

A. Existence of Unified Flow Operator

Theorem 2.3. *For the optimal control problem in Eq. (13) and Eq. (14), there exists an open-loop control $w^* = w^*(q(\mathbf{x}, 0), t)$ such that the induced state $q^*(\mathbf{x}, t)$ satisfies $q^*(\mathbf{x}, \infty) = p(\mathbf{x}|\mathcal{O}_{m+1})$. Moreover, w^* has a fixed expression with respect to $p(\mathbf{x}|\mathcal{O}_m)$ and $p(o_{m+1}|\mathbf{x})$ across different m .*

Proof. By Theorem 2.2, $\tilde{w}^*(q(\mathbf{x}, t), t) := \nabla_x \log q(\mathbf{x}, t)$ can induce the optimal state $\tilde{q}^*(\mathbf{x}, \infty) = p(\mathbf{x}|\mathcal{O}_{m+1})$ and achieve a zero loss, $d = 0$. Hence, \tilde{w}^* is an optimal closed-loop control for this problem.

Although in general closed-loop control has a stronger characterization to the solution, in a deterministic system like Eq. (14), the optimal closed-loop control and the optimal open-loop control will give the same control law and thus the same optimality to the loss function (Dreyfus, 1964). Hence, there exists an optimal open-loop control $w^* = w^*(q(\mathbf{x}, 0), t)$ such that the induced optimal state also gives a zero loss and thus $q^*(\mathbf{x}, \infty) = p(\mathbf{x}|\mathcal{O}_{m+1})$.

More specifically, when the system is deterministic, a state $q(\mathbf{x}, t)$ is just a deterministic result of the initial state $q(\mathbf{x}, 0)$ and the dynamics. The optimal flow determined by $\tilde{w}^*(q(\mathbf{x}, t), t)$ is

$$f = \nabla_x \log p(\mathbf{x}|\mathcal{O}_m)p(o_{m+1}|\mathbf{x}) - \nabla_x \log q(\mathbf{x}, t).$$

The continuity equation gives

$$\begin{aligned} \frac{\partial q(\mathbf{x}, t)}{\partial t} &= -\nabla_x \cdot (q \nabla_x \log p(\mathbf{x}|\mathcal{O}_m)p(o_{m+1}|\mathbf{x})) \\ &\quad + \Delta_x q(\mathbf{x}, t) \\ &:= g(p(\mathbf{x}|\mathcal{O}_m)p(o_{m+1}|\mathbf{x}), q(\mathbf{x}, t)) \end{aligned}$$

Hence, for any \mathbf{x} ,

$$q(\mathbf{x}, t) = q(\mathbf{x}, 0) + \int_0^t g(p(\mathbf{x}|\mathcal{O}_m)p(o_{m+1}|\mathbf{x}), q(\mathbf{x}, \tau)) d\tau.$$

The dynamcis g is a fixed function of $p(\mathbf{x}|\mathcal{O}_m)$, $p(o_{m+1}|\mathbf{x})$ and $q(\mathbf{x}, t)$, so the solution of this initial value problem (IVP) $q(\mathbf{x}, t)$ is a fixed function of $p(\mathbf{x}|\mathcal{O}_m)$, $p(o_{m+1}|\mathbf{x})$, $q(\mathbf{x}, 0)$ and t , which can be written as

$$q(\mathbf{x}, t) = \text{Solve-IVP}(p(\mathbf{x}|\mathcal{O}_m), p(o_{m+1}|\mathbf{x}), q(\mathbf{x}, 0), t).$$

Finally, we can write the optimal open-loop control as

$$\begin{aligned} w^*(q(\mathbf{x}, 0), t) \\ = \nabla_x \log(\text{Solve-IVP}(p(\mathbf{x}|\mathcal{O}_m), p(o_{m+1}|\mathbf{x}), q(\mathbf{x}, 0), t)). \end{aligned}$$

Hence, $w^*(q(\mathbf{x}, 0), t)$ has a fixed form across different m .

□

B. Adjoint Method

To explain it more clearly, let us denote the evolution of the n -th particles at the m -th stage by $\mathbf{x}_m^n(t)$ for $t \in [0, T]$. Note that $\mathbf{x}_m^n(T) = \mathbf{x}_{m+1}^n(0)$. (Then the notation \mathbf{x}_m^n in the main text will become $\mathbf{x}_m^n(T)$.)

