Hamiltonian Score Matching and Generative Flows

Conference on Neural Information Processing Systems (NeurIPS)

Vancouver, 2024

Peter Holderrieth **Yilun Xu** **Tommi Jaakkola**

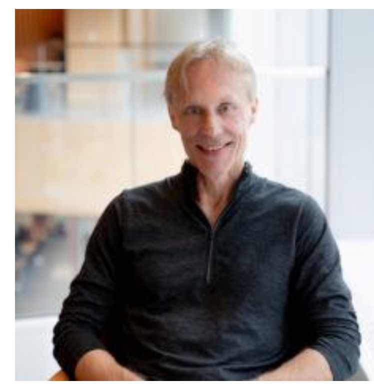

Overview

- **Previous work**: Classical Hamiltonian mechanics has been widely used in ML via Hamiltonian Monte Carlo in applications with *pre-determined* force fields

- **Motivation of this work**: *Design* force fields for Hamiltonian ODEs for the purpose of generative modeling

Contributions

- (1) **Hamiltonian Score Matching (HSM):** learn score functions via Hamiltonian trajectories
- (2) **Hamiltonian Generative Flows (HGFs)**: a generative model framework building on Hamiltonian velocity predictors

Background: Score Matching and Hamiltonian Dynamics

Score Matching:

Given data distribution π score is defined as

$$\nabla \log \pi$$

Score Matching (SM) aims to learn score.

Previous methods:

- Denoising SM
- Implicit SM

Hamiltonian: Energy of system

$$H(x, v) = U(x) + \frac{1}{2}\|v\|^2$$

Potential energy *(neg log-likelihood of data)*

Kinetic energy *(neg log-likelihood of Gaussian)*

Hamiltonian dynamics

$$\left(\frac{d}{dt}x(t), \frac{d}{dt}v(t)\right) = (v(t), -\nabla U(x(t)))$$
$$= (v(t), \nabla \log \pi(x(t)))$$

Score function

Can we use the connection between Hamiltonian dyn. and score function to learn score functions from data?

Characterize Score via Hamiltonian Dynamics

- Given data distribution π, define Boltzmann-Gibbs distribution:
  $$\pi_{BG} = \pi \otimes \mathcal{N}(0, \mathbf{I}_d), \quad \pi_{BG}(x, v) = \exp(-H(x, v))/Z = \pi(x)\mathcal{N}(v; 0, \mathbf{I}_d)$$
- Hamiltonian dynamics preserve Boltzmann-Gibbs distribution – this is used in Hamiltonian Monte Carlo.
- **Idea**: Can we use this fact to uniquely characterize the score function?

Theorem 1. *Let $T > 0$ and $F_\theta(x)$ a force field. Let $\Pi = \pi_{BG} = \pi \otimes \mathcal{N}(0, \mathbf{I}_d)$. The following statements are equivalent:*

1. ***Score vector field:*** *The force field F_θ equals the score, i.e. $F_\theta(x) = \nabla_x \log \pi(x)$ for π-almost every $x \in \mathbb{R}^d$.*
2. ***Preservation of Boltzmann-Gibbs:*** *The PH-ODE with F_θ preserves the Boltzmann-Gibbs distribution π_{BG}.*
3. ***Conditional velocity is zero:*** *The velocity given the location after running the PH-ODE with F_θ is zero if starting conditions $z = (x_0, v_0)$ are sampled from π_{BG}:*
   $$z \sim \pi_{BG} \quad \Rightarrow \quad \mathbb{E}[v_t^\theta(z) | x_t^\theta(z)] = 0 \quad \textit{for all } 0 \leq t < T \tag{13}$$

Idea: Use characterization of score via velocity predictors for score matching

Hamiltonian Score Discrepancy

Define Loss Function:

$$L_{\text{hsm}}(\phi|\theta,t) = \mathbb{E}_{z\sim\pi_{BG}}\left[\|V_\phi(x_t^\theta,t)\|^2 - 2V_\phi(x_t^\theta,t)^T v_t^\theta\right]$$

