

Trajectory Optimization for Unknown Constrained Systems using Reinforcement Learning

Ota, Kei; Jha, Devesh K.; Oiki, Tomohiro; Miura, Mamoru; Nammoto, Takashi; Nikovski, Daniel;
Mariyama, Toshisada

TR2019-129 November 07, 2019

Abstract

In this paper, we propose a reinforcement learning-based algorithm for trajectory optimization for constrained dynamical systems. This problem is motivated by the fact that for most robotic systems, the dynamics may not always be known. Generating smooth, dynamically feasible trajectories could be difficult for such systems. Using samplingbased algorithms for motion planning may result in trajectories that are prone to undesirable control jumps. However, they can usually provide a good reference trajectory which a model-free reinforcement learning algorithm can then exploit by limiting the search domain and quickly finding a dynamically smooth trajectory. We use this idea to train a reinforcement learning agent to learn a dynamically smooth trajectory in a curriculum learning setting. Furthermore, for generalization, we parameterize the policies with goal locations, so that the agent can be trained for multiple goals simultaneously. We show result in both simulated environments as well as real experiments, for a 6-DoF manipulator arm operated in position-controlled mode to validate the proposed idea. We compare the proposed ideas against a PID controller which is used to track a designed trajectory in configuration space. Our experiments show that our RL agent trained with a reference path outperformed a model-free PID controller of the type commonly used on many robotic platforms for trajectory tracking.

IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)

© 2019 MERL. This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Research Laboratories, Inc.; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Research Laboratories, Inc. All rights reserved.

Trajectory Optimization for Unknown Constrained Systems using Reinforcement Learning

Kei Ota¹, Devesh K. Jha², Tomohiro Oiki¹, Mamoru Miura¹, Takashi Nammoto¹,
Daniel Nikovski² and Toshisada Mariyama¹

Abstract—In this paper, we propose a reinforcement learning-based algorithm for trajectory optimization for constrained dynamical systems. This problem is motivated by the fact that for most robotic systems, the dynamics may not always be known. Generating smooth, dynamically feasible trajectories could be difficult for such systems. Using sampling-based algorithms for motion planning may result in trajectories that are prone to undesirable control jumps. However, they can usually provide a good reference trajectory which a model-free reinforcement learning algorithm can then exploit by limiting the search domain and quickly finding a dynamically smooth trajectory. We use this idea to train a reinforcement learning agent to learn a dynamically smooth trajectory in a curriculum learning setting. Furthermore, for generalization, we parameterize the policies with goal locations, so that the agent can be trained for multiple goals simultaneously. We show result in both simulated environments as well as real experiments, for a 6-DoF manipulator arm operated in position-controlled mode to validate the proposed idea. We compare the proposed ideas against a PID controller which is used to track a designed trajectory in configuration space. Our experiments show that our RL agent trained with a reference path outperformed a model-free PID controller of the type commonly used on many robotic platforms for trajectory tracking.

I. INTRODUCTION

In this paper, we present a trajectory optimization algorithm for robots with unknown dynamics operating under state and control constraints using Reinforcement Learning (RL). Trajectory optimization is a procedure that produces state and control sequences for a dynamical system under relevant constraints for the system. Most of the state-of-the-art motion planning algorithms generate a plan in the configuration space of the robot, which is then followed by the robot using a trajectory tracking controller [1]. With unknown robot dynamics, the use of traditional PID controllers for trajectory tracking is commonplace. Most of the planning algorithms ignore the dynamics of the robot and hence return a trajectory which is often unsuitable for a lot of target applications [1]. For example, many industrial robots can achieve very high accelerations and torques that can potentially damage the robot as well as the components being manipulated. As an example, in Figure 1 we show a manipulator arm trying to assemble parts of a computer. For such applications, the robot is trying to assemble delicate

