
Machine-Learned Sampling of Conditioned Path Measures

Qijia Jiang
UC Davis
qjiang@ucdavis.edu

Reuben Cohn-Gordon
UC Berkeley
reubenharry@gmail.com

Abstract

We propose algorithms for sampling from posterior path measures $\mathcal{P}(\mathcal{C}([0, T], \mathbb{R}^d))$ under a general prior process. This leverages ideas from (1) controlled equilibrium dynamics, which gradually transport between two path measures, (2) optimization in ∞ -dimensional probability space endowed with a Wasserstein metric, which can be used to evolve a density curve under the specified likelihood. The resulting algorithms are theoretically grounded and can be integrated seamlessly with neural networks for learning the target trajectory ensembles, without access to data.

1 Introduction

Simulating trajectories through a space, according to a given distribution, is a problem which appears in a broad range of scientific applications. A general formulation of the problem is given by considering a prior path measure that admits a representation as a SDE:

$$dX_t = u_t^{\text{ref}}(X_t)dt + \sqrt{2}dW_t, X_0 \sim \rho_0 \quad \text{Denote this path measure as } (\mathbb{P}_t^{\text{ref}})_{t \leq T}, \quad (1)$$

where the u_t^{ref} part of the drift is assumed known (e.g., with $u_t^{\text{ref}} = -\nabla V_t$, for V_t a double-well potential). Suppose we then observe, at multiple time points t_1, \dots, t_K , potentially noisy observations of the dynamics of the form

$$y_k = X_{t_k} + \sigma z_k, \quad k = 1, \dots, K,$$

which is an additive Gaussian noise model. Given such a prior and data $\{y_k\}_k$, the goal is to infer the posterior path measure q , which is known to take the form [30]:

$$dX_t = [u_t^{\text{ref}}(X_t) + b_t(X_t)]dt + \sqrt{2}dW_t, X_0 \sim \rho_0 \quad \text{Denote this path measure as } (q_t)_{t \leq T}. \quad (2)$$

Equivalently, given methods to sample from $(\mathbb{P}_t^{\text{ref}})_{t \leq T}$, we wish to design algorithms to sample from $(q_t)_{t \leq T}$. Such problems find applications in trajectory inference in biology [14], time series Bayesian analysis, among many others. Notably, this is a generalization of Transition Path Sampling (TPS) from chemistry [8], since one can simply take $t_K = T$ to be a single observation at time $t = T$ with $y_K = B$, $\rho_0 = \delta_A$ and σ small, i.e., we aim to sample from the “bridge” $\mathbb{P}(x_{T-h}, x_{T-2h}, \dots, x_h | x_0 = A, x_T = B)$ in TPS with $J(x; y) \approx \delta(x_T \neq B)$.

The fact that under mild regularity conditions, the posterior path measure (conditioned on the observations $\{y_k\}$) can always be represented with an additional drift b_t^θ [30] means that it can be understood as a “controlled diffusion process”. Thus, the task amounts to inferring the drift of the posterior SDE given a prior, assuming both (1), (2) have the same initial ρ_0 that is perfectly known.

Mathematically, we write the conditional density in path space for $\{x(t)\}_{t \leq T}$ as [3, Section 4.1]

$$Q(x) \propto \rho_0(x_{t_0}) \exp(-I(x) - J(x; y)) \quad (3)$$

where $I(x)$ is the Onsager-Machlup action functional for the unconditional density of path

$$I(x) = \frac{1}{2} \int_0^T \left(\frac{1}{2} \left\| \frac{dX_t}{dt} - u_t^{\text{ref}}(X_t) \right\|_2^2 + \frac{1}{2} \nabla \cdot u_t^{\text{ref}}(X_t) \right) dt$$

and $J(x; y)$ is the likelihood term

$$J(x; y) = \frac{1}{2\sigma^2} \sum_{k=1}^K \|y_k - x_{t_k}\|_2^2 = \frac{1}{2\sigma^2} \int_0^T \sum_{k=1}^K \|y_k - x_t\|_2^2 \cdot \delta(t - t_k) dt =: \int_0^T J_t(x_t) dt. \quad (4)$$

The observations y parameterize the posterior $Q(x)$ we are interested in sampling from, and it is a *global* conditioning. We assume that (1) the observation at time t_k , y_k , depends only on x_{t_k} , and not other time points; (2) $(\mathbb{P}_t^{\text{ref}})_{t \leq T}$ in (1) is a Markov process. Note that the resulting quantity of interest is not a deterministic trajectory but rather a probability measure over path ensembles.

Contributions In this work, we are interested in identifying the optimal b_t^* , which provides a convenient way to draw additional samples / infer the trajectory, rather than just producing a set of sampled trajectories from the target. In particular, we explore novel ways to adapt recent ML ideas from $\mathcal{P}(\mathbb{R}^d)$ sampling to $\mathcal{P}(\mathcal{C}([0, T], \mathbb{R}^d))$ sampling. Our two proposals are presented in Section 2 and 3. Although both methods rely on the posterior nature of the measure rather than just sampling from a generic path measure, they take different viewpoints on the problem. Thus, we expect them to perform well under different likelihood models. Experimental results can be found in Appendix H.

2 Controlled Transport from Prior to Posterior: Tracking b_t

Building on [25] and [1], we introduce a sequence of path measures as

$$\pi_s(x) = Z_s^{-1} \cdot (\rho_0(x_{t_0}) \exp(-I(x) - s \cdot J(x; y))) \quad s \in [0, 1]$$

where $Z_s = \int \rho_0(x_{t_0}) \exp(-I(x) - s \cdot J(x; y)) dx$. It is clear from the definition that π_s interpolate between the prior and the posterior. Similarly, we introduce a family of SDEs parametrized as

$$dX_t^s = b_t^s(X_t^s) dt + \sqrt{2} dW_t, \quad X_0^s \sim \rho_0 \quad \text{and} \quad b_t^0 = u_t^{\text{ref}}$$

inducing corresponding path measures (using the OM functional)

$$\rho_s(x) = Y_s^{-1} \cdot \left[\rho_0(x_{t_0}) \exp \left(-\frac{1}{2} \int_0^T \left(\frac{1}{2} \left\| \frac{dx_t}{dt} - b_t^s(x_t) \right\|_2^2 + \frac{1}{2} \nabla \cdot b_t^s(x_t) \right) dt \right) \right].$$

If we can learn the drift change in b at each step s as the solution of (holding for all path x)

$$\frac{\partial \pi_s(x)}{\partial s} = \pi_s(x) (-J(x; y) + \mathbb{E}_{\pi_s(x)}[J(x; y)]) \stackrel{!}{=} \frac{\partial \rho_s(x)}{\partial s} = \rho_s(x) (h_s(x) - \mathbb{E}_{\rho_s(x)}[h_s(x)]) \quad (5)$$

$$\text{where } h_s(x) = -\frac{1}{2} \int_0^T \left(b_t^s(x_t) - \frac{dx_t}{dt} \right)^\top \frac{\partial b_t^s}{\partial s}(x_t) + \frac{1}{2} \nabla \cdot \frac{\partial b_t^s}{\partial s}(x_t) dt$$

then the resulting drift $\{b_t^s\}_t$ will follow the path measure $\pi_s(x) \forall s \in [0, 1]$ – thus *exactly* hitting target at time $s = 1$. Unlike MCMC (16), this is “controlled”, implying we don’t need to run till $s = \infty$. Operationally we start by drawing samples from the prior $\pi_0(x)$, estimate the expectation in

$$h_s(x) - \mathbb{E}_{\rho_s(x)}[h_s(x)] = -J(x; y) + \mathbb{E}_{\pi_s(x)}[J(x; y)] \quad (6)$$

using empirical samples, and solve for $\partial b^s / \partial s$ on the LHS to update the drift. This is followed by re-simulate paths using this new drift $b^{s+\delta s} \approx b^s + (\partial b^s / \partial s) \cdot \delta s$ for again imposing the consistency equation $\pi_{s+\delta s}(x) = \rho_{s+\delta s}(x)$ in (6) and repeat until $s \approx 1$ since $\pi_1(x) = Q(x)$. Note that estimating (6) does not require knowing the normalizing constant. In practice, one can parametrize $\{\partial b_t^s / \partial s\}_t$ as a NN taking $(x_t^s, t, s) \in \mathbb{R}^d \times [0, T] \times [0, 1]$ as input and outputting a \mathbb{R}^d -valued vector.

To correct for the possible lag $\rho_s \neq \pi_s$, one may optionally simulate the trajectories using the current drift b^s , followed by a few steps of MCMC w.r.t the density π_s , after which we learn the update $\partial b^s / \partial s$ using the latest trajectories, and repeat – this is somewhat similar to Algorithm 2 from [25] where SVGD is used in between for $\mathcal{P}(\mathbb{R}^d)$ sampling. The intuition behind this approach is that it is easier to conceptualize the algorithmic dynamics on the path measure space $q^n \rightarrow q^{n+1} \rightarrow \dots \rightarrow q^*$ and work out the corresponding drift update for the SDE $b^n \rightarrow b^{n+1} \rightarrow \dots \rightarrow b^*$ for implementation purpose. At a high level, this is a transport idea in the path measure setting.

Additional details for this section comparing equilibrium vs. non-equilibrium approaches, and alternative NN-free way to solve for the optimal vector field, can be found in Appendix B.

3 Reformulating Gibbs Variational Principle via Wasserstein Dynamics

There is a variational formulation for the posterior path measure (that can be obtained by generalizing from $\mathcal{P}(\mathbb{R}^d)$) so in our case we can rewrite the problem as

$$\arg \min_{\mathbb{Q}^\theta: \mathbb{Q}_0^\theta = \rho_0} \text{KL}(\mathbb{Q}^\theta \| \mathbb{Q}^*) \Leftrightarrow \arg \min_{\mathbb{Q}^\theta: \mathbb{Q}_0^\theta = \rho_0} \mathbb{E}_{\mathbb{Q}^\theta} [J(x; y)] + \text{KL}(\mathbb{Q}^\theta \| \mathbb{P}^{\text{ref}}) \quad (7)$$

where \mathbb{Q}^* is the target from (3) and the solution is unique under mild assumptions. Equivalently in terms of the drift, using Girsanov's theorem we have

$$\arg \min_{b^\theta} \int_0^T \mathbb{E}_{x_t^\theta \sim q_t^\theta} \left[J_t(x_t^\theta; y) + \frac{1}{4} \|b_t^\theta(x_t^\theta)\|^2 \right] dt. \quad (8)$$

While the loss (8) can in principle be estimated via empirical samples and used to optimize over b^θ , there are a few challenges: (1) the loss is complicated as a function of the drift b^θ (even though in terms of \mathbb{Q}^θ in (7) it's convex); (2) one needs to calculate $\partial x_t^\theta / \partial \theta$ generated with SDE (2) for updating b^θ , which requires back-propagating through the solver. This approach is related to some of the Neural SDE methods for fitting SDEs to data [21, 26].

Taking a different route based on convex optimization, we first leverage the Benamou-Brenier [4] reformulation to pose (7) as joint learning over $(q_t, b_t)_t$, which yields insights that will allow us to simplify further later. For our specific case, it becomes

$$\min_{q_t, b_t} \int_0^T \int_{\mathbb{R}^d} \left[\frac{1}{4} \|b_t(x_t)\|^2 + J_t(x_t; y) \right] q_t(x_t) dx dt \quad \text{s.t.} \quad \partial_t q_t = -\nabla \cdot (q_t(b_t + u_t^{\text{ref}})) + \Delta q_t, q_0 = \rho_0 \quad (9)$$

Now if one follows the augmented Lagrangian approach in [4] that via reparameterization turns the problem into a convex optimization, the Laplacian Δq_t from the diffusion in Fokker-Planck introduces a challenge as we have to deal with a Bilaplacian type of quantity in one of the updates.

3.1 Approximation

For (9), there is a way to move the Laplacian term into the objective that is used in various works (e.g., Eqn (4.25) in [13]). Followed by the similar $(q_t, b_t) \mapsto (q_t, m_t)$ for $m_t = q_t \cdot b_t$ transformation as done in [4], we end up with our *jointly convex* objective with an extra (relative) Fisher information in the objective but a continuity equation as the constraint.

Proposition 1 (Convex Reformulation). *For a general reference $u_t^{\text{ref}} \neq 0$, (9) is equivalent to*

$$\begin{aligned} \min_{q_t, m_t} \int_0^T \int_{\mathbb{R}^d} \left[\frac{1}{4} \left\| \frac{m_t}{q_t}(x_t) - (u_t^{\text{ref}} - \nabla \log \mathbb{P}_t^{\text{ref}})(x_t) \right\|^2 + \frac{1}{4} \left\| \nabla \log \frac{q_t}{\mathbb{P}_t^{\text{ref}}}(x_t) \right\|^2 + J_t(x_t) \right] q_t(x_t) dx dt \\ + \frac{1}{2} \int_{\mathbb{R}^d} q_T(x_T) \log \frac{q_T}{\mathbb{P}_T^{\text{ref}}}(x_T) dx_T \\ \text{s.t.} \quad \partial_t q_t = -\nabla \cdot m_t, q_0 = \rho_0. \end{aligned} \quad (10)$$

When the reference process \mathbb{P}^{ref} is in equilibrium (i.e., $\mathbb{P}_t^{\text{ref}} = \nu \propto e^{-V} \forall t$) the problem reduces to

$$\begin{aligned} \min_{q_t, m_t} \int_0^T \int_{\mathbb{R}^d} \left[\frac{1}{4} \left\| \frac{m_t}{q_t}(x_t) \right\|^2 + \frac{1}{4} \left\| \nabla \log \frac{q_t}{\nu}(x_t) \right\|^2 + J_t(x_t) \right] q_t(x_t) dx dt + \frac{1}{2} \int_{\mathbb{R}^d} q_T(x_T) \log \frac{q_T}{\nu}(x_T) dx_T \\ \text{s.t.} \quad \partial_t q_t = -\nabla \cdot m_t, q_0 = \rho_0 \end{aligned} \quad (11)$$

In this case, $\rho_0 = \nu = \mathbb{P}_0^{\text{ref}}$ for the initial distribution.

Reformulation in terms of marginal density: Compared to the dynamic formulation of \mathcal{W}_2 optimal transport from [4], in Proposition 1 we have an extra (relative) Fisher information and potential energy term involving J , and removal of the terminal ρ_T constraint for a penalty on q_T . In what follows, we adopt the following approximation, which crucially only involves the evolution of a *single* variable, instead of two. Since the objective (11) is (jointly) convex and (almost) separable over

time $t \in [0, T]$ we propose to solve a sequence of ‘‘Wasserstein Gradient Flow’’/JKO steps between adjacent marginals as (initialized at $q_0 = \rho_0$)

$$q_{t+h} \leftarrow \arg \min_{q \in \mathcal{P}(\mathbb{R}^d)} \frac{1}{2h^2} \mathcal{W}_2^2(q, q_t) + F(q) \quad (12)$$

for advancing the solution over $t \in [0, T]$ with stepsize h , where $F(q)$ is the convex relative Fisher information functional $I(q; \nu) := \frac{1}{2} \int \|\nabla \log q(x)/\nu(x)\|^2 q(x) dx +$ term linear in $q : \frac{2}{h} \int J_{t+h}(x)q(x) dx$. To see one can work with $\{q_t\}_t$ only, observe that in (11), the drift term v_t (or equivalently m_t) only appears in the first part, and the rest is a function of the density q_t only. Therefore given a curve of optimal measure $\{q_t\}_{t \leq T}$ one can solve classical OT between adjacent pairs to get the corresponding vector fields $\{v_t\}_t$ that follow the trajectory with *minimum* kinetic energy (hence divergence-free). This is an Action Matching problem [23] if provided with samples from the $\{q_t\}_{t \leq T}$ curve. Conceptually, this is a dynamical system perspective on the problem as a sequence of transformations $q_0 \xrightarrow{b_h} q_h \rightarrow \dots \rightarrow q_{T-h} \xrightarrow{b_T} q_T$, and we study the joint distribution over $\mathcal{P}(\mathbb{R}^d \times \mathbb{R}^d)$ at a time, as the Markovian representation (2) suggests. This reduces the problem from a joint $T \times d$ distribution to a sequence of $2d$ problems. Unlike the more global method in Section 2 that repeatedly updates the path measure $\mathcal{P}(\mathcal{C}([0, T], \mathbb{R}^d))$, it is a one-pass method evolving $\mathcal{P}(\mathbb{R}^d)$ from $t = 0$ to $t = T$, local in nature. Alternative approximations on (11) and more justification for our approach based on (12) are discussed in Section D.0.2, along with an Eulerian solver for low-dimensional case based on (10) in Section C.1.

