Task Transfer with Stability Guarantees via Elastic Dynamical System Motion Policies

Anonymous Author(s)
Affiliation
Address
email

Appendix

A LPV-DS Parameter Optimization

GMM-based LPV-DS formulation It is described in Section 3 Preliminaries: $\mathcal{PC}$-GMM and LPV-DS Motion Policy.

$$\dot{\xi} = f(\xi) = \sum_{k=1}^{K} \gamma_k(\xi) \left(A^k \xi + b^k\right) \quad \text{s.t.} \begin{cases} \left(A^k\right)^T P + P A^k = Q^k, Q^k = \left(Q^k\right)^T \prec 0 \\ b^k = -A^k \xi^* \end{cases} \tag{1}$$

DS Estimation The set of DS parameters $\theta_{DS} = \{A^k, b^k\}$ for $f(\xi)$ is estimated with LPV-DS by minimizing the Mean Square Error (MSE) against the demonstrations [1] subject to stability constraints in Equation 1.

$$\begin{aligned} \min_{\theta_{DS}} J\left(\theta_{DS}\right) = \sum_{n=1}^{N_{\text{ref}}} \sum_{t=1}^{T_N} \left\| \dot{\xi}_{t,n}^{\text{ref}} - f\left(\xi_{t,n}^{\text{ref}}\right) \right\|^2 \\ \text{s. t.} \begin{cases} \left(A^k\right)^T P + P A^k = Q^k, Q^k = \left(Q^k\right)^T \prec 0 \\ b^k = -A^k \xi^* \qquad \forall k = 1, \ldots, K \end{cases} \end{aligned} \tag{2}$$

which is a constrained non-convex semi-definite program (SDP). Further, when P is known (or estimated beforehand as in [1]) the problem becomes a convex SDP that can be solved highly efficiently with off-the-shelf QP solvers [1, 2].

Submitted to the 7th Conference on Robot Learning (CoRL 2023). Do not distribute.

B 2D Experiments

B.1 Via Point

By specifying an interesting point in the demonstration (usually by human specification or upstream computer vision method), the trajectory will be split into separate components. Each individual segment will be processed with Elastic-GMM to meet the new interesting via-point geometric constraints as depicted in green polygons. By performing such steps, we can pose constraints not just on the endpoints of the demonstration but also on the intermediate points of the demonstration. This via-point experiment shows that with changes in the via-point constraints, Elastic-DS can adapt to them.

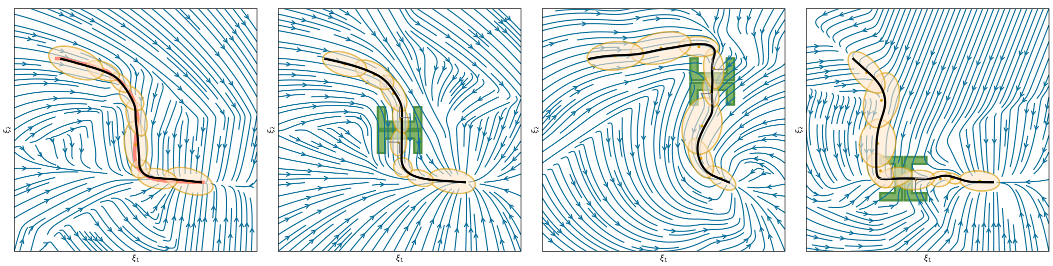

(a) The original DS (b) Original with a via-point constraint (c) Shift the via-point (d) Shift and rotate the via-point

Figure 1: Modified DS based on the via point as a single DS without new demonstration

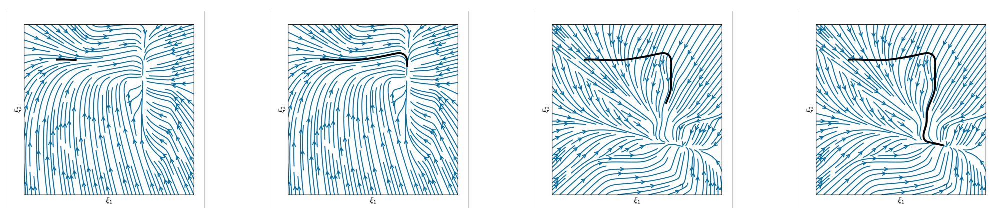

(a) t=1 in the first DS (b) t=2 in the first DS (c) t=1 in the second DS (d) t=2 in the second DS

Figure 2: Modified DS based on the via point as multi-segment DS without new demonstration. The plots here show the rollout trajectory of the switch DS. It uses the same GMM as in 1c.

