Contractive Diffusion Policies

Robust Action Diffusion via Contractive Sampling with Differential Equations

Amin Abyaneh et al.

McGill ICLR Mila

The Intuition Behind Diffusion Modeling

There is a long way from MuJoCo baselines to real-world learning!

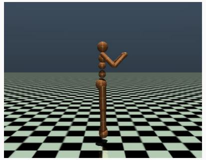


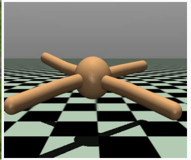

Reliable learning in the real-world is much more complex!

www.unitree.com/G1
www.anybotics.com/robotics/anymal/

The Intuition Behind Diffusion Modeling

Distribution learning by turning noise samples into meaningful samples.

Robotics. Learning conditional action distributions, aka **policies**: $\pi_\theta(\mathbf{a} \mid \mathbf{s})$

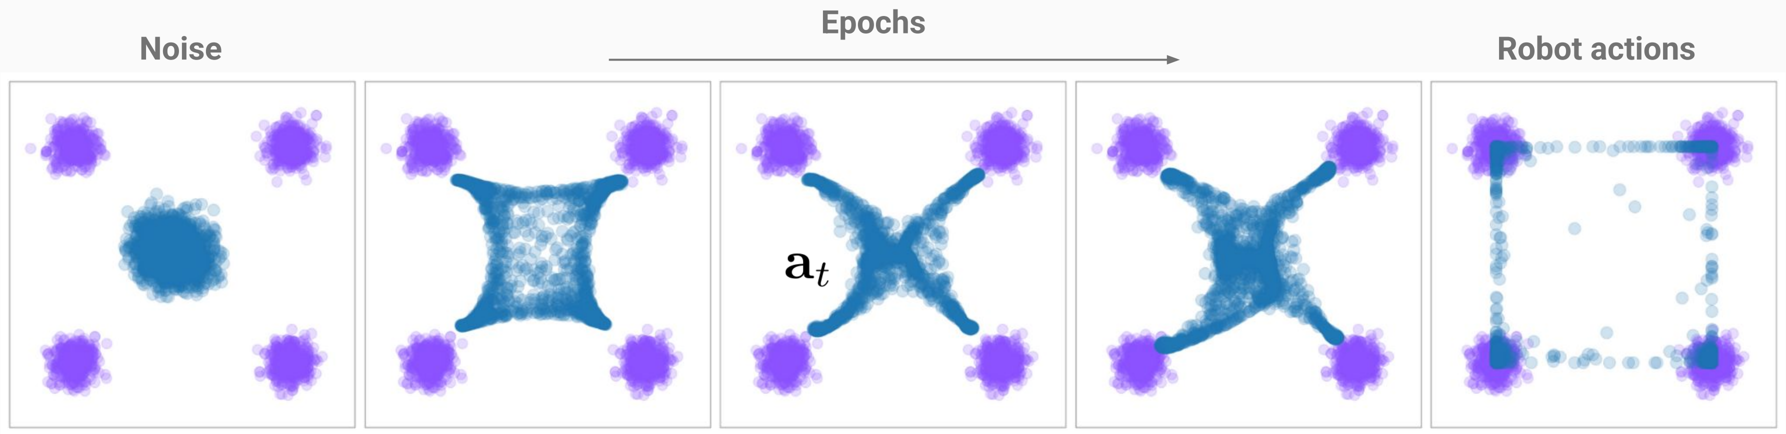

Example of 2D diffusion sampling trained for a fixed condition.

What can go wrong?

Diffusion sampling suffers from **score-matching** and **integration** errors.

Image generation. Slight disturbances, often not critical, at times even beautiful?

Continuous control. Out-of-distribution and critical **failure**.

What can go wrong?

Diffusion sampling suffers from **score-matching** and **integration** errors.

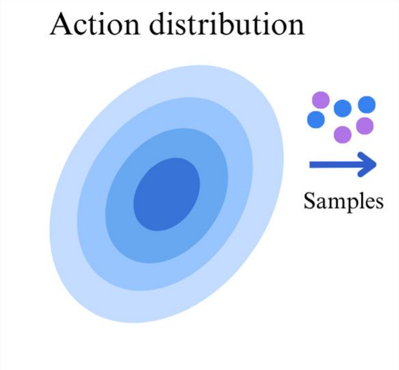

What can go wrong?

Diffusion sampling suffers from **score-matching** and **integration** errors.

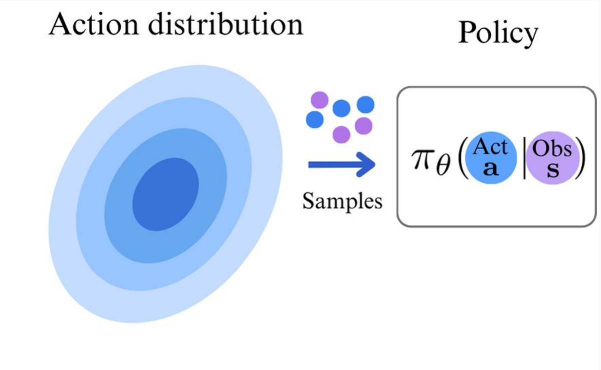

What can go wrong?

Diffusion sampling suffers from **score-matching** and **integration** errors.

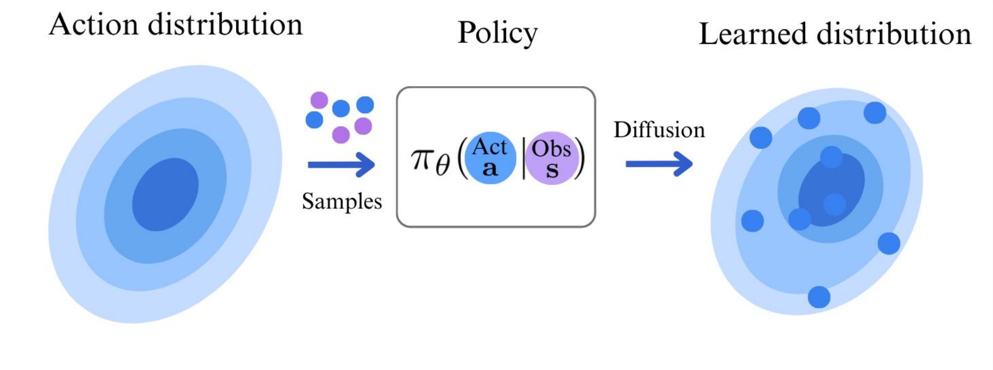

What can go wrong?

Diffusion sampling suffers from **score-matching** and **integration** errors.

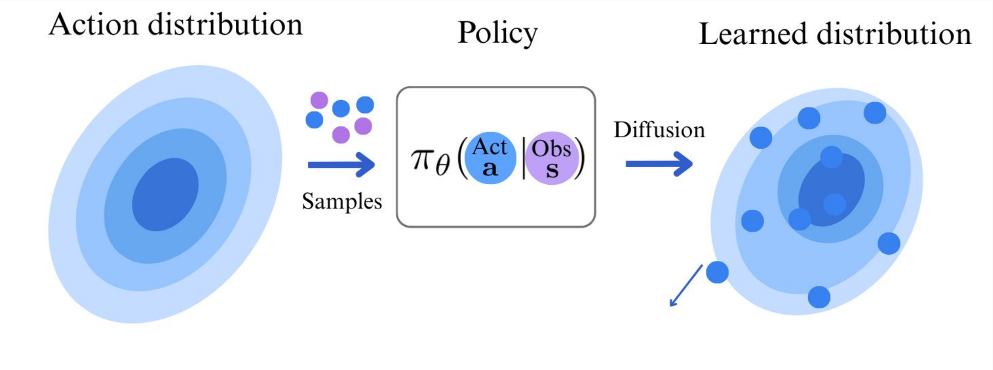

What can go wrong?

Diffusion sampling suffers from **score-matching** and **integration** errors.

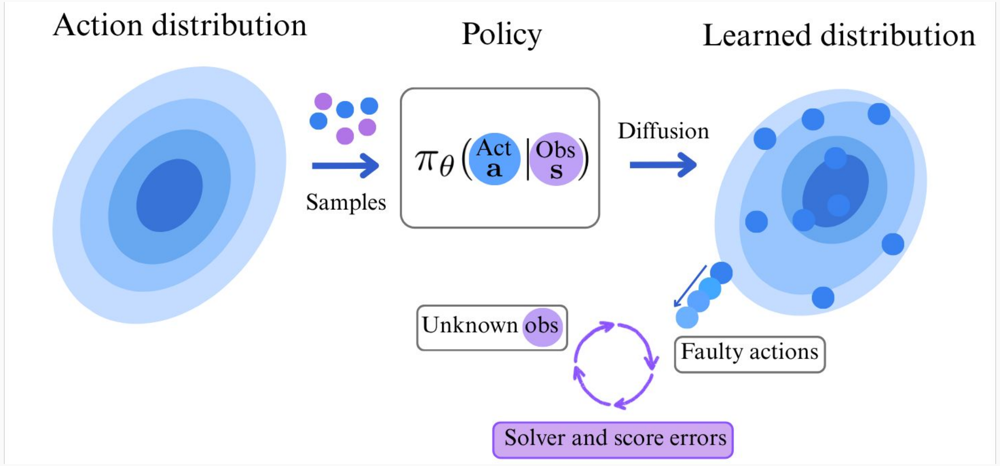

Diffusion Modeling for Action Generation

The process consists of forward and reverse differential equations.

The forward process adds noise to data for a fixed number of steps.

The reverse process uses a learned **score function** to denoise the distorted samples back to actions.