Recall the loss for each task:

$$\mathcal{L}(T) = \frac{1}{MN} \sum_{m=1}^M \sum_{n=1}^N (\log q_m^n(\mathbf{x}_m^n(T), T) - \log p(\mathbf{x}_m^n(T), \mathcal{O}_m)).$$

The loss of one particle \mathbf{x}^n is

$$L^n := \frac{1}{M} \sum_{m=1}^M L_m^n,$$

where

$$L_m^n := -y_m^n(T) - \log p(\mathbf{x}_m^n(T), \mathcal{O}_m)$$

and $y_m^n(t) := -\log q_m^n(x_m^n(t), t)$.

First, an adjoint process is defined as

$$\mathbf{p}_m(t) := \frac{\partial L^n}{\partial [\mathbf{x}_m^n(t), y_m^n(t)]}.$$

Denote $f_m(\mathbf{x}(t), \theta) = f_\theta(\mathcal{X}_m, o_{m+1}, \mathbf{x}(t), t)$. During the m -th stage, the adjoint process follows the following differential equation

$$\frac{d\mathbf{p}_m}{dt} = -\frac{\partial}{\partial [\mathbf{x}_m^n(t), y_m^n(t)]} \left[\begin{array}{c} f_m(\mathbf{x}_m^n(t), \theta) \\ \nabla_{\mathbf{x}} \cdot f_m(\mathbf{x}_m^n(t), \theta) \end{array} \right]^\top \mathbf{p}_m(t). \quad (21)$$

Note that

$$\mathbf{p}_m(T) = \sum_{m' \geq m} \frac{1}{M} \frac{\partial L_{m'}^n}{\partial [\mathbf{x}_m^n(T), y_m^n(T)]}. \quad (22)$$

Claim: The gradient of the loss is the solution of a backward ODE. That is to say, $\frac{\partial L^n}{\partial \theta} = \mathbf{z}_1(0)$, if $\mathbf{z}_M(T) = \mathbf{0}$ and

$$\frac{d\mathbf{z}_m(t)}{dt} = -\left[\begin{array}{c} \frac{\partial f_m}{\partial \theta}(\mathbf{x}_m^n(t), \theta) \\ \frac{\partial}{\partial \theta} [\nabla_{\mathbf{x}} \cdot f_m(\mathbf{x}_m^n(t), \theta)] \end{array} \right]^\top \mathbf{p}_m(t), \quad (23)$$

and $\mathbf{z}_m(T) = \mathbf{z}_{m+1}(0)$, for $m = 0, \dots, M-1$.

Proof. First, we can compute $\frac{d}{dt} \frac{\partial L^n}{\partial \theta}$:

$$\begin{aligned} \frac{d}{dt} \frac{\partial L^n}{\partial \theta} &= \frac{\partial}{\partial \theta} \sum_{m=1}^M \left(\frac{\partial L^n}{\partial \mathbf{x}_m^n(t)} \frac{d\mathbf{x}_m^n(t)}{dt} + \frac{\partial L^n}{\partial y_m^n(t)} \frac{dy_m^n(t)}{dt} \right) \\ &= \frac{\partial}{\partial \theta} \sum_{m=1}^M \left[\mathbf{p}_m(t)^\top \left[\begin{array}{c} f_m(\mathbf{x}_m^n(t), \theta) \\ \nabla_{\mathbf{x}} \cdot f_m(\mathbf{x}_m^n(t), \theta) \end{array} \right] \right] \\ &= \sum_{m=1}^M \left[\begin{array}{c} \frac{\partial f_m}{\partial \theta}(\mathbf{x}_m^n(t), \theta) \\ \frac{\partial}{\partial \theta} [\nabla_{\mathbf{x}} \cdot f_m(\mathbf{x}_m^n(t), \theta)] \end{array} \right]^\top \mathbf{p}_m(t) \end{aligned}$$

Next, we have

$$0 - \frac{\partial L^n}{\partial \theta} = -\int_{t=0}^T \frac{d}{dt} \frac{\partial L^n}{\partial \theta} = \sum_{m=1}^M \int_{t=0}^T -\left[\begin{array}{c} \frac{\partial f_m}{\partial \theta}(\mathbf{x}_m^n(t), \theta) \\ \frac{\partial}{\partial \theta} [\nabla_{\mathbf{x}} \cdot f_m(\mathbf{x}_m^n(t), \theta)] \end{array} \right]^\top \mathbf{p}_m(t) = \mathbf{z}_M(T) - \mathbf{z}_1(0).$$

Hence, $\frac{\partial L^n}{\partial \theta} = \mathbf{z}_1(0)$ if $\mathbf{z}_M(T) = \mathbf{0}$. □

An algorithm for computing $\frac{\partial L}{\partial \theta}$ is summarized in Algorithm 2. A nice python package of realizing this algorithm is provided by [Chen et al. \(2018\)](#).

