
Control, Transport and Sampling: The Benefit of a Reference Process

Qijia Jiang David Nabergoj

Abstract

Leveraging connections between diffusion-based sampling, optimal transport, and optimal stochastic control through their shared links to the Schrödinger bridge problem, we propose novel objective functions that can be used to transport ν to μ , consequently sample from the target μ , via (optimally) controlled dynamics. Throughout, we highlight the importance of having a reference measure on the path space for the design of valid training losses.

1. Introduction

Traditionally, the task of sampling from un-normalized density is largely delegated to MCMC methods. However, modern machine learning developments in optimal transport and generative modeling have greatly expanded the toolbox we have available for performing such tasks, which can further benefit from powerful advancements in deep learning. By reducing the problem to performing ERM with neural networks, they hold promise especially for sampling from high-dimensional and multimodal distributions compared to MCMC-based alternatives. In this work, we propose control objectives that are amenable to tractable approximation (without access to data from the target, as classically considered in the MCMC literature) for sampling purpose, which offer numerical advantages compared to several existing methods. The central object for this synthesis will be forward-backward SDEs via time-reversals.

2. General Framework

In this section we put several recently proposed (and diffusion-inspired) methods in context, rendering them as special instantiations of a more unifying *path-wise* picture. We loosely follow the framework put forth in (Vargas & Nüsken, 2023) in the exposition.

2.1. Goal and Setup

We are interested in sampling from $p_{\text{target}}(x) = \mu(x)/Z$ by minimizing certain tractable loss, assuming sampling from $\nu(z)$ is easy, but unlike in generative modeling, even though analytical expression for μ is readily available, we do not

have access to data from it that can be used to learn the score function. In this sense, in terms of “transport mapping”, the two sides are crucially not symmetric. More specifically, the following represents a coupling between $p_{\text{prior}}(z) = \nu(z)$ and $p_{\text{target}}(x)$, our desired target distribution, where the transitions are parametrized by γ and θ :

$$\pi(x, z) := q^\gamma(z|x)p_{\text{target}}(x) = p^\theta(x|z)p_{\text{prior}}(z).$$

Minimizing objective of the type (D_{KL} below denotes KL divergence)

$$\mathcal{L}(\gamma, \theta) := D_{KL}(q^\gamma(z|x)p_{\text{target}}(x) || p^\theta(x|z)p_{\text{prior}}(z))$$

to 0 will ensure the marginals are satisfied, although the minimizers of the problem will not be unique. More generally, one can make the transitions more hierarchical (yet still Markovian), and for this we introduce forward-backward SDEs. Pictorially, given a base drift f , we have the “sampling process” (the two processes are time reversals of each other):

$$\nu(z) \begin{array}{c} \xrightarrow{\mathbb{P}^{\nu, f + \sigma u}} \\ \xleftarrow{\mathbb{P}^{\mu, f + \sigma v}} \end{array} \mu(x) \quad (1)$$

for tunable controls u, v and terminal marginals ν, μ . Written as SDEs, they become (the drifts u, v here are not independent)

$$dX_t = [f_t(X_t) + \sigma u_t(X_t)]dt + \sigma d\overrightarrow{W}_t, \quad X_0 \sim \nu \quad (2)$$

$$\Rightarrow (X_t)_{t \in [0, T]} \sim \overrightarrow{\mathbb{P}}^{\nu, f + \sigma u},$$

$$dX_t = [f_t(X_t) + \sigma v_t(X_t)]dt + \sigma d\overleftarrow{W}_t, \quad X_T \sim \mu \quad (3)$$

$$\Rightarrow (X_t)_{t \in [0, T]} \sim \overleftarrow{\mathbb{P}}^{\mu, f + \sigma v}.$$

Operationally, (2)-(3) denote (picking $f = 0$, and $\{z_t\}$ a sequence of i.i.d standard Gaussians for illustration)

$$X_t = X_0 + \int_0^t \sigma u_s(X_s)ds + \int_0^t \sigma d\overrightarrow{W}_s \quad (4)$$

$$\Rightarrow X_{t+h} \approx X_t + h\sigma u_t(X_t) + \sqrt{h}\sigma z_t, \quad X_0 \sim \nu$$

$$X_t = X_T - \int_t^T \sigma v_s(X_s)ds - \int_t^T \sigma d\overleftarrow{W}_s \quad (5)$$

$$\Rightarrow X_{t-h} \approx X_t + h\sigma v_t(X_t) + \sqrt{h}\sigma z_t, \quad X_T \sim \mu$$

where the forward (the usual Itô’s) and the backward integrals indicate different endpoints at which we make the

approximation. For any process $(Y_t)_t, (Z_t)_t$,

$$\int_0^T a_t(Y_t) d\overrightarrow{Z}_t \approx \sum_i a_{t_i}(Y_{t_i})(Z_{t_{i+1}} - Z_{t_i}), \quad (6)$$

$$\int_0^T a_t(Y_t) d\overleftarrow{Z}_t \approx \sum_i a_{t_{i+1}}(Y_{t_{i+1}})(Z_{t_{i+1}} - Z_{t_i}), \quad (7)$$

which in particular implies the martingale property

$$\mathbb{E}_{X \sim \overrightarrow{\mathbb{P}}^{\nu, f+\sigma u}} \left[\int_0^t a_s(X_s) d\overrightarrow{W}_s \right] = 0, \quad (8)$$

$$\mathbb{E}_{X \sim \overleftarrow{\mathbb{P}}^{\mu, f+\sigma v}} \left[\int_t^T a_s(X_s) d\overleftarrow{W}_s \right] = 0. \quad (9)$$

Note one can also view the backward process (3) as the time reversal of the following forward SDE with the additional score function: ($\hat{\rho}_t$ is the density of $\overleftarrow{\mathbb{P}}^{\mu, f+\sigma v}$ at time t)

$$d\hat{X}_t = (-\hat{f}_t(\hat{X}_t) - \sigma \hat{v}_t(\hat{X}_t) - \sigma^2 \nabla \log \hat{\rho}_t(\hat{X}_t)) dt + \sigma d\overrightarrow{W}_t, \quad (10)$$

for $\hat{X}_0 \sim \mu$, with the hat denoting $t \rightarrow T - t$. Now time reversing while keeping the same path measure will give the following SDE:

$$dX_t = [f_t(X_t) + \sigma v_t(X_t)] dt + \sigma d\overleftarrow{W}_t, \quad X_T \sim \mu.$$

Therefore both (3) and (10) have the same path measure $\overleftarrow{\mathbb{P}}^{\mu, f+\sigma v}$. Forward/backward integral can be converted through

$$\int_0^T a_t(Y_t) d\overrightarrow{Z}_t = \int_0^T a_t(Y_t) d\overleftarrow{Z}_t - \sigma^2 \int_0^T (\nabla \cdot a_t)(Y_t) dt, \quad (11)$$

which will be used repeatedly throughout.

Remark 2.1 (Nelson's identity). The following relationship between time-reversal drifts for the SDE (2)-(3) is well known: $\overrightarrow{\mathbb{P}}^{\nu, f+\sigma u} = \overleftarrow{\mathbb{P}}^{\mu, f+\sigma v}$ iff $\overrightarrow{\mathbb{P}}_T^{\nu, f+\sigma u} = \mu$, and $\sigma v_t(X_t) = \sigma u_t(X_t) - \sigma^2 \nabla \log(\overrightarrow{\mathbb{P}}_t^{\nu, f+\sigma u}) = \sigma u_t(X_t) - \sigma^2 \nabla \log(\overleftarrow{\mathbb{P}}_t^{\mu, f+\sigma v})$ for all $t \in [0, T]$.

2.2. Related Approaches

We would like the two path measures (with the specified two end-point marginals) to agree progressing in either direction. The methods in (Vargas et al., 2022; Zhang & Chen, 2021) propose to set up a reference process with a similar structure as (1) with $f = 0$:

$$\nu(z) \xrightleftharpoons[\mathbb{P}^{\eta, \sigma v}]{\mathbb{P}^{\nu, \sigma r}} \eta(x) \quad (12)$$

where r is the drift for the reference process:

$$dX_t = \sigma r_t(X_t) dt + \sigma d\overrightarrow{W}_t, \quad X_0 \sim \nu \Rightarrow (X_t)_{t \in [0, T]} \sim \overrightarrow{\mathbb{P}}^{\nu, \sigma r} \quad (13)$$

and $\eta = \overrightarrow{\mathbb{P}}_T^{\nu, \sigma r}$. The part worth attention is that we fix the backward process such that $\overleftarrow{\mathbb{P}}^{\mu, \sigma v}$ and $\overleftarrow{\mathbb{P}}^{\eta, \sigma v}$ differ only in initialization with the same ensuing dynamics driven by v . The method amounts to loss of the following type:

$$\begin{aligned} \mathcal{L}_{KL}(u) &= \mathbb{E}_{\overrightarrow{\mathbb{P}}^{\nu, \sigma u}} \left[\log \left(\frac{d\overrightarrow{\mathbb{P}}^{\nu, \sigma u}}{d\overleftarrow{\mathbb{P}}^{\mu, \sigma v}} \right) \right] \\ &= \mathbb{E}_{\overrightarrow{\mathbb{P}}^{\nu, \sigma u}} \left[\log \left(\frac{d\overrightarrow{\mathbb{P}}^{\nu, \sigma u}}{d\overrightarrow{\mathbb{P}}^{\nu, \sigma r}} \frac{d\overleftarrow{\mathbb{P}}^{\eta, \sigma v}}{d\overleftarrow{\mathbb{P}}^{\mu, \sigma v}} \right) \right] \\ &= \mathbb{E}_{\overrightarrow{\mathbb{P}}^{\nu, \sigma u}} \left[\log \left(\frac{d\overrightarrow{\mathbb{P}}^{\nu, \sigma u}}{d\overrightarrow{\mathbb{P}}^{\nu, \sigma r}} \frac{d\eta}{d\mu} \right) + \log Z \right] \\ &= \mathbb{E}_{X \sim \overrightarrow{\mathbb{P}}^{\nu, \sigma u}} \left[\int_0^T \frac{1}{2} \|u_s(X_s) - r_s(X_s)\|^2 ds + \log \left(\frac{d\eta}{d\mu} \right) (X_T) + \log Z \right]. \end{aligned} \quad (14)$$

This suggests initializing from ν to estimate the loss (14) with the current control $u^{\hat{\theta}}$ by simulating (2), followed by gradient descent to optimize the $\hat{\theta}$ -parameterized control policy and iterating between the two steps can be a viable strategy, that could identify u^{θ^*} eventually and used to run (2) to draw samples from μ . Note η, ν are simple distributions we have the freedom to pick, and the corresponding $r(\cdot)$ as well in (12). The last step of (14) is a simple application of Girsanov's theorem for two SDEs with different drifts. This is highly suggestive of a control-theoretic interpretation (Berner et al., 2022), where the cost (14) is composed of a running control cost and a terminal cost.

Common choice for ν include δ_0 a fixed Dirac delta (this is technically speaking a half-bridge) or $\mathcal{N}(0, \sigma^2 I)$. The r -controlled reference process (13) can be either Brownian motion ($r_t = 0$) or an OU process ($r_t(x) = -\frac{x}{2\sigma}$). Both combinations will give η as normally distributed. In the former case, the optimal drift u^* with $f = 0$ is known to be the Föllmer drift (Tzen & Raginsky, 2019) and can be computed explicitly as (picking $T = 1$ for simplicity below)

$$\begin{aligned} u_t^*(x) &= \nabla \log \mathbb{E}_X \left[\frac{d\mu}{d\eta}(X_1) | X_t = x \right] \\ &= \nabla \log \mathbb{E}_{z \sim \mathcal{N}(0, \sigma^2 I)} \left[\frac{d\mu}{d\mathcal{N}(0, \sigma^2 I)}(x + \sqrt{1-t}z) \right], \end{aligned}$$

where the first expectation above is taken w.r.t the reference measure – Wiener process in this case; and it follows that $v_t^*(x) = x/t$ by Doob's h -transform. In the latter case, process (12) is simply in equilibrium. The purpose of introducing the reference process (12) is to cancel out terms that we don't have convenient estimate for when calculating the likelihood ratio, while fixing the drift v in (1) so that $\min_u \mathcal{L}_{KL}(u)$ from (14) is unique.

Remark 2.2. The approach in (Chen et al., 2021a) that relies on training two drifts essentially amounts to solving

$\min_{u,v} D_{KL}(\overrightarrow{\mathbb{P}}^{\nu, f+\sigma u} \parallel \overleftarrow{\mathbb{P}}^{\mu, f+\sigma v})$ jointly with (63)-(64), and as shown in (Richter et al., 2023), is non-unique.

In diffusion generative modeling applications, score-matching-based loss (Hyvärinen & Dayan, 2005) can be seen as minimizing the reverse KL over $s = (u - v)/\sigma$ by running the backward process (3) and using the Radon-Nikodym derivative in (60)-(61) with $f = 0, r^+ = r^- = 0$. It gives (C is independent of s below)

$$\begin{aligned} \mathcal{L}_{KL}(s) &= \mathbb{E}_{\overleftarrow{\mathbb{P}}^{\mu, \sigma v}} \left[\log \left(\frac{d\overrightarrow{\mathbb{P}}^{\mu, \sigma v}}{d\overleftarrow{\mathbb{P}}^{\nu, \sigma v + \sigma^2 s}} \right) \right] \\ &= \mathbb{E}_{\overleftarrow{\mathbb{P}}} \left[\int_0^T \frac{\sigma^2}{2} \|s_t(X_t)\|^2 dt + \sigma^2 \int_0^T \nabla \cdot s_t(X_t) dt \right] + C \\ &= \mathbb{E}_{\overleftarrow{\mathbb{P}}} \left[\int_0^T \frac{\sigma^2}{2} \|s_t(X_t) - \nabla_x \log p_t^{\mu, \sigma v}(X_t | X_T = x_T)\|^2 dt \right] \end{aligned} \quad (15)$$

up to a constant $+C'$, while fixing the backward drift v for which ν is easy to sample from (e.g., an OU process) in (1). In the last transition above, we used the integration by parts identity

$$\mathbb{E}_{\rho_t} \left[\int_0^T s_t^\top \nabla \log \rho_t dt \right] = -\mathbb{E}_{\rho_t} \left[\int_0^T \nabla \cdot s_t dt \right]. \quad (16)$$

In practice there will be an ir-reducible loss since the terminals ν and $\overleftarrow{\mathbb{P}}_0^{\mu, \sigma v}$ don't match exactly for any finite T , but the dynamics $\overrightarrow{\mathbb{P}}^{\nu, \sigma v + \sigma^2 s}$ still makes sense as for two processes $(p_t)_t, (q_t)_t$ with different initialization, but same drift $\sigma v + \sigma^2 s^*$, $\partial_t D_{KL}(p_t \| q_t) = -\mathbb{E}_{p_t} [\|\nabla \log \frac{p_t}{q_t}\|^2] \leq 0$ contracts, although we have not been quantitative about the decay rate.

We therefore see that either the approach of (14) or (15) relies on fixing some aspect of the process involving (1) to restore uniqueness of the loss \mathcal{L}_{KL} . In both cases, it results in an one-parameter loss, and the minimizer is not affected by the unknown constant Z from the loss. Successfully optimizing $D_{KL}(\overrightarrow{\mathbb{P}}^{\nu, \sigma u} \parallel \overleftarrow{\mathbb{P}}^{\mu, \sigma v})$ or $D_{KL}(\overleftarrow{\mathbb{P}}^{\mu, \sigma v} \parallel \overrightarrow{\mathbb{P}}^{\nu, \sigma u})$ to zero error will imply that u pushes ν to μ and v vice versa, mimicking a noising/denoising reversible procedure.

3. Methodology

As we saw from Section 2, there are many degrees of freedom in transporting ν to μ , which is not desirable for training purpose. This gives us the motivation to turn to losses that can enforce a canonical choice. We will adopt an ‘‘optimal control’’ perspective (which necessitates picking some reference) and leverage special properties of the SB problem (introduced below) to come up with valid control objectives for this effort.

Consider over path space $C([0, T]; \mathbb{R}^d)$, given a base measure Q , the constrained optimization problem

$$P^* = \arg \min_{P_0=\nu, P_T=\mu} D_{KL}(P \| Q) \quad (17)$$

where we assume Q admits the SDE representation

$$dX_t = f_t(X_t)dt + \sigma dW_t, \quad X_0 \sim \nu, \quad (18)$$

which is a slight generalization of the classical case where $f = 0$ is typically assumed. Additional perspectives on the Schrödinger Bridge problem (17)-(18), which we heavily leverage in the next part for designing our losses, are included in Appendix B, along with related methods for solving the problem in different setup than what we consider in Appendix A.

3.1. Training Loss Proposal

Recall our focus is on solving the regularized optimal transport problem (17) between ν and μ , by learning the controls on the basis of samples from ν . Out of many joint couplings with correct marginals (i.e., transport maps), the choice of a reference path process will select a particular trajectory ρ_t between ν and μ .