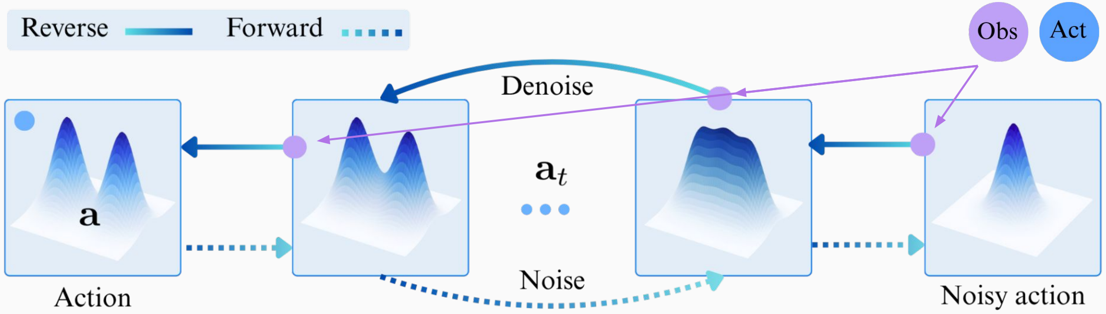

$$d\mathbf{a}_t = f(\mathbf{a}_t, t)dt + g(t)d\mathbf{w}_t, \quad \mathbf{a}_0 \sim p(\mathbf{a})$$

$$d\mathbf{a}_t = \left[f(\mathbf{a}_t, t) - \frac{1}{2}g(t)^2 \nabla_{\mathbf{a}_t} \log p_t(\mathbf{a}_t \mid \mathbf{s})\right] dt$$

Diffusion Modeling for Action Generation

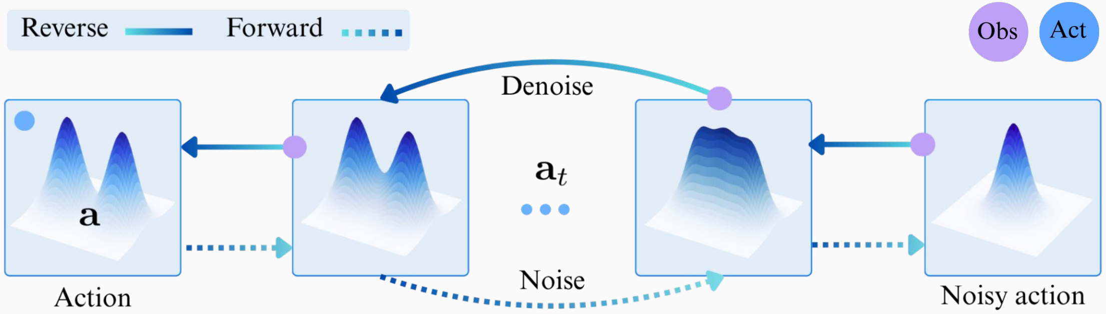

The loss enables learning a score function from data.

$$\mathcal{L}_{\mathrm{d}}(\theta) := \mathbb{E}_{t\sim\mathcal{U}(0,1),\, \mathbf{a}_t\sim p_t(\cdot|\mathbf{s}_t)} \left\|\epsilon_\theta(\mathbf{a}_t, \mathbf{s}, t) + \sigma_t \nabla_{\mathbf{a}_t} \log p_t(\mathbf{a}_t \mid \mathbf{s})\right\|_2^2$$

The **goal** is to deploy the learned score function to generate actions.

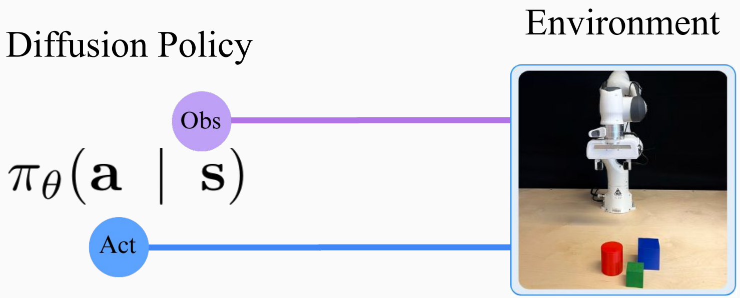

And what is contraction?

Contraction theory provides sufficient conditions to pull nearby ODE flows closer.

⟶ Diffusion sampling can be modeled by an ODE.

Contraction can be achieved by **constraints on the Jacobian of the ODE function**.

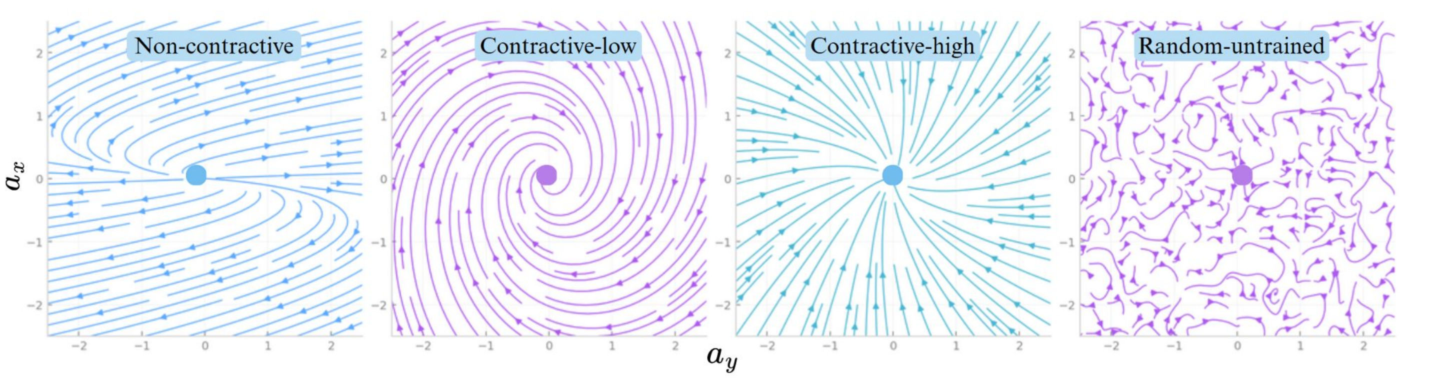

Contraction in Distribution Modeling

Contractive Autoencoders [ICML 2011]

Localized space contraction yields robust features .

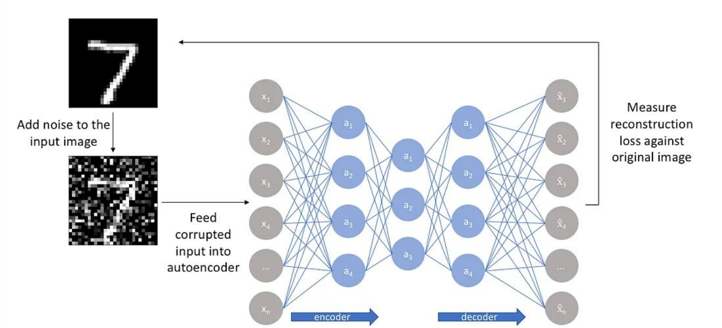

Contractive Diffusion Probabilistic Models [2024]

SDE-style contraction: high compute, and marginal benefits.

Contraction in Diffusion Process

Contraction mitigates solver and score-matching errors in diffusion sampling, while reducing unwanted action variance.

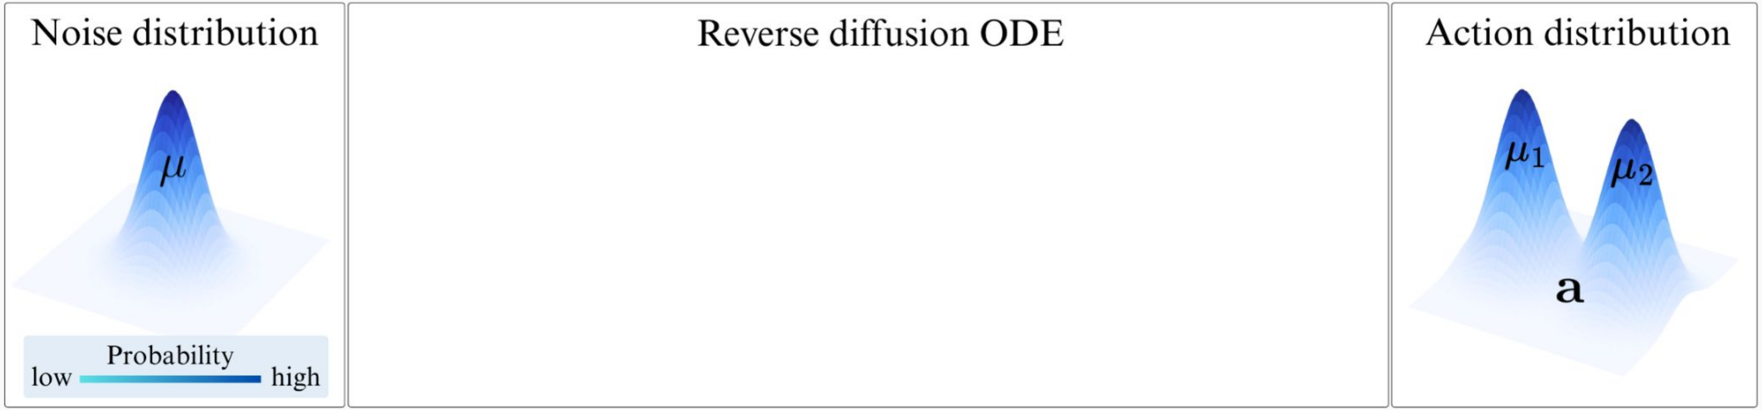

Contraction in Diffusion Sampling

Contraction mitigates solver and score-matching errors in diffusion sampling, while reducing unwanted action variance.