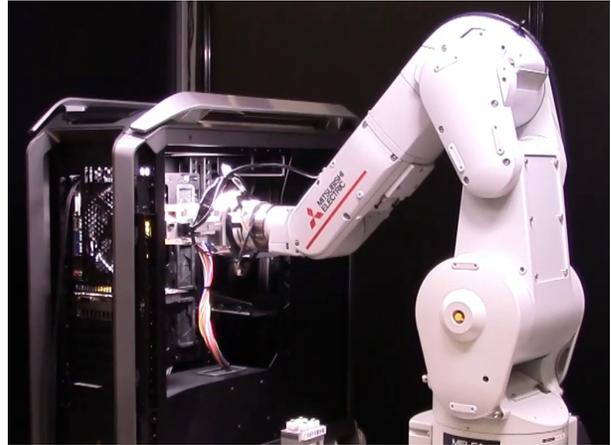


Fig. 1: The proposed method is used to compute smooth trajectories for assembling computer parts using a 6-DoF manipulator arm.

parts connected with wires with fixed length that can easily be broken if high torques are applied. Furthermore, many users deploying such robots have access only to a high-fidelity simulator for the system and its kinematic equations, but not to the true robot dynamics. As a result, a lot of existing results from model-based trajectory optimization cannot be used. Consequently, for many applications, a lot of time is spent in manually designing trajectory tracking controllers. However, high-fidelity simulators provide a good resource that can be used for training an RL agent in the simulated environment, to be used on the real system later.

A model-free RL agent [2] can learn feasible trajectories under dynamic constraints; however, it requires a lot of time for convergence to good solutions. In light of this, we propose to train an RL agent to generate dynamically feasible trajectories, where a sampling-based motion planning algorithm is used to guide the RL agent for faster learning and convergence. We use a RRT-based [1] trajectory to guide an RL agent to learn dynamically-feasible trajectories in the presence of state and input constraints. The RL agent is trained in a curriculum learning setting for faster convergence, and we show that we can achieve good generalization when conditioning the policy on the goal. The proposed algorithm is tested in several simulated environments, as well as with a real robot, by transferring the learned policy from the simulator to the robot. In Figure 1, we show a manipulator arm assembling parts in an environment cluttered with obstacles (a desktop computer) using the proposed algorithm.

¹Kei Ota, Tomohiro Oiki, Mamoru Miura, Takashi Nammoto, and Toshisada Mariyama are with Information Technology R&D Center, Mitsubishi Electric Corporation, Japan. Ota.Kei, Oiki.Tomohiro, Nammoto.Takashi, Mariyama.Toshisada@ds.MitsubishiElectric.co.jp

²Devesh K. Jha and Daniel Nikovski are with Mitsubishi Electric Research Labs, Cambridge, MA, USA. jha, nikovski@merl.com

Contributions. Our paper has the following contributions.

- 1) We propose an RL-based algorithm to efficiently generate control-smooth trajectories for unknown constrained dynamical systems using a goal-directed reference trajectory.
- 2) The proposed algorithm can generalize to local perturbations in goal position in the presence of obstacles.
- 3) We demonstrate the proposed algorithm in both simulation and experimental environments, where the proposed method outperforms a baseline method where a trajectory generated by RRT is tracked by a fine-tuned PID controller.

II. RELATED WORK

Reinforcement learning has recently made huge advances, based on the success of deep learning techniques. Recent RL algorithms have achieved very impressive performance in learning in computer games [3], [4], robotics [5], etc. Broadly, RL algorithms can be classified as model-based or model-free [2], [6]. Model-based algorithms can achieve good sample complexity and generalization, but are generally known to be harder to train for nonlinear dynamical systems. Model-free algorithms, on the other hand, can achieve good asymptotic performance, but suffer from high sample complexity. A lot of recent research has focused on leveraging ideas from control and optimization theory for faster learning [7]–[9]. In a lot of robotics applications, it is generally advantageous to initialize RL agents with demonstrations which can provide them with an initial reference solution [10]. Learning policies from reference trajectories has been studied in [11]–[13]. Motivated by this idea, our work mainly focuses on using reference trajectories that can be provided by off-the-shelf planning algorithms to speed up learning for our RL agent. The planning algorithm serves as a demonstrator for the learning algorithm. The closest work similar to ours is [13]. In [13], a model-based RL agent is learned using a trajectory-centric RL (Guided Policy Search) approach to learn a trajectory-tracking controller for a trajectory provided by RRT [1]. However, using a model-based RL in constrained state and control settings could be difficult, because it is not clear how the underlying trajectory optimization algorithm [14] can account for arbitrary state constraints for manipulator-like systems.

