

Supporting Information for “Learning Assembly Tasks in a Few Minutes by Combining Impedance Control and Residual Recurrent Reinforcement Learning”

Padmaja Kulkarni¹ | Jens Kober² | Robert Babuska³ |
Cosimo Della Santina⁴

¹Affiliation not available

²Affiliation not available

³Affiliation not available

⁴Affiliation not available

This Supporting information includes interactive plots, videos, and data captured while performing evaluation and validation experiments for our paper.

Correspondence

Email:

Funding information

Corresponding author Email: *P. V. Kulkarni@tudelft.nl*

Goal

This paper aims to learn assembly tasks in the presence of goal position uncertainties in real-world settings.

Contributions

- We present a novel approach that combines Cartesian impedance control with recurrent reinforcement learning methods to learn assembly tasks in the presence of position uncertainties.
 - This approach learns to deal with position uncertainties when inserting an object in a frame/hanger within a few minutes of training on a robotic platform.
 - Moreover, the learned policy is robust against the variation in object properties and changes in the frame position and orientation.
- We evaluate/validate our approach in five real-world settings.

1. Supporting Plots

In this scenario, we introduce a Gaussian noise of 0 mean and 15 mm standard deviation along the X , Y , and Z axes of the frame. The setup for this experiment is shown in Fig. 1. We perform five experiments using different random seeds to initialize our RL algorithm. We perform an ablation study for our method, which uses an impedance controller combined with residual recurrent TD3 with decaying nominal controller policy (this approach is explained in detail in our paper). Our method is compared with combinations resulting from removing recurrence and decay and using only an impedance controller.

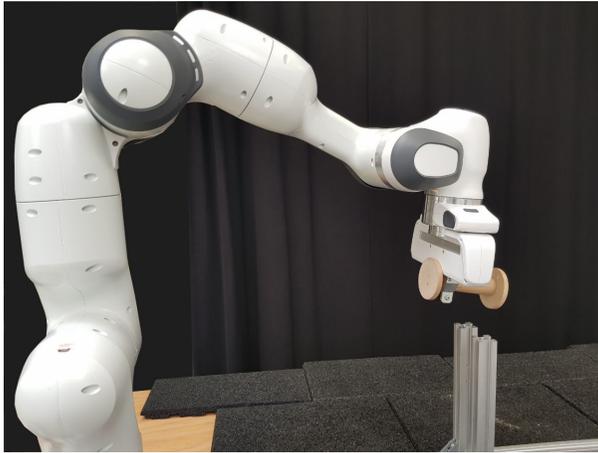


FIGURE 1 Setup for the experiment of inserting the object in the frame.

The resulting plots for average success rate, average episode return, average episode time steps are shown below.

Average success rates

Average episode length needed to finish the task

2. 3-D plots for average success rate, average episode length, and number of training time steps

3. Supporting Videos

Experiment 1: Inserting an object into the frame

Rich media available at <https://youtu.be/a2TYoUtYo5Y>

Experiment 2: Inserting an object with enveloped in foam into the frame

Rich media available at <https://youtu.be/QExQwE31gNI>

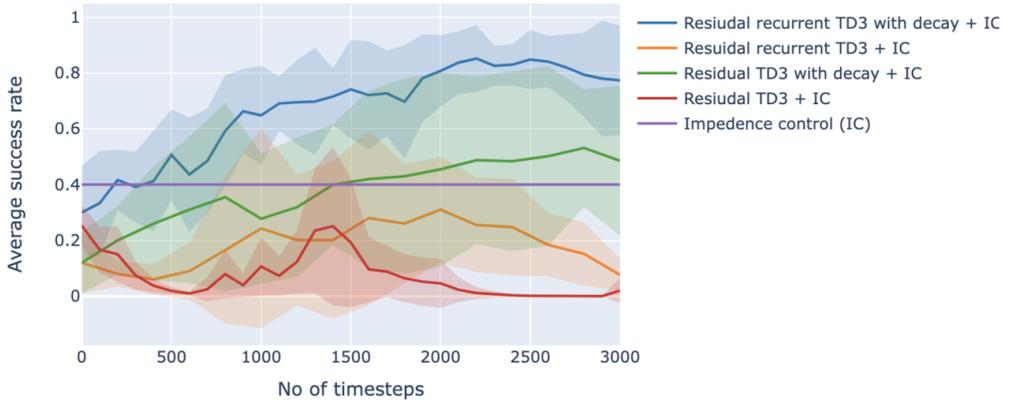


FIGURE 2 Average success rate for five real-world trials. It can be seen that our method, residual recurrent TD3 with impedance controller, learns the task faster than others.

