Learning Globally Contracting Dynamics from Demonstrations

Alexander (Sasha) Davydov

Department of Mechanical Engineering,
Center for Control,
Dynamical Systems & Computation,
University of California, Santa Barbara
davydovalexander.github.io

Workshop on "Contraction Theory for Systems, Control, Optimization, and Learning"
IEEE CDC, Milan, Italy,
December 15, 2024

Motivation – Learning motion plans from demonstrations

- Collection of expert demonstrations – $\mathcal{D} = \{(x_i, \dot{x}_i)\}_{i=1}^{N}$
- Learn dynamics $\dot{x} = F(x)$
- Use low-level controller to track

Challenges:

- Ensure trajectories converge
- Robustness to initial condition

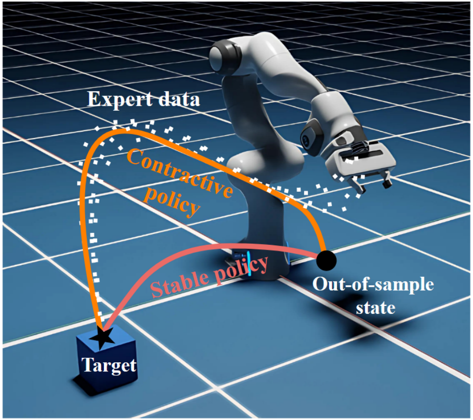

Motivation – Learning motion plans from demonstrations

- Collection of expert demonstrations – $\mathcal{D} = \{(x_i, \dot{x}_i)\}_{i=1}^{N}$
- Learn dynamics $\dot{x} = F(x)$
- Use low-level controller to track

Challenges:

- Ensure trajectories converge
- Robustness to initial condition

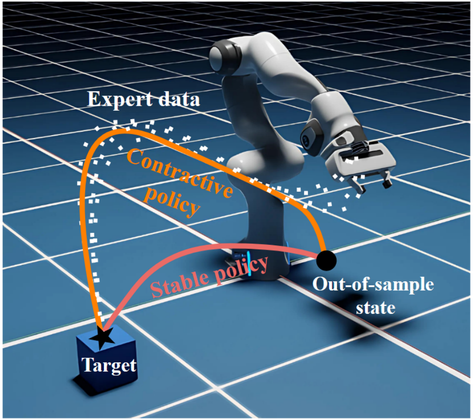

Enforce contraction globally!

Anonymous. Contractive dynamical imitation policies for efficient out-of-sample recovery. *International Conference on Learning Representations (ICLR)*, 2025. URL **https://openreview.net/forum?id=lILEtkWOXD**. Under review

Parametrization of contracting dynamics. Consider

$$\dot{x} = \mathsf{F}(x) = A(x, x^\star)(x - x^\star), \qquad \textbf{(ELCD)}$$

Learning contracting dynamics

Parametrization of contracting dynamics. Consider

$$\dot{x} = \mathsf{F}(x) = A(x, x^\star)(x - x^\star), \qquad \textbf{(ELCD)}$$

- **parametrization**:

$$A(x, x^\star) = -P_s(x, x^\star)^\top P_s(x, x^\star) + P_a(x, x^\star) - P_a(x, x^\star)^\top - cI_d$$

- P_s, P_a neural networks, $x^\star$ learnable equilibrium

Learning contracting dynamics

Parametrization of contracting dynamics. Consider

$$\dot{x} = \mathsf{F}(x) = A(x, x^\star)(x - x^\star), \qquad \textbf{(ELCD)}$$

- **parametrization**:

$$A(x, x^\star) = -P_s(x, x^\star)^\top P_s(x, x^\star) + P_a(x, x^\star) - P_a(x, x^\star)^\top - cI_d$$

- P_s, P_a neural networks, $x^\star$ learnable equilibrium

Then **exp. convergence in ℓ_2 and contraction wrt Riemannian metric**:

$$\|x(t) - x^\star\|_2 \;\leq\; \mathrm{e}^{-ct}\|x(0) - x^\star\|_2 \quad \text{and} \quad \|x_1(t) - x_2(t)\|_2 \leq \kappa \mathrm{e}^{-ct}\|x_1(0) - x_2(0)\|_2$$

S. Jaffe, A. Davydov, D. Lapsekili, A. K. Singh, and F. Bullo. Learning neural contracting dynamics: Extended linearization and global guarantees. In *Advances in Neural Information Processing Systems*, 2024.

