Ai-Sampler: *A*dversarial Learning of Markov kernels with *i*nvolutive maps

Evgenii Egorov*, Riccardo Valperga*, Efstratios Gavves

University of Amsterdam

arXiv:2406.02490

Problem: how to *learn to sample* with better mixing and ESS?

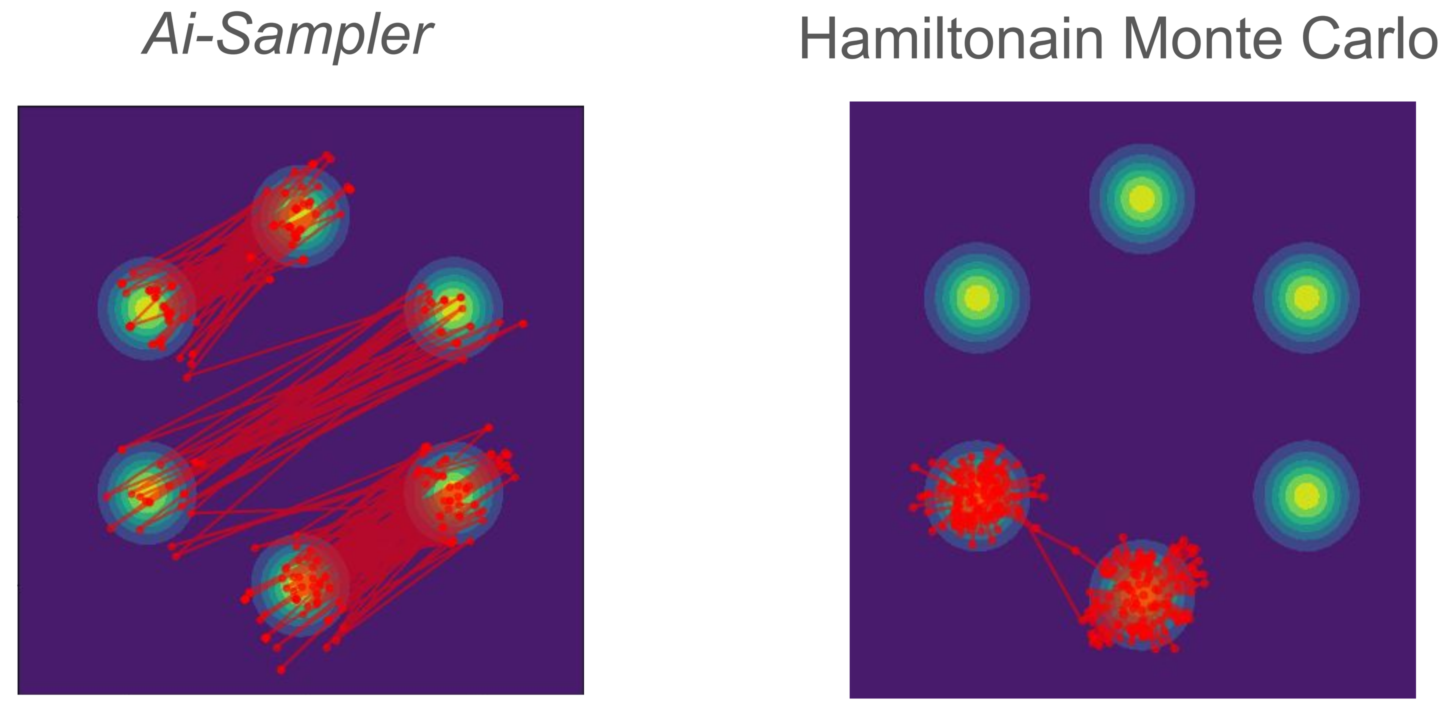

- Markov kernel instead of independent proposal
- Use adversarial loss instead of hand-crafted
  - Derived from detailed-balance condition
- Use time-reversible dynamics as deterministic map

The Markov kernel:

$$t_D(x'|x) = \delta(x' - RL_\theta(x))r\left[D(x)\right] + \delta(x' - x)(1 - r\left[D(x)\right]).$$

$$D^*(x) = $$

$$= \arg\min \int p(x) r\left[D_{\phi,RL_\theta}(x)\right] \log r\left[D_{\phi,RL_\theta}(x)\right] dx$$

$$= \log \frac{p(RL_\theta)}{p(x)} J_x^{RL_\theta(x)}.$$

For such *discriminator* $\quad TV^2[p; t_{D^*} \circ p] = 0$

Our adversarial objective:

$$\max_\theta A_\theta = \max_\theta \mathbb{E}_{p(x)}\left(r\left[D_{\phi,RL_\theta}(x)\right]\right), \text{with fixed } \phi$$

$$\min_\phi \mathbb{E}_{p(x)}\left(r\left[D_{\phi,RL_\theta}(x)\right] \log r\left[D_{\phi,RL_\theta}(x)\right]\right), \text{with fixed } \theta.$$

How to parametrize the discriminator?

