance research in scalable robot learning. In this experiment, our goal is to reconstruct  
 1948 the full video of a robot performing manipulation tasks while training the network only on images  
 1949 from the beginning and end of the sequence, interpreted as samples from the endpoint distributions  
 1950  $\rho_a$  and  $\rho_b$ . No intermediate marginals  $\rho_m$  are used. Following the Sea Temperature Prediction  
 1951 experiment, we introduce a potential function  $U$  that penalizes deviations from the learned data  
 1952 manifold. This manifold is learned from the initial and final frames of the videos, and the penalty  
 1953 encourages plausible intermediate frames consistent with these distributions. We employ a **state-**  
 1954 **of-the-art VAE architecture**, with parameters listed in Table 33, to model the image manifold. The  
 1955 potential function  $U(x^t)$  for samples  $x^t \sim \rho^t$  is defined as the squared distance between a bridge  
 1956 sample  $x^t$  and its VAE reconstruction  $\tilde{x}^t = \text{VAE}(x^t)$ :  $U(x^t) = \|x^t - \tilde{x}^t\|^2$ . Fig. 7 presents  
 1957 snapshots of the reconstructions produced by our CWG method compared to the baselines. In this  
 1958 experiment, the Hamiltonian function  $H^{t_k}$  is held constant, as varying it yields no apparent benefit.  
 1959

In the guided generation setting, we define the feature function  $f$  as

$$f(x_b) = \text{ReLU}(w_c c_b^{(x)} - b_c), \quad (71)$$

1960 where  $c_b^{(x)}$  denotes the  $(x)$ -coordinate pixel position of the centroid corresponding to the target  
 1961 location of the item placed by the robot, extracted from the image  $x_b$  sampled from the reference  
 1962 marginal  $\rho_b$ . We impose a penalty  $f$  on this image (parameters of this function are available in Table  
 1963 34) so that samples corresponding to placements in undesired locations are discouraged, while those  
 1964 leading to desirable targets are favored.  
 1965

1966 The centroid extraction is performed by applying morphological opening and closing operations  
 1967 from the **OpenCV library** to remove noise and refine object boundaries, followed by color-based  
 1968 masking to isolate the object of interest.  
 1969



1970  
 1971  
 1972  
 1973  
 1974  
 1975  
 1976  
 1977  
 1978  
 1979  
 1980  
 1981  
 1982  
 1983  
 1984  
 1985  
 1986  
 1987  
 1988  
 1989  
 1990  
 1991  
 1992  
 1993  
 1994  
 1995  
 1996  
 1997

Figure 14: Reconstructions from CWG (top), GSBM (second), DSBM (third), and SB-Flow (bottom) in the Robot Task Reconstruction experiment. Red row shows the reference.

Table 32: FID score (↓) and training time (tt) (↓) in Robotic Task Reconstruction.

Metric	CWG (ours)	GSBM	DSBM	SB-Flow
FID	<b>18.83<math>\pm</math>0.66</b>	40.23 $\pm$ 1.95	149.78 $\pm$ 0.81	73.46 $\pm$ 0.52
tt (s)	<b>1090<math>\pm</math>40</b>	91100 $\pm$ 8000	27400 $\pm$ 2500	4900 $\pm$ 300

1998 Table 33: Architecture of the ConvVAE used in the Robot Task Reconstruction experiment. All Conv2D and  
 1999 ConvTranspose2D layers use kernel size 4, stride 2, and padding 1, followed by ReLU activations (except the  
 2000 last layer, which uses Sigmoid).

Stage	Layer (channels)	Output size
Input	Single-channel image	$3 \times 64 \times 64$
Encoder	Conv2D (3→32)	$32 \times 32 \times 32$
	Conv2D (32→64)	$64 \times 16 \times 16$
	Conv2D (64→128)	$128 \times 8 \times 8$
	Flatten	8192
Latent space	Linear $\rightarrow \mu$	2
	Linear $\rightarrow \log \sigma^2$	2
Decoder	Linear $\rightarrow$ reshape	$128 \times 8 \times 8$
	ConvT2D (128→64)	$64 \times 16 \times 16$
	ConvT2D (64→32)	$32 \times 32 \times 32$
	ConvT2D (32→3), Sigmoid	$3 \times 64 \times 64$

2008 Table 34: Parameters for the feature function  $f$  (71), used for guided generation in the Robot Task Reconstruction  
 2009 experiment.