Squared norm of predicted velocity

Similarity to actual velocity

Loss at optimal velocity:

$$\mathbb{D}_{hsm}(\theta|t,\pi) := -\min_{\phi\in I} L_{hsm}(\phi|\theta,t) = \mathbb{E}_{z\sim\pi_{BG}}\left[\|\mathbb{E}[v_t^\theta|x_t^\theta]\|^2\right]$$

We want this to be zero

Theorem: Minimization of the HSD results in learning the score:

$$\theta^* = \arg\min_\theta \mathbb{D}_{hsm}(\theta|\pi) \Rightarrow s_{\theta^*} = \nabla\log\pi$$

Hamiltonian Score Matching

Minimize Hamiltonian Score Discrepancy by jointly training velocity predictor and score network

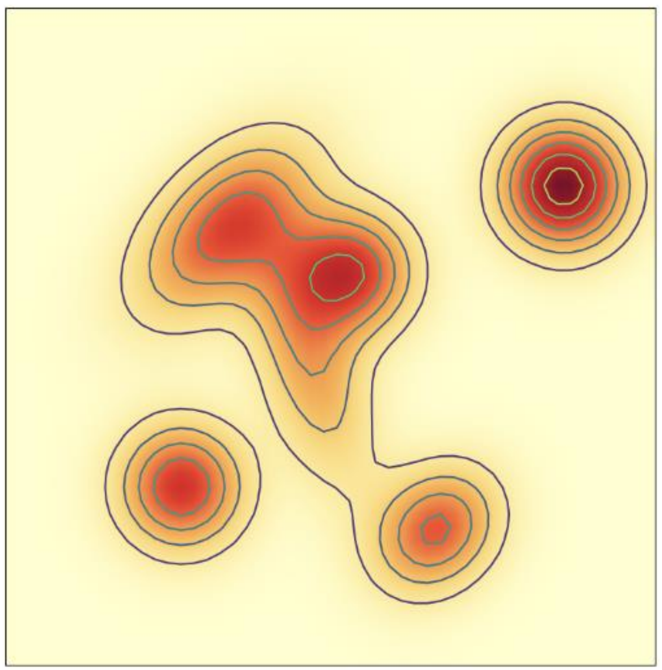

(a) Density $\pi(x)$

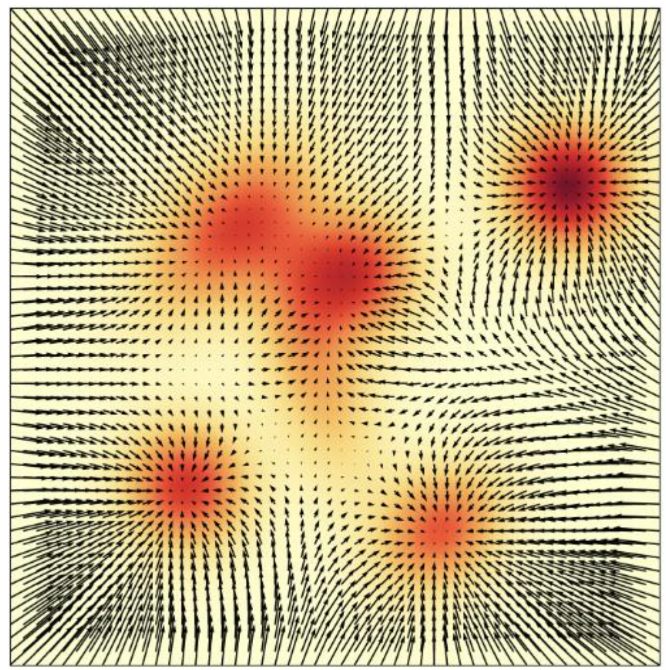

(b) Learnt score F_θ

Vectors closely approximate gradient of log-likelihood

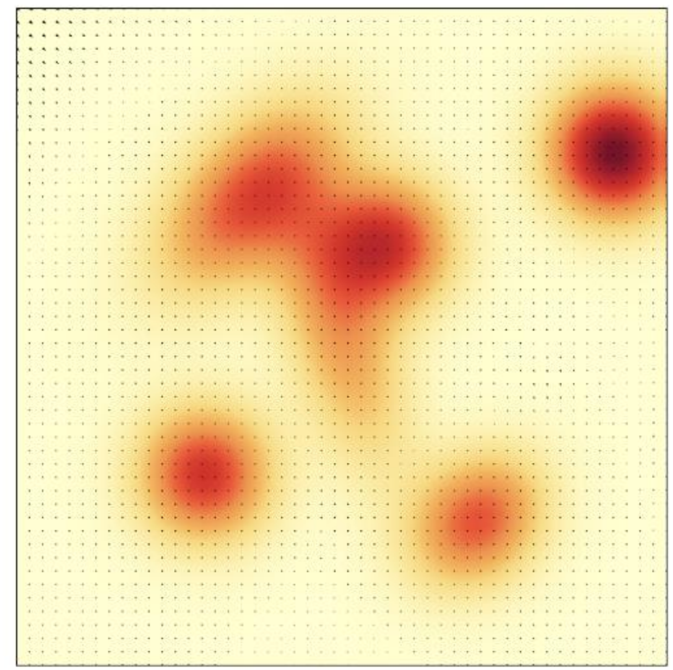

(c) Learnt velocity predictor V_ϕ

Optimal velocity is close to zero everywhere

Hamiltonian Generative Flows

Idea: Can we use Hamiltonian velocity predictors also for suboptimal force fields?

Yes, by **simulating a CNF** with the optimal velocity predictor backwards in time:

$$x_T \sim \pi_T \quad \frac{d}{dt}x(t) = V_{\phi^*}(x,t) \Rightarrow x(0) \sim \pi$$

Resulting **generative model** is similar to FM and diffusion.

2 design choices:

- Force field
- Coupling of distribution over phase space (x,v)

2 requirements:

- Simulation of dynamics with force field have to be efficient
- Tractable distribution at T>0

Example of Hamiltonian Generative Flows - 1

- **Diffusion Models**: zero force field, independent coupling of location and velocity.
- (CondOT) **Flow Matching**: zero force field, coupled velocity and location.
- **Oscillation HGFs**: force field corresponding to simple pendulum

