Controlled Monte Carlo Diffusions

Nikolas Nüsken

King's College London

June 12, 2025

Based on joint work with **F. Vargas** (Cambridge), **S. Padhy** (Cambridge), **D. Bessel** (Karlsruhe)

Goal: Sample from $\pi = \frac{1}{Z} e^{-V} \, \mathrm{d}x$, where

- $V : \mathbb{R}^d \to \mathbb{R}$ is a potential,
- $Z = \int_{\mathbb{R}^d} e^{-V(x)} \, \mathrm{d}x$ is the normalising constant.

Goal: Sample from $\pi = \frac{1}{Z} e^{-V} \, \mathrm{d}x$, where

- $V : \mathbb{R}^d \to \mathbb{R}$ is a potential,
- $Z = \int_{\mathbb{R}^d} e^{-V(x)} \, \mathrm{d}x$ is the normalising constant.

Standard approach: Overdamped Langevin dynamics

$$\mathrm{d}\boldsymbol{Y}_t = -\nabla V(\boldsymbol{Y}_t) \, \mathrm{d}t + \sqrt{2} \, \mathrm{d}\boldsymbol{W}_t,$$

Goal: Sample from $\pi = \frac{1}{Z} e^{-V} \, \mathrm{d}x$, where

- $V : \mathbb{R}^d \to \mathbb{R}$ is a potential,
- $Z = \int_{\mathbb{R}^d} e^{-V(x)} \, \mathrm{d}x$ is the normalising constant.

Standard approach: Overdamped Langevin dynamics

$$\mathrm{d}\boldsymbol{Y}_t = -\nabla V(\boldsymbol{Y}_t)\,\mathrm{d}t + \sqrt{2}\,\mathrm{d}\boldsymbol{W}_t,$$

Key observation:

$\mathrm{Law}\, \boldsymbol{Y}_t \to \pi$ as $t \to \infty$.

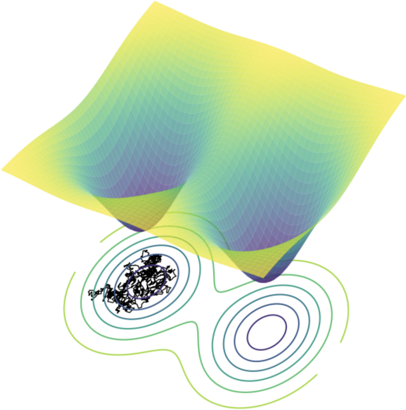

Controlled dynamics

Annealing/tempering: Fix a curve $(\pi_t)_{t\in[0,1]}$ with $\pi_1 = \pi$, e.g.

$$\pi_t = \frac{e^{-th}\pi_0}{Z_t}, \qquad t \in [0,1]$$

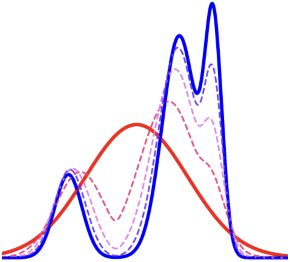

prior $t = 0$ posterior $t = 1$

Controlled dynamics

Annealing/tempering: Fix a curve $(\pi_t)_{t\in[0,1]}$ with $\pi_1 = \pi$, e.g.

$$\pi_t = \frac{e^{-th}\pi_0}{Z_t}, \qquad t \in [0,1]$$

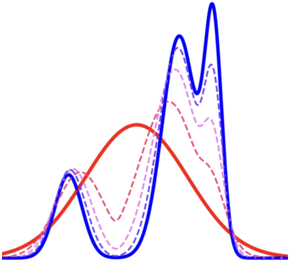

prior $t = 0$ posterior $t = 1$

Idea: With $\nabla V_t := -\nabla \ln \pi_t$, find a control $\nabla\phi_t$ such that

$$\mathrm{d}\boldsymbol{Y}_t = -\nabla V_t(\boldsymbol{Y}_t)\,\mathrm{d}t + \nabla\phi_t(\boldsymbol{Y}_t)\,\mathrm{d}t + \sqrt{2}\,\mathrm{d}\boldsymbol{W}_t, \qquad \boldsymbol{Y}_0 \sim \pi_0$$

reproduces $(\pi_t)_{t\in[0,1]}$, i.e. $\mathrm{Law}(\boldsymbol{Y}_t) = \pi_t$.

Controlled dynamics

Annealing/tempering: Fix a curve $(\pi_t)_{t\in[0,1]}$ with $\pi_1 = \pi$, e.g.

$$\pi_t = \frac{e^{-th}\pi_0}{Z_t}, \qquad t \in [0,1]$$

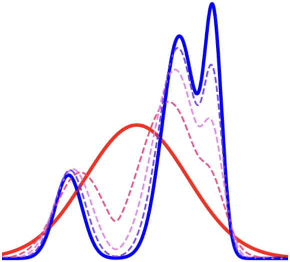

prior $t = 0$ posterior $t = 1$

Idea: With $\nabla V_t := -\nabla \ln \pi_t$, find a control $\nabla \phi_t$ such that

$$\mathrm{d}\boldsymbol{Y}_t = -\nabla V_t(\boldsymbol{Y}_t)\,\mathrm{d}t + \nabla\phi_t(\boldsymbol{Y}_t)\,\mathrm{d}t + \sqrt{2}\,\mathrm{d}\boldsymbol{W}_t, \qquad \boldsymbol{Y}_0 \sim \pi_0$$

reproduces $(\pi_t)_{t\in[0,1]}$, i.e. $\mathrm{Law}(\boldsymbol{Y}_t) = \pi_t$.

- $\nabla\phi_t$ enables transitions between local equilibria.

