
HOTA: Hamiltonian framework for Optimal Transport Advection

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 Optimal transport (OT) has become a natural framework for guiding the probability
2 flows. Yet, the majority of recent generative models assume trivial geometry
3 (e.g., Euclidean) and rely on strong density-estimation assumptions, yielding tra-
4 jectories that do not respect the true principles of optimality in the underlying
5 manifold. We present Hamiltonian Optimal Transport Advection (HOTA), a Hamil-
6 ton–Jacobi–Bellman based method that tackles the dual dynamical OT problem
7 explicitly through Kantorovich potentials, enabling efficient and scalable trajec-
8 tory optimization. Our approach effectively evades the need for explicit density
9 modeling, performing even when the cost functionals are non-smooth. Empirically,
10 HOTA outperforms all baselines in standard benchmarks, as well as in custom
11 datasets with non-differentiable costs, both in terms of feasibility and optimality.

12 1 Introduction

13 *Static (Monge–Kantorovich)* optimal transport was originally considered as the main framework for
14 comparing and finding a cost-minimizing coupling between distributions [Villani et al., 2008], while
15 optimality was mainly measured through the boundary marginals. Development of efficient and
16 scalable OT solvers [Cuturi, 2013, Peyré et al., 2019] popularized OT across different areas, such as
17 generative modeling [Makkuva et al., 2020, Korotin et al., 2022, Buzun et al., 2024], computational
18 biology [Bunne et al., 2022], graphics [Bonneel and Digne, 2023], high-energy physics [Nathan
19 T. Suri, 2024], and reinforcement learning [Klink et al., 2022, Asadulaev et al., 2024, Bobrin et al.,
20 2024, Rupf et al., 2025]. However, one crucial limitation of static formulation is its inability to
21 produce non-straight paths, which completely ignores the underlying geometry of the manifold of the
22 data. In classical OT, the underlying geometric structure is solely determined by the choice of cost
23 function (e.g., Euclidean distance), inherently limiting the capacity for fine-grained control over the
24 trajectories. We refer to [Montesuma et al., 2024, Pereira and Amini, 2025] for recent overview of
25 practical applications of OT and to Villani et al. [2008], Santambrogio [2015], Peyré et al. [2019] for
26 a formal treatment.

27 On the other hand, the *dynamical* optimal transport paradigm, developed by Benamou and Brenier
28 [2000], recasts static OT as a continuous-time variational problem on the space of probability paths,
29 effectively incorporating time variable and enabling more nuanced control over optimal trajectories
30 (e.g., through velocity, acceleration, length, or energy over the paths). Importantly, such formulation
31 enables one to directly operate on manifolds of non-trivial geometry, whenever the underlying space
32 contains curvature, obstacles, or is defined through potentials. This formulation is closely connected
33 to stochastic optimal control (SOC), where trajectories are stochastic yet must still maintain optimality,
34 a problem class known as the generalized Schrödinger bridge (GSB) Liu et al. [2024], Bartosh et al.
35 [2024].

36 A common strategy for GSB involves solving the dual formulation via Hamilton–Jacobi–Bellman
37 (HJB) equations, which provide a flexible and a theoretically grounded framework for deriving
38 optimal trajectories (Liu et al. [2022], Neklyudov et al. [2024]). These methods parameterize the

39 cost through a Lagrangian, enforcing optimality via the preservation of kinetic energy or using other
 40 path-based penalties. While HJB-based approaches yield theoretically sound solutions, they suffer
 41 from critical drawbacks: (1) unstable optimization dynamics, leading to high-variance gradients
 42 and poor sample efficiency in high dimensions, and (2) the absence of a strict terminal distribution
 43 matching criterion, resulting in inexact couplings. Additionally, they typically require differentiable
 44 Lagrangians, restricting applicability to smooth costs only.

45 In the current work, we study the Generalized Schrodinger Bridge problem between two mea-
 46 sures, where the underlying geometry is defined through potentials. We propose a new HJB-based
 47 framework that explicitly solves GSB task, resolves the learning stability problems of the previous
 48 approaches, and has theoretical guarantees. We conduct extensive empirical evaluations on existing
 49 low-dimensional physically-inspired benchmarks, as well as in the high-dimensional generative
 50 setting. In short, our contributions are as follows:

- 51 • Hamiltonian dual reformulation of dynamical OT that binds Kantorovich potentials with an
 52 HJB value function, yielding a density-free objective and providing the performance gain
 53 compared to existing works;
- 54 • Proposed approach is robust to complex geometries and works even with non-smooth cost
 55 functions as the proposed objective explicitly incorporates the potential term;
- 56 • HOTA attains state-of-the-art empirical results in a diverse set of tasks, demonstrating both
 57 better feasibility (exact marginal matching) and optimality (cost along trajectories) compared
 58 to current dynamic OT solvers.

59 2 Related work

60 **Diffusion Models and Matching Algorithms.** Diffusion models have emerged as powerful tools for
 61 generative modeling by prescribing the time evolution of marginal distributions. Matching algorithms,
 62 such as Action Matching (Neklyudov et al. [2024]) and Flow Matching (Lipman et al. [2023]), learn
 63 stochastic differential equations (SDEs) that align with prescribed probability paths [Blessing et al.,
 64 2025]. These methods typically assume explicit or implicit intermediate densities of the flow, whereas
 65 our approach (HOTA) optimizes a complete stochastic path from source to target distributions.

66 **Generalized Schrödinger Bridge.** The GSB problem extends SB by introducing state costs that
 67 penalize or reward specific trajectories (Chen et al., 2015). Prior methods for solving GSB, such as
 68 DeepGSB (Liu et al. [2022]), often relax feasibility constraints or rely on Sinkhorn-based approxima-
 69 tions, which can lead to instability or suboptimal solutions.

70 A recent approach GSBM [Liu et al., 2024] follows an alternating optimization scheme: in the first
 71 stage, it learns the drift field v_t while keeping the marginal distributions $\rho_t(x_t)$ fixed, using a Flow
 72 Matching-style objective. In the second stage, it updates the marginals conditioned on the boundary-
 73 coupled distribution $\rho_t(x_t | x_0, x_1)$, which is defined via the previously learned drift. While GSBM
 74 demonstrates strong empirical performance, it imposes two critical limitations: 1) it requires the state
 75 cost function $U(x_t)$ to be differentiable everywhere, and 2) it assumes that the conditional marginals
 76 $\rho(x_t | x_0, x_1)$ are Gaussian. The first constraint restricts the method’s applicability to domains with
 77 smooth geometries, sometimes mitigated via interpolation [Kapusniak et al., 2024], while the second
 78 can lead to suboptimal solutions, unless U_t function is not quadratic.

79 **Stochastic Optimal Control.** The connection between GSB and stochastic optimal control (SOC) has
 80 been explored in prior works (Theodorou et al. [2010]; Levine [2018]). SOC formulations often relax
 81 hard distributional constraints into soft terminal costs, which can introduce bias or require adversarial
 82 training (Liu et al. [2022]). Recently introduced Adjoint Matching approach [Domingo-Enrich et al.,
 83 2024a] and Stochastic Optimal Control matching (SOCM) [Domingo-Enrich et al., 2024b] address
 84 several existing limitations, but still produce highly unstable variance estimations. Our method
 85 provides a natural way to preserve the feasibility via Kantorovich potential sum.