Parameter	Value
Weight $w_c$	1
Bias $b_c$	30

## G.6 FFHQ TRANSFER

2018 The [Flickr-Faces-HQ \(FFHQ\)](#) dataset is a high-resolution ( $1024 \times 1024$ ) collection of human faces  
 2019 that exhibits a remarkably wide range of visual variations (Karras et al., 2019). We closely follow  
 2020 the experimental protocol adopted by Gushchin et al. (2024), using their publicly available datasets  
 2021 from <https://github.com/ngushchin/LightSB>. We split the dataset into training (first 60K  
 2022 images) and testing (last 10K images) subsets. Each subset is further partitioned into age groups,  
 2023 specifically *adults* and *children*, corresponding to the source distribution  $\rho_a$  and the target distribution  
 2024  $\rho_b$ , respectively. The objective of this experiment is to generate a realistic child image from a  
 2025 given adult image, representing the same individual at a younger age. For each image, we employ  
 2026 the pre-trained ALAE encoder (Pidhorskyi et al., 2020) to extract a 512-dimensional latent vector.  
 2027 The SB solvers are trained directly on these latent representations. During inference, we encode a  
 2028 given image into the latent space, apply the learned bridge using the SB models, and then decode  
 2029 its outcome to obtain the transformed image. Because the probabilistic bridge is already defined  
 2030 within the ALAE latent space, which captures the FFHQ data manifold, the CWG method does  
 2031 not rely on the potential function  $U$  to quantify deviations from image manifold. In contrast, the  
 2032 baseline GSBM approach retains a potential energy term  $U$  that penalizes excessive entropy and  
 2033 overly concentrated probability densities  $\rho^t$ , distinguishing it from DSBM, which does not include  
 2034 any potential term.

2035 Qualitative examples of the images generated by our CWG method and the baseline models are  
 2036 presented in Figures 15–17. A more compact, side-by-side qualitative comparison is provided in  
 2037 Figure 9, while quantitative results are reported in Table 7. The Feasibility metric indicates whether  
 2038 the computed bridges respect the imposed constraints, that is, whether they match the target marginal  
 2039 distribution. It is computed by evaluating the FID score between the terminal distribution generated  
 2040 by the models,  $\rho_{\theta}^{t_K}$ , and the target distribution,  $\rho_b$ , as defined in equation (72b). The starting distribution  
 2041  $\rho_a$  corresponds to the source domain from which the input samples are drawn, and therefore  
 2042 it is naturally satisfied and does not require matching. The Optimality metric measures the geodesic  
 2043 length of the Schrödinger bridges computed by the models, in order to assess the transport cost that  
 2044 characterizes each solution. It is computed by sampling from the probabilistic path  $z^{t_k} \sim \rho^{t_k}$  in the  
 2045 latent space. The pre-trained ALAE decoder is then used to map these latent vectors to images  $x^{t_k}$ ,

2052 which are subsequently encoded into feature representations  $x_{\text{fea}}^{t_k}$  using a pre-trained *Inception-v3*  
 2053 network (the same used to compute the FID score). The curve length is then obtained from equation  
 2054 (72a).

$$\text{Optimality} = \sum_{k=1}^K d_{\mathcal{W}_2}(x_{\text{fea}}^{t_k}, x_{\text{fea}}^{t_{k-1}}); \quad (72\text{a})$$

$$\text{Feasibility} = \text{FID}(\rho_{\theta}^{t_K}, \rho_b). \quad (72\text{b})$$