Bijective layers for more expressive dynamics

Parametrization alone cannot express all contracting dynamics

Bijective layers for more expressive dynamics

Parametrization alone cannot express all contracting dynamics

Idea: use parametrization in latent space and learn a diffeomorphism to map to data space

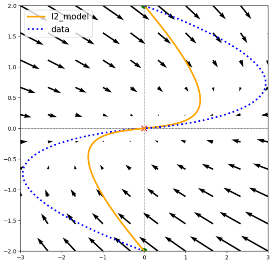

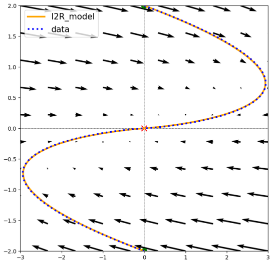

$$\dot{x} = A(x, x^\star)(x - x^\star), \qquad z = g_\theta(x)$$

Contraction is preserved under diffeomorphism

L. Dinh, J. Sohl-Dickstein, and S. Bengio. Density estimation using Real NVP. *International Conference on Learning Representations (ICLR)*, 2017. URL **https://arxiv.org/pdf/1605.08803.pdf**. arXiv preprint arXiv:1605.08803

Numerical results

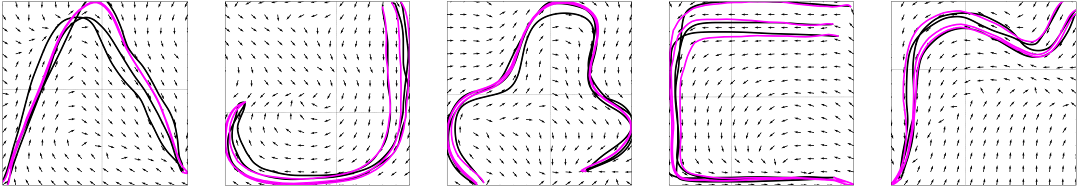

Table: DTWD on LASA, Pendulum, and Rosenbrock datasets

	SDD	EFlow	NCDS	ELCD
LASA-2D	0.37 $\pm$ 0.32	1.05 $\pm$ 0.25	0.59 $\pm$ 0.61	**0.12** $\pm$ **0.11**
LASA-4D	2.49 $\pm$ 2.4	2.24 $\pm$ 0.12	2.19 $\pm$ 1.23	**0.80** $\pm$ **0.54**
LASA-8D	5.26 $\pm$ 0.50	2.66 $\pm$ 0.63	5.04 $\pm$ 0.77	**1.52** $\pm$ **0.61**
Pendulum-4D	0.49 $\pm$ 0.11	0.17 $\pm$ 0.01	1.35 $\pm$ 2.26	**0.03** $\pm$ **0.01**
Pendulum-8D	0.75 $\pm$ 0.08	0.33 $\pm$ 0.01	2.88 $\pm$ 0.69	**0.14** $\pm$ **0.03**
Pendulum-16D	1.86 $\pm$ 0.14	0.45 $\pm$ 0.01	1.65 $\pm$ 0.31	**0.44** $\pm$ **0.09**
Rosenbrock-8D	NaN	1.90 $\pm$ 0.16	2.74 $\pm$ 0.15	**1.22** $\pm$ **0.01**
Rosenbrock-16D	NaN	3.57 $\pm$ 0.66	3.68 $\pm$ 0.12	**2.57** $\pm$ **0.09**