Controlled Monte Carlo Diffusions

$$\pi_t = \frac{e^{-V_t}}{Z_t} \propto \widehat{\pi}_t, \qquad t \in [0, 1].$$

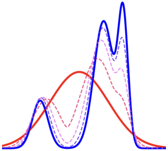

Controlled Monte Carlo Diffusions

$$\pi_t = \frac{e^{-V_t}}{Z_t} \propto \widehat{\pi}_t, \qquad t \in [0,1].$$

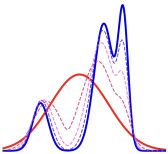

Consider

$$\mathcal{L}(\phi) = \mathbb{E}\left[\int_0^1 |\nabla V_t(\boldsymbol{Y}_t)|^2 \mathrm{d}t - \frac{1}{\sqrt{2}} \int_0^1 (-\nabla V_t + \nabla \phi_t)(\boldsymbol{Y}_t) \cdot \overleftarrow{\mathrm{d}} \boldsymbol{W}_t - \ln \widehat{\pi}_1(\boldsymbol{Y}_t)\right],$$

where the expectation is with respect to

$$\mathrm{d}\boldsymbol{Y}_t = -\nabla V_t(\boldsymbol{Y}_t)\,\mathrm{d}t + \nabla \phi_t(\boldsymbol{Y}_t)\,\mathrm{d}t + \sqrt{2}\,\mathrm{d}\boldsymbol{W}_t, \qquad \boldsymbol{Y}_0 \sim \pi_0. \quad (1)$$

Theorem:

The dynamics (1) with the (unique) minimiser $\phi = \phi^*$ satisfies $\mathrm{Law}\,\boldsymbol{Y}_t = \pi_t$, for all $t \in [0,1]$.

Goals of this talk

1. Explain the construction of the CMCD sampler.

2. More broadly, lay out a general framework for constructing controlled diffusions.

Match forward and backward path measures

$$\mathcal{L}(\phi) = D_{\mathrm{KL}}\left(\overrightarrow{\mathbb{P}}^{\pi_0,a}, \overleftarrow{\mathbb{P}}^{\pi_1,b}\right) + \mathrm{const.},$$

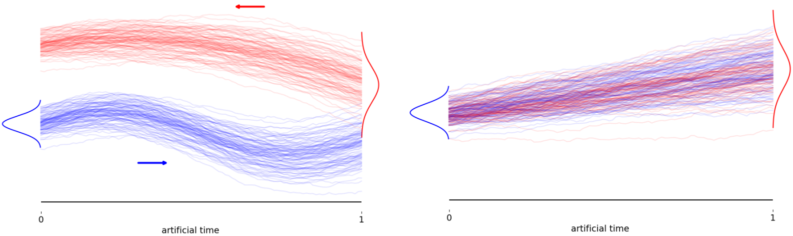

$$\mathrm{d}\boldsymbol{Y}_t = a_t(\boldsymbol{Y}_t)\,\mathrm{d}t + \overrightarrow{\mathrm{d}}\,\boldsymbol{W}_t, \qquad \boldsymbol{Y}_0 \sim \pi_0 \tag{2a}$$

$$\mathrm{d}\boldsymbol{Y}_t = b_t(\boldsymbol{Y}_t)\,\mathrm{d}t + \overleftarrow{\mathrm{d}}\,\boldsymbol{W}_t, \qquad \boldsymbol{Y}_1 \sim \pi_1 \tag{2b}$$

Match forward and backward path measures

$$\mathcal{L}(\phi) = D_{\mathrm{KL}}\left(\overrightarrow{\mathbb{P}}^{\pi_0, a}, \overleftarrow{\mathbb{P}}^{\pi_1, b}\right) + \mathrm{const.},$$

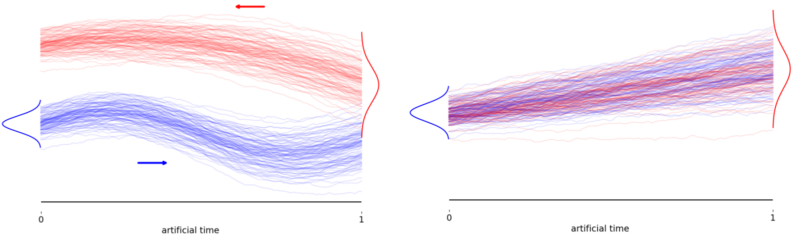

$$\mathrm{d}\boldsymbol{Y}_t = a_t(\boldsymbol{Y}_t)\,\mathrm{d}t + \overrightarrow{\mathrm{d}}\,\boldsymbol{W}_t, \qquad \boldsymbol{Y}_0 \sim \pi_0 \qquad (2a)$$

$$\mathrm{d}\boldsymbol{Y}_t = b_t(\boldsymbol{Y}_t)\,\mathrm{d}t + \overleftarrow{\mathrm{d}}\,\boldsymbol{W}_t, \qquad \boldsymbol{Y}_1 \sim \pi_1 \qquad (2b)$$

- If $D_{\mathrm{KL}}\left(\overrightarrow{\mathbb{P}}^{\pi_0, a}, \overleftarrow{\mathbb{P}}^{\pi_1, b}\right) = 0$, then π_0 is transported to π_1 by (2a), and π_1 is transported to π_0 by (2b).

Match forward and backward path measures

$$\mathcal{L}(\phi) = D_{\mathrm{KL}}\left(\overrightarrow{\mathbb{P}}^{\pi_0, a}, \overleftarrow{\mathbb{P}}^{\pi_1, b}\right) + \text{const.},$$

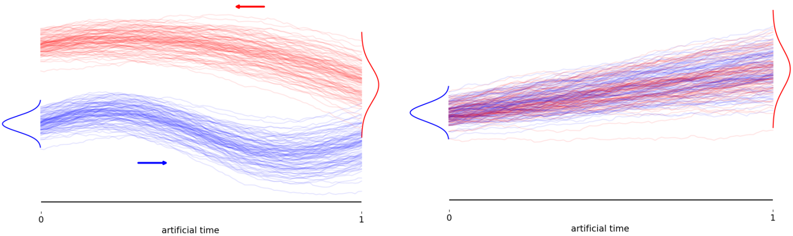

$$\mathrm{d}\boldsymbol{Y}_t = a_t(\boldsymbol{Y}_t)\,\mathrm{d}t + \overrightarrow{\mathrm{d}}\,\boldsymbol{W}_t, \qquad \boldsymbol{Y}_0 \sim \pi_0 \tag{2a}$$

$$\mathrm{d}\boldsymbol{Y}_t = b_t(\boldsymbol{Y}_t)\,\mathrm{d}t + \overleftarrow{\mathrm{d}}\,\boldsymbol{W}_t, \qquad \boldsymbol{Y}_1 \sim \pi_1 \tag{2b}$$

- If $D_{\mathrm{KL}}\left(\overrightarrow{\mathbb{P}}^{\pi_0, a}, \overleftarrow{\mathbb{P}}^{\pi_1, b}\right) = 0$, then π_0 is transported to π_1 by (2a), and π_1 is transported to π_0 by (2b).
- connected to the ELBO in variational inference.

