Using Linearized Optimal Transport to Predict the Evolution of Stochastic Particle Systems

Nicholas Karris[1], **Evangelos A. Nikitopoulos**[2], **Ioannis G. Kevrekidis**[3], **Seungjoon Lee**[4], and **Alexander Cloninger**[1]

[1]University of California, San Diego, [2]University of Michigan, [3]Johns Hopkins University, [4]California State University, Long Beach

Simulating Particle Systems

General Set Up: Have collection of particles $\{x_i(t)\}_{i=1}^N \subset \mathbb{R}^d$ that form an empirical distribution μ_t^N, and have some dynamics that describe their evolution:

$$\mu_{t+h}^N \approx E_h(\mu_t^N)$$

for some noisy E_h. We can numerically simulate the particle system by successively applying E_h for some small h.

Goal: Want to use μ_t^N and μ_{t+h}^N to approximate μ_{t+H}^N without needing to apply E_h repeatedly. If $H \gg h$, this drastically reduces computation time for simulating many steps (e.g., to approximate a steady state distribution).

Idea: Use an analog of Euler's method:

$$\mu_{(n+1)} = (\mathrm{id} + H \cdot F(t_n, \mu_{(n)}))_{\#}\mu_{(n)}$$

where $F(t_n, \mu_{(n)})$ is the tangent vector field describing the flow of mass of $\mu_{(n)}$.

Why Optimal Transport?

Motivating Theorem: [1] Let $t \mapsto \mu_t \in \mathcal{P}_2(\mathbb{R}^d)$ be absolutely continuous in the Wasserstein-2 distance. Then there is a unique minimal-energy vector field $\mathbf{v}_t : \mathbb{R}^d \to \mathbb{R}^d$ such that the continuity equation

$$\partial_t \mu_t + \nabla \cdot (\mathbf{v}_t \mu_t) = 0$$

holds in the sense of distributions. Moreover, for $\mathcal{L}^1$-a.e. t,

$$\lim_{H \to 0} \frac{W_2(\mu_{t+H}, (\mathrm{id} + H\mathbf{v}_t)_{\#}\mu_t)}{|H|} = 0 \qquad (*)$$

and, assuming $\mu_t \ll \mathcal{L}^d$,

$$\lim_{h \to 0} \frac{1}{h}(T_{\mu_t}^{\mu_{t+h}} - \mathrm{id}) = \mathbf{v}_t \qquad \text{in } L^2(\mu_t). \qquad (\dagger)$$

Key Result: We then show that Euler's method on the Wasserstein manifold is first-order accurate:

Theorem (K., Nikitopoulos, Kevrekidis, Lee, Cloninger (2025)) [2]

Let $F(t, \mu)$ be a vector field in the tangent space $T_\mu \mathbb{W}_2$ for all $t \in [0, T]$ and $\mu \in \mathcal{P}_2(\mathbb{R}^d)$, and suppose $t \mapsto \mu_t$ is an absolutely continuous curve satisfying the continuity equation with $\mathbf{v}_t := F(t, \mu_t)$. Under regularity assumptions on F, we can guarantee that the characteristic ODE

$$\gamma_x(0) = x, \qquad \frac{d}{dt}\gamma_x(t) = \mathbf{v}_t(\gamma_x(t))$$

admits a unique C^2 solution for μ_0-a.e. $x \in \mathbb{R}^d$. Furthermore, if

$$M := \left(\int_{\mathbb{R}^d} \max_{t \in [0,T]} \|\gamma_x''(t)\|_2^2 \, d\mu_0(x) \right)^{\frac{1}{2}} < \infty,$$

then for $t_n := nH$, the analog of Euler's method above with $\mu_{(0)} = \mu_0$ satisfies

$$W_2(\mu_{(n)}, \mu_{t_n}) \leq \frac{HM}{2(L_1 + L_2)} \left(e^{nH(L_1+L_2)} - 1 \right),$$

where L_1, L_2 are Lipschitz constants of F.

Experimental Results

Gaussian Diffusion: Results of applying our discrete Euler-type algorithm on a simulation of particles diffusing in two dimensions.

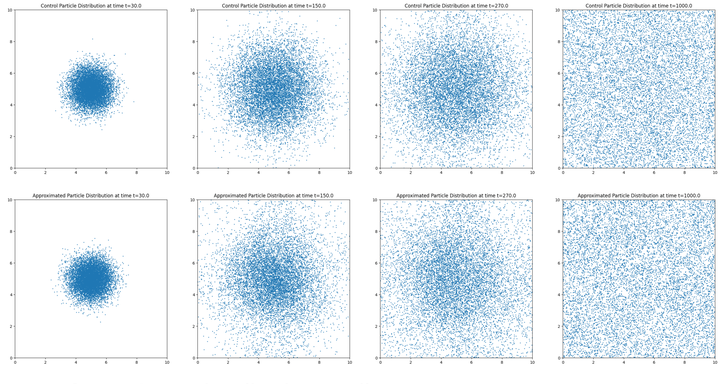

We perform PCA to visualize how well our method approximates the distribution's true evolution.