The execution of this plan crucially hinges on two ingredients: (1) backward / forward change of measure formula established in Lemma C.1; (2) properties of the optimal drifts for SB from Section B, which in turn can be exploited for training the controls. Challenge, as emphasized before, is we need to be able to estimate the resulting loss and optimizing it guarantees convergence to the unique solution dictated by the SB.

Proposition 3.1 below serves as our main result. Both (1) and (4) are guided by ‘‘minimum-action’’ principle w.r.t a reference. (2) bases itself on a reformulation of the HJB PDE involving the optimal control, and (3) is grounded in the FBSDE system for SB optimality (Chen et al., 2021a). In all cases, the objective is a two-parameter loss that also allows us to recover the score $\sigma^2 \nabla \log \rho_t$ from the returned solution. Note that variance is taken w.r.t the uncontrolled process in (2) and w.r.t the controlled process in (3). Losses (1) and (4), being conversion from the constrained problem (and each other), will need λ to be picked relatively small so that the first part of the objective becomes 0 to identify the unique SB solution. All proofs are deferred to Appendix C.

Proposition 3.1 (Regularized Control Objective). *For the problem of (17), the following losses are justified:*

$$(1) \arg \min_{\nabla \phi, \nabla \psi} D_{KL}(\overrightarrow{\mathbb{P}}^{\nu, f+\sigma^2 \nabla \phi} \parallel \overleftarrow{\mathbb{P}}^{\mu, f-\sigma^2 \nabla \psi}) + \lambda \cdot \mathbb{E}_{X \sim \overrightarrow{\mathbb{P}}^{\nu, f+\sigma^2 \nabla \phi}} \left[\int_0^T \frac{\sigma^2}{2} \|\nabla \phi_t(X_t)\|^2 dt \right]$$

$$(2) \arg \min_{\nabla \phi, \nabla \psi} D_{KL}(\overrightarrow{\mathbb{P}}^{\nu, f+\sigma^2 \nabla \phi} \parallel \overleftarrow{\mathbb{P}}^{\mu, f-\sigma^2 \nabla \psi}) +$$

$$\text{Var}_{X \sim \overrightarrow{\mathbb{P}}^{\nu, f}} \left(\psi_T(X_T) - \psi_0(X_0) + \int_0^T \left(-\frac{\sigma^2}{2} \|\nabla \psi_t\|^2 \right) dt \right)$$

$$+ \nabla \cdot f_t - \sigma^2 \Delta \psi_t)(X_t) dt - \sigma \int_0^T \nabla \psi_t(X_t)^\top dW_t)$$

$$\begin{aligned} \text{or } \text{Var}_{X \sim \overrightarrow{\mathbb{P}}^{\nu, f}} \left(\phi_T(X_T) - \phi_0(X_0) + \frac{\sigma^2}{2} \int_0^T \|\nabla \phi_t\|^2(X_t) dt - \sigma \int_0^T \nabla \phi_t(X_t)^\top dW_t \right) \\ (3) \arg \min_{\phi, \psi} \text{Var}_{X \sim \overrightarrow{\mathbb{P}}^{\nu, f + \sigma^2 \nabla \phi}} \left((\phi_T + \psi_T - \log \mu)(X_T) \right) \\ + \text{Var}_{X \sim \overrightarrow{\mathbb{P}}^{\nu, f + \sigma^2 \nabla \phi}} \left((\phi_0 + \psi_0 - \log \nu)(X_0) \right) + \\ \text{Var}_{X \sim \overrightarrow{\mathbb{P}}^{\nu, f + \sigma^2 \nabla \phi}} \left(\phi_T(X_T) - \phi_0(X_0) - \frac{\sigma^2}{2} \int_0^T \|\nabla \phi_t\|^2(X_t) dt - \sigma \int_0^T \nabla \phi_t(X_t)^\top dW_t \right) + \\ \text{Var}_{X \sim \overrightarrow{\mathbb{P}}^{\nu, f + \sigma^2 \nabla \phi}} \left(\psi_T(X_T) - \psi_0(X_0) - \int_0^T \left(\frac{\sigma^2}{2} \|\nabla \psi_t\|^2 + \nabla \cdot (\sigma^2 \nabla \psi_t - f_t) + \sigma^2 \nabla \psi_t^\top \nabla \phi_t \right)(X_t) dt \right. \\ \left. - \sigma \int_0^T \nabla \psi_t(X_t)^\top dW_t \right) \end{aligned}$$

$$(4) \arg \min_{\nabla \phi_t, \nabla \log \rho_t} \lambda \mathbb{E}_{\overrightarrow{\mathbb{P}}^{\nu, f + \sigma^2 \nabla \phi}} \left[\int_0^T \frac{\sigma^2}{2} \|\nabla \phi_t(X_t)\|^2 dt \right] \\ + D_{KL}(\overrightarrow{\mathbb{P}}^{\nu, f + \sigma^2 \nabla \phi} \parallel \overleftarrow{\mathbb{P}}^{\mu, f + \sigma^2 \nabla \phi_t - \sigma^2 \nabla \log \rho_t})$$

In the above, $D_{KL}(\overrightarrow{\mathbb{P}}^{\nu, f + \sigma^2 \nabla \phi} \parallel \overleftarrow{\mathbb{P}}^{\mu, f - \sigma^2 \nabla \psi})$ is

$$\mathbb{E}_{X \sim \overrightarrow{\mathbb{P}}^{\nu, f + \sigma^2 \nabla \phi}} \left[\int_0^T \frac{\sigma^2}{2} \|\nabla \phi_t(X_t) + \nabla \psi_t(X_t)\|^2 - \nabla \cdot (f_t - \sigma^2 \nabla \psi_t)(X_t) dt + \log \frac{\nu(X_0)}{\mu(X_T)} \right] + C \quad (19)$$

and $D_{KL}(\overrightarrow{\mathbb{P}}^{\nu, f + \sigma^2 \nabla \phi_t} \parallel \overleftarrow{\mathbb{P}}^{\mu, f + \sigma^2 \nabla \phi_t - \sigma^2 \nabla \log \rho_t})$ is

$$\mathbb{E}_{\overrightarrow{\mathbb{P}}^{\nu, f + \sigma^2 \nabla \phi}} \left[\log \frac{\nu(X_0)}{\mu(X_T)} + \int_0^T \frac{\sigma^2}{2} \|\nabla \log \rho_t(X_t)\|^2 + \sigma^2 \nabla \cdot (\nabla \log \rho_t - \nabla \phi_t)(X_t) - \nabla \cdot f_t(X_t) dt \right] + C' \quad (20)$$

In all cases, $X_0 = x_0 \sim \nu$ is assumed given as initial condition. In particular, if the loss is 0 for (2) and (3), the resulting $\nabla \phi^*$, $\nabla \psi^*$ solve the Schrödinger Bridge problem from ν to μ , that can in turn be used for sampling from μ .

It is tempting to draw a connection between the variance penalty in Proposition 3.1 and the log-variance divergence between two path measures w.r.t a reference

$$\text{Var}_{\mathbb{P}^w} \left[\log \left(\frac{d\mathbb{P}^u}{d\mathbb{P}^v} \right) \right]$$

studied in (Richter et al., 2023). The loss in (3) is different from log-variance divergence over path space

$$\text{Var}_{\overrightarrow{\mathbb{P}}^{\nu, f + \sigma^2 \nabla \phi}} \left[\log \left(\frac{d\overrightarrow{\mathbb{P}}^{\nu, f + \sigma^2 \nabla \phi}}{d\overleftarrow{\mathbb{P}}^{\mu, f - \sigma^2 \nabla \psi}} \right) \right],$$

which involve the sum of (67) and (68), and will not guarantee finding the optimal path, whereas the objective in (3) incorporates the dynamics of *two* controlled (and coupled) dynamics w.r.t $\overrightarrow{\mathbb{P}}^{\nu, f}$ that we know how to characterize optimality for from results in Section B.

In fact, another way to view the variance regularizers in loss (2) and (3) is through the SDE representation of the controls from Lemma 3.2 below and observe that the variance condition precisely encodes the optimally-controlled dynamical information.

Lemma 3.2 (SDE correspondence to SB optimality). *We have for the optimal forward drift $\nabla \phi_t$ and $X \sim \overrightarrow{\mathbb{P}}^{\nu, f}$ as in (18),*

$$d\phi_t(X_t) = -\frac{\sigma^2}{2} \|\nabla \phi_t(X_t)\|^2 dt + \sigma \nabla \phi_t(X_t)^\top dW_t,$$

analogously for the optimal backward drift $-\nabla \psi_t$, along $X \sim \overrightarrow{\mathbb{P}}^{\nu, f}$,

$$d\psi_t(X_t) = \left[\frac{\sigma^2}{2} \|\nabla \psi_t\|^2 - \nabla \cdot f_t + \sigma^2 \Delta \psi_t \right](X_t) dt + \sigma \nabla \psi_t(X_t)^\top dW_t.$$

Moreover, along the controlled forward dynamics $X \sim \overrightarrow{\mathbb{P}}^{\nu, f + \sigma^2 \nabla \phi}$, the optimal control ϕ, ψ satisfy

$$d\phi_t(X_t) = \frac{\sigma^2}{2} \|\nabla \phi_t(X_t)\|^2 dt + \sigma \nabla \phi_t(X_t)^\top dW_t,$$

$$d\psi_t(X_t) = \left[\frac{\sigma^2}{2} \|\nabla \psi_t\|^2 + \nabla \cdot (\sigma^2 \nabla \psi_t - f_t) + \sigma^2 \nabla \phi_t^\top \nabla \psi_t \right](X_t) dt + \sigma \nabla \psi_t(X_t)^\top dW_t.$$

In the above, $\nabla \phi_t, \nabla \psi_t$ refer to the optimal forward / backward drift in the SDE

$$dX_t = [f_t(X_t) + \sigma^2 \nabla \phi_t(X_t)] dt + \sigma dW_t, X_0 \sim \nu,$$

$$dX_t = [f_t(X_t) - \sigma^2 \nabla \psi_t(X_t)] dt + \sigma dW_t, X_T \sim \mu,$$

and $\phi_T(X_T) + \psi_T(X_T) = \log p_{\text{target}}(X_T)$, $\phi_0(X_0) + \psi_0(X_0) = \log \nu(X_0)$ for $X \sim \overrightarrow{\mathbb{P}}^{\nu, f + \sigma^2 \nabla \phi}$.

It is natural to ask if one could replace the variance regularizer $\text{Var}(\cdot)$ with a moment regularizer $\mathbb{E}[|\cdot|^2]$ in e.g., loss (3). However, while the variance is oblivious to constant shift, the moment loss will require knowledge of the normalizing constant Z of the target μ to make sense. We would also like to add that enforcing the drifts to take the gradient form is not strictly necessary for some parts of the losses.

Remark 3.3. The gradient for either D_{KL} or Log-variance w.r.t. u and v (and the Monte-Carlo estimate thereof) can be found in e.g., (Richter et al., 2023), where the authors show that log-variance has the additional benefit of having variance of gradient = 0 at the optimal u^*, v^* , which is not true for D_{KL} in general, and has consequence for gradient-based

updates such as Algorithm 1. Similar argument applies to the variance regularizer we consider (e.g., loss (2)). For this, we look at the Gâteaux derivative of the variance function V in an arbitrary direction τ since from the chain rule

$$\begin{aligned} \frac{\delta}{\delta u} V(u, v; \tau) &:= \left. \frac{d}{d\epsilon} \right|_{\epsilon=0} V(u + \epsilon\tau, v) \Rightarrow \\ \partial_{\theta_i} V(u_{\theta}, v_{\gamma}) &= \left. \frac{\delta}{\delta u} \right|_{u=u_{\theta}} V(u, v_{\gamma}; \partial_{\theta_i} u_{\theta}). \end{aligned}$$

Now if \hat{V} is the Monte-Carlo estimate of the variance of a random quantity g , it is always the case that (we use $\frac{\delta}{\delta u}(\cdot)_{\tau}$ to denote derivative in the τ direction)

$$\begin{aligned} \frac{\delta}{\delta u} \hat{V}(u, v; \tau) &= \frac{\delta}{\delta u} (\hat{\mathbb{E}}[g(u, v)^2] - \hat{\mathbb{E}}[g(u, v)]^2)_{\tau} \\ &= 2\hat{\mathbb{E}}[g(u, v) \frac{\delta}{\delta u} g(u, v)_{\tau}] - 2\hat{\mathbb{E}}[g(u, v)] \frac{\delta}{\delta u} \hat{\mathbb{E}}[g(u, v)]_{\tau}, \end{aligned}$$

therefore if $g(u^*, v^*) = c$ almost surely for every i.i.d sample, such as the regularizer in loss (2), the derivative w.r.t the control u in direction $\partial_{\theta_i} u_{\theta}$ is 0 at optimality, implying $\text{Var}(\partial_{\theta_i} \hat{V}(u_{\theta^*}, v_{\gamma^*})) = 0$. One can, of course, also replace D_{KL} with log-variance for the first part of the loss.

In Appendix E, we also consider the extension of such training methodologies to the case of second-order dynamics in the augmented (X, V) space à la underdamped Langevin.

3.2. Discretization and Practical Implementation

We discuss practical implementation of our losses and estimators in this part. As in Proposition F.1 of (Vargas & Nüsken, 2023), it is possible to trade the divergence term for a backward integral and estimate $D_{KL}(\overrightarrow{\mathbb{P}}^{\nu, u} \| \overleftarrow{\mathbb{P}}^{\mu, v})$ (up to constant independent of u, v) as

$$\begin{aligned} \frac{1}{N} \sum_{i=1}^N \left[\log \frac{\nu(X_0^i)}{\mu(X_{K+1}^i)} + \sum_{k=0}^{K-1} \frac{1}{2\sigma^2(t_{k+1} - t_k)} \right. \\ \left. \|X_k^i - X_{k+1}^i + v_{k+1}(X_{k+1}^i)(t_{k+1} - t_k)\|^2 \right] \quad (21) \end{aligned}$$

for $z_k^i \sim \mathcal{N}(0, I)$ and

$$X_{k+1}^i = X_k^i + u_k(X_k^i)(t_{k+1} - t_k) + \sigma \sqrt{t_{k+1} - t_k} \cdot z_k^i, \quad (22)$$

the Euler-Maruyama discretization of the forward process using N Monte-Carlo samples, where $K = T/h$ if the stepsize $h = t_{k+1} - t_k$ for all k is kept constant.

Remark 3.4. As observed in (Vargas & Nüsken, 2023), discretizations with backward integral (21)-(22) has the additional benefit of giving ELBO for the normalizing constant Z of p_{target} with the estimator

$$\hat{Z} = \frac{\mu(X_K) q^v(X_{0:K-1} | X_K)}{\nu(X_0) p^u(X_{1:K} | X_0)} \quad (23)$$

in the sense that $\mathbb{E}_{\nu(X_0) p^u(X_{1:K} | X_0)} [\log \hat{Z}] \leq \log(\mathbb{E}_{\nu(X_0) p^u(X_{1:K} | X_0)} [\hat{Z}]) = \log[\int \mu(X_K) dX_K] = \log(Z)$. This aspect is exploited in Appendix D when comparing the $\log Z$ estimator, as our estimator from Lemma 3.6 below can be understood as a ratio of two discrete chains as above.

Putting everything together, we give the proposed algorithm.

Algorithm 1 Control Objective Training for Sampling

Require: Initial draw $(X_0^{i,(0)})_{i=1}^N \in \mathbb{R}^d \sim \nu$ independent, initial controls $u^{(0)}, v^{(0)}$

Require: Un-normalized density μ , drift f , num of time steps K , num of iterations T

for $t = 0, \dots, T - 1$ **do**

Simulate (22) with current control $(u_k^{(t)})_{k=0, \dots, K}$, obtain $(X_k^{n,(t)})_{k=0, \dots, K}$ for $n = 1, \dots, N$

Estimate the loss (21) + discretized regularizer (c.f. Lemma 3.6 below) and the gradient w.r.t the two parameterized controls using the samples $(X_k^{n,(t)})_{k=0, \dots, K, n=1, \dots, N}$

Gradient update on the parameters to obtain $u^{(t+1)}$ and $v^{(t+1)}$

end for

return $X_K^{1,(T)}, \dots, X_K^{N,(T)}$ as N samples from μ with their importance weights $w^{u^{(T)}}(X_K^{n,(T)})$ (34)

In practice, with imperfect control from the training procedure, one can perform importance sampling to correct for the bias / improve on the estimate.