Our combination of RL and reference trajectory tracking can be seen as a form of reward shaping [15]. Reward shaping speeds up learning by creating a more informative reward signal. However, designing shaping rewards requires significant non-trivial reward engineering, and may also alter the optimal solution. To alleviate this problem, automatic reward shaping has been researched [16], [17].

III. BACKGROUND

We consider the standard RL setting that consists of an agent interacting with a stochastic environment. An environment consists of a set of states \mathcal{S} , a set of actions \mathcal{A} , a distribution of initial states $p(s_0)$, a reward function

$r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$, transition probabilities $p(s_{t+1}|s_t, a_t) : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$, and a discount factor $\gamma \in [0, 1]$.

An episode starts with an initial observation s_0 sampled from $p(s_0)$. At each time step t , the agent observes an observation s_t and chooses an action a_t according to a policy $\pi(a_t|s_t)$, which is a mapping from observations to actions: $\pi : \mathcal{S} \rightarrow \mathcal{A}$. Then, the agent obtains a reward $r_t = r(s_t, a_t)$, and the next state s_{t+1} is sampled from $p(s_{t+1}|s_t, a_t)$. The goal of the agent is to maximize the expected discounted sum of rewards $J = \mathbb{E}_\pi[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t)]$. The quality of the agent’s action a_t when receiving an observation s_t can be measured by a Q function $Q(s_t, a_t) = \mathbb{E}_\pi[J|s_t, a_t]$.

DDPG [18] is a model-free Q-learning-based reinforcement learning algorithm for continuous action spaces. It is an extension of the earlier DQN agent [4], using generating distributions over continuous action spaces. In DDPG, we maintain two neural networks: a deterministic policy (called the actor) π and a Q function approximator (called the critic) $Q : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$, parameterized by a set of parameters θ^π and θ^Q . An actor network deterministically maps observations to actions and tries to maximize $Q(s_t, a_t)$. DDPG employs a critic neural network to estimate Q by minimizing the Bellman loss:

$$L_{\text{critic}} = (Q(s_t, a_t) - y_t)^2, \quad (1)$$

where the 1-step target y_t is calculated using target networks Q^* and π^* as

$$y_t = r_t + Q^*(s_{t+1}, \pi^*(s_{t+1})). \quad (2)$$

Each transition of the agent is stored in a replay buffer, from which mini-batches are sampled to train the networks. This stabilizes training by removing temporal correlations, and therefore reduces the changes in the distributions the networks are trying to learn. Additionally, a prioritized replay buffer [19] assigns a priority p_i to each transition, computed as the last temporal difference (TD) error and a small hyperparameter ϵ . For more details, see [19].

However, earlier research has shown that DDPG is prone to overestimating Q-values, and results in sub-optimal policies. TD3 [20] implements three improvements to address the overestimation resulting from approximation errors. First, it maintains two independent critic networks, and always employs the minimum Q-value as the optimization target. Second, it proposes to delay the propagation of weight updates. Finally, it explicitly increases the smoothness of the Q-function prediction by adding a clipped normal noise to the action to the target Q-value. Using these three improvements, we can replace the 1-step target of the critic defined in (2) with

$$y_t = r_t + \min_{i=1,2} Q_i^*(s_{t+1}, \pi^*(s_{t+1}) + \text{clip}(\mathcal{N}(0, \sigma), -c, c)). \quad (3)$$

IV. PROPOSED APPROACH

In this section, we present details of the algorithm and some techniques which allow us to train the algorithm efficiently. We train a TD3 agent (see Section III) using a

reference trajectory provided by RRT. Furthermore, due to the constrained nature of the problem (presence of obstacles), we use curriculum learning to simplify learning for the TD3 agent. These are explained in detail next, and presented as a psuedo code for clarity in Algorithms 1 and 2.