C Reconstruct from Multiple Segments

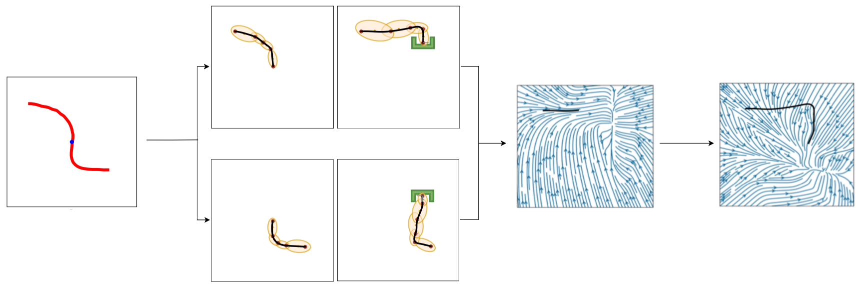

Figure 3: An example of the pipeline for sequential DS. In this case, the single demonstration (in red) is separated at the blue spot. Both segments perform Elastic GMM to meet the geometric descriptor constraints and learn DS individually. The end results are two DS which will run sequentially. After the first DS reaches its attractor, the robot will start to follow the second DS.

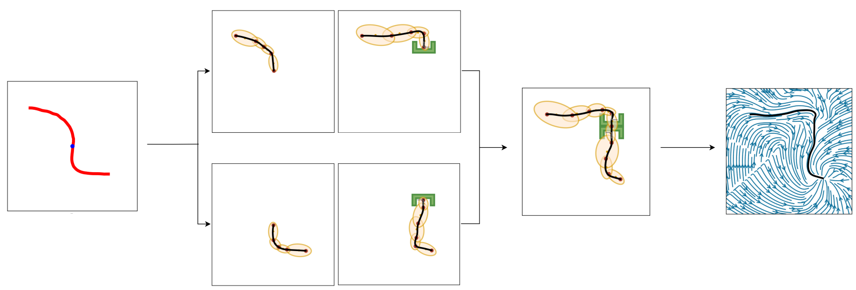

Figure 4: An example of the pipeline for building a single DS. In this case, the single demonstration (in red) is separated at the blue spot. After obtaining the GMM for each segment, they are stitched together and generate a single DS; such an approach will generate smooth velocity but could result in trajectories missing the geometric constraints.

D Robot Experiments Snapshots

D.1 Bookshelf

The goal of this experiment is to teach the robot how to insert a book into a desktop bookshelf. A motion capture marker object is attached to the side of the bookshelf. The geometric descriptor is at an offset from the marker object so that it is inside the bookshelf. We start with the book held by the gripper. It is predefined that the gripper will open once it reaches the attractor.

Figure 5: Demonstration for inserting book into a desktop bookshelf

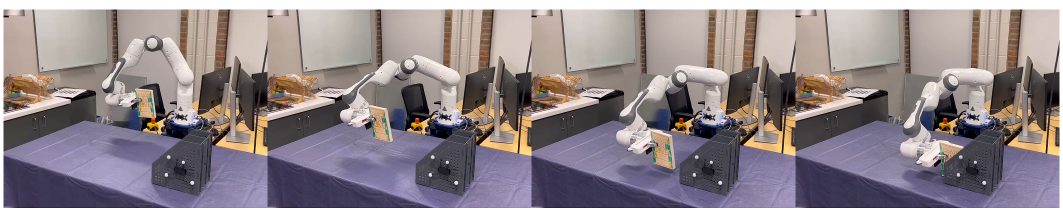

Figure 6: The execution for the learned DS in the original configuration

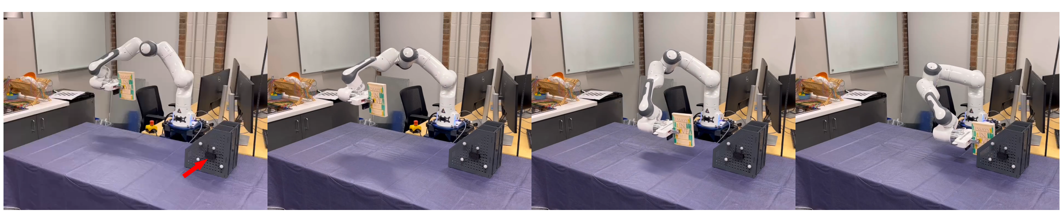

Figure 7: The bookshelf was shifted closer to the robot (The shifting direction is indicated by the red arrow). Without any new demonstration, the learned DS was able to adapt to the new configuration.