3.2 High Dimensional Case: Lagrangian Methods for (11)

In the high-dimensional regime, we will work with the simpler problem of (11) that conditions an equilibrium process with observations, and advocate for parameterizing q_t as either an interacting particle system, or as a parametrized push-forward map [9] that allows us to infer b_t easily. In both cases, we leverage the perspective put forth in (12) so optimization is only done between adjacent marginals to track its evolution, and not joint learning over q_t, b_t . We will heavily rely on the representation of Fisher information as

$$\begin{aligned} R(q_t) &:= \frac{1}{2} \int \|\nabla \log q_t(x)\|^2 q_t(x) dx = 2 \int \|\nabla \sqrt{q_t(x)}\|^2 dx = \frac{1}{2} \int \frac{\|\nabla q_t(x)\|^2}{q_t(x)} dx \\ &= \sup_{\frac{1}{2}\|\beta(x)\|^2 + \alpha(x) \leq 0} \int \nabla q_t(x)^\top \beta(x) + q_t(x) \alpha(x) dx \quad (13) \end{aligned}$$

$$= \sup_{\frac{1}{2}\|\beta(x)\|^2 + \alpha(x) \leq 0} \int (\alpha(x) - \nabla \cdot \beta(x)) \cdot q_t(x) dx \quad (14)$$

and its approximation with empirical samples in Lemma 4. Proofs for this subsection are in App E.

3.2.1 Sequence of Pushforward Maps

The first Lagrangian approach we consider is based on Brenier’s Theorem for \mathcal{W}_2 -OT, which says that the optimal transport map is given by the gradient of a convex function that solves a Monge-Ampère equation. To leverage this push-forward map representation [9], we rewrite (12) as

$$\phi_t \leftarrow \arg \min_{\phi \text{ convex}} \frac{1}{2h^2} \int \|\nabla \phi(x) - x\|^2 \cdot q_t(x) dx + F(\nabla \phi_{\#} q_t); \quad q_{t+h} \leftarrow \nabla \phi_{\#} q_t \quad (15)$$

and propose to use an Input-Convex NN architecture [2] to optimize over a family of convex functions ϕ^θ that takes $t, \{x_t^i\}_{i=1}^n \sim q_t$ and output $\{x_{t+h}^i\}_{i=1}^n \sim q_{t+h}$. The gradient $\nabla \phi^\theta$ gives us the desired transport map. This allows us to evolve an empirical measure \hat{q} in time and space.

In order to evaluate $F(\nabla \phi_{\#} q_t)$ for optimization, the ‘‘internal energy’’ term $\int \|\nabla \log q_t(x_t)\|^2 dq_t(x_t)$ that is nonlinear in q_t seems challenging to estimate with empirical measure \hat{q}_t , and using the change of variables formula involving determinant along the dynamics is computationally intensive. However, here one can linearize via bi-level optimization using (13) as we illustrate in Algorithm 3 (the rest of the terms in F are easily linearizable with possible integration-by-parts). Lemma 4 gives a 0th-order method for approximating the ‘‘gradient’’, and luckily here we only need to draw samples from q_t for implementation, *without* evaluating its density. The particular form (14) is a special feature of Fisher Information – other nonlinear energies may be difficult to re-parameterize to be linearizable.

A second Lagrangian approach based on Interacting Particle System is given in Section D.0.1.

References

- [1] Michael S Albergo and Eric Vanden-Eijnden. NETS: A Non-Equilibrium Transport Sampler. *arXiv preprint arXiv:2410.02711*, 2024.
- [2] Brandon Amos, Lei Xu, and J Zico Kolter. Input convex neural networks. In *International conference on machine learning*, pages 146–155. PMLR, 2017.
- [3] Amit Apte, Martin Hairer, AM Stuart, and Jochen Voss. Sampling the posterior: An approach to non-Gaussian data assimilation. *Physica D: Nonlinear Phenomena*, 230(1-2):50–64, 2007.
- [4] Jean-David Benamou and Yann Brenier. A computational fluid mechanics solution to the Monge-Kantorovich mass transfer problem. *Numerische Mathematik*, 84(3):375–393, 2000.
- [5] Jean-David Benamou, Guillaume Carlier, Simone Di Marino, and Luca Nenna. An entropy minimization approach to second-order variational mean-field games. *Mathematical Models and Methods in Applied Sciences*, 29(08):1553–1583, 2019.
- [6] Jean-David Benamou, Guillaume Carlier, and Maxime Laborde. An augmented Lagrangian approach to Wasserstein gradient flows and applications. *ESAIM: Proceedings and surveys*, 54:1–17, 2016.
- [7] Alexandros Beskos, Gareth Roberts, Andrew Stuart, and Jochen Voss. MCMC methods for diffusion bridges. *Stochastics and Dynamics*, 8(03):319–350, 2008.
- [8] Peter G Bolhuis and David WH Swenson. Transition path sampling as Markov chain Monte Carlo of trajectories: Recent algorithms, software, applications, and future outlook. *Advanced Theory and Simulations*, 4(4):2000237, 2021.
- [9] Charlotte Bunne, Laetitia Papaxanthos, Andreas Krause, and Marco Cuturi. Proximal optimal transport modeling of population dynamics. In *International Conference on Artificial Intelligence and Statistics*, pages 6511–6528. PMLR, 2022.
- [10] Davide Carbone, Mengjian Hua, Simon Coste, and Eric Vanden-Eijnden. Efficient training of energy-based models using jarzynski equality. *Advances in Neural Information Processing Systems*, 36:52583–52614, 2023.
- [11] Guillaume Carlier, Jean-David Benamou, and Daniel Matthes. Wasserstein gradient flow of the Fisher information from a non-smooth convex minimization viewpoint. *Journal of Convex Analysis*, 2024.
- [12] José Antonio Carrillo, Yanghong Huang, Francesco Saverio Patacchini, and Gershon Wolansky. Numerical study of a particle method for gradient flows. *arXiv preprint arXiv:1512.03029*, 2015.
- [13] Yongxin Chen, Tryphon T Georgiou, and Michele Pavon. Stochastic control liaisons: Richard sinkhorn meets gaspard monge on a schrodinger bridge. *Siam Review*, 63(2):249–313, 2021.
- [14] Lénaïc Chizat, Stephen Zhang, Matthieu Heitz, and Geoffrey Schiebinger. Trajectory inference via mean-field langevin in path space. *Advances in Neural Information Processing Systems*, 35:16731–16742, 2022.
- [15] Giovanni Conforti and Luca Tamanini. A formula for the time derivative of the entropic cost and applications. *Journal of Functional Analysis*, 280(11):108964, 2021.
- [16] Yuanqi Du, Michael Plainer, Rob Brekelmans, Chenru Duan, Frank Noe, Carla P Gomes, Alan Apsuru-Guzik, and Kirill Neklyudov. Doob’s Lagrangian: A Sample-Efficient Variational Approach to Transition Path Sampling. *arXiv preprint arXiv:2410.07974*, 2024.
- [17] Ivan Gentil, Christian Léonard, and Luigia Ripani. About the analogy between optimal transport and minimal entropy. *Annales de la Faculté des sciences de Toulouse: Mathématiques*, 26(3):569–600, 2017.

- [18] Ugo Gianazza, Giuseppe Savaré, and Giuseppe Toscani. The Wasserstein gradient flow of the Fisher information and the quantum drift-diffusion equation. *Archive for rational mechanics and analysis*, 194(1):133–220, 2009.
- [19] Jeremy Heng, Valentin De Bortoli, Arnaud Doucet, and James Thornton. Simulating diffusion bridges with score matching. *arXiv preprint arXiv:2111.07243*, 2021.
- [20] Lars Holdijk, Yuanqi Du, Ferry Hooft, Priyank Jaini, Berend Ensing, and Max Welling. Stochastic optimal control for collective variable free sampling of molecular transition paths. *Advances in Neural Information Processing Systems*, 36, 2024.
- [21] Xuechen Li, Ting-Kam Leonard Wong, Ricky TQ Chen, and David Duvenaud. Scalable gradients for stochastic differential equations. In *International Conference on Artificial Intelligence and Statistics*, pages 3870–3882. PMLR, 2020.
- [22] Qiang Liu. Stein variational gradient descent as gradient flow. *Advances in neural information processing systems*, 30, 2017.
- [23] Kirill Neklyudov, Rob Brekelmans, Daniel Severo, and Alireza Makhzani. Action matching: Learning stochastic dynamics from samples. In *International conference on machine learning*, pages 25858–25889. PMLR, 2023.
- [24] Kirill Neklyudov, Rob Brekelmans, Alexander Tong, Lazar Atanackovic, Qiang Liu, and Alireza Makhzani. A computational framework for solving Wasserstein Lagrangian flows. *arXiv preprint arXiv:2310.10649*, 2023.
- [25] Nikolas Nüsken. Stein transport for Bayesian inference. *arXiv preprint arXiv:2409.01464*, 2024.
- [26] Manfred Opper. Variational inference for stochastic differential equations. *Annalen der Physik*, 531(3):1800233, 2019.
- [27] Stefano Peluchetti. Diffusion bridge mixture transports, Schrödinger bridge problems and generative modeling. *Journal of Machine Learning Research*, 24(374):1–51, 2023.
- [28] Gabriel Peyré. Entropic approximation of Wasserstein gradient flows. *SIAM Journal on Imaging Sciences*, 8(4):2323–2351, 2015.
- [29] Jakiw Pidstrigach, Elizabeth Baker, Carles Domingo-Enrich, George Deligiannidis, and Nikolas Nüsken. Conditioning Diffusions Using Malliavin Calculus. *arXiv preprint arXiv:2504.03461*, 2025.
- [30] Maxim Raginsky. A variational approach to sampling in diffusion processes. *arXiv preprint arXiv:2405.00126*, 2024.
- [31] Gabriel Stoltz. Path sampling with stochastic dynamics: Some new algorithms. *Journal of Computational Physics*, 225(1):491–508, 2007.
- [32] Gabriel Stoltz, Mathias Rousset, et al. *Free energy computations: A mathematical perspective*. World Scientific, 2010.
- [33] Andrew M Stuart, Jochen Voss, and Petter Wilberg. Conditional path sampling of SDEs and the Langevin MCMC method. *Communications in Mathematical Sciences*, 2004.
- [34] Tobias Sutter, Arnab Ganguly, and Heinz Koepl. A variational approach to path estimation and parameter inference of hidden diffusion processes. *Journal of Machine Learning Research*, 17(190):1–37, 2016.
- [35] Alexander Tong, Kilian Fatras, Nikolay Malkin, Guillaume Hugué, Yanlei Zhang, Jarrid Rector-Brooks, Guy Wolf, and Yoshua Bengio. Improving and generalizing flow-based generative models with minibatch optimal transport. *arXiv preprint arXiv:2302.00482*, 2023.
- [36] Luke Triplett and Jianfeng Lu. Diffusion methods for generating transition paths. *Journal of Computational Physics*, 522:113590, 2025.

A Related Work

With regards to the infinite-dimensional path-space sampling problem on $\mathcal{C}([0, T], \mathbb{R}^d)$, there exists a longstanding body of research concerning at least two different types of methods.

Existing Method 1: Langevin Stochastic PDE With the target $Q(x)$ defined as above, one can write down a Langevin SPDE on $x(t, s)$:

$$\frac{\partial x(t, s)}{\partial s} = \delta_x \log Q(x) + \sqrt{2} \frac{\partial W}{\partial s} \quad \text{for } (s, t) \in (0, \infty) \times (0, T) \quad (16)$$

where s is the algorithmic time (we expect as $s \rightarrow \infty$ relying on ergodicity of the Markov chain it should converge to the desired equilibrium $Q(x)$) and $t \in [0, T]$ is the physical time that simply index the path. The variational derivative $\delta_x \log Q(x)$ can be worked out analytically (as done in [3]). Together with appropriate boundary conditions, this can serve as a sampler for the problem (assuming observation at endpoint T):

$$\begin{aligned} \delta_x \log Q(x) &= \frac{1}{2} \frac{d^2 x_t}{dt^2} - \frac{1}{2} u(x) \nabla u(x) - \frac{1}{2} \nabla (\nabla \cdot u(x)) + \sum_{k=1}^K \frac{1}{\sigma^2} (y_k - x(t_k, \cdot)) \\ \frac{\partial x(0, s)}{\partial t} - u(x(0, s)) + 2 \nabla \log \rho_0(x(0, s)) &= 0, \quad t = 0 \\ \frac{\partial x(T, s)}{\partial t} - u(x(T, s)) - \frac{2}{\sigma^2} [y_K - x(T, s)] &= 0, \quad t = T \end{aligned} \quad (17)$$

There have been attempts to pursue different variants of MCMC methods from Stuart et al. [33, 7] (possibly with MH adjustment), with a connection to data assimilation generally.

Existing Method 2: learning $b_t(X_t)$ Rather than working directly with paths, which can be memory-intensive, another approach is to directly learn the additional drift $b_t(X_t)$ in (2), which in fact takes a gradient form. This often reduces to an optimal control problem or alternatively a high-dimensional PDE problem. It is known, since the posterior is a reweighting of the prior path measure, the drift admits closed-form in a path integral representation [30, Theorem 1]:

$$b_t(z) = 2 \nabla \log \mathbb{E}_{x \sim \mathbb{P}^{\text{ref}}} \left[\exp \left(- \int_t^T J_s(x_s) ds \right) \middle| x_t = z \right] =: -\nabla v_t(z) \quad (18)$$

for the conditioned SDE, where the expectation is taken over the prior SDE with drift u_t^{ref} . This is something one may hope to simulate to estimate and recovers Doob's h -transform $2 \nabla_x \log p_{T|t}(X_T = B | X_t = x)$ when we only observe at time T [29]. But since the quantity under estimation is typically a rare-event under the reference measure, even with importance weights, it is generally hard to do well; many paths will have negligible weights and the relevant paths are exponentially rare. There is a PDE one can write down for v_t in (18) as well using Feynman-Kac:

$$\frac{\partial}{\partial t} v_t(z) + u_t^{\text{ref}}(z)^\top \nabla v_t(z) + \Delta v_t(z) + J_t(z) = \|\nabla v_t(z)\|^2 \quad (19)$$

but it isn't clear that solving a high-dimensional PDE would be any easier. Some other related work along these lines (with possibly marginal variational approximation) include [20, 16, 34]. A method based on score-matching for learning this drift has also been proposed [19].

In the case of TPS, this is a bona-fide Schrödinger bridge (SB) problem with two delta terminal constraints, for which one can do path-space Sinkhorn / IPF by iteratively minimizing $\text{KL}(\overrightarrow{\mathbb{P}}^{\delta_A, u} \parallel \overleftarrow{\mathbb{P}}^{\delta_B, v})$ over the two slots alternatively. It amounts to setting $u_0 = u^{\text{ref}}$, and perform a sequence of score-matching time reversals to learn the drifts u^*, v^* . However, this is not as general as our setup, and methodology from [27] that is also based on an alternating projection for SB is difficult to generalize when $u^{\text{ref}} \neq 0$ since simulating from the conditional bridge would be challenging. We note that compared to the series of work on flow matching [35], which aim to match the time marginals for a prescribed path (often set up as mixtures of diffusion processes), we aim for the more ambitious goal of sampling from the true target path measure.