86 3 Preliminaries

87 Consider stochastic process with controlled drift and diffusion:

$$dx_t = v(t, x_t) dt + \sigma(t, x_t) dW_t \quad (1)$$

88 where $v : [0, 1] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ is the drift (control), $\sigma : [0, 1] \times \mathbb{R}^d \rightarrow \mathbb{R}$ is the diffusion coefficient, W_t
 89 is d -dimensional Brownian motion. We solve the OT minimization task with marginal distributions $(\alpha,$

90 β) and dynamic cost functions $c(x, \mu)$ and stochastic transport mapping $\mu : \mathbb{R}^d \rightarrow \mathcal{P}(\mathbb{R}^d)$ presented
 91 in paper [Korotin et al. \[2022\]](#)

$$c(x, \mu) = \inf_{v(t, x) : x_0 = x, x_1 \sim \mu} \int_0^1 \mathbb{E} \mathcal{L}(t, x_t, v_t) dt, \quad \mathcal{L}(t, x_t, v_t) = \frac{\|v_t\|^2}{2} + U(x_t). \quad (2)$$

92 This problem is also known as generalized Schrödinger bridge (GSB). It is an extension of the classical
 93 Schrödinger Bridge (SB) problem, which is a distribution-matching task seeking a diffusion model
 94 that transports an initial distribution α to a target distribution β . While the standard SB minimizes the
 95 kinetic energy (L^2 cost in OT), the GSB introduces additional flexibility by incorporating a state cost
 96 $U(x_t)$, allowing for more general optimality conditions beyond just kinetic energy minimization. The
 97 standard SB's reliance on kinetic energy (Euclidean cost) may not be ideal for all applications (e.g.,
 98 image spaces, where distance may not be meaningful). Many scientific domains (population modeling,
 99 robotics, molecular dynamics) require richer optimality conditions, which GSB accommodates via
 100 $U(x_t)$. The potential term usually characterizes the geometry of the space. But in addition, we can
 101 also include some physical properties of the flow, e.g., entropic penalty or “mean-field” interaction
 102 [\[Liu et al., 2022\]](#). Thus, the optimal trajectories are curved to avoid regions with high values of
 103 $U(x_t)$.

104 Neural networks can effectively solve high-dimensional Optimal Transport (OT) problems by learning
 105 the Kantorovich potentials, which maximizes the dual objective [\(Korotin et al. \[2022\], Buzun et al.
 106 \[2024\]\)](#). It is shown in [Villani et al. \[2008\]](#) (Theorem 5.10) that OT task is equivalent to the
 107 maximization of the Kantorovich potentials sum:

$$\sup_{g \in L_1(\beta)} \left[\mathbb{E}_\alpha[g^c(x)] + \mathbb{E}_\beta[g(y)] \right], \quad (3)$$

108 where g^c denotes c -conjugate transform of the potential g :

$$g^c(x) = \inf_{\mu(x) : \mathbb{R}^d \rightarrow \mathcal{P}(\mathbb{R}^d)} \mathbb{E}_{y \sim \mu(x)} [c(x, \mu) - g(y)]. \quad (4)$$

109 Here $\mu(x)$ is the stochastic transport mapping, and in our notation it is the final distribution of
 110 the stochastic process x_1 under condition that $x_0 = x$. The marginality requirement of the final
 111 distribution of x_1 (which must correspond to β) is ensured by the potential difference $\mathbb{E}_\beta g(y)$ and
 112 $\mathbb{E}_\alpha \mathbb{E}_{y \sim \mu(x)} g(y)$, which tends to infinity otherwise.

113 But unlike classical OT, we need to minimize the cost throughout the trajectory $x_t, t \in [0, 1]$ with the
 114 following objective

$$g^c(x) = \inf_{v(x, t)} \mathbb{E} \left[\int_0^1 \left(\frac{\|v(t, x_t)\|^2}{2} + U(x_t) \right) dt - g(x_1) \mid x_0 = x \right]. \quad (5)$$

115 In the last expression, we have united infimums by $\mu(x)$ and control $v(t, x)$ and as a sequence have
 116 removed the right side condition $x_1 \sim \mu(x)$. Based on dynamic programming approach, define the
 117 value function. For any $0 \leq t \leq 1$, the value function satisfies:

$$s(t, x) = \inf_{x_t} \mathbb{E} \left[\int_t^1 \left(\frac{\|v(t, x_t)\|^2}{2} + U(x_t) \right) dt - g(x_1) \mid x_t = x \right], \quad (6)$$

118 such that our objective equals $s(0, x)$ and the boundary condition at time point $t = 1$ is

$$\forall x \in \mathbb{R}^d : s(1, x) = -g(x).$$

119 Function $s(t, x)$ solves the Hamilton-Jacobi-Bellman (HJB) differential equation and it in turn allows
 120 us to find the conjugate potential g^c (4).

$$-\partial_t s(t, x) = \inf_v \{v^T \nabla_x s(t, x) + \mathcal{L}(t, x, v)\} + \frac{\sigma^2}{2} \text{tr}\{\nabla^2 s(t, x)\}. \quad (7)$$

121 Representation of the Lagrange function as a sum of kinetic and potential energy allows us to find
 122 the minimum in velocity (v) in explicit form, such that $v_t = -\nabla_x s(t, x_t)$. Together with potential
 123 optimization (3), we obtain the final GSB objective in dual Kantorovich form. We provide a detailed
 124 proof in Section 6.

125 **Theorem 1** (Dual GSB problem). *Given distributions $\alpha, \beta \in \mathcal{P}(\mathbb{R}^d)$ and stochastic dynamics (1)
126 with cost functional (2), the dynamic optimal transport problem admits the following formulation:*

$$\max_{s(1, \cdot) \in L_1(\beta)} \{ \mathbb{E}_\alpha s(1, x_1) - \mathbb{E}_\beta s(1, y) \} \quad (8)$$

127 where $s(t, x) \in C^{1,2}([0, 1] \times \mathbb{R}^d)$ and satisfies HJB PDE $\forall t \in [0, 1]$ and $\forall x \in \mathbb{R}^d$

$$-\partial_t s(t, x) = -\frac{1}{2} \|\nabla_x s(t, x)\|^2 + U(x) + \frac{\sigma^2}{2} \text{tr}\{\nabla^2 s(t, x)\}. \quad (9)$$

128 The first expression in Theorem 1 plays the role of a discriminator and guarantees matching the
129 target distribution β , and the second one is responsible for the optimality of trajectories. For the
130 HJB equation to have a unique solution (in the viscosity sense), we require *coercivity* (Theorem 4.1
131 [Fleet and Soner, 2006]) of the Hamiltonian for some constants $C_1 > 0$ and $C_2 \geq 0$

$$H(x, \nabla s, \nabla^2 s) = \frac{1}{2} \|\nabla_x s\|^2 - U(x) - \frac{\sigma^2}{2} \text{tr}\{\nabla^2 s\} \quad (10)$$

$$\geq C_1(\|\nabla s\|) - C_2(1 + \|x\| + \|\nabla^2 s\|) \quad (11)$$

132 The term $\|\nabla_x s\|^2$ dominates for large values, so in case $U(x)$ is bounded and $\sigma > 0$ the solution
133 is unique. By means of the optimized function $s(t, x)$ we can generate the OT trajectories using
134 Euler-Maruyama algorithm:

$$x_{t+\Delta t} = -\nabla_x s(t, x_t) \Delta t + \sigma \Delta W, \quad x_0 \sim \alpha. \quad (12)$$

135 Unlike most other methods, here we do not need to model the intermediate density of the x_t
136 ($t \in (0, 1)$) distribution, which greatly simplifies the learning process, but we need to store the
137 generation history in a replay buffer for more stable HJB optimization in high-dimensional spaces.

138 4 Method

139 To find a stable and balanced solution $s(t, x)$ for the given dynamic OT problem (1), we can follow a
140 composite approach that combines optimal control (via HJB PDE constraints) and RL techniques
141 (policy-based trajectory optimization). We approximate the value function using a parametric model
142 $s_\theta(t, x)$. We have to maximize the potential matching functional (8) subject to the constraint that
143 $s_\theta(t, x)$ satisfies the HJB PDE. For that divide the time interval $[0, 1]$ into T time steps and simulate
144 n trajectories $\{t_0^k, x_0^k, \dots, t_T^k, x_T^k\}_{k=1}^n$ using initial α distribution and Euler-Maruyama method (12).
145 Sample also n points y_k from the target distribution β and compute the potential matching loss as

$$L_{\text{pot}}(s_\theta) = \frac{1}{n} \sum_{k=1}^n s_\theta(1, x_T^k) - \frac{1}{n} \sum_{k=1}^n s_\theta(1, y^k). \quad (13)$$