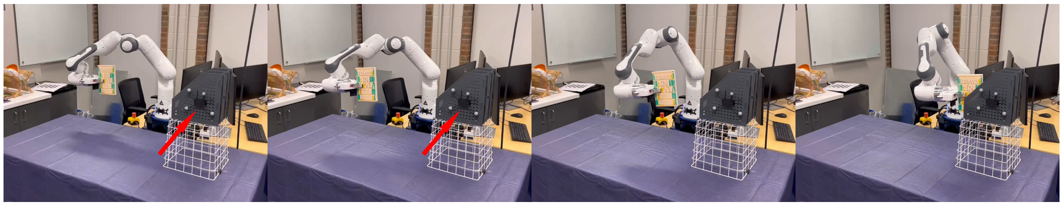

Figure 8: The bookshelf was shifted up (The shifting direction is indicated by the red arrow). Without any new demonstration, the Elastic-DS was able to adapt to the new configuration.

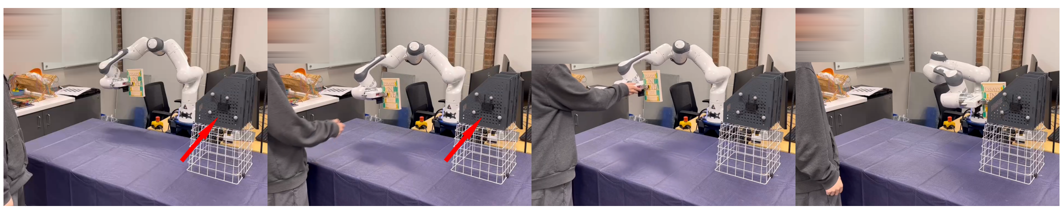

Figure 9: The bookshelf was shifted up (The shifting direction is indicated by the red arrow). Without any new demonstration, the Elastic-DS was able to adapt to the new configuration. The robot was still able to reach the new bookshelf position with human disturbances

Figure 10: The bookshelf is rotated and shifted to the left side of the robot. Without any new demonstration, the Elastic-DS was able to adapt to the new configuration. The robot was still able to reach the new bookshelf configuration. We rotate the end-effector to be parallel with the bookshelf beforehand.

D.1.1 Pick and Place

In this task, we will show the robot how to pick a cube and place it in a bin. The cube position could be changed (labeled by the motion capture marker as the geometric decriptor while the bin position is fixed. During the demonstration, the gripper open/close is done through voice commands (with the microphone at the bottom right of the snapshots). The gripper state is memorized and associated with each segment. At the end of each segment (reaching the attractor), the robot will open/close the gripper depending on the demonstration.

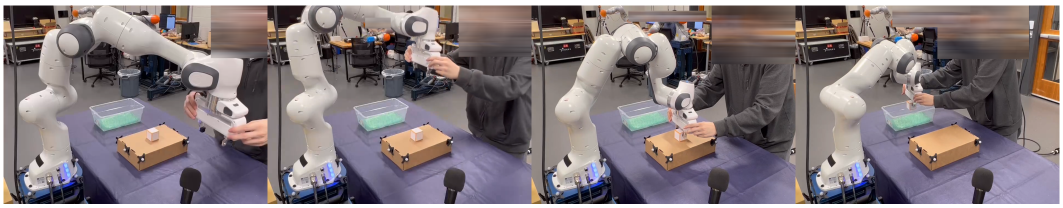

Figure 11: Demonstration for a pick and place task

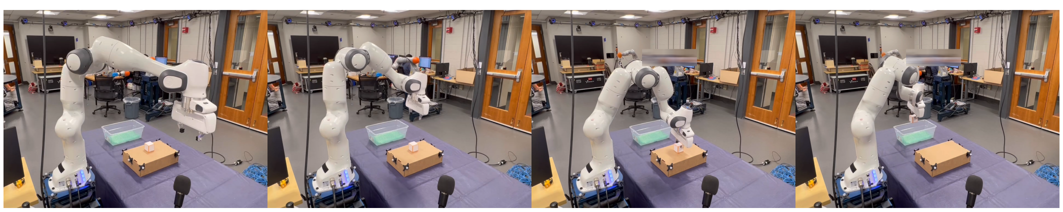

Figure 12: The execution for the learned DS in the original configuration

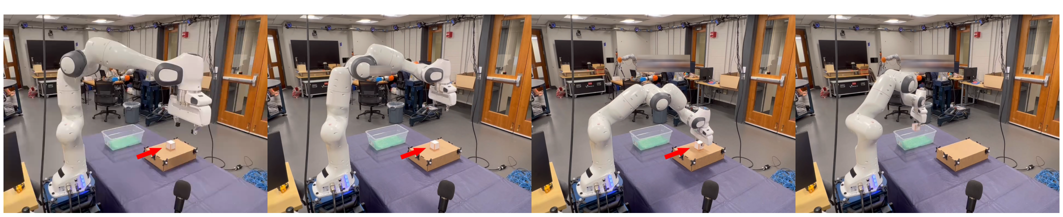

Figure 13: The cube was shifted further away from the robot (The shifting direction is indicated by the red arrow).