A. Reinforcement Learning with Reference Trajectory

We consider the standard RL problem described in Section III with a reference trajectory $z \in \mathcal{Z}$. We include the information about the reference trajectory into a reward function as $r : \mathcal{S} \times \mathcal{A} \times \mathcal{Z} \rightarrow \mathbb{R}$. Therefore, the reward function can be written as

$$r(s_t, a_t, z_t) = f(s_t, a_t) + h(s_t, a_t, z_t). \quad (4)$$

$f(s_t, a_t)$ is the reward that originates from a pure RL setting, and $h(s_t, a_t, z_t)$ is calculated using the reference trajectory. The idea is to accelerate the learning process by the additionally defined term $h(s_t, a_t, z_t)$ in the reward function. This term penalizes search too far from the reference trajectory, and thus limits the search space for the agent.

In prior work, an expert trajectory is generally used to define the function $h(s_t, a_t, z_t)$. In contrast, we use a standard RRT algorithm for generating a reference path, because the computational cost for generating a path is much smaller than that of doing RL. Due to the nature of random-sampling based algorithms, RRT produces a jerky path, and it results in jerky trajectories, because the critic directly optimizes the reward function. In order to mitigate this problem, we investigated two improvements.

First, we reduce the number of vertices that describe the trajectory by randomly short-cutting between them, as described in [21]. To do so, we randomly pick two vertices, and divide the trajectory that connects the two points with a fixed distance. Then, for each vertex, we check if it has contact with obstacles or not, and if none of the vertices collides with obstacles, we short-cut the path, i.e. omit the vertices between the selected two vertices.

Second, we replace the reference trajectory in every episode with a path found by the RL agent during training that satisfies that: 1) the current path reaches the goal without colliding with obstacles; 2) the number of total steps to achieve the goal is the lowest, and 3) the cumulative reward is the highest.

B. Resets to Reference Path

To overcome the problem of exploration, we reset some training episodes to a reference path with a probability of $p_{\text{reset}} \in [0, 1]$. Restarts from them makes the agent explore more efficiently, because the reference trajectory is guaranteed to reach the goal. Prior work [10] employs expert trajectories and resets to a state in them; however, we do not have such expert trajectories. Instead, we utilize a reference path that ensures that the goal is reached. To reset to a reference path, we uniformly sample joint angles θ_t from the set of reference trajectories, and assign the start angle to the sampled value.

C. Curriculum Learning

Generally, RL is harder to train if the reward function is sparse, and if an episode is longer. To simplify learning for the agent, we gradually increase the complexity of the problem in a curriculum learning setting. We use curriculum learning in two different settings. First, we train the RL agent to learn a controller close to the reference trajectory without any sparse penalty for collision (i.e., $w_3 = 0$), with the intuition that state constraints make the problem harder. This in turn provides more dense rewards to reach the goal. Once the agent learns successfully how to learn in the absence of obstacles, we introduce penalty for collision, which we gradually increase, to adapt to the obstacles. Second, to make learning easier, we gradually decrease the goal region for the agent. For our problem, we define an acceptable goal area as r_{goal} , and we declare success when $\theta_{\text{dist}} = \|\theta_t - \theta_{\text{goal}}\|_2 \leq r_{\text{goal}}$. (Note that θ_t is six-dimensional).

The gradual change in penalty for collision w_3 forces the critic to fit to noisy targets. This could lead the actor to converge to a non-optimal local minimum, since the actor learns with gradients computed using the critic network. To avoid this, we store past good experiences in a replay buffer, which is different from the prioritized replay buffer that was described in III, and encourages the actor to choose the same action as the past good experiences in any given state. The additionally prepared buffer stores the past best K episodes as in [22], in the sense that 1) the RL agent reaches the goal without colliding with any obstacle, and 2) the RL agent gets higher episode rewards.