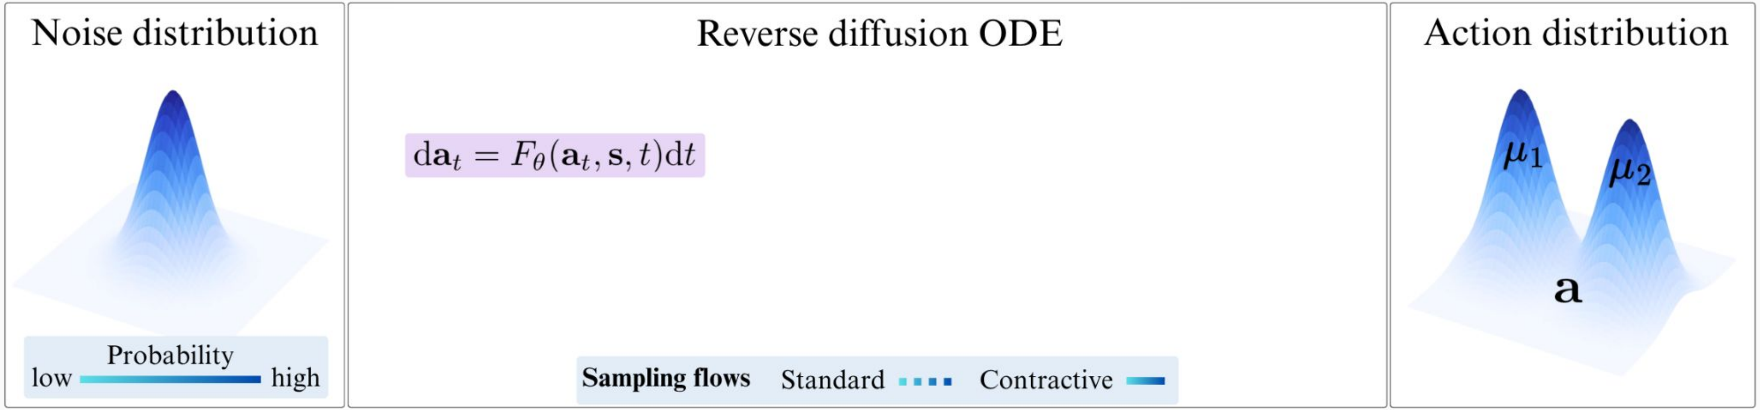

Contraction in Diffusion Sampling

Contraction mitigates solver and score-matching errors in diffusion sampling, while reducing unwanted action variance.

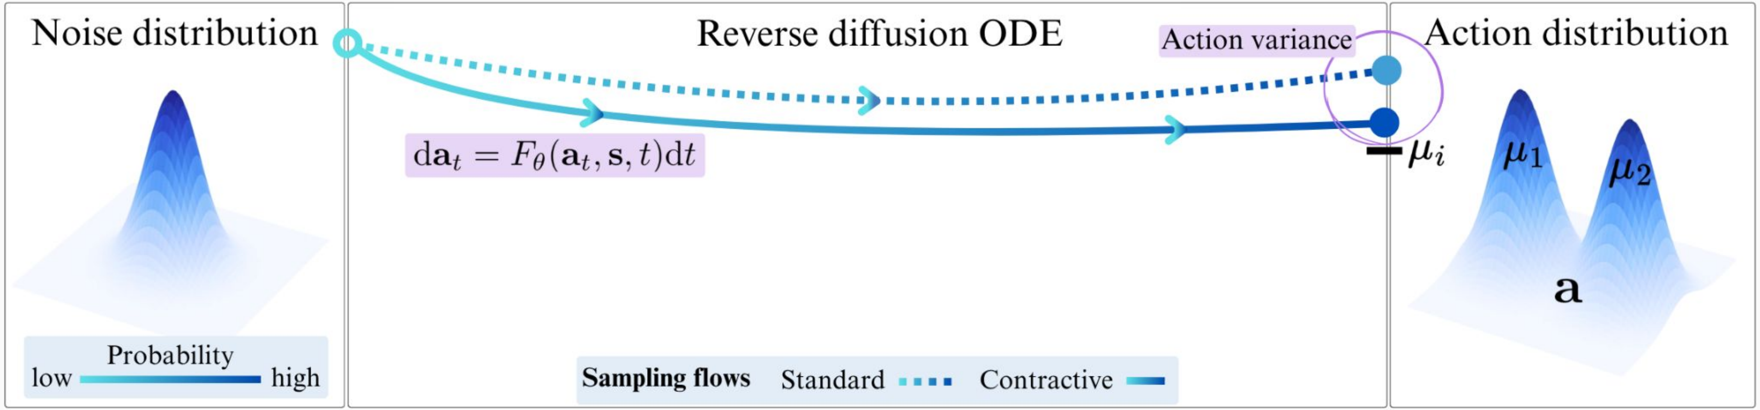

Contraction in Diffusion Sampling

Contraction mitigates solver and score-matching errors in diffusion sampling, while reducing unwanted action variance.

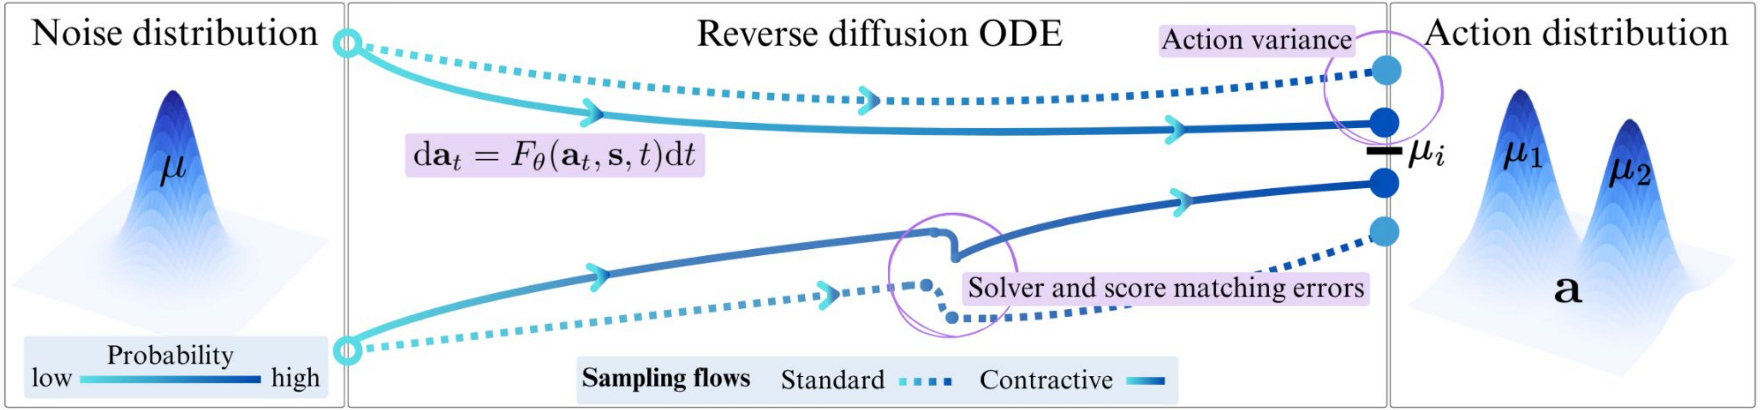

Efficient Control of the Contraction Rate

We draw the connection between the eigenvalues of the **score Jacobian**, and contraction of the reverse process.

Too much contraction → Mode collapse

Too little contraction → No effect!

Theorem 3.1 (Interplay of score Jacobian and contractive sampling). *Given a state* $\mathbf{s} \in \mathcal{S}$ *and a diffusion ODE,* $\mathrm{d}\mathbf{a}_t = F_\theta(\mathbf{a}_t, \mathbf{s}, t)\mathrm{d}t$, F_θ *is contractive, i.e.,* $J_{F_\theta}^{sym} \prec 0$, *iff the score Jacobian satisfies*

$$\lambda_{\max}(J_{\epsilon_\theta}^{sym}) < -f(t)h(t)^{-1}, \quad \forall t \in [0, 1],$$

where $\lambda_{\max}$ *denotes the largest eigenvalue, and* f, h *are defined by the forward process.*

Jacobian of the reverse process

Jacobian of the score function

$$\mathrm{d}\mathbf{a}_t = \left[f(\mathbf{a}_t, t) - \frac{1}{2} g(t)^2 \nabla_{\mathbf{a}_t} \log p_t(\mathbf{a}_t \mid \mathbf{s}) \right] \mathrm{d}t$$

