Learning diffusion at lightspeed

ETH zürich

Antonio Terpin, Nicolas Lanzetti, Martin Gadea, Florian Dörfler

NEURAL INFORMATION PROCESSING SYSTEMS

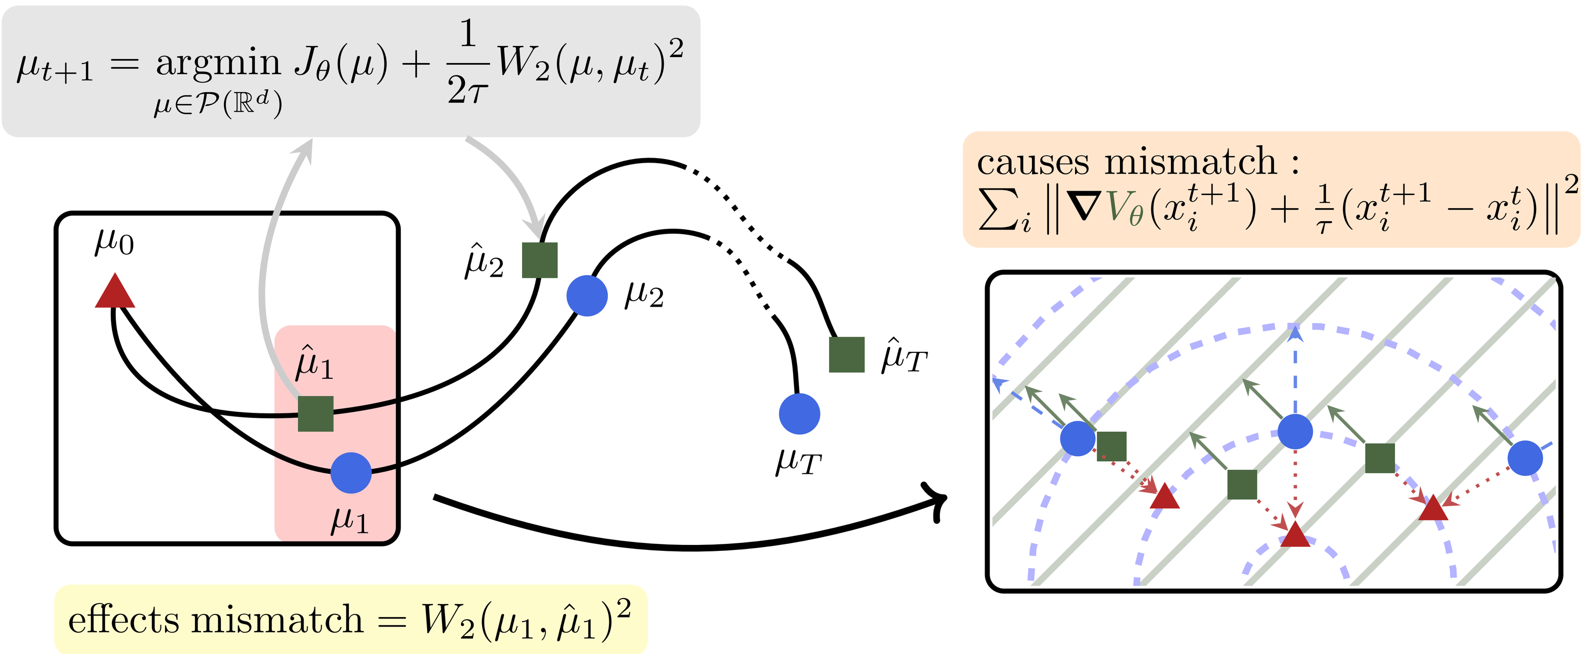

Motivation

- Diffusion processes are widespread in many natural processes.
- Typically characterized by three terms: a drift term due to a potential field, the interaction with other particles, and a stochastic term.
- Existing methods either rely on having access to particles trajectories (often not possible) or are data and compute inefficient.