To imitate from such good experiences, we use the behavioural cloning loss L_{BC} which was proposed in [10] and is defined as:

$$L_{\text{BC}} = \begin{cases} (\pi(s_i) - a_i)^2, & \text{if } Q(s_i, a_i) > Q(s_i, \pi(s_i)) \\ 0 & \text{otherwise} \end{cases}. \quad (5)$$

After sufficient training, the agent might surpass the performance of the past best experiences and thus L_{BC} would then become detrimental to the agent's performance. The Q-filter mitigates this problem by only applying L_{BC} if the critic judges that the action proposed by the actor is worse than the action of the demonstrator, in out setting, past good experiences. Based on adding a behavioral cloning loss, the actor loss results in:

$$L_{\text{actor}} = -L_{\text{critic}}(\pi(s_t)) + \lambda_{\text{BC}} L_{\text{BC}}, \quad (6)$$

where $\lambda_{\text{BC}} \in \mathbb{R}^+$ is a hyperparameter for actor to balance learning from critic or past good experiences.

The whole learning procedure is provided as psuedo code in Algorithm 1 and Algorithm 2. These algorithms are implemented in a curriculum setting, as described above.

D. Goal Parameterization of Policy

To achieve generalization to perturbations in the target state for the agent, we parameterize the policy of the agent on the goal. The idea is that a goal-parameterized policy represented by a network with enough capacity should be able to generalize to perturbations in the goal location. This

is a very desirable property to have in the final policy, because robots are often expected to adapt to some local perturbations in the target state. We assume that the target state g is sampled from the set \mathcal{G} . We train a single network to maximize the expected discounted reward over multiple goal states. The learning problem is to optimize the following expected discounted reward:

$$J(\pi) = \sum_g J(\pi_g) = \sum_g \mathbb{E}_{\pi_g} \left[\sum_{t=0}^{\infty} \gamma^t r_g(s_t, a_t) \right]$$

across all goals $g \sim \mathcal{G}$. The reward r_g is now conditioned on the goal to reflect the fact that rewards depend on the particular goal (or task). This is achieved by increasing the capacity of the network by adding additional input units to the network. In the simplest setting, we achieve this by simply padding extra inputs to the network that contain the goal information.

Algorithm 1 Learning procedure

- 1: Initialize τ' ▷ **Buffer for reference trajectory**
 - 2: Initialize $R_A \leftarrow \{\}$ ▷ **Replay buffer for RL**
 - 3: Initialize $R_K \leftarrow \{\}$ ▷ **Replay buffer for top- K episodes**
 - 4: Compute a reference trajectory τ using RRT
 - 5: Smooth out τ by short-cutting ▷ **Reference trajectory input to RL**
 - 6: **while** Termination condition is False **do**
 - 7: $t_{\text{step}} \leftarrow 0, \tau' \leftarrow \{\}$
 - 8: Initialize s_t with initial state defined in each task
 - 9: Replace the initial state with an uniformly sampled state from τ with probability p_{reset}
 - 10: **while** $t_{\text{step}} < \text{episode length}$ or $\theta_{\text{dist}} \leq r_{\text{goal}}$ **do**
 - 11: Sample using Algorithm 2
 - 12: **end while**
 - 13: **if** Trajectory reaches the goal **then**
 - 14: **if** τ' satisfies the update condition **then**
 - 15: Update reference trajectory, $\tau \leftarrow \tau'$
 - 16: **end if**
 - 17: Update curriculum learning setting
 - 18: **end if**
 - 19: **end while**
-

V. SYSTEM OVERVIEW

In this section, we provide relevant details of the simulator and the real system we used in this paper for our experiments.