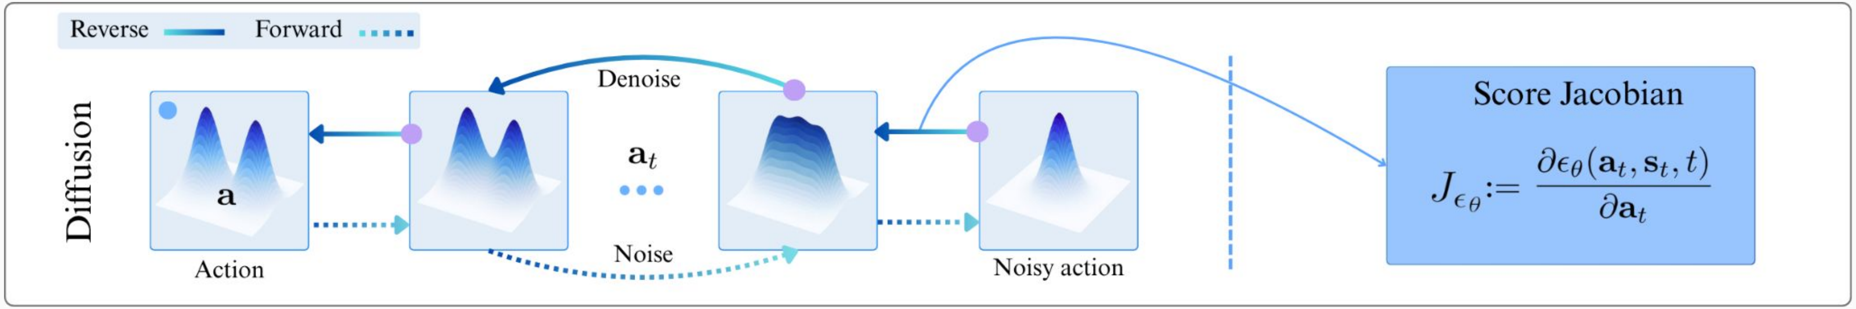

Learning and Deployment Pipeline

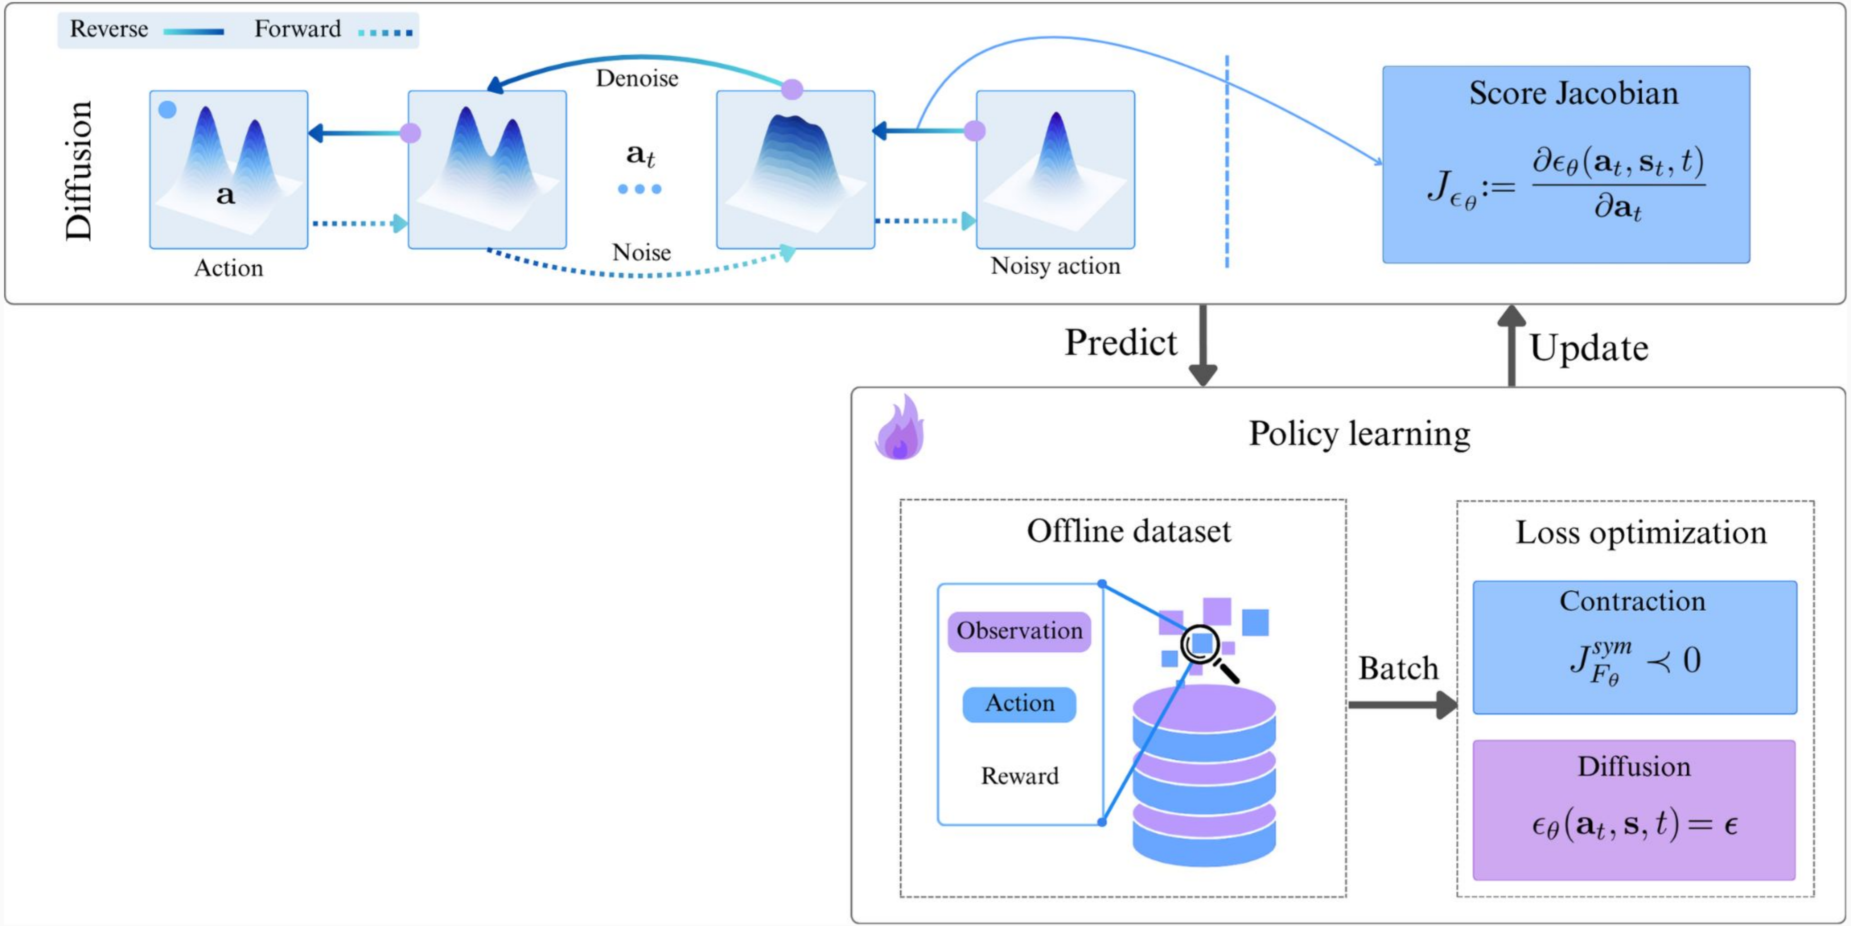

Learning and Deployment Pipeline

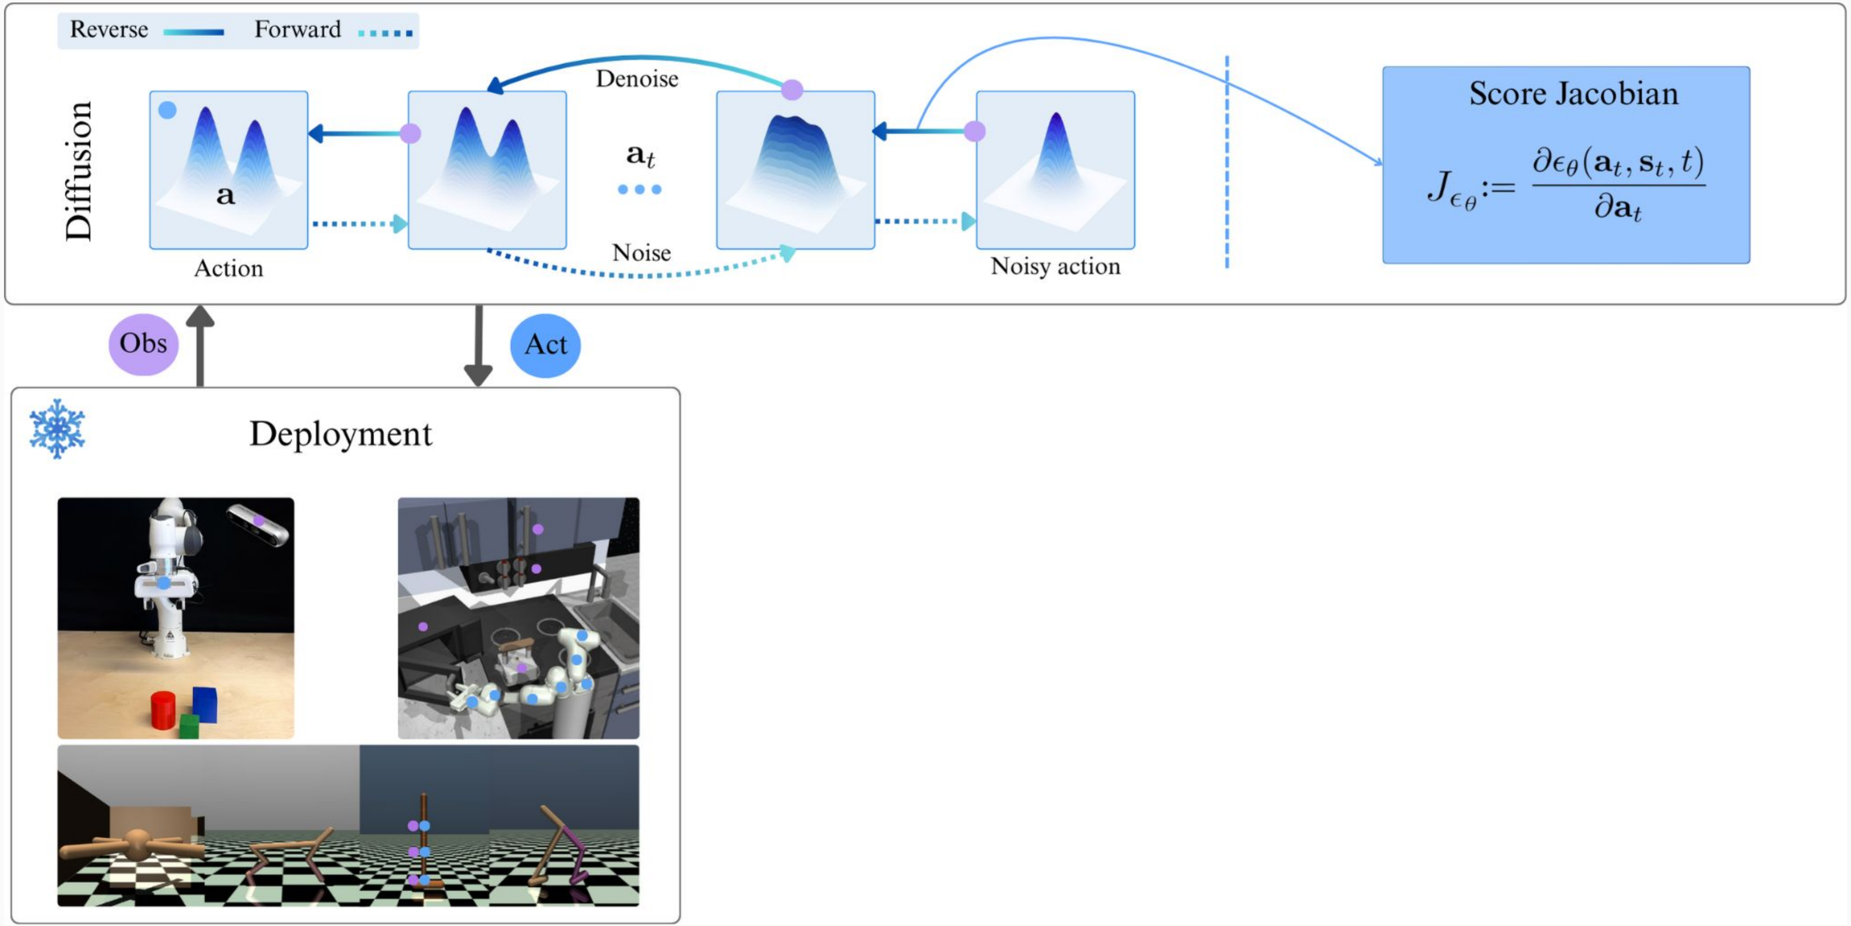

Learning and Deployment Pipeline

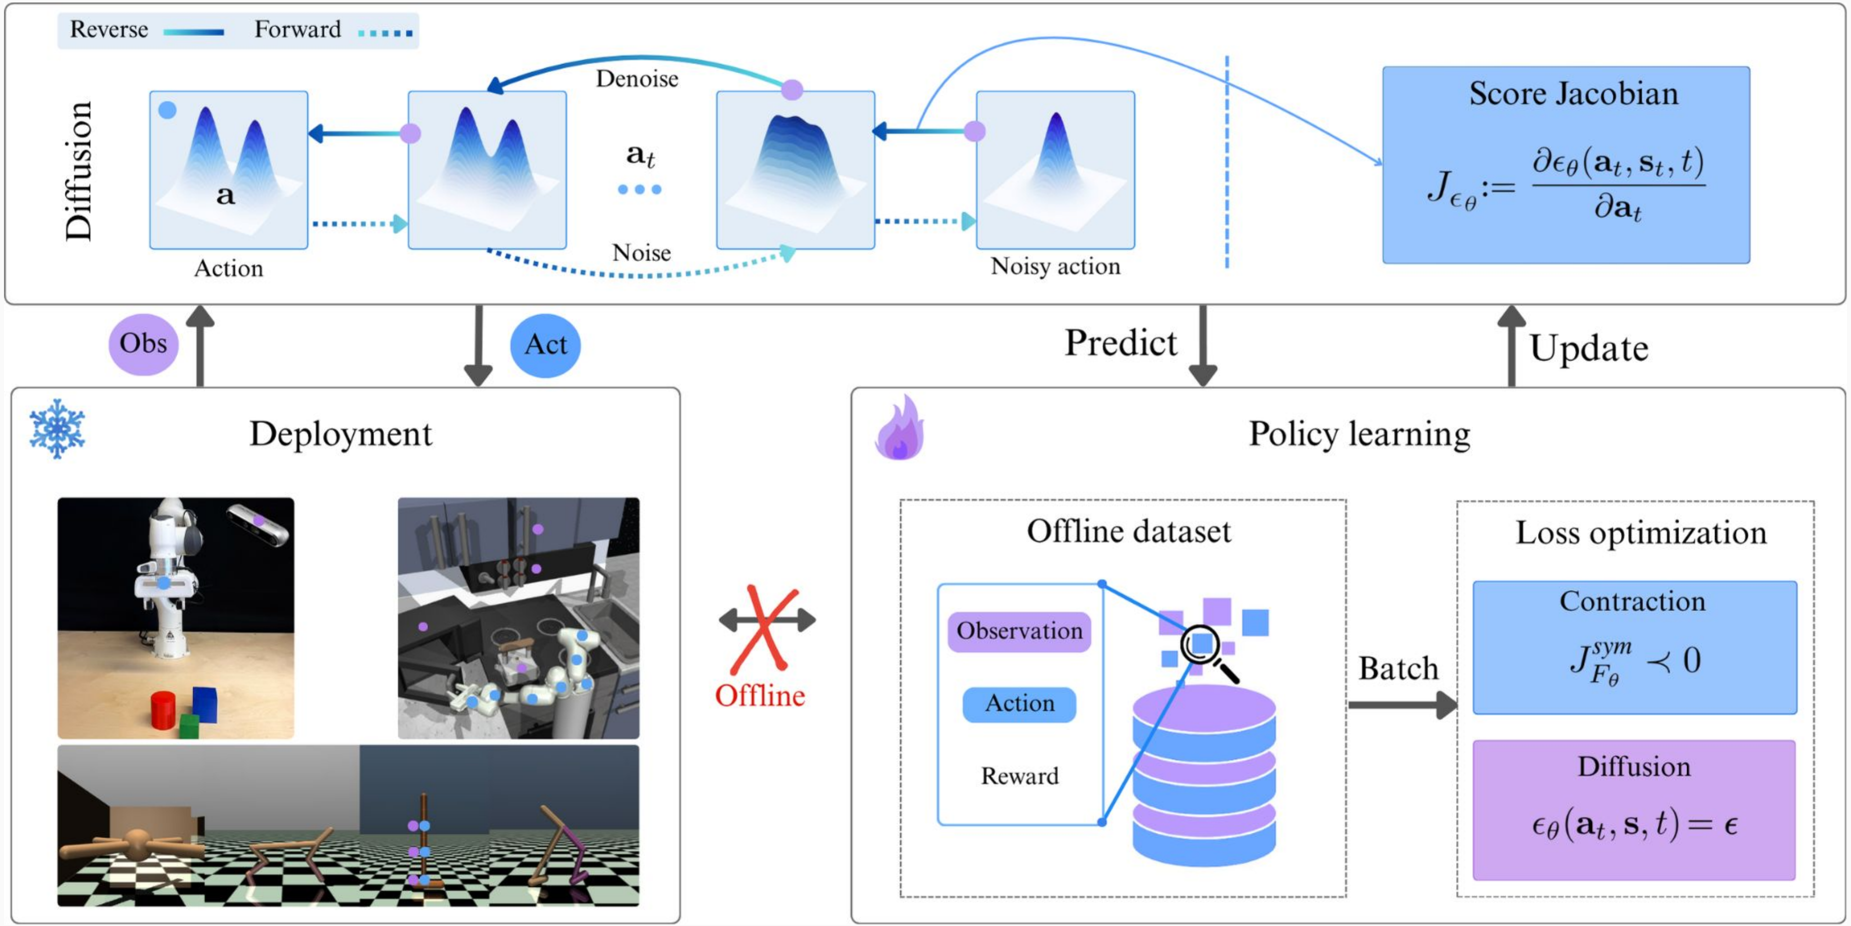

Learning and Deployment Pipeline

Contraction for a Toy Example

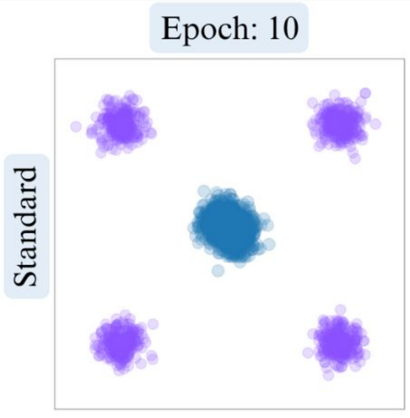

Contraction for a Toy Example