146 The HJB PDE must hold for all $t \in [0, 1]$ and $x \in \mathbb{R}^d$, but in practice, for more effective training, the
147 training data should be sampled in the region of the flow (trajectories) concentration (according to
148 Liu et al. [2022]). We enforce this by linear interpolation between datasets from α and β as a rough
149 estimation of the flow region and subsequently use the replay buffer \mathcal{B} to collect points from the
150 previously obtained trajectories. Using data samples $\{t^k, x^k\}_{k=1}^n$ from \mathcal{B} or the linear interpolation
151 we compute HJB residual loss as

$$L_{\text{hjb}}(s_\theta, \bar{s}) = \frac{1}{n} \sum_{k=1}^n \left(\frac{\partial s_\theta^k}{\partial t} - \frac{1}{2} \|\nabla_x \bar{s}^k\|^2 + U(x^k) + \frac{\sigma^2}{2} \text{tr}\{\nabla^2 \bar{s}^k\} + \lambda_a \|a^k\| \right)^2 \quad (14)$$

$$+ \frac{1}{n} \sum_{k=1}^n \left(\frac{\partial \bar{s}^k}{\partial t} - \frac{1}{2} \|\nabla_x s_\theta^k\|^2 + U(x^k) + \frac{\sigma^2}{2} \text{tr}\{\nabla^2 s_\theta^k\} + \lambda_a \|a^k\| \right)^2, \quad (15)$$

152 where $s_\theta^k = s_\theta(t^k, x^k)$, $\bar{s}^k = \bar{s}(t^k, x^k)$ denotes the target model with EMA parameters, a^k is angular
153 acceleration defined as

$$a^k = \frac{d}{dt} \frac{\nabla s_\theta(t^k, x^k)}{\|\nabla s_\theta(t^k, x^k)\|}. \quad (16)$$

154 The angular acceleration with coefficient λ_a forces the straightening of the trajectories (optionally).
155 We divide the model into s_θ and \bar{s} as it usually done in RL methods to make the optimization problem
156 more similar to regression.

157 In the result, our model is trained on two criteria (L_{pot} and L_{hjb}) simultaneously and to balance both
 158 impacts we scale the gradients of the hjb-loss and sum it with the pot-loss:

$$\nabla_{\theta} L_{\text{pot}}(s_{\theta}) + \lambda_{\text{hjb}} \text{EMA} \left(\frac{\|\nabla_{\theta} L_{\text{pot}}(s_{\theta})\|}{\|\nabla_{\theta} L_{\text{hjb}}(s_{\theta}, \bar{s})\|} \right) \nabla_{\theta} L_{\text{hjb}}(s_{\theta}, \bar{s}). \quad (17)$$

159 The complete method is implemented as shown in Algorithm 1. It effectively combines the theo-
 160 retical guarantees of optimal transport with the flexibility of neural network approximations, while
 161 maintaining numerical stability through careful gradient management. The adaptive balancing of the
 162 potential matching and HJB residual losses ensures stable convergence to a solution that satisfies both
 163 the optimality conditions and the boundary constraints.

Algorithm 1 HOTA: Hamiltonian framework for Optimal Transport Advection

- 1: **Input:** value model s_{θ} , model optimizer s_{opt} , distributions α and β , potential function $U(x)$, diffusion coefficient σ .
- 2: **Hyperparameters:** train steps N , interpolation sample steps N_0 , temporal discretization T , batch size n , hjb-loss weight λ_{hjb} , acceleration coefficient λ_a , learning rate lr, gradients scale EMA coefficient τ .
- 3: **Initialize** target model \bar{s} ; replay buffer $\mathcal{B} \leftarrow \emptyset$; gradients scale $\alpha \leftarrow 1.0$
- 4: **for** iteration $i = 1$ to N **do**
- 5: Sample train data $\{x_0^k\}_{k=1}^n \sim \alpha$; $\{y^k\}_{k=1}^n \sim \beta$
- 6: **if** $i < N_0$ **then**
- 7: Sample times $\{t^k\}_{k=1}^n \sim U(0, 1)$
- 8: For $1 \leq k \leq n$ set $x_0^k = x_0^k \cdot (1 - t^k) + y^k \cdot t^k$
- 9: **else**
- 10: Sample $\{t^k, x^k\}_{k=1}^n \sim \mathcal{B}$
- 11: **end if**
- 12: Generate n trajectories $\{t_0^k, x_0^k, \dots, t_T^k, x_T^k\}_{k=1}^n$ using current policy $v_t = -\nabla s(t, x)$
- 13: Add the 1-st trajectory $\{t_0^0, x_0^0, \dots, t_T^0, x_T^0\}$ to \mathcal{B}
- 14: **Compute gradients:**
- 15: $g_{\text{hjb}} = \nabla_{\theta} L_{\text{hjb}}(s_{\theta}, \bar{s}, \{t^k, x^k\}_{k=1}^n)$
- 16: $g_{\text{pot}} = \nabla_{\theta} L_{\text{pot}}(s_{\theta}, \{x_T^k\}_{k=1}^n, \{y^k\}_{k=1}^n)$
- 17: **Update Parameters:**
- 18: Compute norms $G_{\text{hjb}} = \|g_{\text{hjb}}\|_2$ and $G_{\text{pot}} = \|g_{\text{pot}}\|_2$
- 19: EMA update of gradients scale $\alpha = \tau G_{\text{pot}} / G_{\text{hjb}} + (1 - \tau) \alpha$
- 20: Sum the gradients $g = g_{\text{pot}} + \lambda_{\text{hjb}} \alpha g_{\text{hjb}}$
- 21: Update model parameters θ with $s_{\text{opt}}(g)$
- 22: EMA update of target model \bar{s}
- 23: **end for**

164

5 Experiments

165 In this section, we evaluate our method on a series of distribution matching tasks with non-trivial
 166 geometries. In Section 5.2, we compare HOTA with state-of-the-art baselines, demonstrating its
 167 superior performance on both standard benchmarks including datasets with almost non-differentiable
 168 potentials. In Section 5.3, we demonstrate the scalability of our approach by showcasing its ef-
 169 fectiveness in high-dimensional settings. Finally, in Section 5.4, we ablate key components of our
 170 method.

171

5.1 Experimental Setup

172 **Evaluation** We assess performance using two metrics: *feasibility* and *optimality*. Feasibility reflects
 173 how well the method matches the target distribution, evaluated via Wasserstein distance with squared
 174 Euclidean cost ($W_2(T_{\#} \alpha, \beta)$), where the transport mapping T uses optimized value function s_{θ} and
 175 samples x_1 by procedure (12). Optimality measures the quality of the resulting mapping, estimated
 176 through the integral trajectory cost: $\int_0^1 \left[\mathbb{E}_{\rho_t} \frac{\|v_t(x_t)\|^2}{2} + U(x_t) \right] dt$, where x_t follows (1).

177 **Network** In all our experiments, we employ a simple MLP augmented with Fourier feature encoding
 178 of the time component. For general time embeddings of the form $\text{emb}(t) = \sin(f \cdot t + \varphi)$, the time

179 derivative is given by $\partial_t \text{emb}(t) = f \cdot \cos(f \cdot t + \varphi)$. As the frequency f increases, the magnitude
 180 of this derivative also grows, potentially leading to numerical instability—especially when the time
 181 derivative of the network is explicitly involved in the objective. This issue has been previously
 182 discussed in [Lu and Song \[2024\]](#). To address this, we restrict the frequency range to $[1, 20]$ and
 183 normalize the resulting Fourier features by dividing by the corresponding frequencies.

184 **Baselines** We use source code from [GSBM repository](#) for running it in our experiments on BabyMaze,
 185 Slit and Box datasets. Other results were taken from the original papers [Liu et al. \[2024\]](#), [Pooladian](#)
 186 [et al. \[2024\]](#) where dataset were previously introduced.

187 All experiments are conducted on a GeForce RTX 3090 GPU and take less than ten minutes for
 188 training. Additional experimental details are provided in Appendix [A](#).

