An optimal control perspective on diffusion-based generative modeling

Julius Berner[⋆,1], Lorenz Richter[⋆,2,3], Karen Ullrich[4]

[⋆] Equal contribution, [1] University of Vienna, [2] dida Datenschmiede GmbH, [3] Zuse Institute Berlin, [4] Meta AI

Overview

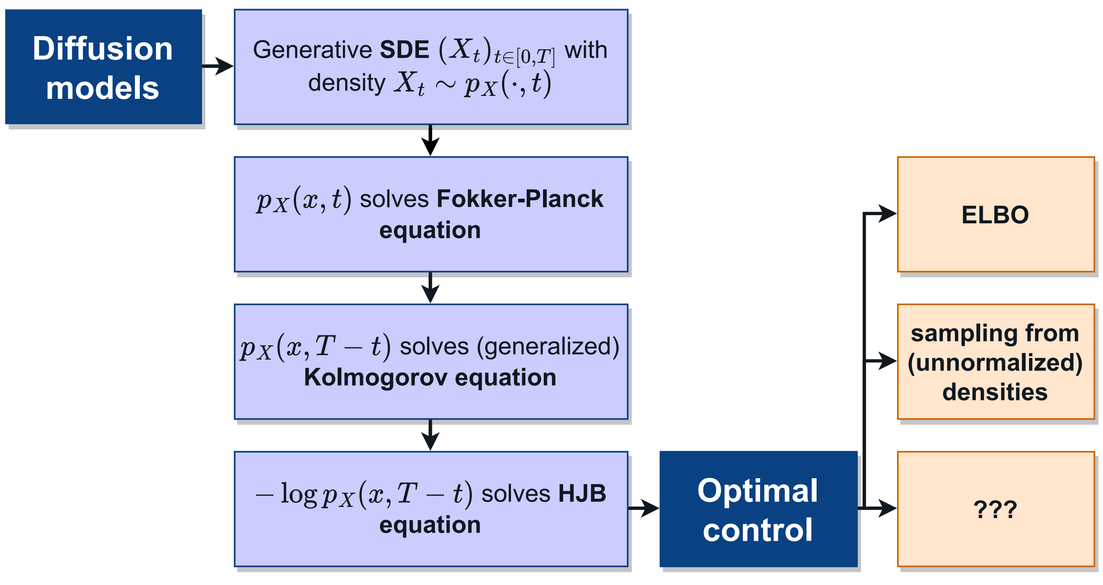

Introduction

- **Goal:** Sample from a data distribution $\mathcal{D}$.
- **SDE-based modeling:** Consider a stochastic differential equation (SDE)

$$\mathrm{d}X_s = \overleftarrow{\mu}(X_s, s)\,\mathrm{d}s + \overleftarrow{\sigma}(s)\,\mathrm{d}B_s, \tag{1}$$

  where $\overleftarrow{\mu}(\cdot,s) := \mu(\cdot, T-s)$ denotes the time-reversal.
- **Problem:** Learn an initial condition X_0 as well as coefficient functions μ and σ such that the distribution of X_T approximates the distribution $\mathcal{D}$.
- **Strategy:** Let us fix X_0 and σ and consider μ as a control function.

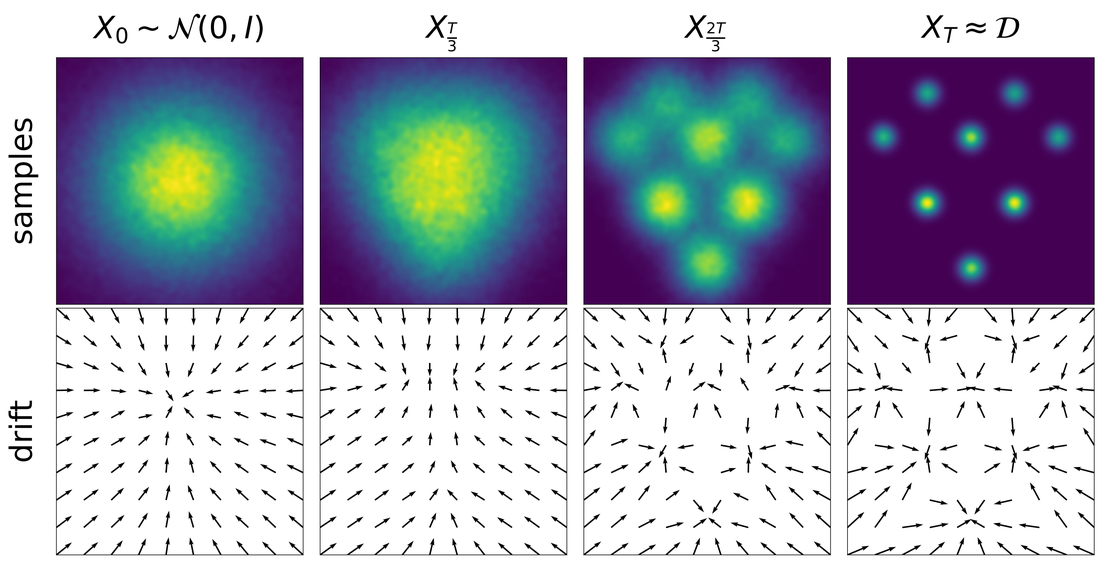

Score-based generative modeling

- **Time-reversal:** One can achieve $X_T \sim \mathcal{D}$ by setting

$$X_0 \sim Y_T \quad \text{and} \quad \mu := \sigma\sigma^\top \nabla \log p_Y - f,$$