Algorithm 2 Adjoint Method of Computing the Gradient

Function $\text{Grad}(\theta, \mathcal{X}_0, p(o|\mathbf{x}), \mathcal{O}_M)$:

 Denote $f_\theta^m = f_\theta(\mathcal{X}_m, o_{m+1}, \mathbf{x}(t), t)$

▶ notation

 Set $y_0^n = -\log p(\mathbf{x}_0^n)$ for each $\mathbf{x}_0^n \in \mathcal{X}_0$
For all $n = 1$ to N **do**
For $m = 0$ to $M - 1$ **do**

$$\begin{bmatrix} \mathbf{x}_{m+1}^n \\ y_{m+1}^n \end{bmatrix} \leftarrow \begin{bmatrix} \mathbf{x}_m^n \\ y_m^n \end{bmatrix} + \int_0^T \begin{bmatrix} f_\theta^m \\ \nabla \cdot f_\theta^m \end{bmatrix} dt$$

 Set $\mathbf{p}_M^n(T) = \mathbf{0}$ and $\mathbf{z}_M^n(T) = \mathbf{0}$
For $m = M$ to 1 **do**

$$\mathbf{p}_m^n(T) \leftarrow \mathbf{p}_m^n(T) + \frac{1}{M} \frac{\partial L_m^n}{\partial [\mathbf{x}_m^n, y_m^n]}$$

 Solve ODEs in Eq. (4), Eq. (8), Eq. (21) and Eq. (23) for $\mathbf{x}_m^n(t)$, $\mathbf{p}_m^n(t)$ and $\mathbf{z}_m^n(t)$ backwardly from T to 0

 Set $\mathbf{x}_{m-1}^n(T) = \mathbf{x}_m^n(0)$, $\mathbf{p}_{m-1}^n(T) = \mathbf{p}_m^n(0)$ and $\mathbf{z}_{m-1}^n(T) = \mathbf{z}_m^n(0)$
return $\frac{1}{N} \sum_{n=1}^N \frac{\partial L^n}{\partial \theta} = \frac{1}{N} \sum_{n=1}^N \mathbf{z}_1^n(0)$

C. Experiment Details

C.1. Parameterization

Overall we parameterize the flow velocity as

$$f = \mathbf{h} \left(\frac{1}{N} \sum_{n=1}^N \phi(\mathbf{x}_m^n), o_{m+1}, \mathbf{x}(t), t \right),$$

 where both ϕ and \mathbf{h} are neural networks. For instance, let $\text{ctx} = [\frac{1}{N} \sum_{n=1}^N \phi(\mathbf{x}_m^n)^\top, o_{m+1}^\top]$ be the context of this conditional flow, where ϕ is a dense feed-forward neural network, a specific neural architecture we use in the experiment is

$$f = \text{Gated}_k(\dots [\text{ctx}, \text{Gated}_2([\text{ctx}, \text{Gated}_1([\text{ctx}, \mathbf{x}(t)^\top]^\top, t)]^\top, t)]^\top \dots, t), \quad (24)$$

$$\text{where } \text{Gated}_j(\mathbf{y}, t) = (W_j \mathbf{y} + \mathbf{b}_j) * \sigma(t \mathbf{v}_j + \mathbf{c}_j) + t \mathbf{c}_j, \quad (25)$$

 where $*$ is element-wise multiplication. The number of layers k can be tuned, but in general \mathbf{h} is a shallow network.

C.2. Evaluation Metric

MMD² The maximum mean discrepancy (MMD) of the true posterior p and the estimated posterior q is defined as

$$\text{MMD}[\mathcal{F}, p, q] := \sup_{f \in \mathcal{F}} (\mathbb{E}_{x \sim p}[f(x)] - \mathbb{E}_{y \sim q}[f(y)]).$$

 When \mathcal{F} is a unit ball in a characteristic RKHS, [Gretton et al. \(2012\)](#) showed that the squared MMD is

$$\text{MMD}^2[\mathcal{F}, p, q] = \mathbb{E}[k(x, x')] - 2\mathbb{E}[k(x, y)] + \mathbb{E}[k(y, y')],$$

 where $x, x' \sim p$ and $y, y' \sim q$.