189 **5.2 Comparative Evaluation on Two-Dimensional Data**

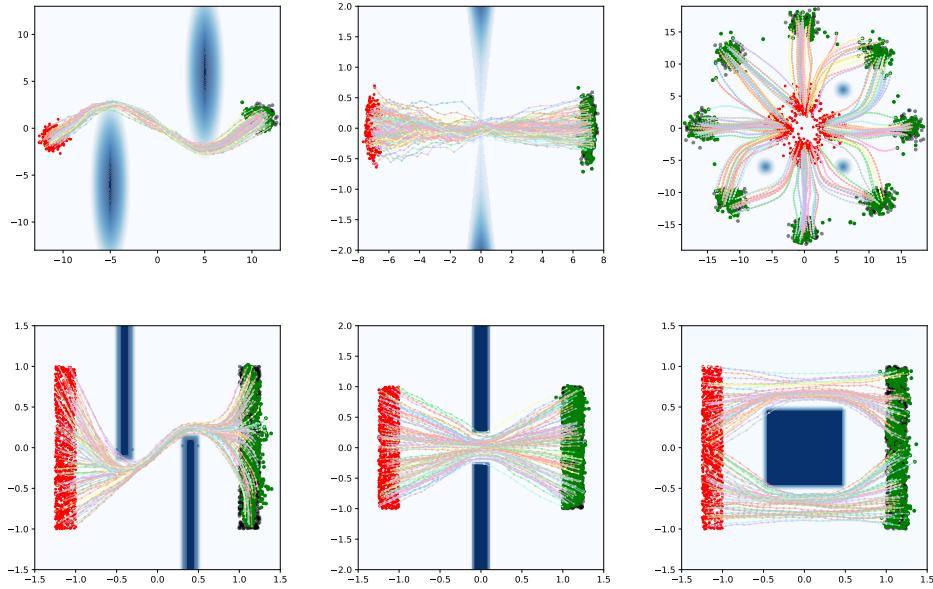


Figure 1: Evaluation of HOTA method on smooth (top) and non-smooth datasets (bottom): Stunnel, Vneck, GMM, BabyMaze, Slit, Box. Blue regions indicate high values of potential $U(x)$. Distributions α (red), β (black) and the mapped $T_{\#}\alpha$ (green).

190 In this section, we compare our method to previous state-of-the-art approaches on the standard benchmarks
 191 including datasets that feature almost non-differentiable potential functions. Visualizations
 192 of the datasets are provided in Figure 1. The first three datasets—Stunnel, Vneck, and GMM—are
 193 adopted from [Liu et al. \[2024\]](#). These benchmarks incorporate state cost functions $U(x_t)$ that en-
 194 courage the optimal solution to respect complex geometric constraints. Each dataset is designed to
 195 highlight specific capabilities of the evaluated algorithms. *Stunnel* assesses whether a method can
 196 capture drift fields that undergo rapid and localized changes. *Vneck* evaluates the ability to model
 197 drift that compresses and expands the support of marginal distributions. *GMM* tests whether the
 198 method can disambiguate closely situated points and assign them to distinct trajectories. The re-
 199 maining datasets—BabyMaze, Slit, and Box ([Pooladian et al. \[2024\]](#))—are constructed using similar
 200 underlying principles but pose additional difficulties due to the presence of almost non-differentiable
 201 state cost functions. A summary of the quantitative results across all datasets is provided in Table 1.
 202 Our method, HOTA, consistently outperforms existing approaches in terms of both feasibility and
 203 optimality. In particular, HOTA achieves a substantial performance gain on the GMM dataset, which
 204 may refer to its superior capability in trajectory separation for closely situated points.

205 **5.3 Scalability to High-Dimensional Spaces**

206 In this section, we test the scalability of our method, demonstrating its stable performance in higher-
 207 dimensional settings. For this purpose, we use *Sphere* datasets parameterized by data dimensionality

Table 1: Quantitative comparison between recent state-of-the-art methods and our approach, HOTA. Performance is evaluated using two criteria: *Feasibility* (how well the target distribution is covered) and *Optimality* (efficiency of the learned mapping). Our method consistently outperforms existing approaches, with significantly better results in certain tasks, such as GMM. N/A cells indicate that original authors of particular method did not include results for those tasks. The mean and the standard deviations of our method are computed across 5 different seeds. Best values are highlighted by **bold** font (lower is better). Gray values correspond to the method’s divergence.

| | Feasibility $W_2(T_\#(\alpha), \beta)$ | | | Optimality (integral cost) | | |
|-------------|--|---------------------------------------|-------------------------------------|-------------------------------------|------------------------------------|-----------------------------------|
| | Stunnel | Vneck | GMM | Stunnel | Vneck | GMM |
| NLSB | 30.54 | 0.02 | 67.76 | 207.06 | 147.85 | 4202.71 |
| GSBM | 0.03 | 0.01 | 4.13 | 460.88 | 155.53 | 229.12 |
| HOTA | 0.006 ± 0.003 | 0.002 ± 0.001 | 0.19 ± 0.05 | 320.90 ± 12.5 | 115.09 ± 8.9 | 80.44 ± 2.6 |
| BabyMaze | | | | | | |
| NLSB | > 1 | 0.013 | 0.024 | N/A | N/A | N/A |
| NLOT | > 1 | 0.013 | 0.016 | N/A | N/A | N/A |
| GSBM | 0.01 | 0.01 | 0.02 | 6.5 | 4.9 | 3.8 |
| HOTA | 0.004 ± 0.003 | 0.0004 ± 0.0001 | 0.002 ± 0.001 | 4.87 ± 0.14 | 3.06 ± 0.09 | 2.84 ± 0.11 |

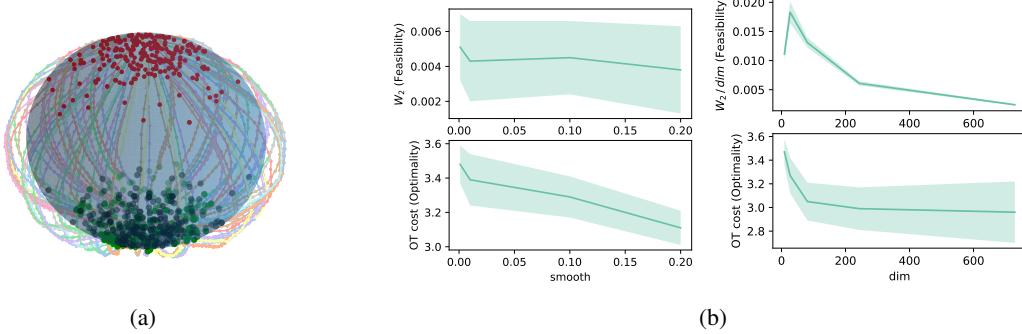


Figure 2: (a) Visualization of *Sphere* dataset for $N = 3$. (b) *Feasibility* and *Optimality* trends with respect to 3D unit sphere smoothness (left) and unit sphere dimensionality (right). Our method maintains robust performance across both non-differentiable potentials and high-dimensional settings.

208 N . Specifically, we define an N -dimensional unit sphere as a potential barrier inducing corresponding
 209 state cost function $U(x_t)$. The source and target distributions are samples from a standard distribution
 210 located at the poles, projected onto the unit sphere. The three-dimensional case is visualized in
 211 Figure 2a. The performance of our method across varying data dimensions is shown in Figure 2b
 212 (right). Notably, HOTA demonstrates robust and stable performance as the dimensionality N
 213 increases.

214 5.4 Ablation study

215 Table 2 presents comparison of the full HOTA model against variants without the replay buffer \mathcal{B}
 216 that stores simulation history or the adaptive gradient balancing by means of α (17), evaluating as
 217 previously feasibility and optimality metrics across Stunnel, Vneck, and GMM datasets. The full
 218 HOTA achieves strong metric scores, while removing the buffer severely degrades feasibility in
 219 Vneck and GMM and increases costs in Stunnel. Disabling gradient balancing harms feasibility in
 220 Stunnel and GMM. The results highlight the buffer’s critical role in maintaining feasibility and the
 221 nuanced trade-offs between gradient balancing and transport efficiency across different scenarios.

222 Additionally we have evaluated the influence of acceleration term $\lambda_a \|a\|$ used in loss L_{hjb} depending
 223 on λ_a (Figure 3). It performs the function of straightening trajectories by penalizing the change
 224 in angular velocity. It follows from the results that increasing λ_a improves the optimality of the
 225 transportation trajectories while introducing a small bias in the matching of the target distribution
 226 β , which is reflected in the feasibility metric. In the GMM task, due to the specificity of the dataset

227 and the divergence of trajectories in different directions, a small penalization of acceleration also improves feasibility.