  where $Y = (Y_s)_{s\in[0,T]}$ is a solution to the SDE

$$\mathrm{d}Y_s = f(Y_s,s)\,\mathrm{d}s + \sigma(s)\,\mathrm{d}B_s, \quad Y_0 \sim \mathcal{D}.$$

  In fact, it holds that $p_Y = \overleftarrow{p}_X$, where $p_Y(\cdot,t)$ denotes the density of Y_t.
- **Strategy**: Choose f such that Y_T is approximately Gaussian, set $X_0 \sim \mathcal{N}(0,I)$, and learn the **optimal control/score** $u^* := \sigma^\top \nabla \log p_Y$.
- Recall the **Fokker-Planck equation** which describes the evolution of a density of a solution X to an SDE via the PDE

$$\partial_t p_X = \operatorname{div}\left(\operatorname{div}\left(\overleftarrow{D} p_X\right) - \overleftarrow{\mu} p_X\right),$$

  where $D := \frac{1}{2}\sigma\sigma^\top$. This implies that the time-reversed density $\overleftarrow{p}_X$ satisfies the **Kolmogorov backwards equation**

$$\partial_t \overleftarrow{p}_X = -\operatorname{Tr}\left(D\nabla^2 \overleftarrow{p}_X\right) + \mu\cdot\nabla\overleftarrow{p}_X + \operatorname{div}(\mu)\overleftarrow{p}_X. \tag{2}$$

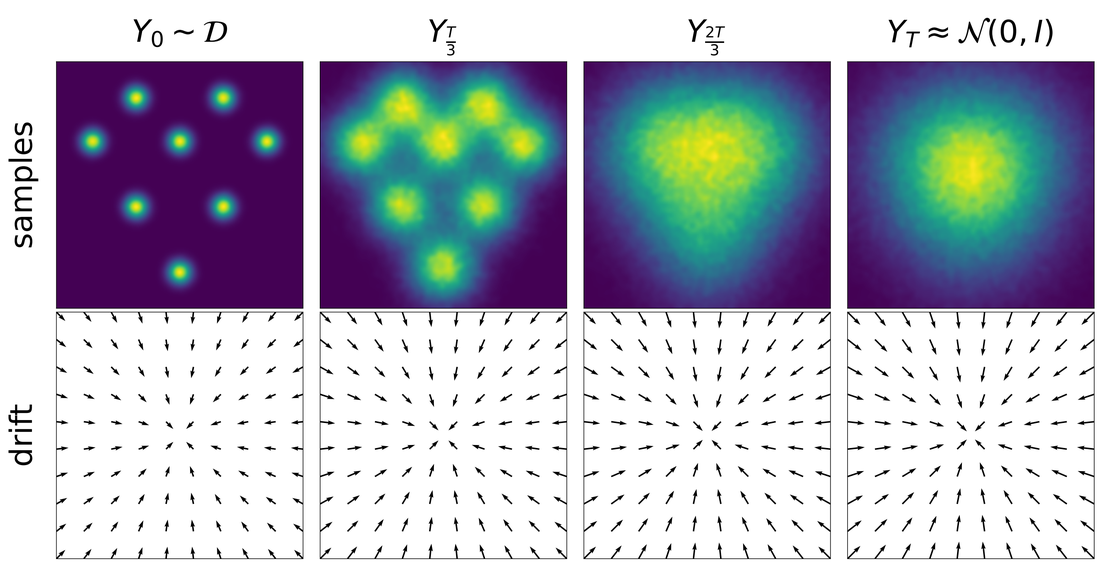

Generative modeling as control problem

- We use the **Hopf–Cole transformation** to transform the linear PDE (2) to an **HJB equation** prominent in control theory.

Lemma (HJB for log-density). *Define* $V := -\log \overleftarrow{p}_X$. *Then*

$$\partial_t V = -\operatorname{Tr}\left(D\nabla^2 V\right) + \mu\cdot\nabla V - \operatorname{div}(\mu) + \frac{1}{2}\left\|\sigma^\top \nabla V\right\|^2$$

with terminal condition $V(\cdot,T) = -\log p_{X_0}$.

- The HJB equation now allows to identify our **control problem**.