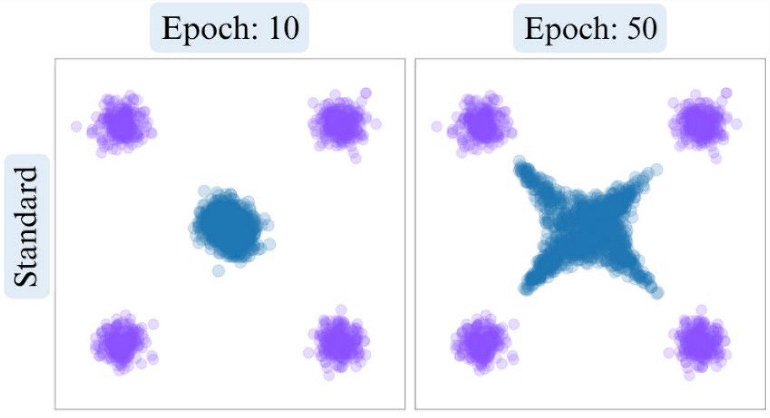

Contraction for a Toy Example

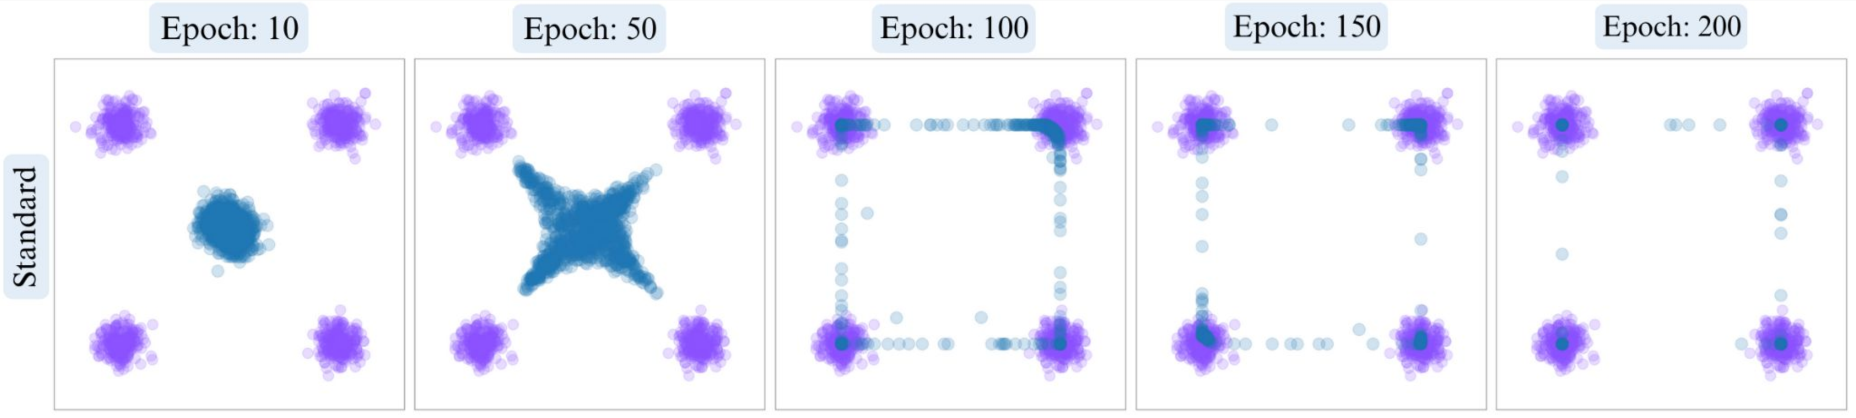

Contraction for a Toy Example

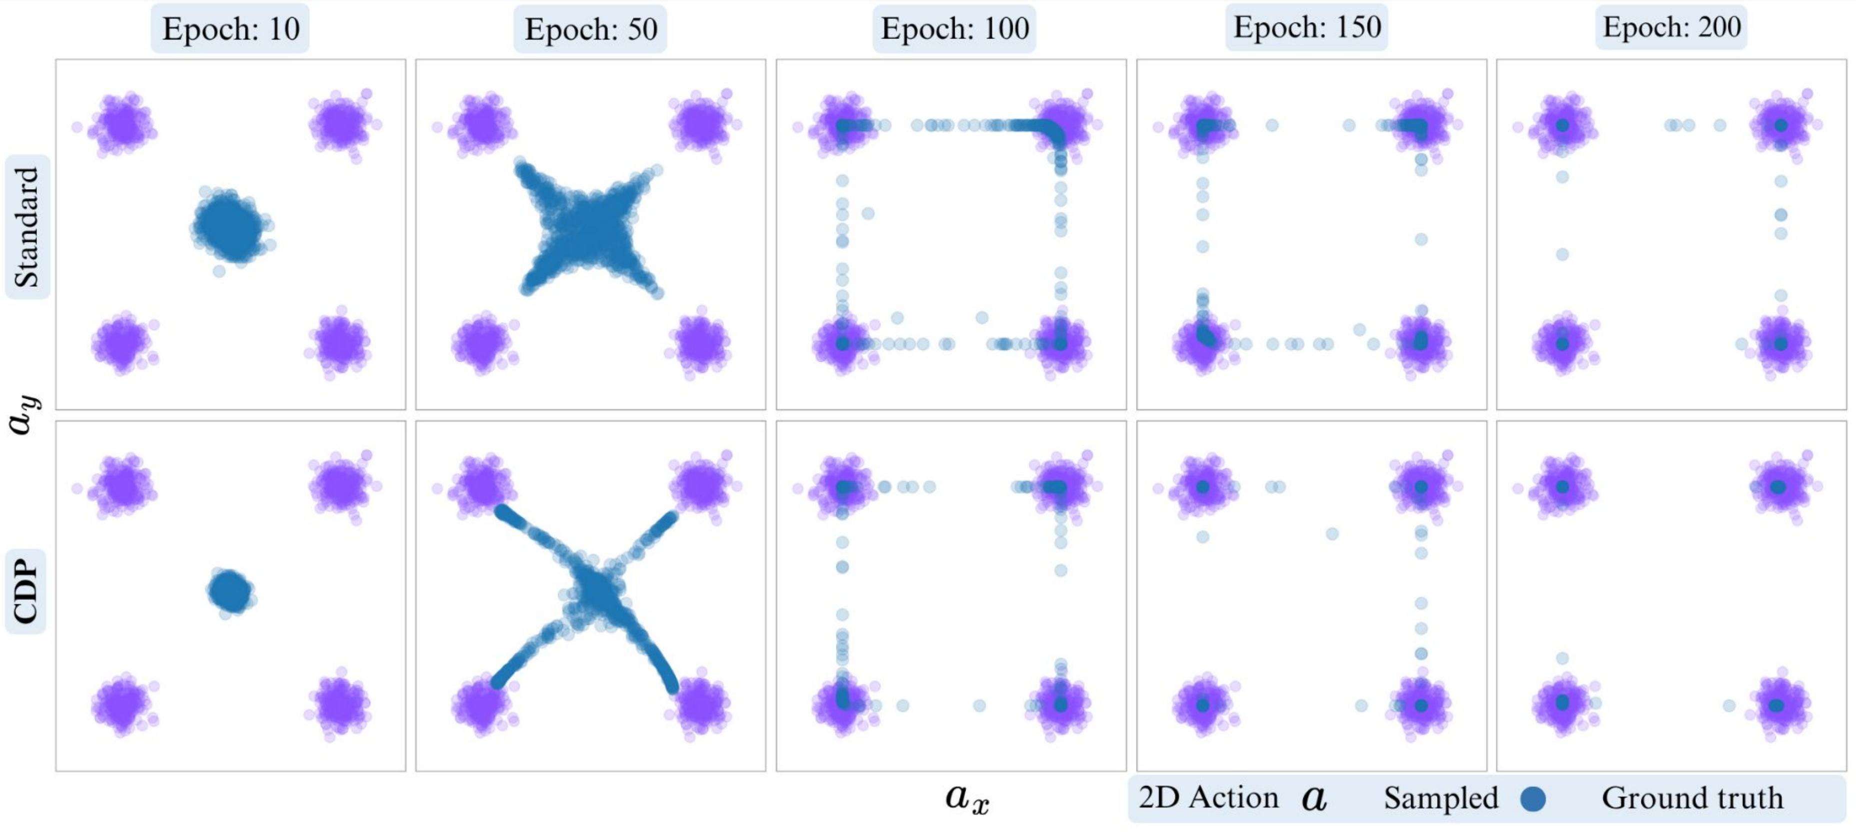

Contraction for a Toy Example

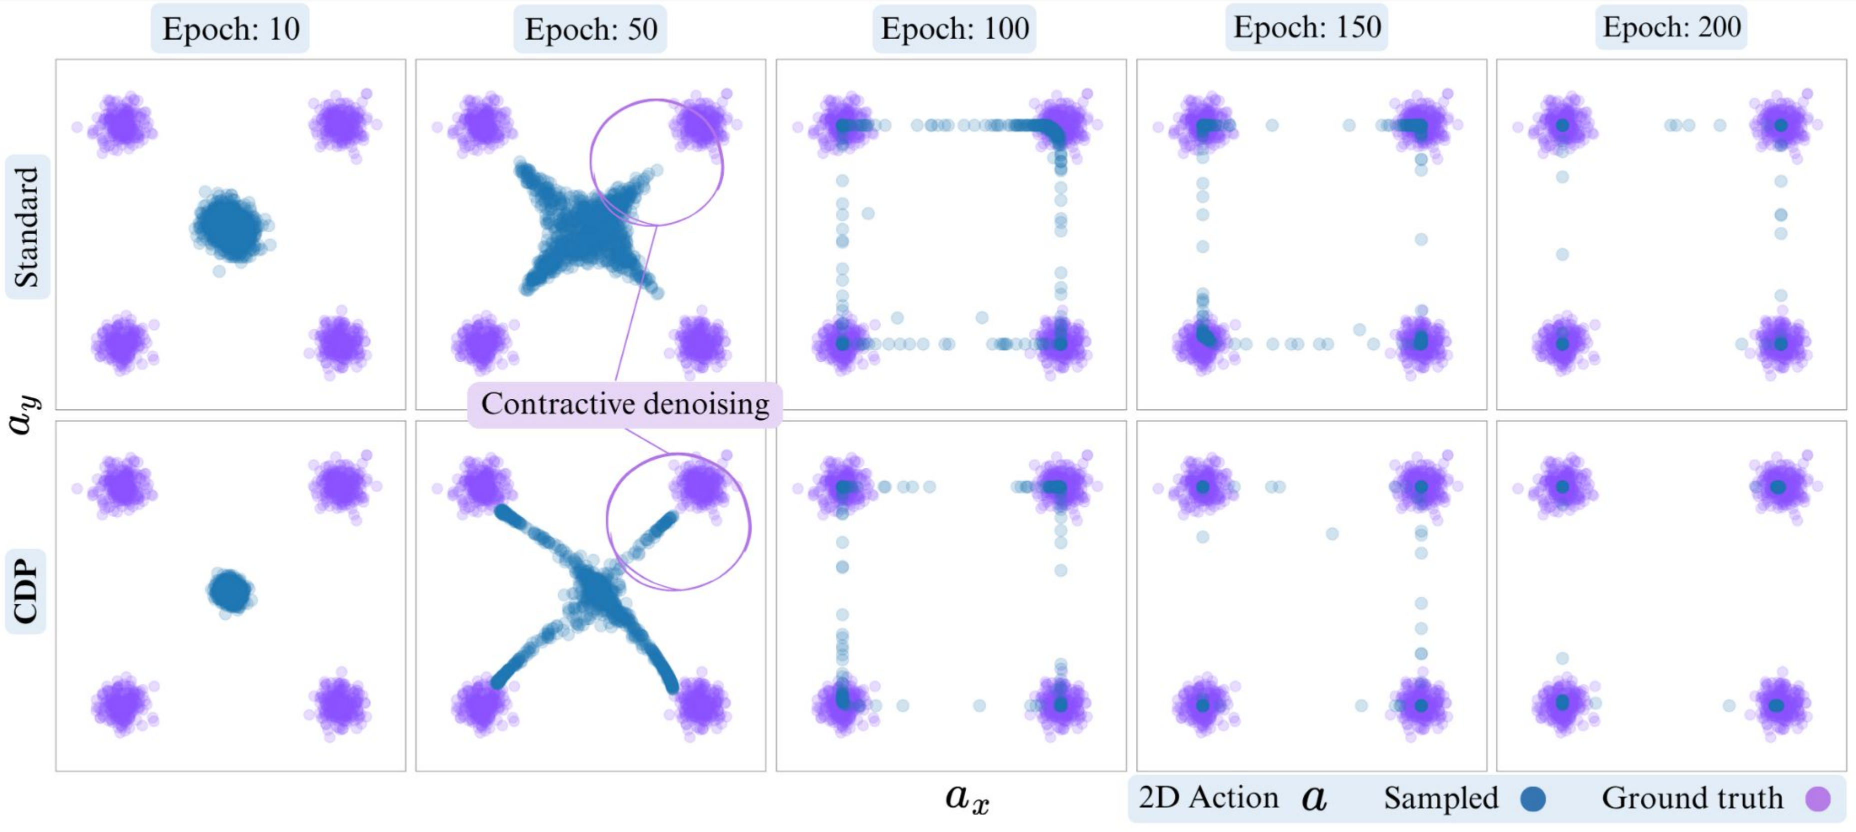

Contraction for a Toy Example

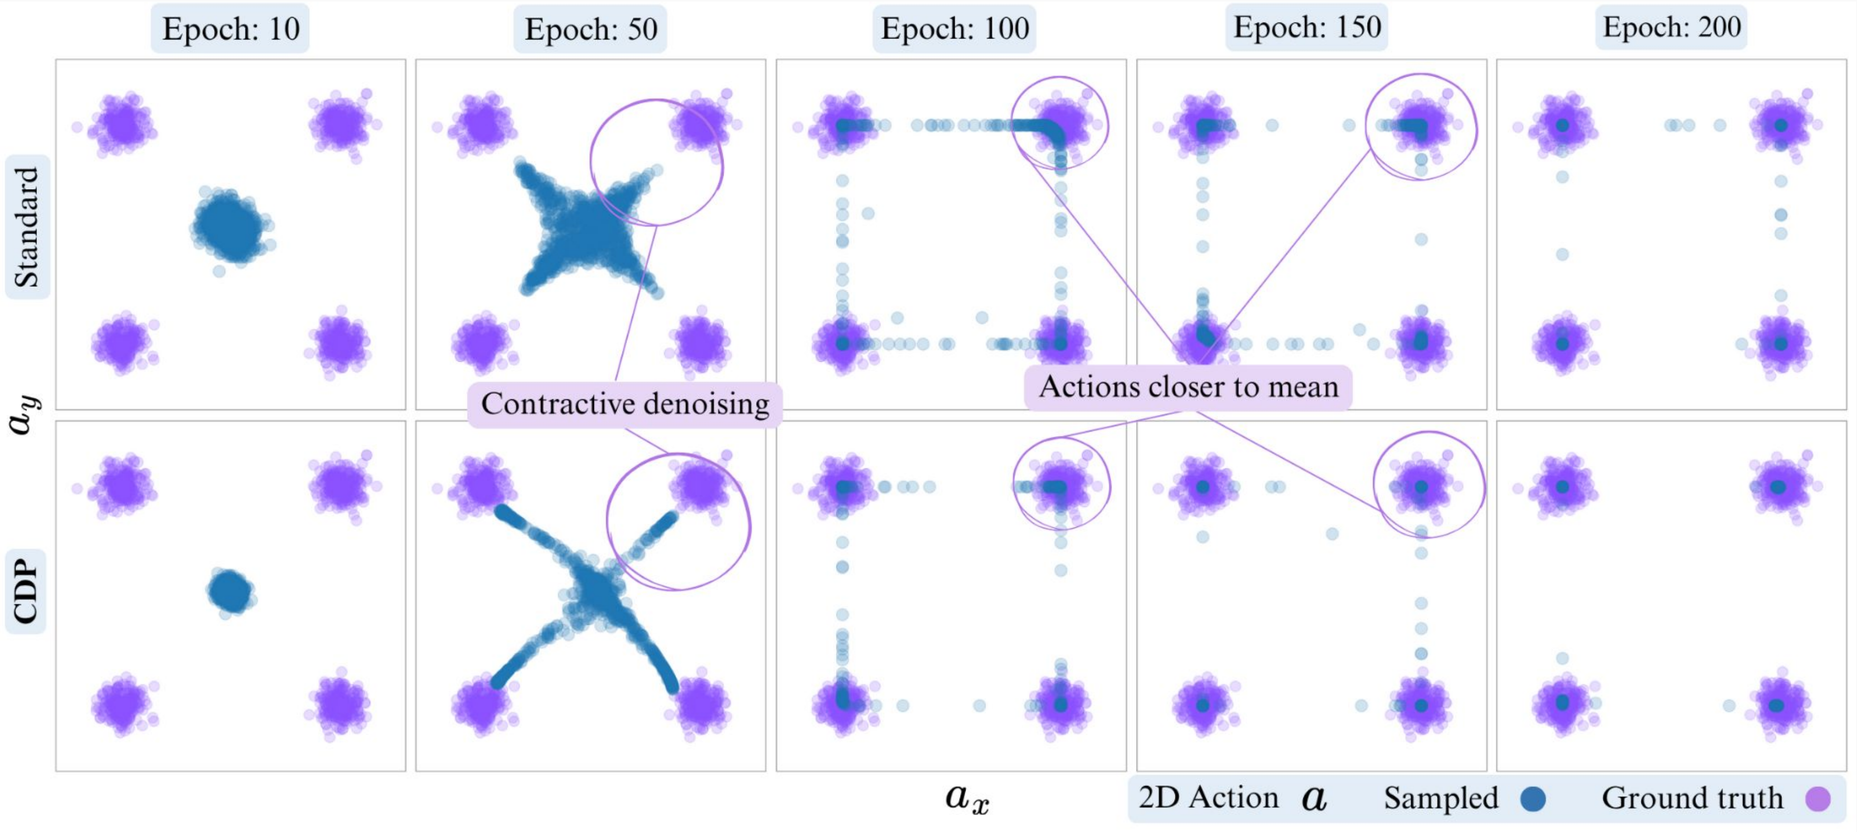

Contraction for a Toy Example

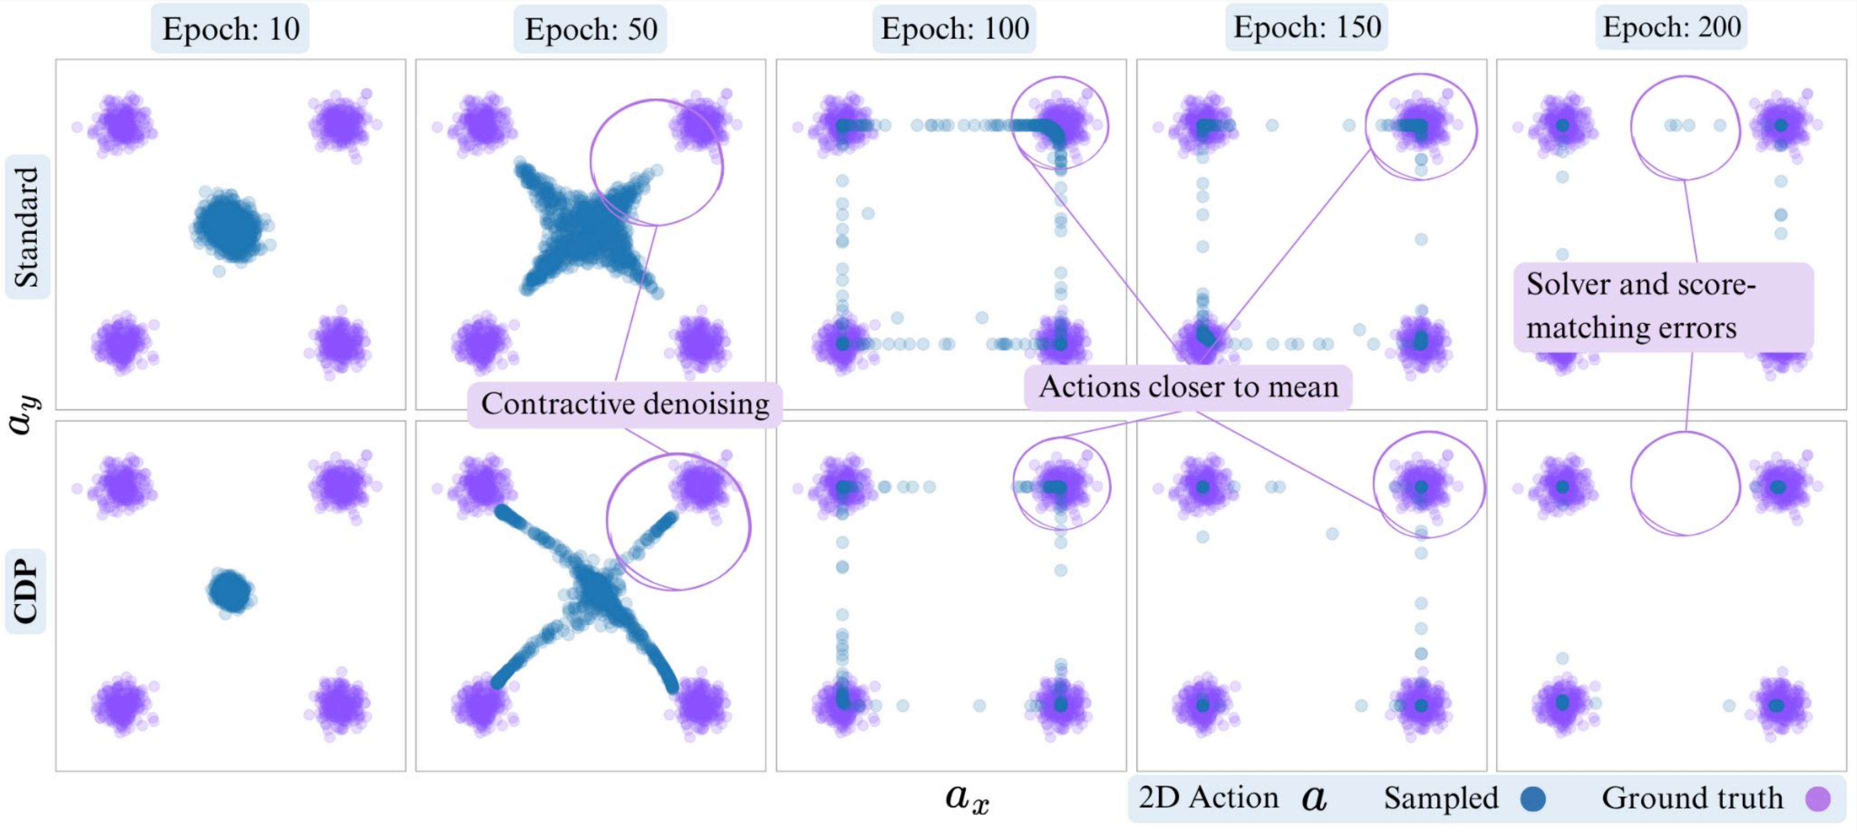