The JKO scheme

The Fokker-Planck equation,

$$\frac{\partial \rho(t,x)}{\partial t} = \nabla \cdot (\nabla V(x)\rho(t,x)) + \beta \nabla^2 \rho(t,x),$$

describes the time evolution of the distribution ρ of a set of particles undergoing drift and diffusion,

$$\mathrm{d}X(t) = -\nabla V(X(t))\mathrm{d}t + \sqrt{2\beta}\mathrm{d}W(t),$$

where $X(t)$ is the state of the particle, $V(x)$ the driving potential, and $W(t)$ the Wiener process. It turns out the resulting particles trajectory can be described via the JKO scheme:

$$\mu_{t+1} = \underset{\mu \in \mathcal{P}(\mathbb{R}^d)}{\operatorname{argmin}} J(\mu) + \frac{1}{2\tau} W_2(\mu, \mu_t)^2, \qquad (1)$$

where J is an energy functional and $\tau > 0$ is the time discretization.

A general energy functional:

$$J(\mu) = \underbrace{\int_{\mathbb{R}^d} V(x)\mathrm{d}\mu(x)}_{\text{potential}} + \underbrace{\int_{\mathbb{R}^d \times \mathbb{R}^d} U(x-y)\mathrm{d}(\mu \times \mu)(x,y)}_{\text{interaction}} + \beta \underbrace{\int_{\mathbb{R}^d} \rho(x)\log(\rho(x))\mathrm{d}x}_{\text{internal}}.$$

Intuition in $\mathbb{R}^d$

Consider the analog of (1) in the $\mathbb{R}^d$,

$$x_{t+1} = \underset{x \in \mathbb{R}^d}{\operatorname{argmin}} J(x) + \frac{1}{2\tau}\|x - x_t\|^2.$$

We replace the above by its first-order optimality condition

$$\nabla J(x_{t+1}) + \frac{1}{\tau}(x_{t+1} - x_t) = 0.$$

Given a dataset $(x_0, x_1, \ldots, x_T)$, we find J as:

$$\min_J \sum_{t=0}^{T-1} \left\| \nabla J(x_{t+1}) + \frac{1}{\tau}(x_{t+1} - x_t) \right\|^2.$$

Our work enables an analogous result in the Wasserstein space.

First-order optimality conditions for the JKO scheme

Our analysis suggest that given a populations dataset $(\mu_0, \mu_1, \ldots, \mu_T)$ we seek the parameters θ minimizing

$$\sum_{t=0}^{T-1} \int_{\mathbb{R}^d \times \mathbb{R}^d} \left\| \nabla V_\theta(x_{t+1}) + \int_{\mathbb{R}^d} \nabla U_\theta(x_{t+1} - x'_{t+1})\mathrm{d}\mu_{t+1}(x'_{t+1}) + \beta_\theta \frac{\nabla \rho_{t+1}(x_{t+1})}{\rho_{t+1}(x_{t+1})} + \frac{1}{\tau}(x_{t+1} - x_t) \right\|^2 \mathrm{d}\gamma_t(x_t, x_{t+1}), \qquad (2)$$

where $\gamma_0, \gamma_1, \ldots, \gamma_T$ are optimal transport plans between μ_t and μ_{t+1}.

- γ_t can be computed once, before-hand, and efficiently;
- when J_θ is a neural network, we minimize (2) via gradient descent;
- when J_θ is parametrized linearly, we have a closed-form solution.