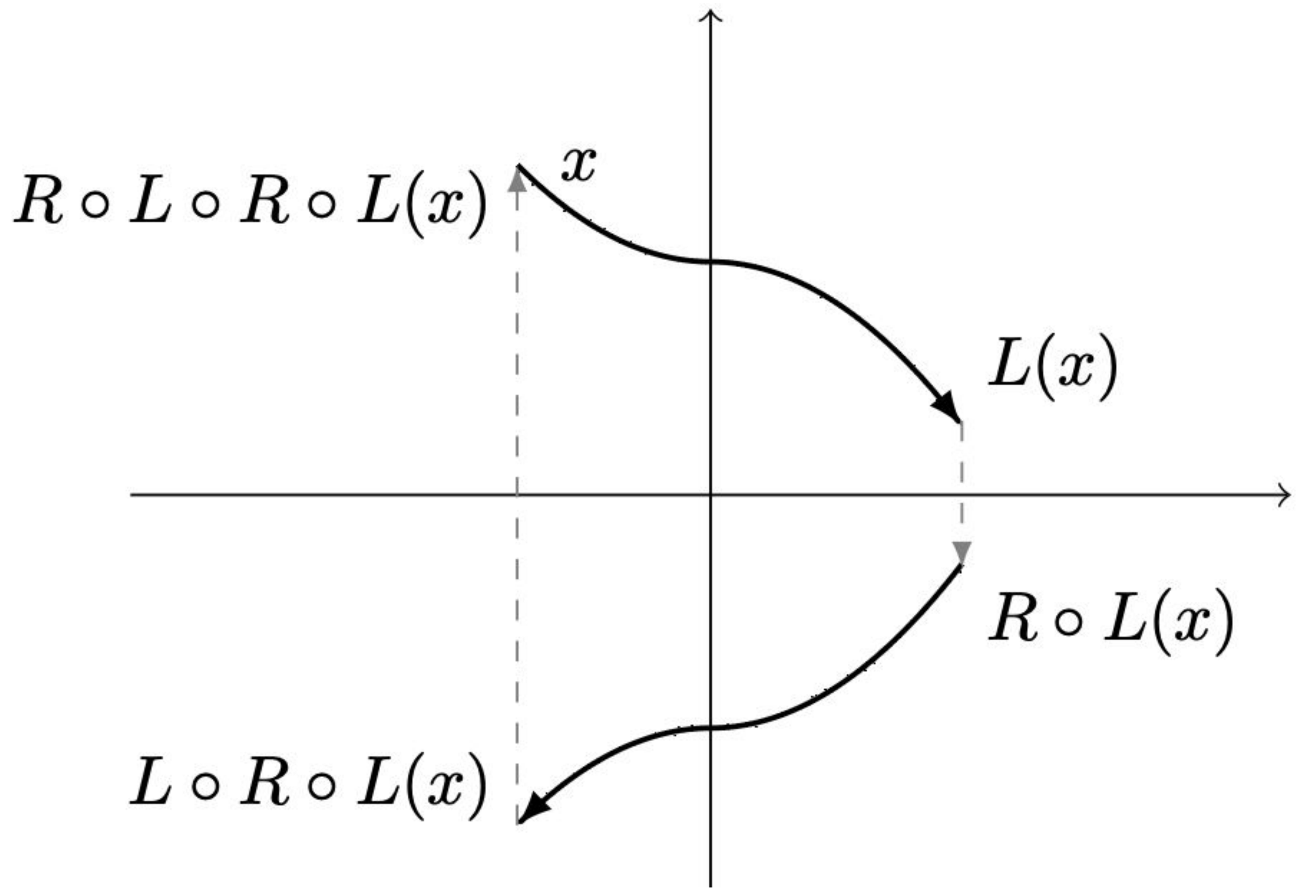

For an R-reversible flow, the density ratio between image and preimage is symmetric:

$$\lambda(x) = \log \frac{p(RLx)}{p(x)} J_x^{RL} = -\log \frac{p(RLRLx)}{p(RLx)} J_{RLx}^{RL} = -\lambda(RLx)$$

Equivariance with respect to:

$$\rho_{2n} : C_2 \to GL(\mathbb{R}^{2n} \oplus \mathbb{R}^{2n}), \ \rho_{2n}(g) = \begin{bmatrix} 0 & I_{2n} \\ I_{2n} & 0 \end{bmatrix}$$