Minimisers of $D_{\mathrm{KL}}\left(\overrightarrow{\mathbb{P}}^{\pi_0,a}, \overleftarrow{\mathbb{P}}^{\pi_1,b}\right)$ are not unique:

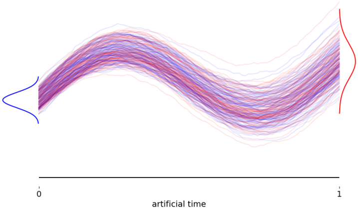

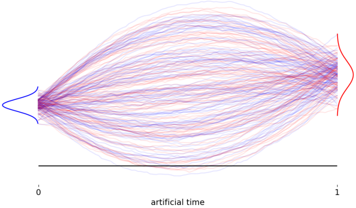

Minimisers of $D_{\mathrm{KL}}\left(\overrightarrow{\mathbb{P}}^{\pi_0,a}, \overleftarrow{\mathbb{P}}^{\pi_1,b}\right)$ are not unique:

$$= \mathbb{E}_{\boldsymbol{Y}\sim\overrightarrow{\mathbb{P}}^{\pi_0,a}}\left[\ln\left(\frac{\mathrm{d}\overrightarrow{\mathbb{P}}^{\pi_0,a}}{\mathrm{d}\overleftarrow{\mathbb{P}}^{\pi_1,b}}\right)(\boldsymbol{Y})\right]$$

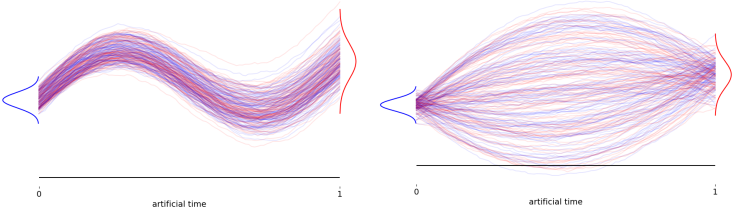

We could

- fix the forward process $\rightsquigarrow$ *diffusion models*
- fix the backward process $\rightsquigarrow$ *optimal control*
- add an additional term $\rightsquigarrow$ *Schrödinger bridges*
- parameterise a and b in a restrictive way $\rightsquigarrow$ *CMCD*

Minimisers of $D_{\mathrm{KL}}\left(\overrightarrow{\mathbb{P}}^{\pi_0,a}, \overleftarrow{\mathbb{P}}^{\pi_1,b}\right)$ are not unique:

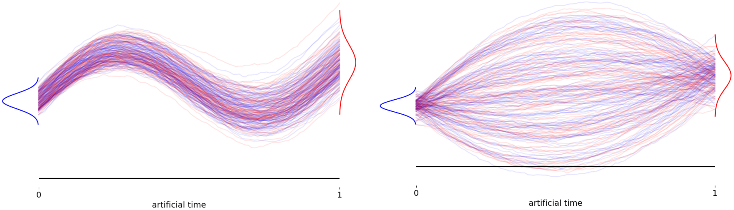

Minimisers of $D_{\mathrm{KL}}\left(\overrightarrow{\mathbb{P}}^{\pi_0,a}, \overleftarrow{\mathbb{P}}^{\pi_1,b}\right)$ are not unique:

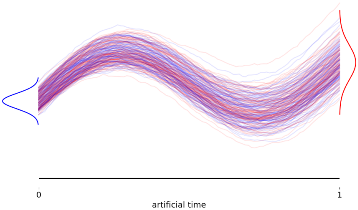

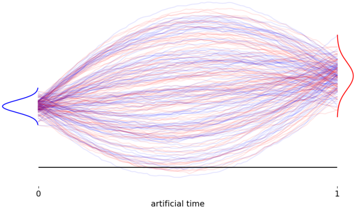

Theorem (Nelson's relation)

Denote the time-marginals by $\overrightarrow{\mathbb{P}}_t^{\pi_0,a} =: \pi_t$. *If* $\overrightarrow{\mathbb{P}}^{\pi_0,a} = \overleftarrow{\mathbb{P}}^{\pi_1,b}$, *then*

$$b_t = a_t - \underbrace{\nabla \ln \pi_t}_{\text{score}}, \qquad \text{for all } t \in (0,1].$$

Minimisers of $D_{\mathrm{KL}}\left(\overrightarrow{\mathbb{P}}^{\pi_0,a}, \overleftarrow{\mathbb{P}}^{\pi_1,b}\right)$ are not unique:

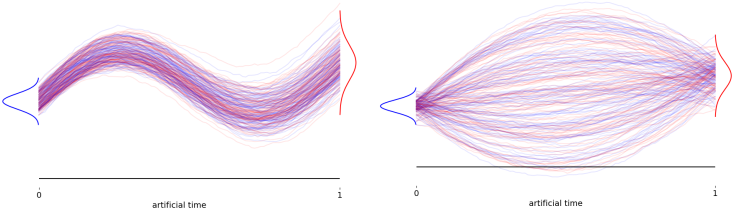

Theorem (Nelson's relation)

Denote the time-marginals by $\overrightarrow{\mathbb{P}}^{\pi_0,a}_t =: \pi_t$. *If* $\overrightarrow{\mathbb{P}}^{\pi_0,a} = \overleftarrow{\mathbb{P}}^{\pi_1,b}$, *then*

$$b_t = a_t - \underbrace{\nabla \ln \pi_t}_{\text{score}}, \qquad \textit{for all } t \in (0,1].$$

$$\implies \qquad \mathcal{L}(\phi) = D_{\mathrm{KL}}\left(\overrightarrow{\mathbb{P}}^{\pi_0,-\frac{1}{2}\nabla V+\nabla\phi}, \overleftarrow{\mathbb{P}}^{\pi_1,\frac{1}{2}\nabla V+\nabla\phi}\right) + \text{const.}$$

fixes all time marginals.