Different HGFs in phase space

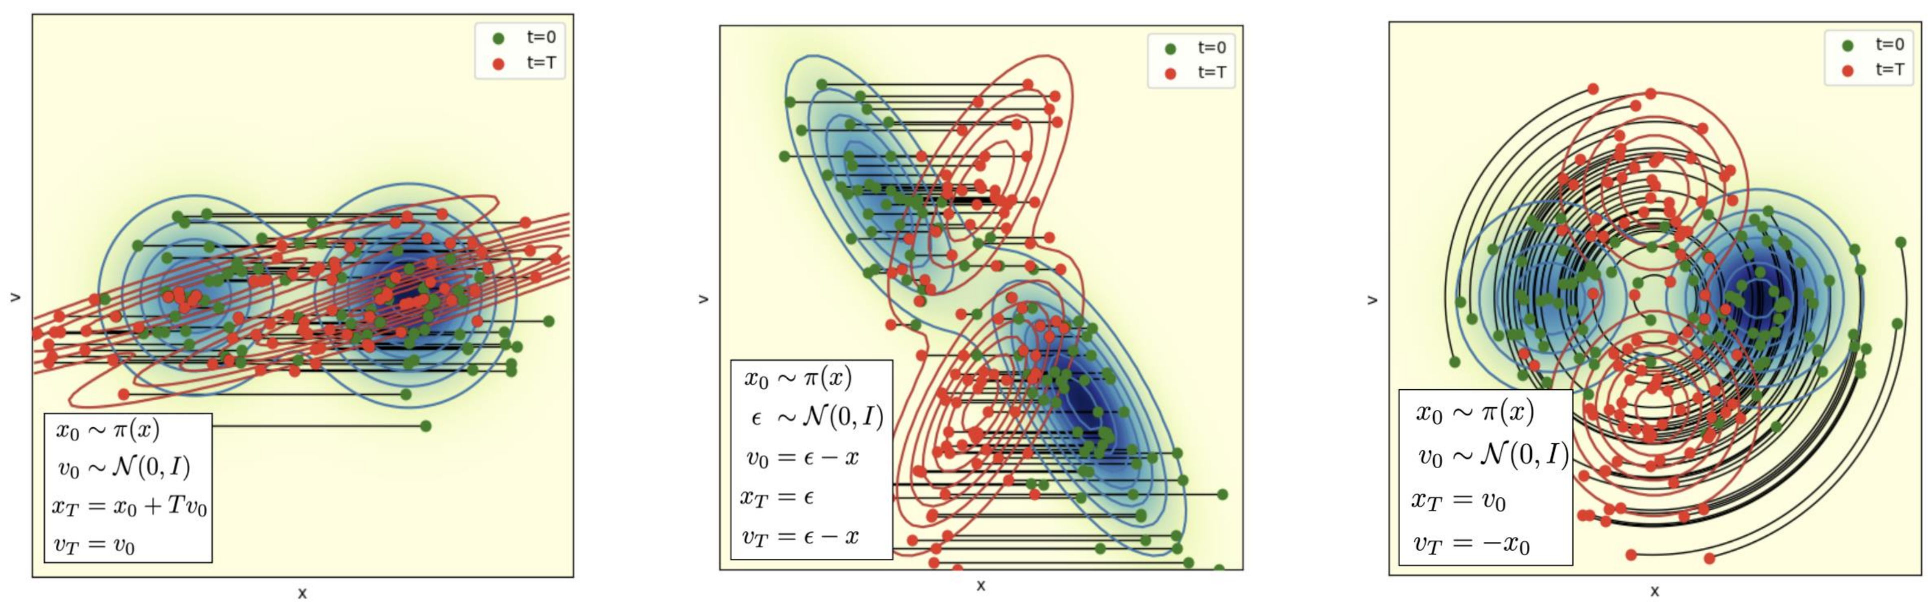

(a) Diffusion (b) Flow matching (c) Oscillation HGFs.

Example of HGFs – 2: Reflection HGFs

Idea: Particles move freely in a box with reflection ("infinite force") at walls

Convergence: Distribution of particles will converge towards a uniform distribution

Training can be done *simulation-free*.