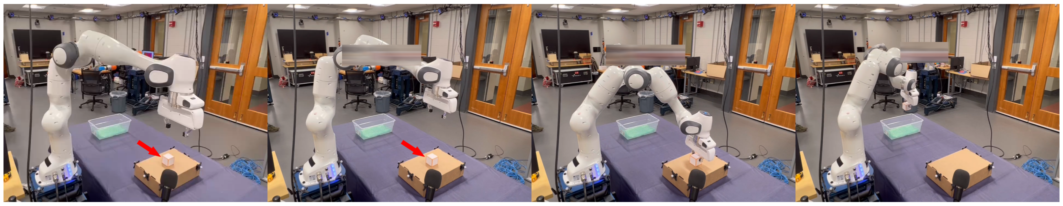

Figure 14: The cube was shifted to the right side of the robot (The shifting direction is indicated by the red arrow).

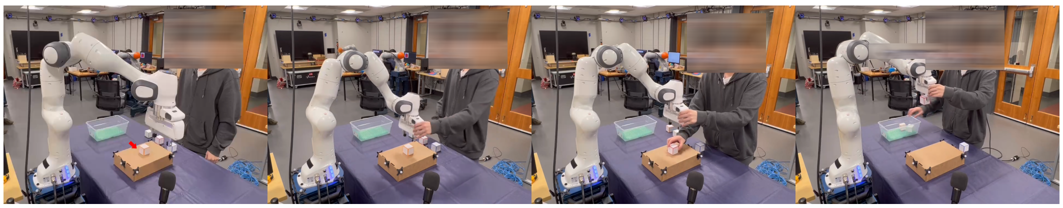

Figure 15: The cube was shifted slightly to the right side of the robot (The shifting direction is indicated by the red arrow). During the execution, the human held the robot and switched to another cube. After that, the robot finished the task

D.1.2 Tunnel

In this experiment, we will show the robot how to pass through a tunnel, mimicking an inspection task. The entrance and the exit are labeled by two different marker objects. There are a total of three segments in this task.

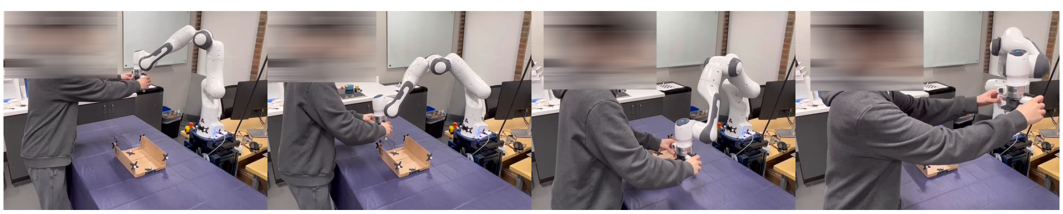

Figure 16: The human guides the end-effector for an inspection task, starting from the left side, passing through a tunnel, and stopping on the right side

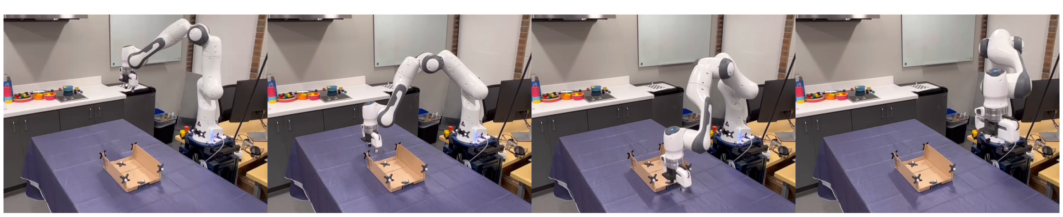

Figure 17: The execution for the learned DS in the original configuration from the demonstration

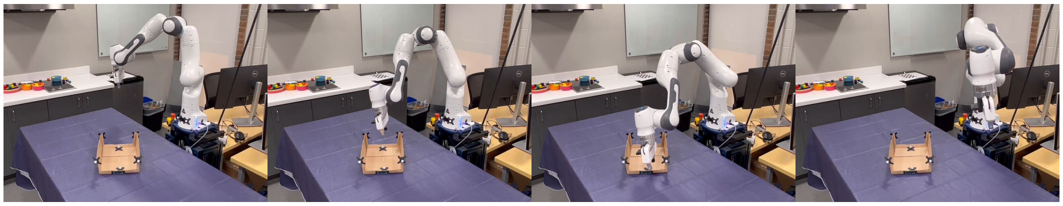

Figure 18: The tunnel is rotated. We rotate the end-effector to be parallel with the tunnel before the execution.