Crucially, our setup does not assume access to data, which distinguishes it from the proposals in [36] that leverage diffusion model for conditional sampling from a function-valued probability density over $\{X_t\}_t$.




Figure 1: On a potential landscape $V(x) = 5(x_1^2 - 1)^2 + 5(x_2^2 - 1)^2$, we condition on the final point of the path being in the top right potential well, and the midpoint being in the bottom right well or top left well. We use Algorithm 1 from Section 2, varying s from 0.0 (fully black path) to 1.0 (fully white path) in increments of 0.1.

B Additional Details for Section 2

We begin with a remark on the choice of the annealing path, and comment on normalizing constant estimator.

Remark 1. In the special case of TPS with $X_0 = A, X_T = B$, there's an alternative path one can anneal against: The measure of interest Q has a density w.r.t the Brownian bridge with the likelihood ratio being

$$J(x) \propto \frac{1}{2} \int_0^T \left(\frac{1}{2} \|\nabla V(X_t)\|^2 - \Delta V(X_t) \right) dt =: \int_0^T J_t(X_t) dt \quad (20)$$

up to a normalizing constant. Assuming one can sample from the Brownian bridge easily, this gives a way to interpolate between the Brownian bridge

$$dX_t = \frac{B - X_t}{T - t} dt + \sqrt{2} dW_t, \quad X_0 = A \quad (21)$$

with

$$\pi_0(x) \propto \exp \left(-\frac{1}{4} \int_0^T \left\| \frac{dX_t}{dt} \right\|^2 dt \right) \delta_A(X_0) \delta_B(X_T) \quad (22)$$

and a more general bridge with $-\nabla V_t$ as the potential. One may again intersperse MCMC with the transport step. With this annealing, we gradually introduce the reference potential instead of the constraint, i.e., we reweight the prior paths not based on $X_T = B$, but the potential V instead. To see that (20)-(22) give the same path probability as the OM functional (3), we simply note that the extra factor of

$$\int_0^T \frac{1}{2} \frac{dX_t}{dt} \top u_t^{\text{ref}}(X_t) dt = - \int_0^T \frac{1}{2} \frac{dX_t}{dt} \top \nabla V(X_t) dt = -\frac{1}{2} \int_0^T \frac{dV(X_t)}{dt} dt = \frac{1}{2} (V(X_0) - V(X_T))$$

is equal to $\frac{1}{2} (V(A) - V(B))$ (i.e., a constant fixed by the boundary condition).

Path-space Normalizing Constant Estimator With the $\rho_s = \pi_s$ interpolation method in Section 2 (i.e., maintaining equilibrium dynamics w.r.t $\pi_s \forall s \in [0, 1]$), one can consider thermodynamic integration in path space and calculate the partition function for Q as

$$\begin{aligned} \frac{\partial Z_s}{\partial s} &= \int -J(x; y) \rho_0(x_{t_0}) \exp(-I(x) - s \cdot J(x; y)) dx \\ &= \frac{\int -J(x; y) \rho_0(x_{t_0}) \exp(-I(x) - s \cdot J(x; y)) dx}{\int \rho_0(x_{t_0}) \exp(-I(x) - s \cdot J(x; y)) dx} Z_s = \mathbb{E}_{\pi_s}[-J(x; y)] \cdot Z_s \end{aligned}$$

therefore $Z_T = Z_0 \cdot \exp \left(\int_0^T \mathbb{E}_{\pi_s}[-J(x; y)] ds \right)$ assuming the initial Z_0 is analytically tractable. Quantity of such type is often of interest in chemistry applications. The end-to-end method is summarized in Algorithm 1 (c.f. Appendix G).

We follow with additional discussion on the controlled path measure methodology put forth in Section 2.

B.1 Kernel Method

Without NN, we can solve (6) with Kernel Ridge Regression (KRR) by assuming b belongs to a RKHS \mathcal{H}_k . The benefit of this is that everything now becomes closed-form. For each s , denote $v^s(x_t, t) := \frac{\partial b^s}{\partial s}(x_t, t)$ as a vector field mapping $\mathbb{R}^d \times \mathbb{R}$ to \mathbb{R}^d , where we have an extra $t \in [T]$ dimension compared to [25] because we are working in path-space. Introduce the linear operator \mathcal{S} mapping the function v^s and path $x^i \in \mathcal{C}([0, T], \mathbb{R}^d)$ to a scalar $h_s(x^i) - \frac{1}{N} \sum_{j=1}^N h_s(x^j)$. Assuming N trajectories, at each $s \in [0, 1]$, one solves

$$\arg \min_{v \in \mathcal{H}_k} \frac{1}{N} \sum_{i=1}^N (\mathcal{S}(v^s)(x^i) - j(x^i))^2 + \lambda \|v^s\|_{\mathcal{H}_k}^2 \quad (23)$$

by letting $j(x^i) := -J(x^i) + \frac{1}{N} \sum_{j=1}^N J(x^j)$ that takes a path $x^i \in \mathcal{C}([0, T], \mathbb{R}^d)$ as input and outputs a scalar.

Lemma 1 (Kernel Method). *KRR method for solving (6) with empirical samples, written in (23) for each $s \in [0, 1]$, can be implemented analytically by solving a linear system of size $N \times N$ at each $t \in [T]$.*

Remark 2. The way we wrote it, it may seem that the T problems to solve are independent from each other, but they are in fact coupled through terms like dx_t^i/dt in the proof.

B.2 Equilibrium vs. Non-Equilibrium Thermodynamics

Our methods from Section 2, B.1, and the alternative ‘‘path-based drifts’’ discussed in Appendix B.4 below, being based on having the dynamics exactly following the interpolation curve, are equilibrium methods. Alternatively, one can use the annealed gradient $\delta_x \log \pi_s(x)$ from the SPDE, together with MH adjustment, as a time-inhomogeneous transition kernel, followed by *weighting* on the samples to track the π_s curve (Algorithm 2). Since the MCMC dynamics alone does not impose $\rho_s = \pi_s$ exactly (a single step of MCMC with changing target will result in a delay), these fall in the category of non-equilibrium methods. The result below bears a resemblance to that in [31, Section 4], and is a generalization of Jarzynski’s equality to path space.

Lemma 2 (Time-discrete Jarzynski). *The analogue of the Jarzynski equality in discrete time in path space is: for $s \in [0, 1]$ along the Algorithm 2 dynamics*

$$\frac{Z_s}{Z_0} = \mathbb{E}_{(X_t \sim \pi_t(x))_{t \leq s}} \left[\exp \left(- \int_0^s J(X_t) dt \right) \right] \approx \frac{1}{M} \sum_{i=1}^M \underbrace{\exp \left(- \sum_{t=0}^{s/\delta_s} \delta_s \cdot J(X_t^i) \right)}_{\exp(-W^{i,s/\delta_s})}$$

where we obtain path X_t from $X_{t-\delta_s}$ using a numerical discretization of a SPDE dynamics leaving $\pi_t(x)$ invariant through MH adjustment, with $X_0 \sim \pi_0 = \mathbb{P}^{ref}$ as initialization. Setting $s = 1$, the average of a statistics $h(\cdot)$ over the path w.r.t the target $Q(x)$ can be computed as

$$\frac{\sum_{j=1}^M h(X_n^j) e^{-W^{j,n}(X^j)}}{\sum_{j=1}^M e^{-W^{j,n}(X^j)}}$$

on the final set of M trajectories $\{X_n^j\}_{j=1}^M$, with $n = 1/\delta_s$ as the number of switching steps.

In practice, we also have to discretize the X_t path along the $[0, T]$ physical time, as shown in Algorithm 2 in detail. Some drawbacks of this non-equilibrium approach: (1) simulating SPDE with MH in general is computationally intensive; (2) the weights may have high variance depending on how fast π_s changes (compared to equilibrium methods, where each path has weight ≈ 1 and the switching can be faster).

We additionally give a discrete-time Jarzynski weighting without Metropolis adjustment (i.e., the Markov kernel at each step doesn’t exactly satisfy detailed balance) in Lemma 3. This is closer to an Annealed Importance Sampling (AIS) derivation, and corresponds to a general definition of a work functional. Moreover, we note that in contrast to Algorithm 2 based on non-equilibrium dynamics, Algorithm 1 is amenable to importance sampling for correcting a potentially suboptimal drift $\hat{b}^1 \neq b^*$ since we always have the current path measure in analytical form, whereas the marginal density along the SPDE dynamics is generally intractable.

B.3 Missing Derivation

We give the two missing proofs from Section B.1 and B.2 below.

Lemma 1 (Kernel Method). *KRR method for solving (6) with empirical samples, written in (23) for each $s \in [0, 1]$, can be implemented analytically by solving a linear system of size $N \times N$ at each $t \in [T]$.*

Proof. Fix a particular s . Using the fact that the operator \mathcal{S} is linear in v if we view it as a function $v_t(x)$ indexed by t , we derive the solution to KRR for each t , following section A.1.2 from [25]. The adjoint satisfies

$$\langle c, \mathcal{S}_t(v_t) \rangle_N = \langle \mathcal{S}_t^* c, v_t \rangle_{\mathcal{H}_k} \quad \forall v_t \in \mathcal{H}_k^d, c \in \mathbb{R}^N$$

implying for $c \in \mathbb{R}^N$, using the definition of h_s ,

$$\mathcal{S}_t^* c = \frac{1}{N} \sum_{i=1}^N \left(\frac{1}{2} \nabla_{x_t^i} k(\cdot, x_t^i) - \frac{1}{2} k(\cdot, x_t^i) \left(b_t(x_t^i) - \frac{dx_t^i}{dt} \right) \right) c_i.$$

for any $c \in \mathbb{R}^N$. Therefore the solution (parametrized by N, λ) is given by

$$v_t^* = \mathcal{S}_t^* (\lambda I_{N \times N} + \mathcal{S}_t \mathcal{S}_t^*)^{-1} j_t =: \mathcal{S}_t^* c, \quad (24)$$

which takes $\{x_t^i\}_{i=1}^N \in \mathbb{R}^d$ as input and output a \mathbb{R}^d -valued vector. Above $j_t[i] = -J_t(x_t^i) + \frac{1}{N} \sum_{j=1}^N J_t(x_t^j)$, $i \in [N]$. Here the linear system to be solved in (24) involving $\mathcal{S}_t \mathcal{S}_t^*$ that maps $\mathbb{R}^d \times \mathbb{R}^d$ to \mathbb{R} can be evaluated to be

$$\begin{aligned} (\mathcal{S}_t \mathcal{S}_t^*)_{ij} &= -\frac{1}{4} \left(b_t(x_t^i) - \frac{dx_t^i}{dt} \right)^\top \nabla_{x_t^j} k(x_t^i, x_t^j) + \frac{1}{4} \left(b_t(x_t^i) - \frac{dx_t^i}{dt} \right)^\top k(x_t^i, x_t^j) \left(b_t(x_t^j) - \frac{dx_t^j}{dt} \right) \\ &\quad + \frac{1}{4} \nabla_{x_t^i} \cdot \nabla_{x_t^j} k(x_t^i, x_t^j) - \frac{1}{4} \nabla_{x_t^i} k(x_t^i, x_t^j)^\top \left(b_t(x_t^j) - \frac{dx_t^j}{dt} \right), \quad i \in [N], j \in [N] \end{aligned}$$

for the corresponding kernel $k(\cdot, \cdot)$. (24) is a linear system of size $N \times N$ given N empirical samples/trajectories $\{x_t^i\}_{i=1}^N$ at each $t \in [T]$. The resulting $(v_t^*)_t$ from above can be used to update the drift $(b_t)_t$ and generate new trajectories for solving (23) again at a new s , with possibly MCMC steps in between. \square

Lemma 2 (Time-discrete Jarzynski). *The analogue of the Jarzynski equality in discrete time in path space is: for $s \in [0, 1]$ along the Algorithm 2 dynamics*

$$\frac{Z_s}{Z_0} = \mathbb{E}_{(X_t \sim \pi_t(x))_{t \leq s}} \left[\exp \left(- \int_0^s J(X_t) dt \right) \right] \approx \frac{1}{M} \sum_{i=1}^M \underbrace{\exp \left(- \sum_{t=0}^{s/\delta_s} \delta_s \cdot J(X_t^i) \right)}_{\exp(-W^{i, s/\delta_s})}$$

where we obtain path X_t from $X_{t-\delta_s}$ using a numerical discretization of a SPDE dynamics leaving $\pi_t(x)$ invariant through MH adjustment, with $X_0 \sim \pi_0 = \mathbb{P}^{ref}$ as initialization. Setting $s = 1$, the average of a statistics $h(\cdot)$ over the path w.r.t the target $Q(x)$ can be computed as

$$\frac{\sum_{j=1}^M h(X_n^j) e^{-W^{j, n}(X^j)}}{\sum_{j=1}^M e^{-W^{j, n}(X^j)}}$$

on the final set of M trajectories $\{X_n^j\}_{j=1}^M$, with $n = 1/\delta_s$ as the number of switching steps.

Proof. Fix a switching with $n := s/\delta_s$ steps, the proof follows closely the derivation in [32, Remark 4.5]. The MH adjustment on top of the SPDE dynamics makes sure the estimator is unbiased, i.e.,

$$\frac{Z_s}{Z_0} = \mathbb{E} \left[\exp \left(- \sum_{t=0}^{s/\delta_s} \delta_s \cdot J(X_t) \right) \right]$$

even with the δ_s algorithmic time discretization. To see this, note the probability to observe a sequence of path X_0, X_1, \dots, X_n under the dynamics is

$$\Pi(dX) = \pi_0(dX_0) \prod_{i=1}^n P_{\pi_i}(X_{i-1}, dX_i)$$

where the transition probability $P_{\pi_i}(X_{i-1}, dX_i)$ leaves the path measure $\pi_i(x) \propto \exp(-I(x) - i \cdot \delta_s J(x)) =: \exp(-V_i(x))$ invariant. Therefore

$$\begin{aligned} \mathbb{E}_{\{X_i\}_{i=0}^n} \left[\exp \left(- \sum_{i=0}^n \delta_s \cdot J(X_i) \right) \right] &= \int \prod_{i=0}^{n-1} \exp(-V_{i+1}(X_i) + V_i(X_i)) \Pi(dX) \\ &= \frac{1}{Z_0} \int \exp(-V_1(X_0)) \prod_{i=1}^{n-1} \exp(-V_{i+1}(X_i) + V_i(X_i)) \prod_{i=1}^n P_{\pi_i}(X_{i-1}, dX_i) dX_0 \\ &= \frac{Z_1}{Z_0} \int \prod_{i=1}^{n-1} \exp(-V_{i+1}(X_i) + V_i(X_i)) \pi_1(dX_1) \prod_{i=2}^n P_{\pi_i}(X_{i-1}, dX_i) \end{aligned}$$

since by the invariance of kernel P_{π_i} on π_i :

$$\int \exp(-V_i(X_{i-1})) P_{\pi_i}(X_{i-1}, dX_i) dX_{i-1} = Z_i \pi_i(dX_i).$$

Inductively using the same argument gives

$$\mathbb{E}_{\{X_i\}_{i=0}^n} \left[\underbrace{\exp \left(- \sum_{i=0}^n \delta_s \cdot J(X_i) \right)}_{e^{-W^n}} \right] = \frac{Z_n}{Z_0},$$

for which we approximate by averaging M empirical samples in practice. To report the average of a statistics h over the path, using importance sampling we have

$$\frac{1}{Z_n} \int h(x) \pi_n(x) dx = \frac{\int h(X) \frac{\pi_n(X)}{\rho_n(X)} d\rho_n(X)}{\int \frac{\pi_n(X)}{\rho_n(X)} d\rho_n(X)} \approx \frac{\sum_{j=1}^M h(X_n^j) e^{-W^{j,n}(X^j)}}{\sum_{j=1}^M e^{-W^{j,n}(X^j)}}$$

with M independent walkers / trajectories. At final step n of the algorithm, each trajectory $j \in [M]$ has individual weights as

$$w_j = e^{-W^{j,n}}$$

where W denotes the work along the MH-adjusted SPDE dynamics. \square

Without the adjustment step in the Markov transition kernel, one can update the weights in discrete-time Jarzynski as follows.