Table 2: Comparison of HOTA method against variants without the replay buffer \mathcal{B} and the adaptive gradient balancing. Best values are highlighted by bold font (lower is better). Gray values correspond to the method’s divergence.

| | Feasibility $W_2(T_\#(\alpha), \beta)$ | | | Optimality (integral cost) | | |
|--------------------------|--|--------------|-------------|----------------------------|---------------|--------------|
| | Stunnel | Vneck | GMM | Stunnel | Vneck | GMM |
| HOTA | 0.006 | 0.002 | 0.19 | 320.90 | 115.09 | 80.44 |
| HOTA w/o buffer | 0.076 | 16.47 | 1.248 | 706.89 | 82.49 | 121.6 |
| HOTA w/o grad. balancing | 3.60 | 0.026 | 2.64 | 325.22 | 109.25 | 72.77 |

228

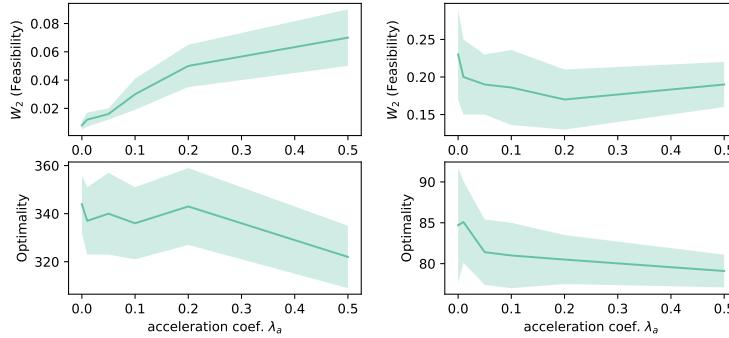


Figure 3: Impact of acceleration coefficient λ_a . Left: Stunnel, right: GMM datasets.

229 6 Proof of Theorem 1 (Dual Formulation of GSB)

230 We prove in the **first step** the equivalence between the GSB (stochastic control formulation) and its
231 dual formulation using Kantorovich-style duality. Remind that we consider the stochastic process
232 x_t (1) with conditions $x_0 \sim \alpha$, $x_1 \sim \beta$, control function $v(t, x_t)$, and Brownian motion $\sigma(t, x_t) dW_t$.
233 The **primary problem** of GSB optimization is:

$$\inf_{v(t,x)} \mathbb{E} \left[\int_0^1 \mathcal{L}(t, x_t, v_t) dt \right] \quad \text{s.t.} \quad x_0 \sim \alpha, x_1 \sim \beta, \quad (18)$$

234 where in the particular case $\mathcal{L}(t, x, v) = v^2/2 + U(x)$. Since the stochastic process x_t starts from
235 $x_0 \sim \alpha$ the primal problem is equivalent to:

$$\inf_{v(t,x)} \left(\mathbb{E} \left[\int_0^1 \mathcal{L}(t, x_t, v_t) dt \right] + \sup_{g \in L_1(\beta)} (-\mathbb{E}[g(x_1)] + \mathbb{E}_\beta[g(y)]) \right), \quad (19)$$

236 where the supremum over g enforces the constraint $x_1 \sim \beta$ (via Lagrange duality). Rewrite the
237 Lagrangian problem as

$$\inf_{v(t,x)} \sup_{g \in L_1(\beta)} \left(\mathbb{E} \left[\int_0^1 \mathcal{L}(t, x_t, v_t) dt - g(x_1) \right] + \mathbb{E}_\beta[g(y)] \right). \quad (20)$$

238 Assuming strong duality holds under mild regularity conditions (e.g., \mathcal{L} convex in v , α, β absolutely
239 continuous), we swap inf and sup:

$$\sup_{g \in L_1(\beta)} \left(\inf_{v(t,x)} \mathbb{E} \left[\int_0^1 \mathcal{L}(t, x_t, v_t) dt - g(x_1) \right] + \mathbb{E}_\beta[g(y)] \right). \quad (21)$$

240 Note that since the optimal $v^*(t, x)$ is Markovian (depends only on current time t and state x) and
241 does not depend on the initial distribution α it holds that

$$\mathbb{E} \left[\int_0^1 \mathcal{L}(t, x_t, v_t^*) dt - g(x_1) \right] = \mathbb{E}_{x \sim \alpha} \mathbb{E} \left[\int_0^1 \mathcal{L}(t, x_t, v_t^*) dt - g(x_1) \mid x_0 = x \right]. \quad (22)$$

242 Buy the definition of c -conjugate transform (5):

$$\mathbb{E}_{x \sim \alpha} \mathbb{E} \left[\int_0^1 \mathcal{L}(t, x_t, v_t^*) dt - g(x_1) \mid x_0 = x \right] = \mathbb{E}_{x \sim \alpha} g^c(x). \quad (23)$$

243 Thus, the **dual problem** becomes: $\sup_{g \in L_1(\beta)} (\mathbb{E}_\alpha[g^c(x)] + \mathbb{E}_\beta[g(y)])$. In the **second step** find the 244 optimal control solution $v^*(t, x)$ by means of dynamic programming principle. Define the value 245 function $s(t, x)$ that for any $0 \leq t \leq \tau \leq 1$ satisfies:

$$s(t, x) = \inf_{v(t, x)} \mathbb{E} \left[\int_t^\tau \mathcal{L}(z, x_z, v_z) dz + s(\tau, x_\tau) \mid x_t = x \right]. \quad (24)$$

246 Applying Ito's formula to $s(\tau, x_\tau)$ we obtain that

$$ds(\tau, x_\tau) = \partial_\tau s d\tau + \nabla s \cdot dx_\tau + \frac{1}{2} \text{tr}(\sigma^2 \nabla^2 s) d\tau \quad (25)$$

$$= \left(\partial_\tau s + \nabla s^T v_\tau + \frac{1}{2} \text{tr}(\sigma^2 \nabla^2 s) \right) d\tau + \nabla s^T \sigma dW_s. \quad (26)$$

247 Consider the evolution of the value between times t and τ :

$$s(\tau, x_\tau) - s(t, x_t) = \int_t^\tau \left(\partial_z s + \nabla s \cdot v_z + \frac{1}{2} \text{tr}(\sigma^2 \nabla^2 s) \right) dz + \int_t^\tau \nabla s^T \sigma dW. \quad (27)$$

248 Basing on the martingale property of Ito integrals ($\mathbb{E}[\int \nabla s \cdot \sigma dW | x_t = x] = 0$) it holds that

$$\mathbb{E}[s(\tau, x_\tau) | x_t = x] = s(t, x) + \mathbb{E} \left[\int_t^\tau \left(\partial_z s + \nabla s \cdot v_z + \frac{1}{2} \text{tr}(\sigma^2 \nabla^2 s) \right) dz \right]. \quad (28)$$

249 Substitute back into dynamic programming and plug the last expression into the equation (24):

$$s(t, x) = \inf_{v(t, x)} \mathbb{E} \left[\int_t^\tau \mathcal{L}(z, x_z, v_z) dz + s(t, x) + \int_t^\tau \left(\partial_z s + \nabla s^T v_\tau + \frac{1}{2} \text{tr}(\sigma^2 \nabla^2 s) \right) dz \right]. \quad (29)$$

250 Cancel $s(t, x)$ from both sides and divide by $(\tau - t)$:

$$0 = \inf_{v(s, t)} \frac{1}{\tau - t} \mathbb{E} \left[\int_t^\tau \left(\mathcal{L}(z, x_z, v_z) + \partial_z s + \nabla s^T v_z + \frac{1}{2} \text{tr}(\sigma^2 \nabla^2 s) \right) dz \right]. \quad (30)$$

251 Take limit $\tau \downarrow t$ to derive the HJB equation for a general Lagrangian \mathcal{L}

$$0 = \inf_v \left\{ \mathcal{L}(t, x, v) + \partial_t s + \nabla s^T v + \frac{1}{2} \text{tr}(\sigma^2 \nabla^2 s) \right\}. \quad (31)$$

252 Identify optimal control for the particular $\mathcal{L}(t, x, v) = v^2/2 + U(x)$. The infimum is attained when 253 $v^* = -\nabla s$, yielding the final result of Theorem 1.

254 7 Limitations and Future Work

255 While HOTA exhibits strong and robust performance, we observed sensitivity to certain network 256 design choices—particularly the Fourier feature encoding of time, a commonly used technique in 257 models that estimate ODE drifts. Additionally, because the value function in our framework must 258 simultaneously support optimal control estimation and serve as a Kantorovich potential, it requires 259 a network architecture capable of aggregating rich temporal and spatial information. The use of a 260 simple MLP, while effective, may not be optimal from an optimization standpoint. Incorporating 261 architectures with stronger inductive biases could further enhance performance. These considerations 262 lie beyond the scope of this work, but we believe they offer promising directions for future research.