Calculational tool: Girsanov theorem

To compute (e.g.) $D_{\mathrm{KL}}(\overrightarrow{\mathbb{P}}^{\pi_0,a} | \overleftarrow{\mathbb{P}}^{\pi_1,b}) = \mathbb{E}_{\boldsymbol{Y} \sim \overrightarrow{\mathbb{P}}^{\pi_0,a}} \left[\ln \left(\frac{\mathrm{d}\overrightarrow{\mathbb{P}}^{\pi_0,a}}{\mathrm{d}\overleftarrow{\mathbb{P}}^{\pi_1,b}} \right) (\boldsymbol{Y}) \right]$:

Calculational tool: Girsanov theorem

To compute (e.g.) $D_{\mathrm{KL}}(\overrightarrow{\mathbb{P}}^{\pi_0,a} | \overleftarrow{\mathbb{P}}^{\pi_1,b}) = \mathbb{E}_{\boldsymbol{Y} \sim \overrightarrow{\mathbb{P}}^{\pi_0,a}} \left[\ln \left(\frac{\mathrm{d}\overrightarrow{\mathbb{P}}^{\pi_0,a}}{\mathrm{d}\overleftarrow{\mathbb{P}}^{\pi_1,b}} \right) (\boldsymbol{Y}) \right]$:

Theorem (Girsanov)

$$\ln \left(\frac{\mathrm{d}\overrightarrow{\mathbb{P}}^{\mu,a}}{\mathrm{d}\overleftarrow{\mathbb{P}}^{\nu,b}} \right) (\boldsymbol{Y}) = \ln \left(\frac{\mathrm{d}\overrightarrow{\mathbb{P}}^{\mu,a}}{\mathrm{d}\overrightarrow{\mathbb{P}}^{\Gamma_0,\gamma^+}} \right) (\boldsymbol{Y}) + \ln \left(\frac{\mathrm{d}\overleftarrow{\mathbb{P}}^{\Gamma_T,\gamma^-}}{\mathrm{d}\overleftarrow{\mathbb{P}}^{\nu,b}} \right) (\boldsymbol{Y})$$

where $\overrightarrow{\mathbb{P}}^{\Gamma_0,\gamma^+} = \overleftarrow{\mathbb{P}}^{\Gamma_T,\gamma^-}$ *is a reference process.*

Calculational tool: Girsanov theorem

To compute (e.g.) $D_{\mathrm{KL}}(\overrightarrow{\mathbb{P}}^{\pi_0,a} | \overleftarrow{\mathbb{P}}^{\pi_1,b}) = \mathbb{E}_{\boldsymbol{Y} \sim \overrightarrow{\mathbb{P}}^{\pi_0,a}} \left[\ln \left(\frac{\mathrm{d}\overrightarrow{\mathbb{P}}^{\pi_0,a}}{\mathrm{d}\overleftarrow{\mathbb{P}}^{\pi_1,b}} \right) (\boldsymbol{Y}) \right]$:

Theorem (Girsanov)

$$\ln\left(\frac{\mathrm{d}\overrightarrow{\mathbb{P}}^{\mu,a}}{\mathrm{d}\overleftarrow{\mathbb{P}}^{\nu,b}}\right)(\boldsymbol{Y}) = \ln\left(\frac{\mathrm{d}\overrightarrow{\mathbb{P}}^{\mu,a}}{\mathrm{d}\overrightarrow{\mathbb{P}}^{\Gamma_0,\gamma^+}}\right)(\boldsymbol{Y}) + \ln\left(\frac{\mathrm{d}\overleftarrow{\mathbb{P}}^{\Gamma_T,\gamma^-}}{\mathrm{d}\overleftarrow{\mathbb{P}}^{\nu,b}}\right)(\boldsymbol{Y})$$
$$= \ln\left(\frac{\mathrm{d}\mu}{\mathrm{d}\Gamma_0}\right)(\boldsymbol{Y}_0) - \ln\left(\frac{\mathrm{d}\nu}{\mathrm{d}\Gamma_T}\right)(\boldsymbol{Y}_T)$$
$$+ \int_0^T (a_t - \gamma_t^+)(\boldsymbol{Y}_t) \cdot \left(\overrightarrow{\mathrm{d}}\,\boldsymbol{Y}_t - \tfrac{1}{2}(a_t + \gamma_t^+)(\boldsymbol{Y}_t)\,\mathrm{d}t\right)$$
$$- \int_0^T (b_t - \gamma_t^-)(\boldsymbol{Y}_t) \cdot \left(\overleftarrow{\mathrm{d}}\,\boldsymbol{Y}_t - \tfrac{1}{2}(b_t + \gamma_t^-)(\boldsymbol{Y}_t)\,\mathrm{d}t\right),$$

where $\overrightarrow{\mathbb{P}}^{\Gamma_0,\gamma^+} = \overleftarrow{\mathbb{P}}^{\Gamma_T,\gamma^-}$ *is a reference process.*

Calculational tool: Girsanov theorem

To compute (e.g.) $D_{\mathrm{KL}}(\overrightarrow{\mathbb{P}}^{\pi_0,a} | \overleftarrow{\mathbb{P}}^{\pi_1,b}) = \mathbb{E}_{\boldsymbol{Y} \sim \overrightarrow{\mathbb{P}}^{\pi_0,a}} \left[\ln \left(\frac{\mathrm{d}\overrightarrow{\mathbb{P}}^{\pi_0,a}}{\mathrm{d}\overleftarrow{\mathbb{P}}^{\pi_1,b}} \right) (\boldsymbol{Y}) \right]$:

Theorem (Girsanov)

$$\ln\left(\frac{\mathrm{d}\overrightarrow{\mathbb{P}}^{\mu,a}}{\mathrm{d}\overleftarrow{\mathbb{P}}^{\nu,b}}\right)(\boldsymbol{Y}) = \ln\left(\frac{\mathrm{d}\overrightarrow{\mathbb{P}}^{\mu,a}}{\mathrm{d}\overrightarrow{\mathbb{P}}^{\Gamma_0,\gamma^+}}\right)(\boldsymbol{Y}) + \ln\left(\frac{\mathrm{d}\overleftarrow{\mathbb{P}}^{\Gamma_T,\gamma^-}}{\mathrm{d}\overleftarrow{\mathbb{P}}^{\nu,b}}\right)(\boldsymbol{Y})$$

$$= \ln\left(\frac{\mathrm{d}\mu}{\mathrm{d}\Gamma_0}\right)(\boldsymbol{Y}_0) - \ln\left(\frac{\mathrm{d}\nu}{\mathrm{d}\Gamma_T}\right)(\boldsymbol{Y}_T)$$

$$+ \int_0^T (a_t - \gamma_t^+)(\boldsymbol{Y}_t) \cdot \left(\overrightarrow{\mathrm{d}}\, \boldsymbol{Y}_t - \tfrac{1}{2}(a_t + \gamma_t^+)(\boldsymbol{Y}_t)\,\mathrm{d}t \right)$$

$$- \int_0^T (b_t - \gamma_t^-)(\boldsymbol{Y}_t) \cdot \left(\overleftarrow{\mathrm{d}}\, \boldsymbol{Y}_t - \tfrac{1}{2}(b_t + \gamma_t^-)(\boldsymbol{Y}_t)\,\mathrm{d}t \right),$$

where $\overrightarrow{\mathbb{P}}^{\Gamma_0,\gamma^+} = \overleftarrow{\mathbb{P}}^{\Gamma_T,\gamma^-}$ *is a reference process.*