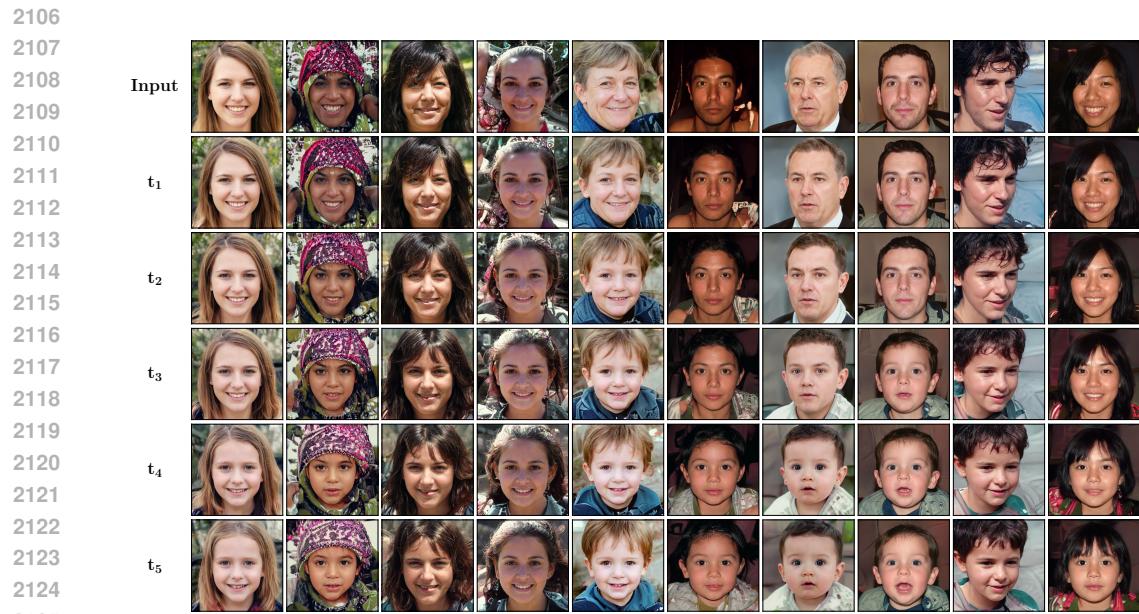
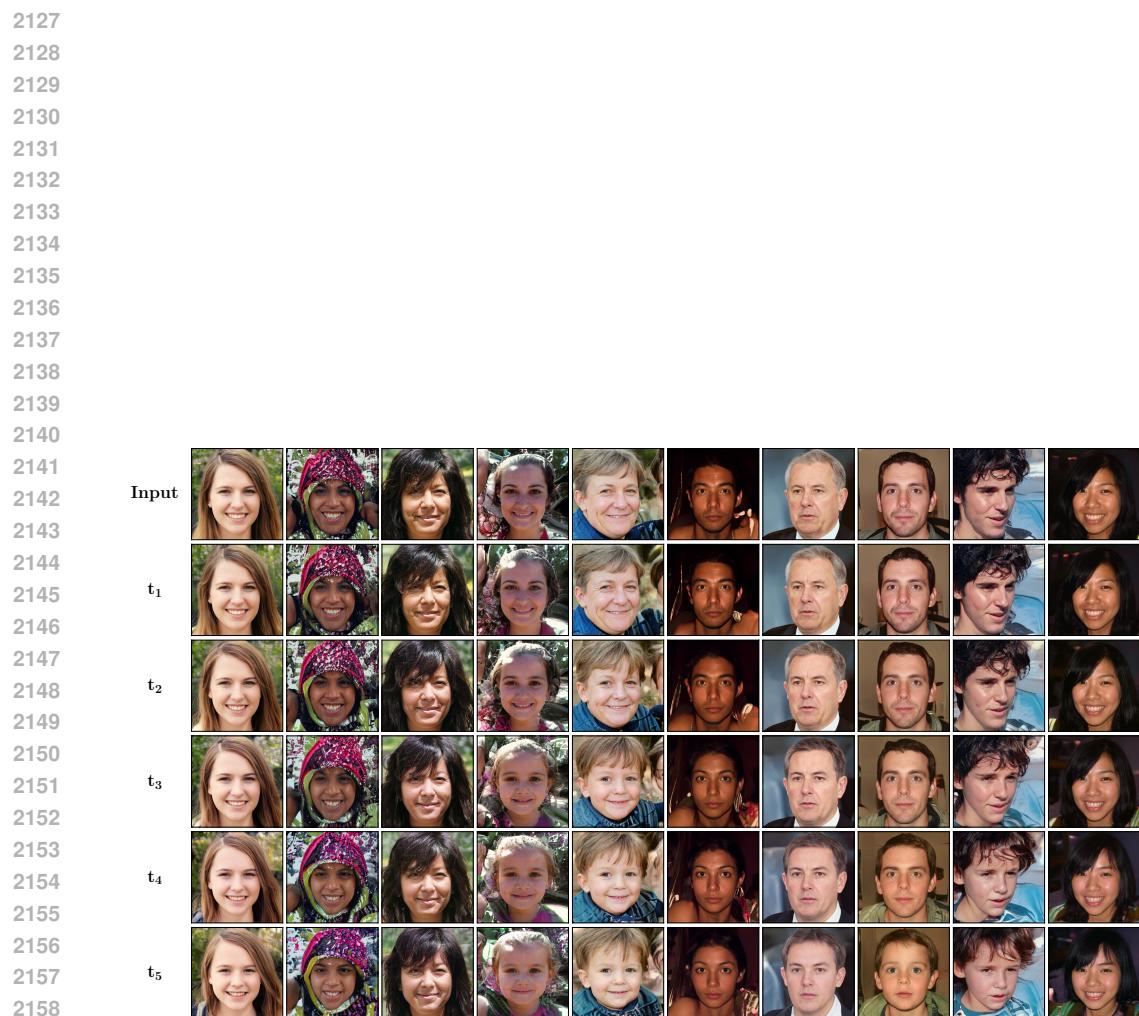
2060 Hence, the length of the geodesic computed from the bridge in the latent space is evaluated according  
 2061 to the Wasserstein metric, pulled back from the feature space to the image space, and from the im-  
 2062 age space to the latent space. Although the baseline models achieve better Optimality scores (lower  
 2063 values), this improvement comes at the cost of failing to satisfy the marginal constraints, as reflected  
 2064 by their poorer Feasibility scores. Since constraint satisfaction is of critical importance, our CWG  
 2065 method exhibits superior overall behavior. This conclusion is further supported by Figure 19, where  
 2066 three facial age estimation networks,  $h_1$  (prithivMLmods/facial-age-detection),  $h_2$  (nateraw/vit-age-  
 2067 classifier), and  $h_3$  (abhilash88/age-gender-prediction), were used to estimate the ages of the gen-  
 2068 erated faces at each generation step. The different curves in the plot represent the predicted ages for  
 2069 the images shown in Figures 15–17. While the CWG model exhibits a more decisive transformation  
 2070 across time steps, the baseline curves remain comparatively flat, often producing outputs that fail to  
 2071 convincingly resemble children images. The age prediction models ( $h_1, h_2, h_3$ ) provide stochastic  
 2072 metrics that quantify their predictive uncertainty, which we use to estimate their variance  $\sigma^2$  under  
 2073 the assumption of a uniform error distribution. The average predicted age and the corresponding  
 2074 variance for the plots represented in Figure 19 are computed as,