Proposition 3.5 (Importance Sampling). *The following can be used to estimate the normalizing constant of the target density p_{target} :*

(1) For the optimal ϕ^*, ψ^* , with $X_t \sim \overrightarrow{\mathbb{P}}^{\nu, f + \sigma^2 \nabla \phi^*}$,

$$\begin{aligned} Z &= \frac{\mu(X_T)}{\nu(X_0)} \exp \left(\frac{\sigma^2}{2} \int_0^T \nabla \cdot (\nabla \phi_t^* - \nabla \psi_t^*)(X_t) dt \right. \\ &\quad \left. + \int_0^T \nabla \cdot f_t(X_t) dt \right). \end{aligned}$$

(2) For any potentially suboptimal ϕ, ψ , with $X_t \sim \overrightarrow{\mathbb{P}}^{\nu, f + \sigma^2 \nabla \phi}$,

$$\begin{aligned} Z &= \mathbb{E}_{\overrightarrow{\mathbb{P}}^{\nu, f + \sigma^2 \nabla \phi}} \left[\exp \left(- \frac{\sigma^2}{2} \int_0^T \|\nabla \phi_t + \nabla \psi_t\|^2 + \nabla \cdot \right. \right. \\ &\quad \left. \left. (f_t - \sigma^2 \nabla \psi_t) dt - \sigma \int_0^T \nabla \phi_t + \nabla \psi_t dW_t - \log \frac{\nu(X_0)}{\mu(X_T)} \right) \right]. \end{aligned}$$

With a sub-optimal control $\nabla \phi$ and $X_t \sim \overrightarrow{\mathbb{P}}^{\nu, f + \sigma^2 \nabla \phi}$, re-weighting can be used to get an unbiased estimator of a

statistics g as $\mathbb{E}_\phi[g(X_T)w^\phi(X_T)] = \mathbb{E}_{\phi^*}[g(X_T)] = \mathbb{E}_\mu[g]$ for

$$w^\phi(X) = \exp\left(\log \frac{\mu}{\nu * \mathcal{N}(0, \sigma^2 T \cdot I)}(X_T) - \frac{1}{2\sigma^2} \int_0^T \|f_t + \sigma^2 \nabla \phi_t\|^2 dt - \frac{1}{\sigma} \int_0^T (f_t + \sigma^2 \nabla \phi_t) dW_t\right).$$

The following lemma provides a recipe for estimating the regularizer and the normalizing constant with discrete-time updates. We focus on loss (3) from Proposition 3.1, since most of the other parts are straightforward to adapt.

Lemma 3.6 (Discretized Loss and Estimator). *For $X \sim \overrightarrow{\mathbb{P}}^{\nu, f + \sigma^2 \nabla \phi}$,*

$$\text{Var}\left(\psi_T(X_T) - \psi_0(X_0) - \sigma \int_0^T \nabla \psi_t(X_t) dW_t - \int_0^T \left(\frac{\sigma^2}{2} \|\nabla \psi_t\|^2 + \nabla \cdot (\sigma^2 \nabla \psi_t - f_t) + \sigma^2 \nabla \psi_t^\top \nabla \phi_t\right)(X_t) dt\right)$$

can be estimated as (Var_N denotes empirical estimate of variance using N samples)

$$\text{Var}_N\left(\psi_K(X_{K+1}^i) - \psi_0(X_0^i) - \sum_{k=0}^{K-1} \frac{1}{2\sigma^2 h} \|X_k^i - X_{k+1}^i + (f_{k+1} - \sigma^2 \nabla \psi_{k+1})(X_{k+1}^i)h\|^2 + \sum_{k=0}^{K-1} \frac{1}{2\sigma^2 h} \|X_{k+1}^i - X_k^i - f_k(X_k^i)h\|^2\right),$$

which compared to the KL estimator (21) has two terms. The importance-weighted Z -estimator from Proposition 3.5 can be approximated as

$$\frac{1}{N} \sum_{i=1}^N \exp\left(\log \frac{\mu(X_K^i)}{\nu(X_0^i)} + \frac{1}{2} \sum_{k=0}^{K-1} \|z_k^i\|^2 - \sum_{k=0}^{K-1} \frac{1}{2\sigma^2 h} \|X_k^i - X_{k+1}^i + (f_{k+1} - \sigma^2 \nabla \psi_{k+1})(X_{k+1}^i)h\|^2\right).$$

In both cases for $i \in [N]$ independently, $z_k^i \sim \mathcal{N}(0, I)$,

$$X_{k+1}^i = X_k^i + (f_k(X_k^i) + \sigma^2 \nabla \phi_k(X_k^i))h + \sigma\sqrt{h} \cdot z_k^i.$$

Remark 3.7. In (Vargas & Nüsken, 2023), the authors propose $\int_0^T \mathbb{E}|\partial_t \phi + f^\top \nabla \phi + \frac{\sigma^2}{2} \Delta \phi + \frac{\sigma^2}{2} \|\nabla \phi\|^2|(X_t) dt$ as the HJB regularizer (c.f. (46)) on top of $D_{KL}(\overrightarrow{\mathbb{P}}^{\nu, f + \sigma^2 \nabla \phi} \parallel \overrightarrow{\mathbb{P}}^{\mu, f - \sigma^2 \nabla \psi})$ as the training loss. As is clear from the proof in Proposition 3.1, it is equivalent to our loss (2). There is, of course, a similar HJB for the backward drift $\nabla \psi$, which in this case will read as $\partial_t \psi + f^\top \nabla \psi - \frac{\sigma^2}{2} \Delta \psi - \frac{\sigma^2}{2} \|\nabla \psi\|^2 + \nabla \cdot f = 0$. By trading a PDE constraint with a SDE one (based on likelihood ratio of path measures), we can avoid evaluating the divergence term for the discrete updates. It remains to note the additional benefit of vanishing variance of the gradient at the optimal control, as remarked earlier.

4. Numerics and Comparisons

In this section, we instantiate the main contributions of the work (Proposition 3.1 and Lemma 3.6) and offer numerical evidence on their advantages compared to existing proposals for solving an optimal trajectory problem in a typical MCMC setup.

4.1. Algorithm specification

Below for simplicity we pick the reference process to be a Brownian motion with $f = 0$ and λ is a parameter that we tune for best performance, but in theory any $\lambda > 0$ would work. This aspect deviates from other constrained formulation of the problem.

(1) PINN-regularization (Vargas & Nüsken, 2023): for $i = 1, \dots, n$, and $Z \sim \mathcal{N}(0, I)$ independently draw in parallel

$$x_{k+1}^i = x_k^i + \sigma^2 h \nabla \phi(x_k^i, kh) + \sigma\sqrt{h} Z_k^i, x_0^i \sim \nu \quad (24)$$

for $k = 0, \dots, K$ with stepsize $h = c/(K+1)$ for some $c \geq 1$. Using these trajectories $\{x_k^i\}$, minimize over ϕ and s ,

$$\begin{aligned} & \frac{1}{n} \sum_{i=1}^n \left[\log \frac{\nu(x_0^i)}{\mu(x_{K+1}^i)} + \sum_{k=0}^K \frac{1}{2\sigma^2 h} \right. \\ & \left. \|x_k^i - x_{k+1}^i + \sigma^2 h (\nabla \phi(x_{k+1}^i, (k+1)h) - \nabla s(x_{k+1}^i, (k+1)h))\|^2 \right] \\ & + \lambda \cdot \frac{1}{n} \sum_{i=1}^n \left[\sum_{k=0}^K \left| \partial_t \phi(x_k^i, kh) + \right. \right. \\ & \left. \left. \frac{\sigma^2}{2} \Delta \phi(x_k^i, kh) + \frac{\sigma^2}{2} \|\nabla \phi(x_k^i, kh)\|^2 \right| \cdot h \right]. \quad (25) \end{aligned}$$

Above $\phi(x, t)$, $s(x, t)$ are two neural networks that take $t \in \mathbb{R}$ and $x \in \mathbb{R}^d$ as inputs and maps to \mathbb{R} , ∇ is the space derivative w.r.t $x \in \mathbb{R}^d$ (i.e., derivative w.r.t the first input), ∂_t is the time derivative (i.e., derivative w.r.t the second input). With the new ϕ, s we repeat (24) and (25) several times, and compute statistics using the latest samples $\{x_{k+1}^i\}_{i=1}^n$.

Another possibility is to replace first term of (25) with

$$\begin{aligned} & \frac{1}{K+1} \text{Var}_n \left[\log \frac{\nu(x_0^i)}{\mu(x_{K+1}^i)} + \sum_{k=0}^K \frac{1}{2\sigma^2 h} \|x_k^i - x_{k+1}^i + \right. \\ & \left. \sigma^2 h (\nabla \phi(x_{k+1}^i, (k+1)h) - \nabla s(x_{k+1}^i, (k+1)h))\|^2 \right. \\ & \left. - \sum_{k=0}^K \frac{1}{2\sigma^2 h} \|x_{k+1}^i - x_k^i - \sigma^2 h \nabla \phi(x_k^i, kh)\|^2 \right], \quad (26) \end{aligned}$$

which will become a log-variance divergence between the two path measures, and is oblivious to unknown normalizing constant. Similarly for the first term of (28) below.

(2) Variance-regularization (loss (2) from Proposition 3.1): Simulate trajectories (24) as before, but additionally simulate $\{y_k^i\}_{k=0}^K$ as follows and cache them:

$$y_{k+1}^i = y_k^i + \sigma\sqrt{h}Z_k^i, y_0^i \sim \nu. \quad (27)$$

Minimize over ϕ, s the following discretized loss

$$\begin{aligned} & \frac{1}{n} \sum_{i=1}^n \left[\log \frac{\nu(x_0^i)}{\mu(x_{K+1}^i)} + \sum_{k=0}^K \frac{1}{2\sigma^2 h} \|x_k^i - x_{k+1}^i + \right. \\ & \left. \sigma^2 h (\nabla\phi(x_{k+1}^i, (k+1)h) - \nabla s(x_{k+1}^i, (k+1)h)) \|^2 \right] \\ & + \frac{\lambda}{K+1} \cdot \text{Var}_n \left[\phi(y_{K+1}^i, (K+1)h) - \phi(y_0^i, 0) + \sum_{k=0}^K \right. \\ & \left. \frac{1}{2\sigma^2 h} (\|y_{k+1}^i - y_k^i - \sigma^2 \nabla\phi(y_k^i, kh)h\|^2 - \|y_{k+1}^i - y_k^i\|^2) \right]. \end{aligned} \quad (28)$$

Notice the first term is the same as (25). Again repeat (24) and (28) several times. Above Var_n denotes the empirical variance across the n trajectories $\{y_k^i\}_{i=1}^n$ of the quantity inside $[\cdot]$. The loss (28), as the proof of Proposition 3.1 shows, comes from the fact that along the prior $X \sim \overrightarrow{\mathbb{P}}^{\nu, f}$,

$$\begin{aligned} & \log \left(\frac{d\overrightarrow{\mathbb{P}}^{\nu, f + \sigma^2 \nabla\phi_t}}{d\overrightarrow{\mathbb{P}}^{\nu, f}} \right) (X) \\ & = \int_0^T -\frac{\sigma^2}{2} \|\nabla\phi_t\|^2 dt + \int_0^T \sigma \nabla\phi_t^\top dW_t \\ & = \phi_T(X_T) - \phi_0(X_0), \end{aligned} \quad (29)$$

and we discretized the Radon-Nikodym derivative similar to how it was done in the KL divergence $D_{KL}(\overrightarrow{\mathbb{P}}^{\nu, \sigma^2 \nabla\phi} \parallel \overleftarrow{\mathbb{P}}^{\mu, -\sigma^2 \nabla\psi})$ (see Lemma 3.6 for a similar derivation).

Instead of the regularizer (28), another discretization of condition (29) can be a TD-like regularizer similar to (Liu et al., 2022):

$$\begin{aligned} & \dots + \lambda \cdot \frac{1}{n} \sum_{i=1}^n \sum_{k=0}^K h \cdot \left| \phi(y_{k+1}^i, (k+1)h) - \phi(y_k^i, kh) + \right. \\ & \left. \frac{\sigma^2 h}{2} \|\nabla\phi(y_k^i, kh)\|^2 - \sigma\sqrt{h} \nabla\phi(y_k^i, kh)^\top Z_k^i \right|. \end{aligned} \quad (30)$$

The loss above can also be justified with Lemma 3.2. In all (25), (28) and (30), we parameterize the forward / backward drift as $\sigma^2 \nabla\phi, \sigma^2 \nabla\phi - \sigma^2 \nabla s$. A regularization on the backward drift involving s is also possible for (25):

$$\begin{aligned} & \dots + \lambda \cdot \frac{1}{n} \sum_{i=1}^n \left[\sum_{k=0}^K \left| \partial_t s(x_k^i, kh) - \partial_t \phi(x_k^i, kh) - \frac{\sigma^2}{2} \right. \right. \\ & \left. \left. \Delta s(x_k^i, kh) + \frac{\sigma^2}{2} \Delta \phi(x_k^i, kh) - \frac{\sigma^2}{2} \right. \right. \end{aligned}$$

$$\left. \|\nabla s(x_k^i, kh) - \nabla \phi(x_k^i, kh)\|^2 \right] \cdot h \Big].$$

and similarly for (28) and (30), which will read as (notice there is no divergence term in the 1st one below in contrast to the 2nd)

$$\begin{aligned} & \dots + \frac{\lambda}{K+1} \cdot \text{Var}_n \left[(s - \phi)(y_{K+1}^i, (K+1)h) - (s - \phi)(y_0^i, 0) + \right. \\ & \left. \sum_{k=0}^K \frac{1}{2\sigma^2 h} (\|y_{k+1}^i - y_k^i\|^2 - \right. \\ & \left. \|y_k^i - y_{k+1}^i + \sigma^2 h (\nabla\phi - \nabla s)(y_{k+1}^i, (k+1)h)\|^2) \right] \end{aligned} \quad (31)$$

and

$$\begin{aligned} & \dots + \lambda \cdot \frac{1}{n} \sum_{i=1}^n \sum_{k=0}^K h \cdot \left| (s - \phi)(y_{k+1}^i, (k+1)h) - (s - \phi)(y_k^i, kh) \right. \\ & \left. - \frac{\sigma^2 h}{2} \|\nabla s(y_k^i, kh) - \nabla \phi(y_k^i, kh)\|^2 \right. \\ & \left. - \sigma^2 \Delta (s - \phi)(y_k^i, kh) + \sigma\sqrt{h} \nabla(\phi - s)(y_k^i, kh)^\top Z_k^i \right|. \end{aligned} \quad (32)$$

respectively, but should witness similar comparison as the regularizer on the forward control that we experiment below.

(3) Separately-controlled loss (loss (3) from Proposition 3.1): Simulate (24) as before, with the n trajectories $\{x_k^i\}$, minimize over ϕ, ψ the following discretized loss (c.f. Lemma 3.6):

$$\begin{aligned} & \text{Var}_n \left[\psi(x_{K+1}^i, (K+1)h) + \phi(x_{K+1}^i, (K+1)h) - \log \mu(x_{K+1}^i) \right] \\ & + \text{Var}_n \left[\psi(x_0^i, 0) + \phi(x_0^i, 0) - \log \nu(x_0^i) \right] + \\ & \frac{\lambda}{K+1} \cdot \text{Var}_n \left[\psi(x_{K+1}^i, (K+1)h) - \psi(x_0^i, 0) + \sum_{k=0}^K \frac{1}{2\sigma^2 h} \right. \\ & \left. (\|x_{k+1}^i - x_k^i\|^2 - \|x_k^i - x_{k+1}^i - \sigma^2 h \nabla\psi(x_{k+1}^i, (k+1)h)\|^2) \right] + \\ & \frac{\lambda}{K+1} \cdot \text{Var}_n \left[\phi(x_0^i, 0) - \phi(x_{K+1}^i, (K+1)h) + \sum_{k=0}^K \frac{1}{2\sigma^2 h} \right. \\ & \left. (\|x_{k+1}^i - x_k^i\|^2 - \|x_{k+1}^i - x_k^i - \sigma^2 h \nabla\phi(x_k^i, kh)\|^2) \right]. \end{aligned} \quad (33)$$

Notice that it consists of 4 variance terms, and we alternate between simulating (24) and updating ϕ, ψ from (33) several times. For this loss, we parameterize the forward / backward drift as $\sigma^2 \nabla\phi, -\sigma^2 \nabla\psi$. One can also consider the SDE-based discretization for the last 2 terms but will incur additional Laplacian and divergence terms as suggested by Lemma 3.2.