Training at lightspeed

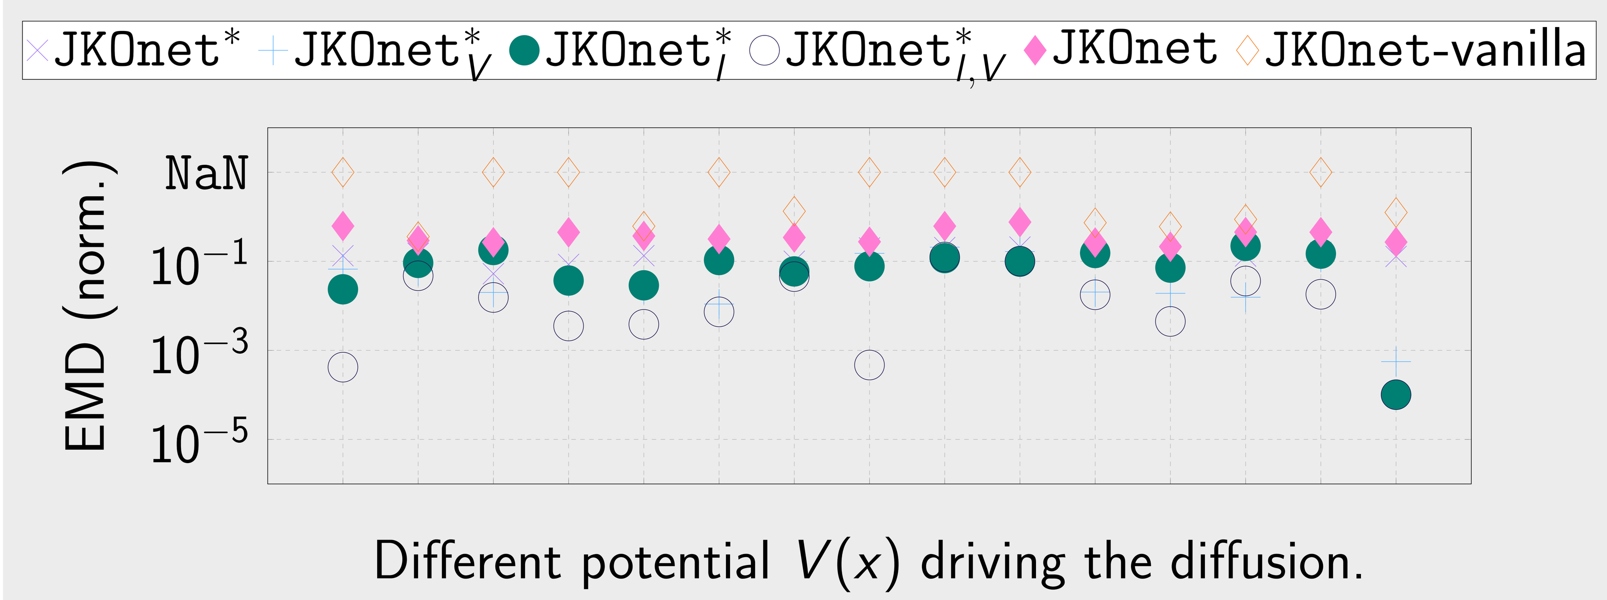

Different potential $V(x)$ driving the diffusion.