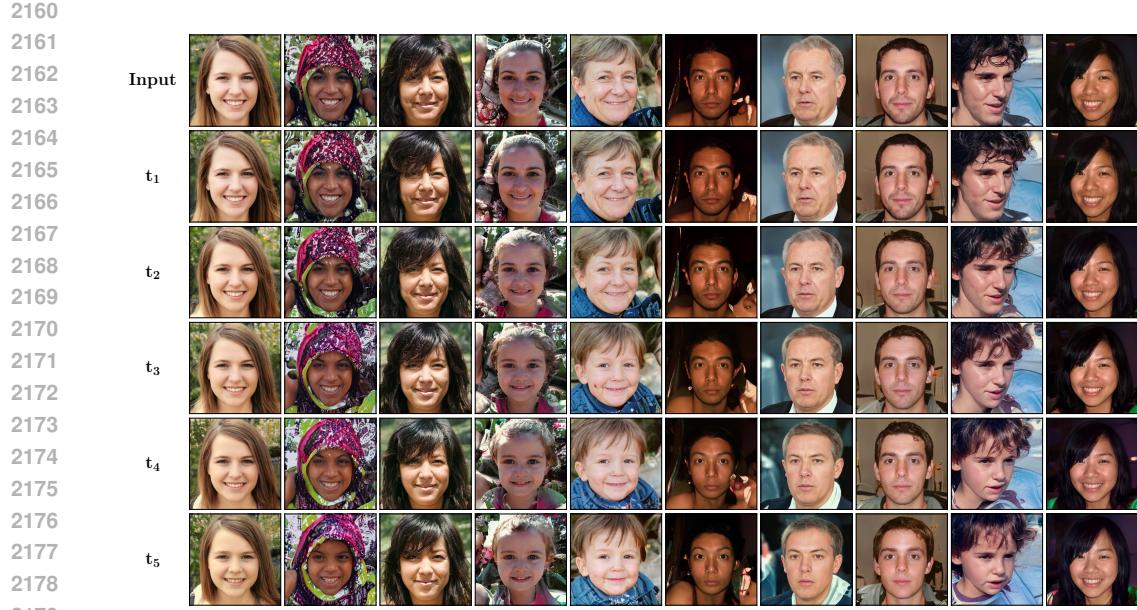
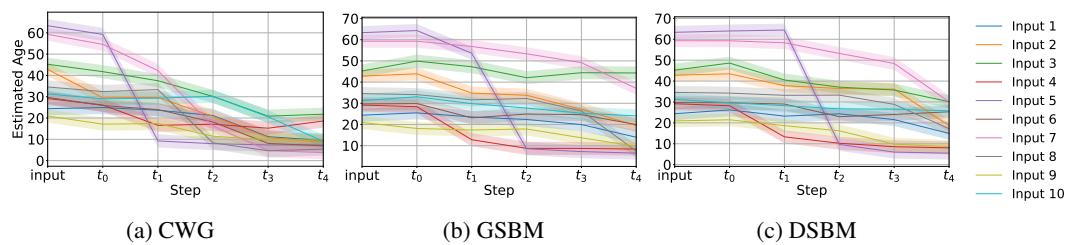
$$\mu_{\text{age}} = \frac{\frac{h_1(x^{t_k})}{\sigma_1^2} + \frac{h_2(x^{t_k})}{\sigma_2^2} + \frac{h_3(x^{t_k})}{\sigma_3^2}}{\frac{1}{\sigma_1^2} + \frac{1}{\sigma_2^2} + \frac{1}{\sigma_3^2}}, \quad \sigma_{\text{age}}^2 = \frac{1}{\frac{1}{\sigma_1^2} + \frac{1}{\sigma_2^2} + \frac{1}{\sigma_3^2}}. \quad (73)$$