Results Overview

1. **General performance improvement in RL**

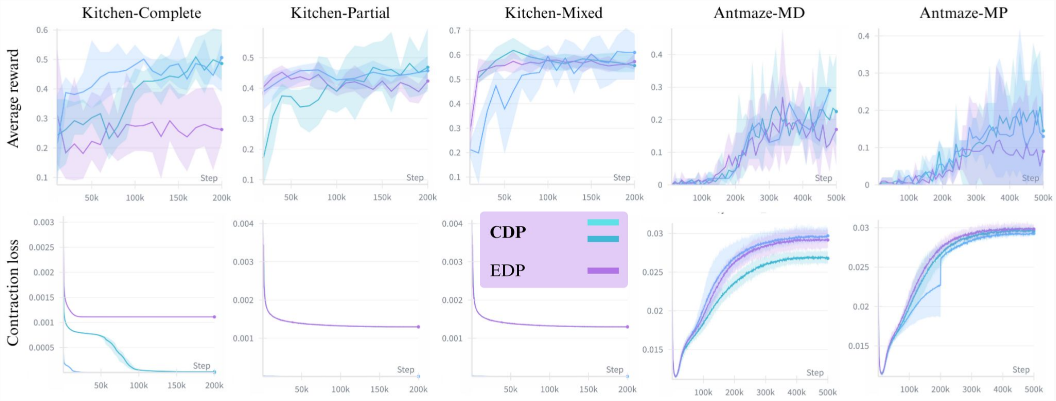

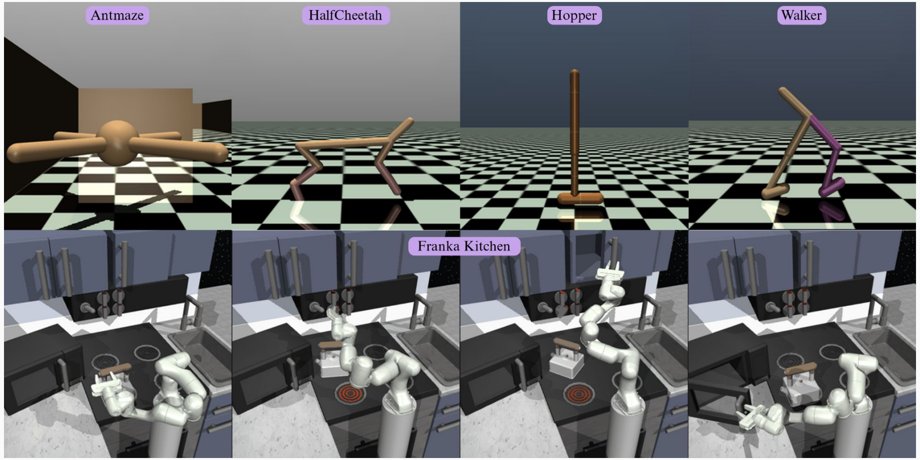

Results Overview

1. General performance improvement in RL

2. **Relative performance improvement in IL**

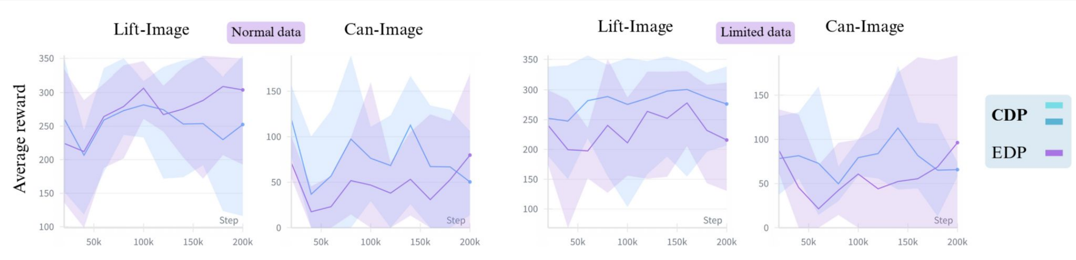

Remarkable improvement in high-dimensional environments!

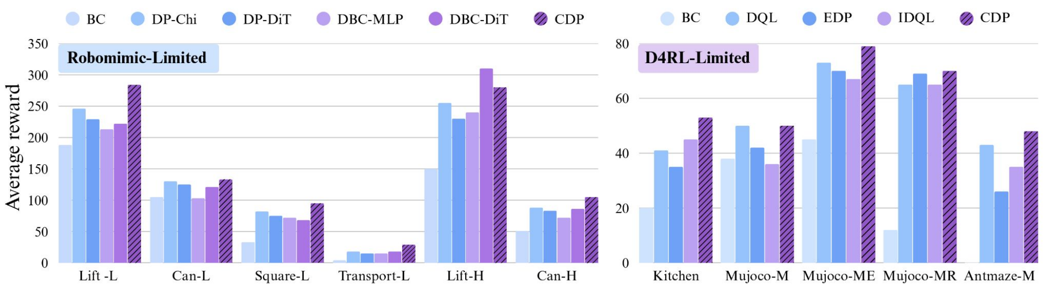