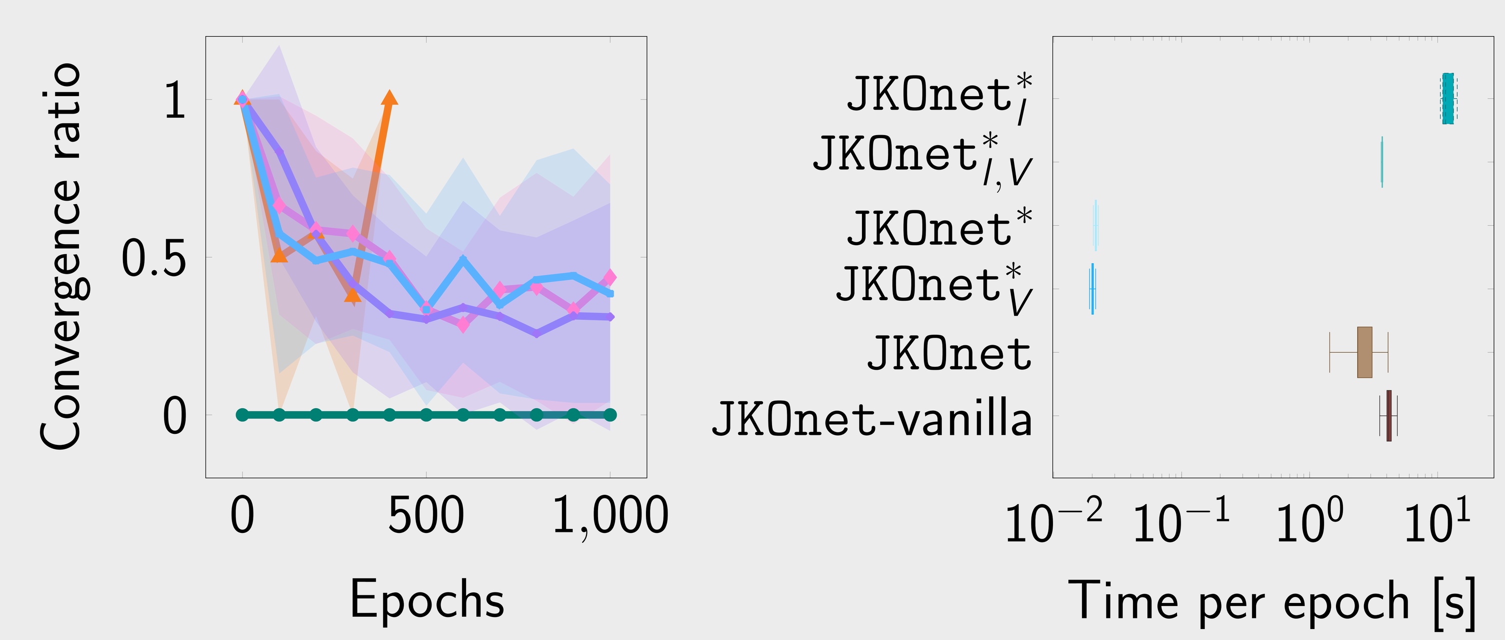

Eye-candies

Level curves of different potentials (green: real; blue: estimated).

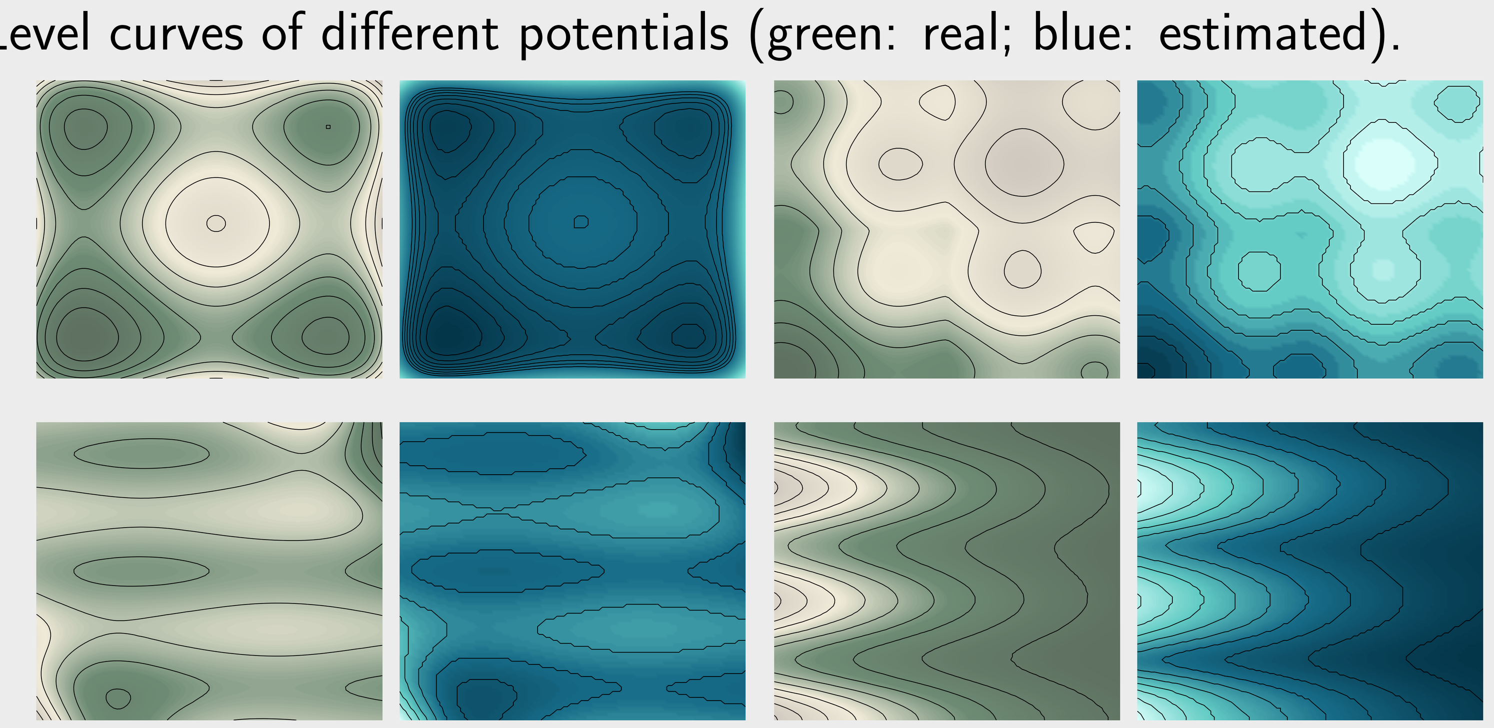

Scaling to high-dimensions

- Sub-linear error growth.
- Negligible increase in training time.