2075 As shown in Figure 10, the energy of the bridge fitted in the latent space affects the age trends of  
 2076 the predictions. Decreasing the energy encourages smaller distances between the discretized nodes  
 2077 of the bridge at its beginning, resulting in images that are closer in age at early steps and gradually  
 2078 diverge toward later steps. Conversely, increasing the energy produces the opposite behavior. Im-  
 2079 portantly, in all cases, the Feasibility is preserved, as the target distribution and mean remain well  
 2080 below the 18-year-old threshold. This contrasts with the baseline models, which in average does not  
 2081 respect this constraint.

2084  
 2085  
 2086  
 2087  
 2088  
 2089  
 2090  
 2091  
 2092  
 2093  
 2094  
 2095  
 2096  
 2097  
 2098  
 2099  
 2100  
 2101  
 2102  
 2103  
 2104  
 2105  
 Table 35: FID scores (↓) measuring the distance to the initial adult distribution  $\rho_a$  and the target child distribu-  
 tion  $\rho_b$  for samples generated at different time steps using the CWG method in the FFHQ transfer.

Metric	$\rho_a$	$\rho_b$
$\text{FID}(x^{t_1})$	<b>3.143<math>\pm</math>0.447</b>	$33.953\pm0.936$
$\text{FID}(x^{t_2})$	$7.596\pm0.743$	$27.413\pm0.879$
$\text{FID}(x^{t_3})$	$25.735\pm0.893$	$7.164\pm0.806$
$\text{FID}(x^{t_4})$	$30.636\pm0.843$	$4.375\pm0.604$
$\text{FID}(x^{t_5})$	$36.001\pm0.856$	<b>4.3316<math>\pm</math>0.526</b>

Figure 15: [Adult  \$\rightarrow\$  Child image generation in the FFHQ transfer using the CWG method.](#)Figure 16: [Adult  \$\rightarrow\$  Child image generation FFHQ transfer using the GSBM method.](#)

Figure 17: [Adult  \$\rightarrow\$  Child image generation FFHQ transfer using the DSBM method.](#)Figure 18: [Adult  \$\rightarrow\$  Child image generation FFHQ transfer using the SB-Flow method.](#)Figure 19: [Age prediction for the samples generated in the FFHQ transfer experiment.](#)

## 2214 2215 2216 2217 2218 2219 2220 2221 2222 2223 2224 2225 2226 2227 2228 2229 2230 2231 2232 2233 2234 2235 2236 2237 2238 2239 2240 2241 2242 2243 2244 2245 2246 2247 2248 2249 2250 2251 2252 2253 2254 2255 2256 2257 2258 2259 2260 2261 2262 2263 2264 2265 2266 2267 G.7 MNIST-TO-EMNIST