Lemma 3 (Time-discrete Jarzynski without detail balance). *With the discretized Gaussian SPDE proposal as in step 5 of Algorithm 2, the weights can be calculated along the dynamics as*

$$\exp(A_1(\{X_s\}_{s \in [0,1]})) = \exp(-\tilde{I}(X_1) - \tilde{J}(X_1) + \tilde{I}(X_0)) \prod_{s=0}^1 \frac{T_s(X_s, X_{s-\delta_s})}{T_{s-\delta_s}(X_{s-\delta_s}, X_s)} \quad (25)$$

such that $\mathbb{E}[\exp(A_1(\{X_s\}_{s \in [0,1]}))] = \tilde{Z}_1 / \tilde{Z}_0$ when initialized at $X_0 \sim \mathbb{P}^{ref}$. Here \tilde{Z}_0, \tilde{Z}_1 are the normalizing constants for the discretized prior $\exp(-\tilde{I}(X))$ and posterior $\exp(-\tilde{I}(X) - \tilde{J}(X))$ in the physical time $t = 0, \delta_t, \dots, T$, as defined in step 6 of Algorithm 2. Moreover, (25) implies that statistics $h(\cdot)$ over the path space can be estimated as

$$\frac{\sum_{i=1}^M h(X_1^i) \exp(A_1(\{X_s^i\}_{s \in [0,1]}))}{\sum_{i=1}^M \exp(A_1(\{X_s^i\}_{s \in [0,1]}))}$$

with M trajectories.

Proof. This is a generalization of Proposition 1 from [10]. Define the backward and forward Markov transition kernels of Algorithm 2 over the path space as

$$T_s(X_s, X_{s-\delta_s}) \propto \exp\left(-\frac{\delta_t}{4\delta_s} \|LX_{s-\delta_s} - RX_s + M_s(X_s)\delta_s\|^2\right)$$

and

$$T_{s-\delta_s}(X_{s-\delta_s}, X_s) \propto \exp\left(-\frac{\delta_t}{4\delta_s} \|LX_s - RX_{s-\delta_s} + M_{s-\delta_s}(X_{s-\delta_s})\delta_s\|^2\right)$$

where $u = -s\nabla V$ in the definition of M_s . Then

$$\mathbb{E}[\exp(A_1(\{X_s\}_{s \in [0,1]}))]$$

$$\begin{aligned} &= \tilde{Z}_0^{-1} \int_{\mathbb{R}^{dT_n}} \exp(-\tilde{I}(X_1) - \tilde{J}(X_1) + \tilde{I}(X_0)) \prod_{s=\delta_s}^1 \frac{T_s(X_s, X_{s-\delta_s})}{T_{s-\delta_s}(X_{s-\delta_s}, X_s)} \exp(-\tilde{I}(X_0)) \prod_{s=\delta_s}^1 T_{s-\delta_s}(X_{s-\delta_s}, X_s) dX_0 \cdots dX_1 \\ &= \tilde{Z}_0^{-1} \int_{\mathbb{R}^{dT_n}} \exp(-\tilde{I}(X_1) - \tilde{J}(X_1)) \prod_{s=\delta_s}^1 T_s(X_s, X_{s-\delta_s}) dX_0 \cdots dX_1 \\ &= \tilde{Z}_0^{-1} \int_{\mathbb{R}^{dT}} \exp(-\tilde{I}(X_1) - \tilde{J}(X_1)) dX_1 = \tilde{Z}_0^{-1} \tilde{Z}_1 \end{aligned}$$

with $n = 1/\delta_s$ as the number of switching steps. Above we used the fact that T_s is a normalized transition density for all s . (25) is a discretization in both the algorithmic and physical time (captured by \tilde{I}, \tilde{J} rather than I, J that's defined over continuous $t \in [0, T]$). \square

B.4 Alternative Equilibrium Schemes

For *exactly* following the π_s curve, one may be tempted to learn the drift term in the SPDE (16) as a replacement for $\delta_x \log Q(x)$, in order to enforce this requirement. This will give a method similar to [25], but here with v_s mapping $\mathcal{C}([0, T], \mathbb{R}^d)$ to $\mathcal{C}([0, T], \mathbb{R}^d)$:

$$-\nabla \cdot (\pi_s(x)v_s(x)) \stackrel{!}{=} \pi_s(x) (J(x) - \mathbb{E}_{\pi_s(x)}[J(x)]) = \partial_s \pi_s \quad (26)$$

or equivalently written as the Stein equation (which is oblivious to normalizing constant)

$$-\nabla \log \pi_s(x) \cdot v_s(x) + \nabla \cdot v_s(x) \stackrel{!}{=} J(x) - \int J(x) d\pi_s \quad (27)$$

and update trajectories as

$$dx_s/ds = v_s(x_s), \quad x_0 \sim \pi_0.$$

This method will directly move the trajectories instead of updating the drifts b^s (as we do in Section 2), and is closer to an interacting particle system, if we heuristically work with an implementation on (27) as (the density $\pi_s(x)$ here should be thought of as a joint distribution over the $\mathbb{R}^{d \times T}$ space, flattened):

$$\arg \min_{v_s} \sum_{i=1}^N \left[\nabla_x \log \pi_s(x^i) \cdot v_s(x^i) + \nabla \cdot v_s(x^i) - \left(J(x^i) - \frac{1}{N} \sum_{n=1}^N J(x^n) \right) \right]^2 \quad \text{for } x^i \sim \pi_s \quad (28)$$

$$x^i \leftarrow x^i + \delta_s \cdot v_s(x^i) \quad \forall i \in [N]$$

that moves trajectories so that the updated $\{x^i\}_{i=1, \dots, N} \sim \pi_{s+\delta_s}$. Note that the vector field $v_s(x^i): \mathbb{R}^{d \times T} \mapsto \mathbb{R}^{d \times T}$ means we are modeling the joint distribution over the temporal dimension, as well as the spatial one. However, solving (28) means we need to keep $N \times \mathbb{R}^d \times T$ in memory, compared to our NN approach in Section 2 that keeps $\#\text{NN}$ parameters in memory. If one considers Fisher-Rao/Birth-Death on the trajectories with (26) (which attach weights to the trajectories based on J), the memory footprint is similar. In Algorithm 1 we pay for compute in terms of re-generating trajectories after each annealing step, but recover the b_t^* drift prescribing the behavior of the dynamics as an output. The proposal above, on the other hand, does not keep the intermediate path measures in the SDE form.

On the other hand, using [1] with a PINN objective imposing (27) does not need interacting particle system, but will require re-simulating paths with the updated v_s at each step and solve a *two-parameter* loss evaluated on the trajectories. We remark that Algorithm 1 from [1] also leverages explicit ‘annealing’ w.r.t the changing potential (assumed known), in addition to the learned transport v_s part.

C General Reduction and Eulerian Solver

We begin with a remark on the general approach adopted in Section 3, followed by the proof of our jointly convex reformulation involving q_t, m_t from Section 3.1.

Remark 3. There are some degrees of freedom in the objective (7)/(8) by picking different decomposition of Q^* , e.g., using Brownian Bridge as \mathbb{P}^{ref} with J from (20) and $u_t^{\text{ref}}(x_t)$ from (21), in addition to the choice of \mathbb{P}^{ref} as a prior, and J the likelihood given by observations as in (4).

Proposition 1 (Convex Reformulation). *For a general reference $u_t^{\text{ref}} \neq 0$, (9) is equivalent to*

$$\begin{aligned} \min_{q_t, m_t} \int_0^T \int_{\mathbb{R}^d} & \left[\frac{1}{4} \left\| \frac{m_t}{q_t}(x_t) - (u_t^{\text{ref}} - \nabla \log \mathbb{P}_t^{\text{ref}})(x_t) \right\|^2 + \frac{1}{4} \left\| \nabla \log \frac{q_t}{\mathbb{P}_t^{\text{ref}}}(x_t) \right\|^2 + J_t(x_t) \right] q_t(x_t) dx dt \\ & + \frac{1}{2} \int_{\mathbb{R}^d} q_T(x_T) \log \frac{q_T}{\mathbb{P}_T^{\text{ref}}}(x_T) dx_T \\ \text{s.t. } \partial_t q_t &= -\nabla \cdot m_t, q_0 = \rho_0. \end{aligned} \quad (10)$$

When the reference process \mathbb{P}^{ref} is in equilibrium (i.e., $\mathbb{P}_t^{\text{ref}} = \nu \propto e^{-V} \forall t$) the problem reduces to

$$\begin{aligned} \min_{q_t, m_t} \int_0^T \int_{\mathbb{R}^d} & \left[\frac{1}{4} \left\| \frac{m_t}{q_t}(x_t) \right\|^2 + \frac{1}{4} \left\| \nabla \log \frac{q_t}{\nu}(x_t) \right\|^2 + J_t(x_t) \right] q_t(x_t) dx dt + \frac{1}{2} \int_{\mathbb{R}^d} q_T(x_T) \log \frac{q_T}{\nu}(x_T) dx_T \\ \text{s.t. } \partial_t q_t &= -\nabla \cdot m_t, q_0 = \rho_0 \end{aligned} \quad (11)$$

In this case, $\rho_0 = \nu = \mathbb{P}_0^{\text{ref}}$ for the initial distribution.

Proof. We give an argument here that is slightly different from Eqn (4.31) in [13] and [17]: let

$$dX_t = u_t^{\text{ref}}(X_t) dt + \sqrt{2} dW_t$$

with any fixed, arbitrary initial distribution, where we denote its marginal density as \bar{q}_t . Now rewriting the constraint in (9), we have (introduce the variable $v_t = b_t + u_t^{\text{ref}} - \nabla \log q_t$)

$$\partial_t q_t = -\nabla \cdot (q_t v_t), q_0 = \rho_0.$$

Similarly for the objective in (9), a change of variables gives

$$\begin{aligned} & \int_{\mathbb{R}^d} \left[\frac{1}{4} \|b_t(x_t)\|^2 + J_t(x_t) \right] q_t(x_t) dx \\ &= \int_{\mathbb{R}^d} \left[\frac{1}{4} \|v_t(x_t) - u_t^{\text{ref}}(x_t) + \nabla \log \bar{q}_t(x_t) + \nabla \log q_t(x_t) - \nabla \log \bar{q}_t(x_t)\|^2 + J_t(x_t) \right] q_t(x_t) dx \end{aligned}$$

Let $\bar{v}_t = u_t^{\text{ref}} - \nabla \log \bar{q}_t$. The cross term organizes into

$$\frac{1}{2} \int \left[(v_t - \bar{v}_t)^\top(x) \nabla \log \frac{q_t}{\bar{q}_t}(x) \right] q_t(x) dx = \frac{1}{2} \frac{d}{dt} \text{KL}(q_t \| \bar{q}_t)$$

which upon time integrating over $[0, T]$ gives $\frac{1}{2} \text{KL}(q_T \| \bar{q}_T)$ if q_T is not fixed and $\bar{q}_0 = q_0 = \rho_0$. After identifying $m_t = q_t v_t$, this gives (10).

If $u_t^{\text{ref}} = -\nabla V$ is not time-varying, and assuming we initialize $\bar{q}_0 \propto e^{-V} =: \nu$ at stationary, then $\bar{q}_t = \nu$ is fixed, and we have

$$\frac{1}{2} \int \left[v_t(x)^\top \nabla \log \frac{q_t}{\nu}(x) \right] q_t(x) dx = \frac{1}{2} \frac{d}{dt} \text{KL}(q_t \| \nu),$$

therefore time-integrating the cross term gives us $\frac{1}{2} \text{KL}(q_T \| \nu) - \frac{1}{2} \text{KL}(\rho_0 \| \nu)$, where the second part vanishes, which gives (11). \square

Before offering the details of our Eulerian solver, we begin with some motivations.

If we split (10) as follows, one can view the highlighted red part as encoding the geometry (i.e., defining appropriate notion of shortest distance between 2 distributions over the probability measure space) and the rest as the “objective” function for the $(q_t)_t$ dynamics.

$$\begin{aligned} & \min_{q_t, m_t} \int_0^T \int_{\mathbb{R}^d} \left[\frac{1}{4} \left\| \frac{m_t}{q_t}(x_t) - (u_t^{\text{ref}} - \nabla \log \mathbb{P}_t^{\text{ref}})(x_t) \right\|^2 \right] q_t(x_t) dx dt \\ & + \int_0^T \int_{\mathbb{R}^d} \left[\frac{1}{4} \left\| \nabla \log \frac{q_t}{\mathbb{P}_t^{\text{ref}}}(x_t) \right\|^2 + J_t(x_t) \right] q_t(x_t) dx dt + \frac{1}{2} \int_{\mathbb{R}^d} q_T(x_T) \log \frac{q_T}{\mathbb{P}_T^{\text{ref}}}(x_T) dx_T \\ & \text{s.t. } \partial_t q_t = -\nabla \cdot m_t, q_0 = \rho_0 \end{aligned}$$

More specifically, since the first part depends on both $q_t, \dot{q}_t = m_t/q_t = v_t$, while the remaining part depends on q_t only, for a given ground cost $L(\cdot, \cdot)$ over Euclidean space, if we define

$$W_L(q_0, q_h) := \inf_{\rho_s} \inf_{v_s} \int_0^h \int_{\mathbb{R}^d} L(x_s, v_s) \rho_s(x_s) dx_s ds \quad \text{s.t. } \dot{\rho}_s = -\nabla \cdot (\rho_s v_s), \rho_0 = q_0, \rho_h = q_h \quad (29)$$

then we may propose to iteratively solve

$$q_{t+h} \leftarrow \arg \min_{q \in \mathcal{P}(\mathbb{R}^d)} W_L(q, q_t) + h \cdot \underbrace{\int_{\mathbb{R}^d} \left[\frac{1}{4} \left\| \nabla \log \frac{q}{\mathbb{P}_{t+h}^{\text{ref}}}(x) \right\|^2 + J_{t+h}(x) \right] q(x) dx}_{F(q)}. \quad (30)$$

for time stepping across $t \in [0, T]$ with time step h .

Choosing $L(x_s, v_s) = \frac{1}{2} \|v_s(x_s)\|^2$ in (29) recovers the \mathcal{W}_2 OT distance used in the reduction (12) for (11). One can view (29) as a dynamic definition of a new metric. Here, again, we use the interpretation that given a curve of optimal densities for (10), the optimization over v_t (equivalently m_t), done in the red part, is simply solving an OT problem between the neighboring marginals, using a special cost. This elimination of v_t (equivalently m_t) turns the problem into a single-variable optimization over the $(q_t)_t$ curve only and provides a good approximation when h is small (we elaborate on this in Appendix F).

C.1 Low Dimensional Case: Composite Optimization for (10)

For the most general case (10), since the resulting optimization is jointly convex in the variables – the system in (10) should be contrasted with Theorem 1 in [16] – one may discretize the space and apply convex optimization methods on the resulting objective.

The formulation in (10) requires $\nabla \log \mathbb{P}_t^{\text{ref}}$ for the reference process, which we typically don’t know analytically. But assuming (1) one can sample from the prior trajectories; (2) $u_t^{\text{ref}} = -\nabla V_t$ takes gradient form, an Action Matching loss [23, Algorithm 1] can be used to learn this, where we expect $\nabla s_t^*(x) \approx -\nabla V_t(x) - \nabla \log \mathbb{P}_t^{\text{ref}}(x)$. In the case of TPS with the Brownian bridge as reference, we have $\nabla \log \mathbb{P}_t^{\text{ref}}(x_t) = -\frac{x_t - A}{2t}$ analytically available. [29, Lemma B.1] gives another way to compute this score, which also only requires simulating from the reference process and fitting a least squares objective.