We emphasize that the discretized regularizers (28) and (31) wouldn't be available without the path-wise stochastic

process perspective. The FBSDE view (Lemma 3.2) will naturally lend to TD-like regularizers (30) and (32) as in (Liu et al., 2022), although the actual form there differs as they consider dynamics with mean-field interaction. Similar comment applies to (33), where we *separately* impose optimality condition on

$$\log \left(\frac{d\overrightarrow{\mathbb{P}}^{\nu, f + \sigma^2 \nabla \phi_t}}{d\overrightarrow{\mathbb{P}}^{\nu, f}} \right) (X) \text{ and } \log \left(\frac{d\overrightarrow{\mathbb{P}}^{\nu, f}}{d\overleftarrow{\mathbb{P}}^{\mu, f - \sigma^2 \nabla \psi_t}} \right) (X),$$

for $X \sim \overrightarrow{\mathbb{P}}^{\nu, f + \sigma^2 \nabla \phi_t}$ to take factorized forms, as opposed to looking at their sum

$$\log \left(\frac{d\overrightarrow{\mathbb{P}}^{\nu, f + \sigma^2 \nabla \phi}}{d\overleftarrow{\mathbb{P}}^{\mu, f - \sigma^2 \nabla \psi}} \right)$$

only, effectively erasing the ‘‘SB optimality enforcement’’ part. The two discretized losses (28) and (33) both have the variance reduction (compared to TD (30) as elaborated in Remark 3.3), without evaluating expensive Laplacian terms (as the PDE-based PINN approach (25) would require). We also highlight that the algorithm does not use gradient information from the target $\nabla \log \mu$ as e.g., Langevin would.

4.2. Importance weighting

For the very last sampling SDE, we simulate (24) with the latest $\nabla \phi$, but re-weight the n samples $\{x_{K+1}^i\}_{i=1}^n$ each with individual weights

$$w(x_{K+1}^i) = \frac{\mu(x_{K+1}^i)}{\rho(x_{K+1}^i)} \exp \left(- \sum_{k=0}^K \frac{1}{2\sigma^2 h} \|x_{k+1}^i - x_k^i\|^2 + \sum_{k=0}^K \frac{1}{2\sigma^2 h} \|x_{k+1}^i - x_k^i - \sigma^2 h \nabla \phi(x_k^i, kh)\|^2 \right) \quad (34)$$

before taking their average, i.e.,

$$\hat{\mathbb{E}}_{\mu}[g] = \frac{\sum_{i=1}^n g(x_{K+1}^i) w(x_{K+1}^i)}{\sum_{i=1}^n w(x_{K+1}^i)} \quad (35)$$

for a summary statistics $g : \mathbb{R}^d \mapsto \mathbb{R}$ we are interested in. Above ρ is the Gaussian pdf of $\mathcal{N}(b_1, \sigma^2 c \cdot I + b_2)$ if $\nu = \mathcal{N}(b_1, b_2)$ and $h = c/(K+1)$. This estimator follows from Proposition 3.5 and is used for post-processing with a potentially suboptimal control $\nabla \phi$.

4.3. Result

In our experiments, we picked the following 4 targets as benchmark as they capture different properties of distributions that can affect the training of the controls. These properties (multimodality, spatially varying curvature) pose challenges for MCMC methods and are ideal examples to demonstrate the differences between the regularizers.

	PINN	Variance	TD	SC
Standard normal	6.318	2.883	2.014	1.305
Funnel	8.854	3.560	3.550	3.391
GMM	5.803	8.347	4.097	4.393
Double well	12.274	5.392	4.437	1.761

Table 1. $-\log Z$ estimator for different regularization methods on 2D benchmarks (lower is better)




Figure 1. Weighted marginal histogram for double well

Additional supporting numerical experiments and details can be found in Appendix D. Across our experiments, we observe that our separately controlled loss (33) is generally much better compared to the PINN loss (25), while TD (30) and variance regularizer (28) can sometimes be comparable (and tend to exhibit mode-seeking behavior due to reverse KL). In practice, we also observe that the separately controlled loss is less sensitive to tuning parameters, as can also be observed by the smoother training loss trajectory.

5. Conclusion

We detailed various connections between the seemingly disjoint areas of diffusion generative-modeling, optimal stochastic control and optimal transport, with the goal of sampling from high-dimensional, complex distributions in mind. This is orthogonal to MCMC-based approaches, and is accomplished by a more ‘‘learning-driven’’ methodology on the optimal control / drift that can be trained with a suitable control objective. There are a lot of flexibility in the choice of the reference, optimality criteria, divergence metric etc. (that also makes the case for SB more compelling than existing diffusion-based samplers mentioned in Section 2.2), which we hope to report on in the future.

Impact Statement

This paper presents work whose goal is to advance the field of machine learning. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here.

References

- Benamou, J.-D. and Brenier, Y. A computational fluid mechanics solution to the Monge-Kantorovich mass transfer problem. *Numerische Mathematik*, 84(3):375–393, 2000.
- Berner, J., Richter, L., and Ullrich, K. An optimal control perspective on diffusion-based generative modeling. *arXiv preprint arXiv:2211.01364*, 2022.
- Caluya, K. F. and Halder, A. Wasserstein proximal algorithms for the Schrödinger bridge problem: Density control with nonlinear drift. *IEEE Transactions on Automatic Control*, 67(3):1163–1178, 2021.
- Chen, R. T., Rubanova, Y., Bettencourt, J., and Duvenaud, D. K. Neural ordinary differential equations. *Advances in neural information processing systems*, 31, 2018.
- Chen, T., Liu, G.-H., and Theodorou, E. Likelihood Training of Schrödinger Bridge using Forward-Backward SDEs Theory. In *International Conference on Learning Representations*, 2021a.
- Chen, Y., Georgiou, T. T., and Pavon, M. Stochastic control liaisons: Richard Sinkhorn meets Gaspard Monge on a Schrödinger Bridge. *SIAM Review*, 63(2):249–313, 2021b.
- De Bortoli, V., Thornton, J., Heng, J., and Doucet, A. Diffusion Schrödinger bridge with applications to score-based generative modeling. *Advances in Neural Information Processing Systems*, 34:17695–17709, 2021.
- Dockhorn, T., Vahdat, A., and Kreis, K. Score-Based Generative Modeling with Critically-Damped Langevin Diffusion. In *International Conference on Learning Representations (ICLR)*, 2022.
- Hyvärinen, A. and Dayan, P. Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research*, 6(4), 2005.
- Janati, H., Muzellec, B., Peyré, G., and Cuturi, M. Entropic optimal transport between unbalanced Gaussian measures has a closed form. *Advances in neural information processing systems*, 33:10468–10479, 2020.
- Léonard, C. A survey of the Schrödinger problem and some of its connections with optimal transport. *Discrete and Continuous Dynamical Systems-Series A*, 34(4):1533–1574, 2014.
- Liu, G.-H., Chen, T., So, O., and Theodorou, E. Deep generalized Schrödinger bridge. *Advances in Neural Information Processing Systems*, 35:9374–9388, 2022.
- Pavon, M., Trigila, G., and Tabak, E. G. The Data-Driven Schrödinger Bridge. *Communications on Pure and Applied Mathematics*, 74(7):1545–1573, 2021.
- Peluchetti, S. Diffusion Bridge Mixture Transports, Schrödinger Bridge Problems and Generative Modeling. *arXiv preprint arXiv:2304.00917*, 2023.
- Peyré, G. and Cuturi, M. Computational optimal transport: With applications to data science. *Foundations and Trends in Machine Learning*, 11(5-6):355–607, 2019.
- Richter, L., Berner, J., and Liu, G.-H. Improved sampling via learned diffusions. *arXiv preprint arXiv:2307.01198*, 2023.
- Schrödinger, E. *Über die umkehrung der naturgesetze*. 1931.
- Song, Y., Sohl-Dickstein, J., Kingma, D. P., Kumar, A., Ermon, S., and Poole, B. Score-Based Generative Modeling through Stochastic Differential Equations. In *International Conference on Learning Representations*, 2020.
- Tzen, B. and Raginsky, M. Theoretical guarantees for sampling and inference in generative models with latent diffusions. In *Conference on Learning Theory*, pp. 3084–3114. PMLR, 2019.
- Vargas, F. and Nüsken, N. Transport, Variational Inference and Diffusions: with Applications to Annealed Flows and Schrödinger Bridges. *arXiv preprint arXiv:2307.01050*, 2023.
- Vargas, F., Thodoroff, P., Lamacraft, A., and Lawrence, N. Solving Schrödinger bridges via maximum likelihood. *Entropy*, 23(9):1134, 2021.
- Vargas, F., Grathwohl, W. S., and Doucet, A. Denoising Diffusion Samplers. In *The Eleventh International Conference on Learning Representations*, 2022.
- Zhang, Q. and Chen, Y. Path Integral Sampler: A Stochastic Control Approach For Sampling. In *International Conference on Learning Representations*, 2021.

A. Additional Related Work: Natural attempts for solving SB

In this section we discuss several natural attempts with the intention of adapting the SB formalism for sampling from un-normalized densities.

(1) Iterative proportional fitting / Sinkhorn relies on factorized form (55) of the optimal solution $\pi^*(z, x) = P_{0T}^*$ for iterative projection as

$$P^{(1)} = \arg \min_{Q \in \mathcal{P}(\nu, \cdot)} D_{KL}(Q \| P^{(0)}) = \frac{P^{(0)} \nu}{P_0^{(0)}}, \quad P^{(0)} = \arg \min_{Q \in \mathcal{P}(\cdot, \mu)} D_{KL}(Q \| P^{(1)}) = \frac{P^{(1)} \mu}{P_T^{(1)}} \quad (36)$$

i.e., only the end point differ and one solves half-bridges using drifts learned from the last trajectory transition rollout, initializing from either samples from $\mu(x)$ or $\nu(z)$. But since we don't have samples from μ , neither the score (i.e., the drift) nor the proceeding IPF updates/refinements can be implemented. In fact, the first iteration of the IPF proposal in (De Bortoli et al., 2021; Vargas et al., 2021) precisely corresponds to the score-based diffusion proposal in (Song et al., 2020). However, there is a connection between IPF and Path space EM, which means that coordinate descent on the objective $\min_{\phi, \psi} D_{KL}(\overrightarrow{\mathbb{P}}^{\nu, f + \sigma^2 \nabla \phi} \| \overleftarrow{\mathbb{P}}^{\mu, f + \sigma^2 \nabla \psi})$, when initializing at $\phi = 0$ (i.e., the Schrödinger prior), is a valid strategy for solving (17).

Lemma A.1 (Optimization of drifts of EM). *The alternating scheme initializing with $\phi_0 = 0$ converge to $\pi^*(z, x)$:*

$$\psi_n = \arg \min_{\psi} D_{KL}(\overrightarrow{\mathbb{P}}^{\nu, f + \sigma^2 \nabla \phi_{n-1}} \| \overleftarrow{\mathbb{P}}^{\mu, f + \sigma^2 \nabla \psi}) \quad (37)$$

$$\phi_n = \arg \min_{\phi} D_{KL}(\overrightarrow{\mathbb{P}}^{\nu, f + \sigma^2 \nabla \phi} \| \overleftarrow{\mathbb{P}}^{\mu, f + \sigma^2 \nabla \psi_n}). \quad (38)$$

Moreover, both updates are implementable assuming samples from ν is available, which resolves non-uniqueness in an algorithmic manner (fixing one direction of the drift at a time).

However, Lemma A.1 above also shows that the prior only enters in the first step, therefore as it proceeds, the prior influence tends to be ignored as error accumulates – this aspect is different from our proposal in Proposition 3.1 below. Another advantage of our joint minimization procedure compared to an iterative IPF method is that while the equivalence of EM and IPF will no longer hold for neural network models not expressive enough (therefore the convergence does not immediately carry over), one could still estimate the normalizing constant / sample from the target with a sub-optimal control learned from the losses proposed in Proposition 3.1 with e.g., important weights (c.f. (35)).

(2) Some alternatives to solve SB do not require analytical expression for μ : diffusion mixture matching (Peluchetti, 2023) tilts the product measure $\nu \otimes \mu$ towards optimality gradually by learning a slightly different term than the score; data-driven bridge (Pavon et al., 2021) aims at setting up a fixed point recursion on the SB system (56) for finding the optimal ϕ^*, ψ^* , but both rely on the availability of samples from μ (as well as ν) to estimate various quantities for implementation.

(3) Naively one may hope to reparametrize (52)-(53) as (with the time-reversal condition baked in)

$$\min_u D_{KL}(\overrightarrow{\mathbb{P}}^{\nu, \frac{\sigma^2}{2} \nabla \log \rho_t + \nabla u_t} \| \overleftarrow{\mathbb{P}}^{\mu, \nabla u_t - \frac{\sigma^2}{2} \nabla \log \rho_t})$$

to learn the drift ∇u (this is somewhat akin to annealed flow, and expected to have unique solution). This way, when the two path measures are optimized to agree, the resulting forward & backward drift necessarily obey Nelson's identity for the desired curve ρ_t . However, this only applies if we have access to the score $\nabla \log \rho_t$, for $(\rho_t)_{t \in [T]}$ the SB curve of density we wish to follow interpolating from ν to μ . Learning both ρ and ∇u will become similar to learning ϕ and ψ , as the approaches we study in Section 3.1. In a similar vein, if one chooses $\Gamma_0 = \nu$, $\Gamma_T = \mu$ and $r_t^+ = \sigma \nabla u_t^*$, $r_t^- = \sigma \nabla u_t^* - \sigma^2 \nabla \log \rho_t$ the optimal SB process as reference, Lemma C.1 asserts that the loss becomes

$$\arg \min_{\nabla \phi} D_{KL}(\overrightarrow{\mathbb{P}}^{\nu, \sigma \nabla \phi_t} \| \overleftarrow{\mathbb{P}}^{\mu, \sigma \nabla \phi_t - \sigma^2 \nabla \log \rho_t}) = \arg \min_{\nabla \phi} \mathbb{E} \left[\int_0^T \|\nabla \phi_t - \nabla u_t^*\|^2 dt \right],$$

something we cannot estimate and optimize for without additional information.

(4) The work of (Caluya & Halder, 2021) investigated the case where the reference process has a gradient drift (i.e., $f = -\nabla U$) and reduce the optimal control task to solving a high-dimensional PDE subject to initial-value constraint (c.f. Eqn (33) and (47) therein). However, solving PDEs is largely regarded to be computationally more demanding than simulating SDEs.

B. Various Perspectives on Schrödinger Bridge

Most of these can be traced out in (Léonard, 2014; Chen et al., 2021b), which has its roots in statistical mechanics (Schrödinger, 1931) (and in modern terms, closely related to large-deviation results via Sanov’s theorem). For us however, the following perspectives will be more fruitful.

(1) Mixture of pinned diffusions (weights given by π^*) can be seen by disintegration of path measure:

$$D_{KL}(P\|Q) = D_{KL}(P_{0T}\|Q_{0T}) + \int D_{KL}(P^{zx}\|Q^{zx}) dP_{0T}(z, x). \quad (39)$$

Since the only constraints are on the two end-points (17) is therefore equivalent to (40), by choosing $P^{zx} = Q^{zx}$. The optimal solution takes the form $P^* = \pi^* Q^{zx}$, which means one can sample from $(z, x) \sim \pi^* \in \Pi_{\nu, \mu}$, and sample from the bridges conditioned on the end-points at $t = 0, T$.

(2) It has the interpretation of entropy-regularized optimal transport (Peyré & Cuturi, 2019) when $f = 0$: in terms of static formulation because of (39), if the reference Q is simply the Wiener process (accordingly $r(x|z) \propto e^{-\frac{1}{2T}\|x-z\|^2}$),

$$\begin{aligned} \pi^*(z, x) &= \arg \min_{\pi_z=\nu, \pi_x=\mu} D_{KL}(\pi(z, x)\|r(z, x)) \\ &= \arg \min_{\pi_z=\nu, \pi_x=\mu} \mathbb{E}_{z \sim \nu} [D_{KL}(\pi(x|z)\|r(x|z))] + D_{KL}(\nu(z)\|r(z)) \\ &= \arg \min_{\pi_z=\nu, \pi_x=\mu} \int \frac{1}{2} \|x - z\|^2 \pi(z, x) dx dz + T \int \pi(z, x) \log \pi(z, x) dx dz, \end{aligned} \quad (40)$$

where the first term is nothing but the definition of the Wasserstein-2 distance (and optimal transport in the sense of Kantorovich). The second entropy term favors independent coupling. This is a Lagrangian description of the transport. In terms of dynamical formulation, via Girsanov’s theorem on path measure, with the controlled dynamics as $dX_t = [f_t(X_t) + \sigma u_t(X_t)]dt + \sigma dW_t$, (17) can be reformulated as a *constrained* problem over u

$$P^* = \arg \min_P \mathbb{E}_{X \sim \vec{\mathbb{P}}^{\nu, f + \sigma u}} \left[\frac{1}{2} \int_0^T \|u_t(X_t)\|^2 dt \mid \vec{\mathbb{P}}_T^{\nu, f + \sigma u} = \mu \right]. \quad (41)$$

Or a *regularized* Benamou-Brenier fluid-dynamics (Benamou & Brenier, 2000) analogy of the optimal transport

$$\inf_{\rho, v} \int_{\mathbb{R}^d} \int_0^T \left[\frac{1}{2} \|v_t(X_t) - \bar{v}_t(X_t)\|^2 + \frac{\sigma^4}{8} \left\| \nabla \log \frac{\rho_t(X_t)}{\bar{\rho}_t(X_t)} \right\|^2 \right] \rho_t(X_t) dt dx \quad (42)$$

$$\text{s.t. } \frac{\partial \rho_t}{\partial t} + \nabla \cdot (\rho_t v_t) = 0, \rho_0 = \nu, \rho_T = \mu \quad (43)$$

where above $\bar{v}_t = f_t - \frac{\sigma^2}{2} \nabla \log \bar{\rho}_t$ is the current velocity of the prior process, and we see the penalization results in an additional relative Fisher information term. Since v^* is a velocity field in the continuity equation, it means a deterministic evolution (i.e., ODE) as

$$\dot{X}_t = v_t(X_t), X_0 \sim \nu$$

will have $X_t \sim \rho_t$, the optimal entropic interpolation flow, which gives an Eulerian viewpoint.