A. Hardware

We use a MELFA RV-FR robot, which is an industrial robot that has 6 degrees of freedom [23]. The generated trajectories must ensure that joint angles and angular velocities that consist the trajectories are within a known specified range. The robot used in the experiments in the paper is operated in a position control mode where a position command is sent to the robot every $\Delta t = 0.0035$ seconds, which comes from the minimum operational time of the

Algorithm 2 Environment sampling

- 1: Observe s_t and r_t .
 - 2: Store data for t into $R_A: \{s_t, r_t, a_t\}$ and $\tau' \leftarrow \tau' \cup \{\theta_t\}$
 - 3: **if** s_t concludes an episode **then**
 - 4: Perform step of TD3 ▷ **Update actor and critic networks weights**
 - 5: **if** Current episode deserves top- K episodes **then**
 - 6: Update top- K replay buffer R_K
 - 7: **end if**
 - 8: **else**
 - 9: Sample the current policy $a_t \sim \pi(a|s_t)$
 - 10: Advance the environment by performing a_t
 - 11: **end if**
-

industrial robot we used in a real setting. As a result, the control input is the velocity for each joints. We, however, would like to minimize the acceleration (i.e., the derivative of the control signal or the control jumps) during operation. This is a desirable feature for a lot of industrial manipulators where direct torque control is not accessible.

B. Simulator

We utilize a simulator to generate trajectories and then deploy them in a real setting. The simulator is a high-fidelity simulator for the MELFA RV-FR called RT ToolBox3 [24]. The baseline controller we use to compare the RL agent in this work is a PID trajectory-tracking controller that can be designed in the simulator given a reference trajectory. For our experiment, this function is our initial baseline, which is described in detail in VI. The simulator has a built-in function for collision checking between the manipulator and obstacles present in the environment, and we use the same function for collision detection during planning. However, the proposed algorithm is agnostic to the collision checking method and simulation environment.

VI. EXPERIMENTS

In this section, we will describe several different environments in which we test our proposed algorithm. In particular, we test it in two environments in simulation—a *Book-shelf* environment (see Figure 2a) and an *Open-Computer* environment (see Figure 2b). In these environments, the robot is trying to manipulate objects that can be damaged if excessive torque or acceleration is applied. Furthermore, we will show experimental results with a real robot for the *Open-Computer* environment (see Figure 2c). Videos of the learned behavior of the robot could be seen in the supplementary material.

In our experiments, we try to investigate the following questions:

- 1) Does the combination of a reference trajectory and RL improve the performance of each one of them in isolation?
- 2) Does the proposed algorithm generate feasible trajectories in the presence of state and control constraints better than some of the traditional control techniques of trajectory tracking with a reference trajectory?

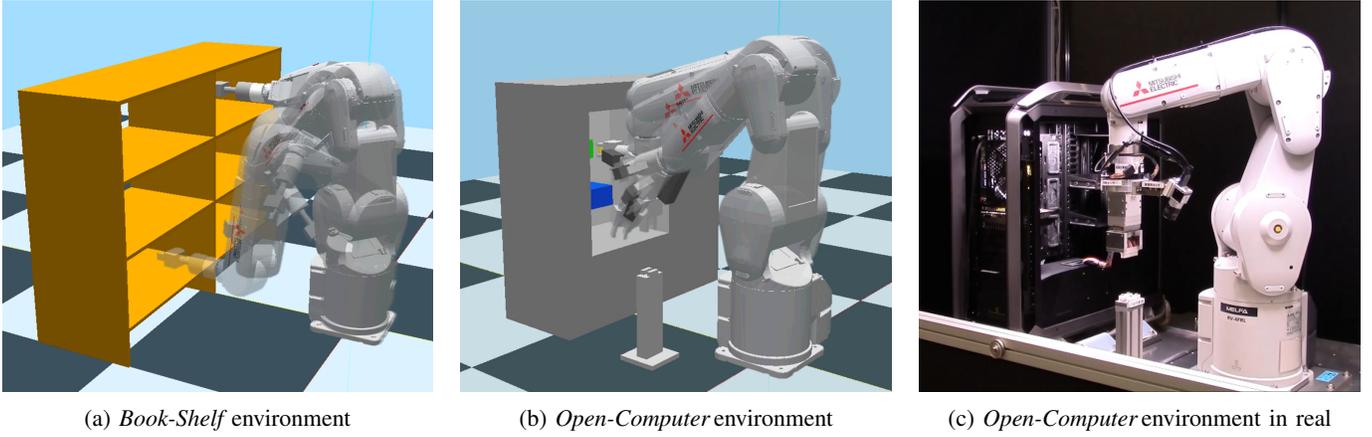


Fig. 2: We show the trajectories obtained by the proposed method in the above figures for the three different settings.