The MNIST and EMNIST datasets are low-resolution ( $28 \times 28$ ) handwritten character datasets that serve as long-standing benchmarks for generative modeling, representation learning, and probabilistic transport (Cohen et al., 2017). In this work, we adopt MNIST as the source distribution  $\rho_a$  and EMNIST (letters) as the target  $\rho_b$  to study unpaired image-to-image translation. We restrict the analysis to ten digits and ten letters. The two datasets are visually related but statistically distinct, making the MNIST-to-EMNIST transport problem sufficiently non-trivial while still permitting clear qualitative and quantitative evaluation. Our goal is to produce terminal samples that match the EMNIST distribution while ensuring that the Schrödinger Bridge generates meaningful intermediate states that remain on the data manifold rather than collapsing into noise or artifacts.

For this purpose, we introduce a potential function  $U$  that penalizes deviations from the learned data manifold. Building on the approach of Song & Itti (2025), where generative models are evaluated by measuring the distance between their outputs and a geometric manifold of real images learned by a VAE, we adopt a similar strategy. Specifically, we use a state-of-the-art VAE architecture, with parameters listed in Table 37, to learn a manifold composed of images from both marginal distributions  $\rho_a$  and  $\rho_b$ . The potential function  $U(x^t)$ , for samples  $x^t \sim \rho^t$ , is then defined as the squared distance between a bridge sample  $x^t$  and its VAE-projected reconstruction  $\tilde{x}^t = \text{VAE}(x^t)$ :  $U(x^t) = \|x^t - \tilde{x}^t\|^2$ . No intermediate samples  $\rho_m$  are used in this experiment.

The results are shown visually in Figure 20 and quantitatively in Table 38. As observed in the other experiments, DSBM fails to regularize the stochastic path between the two distributions, and SB-Flow searches for an interpolation that first removes features of the source digit not present in the target letter, rather than producing intermediate characters that remain close to the data manifold at every time step. GSBM produces only a deterministic trajectory along the data manifold, which it then uses as the mean of a transient Gaussian distribution. However, because the stochastic component is not properly constrained, these Gaussian samples can drift off the manifold, leading to unrealistic intermediate states. This drift appears as increased noise in the midpoint samples and results in higher FID scores. In contrast, the CWG method produces intermediate characters that remain closest to the source and target distributions at every discretized time step. Notably, our method also demonstrates strong computational efficiency, achieving training times significantly lower than those of the baselines.

To evaluate the guided version of the CWG method, we propose a Gaussian deblurring test, a well-established inverse problem commonly used in image restoration (Feng et al., 2025; Wang et al., 2025). In this setup, a reference image from the EMNIST dataset is processed with a Gaussian kernel and provided as the target for restoration. In this test, the guided loss function,  $\|f(x^{t_K}) - y\|$ , uses the blurry image as the target  $y$ , and the feature function  $f$  is defined as,

$$f(x^{t_K}) = \mathbf{K}_{\text{gaussian}}(\sigma, r) \odot x^{t_K}, \quad (74)$$

where  $\mathbf{K}_{\text{gaussian}}(\sigma, r)$  is a Gaussian kernel operator of variance  $\sigma$  and size  $r$ , applied to the sample  $x^{t_K}$  generated at the last time step of the bridge. The specific kernel parameters are detailed in Table 36. This loss measures the similarity between the blurred version of the generated image and the blurred reference, ensuring compatibility between the two. The restoration results are shown visually in Figure 21 and quantitatively in Table 39.

Table 36: Parameters for the feature function  $f$  (74), used for guided generation in the Mnist-to-Emnist experiment.

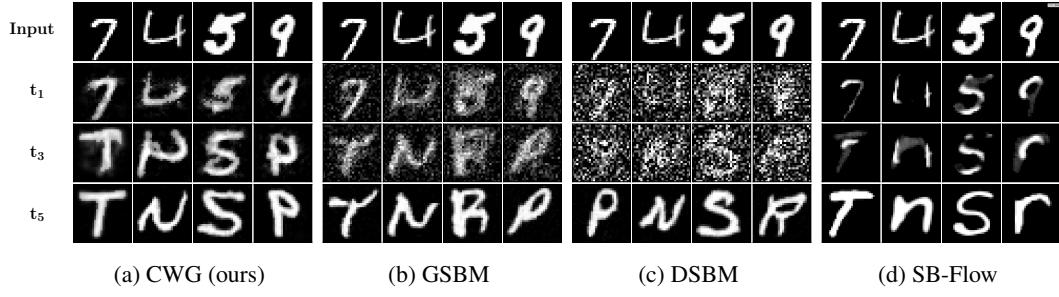
Parameter	Value
Kernel variance $\sigma$	3
Kernel size $r$	9