(3) Optimal control views the problem as steering ν at $t = 0$ to μ at $t = T$ with minimal control effort. The value function (i.e., optimal cost-to-go)

$$V(x, t) := \min_u \mathbb{E}_u \left[\frac{1}{2} \int_t^T \|u_s(X_s)\|^2 ds \mid X_t^u = x, X_T^u \sim \mu \right] \quad (44)$$

with the expectation taken over the stochastic dynamics

$$dX_t^u = [f_t(X_t^u) + \sigma u_t(X_t^u)]dt + \sigma dW_t, X_0^u \sim \nu \quad (45)$$

should satisfy the Hamilton-Jacobi-Bellman equation via the dynamical programming principle

$$\frac{\partial V(x, t)}{\partial t} + f_t(x)^\top \nabla V(x, t) + \frac{\sigma^2}{2} \Delta V(x, t) - \frac{\sigma^2}{2} \|\nabla V(x, t)\|^2 = 0, \quad (46)$$

where the optimal control $u_t^*(X_t) = -\sigma \nabla V(X_t, t)$ and gives the unique solution $(\rho_t^u)_{t \geq 0}$ solving

$$\frac{\partial \rho_t^u}{\partial t} = -\nabla \cdot (\rho_t^u (f_t - \sigma^2 \nabla V_t)) + \frac{\sigma^2}{2} \Delta \rho_t^u, \quad \rho_0^u \sim \nu, \rho_T^u \sim \mu. \quad (47)$$

The correspondence between (44) and (46) can also be seen with the Feynman-Kac formula. Notice that

$$\min_u \mathbb{E}_u \left[\int_0^T \frac{1}{2} \|u_t(X_t)\|^2 dt - \log \frac{\mu}{Q_T}(X_T) \right] = 0,$$

which also gives an intuitive explanation of the optimally-controlled process. This holds since

$$D_{KL}(P \| Q) = \mathbb{E}_u \left[\int_0^T \frac{1}{2} \|u_t(X_t)\|^2 dt \right] = \mathbb{E}_u \left[\frac{1}{2} \int_0^T \|u_t(X_t)\|^2 dt - \log \frac{\mu}{Q_T}(X_T) \right] + D_{KL}(\mu \| Q_T),$$

which gives the claim by data processing inequality. The two coupled PDEs, one run forward in time (Fokker-Planck (47)), one run backward in time (HJB (46)) are also the first-order optimality condition of

$$\inf_{\rho, u} \int_{\mathbb{R}^d} \int_0^T \frac{1}{2} \|u_t(X_t)\|^2 \rho_t(X_t) dt dx \quad (48)$$

$$\text{s.t. } \frac{\partial \rho_t}{\partial t} + \nabla \cdot (\rho_t (f_t + \sigma u_t)) = \frac{\sigma^2}{2} \Delta \rho_t, \quad \rho_0 = \nu, \rho_T = \mu \quad (49)$$

where again the optimal $u^* = -\sigma \nabla V$ is of gradient type. Above the Laplacian is responsible for the diffusion part, and (48)-(49) is related to (42)-(43) via a change of variable. One might try to design schemes by forming the Lagrangian for the above (48)-(49), and solve the resulting saddle-point problem, but this deviates somewhat from our pathwise narrative. *Remark B.1.* Compared to the Langevin SDE $dX_t = -\nabla f(X_t) dt + \sqrt{2} dW_t$, which only involves forward-evolving density characterization and reaches equilibrium as $T \rightarrow \infty$, the controlled SDE (45) is time-inhomogeneous and involves two PDEs (46)-(47). Langevin also has a backward Kolmogorov evolution for the expectation of a function g : let $V(x, t) = \mathbb{E}[g(X_T) | X_t = x]$ we have $\partial_t V(x, t) - \nabla f(x)^\top \nabla V(x, t) + \Delta V(x, t) = 0$ with $V(x, T) = g(x)$, and is de-coupled from the Fokker-Planck equation $\partial_t \rho_t - \nabla \cdot (\rho_t \nabla f) - \Delta \rho_t = 0$ with $\rho_0 = \nu$ in this case.

(4) The dual potentials ϕ, ψ yield optimal forward/backward drifts: It holds that the optimal curve admits representation

$$\log \rho_t = \log \phi_t + \log \psi_t \quad \text{for all } t \quad (50)$$

and solving the boundary-coupled linear PDE system on the control

$$\frac{\partial \phi_t}{\partial t} = -\frac{\sigma^2}{2} \Delta \phi_t - \nabla \phi_t^\top f_t, \quad \frac{\partial \psi_t}{\partial t} = \frac{\sigma^2}{2} \Delta \psi_t - \nabla \cdot (\psi_t f_t) \quad \text{for } \phi_0 \psi_0 = p_{\text{prior}}, \phi_T \psi_T = p_{\text{target}} \quad (51)$$

gives two SDEs for the optimal curve in (17):

$$dX_t = [f_t(X_t) + \sigma^2 \nabla \log \phi_t(X_t)] dt + \sigma d\overrightarrow{W}_t, \quad X_0 \sim \nu \quad (52)$$

$$dX_t = [f_t(X_t) - \sigma^2 \nabla \log \psi_t(X_t)] dt + \sigma d\overleftarrow{W}_t, \quad X_T \sim \mu. \quad (53)$$

Note that equations (52)-(53) are time reversals of each other and obey Nelson's identity thanks to (50). The transformation (50)-(51) that involves $(\rho_t^*, u_t^*) = (\rho_t^*, -\sigma \nabla V_t) \mapsto (\phi_t, \psi_t)$ is a typical log \leftrightarrow exp Hopf-Cole change-of-variable from (46)-(47). These PDE optimality results can be found in (Caluya & Halder, 2021).

(5) Factorization of the optimal coupling: in fact it is always the case that

$$\frac{d\pi^*}{dr}(X_0, X_T) = e^{f(X_0)} e^{g(X_T)} \quad r\text{-a.s.} \quad (54)$$

Moreover under mild conditions if there exists π, f, g for which such decomposition holds and $\pi_0 = \nu, \pi_T = \mu$, π must be optimal – such condition (54) is *necessary and sufficient* for characterizing the solution to the SB problem. (54) together with (39) give that (which can also be thought of as re-weighting on the path space)

$$\frac{d\pi^*}{dr}(X_{0:T}) = e^{f(X_0)} e^{g(X_T)} \quad r\text{-a.s.}, \quad (55)$$

obeying the Schrödinger system

$$\int e^{f(X_0)} r(X_{0T}) e^{g(X_T)} dX_T = p_{\text{prior}}(X_0), \quad \int e^{f(X_0)} r(X_{0T}) e^{g(X_T)} dX_0 = p_{\text{target}}(X_T). \quad (56)$$

Using Doob's h -transform, one can write the transition kernel of π^* as a twisted dynamic of the prior r , and the ϕ, ψ in (52)-(53) can also be expressed as a conditional expectation:

$$\begin{aligned} \phi_t(x) &= \int e^{g(X_T)} r(X_T | X_t = x) dX_T = \int_{\mathbb{R}^d} \phi_T(X_T) r(X_T | X_t = x) dX_T, \\ \psi_t(x) &= \int e^{f(X_0)} r(X_t = x, X_0) dX_0 = \int_{\mathbb{R}^d} \psi_0(X_0) r(X_t = x | X_0) dX_0 \end{aligned}$$

for $t \in [0, T]$, where the consistency of

$$e^{f(X_0)} r(X_0) = \psi_0(X_0), \quad e^{g(X_T)} = \phi_T(X_T) \Rightarrow \pi^*(X_0, X_T) = \psi_0(X_0) r(X_T | X_0) \phi_T(X_T)$$

can be verified by (56) and (50). But in general, the transition kernel of the un-controlled process $r(\cdot | \cdot)$ (or ϕ_T, ψ_0) is not available analytically for implementation as such. But it suggests there's a backward-forward SDE dynamics one could write for ϕ_t, ψ_t with drift f_t initialized at ϕ_T, ψ_0 respectively, as hinted by (51) as well, if we view them as densities.

(6) Equation (52) means the optimal v_t^* in the continuity equation (43) is $f_t + \sigma^2 \nabla \log \phi_t - \frac{\sigma^2}{2} \nabla \log \rho_t = f_t + \frac{\sigma^2}{2} \nabla \log \frac{\phi_t}{\psi_t}$, which gives the ODE probability flow for this dynamics (with $X_0 \sim \nu$)

$$dX_t = f_t(X_t) + \frac{\sigma^2}{2} (\nabla \log \phi_t(X_t) - \nabla \log \psi_t(X_t)) dt, \quad (57)$$

in the sense that the time marginals $\rho_t^{(57)} = \rho_t^{(52)} = \rho_t^{(53)} = P_t^{(17)}$ all agree. Moreover, using the instantaneous change of variables formula (Chen et al., 2018), we also have

$$\log \rho_T(X_T) = \log \rho_0(X_0) - \int_0^T \nabla \cdot f_t(X_t) dt - \frac{\sigma^2}{2} \int_0^T \nabla \cdot (\nabla \log \phi_t(X_t) - \nabla \log \psi_t(X_t)) dt, \quad (58)$$

which will be useful in Proposition 3.5 for estimating the normalizing constant. Note that in (58) both the density and the point at which we are evaluating is changing.

Remark B.2. The stochastic control formulation makes it clear that the trajectory we are trying to recover is a meaningful one in the sense of minimal effort. If one were to switch the order of P and Q in (17), the optimal control problem becomes (e.g., for $f = 0$)

$$\begin{aligned} \inf_u \mathbb{E} \left[\int_0^T \frac{1}{2} \|u_t(X_t)\|^2 dt \right] \\ \text{s.t. } dX_t = \sigma u_t(X_t) dt + \sigma dW_t, X_0 \sim \nu, X_T \sim \mu \end{aligned}$$

for the expectation taken over the reference process, instead of the controlled state density ρ_t^u , which is not very intuitive. The slightly non-conventional aspect of this control problem is the fixed terminal constraint.

To briefly summarize, all these different viewpoints explore the deep connections between PDEs (controls) and SDEs (diffusions) in one way or another.

C. Missing Proofs

We give the proof of the likelihood ratio calculation Lemma C.1 below, followed by Lemma A.1 from Section A.

Lemma C.1 (Forward / Backward path-space likelihood ratio). *In the general case of reference*

$$dX_t = \sigma r_t^+(X_t) dt + \sigma d\overrightarrow{W}_t, X_0 \sim \Gamma_0 \quad dX_t = \sigma r_t^-(X_t) dt + \sigma d\overleftarrow{W}_t, X_T \sim \Gamma_T \quad (59)$$

where $\overrightarrow{\mathbb{P}}^{\Gamma_0, \sigma r^+} = \overleftarrow{\mathbb{P}}^{\Gamma_T, \sigma r^-}$, there are 4 terms in $D_{KL}(\overrightarrow{\mathbb{P}}^{\nu, f+\sigma u} \parallel \overleftarrow{\mathbb{P}}^{\mu, f+\sigma v})$ as

$$\begin{aligned} & \mathbb{E}_{X \sim \overrightarrow{\mathbb{P}}^{\nu, f+\sigma u}} \left[\log\left(\frac{d\nu}{d\Gamma_0}\right)(X_0) - \log\left(\frac{d\mu}{d\Gamma_T}\right)(X_T) + \log Z \right. \\ & + \frac{1}{2\sigma^2} \int_0^T (f_t + \sigma u_t - \sigma r_t^+)(X_t)^\top (f_t + \sigma u_t - \sigma r_t^+)(X_t) dt \\ & \left. - \frac{1}{\sigma^2} \int_0^T (f_t + \sigma v_t - \sigma r_t^-)(X_t)^\top \left(\frac{1}{2}f_t + \sigma u_t - \frac{\sigma}{2}v_t - \frac{\sigma}{2}r_t^-\right)(X_t) dt - \int_0^T \nabla \cdot (f_t + \sigma v_t - \sigma r_t^-)(X_t) dt \right]. \end{aligned}$$

By picking $r^+, r^- = 0$, and Lebesgue base measure for Γ_0, Γ_T , $D_{KL}(\overrightarrow{\mathbb{P}}^{\nu, f+\sigma u} \parallel \overleftarrow{\mathbb{P}}^{\mu, f+\sigma v})$ becomes

$$\mathbb{E}_{X \sim \overrightarrow{\mathbb{P}}^{\nu, f+\sigma u}} \left[\int_0^T \frac{1}{2} \|u_t(X_t) - v_t(X_t)\|^2 - \nabla \cdot (f_t + \sigma v_t)(X_t) dt + \log \frac{d\nu(X_0)}{d\mu(X_T)} \right] + \log Z.$$

The loss (14) enforce uniqueness (and correct marginals ν, μ) but doesn't impose optimality of the trajectory in any sense. The judicious choice of common v and ν in (1) and (12) allow (14) to take on a control-theoretic interpretation.

Proof of Lemma C.1. We give a few equivalent expressions for the Radon-Nikodym derivative, written in terms of the drifts. Starting with Proposition 2.2 of (Vargas & Nüsken, 2023), $\overrightarrow{\mathbb{P}}^{\nu, f+\sigma u}$ almost surely,

$$\begin{aligned} \log \left(\frac{d\overrightarrow{\mathbb{P}}^{\nu, f+\sigma u}}{d\overleftarrow{\mathbb{P}}^{\mu, f+\sigma v}} \right) (X) &= \log\left(\frac{d\nu}{d\Gamma_0}\right)(X_0) - \log\left(\frac{d\mu}{d\Gamma_T}\right)(X_T) + \log Z \\ &+ \frac{1}{\sigma^2} \int_0^T (f_t + \sigma u_t - \sigma r_t^+)(X_t) \left(\overrightarrow{dX}_t - \frac{1}{2}(f_t + \sigma u_t + \sigma r_t^+)(X_t) dt \right) \end{aligned} \quad (60)$$

$$- \frac{1}{\sigma^2} \int_0^T (f_t + \sigma v_t - \sigma r_t^-)(X_t) \left(\overleftarrow{dX}_t - \frac{1}{2}(f_t + \sigma v_t + \sigma r_t^-)(X_t) dt \right). \quad (61)$$

One can use (11) to convert the backward integral to a forward one with an additional divergence term

$$\begin{aligned} D_{KL}(\overrightarrow{\mathbb{P}}^{\nu, f+\sigma u} \parallel \overleftarrow{\mathbb{P}}^{\mu, f+\sigma v}) &= \mathbb{E}_{X \sim \overrightarrow{\mathbb{P}}^{\nu, f+\sigma u}} \left[\log \left(\frac{d\overrightarrow{\mathbb{P}}^{\nu, f+\sigma u}}{d\overleftarrow{\mathbb{P}}^{\mu, f+\sigma v}} \right) (X) \right] \\ &= \mathbb{E}_{X \sim \overrightarrow{\mathbb{P}}^{\nu, f+\sigma u}} \left[\log\left(\frac{d\nu}{d\Gamma_0}\right)(X_0) - \log\left(\frac{d\mu}{d\Gamma_T}\right)(X_T) + \log Z \right. \\ &+ \frac{1}{2\sigma^2} \int_0^T (f_t + \sigma u_t - \sigma r_t^+)(X_t)^\top (f_t + \sigma u_t - \sigma r_t^+)(X_t) dt \\ &- \frac{1}{\sigma^2} \int_0^T (f_t + \sigma v_t - \sigma r_t^-)(X_t)^\top \left(\frac{1}{2}f_t + \sigma u_t - \frac{\sigma}{2}v_t - \frac{\sigma}{2}r_t^-\right)(X_t) dt - \int_0^T \nabla \cdot (f_t + \sigma v_t - \sigma r_t^-)(X_t) dt \left. \right] \quad (62) \\ &+ \mathbb{E}_{X \sim \overrightarrow{\mathbb{P}}^{\nu, f+\sigma u}} \left[\int_0^T (u_t - r_t^+)(X_t) - (v_t - r_t^-)(X_t) d\overrightarrow{W}_t \right] \end{aligned}$$

and the last term vanishes because of (8). Therefore we see that there are 2 boundary terms, and 3 extra terms corresponding to the forward/backward process. The loss of (14) clearly follows since only the second line and one term in the first line survive. By picking $\gamma^+, \gamma^- = 0$, and Lebesgue base measure for Γ_0, Γ_T , we get

$$D_{KL}(\overrightarrow{\mathbb{P}}^{\nu, f+\sigma u} \parallel \overleftarrow{\mathbb{P}}^{\mu, f+\sigma v}) = \mathbb{E}_{X \sim \overrightarrow{\mathbb{P}}^{\nu, f+\sigma u}} \left[\int_0^T \frac{1}{2} \|u_t(X_t) - v_t(X_t)\|^2 - \nabla \cdot (f_t + \sigma v_t)(X_t) dt \right] \quad (63)$$

$$+ \mathbb{E}_{X \sim \overrightarrow{\mathbb{P}}^{\nu, f+\sigma u}} \left[\log \frac{d\nu(X_0)}{d\mu(X_T)} + \int_0^T (u_t - v_t)(X_t) d\overrightarrow{W}_t \right] + \log Z \quad (64)$$

which agrees with Proposition 2.3 in (Richter et al., 2023) up to a conventional sign in v , and is also the same as the ELBO loss in (Chen et al., 2021a, Theorem 4). Note the different expressions in (14) and (63)-(64) by picking distinct references (59) vs. (12).