This induces a constraint on linear layers:

$$\begin{bmatrix} A & B \\ C & D \end{bmatrix} \begin{bmatrix} 0 & I_{2n} \\ I_{2n} & 0 \end{bmatrix} \begin{bmatrix} RL(x) \\ x \end{bmatrix} = \begin{bmatrix} 0 & I_{2s} \\ I_{2s} & 0 \end{bmatrix} \begin{bmatrix} A & B \\ C & D \end{bmatrix} \begin{bmatrix} RL(x) \\ x \end{bmatrix}$$

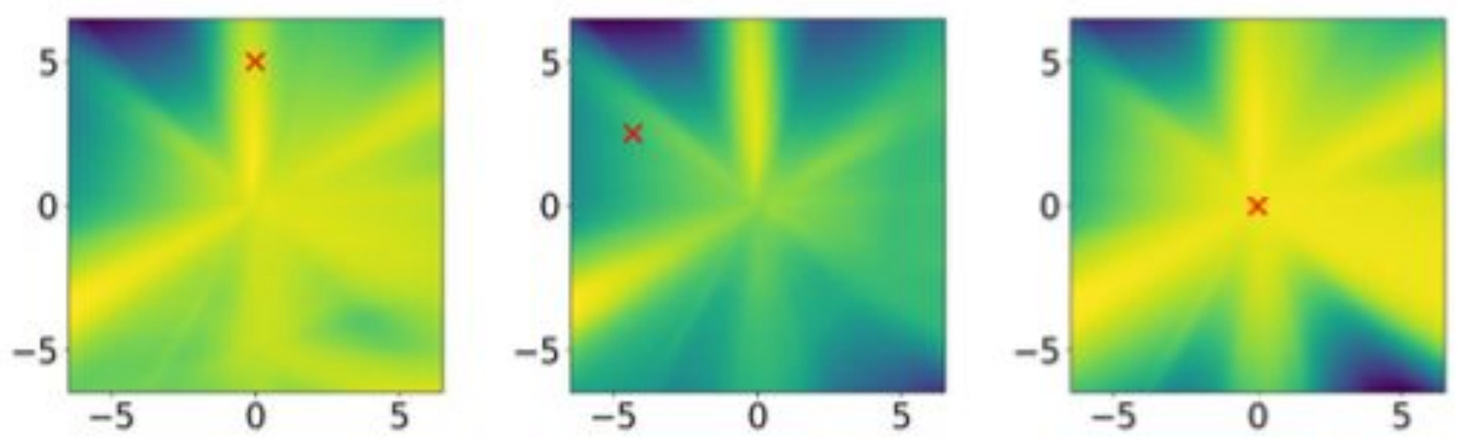

How to parametrize an R-reversible map?

Theorem 4.1. (Valperga et al., 2022) *Let $L : \mathbb{R}^D \to \mathbb{R}^D$ be an R-reversible diffeomorphism[2], with R being a linear involution. Then, there exists a unique diffeomorphism $g : \mathbb{R}^D \to \mathbb{R}^D$, such that $L = R \circ g^{-1} \circ R \circ g$. If L is symplectic, then g can be chosen symplectic.*

Some results:

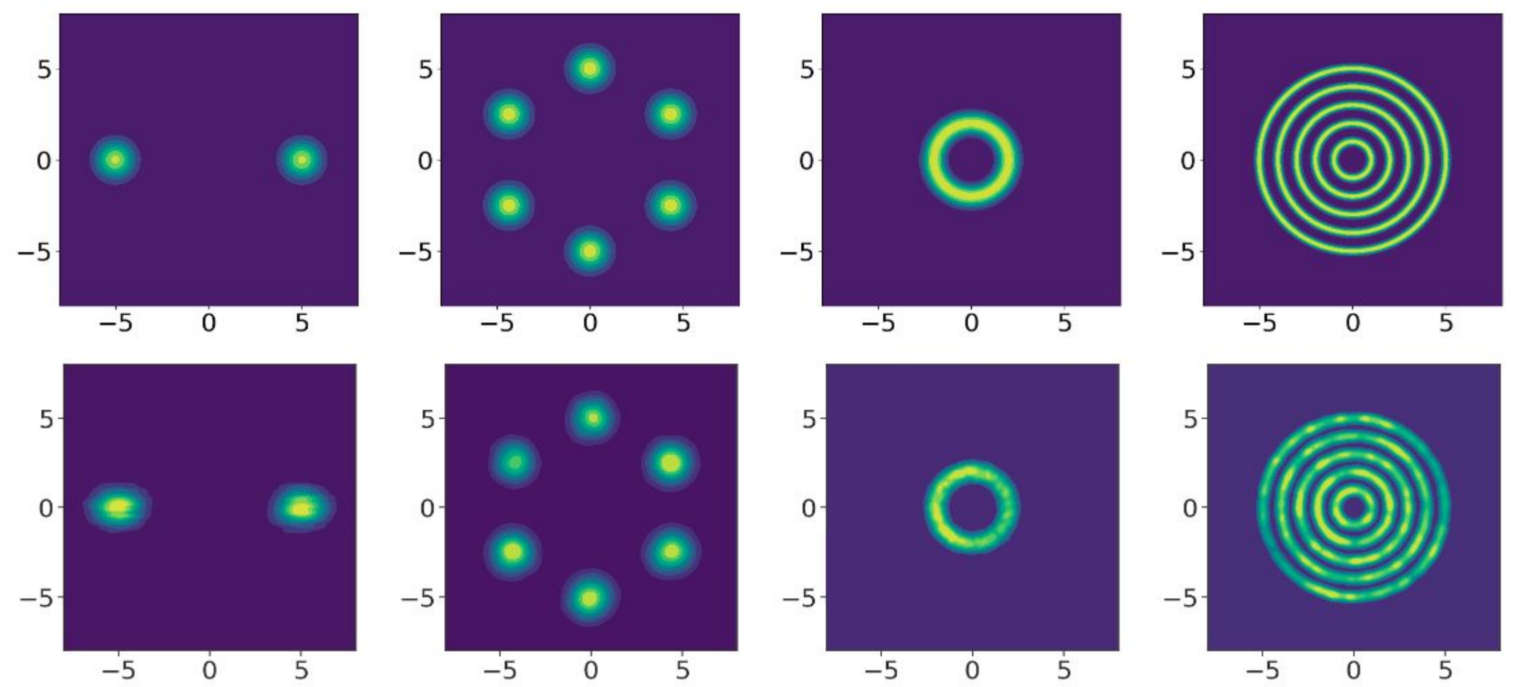

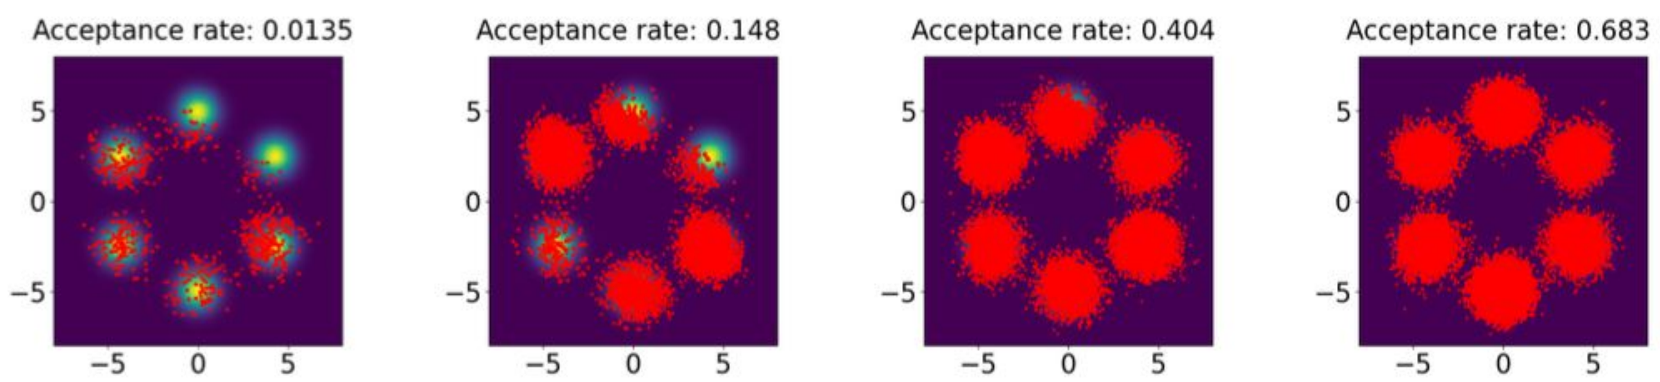

Figure 7. Adversarial objective and acceptance rate during training. Sample quality increasing during training.

Density	ESS		
	HMC	A-NICE-MC	Ai-sampler (ours)
mog2	0.8	355.4	**1000.0**
mog6	2.4	320.0	**1000.0**
ring	981.3	**1000.0**	378.0
ring5	256.6	155.57	**396.5**

Density	ESS/s	
	HMC	Ai-sampler (ours)
mog2	0.4	**1052.6**
mog6	0.98	**1041.7**
ring	**2725.8**	402.1
ring5	333.2	**434.7**