263 8 Conclusion

264 In this work, we introduced HOTA, a new OT method based on the Hamilton–Jacobi–Bellman 265 (HJB) framework for solving the Generalized Schrödinger Bridge problem. We demonstrated that 266 HOTA consistently outperforms recent state-of-the-art methods on standard benchmarks and scales 267 effectively to high-dimensional settings. Remarkably, it works for non-smooth potentials and with 268 non-differentiable cost functions, yielding robust performance gain in terms of strictly defined 269 concepts of feasibility and optimality.

270 **References**

271 Arip Asadulaev, Rostislav Korst, Aleksandr Korotin, Vage Egiazarian, Andrey Filchenkov, and Evgeny Burnaev.
272 Rethinking optimal transport in offline reinforcement learning. *Advances in Neural Information Processing
273 Systems*, 37:123592–123607, 2024.

274 Grigory Bartosh, Dmitry P Vetrov, and Christian Andersson Naesseth. Neural flow diffusion models: Learnable
275 forward process for improved diffusion modelling. *Advances in Neural Information Processing Systems*, 37:
276 73952–73985, 2024.

277 Jean-David Benamou and Yann Brenier. A computational fluid mechanics solution to the monge-kantorovich
278 mass transfer problem. *Numerische Mathematik*, 84(3):375–393, 2000.

279 Denis Blessing, Julius Berner, Lorenz Richter, and Gerhard Neumann. Underdamped diffusion bridges with
280 applications to sampling. In *The Thirteenth International Conference on Learning Representations*, 2025.
281 URL <https://openreview.net/forum?id=Q1QTxFm0Is>.

282 Maksim Bobrin, Nazar Buzun, Dmitrii Krylov, and Dmitry V Dylov. Align your intents: Offline imitation
283 learning via optimal transport. *arXiv preprint arXiv:2402.13037*, 2024.

284 Nicolas Bonneel and Julie Digne. A survey of optimal transport for computer graphics and computer vision.
285 *Computer Graphics Forum*, 42(2):439–460, 2023. doi: <https://doi.org/10.1111/cgf.14778>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.14778>.

287 Charlotte Bunne, Laetitia Papaxanthos, Andreas Krause, and Marco Cuturi. Proximal optimal transport
288 modeling of population dynamics. In Gustau Camps-Valls, Francisco J. R. Ruiz, and Isabel Valera, editors,
289 *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, volume 151 of
290 *Proceedings of Machine Learning Research*, pages 6511–6528. PMLR, 28–30 Mar 2022. URL <https://proceedings.mlr.press/v151/bunne22a.html>.

292 Nazar Buzun, Maksim Bobrin, and Dmitry V Dylov. Expectile regularization for fast and accurate training of
293 neural optimal transport. *Advances in Neural Information Processing Systems*, 37:119811–119837, 2024.

294 Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in neural information
295 processing systems*, 26, 2013.

296 Carles Domingo-Enrich, Michal Drozdzal, Brian Karrer, and Ricky TQ Chen. Adjoint matching: Fine-tuning flow
297 and diffusion generative models with memoryless stochastic optimal control. *arXiv preprint arXiv:2409.08861*,
298 2024a.

299 Carles Domingo-Enrich, Jiequn Han, Brandon Amos, Joan Bruna, and Ricky T. Q. Chen. Stochastic optimal
300 control matching. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024b.
301 URL <https://openreview.net/forum?id=wfU2CdgmWt>.

302 Wendell H. Fleet and H. Mete Soner. Controlled markov processes and viscosity solutions (2nd ed.). *Springer*,
303 2006.

304 Kacper Kausniak, Peter Potapchik, Teodora Reu, Leo Zhang, Alexander Tong, Michael Bronstein, Joey Bose,
305 and Francesco Di Giovanni. Metric flow matching for smooth interpolations on the data manifold. *Advances
306 in Neural Information Processing Systems*, 37:135011–135042, 2024.

307 Pascal Klink, Haoyi Yang, Carlo D’Eramo, Jan Peters, and Joni Pajarinen. Curriculum reinforcement learning
308 via constrained optimal transport. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari,
309 Gang Niu, and Sivan Sabato, editors, *Proceedings of the 39th International Conference on Machine Learning*,
310 volume 162 of *Proceedings of Machine Learning Research*, pages 11341–11358. PMLR, 17–23 Jul 2022.
311 URL <https://proceedings.mlr.press/v162/klink22a.html>.

312 Alexander Korotin, Daniil Selikhanovich, and Evgeny Burnaev. Neural optimal transport. *arXiv preprint
313 arXiv:2201.12220*, 2022.

314 Sergey Levine. Reinforcement learning and control as probabilistic inference: Tutorial and review. *CoRR*,
315 abs/1805.00909, 2018.

316 Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow matching for
317 generative modeling. In *The Eleventh International Conference on Learning Representations*, 2023.

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

320 Guan-Horng Liu, Yaron Lipman, Maximilian Nickel, Brian Karrer, Evangelos Theodorou, and Ricky T. Q.
321 Chen. Generalized schrödinger bridge matching. In *The Twelfth International Conference on Learning
322 Representations*, 2024.

323 Cheng Lu and Yang Song. Simplifying, stabilizing and scaling continuous-time consistency models. 2024.

324 Ashok Makkluva, Amirhossein Taghvaei, Sewoong Oh, and Jason Lee. Optimal transport mapping via input
325 convex neural networks. In *International Conference on Machine Learning*, pages 6672–6681. PMLR, 2020.

326 Eduardo Fernandes Montesuma, Fred Maurice Ngole Mboula, and Antoine Souloumiac. Recent advances in
327 optimal transport for machine learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*,

328 2024.

329 Benjamin Nachman Nathan T. Suri, Vinicius Mikuni. Wotan: Weakly-supervised optimal transport attention-
330 based noise mitigation. *NeurIPS 2024*, 2024.

331 Kirill Neklyudov, Rob Brekelmans, Alexander Tong, Lazar Atanackovic, Qiang Liu, and Alireza Makhzani. A
332 computational framework for solving Wasserstein lagrangian flows. In Ruslan Salakhutdinov, Zico Kolter,
333 Katherine Heller, Adrian Weller, Nuria Oliver, Jonathan Scarlett, and Felix Berkenkamp, editors, *Proceedings
334 of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning
335 Research*, pages 37461–37485. PMLR, 21–27 Jul 2024.

336 Luiz Manella Pereira and M Hadi Amini. A survey on optimal transport for machine learning: Theory and
337 applications. *IEEE Access*, 2025.

338 Gabriel Peyré, Marco Cuturi, et al. Computational optimal transport: With applications to data science.
339 *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, 2019.

340 Aram-Alexandre Pooladian, Carles Domingo-Enrich, Ricky T. Q. Chen, and Brandon Amos. Neural optimal
341 transport with lagrangian costs. In *The 40th Conference on Uncertainty in Artificial Intelligence*, 2024.

342 Thomas Rupf, Marco Bagatella, Nico Görtler, Jonas Frey, and Georg Martius. Zero-shot offline imitation
343 learning via optimal transport, 2025. URL <https://openreview.net/forum?id=vDecbmWf6w>.

344 Filippo Santambrogio. *Optimal transport for applied mathematicians*, volume 87. Springer, 2015.

345 Evangelos Theodorou, Jonas Buchli, and Stefan Schaal. Learning policy improvements with path integrals. In
346 Yee Whye Teh and Mike Titterington, editors, *Proceedings of the Thirteenth International Conference on
347 Artificial Intelligence and Statistics*, volume 9 of *Proceedings of Machine Learning Research*, pages 828–835,
348 Chia Laguna Resort, Sardinia, Italy, 13–15 May 2010. PMLR.

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

350 **A Additional Experimental Details**

351 **Hyperparameters** Table 3 summarizes the hyperparameters used for each dataset presented in the
 352 paper. Note that the Sphere datasets, which are parameterized by data dimensionality, share all
 353 hyperparameters except for the potential weight, which may take value 10 for the low dimensions are
 30 for high ones.