With (10) in mind, we can adapt the iterative procedure in [11] to solve this convex problem via a primal-dual method (a related Augmented Lagrangian method for WGF is also considered in [6]). Our method reduces to minimizing a finite-dimensional composite objective of type

$$\min_y \Phi(y) + \Psi(\Lambda y)$$

for each time stepping from t to $t+h$ – therefore we solve in total T/h problems of such type. Above Φ, Ψ are both convex functions that admit an easy-to-compute proximal operator. Departing from the min-max formulation [24], this is generally simpler, with convergence ensured.

Now we are ready to proceed with the approximation (30) and the definition (29) to design our Eulerian solver. We preface with a definition.

Definition 1. An important convex set \mathcal{K} that will be repeatedly used is

$$\mathcal{K} := \left\{ (a, b) \in \mathbb{R} \times \mathbb{R}^d : a + \frac{1}{2} \|b\|^2 \leq 0 \right\}.$$

The convex conjugate of the characteristic function of the set $\chi_{\mathcal{K}}^*(a, b) = \sup_{(a,b) \in \mathcal{K}} \{ar + b^\top m\} = \frac{\|m\|^2}{2r}$ for $(r, m) \in \mathbb{R} \times \mathbb{R}^d$ is its support function.

Without loss of generality, consider the first $q_0 \rightarrow q_h$ step. We approximate the $W_L(q_0, q_h)$ term, similarly as done in [11] as

$$W_L(q_0, q_h) \approx \inf_{\bar{m}} \left\{ \int \frac{\|\bar{m}\|^2}{2q_0}(x) - \bar{m}^\top(x)(u_0^{\text{ref}} - \nabla \log \mathbb{P}_0^{\text{ref}})(x) + \frac{1}{2} \|u_0^{\text{ref}}(x) - \nabla \log \mathbb{P}_0^{\text{ref}}(x)\|^2 q_0(x) dx : q_h - q_0 = -\nabla \cdot \bar{m} \right\}$$

This relies on the fact that if q_0, q_h are close, the momentum m_s in the definition of $W_L(q_0, q_h)$ becomes almost s -independent, which we denote by \bar{m} above.

Now introduce notations that group known/fixed quantities together: let

$$g(x) = u_0^{\text{ref}}(x) - \nabla \log \mathbb{P}_0^{\text{ref}}(x), \quad f_h(x) = 2 \cdot J_h(x) + \frac{h}{2} \|\nabla \log \mathbb{P}_h^{\text{ref}}(x)\|^2 + h \cdot \Delta \log \mathbb{P}_h^{\text{ref}}(x)$$

and using the representation of Fisher information in (14), the problem reduces to solve the following convex problem (for some constant C)

$$\inf_{q_h, \bar{m}} \int \frac{1}{2} \frac{\|\bar{m}\|^2}{q_0}(x) - \bar{m}^\top(x)g(x) + f_h(x)q_h(x)dx + C + \sup_{\phi_h, (\alpha, \beta) \in \mathcal{K}} \int h \cdot (\alpha(x)q_h(x) - \nabla \cdot \beta(x)q_h(x)) + \phi_h(x)(q_h(x) - q_0(x) + \nabla \cdot \bar{m}(x))dx.$$

It implies from optimality condition on q_h, \bar{m} that

$$\bar{m}(x) = q_0(x)[g(x) + \nabla \phi_h(x)], \quad \alpha(x) = -\frac{1}{h}f_h(x) + \nabla \cdot \beta(x) - \frac{1}{h}\phi_h(x).$$

Plugging it back in and flipping the sign, we end up with solving

$$\begin{aligned} \inf_{\phi_h, \beta} \quad & \frac{1}{2} \int_{\mathbb{R}^d} \|\nabla \phi_h(x) + g(x)\|^2 q_0(x) dx + \int_{\mathbb{R}^d} \phi_h(x) q_0(x) dx \\ \text{s.t.} \quad & \left(\nabla \cdot \beta(x) - \frac{1}{h}f_h(x) - \frac{1}{h}\phi_h(x), \beta(x) \right) \in \mathcal{K} \quad \forall x \in \mathbb{R}^d. \end{aligned} \quad (31)$$

And the resulting updated q_h, v_h can be recovered as

$$q_h(x) \approx q_0(x) - \nabla \cdot (q_0(x)[\nabla \phi_h(x) + g(x)]) \quad \text{and} \quad v_h(x) \approx \nabla \phi_h(x) + g(x). \quad (32)$$

Problem (31) can now be solved using the Chambolle-Pock algorithm similar to what's done in [11]. As a splitting scheme for alternatively updating the 2 parts of the convex objective:

$$\min_z \Phi(z) + \Psi(\Lambda z)$$

we write for $z = (\phi_h, \beta)$ and $\Lambda z = (\nabla \cdot \beta - \frac{1}{h}\phi_h, \beta)$:

$$\begin{aligned} \Phi(z) &= \frac{1}{2} \int_{\mathbb{R}^d} \|\nabla \phi_h(x) + g(x)\|^2 q_0(x) dx + \int_{\mathbb{R}^d} \phi_h(x) q_0(x) dx, \\ \Psi(\Lambda z) &= \Psi(z_1, z_2) = \begin{cases} 0, & \text{if } (z_1(x), z_2(x)) \in \mathcal{K}_{1/h \cdot f_h}, \text{ for each } x \in \mathbb{R}^d \\ \infty, & \text{otherwise} \end{cases} \end{aligned}$$

where we overload the notation so $\mathcal{K}_c = \{(a, b) \in \mathbb{R} \times \mathbb{R}^d : a + \frac{1}{2}\|b\|^2 \leq c\}$ for some known constant c .

In order to run the iterative update prescribed by Chambolle-Pock (for $k \geq 1$, some stepsize $\sigma, \delta > 0$, and initialization $y^0 = u^0 = z^0$)

$$\begin{aligned} y^{k+1} &= \text{prox}_{\sigma\Psi^*}(y^k + \sigma\Lambda u^k) \\ z^{k+1} &= \text{prox}_{\delta\Phi}(z^k - \delta\Lambda^* y^{k+1}) \\ u^{k+1} &= 2z^{k+1} - z^k \end{aligned} \quad (33)$$

we only need to know the prox operator for Φ, Ψ ; since $\text{prox}_{\Psi^*} = I_d - \text{prox}_{\Psi}$ follows easily. The prox operator for $\Phi(z)$ involves solving a linear system on ϕ_h , which returns $\text{prox}_{\delta\Phi}(\phi_0, \beta_0) = (\phi, \beta_0)$ where ϕ solves at each x on the grid:

$$\phi + \delta \cdot D_\tau^*(D_\tau \phi q_0) = \phi_0 - \delta \cdot q_0 - \delta \cdot D_\tau^*(gq_0).$$

Above D_τ is a finite difference operator with grid spacing τ in each direction, and $-D_\tau^*$ approximates the divergence. The other prox operator for Ψ is a projection onto the set \mathcal{K} that is separable across space (turns into a 1D problem) and explicit [4]. It remains to work out Λ^* for running the update (33): since $\langle y, \Lambda z \rangle = \langle \Lambda^* y, z \rangle$, we have $\Lambda^* y = (-1/h \cdot y_1, -D_\tau y_1 + y_2)$ in this case.

We have seen how to advance $t = 0 \rightarrow t = h$ for approximately solving (10) – the whole process can be sequentially repeated until time T since we have the density $q(x)$ explicitly computed (and maintained) on a fixed grid over \mathbb{R}^d at each step (c.f. (32)).

The method above, though NN-free, requires updates on a *fixed grid* in \mathbb{R}^d (since it uses an Eulerian representation of the discretized density over space), so may run into scalability issues in high dimension.

D Additional Details for Section 3.2

Below we give the second Lagrangian approach based on (12) in Section D.0.1, and provide further justification of the reduction from (11) to (12) in Section D.0.2.

D.0.1 Interacting Particle System

To understand the iterative scheme (12), which is better termed acceleration in Wasserstein geometry with $I(q; \nu)$ functional, we see that although one description on the optimal marginal density evolution in $\mathcal{P}(\mathbb{R}^d)$ is through the pair of PDEs Fokker-Planck (evolved till time T):

$$\partial_t q_t = -\nabla \cdot (q_t(b_t + u^{\text{ref}})) + \Delta q_t, \quad q_0 = \rho_0$$

and HJB on b_t (18)-(19) separately, it is more convenient to work with the following.

Proposition 2 (Optimality Condition). *The Euler–Lagrange equations for (11) over $[0, T]$ take the more suggestive coupled PDE form (also implies the optimal $\nabla\theta_t$ takes gradient form)*

$$\begin{aligned} \partial_t q_t &= -\nabla \cdot (q_t \nabla \theta_t), \quad q_0 = \rho_0, \quad \nabla \theta_0 = b_0 \\ \partial_t \theta_t + \frac{1}{2} \|\nabla \theta_t\|^2 &= -2(\Delta \sqrt{q_t})/\sqrt{q_t} + 2\Delta \sqrt{\nu}/\sqrt{\nu} + 2J_t \\ 2J_T(X_T) + \log \frac{q_T}{\nu}(X_T) + \theta_T(X_T) &= 0. \end{aligned}$$

Therefore for $t \in (0, T)$, the solution to (11) obeys

$$\begin{aligned} \dot{X}_t &= \nabla \theta_t(X_t) \\ \ddot{X}_t &= \nabla \left[\frac{\partial \theta_t(X_t)}{\partial t} + \frac{1}{2} \|\nabla \theta_t(X_t)\|^2 \right] = - \left(2\nabla \frac{\Delta \sqrt{q_t}}{\sqrt{q_t}} - 2\nabla \frac{\Delta \sqrt{\nu}}{\sqrt{\nu}} - 2\nabla J_t \right) (X_t) \\ &= -\nabla \left(\Delta \log q_t + \frac{1}{2} \|\nabla \log q_t\|^2 + \Delta V - \frac{1}{2} \|\nabla V\|^2 - 2J_t \right) (X_t) =: -\nabla F(q_t)(X_t). \end{aligned} \quad (34)$$

This implies in (34) we are accelerating with force field F as in Newton’s law $\ddot{X}_t = -\nabla F(q_t)(X_t)$.

Proposition 2 gives a proper meaning to acceleration on the particle level. Compared to OT that gives constant-speed geodesic $\ddot{X}_t = 0$, here we have that acceleration is given by both the gradient of the spatial potential and an extra “potential energy” term from the KL functional $\frac{1}{2} \|\nabla_{\mathcal{W}_2} \text{KL}(q\|\nu)\|_q^2 = I(q; \nu)$. In the absence of J and the terminal constraint, one expects $q_T \rightarrow \nu$ asymptotically.

Written in integral form, the trajectory of the particle is therefore given by $(X_0 \sim \rho_0, \nabla \theta_0(X_0) \approx 0)$

$$X_h = X_0 + h \nabla \theta_0(X_0) - \int_0^h \int_0^s \nabla F(q_t)(X_t) dt ds = X_0 + h \nabla \theta_0(X_0) - \int_0^h \nabla F(q_t)(X_t)(h-t) dt. \quad (35)$$

To numerically simulate (34) to approximately follow the optimal marginal density curve q_t , one can kernelize the gradient field in (34) with an interacting particle system. The nonlinear term, $\nabla(\Delta\sqrt{q}/\sqrt{q})$, which corresponds to $\nabla\delta R(q)$ [18], at first glance, poses a challenge for estimation with samples from q_t only, but similar linearization ideas as earlier can be used here for approximation. In particular, the relation (14) shows that $\delta R(q_t)(x) = \alpha^*(x) - \nabla \cdot \beta^*(x)$ from Danskin's Theorem, where $\alpha^*(x) = -\frac{1}{2}\|\nabla \log q_t(x)\|^2$, $\beta^*(x) = \nabla \log q_t(x)$ – this representation of the first variation can bypass density estimation and becomes amenable to particle implementation via $\mathcal{K}_{q_t} \nabla \delta R(q_t)(z) = \int \mathcal{K}(x-z) \nabla[\alpha^*(x) - \nabla \cdot \beta^*(x)] q_t(x) dx$ that can be used in $\nabla F(q_t)$ for running the 2nd-order ODE dynamics (35). More details can be found in Appendix E.1. Compared to Algorithm 3 from Section 3.2.1, the benefit of this approach is that we don't need to fit a NN since everything is analytical, although we expect the performance to hinge on that of the kernelized gradient field approximation.

D.0.2 Entropic-regularized OT and Alternative Approximation

To better interpret the methodology proposed in Section 3.2.1, D.0.1 within a broader context, re-writing (11), between $t \in [0, h]$ our method solves (assume J consists of δ -functions as in (4))

$$\begin{aligned} & \min_{q_t \in [0, h], v_t} \frac{1}{2h} \left(h \int_0^h \int_{\mathbb{R}^d} q_t(x) \|v_t(x)\|^2 dx dt \right) + \int_0^h \int \left[\frac{1}{2} \|\nabla \log \frac{q_t}{\nu}(x_t)\|^2 + 2J_t(x_t) \right] q_t(x_t) dx dt \\ & \approx \frac{1}{2h} \left(h \int_0^h \int_{\mathbb{R}^d} q_t(x) \|v_t(x)\|^2 dx dt \right) + h \cdot \int \frac{1}{2} \|\nabla \log \frac{q_h}{\nu}(x_h)\|^2 q_h(x_h) dx + 2 \int J_h(x_h) q_h(x_h) dx \end{aligned} \quad (36)$$

subject to continuity equation and $q_0 = \rho_0$. In (12) we used the 2nd line as approximation (after rescaling). But a different grouping on the 1st line leads to (using argument very similar to Proposition 1, also see [15])

$$= \min_{q_h} \left\{ \underbrace{\min_{\pi \in \Pi(\rho_0, q_h)} \{ \text{KL}(\pi \| R_{0,h}) \}}_{\zeta_h(\rho_0, q_h)} + \text{KL}(\rho_0 \| \nu) - \text{KL}(q_h \| \nu) + 2 \int J_h(x_h) q_h(x_h) dx \right\} \quad (37)$$

where $R_{0,h}$ is the joint law at time 0, h of the reference SDE with invariant measure ν . From [15, Eqn (1.10)] we have that in the short-time limit

$$h \cdot \zeta_h(\rho_0, q_h) = \underbrace{\frac{1}{2} \mathcal{W}_2^2(\rho_0, q_h) + h(\text{KL}(q_h \| \nu) - \text{KL}(\rho_0 \| \nu))}_{(*)} + \underbrace{\frac{h^2}{2} \int \int_0^1 \|\nabla \log \frac{q_t^0}{\nu}\|^2 q_t^0 dt dx}_{\approx h^2 \cdot I(q_h; \nu)} + o(h^2) \quad (38)$$

for $(q_t^0)_t$ the (unique) Wasserstein geodesic between ρ_0 and q_h . Dividing by h , this recovers (36) and justifies the derivation in (12), giving a quantitative bound on the approximation.

It is clear that with formulation (37), we only need to solve at each new observation time $\{t_k\}_{k=1}^K$ – since without the J term, its solution is given by $q_h = \nu$ if $\rho_0 = \nu$. For reversible Brownian motion W as reference (i.e., $\nu = \text{Leb}$), (37) amounts to solving between 2 observations (after cancellation):

$$q_{t_{i+1}} \leftarrow \arg \min_q \mathcal{W}_{2, t_{i+1}-t_i}^2(q, q_{t_i}) + \bar{F}(q) \quad (39)$$

with entropy-regularized \mathcal{W}_2 distance

$$\mathcal{W}_{2, t_{i+1}-t_i}^2(q, q_{t_i}) = \min_{\pi^i \in \Pi(q_{t_i}, q)} \int \frac{1}{2(t_{i+1} - t_i)} \|x - y\|^2 d\pi^i(x, y) + \text{KL}(\pi^i \| q_{t_i} \otimes q)$$

as the distance-inducing metric and the likelihood $2 \int J(x) dq(x)$ term as the objective functional \bar{F} . We give methods for implementation with this objective, along with additional discussions in Appendix F. While (39) has a simpler objective than (12), it does not admit a push-forward map solution (15) or modeling by ODE dynamics (34). Nevertheless, the sequential reduction implies that we still track the $\{q_t\}_t$ dynamics as a consecutive optimization over $\mathcal{P}(\mathbb{R}^d \times \mathbb{R}^d)$, instead of a multi-marginal problem on the joint space $\mathcal{P}(\mathbb{R}^d)^K$.