2268 Table 37: Architecture of the ConvVAE used in the MNIST-to-EMNIST experiment. All Conv2D and Con-  
 2269 vTranspose2D layers use ReLU activations unless otherwise specified.

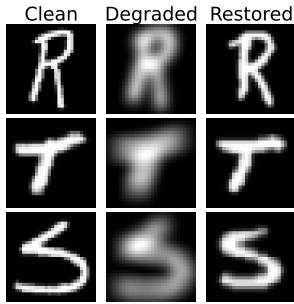
Stage	Layer (channels)	Output size
Input	Single-channel image	$1 \times 28 \times 28$
Encoder	Conv2D ( $1 \rightarrow 32$ , k4, s2, p1)	$32 \times 14 \times 14$
	Conv2D ( $32 \rightarrow 64$ , k4, s2, p1)	$64 \times 7 \times 7$
	Conv2D ( $64 \rightarrow 128$ , k3, s1, p1)	$128 \times 7 \times 7$
	Flatten	$128 \cdot 7 \cdot 7 = 6272$
Latent space	Linear $\rightarrow \mu$	2
	Linear $\rightarrow \log \sigma^2$	2
Decoder	Linear $\rightarrow$ reshape	$128 \times 7 \times 7$
	ConvT2D ( $128 \rightarrow 64$ , k4, s2, p1)	$64 \times 14 \times 14$
	ConvT2D ( $64 \rightarrow 32$ , k4, s2, p1)	$32 \times 28 \times 28$
	ConvT2D ( $32 \rightarrow 1$ , k3, s1, p1)	$1 \times 28 \times 28$

2285 Table 38: FID scores computed w.r.t. the EMNIST distribution ( $\downarrow$ ) and training time (tt) in the MNIST-to-  
 2286 EMNIST experiment.

Metric	CWG	GSBM	DSBM	SB-Flow
$\text{FID}(x^{t_0})$	$77.25 \pm 5.37$	$86.80 \pm 6.11$	$102.94 \pm 7.04$	$79.35 \pm 4.86$
$\text{FID}(x^{t_1})$	$83.33 \pm 3.74$	$114.53 \pm 8.48$	$228.36 \pm 13.35$	$106.28 \pm 5.46$
$\text{FID}(x^{t_2})$	$95.29 \pm 6.39$	$128.69 \pm 7.80$	$195.38 \pm 12.24$	$121.80 \pm 6.32$
$\text{FID}(x^{t_3})$	$56.42 \pm 5.63$	$89.75 \pm 6.25$	$163.69 \pm 11.34$	$107.29 \pm 4.74$
$\text{FID}(x^{t_4})$	$29.73 \pm 3.54$	$58.28 \pm 7.17$	$108.53 \pm 11.99$	$34.37 \pm 2.17$
$\text{FID}(x^{t_5})$	$11.42 \pm 0.32$	$11.75 \pm 0.36$	$11.69 \pm 0.31$	$11.29 \pm 0.27$
tt (s)	$570 \pm 30$	$30100 \pm 1000$	$8750 \pm 300$	$2200 \pm 60$



2307 Figure 20: Snapshots from the bridges computed in the MNIST-to-EMNIST experiment.



Metric	Score
FID	$12.72 \pm 0.36$
LPIPS	$0.28 \pm 0.02$
PSNR	$12.48 \pm 1.11$
SSIM	$0.37 \pm 0.04$
tt (s)	$320 \pm 20$

2309 Table 39: Quantitative metrics for the Gaussian deblurring experiment: FID score ( $\downarrow$ ) between the clean and reconstructed samples, LPIPS ( $\downarrow$ ), PSNR ( $\uparrow$ ), SSIM ( $\uparrow$ ), training time (tt) ( $\downarrow$ ).

2318 Figure 21: Gaussian deblurring test for the Mnist-to-Emnist  
 2319 experiment. Results for the guided CWG method.