Table 3: Hyperparameters used for each dataset presented in the paper.

| Hyperparameter | Stunnel | Vneck | GMM | BabyMaze | Slit | Box | Sphere |
|-----------------------------|---------|-------|-----|--|-------|------|----------|
| MLP hidden layers | | | | [512, 512, 512, 1] | | | |
| Fourier frequencies | | | | {1, ..., 20} | | | |
| optimizer | | | | Adam with cosine annealing ($\alpha = 1e-2$) | | | |
| initial learning rate | | | | 5×10^{-4} | | | |
| Adam $[\beta_1, \beta_2]$ | | | | [0.9, 0.99] | | | |
| # training iterations | | | | 70000 | | | |
| batch size | | | | 1024 | | | |
| EMA decay, τ | | | | 0.9 | | | |
| # control steps | | | | 30 | | | |
| diffusion coef., σ | 0.3 | 0.2 | 0.1 | 0.03 | 0.05 | 0.03 | 0.01 |
| control weight, λ_a | 1.0 | 2.0 | 0.7 | 0.5 | 2.0 | 0.3 | 0.4 |
| acc. weight, λ_a | 0.0001 | 0.001 | 0.2 | 0.05 | 0.001 | 0.01 | 0 |
| potential weight | 25 | 1000 | 25 | 10 | 30 | 700 | {10, 30} |

354

355 **NeurIPS Paper Checklist**

356 **1. Claims**

357 Question: Do the main claims made in the abstract and introduction accurately reflect the
358 paper's contributions and scope?

359 Answer: [\[Yes\]](#)

360 Justification: We provide theoretical derivation of the dual SGB problem. We made a
361 thorough experiments on established benchmarks and compare proposed method with the
362 other GSB solvers. In all extensive tests, HOTA outperforms the competition both in terms
363 of feasibility and optimality.

364 Guidelines:

- 365 • The answer NA means that the abstract and introduction do not include the claims
366 made in the paper.
- 367 • The abstract and/or introduction should clearly state the claims made, including the
368 contributions made in the paper and important assumptions and limitations. A No or
369 NA answer to this question will not be perceived well by the reviewers.
- 370 • The claims made should match theoretical and experimental results, and reflect how
371 much the results can be expected to generalize to other settings.
- 372 • It is fine to include aspirational goals as motivation as long as it is clear that these goals
373 are not attained by the paper.

374 **2. Limitations**

375 Question: Does the paper discuss the limitations of the work performed by the authors?

376 Answer: [\[Yes\]](#)

377 Justification: We provided discussions on limitations in Section 6 of main paper.

378 Guidelines:

- 379 • The answer NA means that the paper has no limitation while the answer No means that
380 the paper has limitations, but those are not discussed in the paper.
- 381 • The authors are encouraged to create a separate "Limitations" section in their paper.
- 382 • The paper should point out any strong assumptions and how robust the results are to
383 violations of these assumptions (e.g., independence assumptions, noiseless settings,
384 model well-specification, asymptotic approximations only holding locally). The authors
385 should reflect on how these assumptions might be violated in practice and what the
386 implications would be.
- 387 • The authors should reflect on the scope of the claims made, e.g., if the approach was
388 only tested on a few datasets or with a few runs. In general, empirical results often
389 depend on implicit assumptions, which should be articulated.
- 390 • The authors should reflect on the factors that influence the performance of the approach.
391 For example, a facial recognition algorithm may perform poorly when image resolution
392 is low or images are taken in low lighting. Or a speech-to-text system might not be
393 used reliably to provide closed captions for online lectures because it fails to handle
394 technical jargon.
- 395 • The authors should discuss the computational efficiency of the proposed algorithms
396 and how they scale with dataset size.
- 397 • If applicable, the authors should discuss possible limitations of their approach to
398 address problems of privacy and fairness.
- 399 • While the authors might fear that complete honesty about limitations might be used by
400 reviewers as grounds for rejection, a worse outcome might be that reviewers discover
401 limitations that aren't acknowledged in the paper. The authors should use their best
402 judgment and recognize that individual actions in favor of transparency play an impor-
403 tant role in developing norms that preserve the integrity of the community. Reviewers
404 will be specifically instructed to not penalize honesty concerning limitations.

405 **3. Theory assumptions and proofs**

406 Question: For each theoretical result, does the paper provide the full set of assumptions and
407 a complete (and correct) proof?

408 Answer: [Yes]

409 Justification: We provided a theoretical proof of dual GSB task formulation.

410 Guidelines:

- 411 • The answer NA means that the paper does not include theoretical results.
- 412 • All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- 413 • All assumptions should be clearly stated or referenced in the statement of any theorems.
- 414 • The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- 415 • Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- 416 • Theorems and Lemmas that the proof relies upon should be properly referenced.

417 4. Experimental result reproducibility

418 Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

419 Answer: [Yes]

420 Justification: Along with detailed hyperparameter specifications in the Appendix, we included easy to follow Jupyter notebook which can be found in supplementary materials, enabling the others to fully reproduce the results in the paper.

421 Guidelines:

- 422 • The answer NA means that the paper does not include experiments.
- 423 • If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- 424 • If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- 425 • Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- 426 • While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - 427 (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - 428 (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - 429 (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - 430 (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

431 5. Open access to data and code

462 Question: Does the paper provide open access to the data and code, with sufficient instruc-
463 tions to faithfully reproduce the main experimental results, as described in supplemental
464 material?

465 Answer: [Yes]

466 Justification: We provided a fully reproducible code, written in JAX framework. Moreover,
467 we provided step-by-step jupyter notebook, showcasing the performance of the proposed
468 algorithm in all discussed tasks.

469 Guidelines:

- 470 • The answer NA means that paper does not include experiments requiring code.
- 471 • Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- 473 • While we encourage the release of code and data, we understand that this might not be
474 possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not
475 including code, unless this is central to the contribution (e.g., for a new open-source
476 benchmark).
- 477 • The instructions should contain the exact command and environment needed to run to
478 reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- 480 • The authors should provide instructions on data access and preparation, including how
481 to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- 482 • The authors should provide scripts to reproduce all experimental results for the new
483 proposed method and baselines. If only a subset of experiments are reproducible, they
484 should state which ones are omitted from the script and why.
- 485 • At submission time, to preserve anonymity, the authors should release anonymized
486 versions (if applicable).
- 487 • Providing as much information as possible in supplemental material (appended to the
488 paper) is recommended, but including URLs to data and code is permitted.

489 6. Experimental setting/details

490 Question: Does the paper specify all the training and test details (e.g., data splits, hyper-
491 parameters, how they were chosen, type of optimizer, etc.) necessary to understand the
492 results?

493 Answer: [Yes]

494 Justification: We provided detailed hyperparameters specifications in the Appendix for each
495 of the tested benchmarks.

496 Guidelines:

- 497 • The answer NA means that the paper does not include experiments.
- 498 • The experimental setting should be presented in the core of the paper to a level of detail
499 that is necessary to appreciate the results and make sense of them.
- 500 • The full details can be provided either with the code, in appendix, or as supplemental
501 material.

502 7. Experiment statistical significance

503 Question: Does the paper report error bars suitably and correctly defined or other appropriate
504 information about the statistical significance of the experiments?

505 Answer: [Yes]

506 Justification: All reported results are statistically significant. We include evaluation error
507 (StDev) for each model and each dataset in the study across different runs and seeds.

508 Guidelines:

- 509 • The answer NA means that the paper does not include experiments.
- 510 • The authors should answer “Yes” if the results are accompanied by error bars, confi-
511 dence intervals, or statistical significance tests, at least for the experiments that support
512 the main claims of the paper.

513 • The factors of variability that the error bars are capturing should be clearly stated (for
 514 example, train/test split, initialization, random drawing of some parameter, or overall
 515 run with given experimental conditions).
 516 • The method for calculating the error bars should be explained (closed form formula,
 517 call to a library function, bootstrap, etc.)
 518 • The assumptions made should be given (e.g., Normally distributed errors).
 519 • It should be clear whether the error bar is the standard deviation or the standard error
 520 of the mean.
 521 • It is OK to report 1-sigma error bars, but one should state it. The authors should
 522 preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis
 523 of Normality of errors is not verified.
 524 • For asymmetric distributions, the authors should be careful not to show in tables or
 525 figures symmetric error bars that would yield results that are out of range (e.g. negative
 526 error rates).
 527 • If error bars are reported in tables or plots, The authors should explain in the text how
 528 they were calculated and reference the corresponding figures or tables in the text.