- Flexibility in $\overrightarrow{\mathbb{P}}^{\Gamma_0,\gamma^+} = \overleftarrow{\mathbb{P}}^{\Gamma_T,\gamma^-}$ allows us to cancel intractable terms.

Up to an additive constant, $\mathcal{L}_{D_{\mathrm{KL}}}^{\mathrm{CMCD}}(\phi)$ is given by

$$\mathbb{E}\left[\int_0^1 |\nabla \ln \pi_t(\boldsymbol{Y}_t)|^2 \mathrm{d}t + \frac{1}{\sqrt{2}} \int_0^1 (\nabla \ln \pi_t - \nabla \phi_t)(\boldsymbol{Y}_t) \cdot \overleftarrow{\mathrm{d}\boldsymbol{W}_t} - \ln \hat{\pi}_1(\boldsymbol{Y}_t)\right],$$

$$\approx \mathbb{E}\left[\ln \frac{\pi_0(\boldsymbol{Y}_0)}{\hat{\pi}(\boldsymbol{Y}_T)} \prod_{k=0}^{K-1} \frac{\mathcal{N}(\boldsymbol{Y}_{t_{k+1}} | \boldsymbol{Y}_{t_k} + (\nabla \ln \pi_{t_k} + \nabla \ln \phi_{t_k})(\boldsymbol{Y}_{t_k})\Delta t_k, 2\sigma^2 \Delta t_k)}{\mathcal{N}(\boldsymbol{Y}_{t_k} | \boldsymbol{Y}_{t_{k+1}} + (\nabla \ln \pi_{t_{k+1}} - \nabla \ln \phi_{t_{k+1}})(\boldsymbol{Y}_{t_{k+1}})\Delta t_k, 2\sigma^2 \Delta t_k)}\right]$$

where the expectation is with respect to

$$\mathrm{d}\boldsymbol{Y}_t = -\nabla V_t(\boldsymbol{Y}_t)\,\mathrm{d}t + \nabla \phi_t(\boldsymbol{Y}_t)\,\mathrm{d}t + \sqrt{2}\,\mathrm{d}\boldsymbol{W}_t, \qquad \boldsymbol{Y}_0 \sim \pi_0.$$

Up to an additive constant, $\mathcal{L}^{\mathrm{CMCD}}_{D_{\mathrm{KL}}}(\phi)$ is given by

$$\mathbb{E}\left[\int_0^1 |\nabla \ln \pi_t(\boldsymbol{Y}_t)|^2 \mathrm{d}t + \frac{1}{\sqrt{2}} \int_0^1 (\nabla \ln \pi_t - \nabla \phi_t)(\boldsymbol{Y}_t) \cdot \overleftarrow{\mathrm{d}\boldsymbol{W}_t} - \ln \hat{\pi}_1(\boldsymbol{Y}_t)\right],$$

$$\approx \mathbb{E}\left[\ln \frac{\pi_0(\boldsymbol{Y}_0)}{\hat{\pi}(\boldsymbol{Y}_T)} \prod_{k=0}^{K-1} \frac{\mathcal{N}(\boldsymbol{Y}_{t_{k+1}} | \boldsymbol{Y}_{t_k} + (\nabla \ln \pi_{t_k} + \nabla \ln \phi_{t_k})(\boldsymbol{Y}_{t_k})\Delta t_k, 2\sigma^2 \Delta t_k)}{\mathcal{N}(\boldsymbol{Y}_{t_k} | \boldsymbol{Y}_{t_{k+1}} + (\nabla \ln \pi_{t_{k+1}} - \nabla \ln \phi_{t_{k+1}})(\boldsymbol{Y}_{t_{k+1}})\Delta t_k, 2\sigma^2 \Delta t_k)}\right]$$

where the expectation is with respect to

$$\mathrm{d}\boldsymbol{Y}_t = -\nabla V_t(\boldsymbol{Y}_t)\,\mathrm{d}t + \nabla \phi_t(\boldsymbol{Y}_t)\,\mathrm{d}t + \sqrt{2}\,\mathrm{d}\boldsymbol{W}_t, \qquad \boldsymbol{Y}_0 \sim \pi_0.$$

Normalising constants:

$$Z = \underbrace{\mathbb{E}\left[\frac{\mathrm{d}\overleftarrow{\mathbb{P}}^{\hat{\pi}_T, -\sigma^2 \nabla \ln \pi + \nabla \phi}}{\mathrm{d}\overrightarrow{\mathbb{P}}^{\pi_0, \sigma^2 \nabla \ln \pi + \nabla \phi}}\right]}_{\text{holds for any } \phi} = \underbrace{\frac{\mathrm{d}\overleftarrow{\mathbb{P}}^{\hat{\pi}_T, -\sigma^2 \nabla \ln \pi + \nabla \phi^*}}{\mathrm{d}\overrightarrow{\mathbb{P}}^{\pi_0, \sigma^2 \nabla \ln \pi + \nabla \phi^*}}(\boldsymbol{Y})}_{\text{zero variance estimator (holds for any } \boldsymbol{Y})},$$

Other stuff: Score-based generative modelling

Set $\pi_0 = \mu_{\text{data}}$, fix an **ergodic drift** a_t so that $\overrightarrow{\mathbb{P}}^{\mu,a}_T \approx \pi_T$ for $T \to \infty$,

$$\mathcal{L}(s) := D_{\text{KL}}(\overbrace{\overrightarrow{\mathbb{P}}^{\mu,a}}^{\textbf{FIX!}} || \overleftarrow{\mathbb{P}}^{\nu,a-s})$$