Remark 4. The more serious concern for this regularized-OT approach (39), compared to Algorithm 3 and Proposition 2 that only require ∇V , is its generalization to $u^{\text{ref}} \neq 0$. Importantly, we show both \mathcal{W}_2^2 and $I(q; \nu)$ are easily estimable in Algorithm 3 with empirical samples. In (39), although we can time step across $t \in [0, T]$ with a larger interval $t_{i+1} - t_i$ and it's expected to be more accurate, the inner optimization $\mathcal{W}_{2, t_{i+1}-t_i}^2$ is not easily solvable without knowledge of the analytical transition kernel of the reference process. And the only asymptotic one may resort to is $R_{t_i, t_{i+1}} \rightarrow \nu \otimes \nu$ as $t_{i+1} - t_i \rightarrow \infty$, which is too loose for most practical purpose. In contrast, in (12), we pay the price of having to do smaller step-size h between adjacent marginals because of the $o(h^2)$ approximation resulting from evaluating $I(q_t; \nu)$ at a single marginal instead of along the entropic interpolation, but if h is small, even the first term (\star) in (38) involving $\frac{1}{2}\mathcal{W}_2^2$ furnishes a $o(h)$ approximation [15].

E Missing Derivation for the Two Lagrangian Methods from Section 3.2

Below we give the details for our two Lagrangian methods from Section 3.2 – we begin with the Fisher Information Functional estimator used in Section 3.2.1 and Algorithm 3.

Lemma 4. *Given n samples $\{x_i\}_{i=1}^n \sim q_t$, one can efficiently estimate the Fisher Information functional as $\hat{R}(q_t) \approx \frac{1}{n} \sum_{i=1}^n \left(\frac{1}{m} \sum_{j=1}^m \frac{\|y_i^j - x_i\|^2}{2\sigma^4} \right)$ for $y_i^j \sim \mathcal{N}(x_i, \sigma^2 I)$, $j = 1 \dots, m$, and σ a bandwidth parameter.*

Proof. Notice the optimization (14) becomes essentially 1D and separable across grid points, if we approximate

$$\hat{q}_t(x) = \frac{1}{n} \sum_{i=1}^n \delta_{x_i} \approx \frac{1}{n} \sum_{i=1}^n \delta_{x_i}(x) * \mathcal{N}(0, \sigma^2 I) = \frac{1}{n} \sum_{i=1}^n \mathcal{N}(x_i, \sigma^2 I)(x) \text{ for } \{x_i\}_{i=1}^n \sim q_t$$

with a bandwidth σ , so

$$\nabla \hat{q}_t(x) \approx \frac{1}{n} \sum_{i=1}^n -\frac{x - x_i}{\sigma^2} \mathcal{N}(x_i, \sigma^2 I)(x).$$

This makes $\nabla \hat{q}_t$, and therefore $R(\hat{q}_t)$ amenable to estimation given i.i.d samples from q_t . For solving (13), we sample y_j 's from Gaussians centered at $\{x_i\}_{i=1}^n$ where $x_i \sim q_t(x)$, and return

$$\begin{aligned} \int \nabla \hat{q}_t(y) \beta(y) dy &\approx \frac{1}{n} \sum_{i=1}^n \int -\frac{y - x_i}{\sigma^2} \mathcal{N}(x_i, \sigma^2 I)(y) \beta(y) dy \approx \frac{1}{n} \sum_{i=1}^n -\frac{y_i - x_i}{\sigma^2} \beta(y_i), \\ \int \alpha(y) \hat{q}_t(y) dy &= \frac{1}{n} \sum_{i=1}^n \int \alpha(y) \mathcal{N}(x_i, \sigma^2 I)(y) dy \approx \frac{1}{n} \sum_{i=1}^n \alpha(y_i) \end{aligned}$$

so estimating $R(\hat{q}_t)$ in turn becomes solving on the (random) n -point grid $y_i \sim \mathcal{N}(x_i, \sigma^2 I)$

$$R(\hat{q}_t) \approx \sup_{(\alpha_i, \beta_i) \in \mathbb{R} \times \mathbb{R}^d \in \mathcal{K}, \forall i} \frac{1}{n} \sum_{i=1}^n \left\{ \left(-\frac{y_i - x_i}{\sigma^2} \right)^\top \beta_i + \alpha_i \right\} = \frac{1}{n} \sum_{i=1}^n \frac{\|y_i - x_i\|^2}{2\sigma^4}. \quad (40)$$

This can also be seen as a consequence of Stein's lemma: $\mathbb{E}[(X - \mu)^\top \beta(X)] = \mathbb{E}[\nabla \cdot \beta(X)]$ for $X \sim \mathcal{N}(\mu, I)$ on (14). In practice, a scheme that involves multiple samples at each x_i point as $y_i^j \sim \mathcal{N}(x_i, \sigma^2 I)$, $j \in [m]$

$$R(\hat{q}_t) \approx \frac{1}{n} \sum_{i=1}^n \left(\frac{1}{m} \sum_{j=1}^m \frac{\|y_i^j - x_i\|^2}{2\sigma^4} \right)$$

maybe more accurate. □

We'd like to emphasize that the estimator given in Lemma 4 is much simpler than what one would get by plugging in a density estimator \hat{q}_t for q_t , and compute

$$R(q_t) \approx R(\hat{q}_t) = \frac{1}{2} \int \|\nabla \log \hat{q}_t(x)\|^2 \hat{q}_t(x) dx$$

directly. And this convenience crucially hinges on the variational representation (14).

E.1 Missing Proofs for Section D.0.1

The core result on which our interacting particle system is based off is given below.

Proposition 2 (Optimality Condition). *The Euler–Lagrange equations for (11) over $[0, T]$ take the more suggestive coupled PDE form (also implies the optimal $\nabla\theta_t$ takes gradient form)*

$$\begin{aligned}\partial_t q_t &= -\nabla \cdot (q_t \nabla \theta_t), \quad q_0 = \rho_0, \quad \nabla \theta_0 = b_0 \\ \partial_t \theta_t + \frac{1}{2} \|\nabla \theta_t\|^2 &= -2(\Delta \sqrt{q_t})/\sqrt{q_t} + 2\Delta \sqrt{\nu}/\sqrt{\nu} + 2J_t \\ 2J_T(X_T) + \log \frac{q_T}{\nu}(X_T) + \theta_T(X_T) &= 0.\end{aligned}$$

Therefore for $t \in (0, T)$, the solution to (11) obeys

$$\begin{aligned}\dot{X}_t &= \nabla \theta_t(X_t) \\ \ddot{X}_t &= \nabla \left[\frac{\partial \theta_t(X_t)}{\partial t} + \frac{1}{2} \|\nabla \theta_t(X_t)\|^2 \right] = - \left(2\nabla \frac{\Delta \sqrt{q_t}}{\sqrt{q_t}} - 2\nabla \frac{\Delta \sqrt{\nu}}{\sqrt{\nu}} - 2\nabla J_t \right) (X_t) \\ &= -\nabla \left(\Delta \log q_t + \frac{1}{2} \|\nabla \log q_t\|^2 + \Delta V - \frac{1}{2} \|\nabla V\|^2 - 2J_t \right) (X_t) =: -\nabla F(q_t)(X_t).\end{aligned}\tag{34}$$

This implies in (34) we are accelerating with force field F as in Newton's law $\ddot{X}_t = -\nabla F(q_t)(X_t)$.

Proof. We multiply through by 2 to make the final expression cleaner. Writing out the Lagrangian for (11), we have using integration by parts:

$$\begin{aligned}\min_{q_t, v_t} \max_{\lambda_0, \theta_t} \int_0^T \int \left[\frac{1}{2} \|v_t(X_t)\|^2 + \frac{1}{2} \|\nabla \log \frac{q_t}{\nu}(X_t)\|^2 + 2J_t(X_t) \right] q_t(X_t) dX_t dt \\ - \int_0^T \int (\partial_t \theta_t(X_t) \cdot q_t(X_t) + \nabla \theta_t(X_t)^\top v_t(X_t) q_t(X_t)) dX_t dt \\ + \int \theta_T(X_T) q_T(X_T) dX_T - \int \theta_0(X_0) q_0(X_0) dX_0 \\ + \int \lambda_0(X_0) (q_0(X_0) - \rho_0(X_0)) dX_0 + \int q_T(X_T) \log \frac{q_T}{\nu}(X_T) dX_T\end{aligned}$$

Now the optimality condition for v_t implies

$$v_t(X_t) = \nabla \theta_t(X_t)$$

and optimality condition for q_t can be worked out using the first variation of $\delta I(q_t)$ from [18], where $I(q_t; \nu)$ is regarded as a potential energy involving the density:

$$\delta q I(q_t; \nu)(X_t) + \frac{1}{2} \|v_t(X_t)\|^2 + 2J_t(X_t) - \partial_t \theta_t(X_t) - \nabla \theta_t(X_t)^\top v_t(X_t) = 0\tag{41}$$

which putting together with the previous equation yields the claim. Other parts are straightforward.

The continuity equation, which comes from the optimality condition for θ_t , implies that along the optimal dynamics, particle moves with velocity field $\nabla \theta_t(X_t)$, which upon taking another time derivatives suggests the particles accelerate as

$$\ddot{X}_t = \nabla \frac{\partial \theta_t}{\partial t}(X_t) + \nabla^2 \theta_t(X_t) \nabla \theta_t(X_t).$$

Direct calculation gives the equivalence to (34) using the derived PDE optimality condition (41). \square

Armed with this, we can simply run a deterministic 2nd order ODE dynamics for simulating particle trajectories with the force field $-\nabla F(q_t)$ from Proposition 2, which consists of a linear part involving J and a nonlinear part involving $I(q; \nu)$:

$$\begin{aligned}X_h &= X_0 + h \nabla \theta_0(X_0) - \int_0^h \int_0^s \nabla F(q_t)(X_t) dt ds = X_0 + h \nabla \theta_0(X_0) - \int_0^h \nabla F(q_t)(X_t)(h-t) dt \\ \nabla \theta_h(X_h) &= \nabla \theta_0(X_0) - \int_0^h \nabla F(q_t)(X_t) dt.\end{aligned}$$

Upon time discretization it allows us to trace the optimal density curve.

Kernelizing the vector field With the force field depending on q_t , we resort to an interacting particle system for executing this proposal. This is again, where the relation (14) for the Fisher Information functional becomes extremely useful, since naïvely it is not clear how one may estimate terms such as $\nabla(\Delta\sqrt{q}/\sqrt{q})$. In particular (14) shows that $\delta R(q_t)(x) = \alpha^*(x) - \nabla \cdot \beta^*(x)$ from Danskin's Theorem, where $\alpha^*(x) = -\frac{1}{2}\|\nabla \log q_t(x)\|^2$, $\beta^*(x) = \nabla \log q_t(x)$. Therefore the kernelized vector field evaluated at any $z \in \mathbb{R}^d$ is

$$\begin{aligned}
[\mathcal{K}_{q_t} \nabla \delta R(q_t)](z) &= \int \mathcal{K}(x-z) \nabla[\alpha^*(x) - \nabla \cdot \beta^*(x)] q_t(x) dx \\
&= - \int \alpha^*(x) \nabla(\mathcal{K}(x-z) q_t(x)) dx + \int \nabla \cdot \beta^*(x) \nabla(\mathcal{K}(x-z) q_t(x)) dx \\
&= - \int \alpha^*(x) (\nabla \mathcal{K}(x-z) q_t(x) + \mathcal{K}(x-z) \nabla q_t(x)) dx - \int \beta^*(x) \nabla \cdot [\nabla(\mathcal{K}(x-z) q_t(x))] dx \\
&= - \int \alpha^*(x) (\nabla \mathcal{K}(x-z) q_t(x) + \mathcal{K}(x-z) \nabla q_t(x)) dx \tag{42} \\
&\quad - \int \beta^*(x) \nabla \cdot \nabla \mathcal{K}(x-z) q_t(x) + 2\beta^*(x) \nabla \mathcal{K}(x-z)^\top \nabla q_t(x) dx - \int \beta^*(x) \mathcal{K}(x-z) \nabla \cdot (\nabla q_t(x)) dx
\end{aligned}$$

This suggests an implementation as

$$\begin{aligned}
X_{t+h}^i &= X_t^i + hV_t(X_t^i) + \frac{h^2}{2} \cdot (\mathcal{K} \nabla \delta R(\hat{q}_t)(X_t^i) + 2\nabla J_t(X_t^i)), \quad i \in [n] \\
V_{t+h}(X_{t+h}^i) &= V_t(X_t^i) + h \cdot (\mathcal{K} \nabla \delta R(\hat{q}_t)(X_t^i) + 2\nabla J_t(X_t^i))
\end{aligned}$$

for

$$\hat{q}_t(x) = \frac{1}{n} \sum_{i=1}^n \delta_{x_i} \approx \frac{1}{n} \sum_{i=1}^n \mathcal{N}(x_i, \sigma^2 I)(x) \quad \text{with } x_i \sim q_t$$

so the particle system with n particles will approximately follow the optimal $(q_t)_{t \leq T}$ curve with stepsize h .

In (42), earlier calculation from (40) shows that e.g., $\int \alpha^*(x) \nabla \mathcal{K}(x-z) q_t(x) dx$ can be approximated as

$$\frac{1}{n} \sum_{i=1}^n \underbrace{\int \alpha^*(x) \nabla \mathcal{K}(x-z) \mathcal{N}(x_i, \sigma^2 I)(x) dx}_{\mathbb{E}_{y \sim \mathcal{N}(x_i, \sigma^2 I)}[\alpha^*(y) \nabla \mathcal{K}(y-z)]} \approx \frac{1}{n} \sum_{i=1}^n \left(\frac{1}{m} \sum_{j=1}^m -\frac{\|y_i^j - x_i\|^2}{2\sigma^4} \nabla \mathcal{K}(y_i^j - z) \right) \tag{43}$$

since on a random grid $\{y_i^j\}_j \sim \mathcal{N}(x_i, \sigma^2 I)$ with $x_i \sim q_t$, we showed

$$\alpha^*(y_i^j) = -\frac{\|y_i^j - x_i\|^2}{2\sigma^4}, \quad \beta^*(y_i^j) = \frac{x_i - y_i^j}{\sigma^2}.$$

The only term that's not in terms of $q_t, \nabla q_t$ in (42) is the last one – the rest can be similarly estimated as above. However, one can show that the divergence term also takes a rescaled Gaussian form:

$$\begin{aligned}
\nabla \cdot \nabla \hat{q}_t(x) &= \frac{1}{n} \sum_{i=1}^n -\nabla \cdot \left(\frac{x - x_i}{\sigma^2} \mathcal{N}(x_i, \sigma^2 I)(x) \right) \\
&= \frac{1}{n} \sum_{i=1}^n -\frac{d}{\sigma^2} \mathcal{N}(x_i, \sigma^2 I)(x) + \frac{1}{n} \sum_{i=1}^n \frac{\|x - x_i\|^2}{\sigma^4} \mathcal{N}(x_i, \sigma^2 I)(x),
\end{aligned}$$

which makes the last term expressible as an expectation w.r.t a Gaussian mixture density – hence readily estimable provided samples from $\{y_i^j\}_j \sim \mathcal{N}(x_i, \sigma^2 I)$ for $x_i \sim q_t$, as done in (43).

The additional KL part needed for the update at T can be kernelized as

$$\int \mathcal{K}(x_T - z) \nabla \log \frac{q_T}{\nu}(x_T) q_T(x_T) dx_T = \int -\nabla_1 \mathcal{K}(x_T - z) q_T(x_T) + \mathcal{K}(x_T - z) \nabla V(x_T) q_T(x_T) dx_T,$$

followed by replacing q_T with \hat{q}_T . This is similar to what's done in SVGD [22].