One could also start with the forward SDE (10) and apply Girsanov's theorem with (2),(10) together with the chain rule for the KL to get

$$\begin{aligned} D_{KL}(\overrightarrow{\mathbb{P}}^{\nu, f+\sigma u} \parallel \overleftarrow{\mathbb{P}}^{\mu, f+\sigma v}) &= \mathbb{E}_{\overrightarrow{\mathbb{P}}^{\nu, f+\sigma u}} \left[\log \left(\frac{d\overrightarrow{\mathbb{P}}^{\nu, f+\sigma u}}{d\overleftarrow{\mathbb{P}}^{\mu, f+\sigma v}} \right) (X) \right] \\ &= \mathbb{E}_{\overrightarrow{\mathbb{P}}^{\nu, f+\sigma u}} \left[\log \left(\frac{d\nu}{d\overleftarrow{\mathbb{P}}_0^{\mu, f+\sigma v}} \right) (X_0) + \frac{1}{2} \int_0^T \|u_t(X_t) - (v_t + \sigma \nabla \log \rho_t^{\mu, f+\sigma v})(X_t)\|^2 dt \right]. \end{aligned} \quad (65)$$

To see this is equivalent to (63)-(64), we use Fokker-Planck on the process (10) to reach

$$\begin{aligned} \mathbb{E}_{\overrightarrow{\mathbb{P}}^{\nu, f+\sigma u}} \left[\log \left(\frac{d\mu}{d\overleftarrow{\mathbb{P}}_0^{\mu, f+\sigma v}} \right) (X) - \log Z \right] &= \mathbb{E}_{\overrightarrow{\mathbb{P}}^{\nu, f+\sigma u}} \left[\log \left(\frac{d\overleftarrow{\mathbb{P}}_T^{\mu, f+\sigma v}}{d\overleftarrow{\mathbb{P}}_0^{\mu, f+\sigma v}} \right) (X) \right] \\ &= \mathbb{E}_{\overrightarrow{\mathbb{P}}^{\nu, f+\sigma u}} \left[\int_0^T -\nabla \cdot (f_t + \sigma v_t)(X_t) + \sigma(u_t - v_t)(X_t)^\top \nabla \log \rho_t^{\mu, f+\sigma v}(X_t) - \frac{\sigma^2}{2} \|\nabla \log \rho_t^{\mu, f+\sigma v}(X_t)\|^2 dt \right] \end{aligned} \quad (66)$$

as by Itô's lemma for the process (2) with drift $f + \sigma u$,

$$\begin{aligned} \int_0^T \partial_t \log \overleftarrow{\mathbb{P}}_t^{\mu, f+\sigma v}(X_t^{f+\sigma u}) dt &= \int_0^T \frac{1}{\rho_t} \left(-\nabla \cdot (\rho_t(f + \sigma v)) + \frac{\sigma^2}{2} \Delta \rho_t + \nabla \rho_t^\top (f + \sigma u) \right) dt \\ &\quad + \sigma \int_0^T \frac{\nabla \rho_t^\top}{\rho_t} dW_t + \int_0^T \frac{\sigma^2}{2} \Delta \log \rho_t dt \\ &= \int_0^T -\frac{\nabla \rho_t^\top}{\rho_t} (f_t + \sigma v_t + \sigma^2 \nabla \log \rho_t - f_t - \sigma u_t) \\ &\quad - \nabla \cdot (f_t + \sigma v_t + \sigma^2 \nabla \log \rho_t) + \frac{\sigma^2}{2} \Delta \log \rho_t dt + \int_0^T \sigma \frac{\nabla \rho_t^\top}{\rho_t} dW_t, \end{aligned}$$

which upon simple arranging, using (69) and taking expectation over $X^{f+\sigma u} \sim \overrightarrow{\mathbb{P}}^{\nu, f+\sigma u}$ give (66). Now adding up the previous two displays (65) and (66) using

$$\mathbb{E}_{\overrightarrow{\mathbb{P}}^{\nu, f+\sigma u}} \left[\log \left(\frac{d\nu}{d\overleftarrow{\mathbb{P}}_0^{\mu, f+\sigma v}} \right) \right] = \mathbb{E}_{\overrightarrow{\mathbb{P}}^{\nu, f+\sigma u}} \left[\log \left(\frac{d\nu}{d\mu} \right) + \log \left(\frac{d\mu}{d\overleftarrow{\mathbb{P}}_0^{\mu, f+\sigma v}} \right) \right]$$

finishes the proof. \square

Proof of Lemma A.1. We have up to a constant independent of ϕ, ψ , using Lemma C.1, (37)-(38) give the loss

$$\begin{aligned} \psi_n &\leftarrow \arg \min_{\psi} \int_0^T \mathbb{E}_{X \sim \overrightarrow{\mathbb{P}}^{\nu, f+\sigma^2 \nabla \phi_{n-1}}} \left[\frac{\sigma^2 \|\nabla \psi\|^2}{2} - \sigma^2 \nabla \phi_{n-1}^\top \nabla \psi - \nabla \cdot (\sigma^2 \nabla \psi) \right] (X_t) dt \\ \phi_n &\leftarrow \arg \min_{\phi} \int_0^T \mathbb{E}_{X \sim \overrightarrow{\mathbb{P}}^{\nu, f+\sigma^2 \nabla \phi}} \left[\frac{\sigma^2}{2} \|\nabla \phi - \nabla \psi_n\|^2 - \nabla \cdot (f + \sigma^2 \nabla \psi_n) \right] (X_t) dt + \mathbb{E} \left[\log \frac{\nu(X_0)}{\mu(X_T)} \right] \end{aligned}$$

where updating ϕ bears resemblance to (14), and updating ψ is akin to the score matching loss (15). The equivalence to (36) assuming ϕ, ψ is expressive enough (i.e., $(\phi^*, \psi^*) = (\phi_{\theta^*}, \psi_{\gamma^*})$) and $\phi_0 = 0$ is shown in (Vargas & Nüsken, 2023) and relies on the fact that ϕ, ψ update (37)-(38) solve respectively

$$\nabla \psi_n = \nabla \phi_{n-1} - \nabla \log \rho_t^{\nu, f+\sigma^2 \nabla \phi_{n-1}}, \quad \nabla \phi_n = \nabla \psi_n - \nabla \log \rho_t^{\mu, f+\sigma^2 \nabla \psi_n},$$

which fix one side of the marginal alternatively, while reversing the path-wise transition from the previous one. It is known (36) converges to π^* as $n \rightarrow \infty$. \square

In Section 2.2 we also claimed that for two processes $(p_t)_t, (q_t)_t$ with different initialization, but same drift $\sigma v + \sigma^2 s$ where s is the score function, $\partial_t D_{KL}(p_t \| q_t) = -\sigma^2/2 \cdot \mathbb{E}_{p_t}[\|\nabla \log \frac{p_t}{q_t}\|^2] \leq 0$ contracts towards each other, when we run the generative process

$$dX_t = [\sigma v_t(X_t) + \sigma^2 s_t(X_t)]dt + \sigma d\overrightarrow{W}_t, \quad X_0 \sim \nu \neq \overleftarrow{\mathbb{P}}_0^{\mu, \sigma v},$$

although the actual rate maybe slow. We give this short calculation here for completeness:

$$\begin{aligned} & \partial_t D_{KL}(p_t \| q_t) \\ &= \int (\partial_t p_t(x)) \log \frac{p_t(x)}{q_t(x)} dx + \int q_t(x) \left[\frac{\partial_t p_t(x)}{q_t(x)} - \frac{p_t(x) \partial_t q_t(x)}{q_t^2(x)} \right] dx \\ &= \int p_t \left[\sigma v_t + \sigma^2 s_t - \frac{\sigma^2}{2} \nabla \log p_t(x) \right]^\top \nabla \log \frac{p_t(x)}{q_t(x)} dx - \int \frac{p_t(x)}{q_t(x)} \partial_t q_t(x) dx \\ &= \int p_t \left[\sigma v_t + \sigma^2 s_t - \frac{\sigma^2}{2} \nabla \log p_t(x) \right]^\top \nabla \log \frac{p_t(x)}{q_t(x)} dx - \int \nabla \log \frac{p_t(x)}{q_t(x)} \left[\sigma v_t + \sigma^2 s_t - \frac{\sigma^2}{2} \nabla \log q_t(x) \right] q_t(x) dx \\ &= \int \left[\sigma v_t + \sigma^2 s_t - \frac{\sigma^2}{2} \nabla \log p_t(x) - \sigma v_t - \sigma^2 s_t + \frac{\sigma^2}{2} \nabla \log q_t(x) \right]^\top \nabla \log \frac{p_t(x)}{q_t(x)} \cdot p_t(x) dx \\ &= -\frac{\sigma^2}{2} \int \left\| \nabla \log \frac{p_t(x)}{q_t(x)} \right\|^2 p_t(x) dx \leq 0. \end{aligned}$$

Immediately follows is the proof of our main result: Proposition 3.1, along with Lemma 3.2 from Section 3.1.

Proof of Proposition 3.1. The first term (19)-(20) simply involves writing D_{KL} in terms of the (current) controls using Lemma C.1, up to constants. One can regularize to remove the constraints in (41) while still ensuring the right marginals via

$$\mathcal{L}(\nabla \phi, \nabla \psi) := D_{KL}(\overrightarrow{\mathbb{P}}^{\nu, \sigma^2 \nabla \phi} \| \overleftarrow{\mathbb{P}}^{\mu, -\sigma^2 \nabla \psi}) + \lambda R(\nabla \psi) \quad \text{or} \quad D_{KL}(\overrightarrow{\mathbb{P}}^{\nu, \sigma^2 \nabla \phi} \| \overleftarrow{\mathbb{P}}^{\mu, -\sigma^2 \nabla \psi}) + \lambda' R'(\nabla \phi)$$

where $R(\cdot)$ can be either from reference measure as in (1) or from optimality condition on the control as in (2). In both cases, the first term is a ρ constraint (time reversal consistency with correct marginals), and the second one a ϕ -constraint (enforce optimality).

For (2), we give one direction of the argument for $\nabla \psi$ first. Let $\overrightarrow{\mathbb{P}}^{\nu, f}$ denote the path measure associated with $dX_t = f_t(X_t)dt + \sigma dW_t, X_0 \sim \nu$, then $\overrightarrow{\mathbb{P}}^{\nu, f}$ almost surely, the Radon-Nikodym derivative between $\overrightarrow{\mathbb{P}}^{\nu, f}$ and the controlled backward process is

$$\log \left(\frac{d\overrightarrow{\mathbb{P}}^{\nu, f}}{d\overleftarrow{\mathbb{P}}^{\mu, f - \sigma^2 \nabla \psi_t}} \right) (X) = \log \frac{d\nu(X_0)}{d\mu(X_T)} + \int_0^T \frac{\sigma^2}{2} \|\nabla \psi_t\|^2 - \nabla \cdot (f - \sigma^2 \nabla \psi_t) dt + \sigma \int_0^T \nabla \psi_t dW_t.$$

In the case when the variance (taken along the prior $X \sim \overrightarrow{\mathbb{P}}^{\nu, f}$)

$$\text{Var} \left(\psi_T(X_T) - \psi_0(X_0) - \int_0^T \frac{\sigma^2}{2} \|\nabla \psi_t\|^2 + \nabla \cdot f - \sigma^2 \Delta \psi_t(X_t) dt - \int_0^T \sigma \nabla \psi_t dW_t \right) = 0,$$

it implies that the random quantity is almost surely a constant independent of the realization, and

$$\log \left(\frac{d\overrightarrow{\mathbb{P}}^{\nu, f}}{d\overleftarrow{\mathbb{P}}^{\mu, f - \sigma^2 \nabla \psi_t}} \right) (X) = \log \frac{d\nu(X_0)}{d\mu(X_T)} + \psi_T(X_T) - \psi_0(X_0) + \log Z,$$

from which using the factorization characterization (55), and the terminal constraint $\overleftarrow{\mathbb{P}}_0^{\mu, f - \sigma^2 \nabla \psi_t} = \nu, \overleftarrow{\mathbb{P}}_T^{\mu, f - \sigma^2 \nabla \psi_t} = \mu$ imposed by the first KL condition, concludes that $\overleftarrow{\mathbb{P}}^{\mu, f - \sigma^2 \nabla \psi_t}$ must be the unique solution to the SB problem. The other direction on $\nabla \phi$ is largely similar, and one can show using Girsanov's theorem and the variance condition that along the prior $X \sim \overrightarrow{\mathbb{P}}^{\nu, f}$,

$$\log \left(\frac{d\overrightarrow{\mathbb{P}}^{\nu, f + \sigma^2 \nabla \phi_t}}{d\overleftarrow{\mathbb{P}}^{\nu, f}} \right) (X) = \int_0^T -\frac{\sigma^2}{2} \|\nabla \phi_t\|^2 dt + \int_0^T \sigma \nabla \phi_t dW_t = \phi_T(X_T) - \phi_0(X_0),$$

to conclude in the same way. Note that evaluating the variance regularizer only requires simulating from the base process (18).

In (3), for any ϕ_t , the likelihood ratio along $dX_t = (f + \sigma^2 \nabla \phi_t) dt + \sigma dW_t$, $X_0 \sim \nu$ is

$$\log \left(\frac{d\overrightarrow{\mathbb{P}}^{\nu, f + \sigma^2 \nabla \phi_t}}{d\overrightarrow{\mathbb{P}}^{\nu, f}} \right) (X) = \int_0^T \frac{\sigma^2}{2} \|\nabla \phi_t\|^2 dt + \int_0^T \sigma \nabla \phi_t dW_t, \quad (67)$$

which according to (55), has to be equal to $\psi_0(X_0) + \phi_T(X_T) - \log \nu(X_0) = -\phi_0(X_0) + \phi_T(X_T) \mathbb{P}^{\nu, f} a.s. \Rightarrow \mathbb{P}^{\nu, f + \sigma^2 \nabla \phi} a.s.$ to be optimal, justifying

$$\text{Var}_{X \sim \overrightarrow{\mathbb{P}}^{\nu, f + \sigma^2 \nabla \phi}} \left(\phi_T(X_T) - \phi_0(X_0) - \frac{\sigma^2}{2} \int_0^T \|\nabla \phi_t\|^2(X_t) dt - \sigma \int_0^T \nabla \phi_t(X_t) dW_t \right).$$

In a similar spirit, for any ψ_t , along the same process $X \sim \overrightarrow{\mathbb{P}}^{\nu, f + \sigma^2 \nabla \phi_t}$,

$$\begin{aligned} \log \left(\frac{d\overrightarrow{\mathbb{P}}^{\mu, f}}{d\overrightarrow{\mathbb{P}}^{\mu, f - \sigma^2 \nabla \psi_t}} \right) (X) &= \int_0^T \sigma^2 \nabla \psi_t^\top \nabla \phi_t + \frac{\sigma^2}{2} \|\nabla \psi_t\|^2 - \nabla \cdot (f_t - \sigma^2 \nabla \psi_t) dt + \int_0^T \sigma \nabla \psi_t dW_t \\ &+ \log \left(\frac{d\nu(X_0)}{d\mu(X_T)} \right) + \log Z, \end{aligned} \quad (68)$$

which again using (55), has to be equal to

$$-\psi_0(X_0) - \phi_T(X_T) + \log \nu(X_0) = -\psi_0(X_0) - \log p_{\text{target}}(X_T) + \psi_T(X_T) + \log \nu(X_0)$$

to be optimal, yielding the anticipated variance regularizer. Additionally, in order to enforce the terminal constraint along $\overrightarrow{\mathbb{P}}^{\nu, f + \sigma^2 \nabla \phi_t}$, it suffices to sum up (67)-(68) to impose the time-reversal consistency, which means $\phi_T(X_T) + \psi_T(X_T) = \log p_{\text{target}}(X_T)$ and $\phi_0(X_0) + \psi_0(X_0) = \log p_{\text{prior}}(X_0) = \log \nu(X_0)$ necessarily.