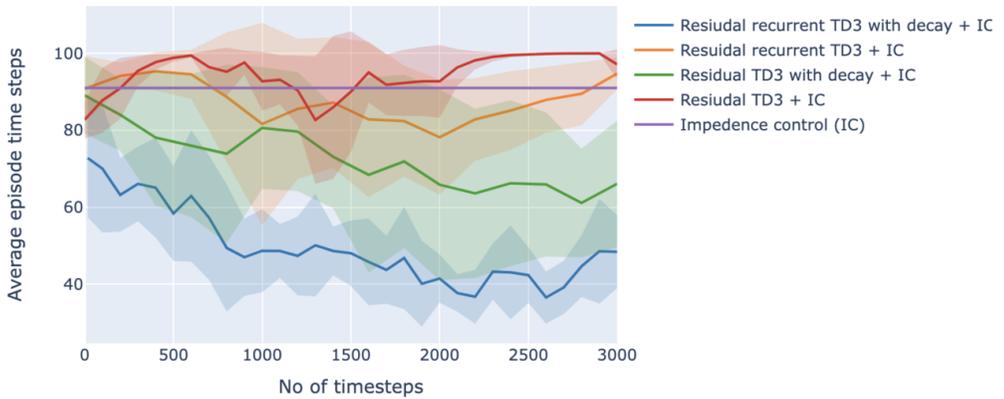


FIGURE 3 Average episode time steps for five real-world trials. It can be seen that our method, residual recurrent TD3 with impedance controller, learns to complete the task in fewer time steps than other methods.

Experiment 3: Hanging an object in the hanger

Rich media available at <https://youtu.be/4yAgNxqefrE>

Experiment 4: Stacking discs on the wooden stick (tower)

Rich media available at https://youtu.be/6cgU_J1r_-w

Experiment 4: Hanging wine glasses in the rack

Rich media available at <https://youtu.be/fscCo-7Qcq0>

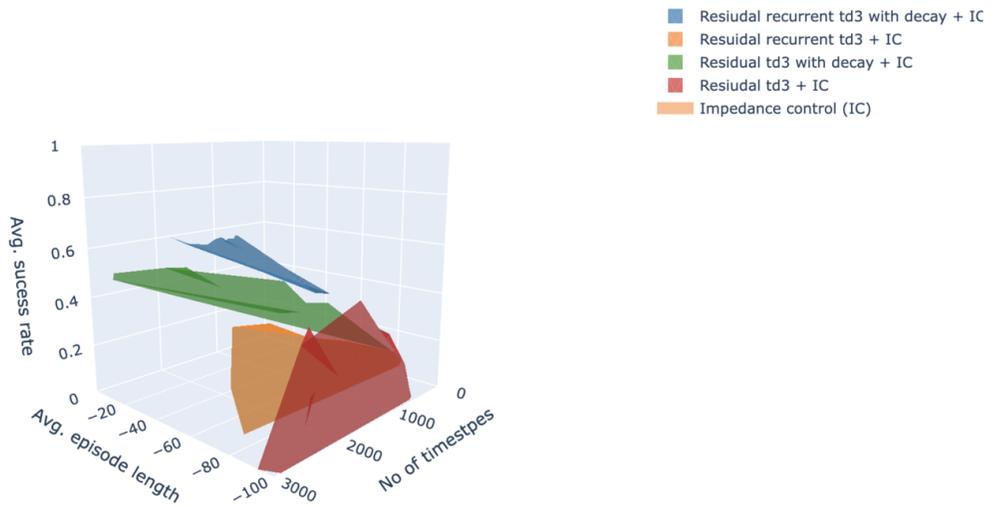


FIGURE 4 This is the 3-D plot for average success rate, average episode length vs. a number of time steps. It can be seen as the number of training steps increase, the average success rate increases, and the number of time-steps needed to finish the task decrease. Our method, residual recurrent TD3 with impedance controller, significantly performs better than other methods.