- Color: Control (true) evolution using an Euler-Maruyama microscale simulation of Gaussian diffusion
- Black: Approximated evolution using our discrete approximation algorithm

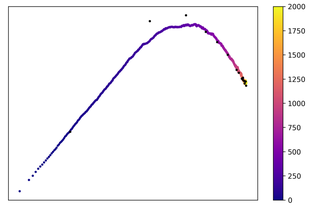

Bacterial Chemotaxis: We also apply the algorithm on a bacterial chemotaxis simulation that models bacteria responding to a chemoattractant, as described in [3].

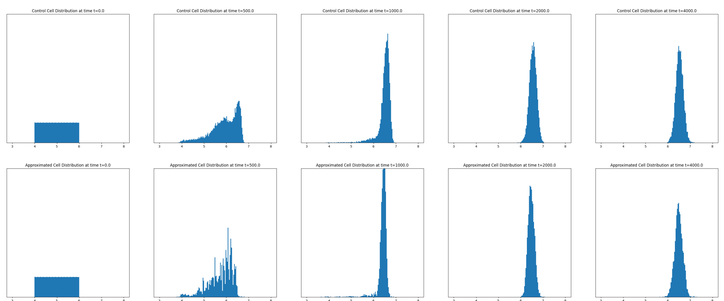

PCA plot to compare the approximations with the control cell distributions. The approximations appear slightly behind the true curve because some underlying cell variables take longer to equilibriate after each Euler-step. To accurately approximate the evolution, we must accurately predict *all* simulation data, not just bacteria locations.

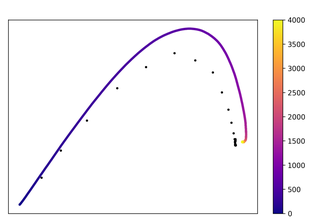

Choosing an Optimal Step Size

We approximate the velocity field

$$\mathbf{v}_t(x_i) \approx \frac{1}{h}(T_{\mu_t}^{\mu_{t+h}}(x_i) - x_i) =: \mathbf{v}_t^h(x_i)$$

for an appropriate h. Why not take $h \to 0$?

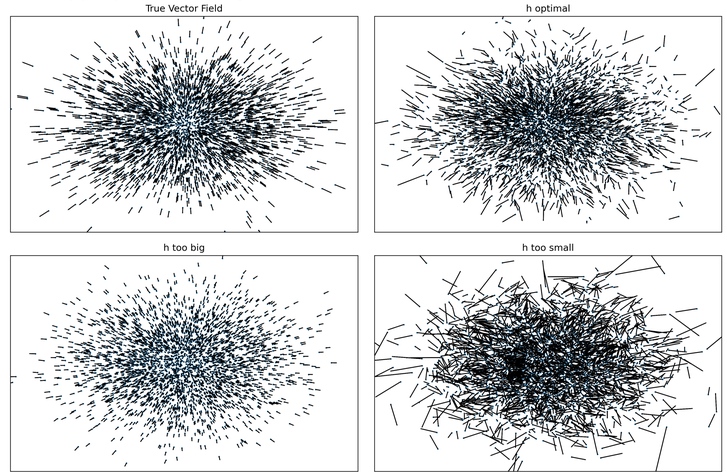

For Gaussian diffusion, we know the true underlying vector field $\mathbf{v}_t$. The approximation error $\|\mathbf{v}_t - \mathbf{v}_t^h\|_{L^2(\mu_t)}$:

- is large for small h,
- approaches the theoretical finite difference error for large h.

Empirically, there is an optimal choice for h that decreases as N increases.

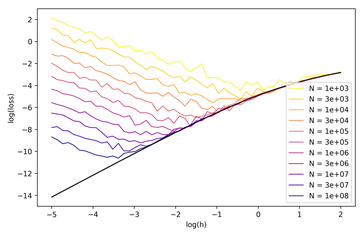

Connections to Fast-Slow Systems

Application: We get a "time stepper" that sits between the two time scales of particle systems. We can numerically integrate on the long scale, which reduces the number of microscale steps needed to observe the long-term behavior of interest.

Ongoing Work: (1) How do we choose an appropriate h without knowing the underlying (theoretical) vector field? (2) Can we numerically estimate the second-derivative bound M based on simulation data? (3) Is there an analogous Newton-type method for quickly approximating the steady state distribution?

References

[1] Luigi Ambrosio, Nicola Gigli, and Giuseppe Savaré. *Gradient Flows in Metric Spaces and in the Space of Probability Measures*. Lectures in Mathematics ETH Zürich. Birkhäuser, 2. ed edition, 2008. OCLC: 254181287.

[2] Nicholas Karris, Evangelos A. Nikitopoulos, Ioannis G. Kevrekidis, Seungjoon Lee, and Alexander Cloninger. Using linearized optimal transport to predict the evolution of stochastic particle systems, 2025.

[3] S. Setayeshgar, C. W. Gear, H. G. Othmer, and I. G. Kevrekidis. Application of coarse integration to bacterial chemotaxis. *Multiscale Modeling & Simulation*, 4(1):307–327, 2005.