In (4), we use $\nabla \phi_t$ and $\nabla \log \rho_t$ as optimization variables instead of the two drifts, and it follows from the dynamical formulation (48)-(49). One can also read (20) as

$$\begin{aligned} \partial_t \log \rho_t(X_t) &= \frac{1}{\rho_t} (\partial_t \rho_t + \nabla \rho_t^\top \dot{X}_t) \\ &= \frac{1}{\rho_t} [-\nabla \cdot (\rho_t (f_t + \sigma^2 \nabla \phi_t)) + \frac{\sigma^2}{2} \Delta \rho_t] + \frac{1}{\rho_t} \nabla \rho_t^\top (f_t + \sigma^2 \nabla \phi_t - \frac{\sigma^2}{2} \nabla \log \rho_t) \\ &= -\sigma^2 \Delta \phi_t + \frac{\sigma^2}{2} \Delta \log \rho_t - \nabla \cdot f_t, \end{aligned}$$

which upon integrating and using integration by parts (16) give the loss. This first part establishes a particular relationship between the two variables, and the second part enforces optimality among all curves transporting between ν and μ driven by such dynamics – one can also switch this part to (42) if \bar{v}_t is known. \square

Proof of Lemma 3.2. Using Itô's lemma, we have along the reference SDE $dX_t = f_t(X_t) dt + \sigma dW_t$, $X_0 \sim \nu$,

$$d\phi_t = \left[\frac{\partial \phi_t}{\partial t} + \nabla \phi_t^\top f_t + \frac{\sigma^2}{2} \Delta \phi_t \right] dt + \sigma \nabla \phi_t^\top dW_t.$$

Now deducing from (51), since the optimal ϕ_t solves the PDE for all $(t, x) \in [0, T] \times \mathbb{R}^d$

$$\partial_t \phi_t = -f_t^\top \nabla \phi_t - \frac{\sigma^2}{2} \Delta \phi_t - \frac{\sigma^2}{2} \|\nabla \phi_t\|^2,$$

substituting the last display into the previous one gives the result. Analogously, along the same reference process with Itô's lemma,

$$d\psi_t = \left[\frac{\partial \psi_t}{\partial t} + \nabla \psi_t^\top f_t + \frac{\sigma^2}{2} \Delta \psi_t \right] dt + \sigma \nabla \psi_t^\top dW_t$$

and using the fact (51) that the optimal ψ_t solves for all $(t, x) \in [0, T] \times \mathbb{R}^d$

$$\partial_t \psi_t = -\nabla \psi_t^\top f_t - \nabla \cdot f_t + \frac{\sigma^2}{2} (\|\nabla \psi_t\|^2 + \Delta \psi_t)$$

and plugging into the previous display yields the claim. In both of the PDE derivations above, we used the fact that for any $g: \mathbb{R}^d \rightarrow \mathbb{R}$,

$$\frac{1}{g} \nabla^2 g = \nabla^2 \log g + \frac{1}{g^2} \nabla g \nabla g^\top \Rightarrow \frac{1}{g} \Delta g = \Delta \log g + \|\nabla \log g\|^2 \quad (69)$$

by taking trace on both sides.

The second part of the lemma statement, where X_t evolves along the optimally controlled SDE, follows from (Chen et al., 2021a, Theorem 3) up to a change of variable. Notice the sign change and the absence of the cross term in the dynamics for ϕ_t, ψ_t when $X \sim \overrightarrow{\mathbb{P}}^{\nu, f}$ vs. $X \sim \overrightarrow{\mathbb{P}}^{\nu, f + \sigma^2 \nabla \phi}$. \square

Proposition 3.5 and Lemma 3.6 are stated in Section 3.2, whose proof we give below.

Proof of Proposition 3.5. The optimal $\nabla \phi_t, \nabla \psi_t$ allow us to calculate $\log Z$ for p_{target} using (19) as either (can be used without expectation, or with expectation and ignore the last term)

$$\begin{aligned} \log Z &= \mathbb{E}_{\overrightarrow{\mathbb{P}}^{\nu, f + \sigma^2 \nabla \phi}} \left[-\frac{\sigma^2}{2} \int_0^T \|\nabla \phi_t(X_t) + \nabla \psi_t(X_t)\|^2 dt + \int_0^T \nabla \cdot (f_t(X_t) - \sigma^2 \nabla \psi_t(X_t)) dt - \log \frac{d\nu(X_0)}{d\mu(X_T)} \right] \\ &+ \mathbb{E}_{\overrightarrow{\mathbb{P}}^{\nu, f + \sigma^2 \nabla \phi}} \left[-\sigma \int_0^T (\nabla \phi_t + \nabla \psi_t)(X_t) d\overrightarrow{W}_t \right] =: \mathbb{E}[-S], \end{aligned}$$

or using (58), with the optimal $\nabla \phi^*, \nabla \psi^*$, since Z is independent of X_T ,

$$-\log Z = \log \nu(X_0) - \frac{\sigma^2}{2} \int_0^T \nabla \cdot (\nabla \log \phi_t^* - \nabla \log \psi_t^*)(X_t) dt - \int_0^T \nabla \cdot f_t(X_t) dt - \log \mu(X_T), \quad (70)$$

in which case the estimator is exact with X_t following (52), or equivalently, (57). Notice (70) bears resemblance to Jarzynski's identity, since the term that's been integrated is essentially $\partial_t \log \rho_t(X_t)$.

In general for imperfect control, since $D_{KL}(\overrightarrow{\mathbb{P}}^{\nu, f + \sigma^2 \nabla \phi} \parallel \overleftarrow{\mathbb{P}}^{\mu, f - \sigma^2 \nabla \psi}) > 0$, $\log Z$ will only be lower bounded by $\mathbb{E}[-S]$. Using Lemma C.1 however,

$$\begin{aligned} 1 &= \mathbb{E}_{\overrightarrow{\mathbb{P}}^{\nu, f + \sigma^2 \nabla \phi}} \left[\left(\frac{d\overrightarrow{\mathbb{P}}^{\nu, f + \sigma^2 \nabla \phi}}{d\overleftarrow{\mathbb{P}}^{\mu, f - \sigma^2 \nabla \psi}} \right)^{-1} \right] \\ &= \mathbb{E}_{\overrightarrow{\mathbb{P}}^{\nu, f + \sigma^2 \nabla \phi}} \left[\exp \left(-\frac{\sigma^2}{2} \int_0^T \|\nabla \phi_t + \nabla \psi_t\|^2 + \nabla \cdot (f_t - \sigma^2 \nabla \psi_t) dt - \sigma \int_0^T \nabla \phi_t + \nabla \psi_t dW_t - \log \frac{\nu}{\mu} \right) \frac{1}{Z} \right] \quad (71) \\ &=: \mathbb{E}[\exp(-S')/Z], \end{aligned}$$

giving $Z = \mathbb{E}[\exp(-S')]$ using any (potentially sub-optimal) control $\nabla \phi, \nabla \psi$.

For the importance sampling, we use path weights suggested by the terminal requirement that $X_T^{\phi^*} \sim \mu$, therefore for any ϕ and $X \sim \overrightarrow{\mathbb{P}}^{\nu, f + \sigma^2 \nabla \phi}$,

$$\begin{aligned} w^\phi(X) &= \frac{d\mathbb{P}_{X^{\phi^*}}}{d\mathbb{P}_{X^\phi}}(X) = \frac{d\mathbb{P}_{X^{\phi^*}}}{d\mathbb{P}_{X^r}}(X) \frac{d\mathbb{P}_{X^r}}{d\mathbb{P}_{X^\phi}}(X) \\ &= \frac{d\mu}{d\mathbb{P}_{X_T^r}}(X_T) \frac{d\mathbb{P}_{X^r}}{d\mathbb{P}_{X^\phi}}(X) \\ &= \exp \left(\log \frac{d\mu}{d\mathbb{P}_{X_T^r}}(X_T) - \frac{1}{2\sigma^2} \int_0^T \|f_t + \sigma^2 \nabla \phi_t\|^2 dt - \frac{1}{\sigma} \int_0^T (f_t + \sigma^2 \nabla \phi_t) dW_t \right), \quad (72) \end{aligned}$$

where we assumed the reference r is drift-free:

$$dX_t = \sigma dW_t, \quad X_0 \sim \nu$$

hence $X_T^r \sim \nu * \mathcal{N}(0, \sigma^2 T \cdot I) \stackrel{d}{=} y + \sigma\sqrt{T}z$ for $y \sim \nu$, and $\mathbb{P}_{X_T^r}$ is easy to evaluate. Indeed, this choice guarantees

$$w^\phi(X_T^\phi) = \frac{d\mu}{d\mathbb{P}_{X_T^\phi}}(X_T^\phi)$$

therefore $\mathbb{E}_\phi[g(X_T^\phi)w^\phi(X_T^\phi)] = \int g(x)\mu(x)dx$ for any function g on the terminal variable X_T generated with the control ϕ . \square

Proof of Lemma 3.6. Similar to (22), we obtain the forward trajectory

$$X_{k+1}^i = X_k^i + (f_k(X_k^i) + \sigma^2 \nabla \phi_k(X_k^i))(t_{k+1} - t_k) + \sigma \sqrt{t_{k+1} - t_k} \cdot z_k^i, \quad z_k^i \sim \mathcal{N}(0, I) \quad (73)$$

with Euler-Maruyama for each of the $i \in [N]$ samples. Rewriting the term (68) using (11), (60)-(61) and proceeding with the approximation (6) give

$$\begin{aligned} & \frac{1}{\sigma^2} \int_0^T f_t(X_t) d\overrightarrow{X}_t - \frac{1}{2\sigma^2} \int_0^T \|f_t(X_t)\|^2 dt - \frac{1}{\sigma^2} \int_0^T (f_t - \sigma^2 \nabla \psi_t)(X_t) d\overleftarrow{X}_t + \frac{1}{2\sigma^2} \int_0^T \|(f_t - \sigma^2 \nabla \psi_t)(X_t)\|^2 dt \\ & \approx \frac{1}{\sigma^2} \sum_{k=0}^{K-1} f_k(X_k)^\top (X_{k+1} - X_k) - \frac{1}{2\sigma^2} \|f_k(X_k)\|^2 (t_{k+1} - t_k) \\ & - \frac{1}{\sigma^2} (f_{k+1}(X_{k+1}) - \sigma^2 \nabla \psi_{k+1}(X_{k+1}))(X_{k+1} - X_k) + \frac{1}{2\sigma^2} \|f_{k+1}(X_{k+1}) - \sigma^2 \nabla \psi_{k+1}(X_{k+1})\|^2 (t_{k+1} - t_k) \\ & = \sum_{k=0}^{K-1} \frac{1}{2\sigma^2 (t_{k+1} - t_k)} \|X_k - X_{k+1} + (f_{k+1} - \sigma^2 \nabla \psi_{k+1})(X_{k+1})(t_{k+1} - t_k)\|^2 \\ & - \sum_{k=0}^{K-1} \frac{1}{2\sigma^2 (t_{k+1} - t_k)} \|X_{k+1} - X_k - f_k(X_k)(t_{k+1} - t_k)\|^2, \end{aligned}$$

which is the estimator for the variance regularizer with (X_k) following (73).

It is also possible to avoid the divergence (i.e., Hutchinson's trace estimator) in the estimator (71) when discretizing by leveraging similar ideas. Direct computation using (60)-(61) tell us

$$Z \approx \frac{1}{N} \sum_{i=1}^N \exp \left(\log \frac{\mu}{\nu} - \sum_{k=0}^{K-1} \frac{1}{2\sigma^2 (t_{k+1} - t_k)} \|X_k^i - X_{k+1}^i + (f_{k+1} - \sigma^2 \nabla \psi_{k+1})(X_{k+1}^i)(t_{k+1} - t_k)\|^2 + \frac{1}{2} \|z_k\|^2 \right)$$

for the same Euler-Maruyama trajectory (73) as the estimator for the normalizing constant, where we used that the additional term

$$\sum_{k=0}^{K-1} \frac{1}{2\sigma^2 (t_{k+1} - t_k)} \|X_{k+1} - X_k - (f_k + \sigma^2 \nabla \phi_k)(X_k)(t_{k+1} - t_k)\|^2 = \frac{1}{2} \sum_{k=0}^{K-1} \|z_k\|^2$$

from the update. \square

D. Numerical Experiments

We provide additional details on the numerical experiments in this section.

D.1. Benchmark and Metrics

We consider 4 targets with different priors ν :

- 2D standard Gaussian: $\mathcal{N}(x; 0, I)$ with prior $\nu(x) = \mathcal{N}(x; 0, 2I)$

- Funnel: $x_0 \sim \mathcal{N}(0, 3^2)$, $x_1|x_0 \sim \mathcal{N}(0, \exp(x_0))$ with prior $\nu(x) = \mathcal{N}(x; 0, 2I)$
- Gaussian mixture model with 9 modes: $\frac{1}{9} \sum_{i=1}^9 \mathcal{N}(x; \mu_i, I)$ where $\{\mu_i\}_{i=1}^9 = \{-5, 0, 5\} \times \{-5, 0, 5\}$ with prior $\nu(x) = \mathcal{N}(x; 0, 3.5^2 I)$
- Double well: $\mu(x) \propto -(x_0^2 - 2)^2 - (x_1^2 - 2)^2$ with prior $\nu(x) = \mathcal{N}(0, 2I)$

For each of the benchmarks and 4 losses, we plot the marginals and report the following:

- Absolute error in mean and relative error in standard deviation compared to the ground truth using importance-weighted Monte-Carlo estimates (35)
- The log normalizing constant log Z estimator (c.f. Lemma 3.6):

$$\log \left(\frac{1}{n} \sum_{i=1}^n \frac{\mu(x_{K+1}^i)}{\nu(x_0^i)} \exp \left[\sum_{k=0}^K \frac{1}{2} \|Z_k^i\|^2 - \frac{1}{2\sigma^2 h} \|x_k^i - x_{k+1}^i + \sigma^2 h (\nabla \phi(x_{k+1}^i, (k+1)h) - \nabla s(x_{k+1}^i, (k+1)h))\|^2 \right] \right)$$

for the PINN, Variance and TD losses, and

$$\log \left(\frac{1}{n} \sum_{i=1}^n \frac{\mu(x_{K+1}^i)}{\nu(x_0^i)} \exp \left[\sum_{k=0}^K \frac{1}{2} \|Z_k^i\|^2 - \sum_{k=0}^K \frac{1}{2\sigma^2 h} \|x_k^i - x_{k+1}^i - \sigma^2 h \nabla \psi(x_{k+1}^i, (k+1)h)\|^2 \right] \right)$$

for the separate control loss, both using the final trajectory $\{x_k^i\}$ from (24).

D.2. Result

For the experiments, we parameterize $\phi, s : \mathbb{R}^d \times [0, c] \rightarrow \mathbb{R}$ as residual feed-forward neural networks where c is a hyperparameter. Each network first maps the input vector to a hidden state with a linear transformation, which is then propagated through several residual layers that maintain hidden state dimensionality. The network maps the final hidden state to the output vector. Each residual layer f receives as input a state-time pair $(x, t) \in \mathbb{R}^{d+1}$ and outputs $f(x, t) = \sigma(W(x, t) + b) + x$, where σ is the ReLU activation function. The number of hidden layers and their sizes vary by target distribution. Adam optimizer was used to train the models with $\beta_1 = 0.9, \beta_2 = 0.999$, and weight decay 0.01, where batches of trajectories are used for several steps of gradient updates in each epoch, before regenerating the n trajectories and estimating the objective for the next round of updates on the NN parameters. Across all experiments, we initialize $\phi(x_0^i, 0) \approx \log \mu(x_0^i)$. For fair comparison, the training process is stopped when the loss stops noticeably decreasing.

In practice we find the log-variance divergence (26) to perform better than the KL divergence, therefore proceed with this choice for the losses (25), (28), and (30).

For a comparison on the computation speed of the four regularizers: for the standard normal target it took 624 seconds to train with PINN, 108 seconds with the Variance regularizer, 106 seconds with TD, and 85 seconds with Separate Control. Generating trajectories requires much less time than evaluating the loss and its gradient. Computing each loss requires processing all K states across n trajectories. Combined with the number of training epochs and the number of updates per batch, we estimate the processing time per trajectory state as approximately

$$\frac{\text{training_time}}{\text{epochs} \times \text{trajectories} \times \text{updates_per_batch} \times K}.$$

The stated processing time is therefore roughly $6.2 \cdot 10^{-3}$ s for PINN, $3.6 \cdot 10^{-7}$ s for Variance, $3.5 \cdot 10^{-7}$ s for TD, and $2.8 \cdot 10^{-7}$ s for Separate Control. The latter three are comparable and much faster than PINN whose Laplacian computation adds an order of magnitude processing time. This efficiency comparison is also independent of the target distribution.