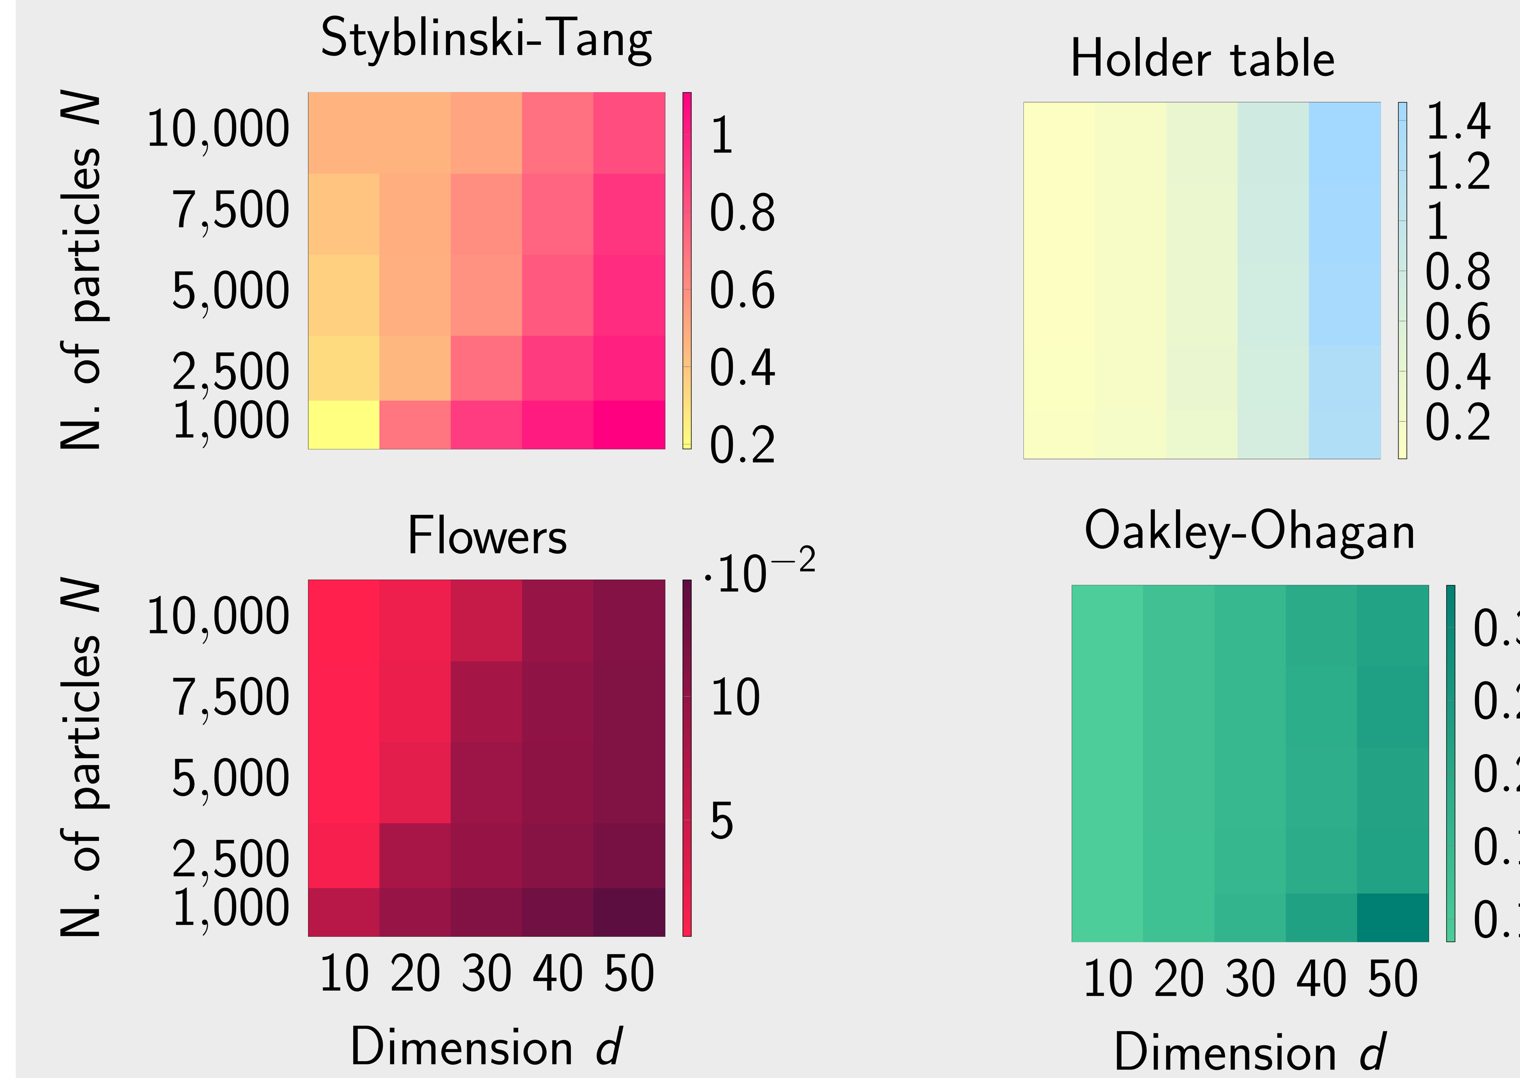

Learning general energy terms

Additional error sources:

- Sampling error (internal energy)
- Estimation of the densities

Nonetheless: our method (JKOnet*, non-linear and linearly parametrized) recovers all the energy terms.

Open question: observability of the different energy terms?

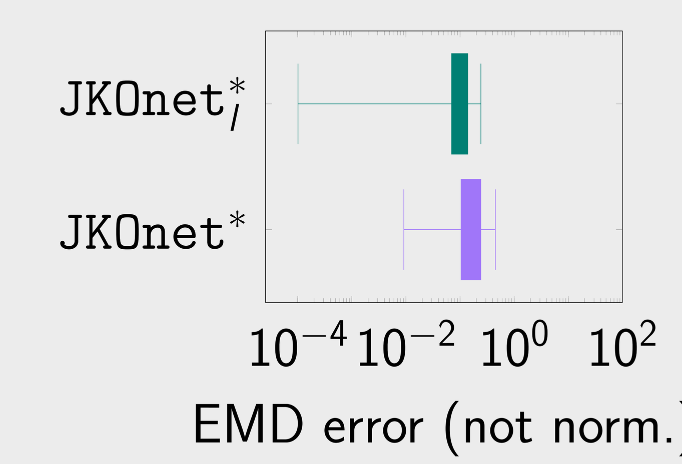

Learning single-cell diffusion dynamics

- Intuition: cellular evolution minimizes (some) energy;