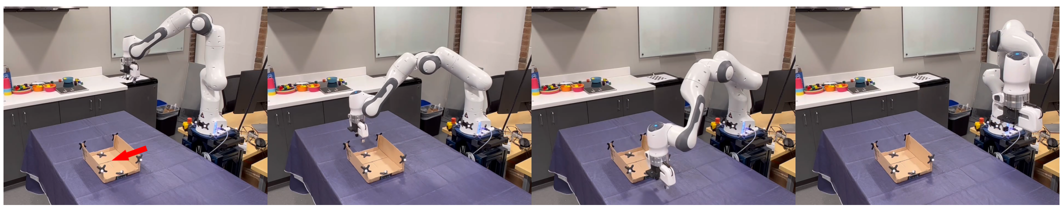

Figure 19: The tunnel is shifted further away from the robot, indicated by the red arrow.

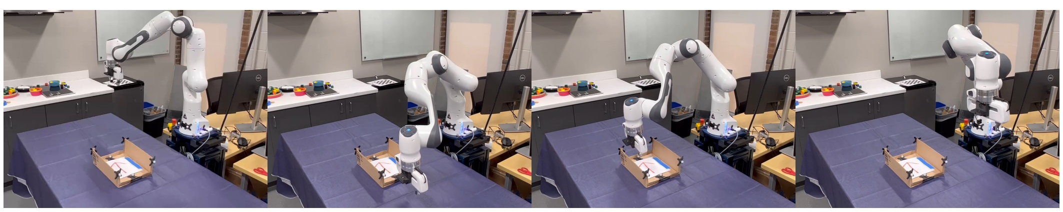

Figure 20: The tunnel is flipped, indicated by the arrow in the tunnel (as opposed to the direction from the demonstration). The robot needs to move to the right side to enter the tunnel and exit to the left side before reaching the end pose.

D.1.3 Combine

There is no demonstration in this task. In the previous experiments, we generate different motion policies for different tasks. We can reuse them (the segments being separated by geometric descriptors) to compose new sequences which perform new tasks.

Figure 21: Composing the learned DSs with task transfer parameters from the ”pick and place” and ”tunnel” tasks. The robot is able to pick up a block, pass through the tunnel for scanning, and place the block in the bin. The entire motion does not require extra demonstration. Note the positions of the objects are not the same as in the original demonstrations.

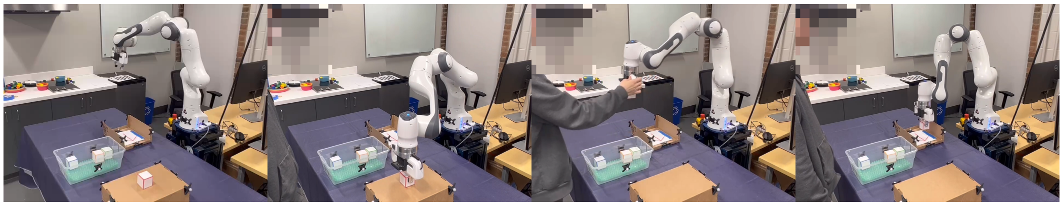

Figure 22: We shift the cube starting platform, the tunnel, and the bin. By reusing and composing the previously learned DS with task transfer, the robot is able to finish the tasks of picking, scanning, and placing without new demonstration in this new environment configuration. Also, there is human disturbance involved during the task execution.

References

[1] N. Figueroa and A. Billard. A physically-consistent bayesian non-parametric mixture model for dynamical system learning. In A. Billard, A. Dragan, J. Peters, and J. Morimoto, editors, *Proceedings of The 2nd Conference on Robot Learning*, volume 87 of *Proceedings of Machine Learning Research*, pages 927–946. PMLR, 29–31 Oct 2018. URL https://proceedings.mlr.press/v87/figueroa18a.html.

[2] A. Billard, S. Mirrazavi, and N. Figueroa. *Learning for Adaptive and Reactive Robot Control: A Dynamical Systems Approach*. MIT Press, 2022.