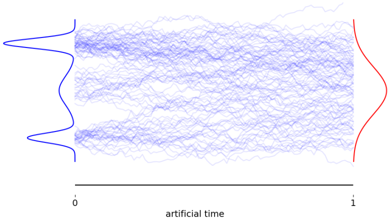

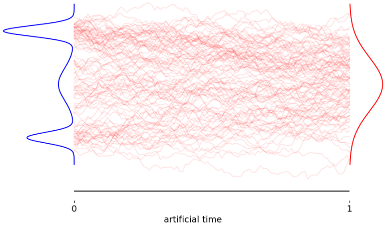

Other stuff: Score-based generative modelling

Set $\pi_0 = \mu_{\mathrm{data}}$, fix an **ergodic drift** a_t so that $\overrightarrow{\mathbb{P}}^{\mu,a}_T \approx \pi_T$ for $T \to \infty$,

$$\mathcal{L}(s) := D_{\mathrm{KL}}(\overbrace{\overrightarrow{\mathbb{P}}^{\mu,a}}^{\textbf{FIX!}} \,||\, \overleftarrow{\mathbb{P}}^{\nu,a-s})$$

$$= \mathbb{E}_{\boldsymbol{Y} \sim \overrightarrow{\mathbb{P}}^{\mu,a}} \left[\int_0^T \underbrace{\left(\tfrac{1}{2} s_t^2(\boldsymbol{Y}_t) + (\nabla \cdot s_t)(\boldsymbol{Y}_t)\right)}_{\text{score-matching objective}} \mathrm{d}t \right] + \mathrm{const.}$$

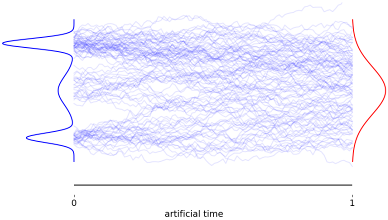

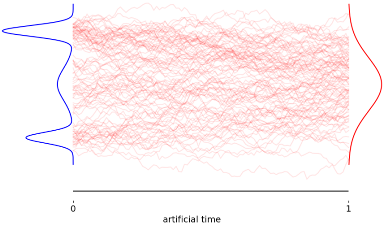

Other stuff: optimal control and guidance

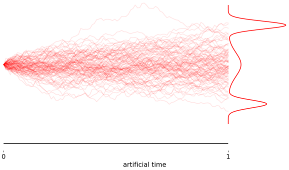

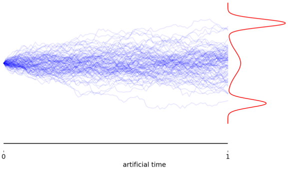

Singular backward drift $b_t(x) = x/t$ enforces $\pi_0 = \delta_0$.

$$\underbrace{D_{\mathrm{KL}}(\overrightarrow{\mathbb{P}}^{\delta_0, a} | \overbrace{\overleftarrow{\mathbb{P}}^{\pi_1, b}}^{\textbf{FIX!}})}_{\mathcal{L}(a)} =$$

Other stuff: optimal control and guidance

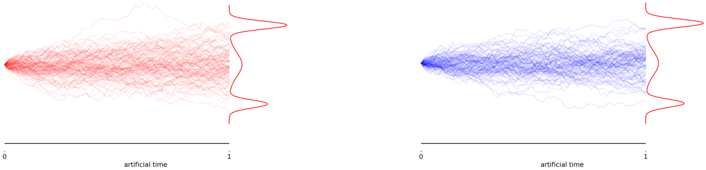

Singular backward drift $\color{red}{b_t}(x) = x/t$ enforces $\pi_0 = \delta_0$.

$$\underbrace{D_{\mathrm{KL}}\big(\overrightarrow{\mathbb{P}}^{\delta_0,\color{blue}{a}} \,|\, \overbrace{\overleftarrow{\mathbb{P}}^{\pi_1,b}}^{\textbf{FIX!}}\big)}_{\mathcal{L}(\color{blue}{a})} = \mathbb{E}_{\boldsymbol{Y}\sim\overrightarrow{\mathbb{P}}^{\pi_0,a}}\left[\underbrace{\int_0^1 a^2(\boldsymbol{Y}_t)\,\mathrm{d}t}_{\text{running cost}} + \underbrace{\ln\left(\frac{\mathrm{d}\mathcal{N}(0,1\sigma^2)}{\mathrm{d}\pi_1}\right)}_{\text{terminal cost}}(\boldsymbol{Y}_1)\right] + \mathrm{c.},$$

Other stuff: optimal control and guidance

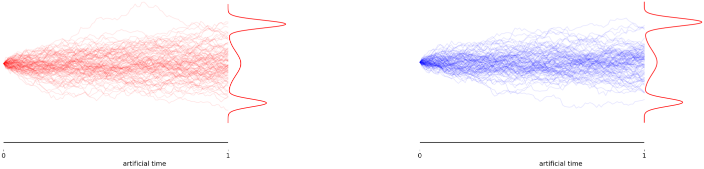

Singular backward drift $\color{red}{b_t}(x) = x/t$ enforces $\pi_0 = \delta_0$.

$$\underbrace{D_{\mathrm{KL}}(\overrightarrow{\mathbb{P}}^{\delta_0,\color{blue}{a}} | \overbrace{\overleftarrow{\mathbb{P}}^{\pi_1,b}}^{\textbf{FIX!}})}_{\mathcal{L}(\color{blue}{a})} = \mathbb{E}_{\boldsymbol{Y}\sim\overrightarrow{\mathbb{P}}^{\pi_0,a}}\left[\underbrace{\int_0^1 a^2(\boldsymbol{Y}_t)\,\mathrm{d}t}_{\text{running cost}} + \underbrace{\ln\left(\frac{\mathrm{d}\mathcal{N}(0,1\sigma^2)}{\mathrm{d}\pi_1}\right)}_{\text{terminal cost}}(\boldsymbol{Y}_1)\right] + \mathrm{c.},$$

Similar ideas: denoising diffusion samplers (Vargas et al, 2023), action matching (Neklykudov, 2023), iterative proportional fitting (Vargas et al, 2021), other divergences (Richter and N., 2021).