3) Does curriculum learning helps the agent learn faster?

In the following text, we answer the above-mentioned questions, and demonstrate that we can generate smooth trajectories and the agent can generalize to unseen goal conditions upon conditioning the policy on goal position.

A. Environment

1) *States*: The states of the system consist of current angles θ_t and angular velocities $\dot{\theta}_t$. Therefore, the state set is represented in $\mathcal{S} \in \mathbb{R}^{12}$. The initial angles and angular velocities are deterministically reset to $\theta_0 = \theta_{\text{start}}$, and $\dot{\theta}_0 = \mathbf{0}$.

2) *Actions*: The action of the agent \mathbf{a}_t is the vector of angular velocities $\dot{\theta}$ for the next step. Since we consider a six dimensional configuration space environment, the action set is described as $\mathcal{A} \in \mathbb{R}^6$. We define a time step described as Δt . Therefore, the angles of the next step θ_{t+1} can be calculated as

$$\theta_{t+1} = \theta_t + \mathbf{a}_t \times \Delta t. \quad (7)$$

3) *Rewards*: As described in IV-A, we add $h(s_t, a_t, z_t)$, which is calculated from the reference trajectory, to the conventional reward term $f(s_t, a_t)$. First, referring to [25], we define the conventional reward term as

$$f(s_t, a_t) = w_1 d_{\text{goal}} + w_2 \mathbb{I}_{\text{goal}} + w_3 \mathbb{I}_{\text{collision}} + w_4 \mathbb{I}_{\text{anglelimit}} + w_5 \|\ddot{\theta}\| + w_6, \quad (8)$$

where d_{goal} is an Euclidean distance to the goal, i.e., $d_{\text{goal}} = \|\theta_t - \theta_{\text{goal}}\|$. \mathbb{I}_{goal} , $\mathbb{I}_{\text{collision}}$, and $\mathbb{I}_{\text{anglelimit}}$ are indicators of whether the agent reaches the goal, and whether collision between the agent and the obstacles occurs, and whether the agent violates the constraint of joint angles respectively. The fifth term encourages the agent to generate a smoother trajectory, which is essential when operating the real system. The final term is negative value, so it encourages the agent to reach goal with smaller steps.

Then, we design an additional term by using a reference path as

$$h(s_t, a_t, z_t) = w_7 d_{\text{path}} + w_8 n_{\text{progress}}, \quad (9)$$

where d_{path} is the distance to the reference path and n_{progress} is the progress along the path. The first term penalizes search too far from reference path, and second term encourages to go towards goal target angles along with reference path.

In order to calculate d_{path} and n_{progress} , we divide the reference path and agent's path at regular intervals, as shown in Fig. 3. By dividing the path, we obtain the subsampled vertices $\mathbf{p}'_0, \mathbf{p}'_1, \dots, \mathbf{p}'_{N'_p-1}$ for the reference path, and $\theta_t^0, \theta_t^1, \dots, \theta_t^{N_t-1}$ for the agent's path, where N'_p and N_t are the numbers of vertices in each divided path. We can then define the distance to the given path as $d_{\text{path}} = \max_i D(\theta_t^i)$, where $D(\theta)$ is the distance to the path calculated as $\min_i \|\theta - \mathbf{p}'_i\|$. We can also observe the progress along the path as $n_{\text{progress}} = \text{NNI}(\theta_t) - \text{NNI}(\theta_{t-1})$, where $\text{NNI}(\theta)$ is the vertex index of the nearest neighbor to θ , i.e., $\text{NNI}(\theta) = \arg \min_i \|\theta - \mathbf{p}'_i\|$.