Cross-entropy Evaluating the KL divergence is equivalent to evaluating the cross-entropy.

$$\mathbb{E}_{x \sim p} - \log q(x) \approx \frac{1}{n} \sum_{n=1}^N (-\log q(x^n)), \quad (26)$$

 where $q(x)$ is approximated by kernel density estimation on the set of particles obtained from different sampling methods.

Integral Evaluation When the true posterior is a Gaussian distribution $\mathcal{N}(\mu, \Sigma)$, the expectation of the following test functions have closed-form expressions.

- $\mathbb{E}[\mathbf{x}] = \mu$
- $\mathbb{E}[\mathbf{x}^\top A \mathbf{x}] = \text{tr}(A \Sigma) + \mu^\top A \mu$
- $\mathbb{E}[(A \mathbf{x} + \mathbf{a})^\top (B \mathbf{x} + \mathbf{b})] = \text{tr}(A \Sigma B^\top) + (A \mu + \mathbf{a})^\top (B \mu + \mathbf{b})$

D. More Experimental Results

D.1. Multivariate Gaussian Model




Figure 7: Experimental results on 2 dimensional multivariate Gaussian model.

D.2. LDS Model

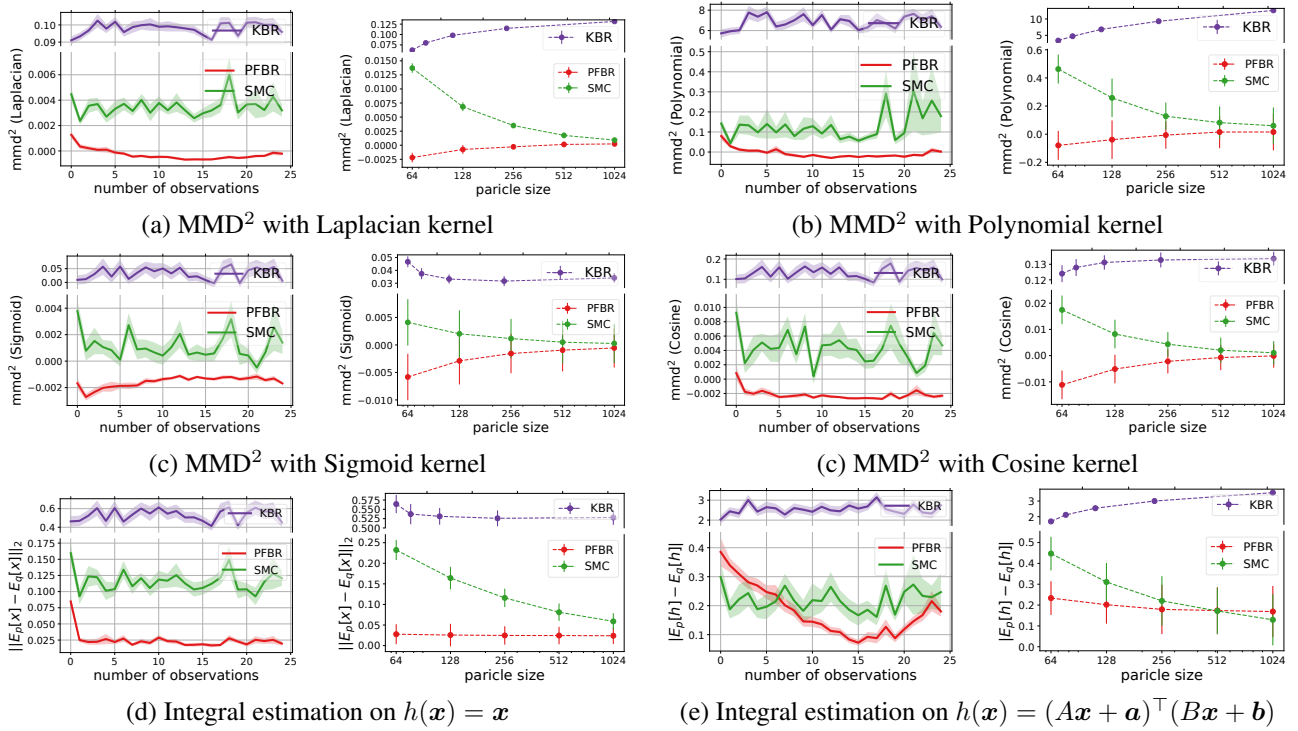


Figure 8: Experimental results on LDS model.