Schrödinger bridges

$$\mathbb{P}^* \in \underset{\overrightarrow{\mathbb{P}}^{\mu,a}_T = \nu}{\arg\min}\, \mathbb{E}_{\boldsymbol{Y} \sim \overrightarrow{\mathbb{P}}^{\mu,a}} \left[\frac{1}{2} \int_0^T \|a_t - f_t\|^2(\boldsymbol{Y}_t)\, \mathrm{d}t \right].$$

The drift a_t should be determined such that

- the diffusion reaches ν at time T.
- a_t remains close to the prescribed (physically motivated) f_t.

trajectories

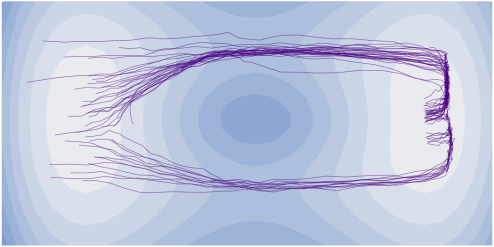

Schrödinger bridges

$$\mathbb{P}^* \in \underset{\overrightarrow{\mathbb{P}}_T^{\mu,a} = \nu}{\arg\min}\ \mathbb{E}_{\boldsymbol{Y} \sim \overrightarrow{\mathbb{P}}^{\mu,a}} \left[\frac{1}{2} \int_0^T \|a_t - f_t\|^2(\boldsymbol{Y}_t)\,\mathrm{d}t \right]. \tag{4}$$

Theorem (Mean-field game formulation)

Assume that ϕ satisfies the following conditions:

1. $\mathrm{d}\boldsymbol{Y}_t = f_t(\boldsymbol{Y}_t)\,\mathrm{d}t + \nabla\phi_t(\boldsymbol{Y}_t)\,\mathrm{d}t + \overrightarrow{\mathrm{d}}\boldsymbol{W}_t, \qquad \boldsymbol{Y}_0 \sim \mu$

   satisfies the terminal constraint $\boldsymbol{Y}_T \sim \nu$.
2. *HJB-equation:* $\partial_t\phi + f \cdot \nabla\phi + \frac{1}{2}\Delta\phi + \frac{1}{2}|\nabla\phi|^2 = 0$.

Then $a_t = \nabla\phi_t$ provides the unique solution to (4).

Schrödinger bridges

$$\mathbb{P}^* \in \underset{\overrightarrow{\mathbb{P}}_T^{\mu,a} = \nu}{\arg\min} \, \mathbb{E}_{\boldsymbol{Y} \sim \overrightarrow{\mathbb{P}}^{\mu,a}} \left[\frac{1}{2} \int_0^T \|a_t - f_t\|^2(\boldsymbol{Y}_t) \, \mathrm{d}t \right]. \qquad (4)$$

Theorem (Mean-field game formulation)

Assume that ϕ satisfies the following conditions:

1. $\mathrm{d}\boldsymbol{Y}_t = f_t(\boldsymbol{Y}_t)\,\mathrm{d}t + \nabla\phi_t(\boldsymbol{Y}_t)\,\mathrm{d}t + \overrightarrow{\mathrm{d}}\boldsymbol{W}_t, \qquad \boldsymbol{Y}_0 \sim \mu$

   satisfies the terminal constraint $\boldsymbol{Y}_T \sim \nu$.
2. *HJB-equation:* $\partial_t \phi + f \cdot \nabla\phi + \frac{1}{2}\Delta\phi + \frac{1}{2}|\nabla\phi|^2 = 0.$

Then $a_t = \nabla\phi_t$ provides the unique solution to (4).

Corollary

Use $\mathcal{L}_{\mathrm{Schr}}(\phi,\theta) := D_{\mathrm{KL}}(\overrightarrow{\mathbb{P}}^{\mu,f+\nabla\phi}, \overleftarrow{\mathbb{P}}^{\nu,f+\nabla\theta}) + \lambda \mathrm{Reg}(\phi),$

where $\mathrm{Reg}(\phi) = \int_0^T \mathbb{E}\left|\partial_t\phi + f\cdot\nabla\phi + \frac{1}{2}\Delta\phi + \frac{1}{2}|\nabla\phi|^2\right|(\boldsymbol{Y}_t)\,\mathrm{d}t.$

IPF for Schrödinger bridges

Given: Marginals μ and ν, reference drift f_t
Goal: Find a diffusion that connects μ and ν, and that stays *"close to f_t"*

IPF for Schrödinger bridges

Given: Marginals μ and ν, reference drift f_t
Goal: Find a diffusion that connects μ and ν, and that stays *"close to f_t"*

Our approach:

$$\mathcal{L}_{\mathrm{Schr}}(\phi, \theta) := D_{\mathrm{KL}}(\overrightarrow{\mathbb{P}}^{\mu, f+\nabla\phi}, \overleftarrow{\mathbb{P}}^{\nu, f+\nabla\theta}) + \lambda \mathrm{Reg}(\phi),$$

works for every $\lambda > 0$.

IPF for Schrödinger bridges

Given: Marginals μ and ν, reference drift f_t
Goal: Find a diffusion that connects μ and ν, and that stays *"close to f_t"*

Our approach:

$$\mathcal{L}_{\mathrm{Schr}}(\phi, \theta) := D_{\mathrm{KL}}(\overrightarrow{\mathbb{P}}^{\mu, f+\nabla\phi}, \overleftarrow{\mathbb{P}}^{\nu, f+\nabla\theta}) + \lambda \mathrm{Reg}(\phi),$$

works for every $\lambda > 0$.

Iterative proportional fitting (IPF)

- Initialise $\phi_0 = 0$, i.e., start with the Schrödinger reference $\overrightarrow{\mathbb{P}}^{\mu, f}$.
- Alternately, solve

$$\theta_{i+1} \in \arg\min_{\theta} D_{\mathrm{KL}}(\overbrace{\overrightarrow{\mathbb{P}}^{\mu, f+\nabla\phi_i}}^{\textbf{FIX!}}, \overleftarrow{\mathbb{P}}^{\nu, f+\nabla\theta})$$

$$\phi_{i+1} \in \arg\min_{\phi} D_{\mathrm{KL}}(\overrightarrow{\mathbb{P}}^{\mu, f+\nabla\phi}, \underbrace{\overleftarrow{\mathbb{P}}^{\nu, f+\nabla\theta_i}}_{\textbf{FIX!}})$$

IPF for Schrödinger bridges

Our approach:

$$\mathcal{L}_{\mathrm{Schr}}(\phi,\theta) := D_{\mathrm{KL}}(\overrightarrow{\mathbb{P}}^{\mu,f+\nabla\phi}, \overleftarrow{\mathbb{P}}^{\nu,f+\nabla\theta}) + \lambda \mathrm{Reg}(\phi),$$

works for every $\lambda > 0$.