Theorem (Verification theorem). *Let* Y^u *be defined by*

$$\mathrm{d}Y^u_s = (\sigma u - \mu)(Y^u_s, s)\,\mathrm{d}s + \sigma(s)\,\mathrm{d}B_s.$$

Conditioned on Y_0, *it then holds almost surely that*

$$-\log p_{X_T}(Y^u_0) = \min_{u\in\mathcal{U}} \mathbb{E}\left[\int_0^T \left(\operatorname{div}(\mu) + \frac{\|u\|^2}{2}\right)(Y^u_s, s)\,\mathrm{d}s - \log p_{X_0}(Y^u_T)\right],$$

where the unique minimum is attained by $u^* := \sigma^\top \nabla \log p_Y$.

- The control costs equal the negative **ELBO**, see Huang et al. (2021).
- Via the reparametrization $\mu := \sigma u - f$, which yields the SDEs

$$\mathrm{d}Y_s = f(Y_s,s)\,\mathrm{d}s + \sigma(s)\,\mathrm{d}B_s, \quad Y_0 \sim \mathcal{D}, \tag{3}$$

$$\mathrm{d}X^u_s = \left(\overleftarrow{\sigma}\overleftarrow{u} - \overleftarrow{f}\right)(X^u_s, s)\,\mathrm{d}s + \overleftarrow{\sigma}(s)\,\mathrm{d}B_s, \tag{4}$$

  this can be related to the **denoising score matching objective**

$$-\log p_{X^u_T}(Y_0) \le \frac{T}{2}\mathbb{E}\left[\left\|u(Y_\tau,\tau) - \sigma^\top(\tau)\nabla\log p_{Y_\tau|Y_0}(Y_\tau|Y_0)\right\|^2\right] + \text{const.},$$

  where $\tau \sim \mathcal{U}([0,T])$ and $p_{Y_\tau|Y_0}$ denotes the conditional density of Y_τ given Y_0 (which can be explicitly computed for certain SDEs).

Sampling from (unnormalized) densities

- **Application:** Data distribution $\mathcal{D}$ admits the density $\rho/\mathcal{Z}$, where ρ is known, but the normalizing constant $\mathcal{Z}$ is intractable.
- **Sampling method based on diffusion models:** Interchange X and Y^u as in (3) and (4). The verification theorem yields the control objective

$$-\log\mathcal{Z} = \min_{\overleftarrow{u}\in\mathcal{U}} \mathbb{E}\left[\int_0^T \left(\operatorname{div}(\overleftarrow{f}) + \frac{\|\overleftarrow{u}\|^2}{2}\right)(X^u_s, s)\,\mathrm{d}s + \log\frac{p_{Y_T}(X^u_0)}{\rho(X^u_T)}\right].$$

- **Path measure interpretation**: We can rewrite the objective (up to $\log\mathcal{Z}$) as Kullback-Leibler divergence

$$D_{\mathrm{KL}}(\mathbb{P}_{X^u}|\mathbb{P}_{X^{u^*}}) = D_{\mathrm{KL}}(\mathbb{P}_{X^u}|\mathbb{P}_{\overleftarrow{Y}}) - D_{\mathrm{KL}}(\mathbb{P}_{X^u_0}|\mathbb{P}_{Y_T}),$$

  where $\mathbb{P}_{X^u}$ denotes the path measure of X^u and $\overleftarrow{Y}$ and denotes the solution to the reverse-time SDE.
- **Example:** Successful application to tasks in molecular dynamics, e.g.,

$$\rho(x) := \exp\left(-\sum_{i=1}^{5}(x_i^2-2)^2 - \frac{1}{2}\sum_{i=6}^{50} x_i^2\right)$$

  in 50 dimensions with 32 modes.

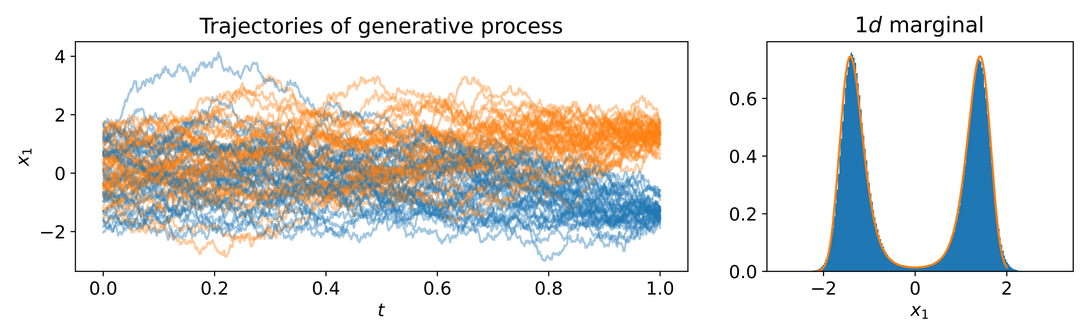

Conclusion and outlook

- Our perspective allows to transfer methods from optimal control theory to generative modeling and vice versa.
- ELBO and diffusion-based sampling methods are direct consequences from the control theory perspective.

arxiv.org/abs/2211.01364