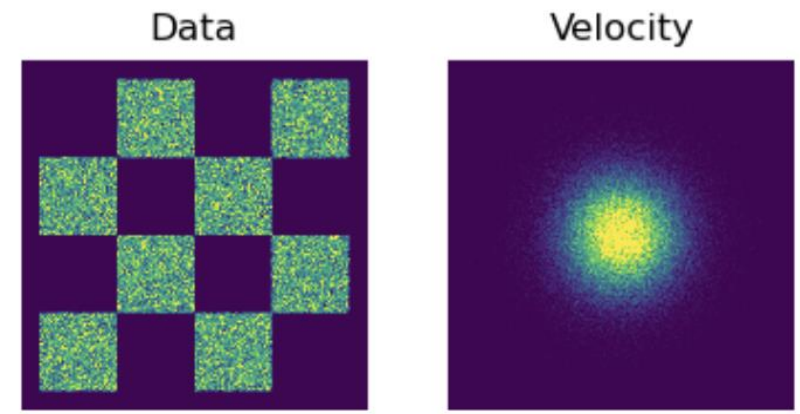

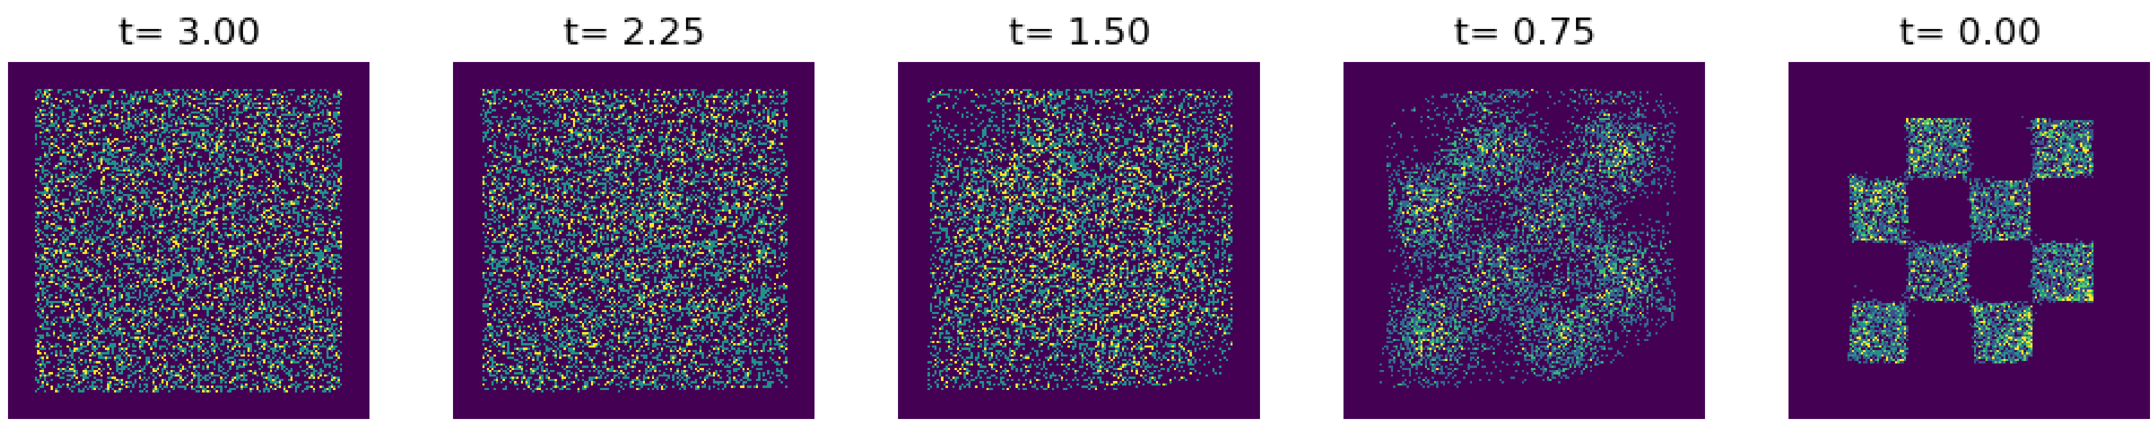

Uniform distribution (max entropy)