For reproducibility, the anonymous Github repository can be found at the following link:

<https://anonymous.4open.science/r/diffusion-sampler-E4F7>

where hyperparameters used for the experiments are listed in the corresponding notebooks.

In Table 2, Table 3 and Figure 2, Figure 3, Figure 4 below we show the simulation results.

	PINN	Variance	TD	Separate control
Standard normal	0.246	0.154	0.156	0.158
Funnel	0.796	0.781	0.786	0.829
GMM	0.026	0.651	0.130	0.025
Double well	0.503	0.404	0.346	0.242

Table 2. Relative error of weighted empirical standard deviation with respect to ground truth standard deviations (lower is better)

	PINN	Variance	TD	Separate control
Standard normal	0.124	0.038	0.014	0.031
Funnel	0.794	0.303	0.262	0.162
GMM	0.785	0.039	0.084	0.103
Double well	0.234	0.736	0.166	0.362

Table 3. Absolute error of empirical mean with respect to ground truth mean (lower is better)




Figure 2. Training loss plot for GMM with TD and Separate Control




Figure 3. GMM marginal before (left) and after (right) importance weighting

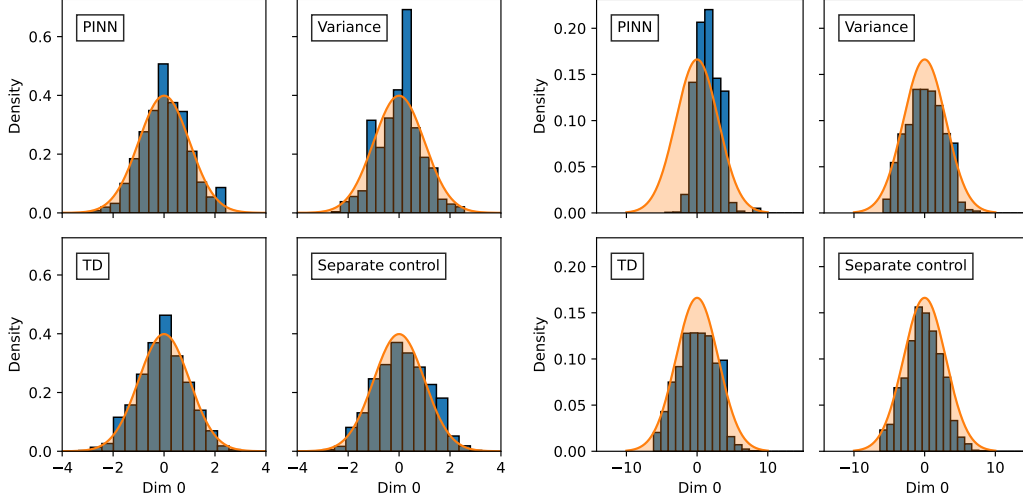


Figure 4. Standard normal marginal (left) and Funnel marginal (right) with importance weighting

E. Second order dynamics

The motivation of this section is to design a smoother trajectory for X_t while still maintaining optimality (and therefore uniqueness) of the dynamics. Suppose the reference Q is given by the following augmented process with a velocity variable $V_t \in \mathbb{R}^d$, and (X_t, V_t) have as the stationary distribution $\mathcal{N}(0, I) \otimes \mathcal{N}(0, I)$:

$$dX_t = V_t dt \quad (74)$$

$$dV_t = -X_t dt - \gamma V_t dt + \sqrt{2\gamma} d\vec{W}_t \quad (75)$$

We assume it is initialized at $Z_0 := (X_0, V_0) \sim \nu \otimes \mathcal{N}(0, I)$ and would like to enforce $Z_T := (X_T, V_T) \sim \mu \otimes \mathcal{N}(0, I)$ at the terminal.

In this underdamped case (Dockhorn et al., 2022), to solve the optimal trajectory problem, we formulate P as the controlled process

$$\inf_{\rho, u} \int_{\mathbb{R}^d} \int_0^T \frac{1}{2} \|u_t(Z_t)\|^2 \rho_t(Z_t) dt dx \quad (76)$$

$$\text{s.t. } dX_t = V_t dt \quad (77)$$

$$dV_t = -[X_t + u_t(Z_t)]dt - \gamma V_t dt + \sqrt{2\gamma} d\vec{W}_t, \quad \rho_0(Z_0) = \nu \otimes \mathcal{N}(0, I), \quad \rho_T(Z_T) = \mu \otimes \mathcal{N}(0, I) \quad (78)$$

One could check via the Girsanov's theorem that by picking

$$\bar{b}_t = \begin{bmatrix} V_t \\ -X_t - \gamma V_t \end{bmatrix}, \quad b_t = \begin{bmatrix} V_t \\ -X_t - u_t(Z_t) - \gamma V_t \end{bmatrix}, \quad G_t = \begin{bmatrix} 0 & 0 \\ 0 & \sqrt{2\gamma} I \end{bmatrix}$$

since $\bar{b}_t - b_t \in \text{image}(G_t)$, the KL divergence between the two path measures is $\mathbb{E}[\int_0^T \frac{1}{4\gamma} \|u_t(Z_t)\|^2 dt]$, which gives the minimum control effort, and agrees with the SB goal of

$$\min_{P_0, X \sim \nu, P_T, X \sim \mu} D_{KL}(P \| Q).$$

From now on we will take $\gamma = 1/2$ for simplicity and denote the corresponding path measure (77)-(78) as $\vec{\mathbb{P}}$.

We derive below a corresponding PINN and a variance regularizer for this dynamics, first on the forward drift followed by the backward drift. Writing out the optimality condition for the optimization problem (76), we have for a Lagrange

multiplier $\lambda(Z_t, t)$, since by continuity equation the dynamics (77)-(78) can be rewritten as

$$\partial_t \rho_t(Z_t) + \nabla_{[X, V]} \cdot \left(\rho_t(Z_t) \left[-X_t - u_t(Z_t) - 1/2 V_t - 1/2 \nabla_V \log \rho_t(Z_t) \right] \right) = 0,$$

we seek to optimize the Lagrangian

$$\int_{\mathbb{R}^d} \int_0^T \frac{1}{2} \|u_t(Z_t)\|^2 \rho_t(Z_t) + \lambda(Z_t, t) \left[\partial_t \rho + \nabla_{[X, V]} \cdot \left(\rho_t(Z_t) \left[-X_t - u_t(Z_t) - 1/2 V_t - 1/2 \nabla_V \log \rho_t(Z_t) \right] \right) \right] dt dx.$$

Performing integration by parts, we have

$$\int_{\mathbb{R}^d} \int_0^T \left\{ \frac{1}{2} \|u_t(Z_t)\|^2 - \partial_t \lambda(Z_t, t) - \nabla_{[X, V]} \lambda(Z_t, t)^\top \left[-X_t - u_t(Z_t) - 1/2 V_t \right] - 1/2 \Delta_V \lambda(Z_t, t) \right\} \rho_t(Z_t) dt dx.$$

Fixing ρ , the optimal control $u_t \in \mathbb{R}^d$ satisfies

$$u_t(Z_t) = -\nabla_V \lambda(Z_t, t). \quad (79)$$

Plugging this choice back in, the optimal $\lambda(Z_t, t)$ must satisfies the following PDE:

$$\frac{1}{2} \|\nabla_V \lambda(Z_t, t)\|^2 - \partial_t \lambda(Z_t, t) - \nabla_X \lambda(Z_t, t)^\top V_t + \nabla_V \lambda(Z_t, t)^\top (X_t - \nabla_V \lambda(Z_t, t) + \frac{1}{2} V_t) - \frac{1}{2} \Delta_V \lambda(Z_t, t) = 0 \quad (80)$$

where the $\Delta_V \lambda(Z_t, t)$ term is interpreted as $\text{Trace}(\nabla_V^2 \lambda(Z_t, t))$. Turning to the variance loss, we use Itô's lemma to get that along the reference SDE (74)-(75),

$$d\lambda_t(Z_t) = \left[\frac{\partial \lambda_t(Z_t)}{\partial t} + \nabla_X \lambda_t(Z_t)^\top V_t - \nabla_V \lambda_t(Z_t)^\top (X_t + 1/2 V_t) + \frac{1}{2} \Delta_V \lambda_t(Z_t) \right] dt + \nabla_V \lambda_t(Z_t)^\top dW_t,$$

which by plugging in the PDE (80), we end up with

$$d\lambda_t(Z_t) = -\frac{1}{2} \|\nabla_V \lambda(Z_t, t)\|^2 dt + \nabla_V \lambda(Z_t, t)^\top dW_t.$$

Compared to the PDE (80) it simplifies tremendously and implies the regularizer on the forward control as

$$\text{Var} \left(\lambda(Z_T, T) - \lambda(Z_0, 0) + \int_0^T \frac{1}{2} \|\nabla_V \lambda(Z_t, t)\|^2 dt - \int_0^T \nabla_V \lambda(Z_t, t)^\top dW_t \right). \quad (81)$$

Once we have the optimal control $u(Z_t, t)$ from (79), one simply simulate (77)-(78) from Z_0 to draw samples from μ .

The Radon-Nikodym likelihood ratio from Lemma C.1 still applies, therefore to impose path measure consistency, we compute for a backward control $o(Z_t, t)$ for the process $\overleftarrow{\mathbb{P}}$:

$$dX_t = V_t dt \quad (82)$$

$$dV_t = -[X_t - o_t(Z_t)] dt - \gamma V_t dt + \sqrt{2\gamma} \overleftarrow{dW}_t \quad (83)$$

the KL divergence $D_{KL}(\overrightarrow{\mathbb{P}} \| \overleftarrow{\mathbb{P}})$ is

$$\mathbb{E}_{X \sim \overrightarrow{\mathbb{P}}} \left[\int_0^T \frac{1}{2} \|b_t(Z_t) + \hat{b}_t(Z_t)\|^2 - \nabla \cdot \hat{b}_t(Z_t) dt + \log \frac{\nu(X_0) \otimes \mathcal{N}(V_0)}{\mu(X_T) \otimes \mathcal{N}(V_T)} \right] + C \quad (84)$$

for $\gamma = 1/2$ and

$$\hat{b}_t = \begin{bmatrix} V_t \\ -X_t + o_t(Z_t) - \gamma V_t \end{bmatrix},$$

that we minimize over $u(Z_t, t), o(Z_t, t)$ taking gradient forms.

This also implies the PDE regularizer on $o(Z_t, t) = -\nabla_V \gamma(Z_t, t)$ should read as (via Hopf-Cole transform)

$$\partial_t \gamma(Z_t, t) = -\nabla_{[X, V]} \gamma(Z_t, t)^\top \bar{b}_t - \nabla \cdot \bar{b}_t + \frac{1}{2} \|\nabla_V \gamma(Z_t, t)\|^2 + \frac{1}{2} \Delta_V \gamma(Z_t, t) \quad (85)$$

and the corresponding SDE regularizer along the reference (74)-(75) becomes (using $\nabla \cdot \bar{b}_t = -d/2$):

$$d\gamma_t(Z_t) = \left[\frac{1}{2} \|\nabla_V \gamma(Z_t, t)\|^2 + \Delta_V \gamma(Z_t, t) + d/2 \right] dt + \nabla_V \gamma(Z_t, t)^\top dW_t.$$

Moreover the boundary condition should still match the marginal densities at $t = 0, T$ in the sense that time reversal imposes $u_t(Z_t) + o_t(Z_t) = \nabla_V \log \rho_t(Z_t)$ for all $t \in [T]$.

The deterministic ODE implementation of the dynamics with given $(u_t(Z_t), o_t(Z_t))$ therefore becomes

$$dZ_t = \begin{bmatrix} dX_t \\ dV_t \end{bmatrix} = \begin{bmatrix} V_t \\ -X_t - 1/2u_t(Z_t) - 1/2V_t + 1/2o_t(Z_t) \end{bmatrix} dt. \quad (86)$$

Remark E.1. We would like to note that the reference process (74)-(75) still gives an entropy regularized optimal transport objective with quadratic cost, since

$$\begin{aligned} & \arg \min_{P_0 \sim \nu \otimes \mathcal{N}, P_T \sim \mu \otimes \mathcal{N}} KL(P_{0T} \| Q_{0T}) \\ &= \arg \min_{P_0 \sim \nu \otimes \mathcal{N}, P_T \sim \mu \otimes \mathcal{N}} \mathbb{E}_{P_{0T}}[-\log Q_{T|0}] - \int P(Z_0, Z_T) \log P(Z_0, Z_T) dZ_0 dZ_T \\ &= \arg \min_{P_0 \sim \nu \otimes \mathcal{N}, P_T \sim \mu \otimes \mathcal{N}} \mathbb{E}_{P_{0T}}[(Z_T - M_{0T} Z_0)^\top D_{T|0}^{-1} (Z_T - M_{0T} Z_0)] - H(P_{0T}) \end{aligned}$$

because the conditional distribution of Z_T given Z_0 is Gaussian, which we denote with mean $M_{0T} Z_0$ and covariance $D_{T|0}$. Using the fact that the reference is a linear SDE that admits analytical solution, assuming we initialize $(X_0, V_0) \sim \mathcal{N}\left(\begin{pmatrix} \mu_X^0 \\ \mu_V^0 \end{pmatrix}, \begin{pmatrix} \sigma_{XX}^0 \cdot I & 0 \\ 0 & \sigma_{VV}^0 \cdot I \end{pmatrix}\right)$ the dynamics will remain Gaussian with

$$\begin{aligned} \mathbb{E}[Z_T] &= \begin{bmatrix} (4T+1)\mu_X^0 + 64T\mu_V^0 \\ -T\mu_X^0 + (1-4T)\mu_V^0 \end{bmatrix} e^{-4T} = \begin{bmatrix} (4T+1)e^{-4T} \cdot I & 64Te^{-4T} \cdot I \\ -Te^{-4T} \cdot I & (1-4T)e^{-4T} \cdot I \end{bmatrix} \begin{pmatrix} \mu_X^0 \\ \mu_V^0 \end{pmatrix} \rightarrow 0_{2d} \\ \text{Cov}[Z_T] &= \begin{bmatrix} \sigma_{XX}^T \cdot I & \sigma_{XV}^T \cdot I \\ \sigma_{VX}^T \cdot I & \sigma_{VV}^T \cdot I \end{bmatrix} e^{-8T} \end{aligned}$$

for

$$\begin{aligned} \sigma_{XX}^T &= (1 + 8T + 16T^2)\sigma_{XX}^0 + 256T^2\sigma_{VV}^0 + e^{8T} - 1 - 8T - 32T^2 \rightarrow 1 \\ \sigma_{XV}^T &= (-T - 4T^2)\sigma_{XX}^0 + (16T - 64T^2)\sigma_{VV}^0 + 8T^2 \rightarrow 0 \\ \sigma_{VV}^T &= T^2\sigma_{XX}^0 + (1 + 16T^2 - 8T)\sigma_{VV}^0 - 2T^2 + T/2 + 1/16(e^{8T} - 1) \rightarrow 1/16 \end{aligned}$$

and plugging in $\mu_X^0 = x_0, \sigma_{XX}^0 = \sigma_{VV}^0 = 0, \mu_V^0 = v_0$ gives the desired M_{0T} and $D_{T|0}$.

Note that the coupling $P_{0T}(X_0, V_0, X_T, V_T) = \pi^*(X_0, X_T)\mathcal{N}(V_0)\mathcal{N}(V_T)$ for $\int \pi^*(X_0, X_T)dX_T = \nu(X_0), \int \pi^*(X_0, X_T)dX_0 = \mu(X_T)$ is always valid, in which case the objective will simplify to a quadratic involving (X_0, X_T) only since $\mathbb{E}_{P_{0T}}[V_0 V_T] = \mathbb{E}_{P_{0T}}[V_0 X_T] = \mathbb{E}_{P_{0T}}[X_0 V_T] = 0$ and the constraints already enforce $\mathbb{E}_{P_{0T}}[\|V_T\|^2]$ and $\mathbb{E}_{P_{0T}}[\|V_0\|^2]$ to be fixed and $\mathbb{E}_{P_{0T}}[X_T V_T] = \mathbb{E}_{P_{0T}}[X_0 V_0] = 0$. Similarly the entropy term will also only involve (X_0, X_T) in this case.

When both ν and μ are Gaussian with arbitrary covariance, one can check using results from (Janati et al., 2020) that the optimal coupling takes the form of $P_{0T}^*(X_0, V_0, X_T, V_T) = \pi^*(X_0, X_T)\bar{\pi}^*(V_0, V_T)$ where π^* and $\bar{\pi}^*$ are respectively the entropy-regularized optimal coupling between the X -variables and the V -variables.