- We can account for unobserved variables via time-varying energies;
- State-of-the-art accuracy at a fraction of the computational cost of the other methods: less than minute vs hours.

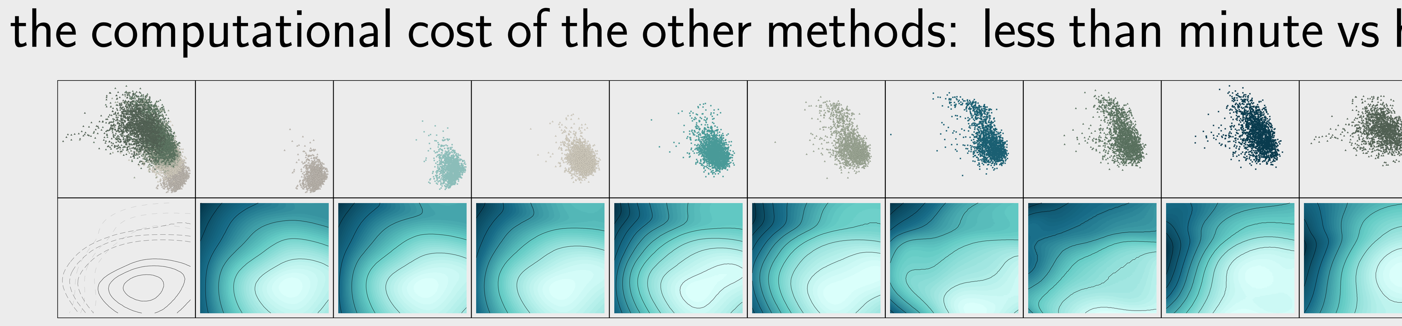

Cool stuff worth looking at next

- Fast distillation of diffusion models for one- or few-steps generation;
- Observability of the different energy terms;
- Can optimality conditions in $\mathcal{P}(\mathbb{R}^d)$ be helpful in your work? Reach out!