Compared to Algorithm 3 that leverages push-forward maps, this particular approach, although also evolves a density over time through space, does not require NN training. Other particle methods for implementing dynamics in Wasserstein space can be found in e.g., [12]. We note that most of the literature focuses on Wasserstein Gradient Flow (à la JKO), whereas here we consider 2nd-order dynamics.

Remark 5. A rewriting of the acceleration field makes it clear what the driving forces behind the dynamics are:

$$\ddot{X}_t = \nabla_{\mathcal{W}_2} \left(\|\nabla_{\mathcal{W}_2} \text{KL}(q\|\nu)\|_q^2 + \mathbb{E}_q[J] \right) (X_t).$$

Since the squared norm of the \mathcal{W}_2 -gradient for the KL divergence is the relative Fisher Information, we expect the particles to undergo rapid speed change when either q differs significantly from ν or if the external potential J is exerting large force.

F Alternative Decompositions & Approximations with Entropic-OT

In this section, we give additional details on Section D.0.2. Our earlier reductions operate in the space of probability measures $\mathcal{P}(\mathbb{R}^d)$ endowed with a Wasserstein \mathcal{W}_2 geometry. We explore other geometries here.

As mentioned in the main text, with the grouping (37), we can afford a larger time-stepping, i.e., solve between each adjacent pair of observation time $\{t_k\}_{k=1}^K$ – this is different from the $o(h^2)$ approximation error we incur when using \mathcal{W}_2 . For reversible Brownian motion W as reference, it becomes iteratively solving between 2 observations:

$$q_{t_{i+1}} \leftarrow \arg \min_{q \in \mathcal{P}(\mathbb{R}^d)} \mathcal{W}_{2,t_{i+1}-t_i}^2(q, q_{t_i}) + 2 \underbrace{\int J_{t_{i+1}}(x)q(x)dx}_{\bar{F}(q)} \quad (44)$$

with entropy-regularized quadratic \mathcal{W}_2 distance

$$\mathcal{W}_{2,t_{i+1}-t_i}^2(q, q_{t_i}) = \min_{\pi^i \in \Pi(q_{t_i}, q)} \int \frac{1}{2(t_{i+1} - t_i)} \|x - y\|^2 d\pi^i(x, y) + \text{KL}(\pi^i \| q_{t_i} \otimes q)$$

as the distance-measuring metric and the linear likelihood term as the objective functional \bar{F} , which is arguably simpler than that based on (12)

$$q_{t+h} \leftarrow \arg \min_{q \in \mathcal{P}(\mathbb{R}^d)} \frac{1}{2h} \mathcal{W}_2^2(q, q_t) + h \cdot \underbrace{\int \frac{1}{2} \|\nabla \log \frac{q}{\nu}(x)\|^2 q(x) dx + 2 \int J_{t+h}(x)q(x) dx}_{F(q)}$$

involving $I(q; \nu)$.

Eulerian Method Using ideas from [28] gives a proximal splitting first order method in low dimension for (44). It results in a sequence of Sinkhorn-like updates on the dual potentials. But being confined to a fixed grid, such methods will run into scalability problem in high dimension.

Lagrangian Method One candidate here for implementing e.g., $t = t_i \rightarrow t_{i+1}$ update in (44) in high dimension can be: let $\hat{q}_{t_i} = \frac{1}{n} \sum_{j=1}^n \delta_{x_j^{t_i}}$ be given as n particles distributed as q_{t_i} , since the first variation of $q \mapsto \mathcal{W}_{2,t_{i+1}-t_i}^2(q, q_{t_i})$ is the optimal dual potential φ for the coupling, for iteration $k \geq 1$ perform

$$\begin{aligned} x_j^{t_{i+1}}[k+1] &\leftarrow x_j^{t_{i+1}}[k] - \eta \cdot \left(\frac{1}{t_{i+1} - t_i} \nabla \varphi^{\hat{q}_{t_{i+1}}^k}(x_j^{t_{i+1}}[k]) + 2 \nabla J_{t_{i+1}}(x_j^{t_{i+1}}[k]) \right), \quad j \in [n] \\ \hat{q}_{t_{i+1}}^k &= \frac{1}{n} \sum_{j=1}^n \delta_{x_j^{t_{i+1}}[k]} \quad \text{with } \hat{q}_{t_{i+1}}^1 = \hat{q}_{t_i}. \end{aligned} \quad (45)$$

In the above we expect $\hat{q}_{t_{i+1}}^k \rightarrow q_{t_{i+1}}$, the optimal solution of (44), as $k \rightarrow \infty$. To get $\varphi^{\hat{q}}$ needed in (45), one can solve an empirical Entropy-regularized OT problem, using its dual formulation: the

optimal dual potential $\varphi^{\hat{q}_{t_{i+1}}^k}$ for the marginal $\hat{q}_{t_{i+1}}^k$ (with the other being fixed at \hat{q}_{t_i}) solves

$$\arg \max_{\varphi, \psi} \sum_{j=1}^n \varphi(X_j) + \psi(Y_j) - (t_{i+1} - t_i) \cdot \exp\left(\frac{\varphi(X_j) + \psi(Y_j) - \frac{\|X_j - Y_j\|^2}{2}}{t_{i+1} - t_i}\right), \quad \text{for } X_j \sim \hat{q}_{t_{i+1}}^k, Y_j \sim \hat{q}_{t_i}$$

From this we take the optimal $\{\psi^*(Y_j)\}_{j=1}^n$, and set

$$\varphi^*(x) = -(t_{i+1} - t_i) \log\left(\frac{1}{n} \sum_{j=1}^n e^{\frac{\psi^*(Y_j) - \|x - Y_j\|^2/2}{t_{i+1} - t_i}}\right)$$

to be the soft c -transform, after which we can take its derivative to get (with this expression $x \in \mathbb{R}^d$ can be evaluated anywhere)

$$\nabla \varphi^{\hat{q}_{t_{i+1}}^k}(x) = \nabla \varphi^*(x) = \frac{\sum_{j=1}^n (x - Y_j) e^{\frac{\psi^*(Y_j) - \|x - Y_j\|^2/2}{t_{i+1} - t_i}}}{\sum_{j=1}^n e^{\frac{\psi^*(Y_j) - \|x - Y_j\|^2/2}{t_{i+1} - t_i}}}$$

in order to iterate the algorithm (45). It is known that the optimal potentials φ^*, ψ^* are related through soft c -transforms at optimality.

The procedure described above requires evolving n interacting particles over the iteration $k \geq 1$ and between each pair of time $(t_i, t_{i+1}) \in [T]$. Each update k requires running the Sinkhorn algorithm to find the optimal dual potential $\{\psi^*(Y_j)\}_{j=1}^n$ for the updated empirical density $\hat{q}_{t_{i+1}}^k$. Repeatedly running such subroutines may not be cheap computationally.

But as already commented in Remark 4, the more serious concern here is that decomposition (44) cannot be generalized when we don't know the transition kernel of the reference process – the cost function is fully determined by it.

Multi-marginal Schrödinger Bridge More generally, using [5, Lemma 3.4] on the $\text{KL}(\mathbb{Q}^\theta \|\mathbb{P}^{\text{ref}})$ term in the second equation of (7):

$$\arg \min_{\mathbb{Q}^\theta: \mathbb{Q}_0^\theta = \rho_0} \mathbb{E}_{\mathbb{Q}^\theta} [J(x; y)] + \text{KL}(\mathbb{Q}^\theta \|\mathbb{P}^{\text{ref}}),$$

we have that the objective over $\mathcal{P}(\mathcal{C}([0, T], \mathbb{R}^d))$ above can equivalently be written as

$$\begin{aligned} & \arg \min_{\{q_{t_i}\}_i} \sum_{i=1}^K \int J_{t_i}(x) q_{t_i}(x) dx + \sum_{i=1}^{K-1} \text{KL}(q_{t_i, t_{i+1}} \| R_{t_i, t_{i+1}}) - \sum_{i=2}^{K-1} \text{KL}(q_{t_i} \| R_{t_i}) \\ & \Leftrightarrow \arg \min_{\{q_{t_i}\}_i} \sum_{i=1}^K \int J_{t_i}(x) q_{t_i}(x) dx + \sum_{i=1}^{K-1} \text{KL}(\pi^*(q_{t_i}, q_{t_{i+1}}) \| R_{t_i, t_{i+1}} q_{t_i} \otimes q_{t_{i+1}}) + H(q_{t_i}) + H(q_{t_{i+1}}) \\ & \quad - \sum_{i=2}^{K-1} H(q_{t_i}) + \sum_{i=2}^{K-1} \mathbb{E}_{q_{t_i}} [\log \nu] \\ & \Leftrightarrow \arg \min_{\{q_{t_i}\}_i} \sum_{i=1}^K \int J_{t_i}(x) q_{t_i}(x) dx \quad (\text{assuming Brownian motion reference}) \\ & \quad + \sum_{i=1}^{K-1} \underbrace{\min_{\pi^i \in \Pi(q_{t_i}, q_{t_{i+1}})} \int \frac{1}{2(t_{i+1} - t_i)} \|y - x\|^2 d\pi^i(x, y) + \text{KL}(\pi^i \| q_{t_i} \otimes q_{t_{i+1}})}_{\text{SB}(q_{t_i}, q_{t_{i+1}})} + \sum_{i=1}^K H(q_{t_i}). \end{aligned}$$

This implies that the posterior sampling problem we consider can be exactly reduced to a multi-marginal SB over the K observation times (i.e., joint minimization over $\{q_{t_i}\}_{i=1}^K$), with additional parts of the loss involving simple functionals on a single marginal only. However, beyond the simple cases of e.g., Ornstein–Uhlenbeck or Brownian motion, it is not possible to specify the cost in $\text{SB}(q_{t_i}, q_{t_{i+1}})$ (transition probability density of the reference measure unknown generally). Such

approach, in the case of the Wiener reference, is adopted in [14, Theorem 3.1] where the authors propose a mean-field Langevin algorithm to solve this reduced problem over the joint $\mathcal{P}(\mathbb{R}^d)^K$ space, after which one samples from the conditional Brownian bridge between each neighboring pair to reconstruct the entire trajectory. Such a simulation is also challenging for the non-linear reference process that we consider.

G Pseudocode for Algorithms

The algorithm we use for learning the SDE drift $\{b_s(X_t, t)\}_{s \in [0,1]}$ from Section 2 is provided below. This is based on controlled equilibrium dynamics since it seeks to impose the marginal density $(\pi_s)_{s \in [0,1]}$ exactly.

Algorithm 1 Controlled transport from prior $\exp(-I(x))$ to posterior $\exp(-I(x) - J(x))$

- 1: **Input:** Initial samples from prior $\{X_t\}_{t \leq T} \in \mathbb{R}^{d \times T/\delta_t} \sim \pi_0 \propto \exp(-I(x))$ with $X_0 \sim \rho_0$, potential/reference drift $-\nabla V$, likelihood J over path, ensemble size K , anneal stepsize δ_s
- 2: **Input:** NN parameterizing $\{\partial b_t^s/\partial s\}_t: \phi^\theta(x_t^s, t, s): \mathbb{R}^d \times [0, T] \times [0, 1] \mapsto \mathbb{R}^d$
- 3: Set $A_0 = 0, b_0(\cdot) = -\nabla V(\cdot)$
- 4: **for** $s = 0$ **to** $s = 1$ **do**
- 5: Simulate until time T for K independent walkers:

$$X_{t+\delta_t} = X_t + \delta_t \cdot b_s(X_t, t) + \sqrt{2\delta_t} \cdot \zeta_t, \quad X_0 \sim \rho_0, \zeta \sim \mathcal{N}(0, I) \quad (46)$$

Optional: Anneal w.r.t $\pi_s(X) \propto \exp(-I(X) - sJ(X))$ using SPDE (16)-(17) for a few steps, starting from trajectories got from (46), output $\{X_t\}_{t \leq T}$ with $X_0 \sim \rho_0$

- 6: Compute quantities evaluated on the trajectory

$$\begin{aligned} \bar{J}^s &= \frac{1}{K} \sum_{k=1}^K J(\{X_t^k\}_{t \leq T}) \\ h_\theta^s(\{X_t^k\}_{t \leq T}) &= \sum_{t=0}^{T/\delta_t} -\frac{\delta_t}{2} \cdot \left(b_s(X_t^k, t) - \frac{X_{t+\delta_t}^k - X_t^k}{\delta_t} \right)^\top \phi^\theta(X_t^k, t, s) - \frac{\delta_t}{4} \nabla \cdot \phi^\theta(X_t^k, t, s) \\ \bar{h}_\theta^s &= \frac{1}{K} \sum_{k=1}^K h_\theta^s(\{X_t^k\}_{t \leq T}) \end{aligned}$$

Solve for

$$\phi^\theta \leftarrow \min_{\theta} \frac{1}{K} \sum_{k=1}^K (h_\theta^s(\{X_t^k\}_{t \leq T}) - \bar{h}_\theta^s + J(\{X_t^k\}_{t \leq T}) - \bar{J}^s)^2$$

- 7: Update the drift (can be evaluated any X_t, t)

$$b_{s+\delta_s}(\cdot) \approx b_s(\cdot) + \phi^\theta(\cdot, s) \cdot \delta s$$

- 8: Calculate

$$A_{s+\delta_s} = A_s - \delta s \cdot \bar{J}^s$$

- 9: **end for**

10: **return** $\exp(A_1) \approx Z_1/Z_0$ as normalizing constant estimator for $Q(x)$ and $b_1(\cdot)$

We use a particular implicit discretization scheme for the s variable advocated in [7] below for the non-equilibrium SPDE dynamics from Section B.2. This is an extension of Jarzynski-style result to path space – the algorithm, for simplicity, is written for TPS in this case. The additional Metropolis-Hastings adjustment can be used to eliminate the time-step error, as shown in Lemma 2.

Algorithm 2 Annealed SPDE from Brownian Bridge + MH adjustment for TPS with weighting (built on [7] with path $x(t, s)$ for $(s, t) \in (0, \infty) \times [0, T]$)

- 1: **Input:** Initial Brownian bridge path $x(t, 0) \in \mathbb{R}^{d \times T/\delta_t} \sim \pi_0$ with $x(0, 0) = A, x(T, 0) = B$, potential/reference drift $-\nabla V$, ensemble size K , anneal stepsize δ_s , physical time stepsize δ_t
- 2: Set $A_0 = 0$
- 3: **for** $s = 0$ **to** $s = 1$ **do**
- 4: Set the current potential to be $u = -s\nabla V$
- 5: Using the discretization in physical time: for A the discrete Laplacian matrix

$$A = \frac{1}{\delta_t^2} \begin{bmatrix} -2 & 1 & & & & \\ 1 & -2 & 1 & & & \\ & & & \ddots & & \\ & & & & 1 & -2 & 1 \\ & & & & & & 1 & -2 \end{bmatrix}$$

$$\frac{x(t, s)}{ds} = \frac{1}{2}Ax(t, s) - \frac{1}{2}u(x)\nabla u(x) - \frac{1}{2}\nabla(\nabla \cdot u(x)) + \sqrt{\frac{2}{\delta_t}} \frac{\partial W}{\partial s}$$

We update with discretization in algorithmic time as (for $\zeta \sim \mathcal{N}(0, I)$):

$$\underbrace{\left(I - \frac{1}{4}\delta_s A\right)}_L x(t, s + \delta_s) = \underbrace{\left(I + \frac{1}{4}\delta_s A\right)}_R x(t, s) - \underbrace{\frac{1}{2}u(x(t, s))\nabla u(x(t, s))\delta_s - \frac{1}{2}\nabla(\nabla \cdot u(x(t, s)))\delta_s}_{M(x(t, s))\delta_s} + \sqrt{\frac{2\delta_s}{\delta_t}}\zeta$$

with boundary condition $x(0, s + \delta_s) = A, x(T, s + \delta_s) = B$.