Iterative proportional fitting (IPF)

- Initialise $\phi_0 = 0$, i.e., start with the Schrödinger reference $\overrightarrow{\mathbb{P}}^{\mu,f}$.
- Alternately, solve

$$\theta_{i+1} \in \arg\min_{\theta} D_{\mathrm{KL}}(\overbrace{\overrightarrow{\mathbb{P}}^{\mu,f+\nabla\phi_i}}^{\textbf{FIX!}}, \overleftarrow{\mathbb{P}}^{\nu,f+\nabla\theta})$$

$$\phi_{i+1} \in \arg\min_{\phi} D_{\mathrm{KL}}(\overrightarrow{\mathbb{P}}^{\mu,f+\nabla\phi}, \underbrace{\overleftarrow{\mathbb{P}}^{\nu,f+\nabla\theta_i}}_{\textbf{FIX!}})$$

IPF for Schrödinger bridges

Our approach:

$$\mathcal{L}_{\mathrm{Schr}}(\phi,\theta) := D_{\mathrm{KL}}(\overrightarrow{\mathbb{P}}^{\mu,f+\nabla\phi}, \overleftarrow{\mathbb{P}}^{\nu,f+\nabla\theta}) + \lambda \mathrm{Reg}(\phi),$$

works for every $\lambda > 0$.

Iterative proportional fitting (IPF)

- Initialise $\phi_0 = 0$, i.e., start with the Schrödinger reference $\overrightarrow{\mathbb{P}}^{\mu,f}$.
- Alternately, solve

$$\theta_{i+1} \in \arg\min_{\theta} D_{\mathrm{KL}}(\overbrace{\overrightarrow{\mathbb{P}}^{\mu,f+\nabla\phi_i}}^{\textbf{FIX!}}, \overleftarrow{\mathbb{P}}^{\nu,f+\nabla\theta})$$

$$\phi_{i+1} \in \arg\min_{\phi} D_{\mathrm{KL}}(\overrightarrow{\mathbb{P}}^{\mu,f+\nabla\phi}, \underbrace{\overleftarrow{\mathbb{P}}^{\nu,f+\nabla\theta_i}}_{\textbf{FIX!}})$$

Selling points:

- One end-to-end training instead of multiple training runs.
- Schrödinger reference is enforced explicitly.
- Historical precedent: Expectation maximisation (EM) vs. variational autoencoders (VAEs).

Other divergences: log-variance

for optimal control, BSDEs

$$\mathrm{d}X_s = \sigma\,\mathrm{d}W_s, \qquad X_0 = x_0,$$

$$\mathrm{d}Y_s^a = \frac{1}{2}|a_s(X_s)|^2\,\mathrm{d}s - a_s(X_s)\cdot\mathrm{d}W_s, \qquad Y_T^a = g(X_T),$$

$$\mathcal{L}_{\mathrm{Var}}^{\ln}(a) = \underbrace{\mathrm{Var}\left(\ln\frac{\mathrm{d}\overrightarrow{\mathbb{P}}^{\delta_0,a}}{\mathrm{d}\overleftarrow{\mathbb{P}}^{\nu,b}}\right)}_{\text{family of divergences}} = \mathrm{Var}(Y_T^a - g(X_T^\nu))$$

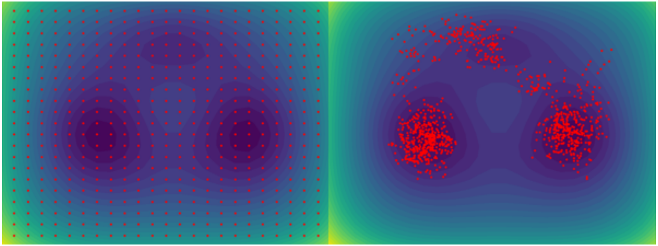

Other divergences: log variance

for Controlled Monte Carlo diffusions, normalising flows

$$\mathcal{L}_{\text{log-Var}}(\phi) = \mathrm{Var}\left(\overbrace{\ln \frac{\pi_T(\boldsymbol{Y}_T)}{\pi_0(\boldsymbol{Y}_0)} + \int_0^T \Delta\phi_t(\boldsymbol{Y}_t)\,\mathrm{d}t}^{\text{normalising flow objective}} \quad + \quad \text{noise} \right),$$

where

$$\text{noise} = -\sqrt{2}\sigma \int_0^T \nabla \ln \pi_t(\boldsymbol{Y}_t) \circ \mathrm{d}\boldsymbol{W}_t - \sigma^2 \int_0^T |\nabla \ln \pi_t(\boldsymbol{Y}_t)|^2\,\mathrm{d}t.$$

Other divergences: log variance

for Controlled Monte Carlo diffusions, normalising flows

$$\mathcal{L}_{\text{log-Var}}(\phi) = \mathrm{Var}\left(\overbrace{\ln \frac{\pi_T(\boldsymbol{Y}_T)}{\pi_0(\boldsymbol{Y}_0)} + \int_0^T \Delta\phi_t(\boldsymbol{Y}_t)\,\mathrm{d}t}^{\text{normalising flow objective}} \quad + \quad \text{noise} \right),$$

where

$$\text{noise} = -\sqrt{2}\sigma \int_0^T \nabla \ln \pi_t(\boldsymbol{Y}_t) \circ \mathrm{d}\boldsymbol{W}_t - \sigma^2 \int_0^T |\nabla \ln \pi_t(\boldsymbol{Y}_t)|^2 \,\mathrm{d}t.$$

Continuous-time normalising flows:

$$\text{If} \qquad \frac{\mathrm{d}\boldsymbol{Y}_t}{\mathrm{d}t} = v_t(\boldsymbol{Y}_t), \qquad X_0 \sim \pi_0,$$

$$\text{then} \qquad \partial_t \ln \pi_t(\boldsymbol{Y}_t) = -\nabla \cdot v_t(\boldsymbol{Y}_t).$$

Contact: nikolas.nusken@kcl.ac.uk

References:

- *Transport meets variational inference: Controlled Monte Carlo diffusions* (with F. Vargas, S. Padhy and D. Blessing).