529 8. **Experiments compute resources**

530 Question: For each experiment, does the paper provide sufficient information on the com-
 531 puter resources (type of compute workers, memory, time of execution) needed to reproduce
 532 the experiments?

533 Answer: [\[Yes\]](#)

534 Justification: We include the exact computer configuration in Appendix and mention GPU
 535 model in the main text.

536 Guidelines:

537 • The answer NA means that the paper does not include experiments.
 538 • The paper should indicate the type of compute workers CPU or GPU, internal cluster,
 539 or cloud provider, including relevant memory and storage.
 540 • The paper should provide the amount of compute required for each of the individual
 541 experimental runs as well as estimate the total compute.
 542 • The paper should disclose whether the full research project required more compute
 543 than the experiments reported in the paper (e.g., preliminary or failed experiments that
 544 didn't make it into the paper).

545 9. **Code of ethics**

546 Question: Does the research conducted in the paper conform, in every respect, with the
 547 NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

548 Answer: [\[Yes\]](#)

549 Justification: We read it and adhered to the ethical guidelines.

550 Guidelines:

551 • The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
 552 • If the authors answer No, they should explain the special circumstances that require a
 553 deviation from the Code of Ethics.
 554 • The authors should make sure to preserve anonymity (e.g., if there is a special consid-
 555 eration due to laws or regulations in their jurisdiction).

556 10. **Broader impacts**

557 Question: Does the paper discuss both potential positive societal impacts and negative
 558 societal impacts of the work performed?

559 Answer: [\[No\]](#)

560 Justification: Our paper does not address the societal impact as we operate with common
 561 datasets and benchmarks for testing GSB solvers.

562 Guidelines:

563 • The answer NA means that there is no societal impact of the work performed.

- 564 • If the authors answer NA or No, they should explain why their work has no societal
565 impact or why the paper does not address societal impact.
- 566 • Examples of negative societal impacts include potential malicious or unintended uses
567 (e.g., disinformation, generating fake profiles, surveillance), fairness considerations
568 (e.g., deployment of technologies that could make decisions that unfairly impact specific
569 groups), privacy considerations, and security considerations.
- 570 • The conference expects that many papers will be foundational research and not tied
571 to particular applications, let alone deployments. However, if there is a direct path to
572 any negative applications, the authors should point it out. For example, it is legitimate
573 to point out that an improvement in the quality of generative models could be used to
574 generate deepfakes for disinformation. On the other hand, it is not needed to point out
575 that a generic algorithm for optimizing neural networks could enable people to train
576 models that generate Deepfakes faster.
- 577 • The authors should consider possible harms that could arise when the technology is
578 being used as intended and functioning correctly, harms that could arise when the
579 technology is being used as intended but gives incorrect results, and harms following
580 from (intentional or unintentional) misuse of the technology.
- 581 • If there are negative societal impacts, the authors could also discuss possible mitigation
582 strategies (e.g., gated release of models, providing defenses in addition to attacks,
583 mechanisms for monitoring misuse, mechanisms to monitor how a system learns from
584 feedback over time, improving the efficiency and accessibility of ML).

585 11. Safeguards

586 Question: Does the paper describe safeguards that have been put in place for responsible
587 release of data or models that have a high risk for misuse (e.g., pretrained language models,
588 image generators, or scraped datasets)?

589 Answer: [NA]

590 Justification: Not applicable to this work.

591 Guidelines:

- 592 • The answer NA means that the paper poses no such risks.
- 593 • Released models that have a high risk for misuse or dual-use should be released with
594 necessary safeguards to allow for controlled use of the model, for example by requiring
595 that users adhere to usage guidelines or restrictions to access the model or implementing
596 safety filters.
- 597 • Datasets that have been scraped from the Internet could pose safety risks. The authors
598 should describe how they avoided releasing unsafe images.
- 599 • We recognize that providing effective safeguards is challenging, and many papers do
600 not require this, but we encourage authors to take this into account and make a best
601 faith effort.

602 12. Licenses for existing assets

603 Question: Are the creators or original owners of assets (e.g., code, data, models), used in
604 the paper, properly credited and are the license and terms of use explicitly mentioned and
605 properly respected?

606 Answer: [Yes]

607 Justification: We properly refer to the original papers and use the open source codes from
608 official repositories, providing the URLs to them.

609 Guidelines:

- 610 • The answer NA means that the paper does not use existing assets.
- 611 • The authors should cite the original paper that produced the code package or dataset.
- 612 • The authors should state which version of the asset is used and, if possible, include a
613 URL.
- 614 • The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- 615 • For scraped data from a particular source (e.g., website), the copyright and terms of
616 service of that source should be provided.

617 • If assets are released, the license, copyright information, and terms of use in the
618 package should be provided. For popular datasets, paperswithcode.com/datasets
619 has curated licenses for some datasets. Their licensing guide can help determine the
620 license of a dataset.
621 • For existing datasets that are re-packaged, both the original license and the license of
622 the derived asset (if it has changed) should be provided.
623 • If this information is not available online, the authors are encouraged to reach out to
624 the asset's creators.

625 **13. New assets**

626 Question: Are new assets introduced in the paper well documented and is the documentation
627 provided alongside the assets?

628 Answer: **[Yes]**

629 Justification: We include all of the details corresponding to train procedures, datasets used,
630 and citations. Moreover, we provide a readme file for the repository details. The released
631 code is legally approved for the publication; no special documentation is needed.

632 Guidelines:

633 • The answer NA means that the paper does not release new assets.
634 • Researchers should communicate the details of the dataset/code/model as part of their
635 submissions via structured templates. This includes details about training, license,
636 limitations, etc.
637 • The paper should discuss whether and how consent was obtained from people whose
638 asset is used.
639 • At submission time, remember to anonymize your assets (if applicable). You can either
640 create an anonymized URL or include an anonymized zip file.

641 **14. Crowdsourcing and research with human subjects**

642 Question: For crowdsourcing experiments and research with human subjects, does the paper
643 include the full text of instructions given to participants and screenshots, if applicable, as
644 well as details about compensation (if any)?

645 Answer: **[NA]**

646 Justification: No crowdsourcing was used in this study.

647 Guidelines:

648 • The answer NA means that the paper does not involve crowdsourcing nor research with
649 human subjects.
650 • Including this information in the supplemental material is fine, but if the main contribu-
651 tion of the paper involves human subjects, then as much detail as possible should be
652 included in the main paper.
653 • According to the NeurIPS Code of Ethics, workers involved in data collection, curation,
654 or other labor should be paid at least the minimum wage in the country of the data
655 collector.

656 **15. Institutional review board (IRB) approvals or equivalent for research with human
657 subjects**

658 Question: Does the paper describe potential risks incurred by study participants, whether
659 such risks were disclosed to the subjects, and whether Institutional Review Board (IRB)
660 approvals (or an equivalent approval/review based on the requirements of your country or
661 institution) were obtained?

662 Answer: **[NA]**

663 Justification: No human studies/IRB was needed for this study.

664 Guidelines:

665 • The answer NA means that the paper does not involve crowdsourcing nor research with
666 human subjects.

667 • Depending on the country in which research is conducted, IRB approval (or equivalent)
668 may be required for any human subjects research. If you obtained IRB approval, you
669 should clearly state this in the paper.
670 • We recognize that the procedures for this may vary significantly between institutions
671 and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the
672 guidelines for their institution.
673 • For initial submissions, do not include any information that would break anonymity (if
674 applicable), such as the institution conducting the review.

675 **16. Declaration of LLM usage**

676 Question: Does the paper describe the usage of LLMs if it is an important, original, or
677 non-standard component of the core methods in this research? Note that if the LLM is used
678 only for writing, editing, or formatting purposes and does not impact the core methodology,
679 scientific rigorousness, or originality of the research, declaration is not required.

680 Answer: [NA]

681 Justification: This research does not involve LLMs.

682 Guidelines:

683 • The answer NA means that the core method development in this research does not
684 involve LLMs as any important, original, or non-standard components.
685 • Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>)
686 for what should or should not be described.