- 6: MH adjustment: compute transition density

$$q(s \rightarrow s + \delta_s) \propto \exp\left(-\frac{\delta_t}{4\delta_s} \|Lx(t, s + \delta_s) - Rx(t, s) + M(x(t, s))\delta_s\|^2\right)$$

Accept with probability (otherwise stay at $x(t, s)$, i.e., $x(t, s + \delta_s) = x(t, s)$)

$$\alpha = \min\left(1, \frac{\pi(x(t, s + \delta_s))q(x(t, s + \delta_s) \rightarrow x(t, s))}{\pi(x(t, s))q(x(t, s) \rightarrow x(t, s + \delta_s))}\right)$$

where $\pi(x(t, s)) \propto$

$$\exp\left(\underbrace{-s \cdot \sum_{t=0}^{T/\delta_t} \left(\frac{1}{4}\|-\nabla V(x(t, s))\|^2 + \frac{1}{2}\nabla \cdot -\nabla V(x(t, s))\right)}_{\tilde{J}(x)} \delta_t + \underbrace{\frac{\delta t}{4} \sum_{i=1}^d (x(t, s)_i^\top A x(t, s)_i)}_{-\tilde{I}(x)}\right)$$

for $x(0, s) = A, x(T, s) = B$. The first part is discretization of (20). Second part for (22).

- 7: Calculate weight

$$A_{s+\delta_s} = A_s - \delta_s \cdot \tilde{J}(x(t, s + \delta_s))$$

- 8: **end for**

- 9: Repeat Line 2 - Line 8 for K times to get K independent estimator $A_1^1, A_1^2, \dots, A_1^K$

- 10: **return** $\frac{1}{K} \sum_{k=1}^K \exp(A_1^k) \approx Z_1/Z_0$ as normalizing constant estimator for $Q(x)$
-

The Lagrangian method from Section 3.2.1 for approximately following the Wasserstein dynamics is parametrized by a sequence of push-forward maps, and outlined below. The randomized $I(q; \nu)$ estimator here is built with Lemma 4. Note that the algorithm only requires knowledge of $\nabla V, \nu$.

Algorithm 3 JKO with pushforward map $\nabla\phi^\theta$ for solving (11) based on the proposal (15)

- 1: **Input:** Initial samples $\{X_0^i\}_{i=1}^n \in \mathbb{R}^d \sim \rho_0$, potential/reference drift $-\nabla V$ with stationary distribution $\nu = \rho_0 = q_0$, likelihood J_t over path, stepsize h , small σ^2, m
- 2: **Input:** Input-convex NN parameterizing pushforward map $\phi^\theta(x_t, t): \mathbb{R}^d \times [0, T] \mapsto \mathbb{R}$
- 3: **for** $t = 0$ **to** $t = T$ **do**
- 4: Given current $\phi^\theta(\cdot, t)$, and samples $\{X_t^i\}_{i=1}^n \sim q_t$, let $X_{t+h}^i = \nabla\phi^\theta(X_t^i, t)$
- 5: Compute the following quantity on the n particles:

$$\mathcal{L}^\theta(\{X_{t+h}^i\}) = \frac{1}{n} \sum_{i=1}^n \left(\Delta \log \nu(X_{t+h}^i) + \frac{1}{2} \|\nabla V(X_{t+h}^i)\|^2 \right)$$

With the current samples $\{X_{t+h}^i\}_{i=1}^n$ generated from $\phi^\theta(\cdot, t)$, draw $Y_i^j \sim \mathcal{N}(X_{t+h}^i, \sigma^2 I)$ for $j \in [m]$ times and compute

$$R^\theta(X_{t+h}^i) = \frac{1}{m} \sum_{j=1}^m \frac{\|Y_i^j - X_{t+h}^i\|^2}{2\sigma^4}$$

- 6: Solve for

$$\phi^\theta(\cdot, t) \leftarrow \arg \min_{\theta} \frac{1}{2h^2} \frac{1}{n} \sum_{i=1}^n (X_{t+h}^i - X_t^i)^2 + \frac{1}{n} \sum_{i=1}^n \left[R^\theta(X_{t+h}^i) + \frac{2}{h} J_{t+h}(X_{t+h}^i) \right] + \mathcal{L}^\theta(\{X_{t+h}^i\}) \quad (47)$$

- 7: **if** $t = T - h$ **then**
- 8: Add additional term $\frac{1}{nh} \sum_{i=1}^n \log \frac{1/n}{\nu(X_T^i)}$ to the loss (47) with samples $X_T^i = \nabla\phi_{T-h}^\theta(X_{T-h}^i)$
- 9: **end if**
- 10: Update the n particles with the output ϕ_t^θ from (47):

$$X_{t+h}^i = \nabla\phi_t^\theta(X_t^i, t) \text{ so } \{X_{t+h}^i\}_{i=1}^n \sim q_{t+h}$$

- 11: **end for**
 - 12: **return** $\nabla\phi_t^\theta(\cdot)$ for all $t \in [0, T]$
-

H Numerics

To test the efficacy of the proposed schemes in practice, we conduct a series of numerical experiments. Our implementation is in Python, using the JAX library.¹

H.1 Implementation of Algorithms from Section 2

We parameterize $\frac{\partial b}{\partial s}$ as a neural net. Concretely, it is chosen to be a MLP with two hidden layers of dimension 20 and 30, and ReLU non-linearities. Recalling that at a fixed s , $\frac{\partial b}{\partial s}$ takes input (X_t, t) of size $\mathbb{R}^d \times [0, T]$, in the experiments below we train a new MLP at each iteration $s \in [0, 1]$ of the algorithm, and use an ADAM optimizer with a learning rate 10^{-3} to minimize the specified loss.

H.1.1 Brownian bridge posterior

We begin with the simple case of sampling from a Brownian bridge pinned at two points $x(0) = -2$ and $x(1) = +2$. Using Algorithm 1, we consider transport from the prior $\exp(-I(x))$, set up as a Brownian motion initialized at $x(0) = -2$, i.e., $u^{\text{ref}} = 0$. By gradually introducing the endpoint constraint, we design a sequence of measures $\pi_s(x) \propto \exp(-I(x) - sJ(x))$, where J takes the form of $\frac{1}{2\sigma^2} \|x(1) - 2\|^2$ with $\sigma = 0.1$. As s goes from 0 to 1 in 10 increments of 0.1, one gradually moves towards the desired Brownian bridge posterior.

An example of some resulting samples of the posterior Brownian bridge paths are shown in Figure 2 and compared to exact samples simulated from the Brownian bridge SDE (21) (which we treat as ground truth). In the experiment below, physical time is evolved in increments of $dt = 0.01$. At each iteration, the NN is trained for 200 steps, with $K = 500$ paths sampled.




Figure 2: Left: samples from Algorithm 1 at $s = 1$. Right: exact samples from a Brownian bridge SDE.

H.1.2 Transition path sampling

We next consider a TPS problem in a double well potential $V(x) = 5(x^2 - 1)^2$ with start and end states $A = -1$ and $B = 1$ respectively between time $t \in [0, 1]$. We conducted two sets of experiments, using two different annealing paths, both based on Algorithm 1.

In the first experiment, we initialize with samples from a double-well potential, treated as the prior $\exp(-I(x))$, and gradually add in the endpoint $B = 1$ constraint through $J(x) = \frac{1}{2\sigma^2} \|x(1) - 1\|^2$ with $\sigma = 0.1$.

In the second experiment, as suggested by the choice of I and J in (20)-(22), samples are initially drawn from a Brownian bridge (therefore the endpoints are already pinned at A, B), and we progressively add in the double-well potential to make sure the trajectory probabilities are weighted correctly. Specifically, we expect the trajectories moving fast in between the low density region between the wells to have higher probability under the posterior. Result are shown below in Figure 3.

In both cases, s is varied between $[0, 1]$ in increment of $\delta_s = 0.1$, i.e., 10 annealing steps. Physical time is evolved with stepsize $dt = 0.01$ with $T = 1$ as the trajectory length. At each iteration, the NN

¹A link to a repository can be found here: https://anonymous.4open.science/r/path_sampling-F861/README.md

is trained for 250 steps in the first experiment and 1000 in the second, with $K = 500$ as the ensemble size in the first experiment, and $K = 250$ in the second.

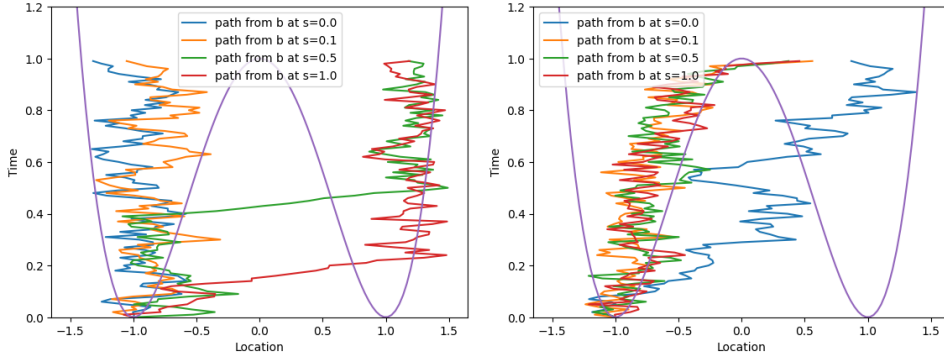


Figure 3: TPS example with paths shown at $s = 0.1$, $s = 0.5$ and $s = 1.0$. Left: evolution from the prior double-well SDE to the posterior with a (soft) constraint at $B = 1$. Right: evolution from the Brownian bridge to the posterior. Potential (purple) is scaled for convenience. MCMC is not used.

We see that when transporting from a Brownian bridge prior, even the paths at small s cross the barrier, while in the other case, the crossing of the barrier happen only when s is sufficiently large.

The left plot in Figure 3 captures both the constraint and the transition probability very well. With the right plot transporting from the Brownian bridge, we explore the option below with SPDE MCMC turned on between consecutive s updates (c.f. Figure 4).

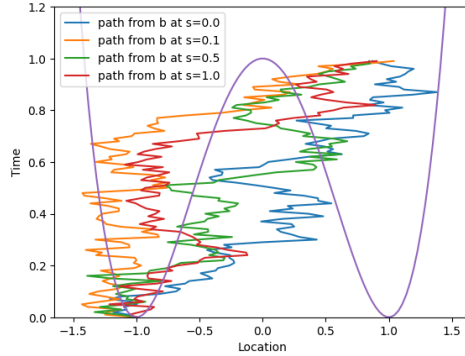


Figure 4: TPS example with paths shown at $s = 0.1$, $s = 0.5$ and $s = 1.0$. Evolution from the Brownian bridge to the posterior with MCMC in between. Potential (purple) is scaled for convenience.

H.1.3 Non-equilibrium dynamics

For the same TPS example as above, we compare what we observe with non-equilibrium dynamics based on Algorithm 2. Here the switching is set to happen very slowly with $s \in [0, 1]$ in 100 steps (i.e., $\delta_s = 0.01$). MH adjustment is applied on a discretization of the SPDE after each step of the s update. The annealing path is based on transporting the Brownian Bridge to a pinned double-well posterior at $A = -1, B = 1$, with ensemble size $K = 100$, stepsize $\delta_t = 10^{-3}$.

We show in Figure 5 the path switch as $s \in [0, 1]$, as well as a final set of trajectories at $s = 1$. We see quite a lot of variability in the final trajectories, which is an indication of the sampler’s performance – high variance reflects that the trajectories are far from equilibrium and dominated by rare events.

H.2 Implementation of Algorithm from Section 3

We set up the experiment below so ν corresponds to the stationary distribution of an Ornstein–Uhlenbeck process

$$dX_t = -\beta X_t dt + \sqrt{2} dW_t$$




Figure 5: Left: One of the trajectories in the ensemble, at each iteration of the algorithm, going from red ($s = 0$) to blue ($s = 1$). Right: 5 trajectories from the ensemble at $s = 1$.

with $\beta = 1/4$ – this is considered to be the prior reference process in this example. Moreover, we initialize the $n = 200$ particles as $X_0 \sim \mathcal{N}(0, 1/\beta)$ at stationarity. The likelihood J_t is chosen so that we impose a soft constraint in the form of $\frac{1}{2\gamma^2} \|x(t) - x_{\text{fixed}}\|^2$ at the midpoint with a big γ , and the same constraint near the endpoint with a small γ , at two different pre-specified locations x_{fixed} .

In the experiment below (Figure 6), stepsize $h = 0.2$, length of trajectory $T = 1$ and the perturbation variance $\sigma^2 = 0.4^2$ with #perturbations $m = 30$ for the $I(q; \nu)$ estimator. This general Bayesian posterior path example using Algorithm 3 is conducted with a 7-layer MLP parameterizing the pushforward map $\nabla\phi^\theta$. With $x_{\text{fixed}} = -1$ at the midpoint, and $x_{\text{fixed}} = +1$ near the endpoint, the likelihood in this case has the effect of slightly nudging trajectories towards -1 in the middle of the path and sharply pin them at $+1$ at the terminal boundary. In between the two observations, we expect the particles to roughly have a Gaussian distribution not too far from stationary.




Figure 6: Density evolution of trajectories (reference is set to be an OU process in equilibrium).

I Conclusion

In this work, we proposed two classes of new methodologies for bringing NN training to the problem of posterior path sampling. The first is based on a *controlled* dynamics over path space in order to progressively learn the optimal b_t^* vector field in a memory-efficient way (Sec 2), and is guaranteed to hit target Q at *finite* time without additional reweighting, or solving high-dimensional PDEs. The

second is based on convex optimization in Wasserstein space and OT, where we gradually evolve a curve of marginal densities over time. We propose grid-free stochastic methods, based either on push-forward maps (Sec 3.2.1), or an interacting particle system (Sec D.0.1), that can avoid the curse of dimensionality, while maintaining generality and good approximation properties (Sec D.0.2).

NeurIPS Paper Checklist

1. Claims

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

Answer: [Yes]

Justification: We summarize our main contributions in the first 2 pages.

Guidelines:

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

2. Limitations

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

Answer: [Yes]

Justification: We discuss alternatives to our proposals and compare their relative advantages / disadvantages.

Guidelines:

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

3. Theory assumptions and proofs

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

Answer: [Yes]

Justification: We provide full proofs to our theoretical results in the Appendix.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- 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.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

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)?

Answer: [Yes]

Justification: These are included in the Appendix.

Guidelines:

- The answer NA means that the paper does not include experiments.
- 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.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- 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.
- 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
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (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).
 - (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.

5. **Open access to data and code**

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We provide a link to the code used for our experiments.

Guidelines:

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

6. Experimental setting/details

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

Answer: [Yes]

Justification: These are included in the experiment section of the appendix.

Guidelines:

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

7. Experiment statistical significance

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

Answer: [NA]

Justification: We don't observe large variability across our runs, and therefore don't find error bars particularly meaning to report.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).

- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

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

Answer: [Yes]

Justification: These are included in the experiment section of the appendix.

Guidelines:

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

9. Code of ethics

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

Answer: [Yes]

Justification: We followed the code of ethics while conducting this research.

Guidelines:

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

10. Broader impacts

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

Answer: [NA]

Justification: We don't anticipate significant societal impacts beyond intellectual inquiry.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

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

Answer: [NA]

Justification: Our paper poses no such risks.

Guidelines:

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

12. Licenses for existing assets

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

Answer: [NA]

Justification: Our paper does not use existing assets.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset’s creators.

13. **New assets**

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

Answer: [NA]

Justification: Our paper does not release new assets.

Guidelines:

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

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

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

Answer: [NA]

Justification: We do not use crowd-sourcing or research with human subjects.

Guidelines:

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

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

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

Answer: [NA]

Justification: We do not use crowd-sourcing or research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

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

Answer: [NA]

Justification: We have not used LLM in the research reported in this paper.

Guidelines:

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