Data distribution

Results

(1) Validated HSD as a novel metric and HSM as a novel score matching method

(2) Achieved near-SOTA results with Oscillation HGFs on image generation

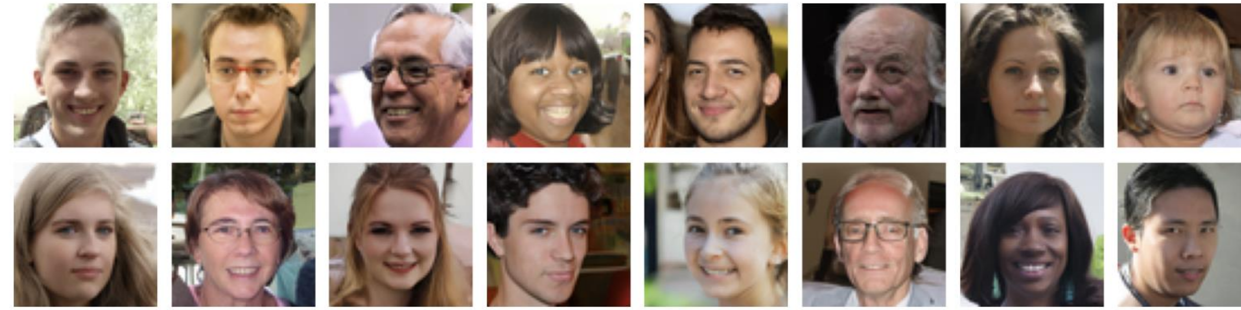

Figure 4: Image generation examples based on Oscillation HGFs for FFHQ.

Table 1: Sample quality (FID) and number of function evaluation (NFE).

METHOD	FID ↓	NFE ↓
CIFAR-10 (unconditional)-32x32		
StyleGAN2-ADA [27]	2.92	1
DDPM [21]	3.17	1000
LSGM [47]	2.10	147
PFGM [50]	2.48	104
VE-SDE [45]	3.77	35
VP-SDE [45]	3.01	35
EDM [26]	1.98	35
FM-OT (BNS) [40]	2.73	8
Oscillation HGF (ours)	2.12	35
CIFAR-10 (class conditional)-32x32		
VE-SDE [45]	3.11	35
VP-SDE [45]	2.48	35
EDM [26]	1.79	35
Oscillation HGF (ours)	1.97	35
FFHQ (unconditional)-64x64		
VE-SDE [45]	25.95	79
VP-SDE [45]	3.39	79
EDM [26]	2.39	79
Oscillation HGF (ours)	2.86	79

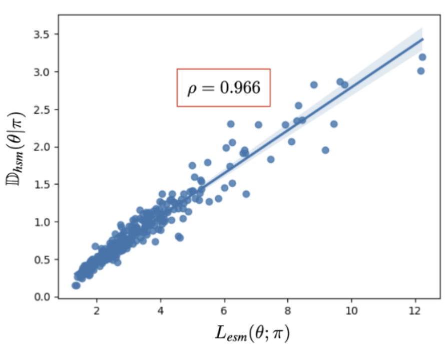

(a) ESM loss vs HSD for networks trained for 1 epoch

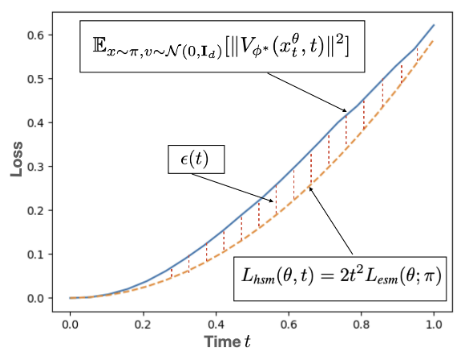

(b) Empirical HSD vs. Taylor approximation (see Proposition 2)

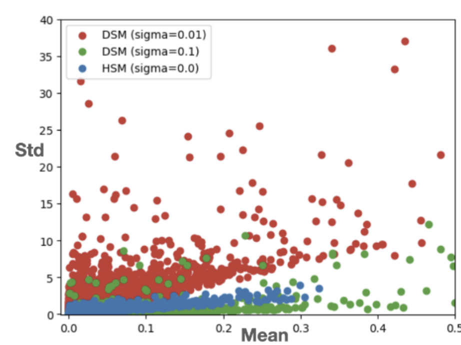

(c) Std vs absolute mean of derivative of param. of score network.

Thank you for your attention!

Contact: phold@mit.edu

Poster session: Fri 13 Dec 4:30 p.m. PST — 7:30 p.m. PST

Visit us at our poster session!

Thank you to my co-authors!

Peter Holderrieth

Yilun Xu

Tommi Jaakkola