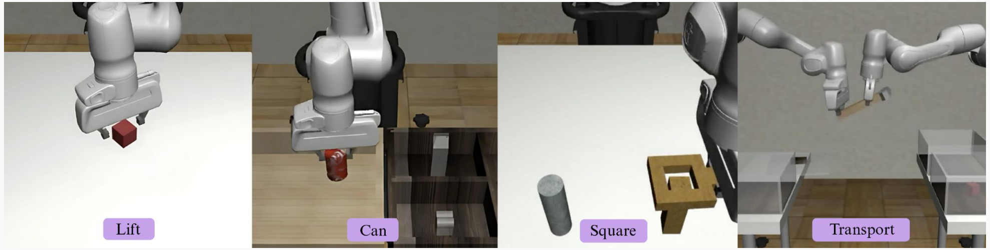

Results Overview

1. General performance improvement in RL

2. Relative performance improvement in IL

3. **Real-world deployment**

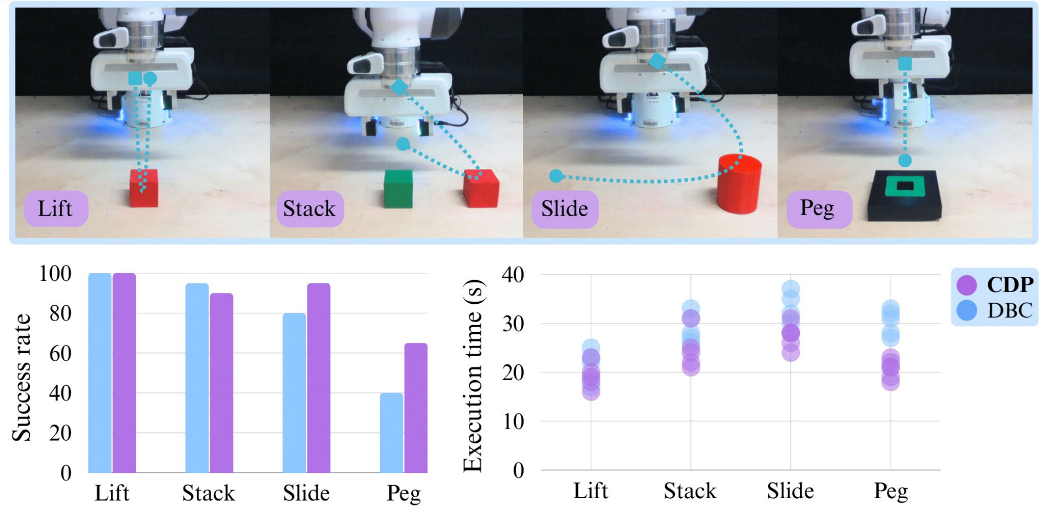

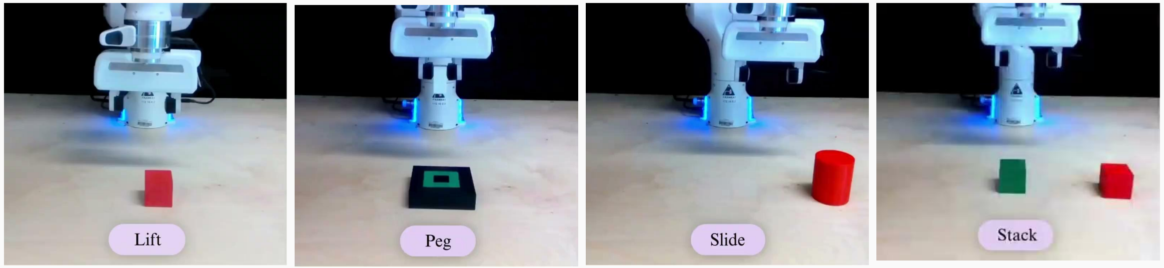

Future Directions

Contraction in diffusion or flow-matching heads of VLAs (Reviewer #2).

Contraction in higher dimensional action and observation spaces.

Integration into consistency policies .

Final point

There is more to embracing foundation models for robot learning than scaling laws.

Thanks!

amin.abyaneh@mail.mcgill.ca

contractive-diffusion.github.io

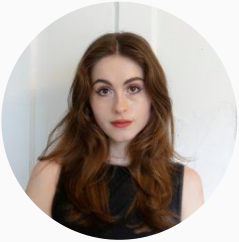

Charlotte Morissette

Mohamad Danesh

Anas Houssaini

David Meger

Gregory Dudek

Hsiu-Chin Lin

Special thanks to my collaborators.