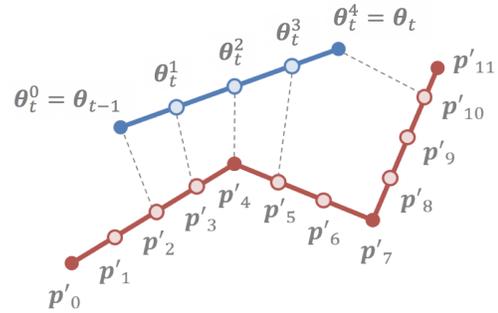


Fig. 3: Path division for calculating rewards. The blue line is the agent's path, and the red line is the reference path. θ_t is the joint angles of the agent at time step t , and \mathbf{p}'_j is the index of the divided path. The dashed lines indicate the correspondence to the nearest neighbor.

4) *Termination Condition*: An episode terminates with following two conditions: the joint angles of the agent θ_t are sufficiently close to the goal state as described in IV-C, or the number of steps of an episode is over a specified threshold.

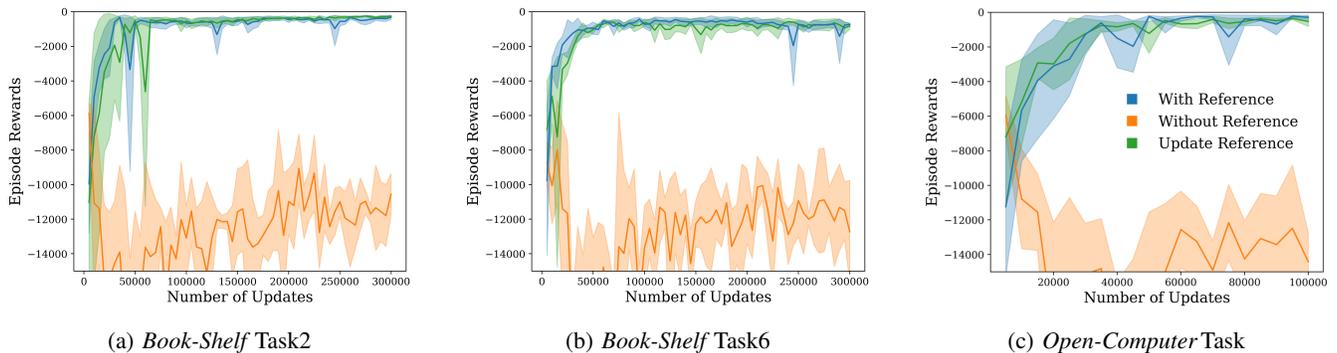


Fig. 4: A comparison between RL without reference paths, RL with reference paths, and RL with reference paths that are updated in the course of learning. The experiments are conducted over 5 random seeds. The bold line shows the average episode rewards, and the shaded region is one standard deviation from the average. The plot shows the faster and stable learning that we achieve using a reference trajectory which is the updated for self imitation in each episode.

B. Book-Shelf Environment

The *Book-Shelf* environment consists of a two-row, three-stage bookcase, simulating a pick and place task. Each of the cube in the bookshelf is 300 mm deep, 200 mm high and 500 mm wide. The manipulator starts from an initial pose denoted by θ_{start} , and has to reach 6 different points specified as $\theta_{\text{BL}}, \theta_{\text{ML}}, \theta_{\text{TL}}, \theta_{\text{BR}}, \theta_{\text{MR}}, \theta_{\text{TR}}$, those are center positions for each cube of the bookshelf, defined as $\theta_{\text{start}} = [0, 8, 131, 0, 41, 180]$, $\theta_{\text{BL}} = [-52, 59, 106, -141, 78, 170]$, $\theta_{\text{ML}} = [-52, 28, 111, -134, 60, 152]$, $\theta_{\text{TL}} = [-52, 13, 95, -111, 42, 117]$, $\theta_{\text{BR}} = [-128, 59, 106, 141, 78, 190]$, $\theta_{\text{MR}} = [-128, 28, 111, 134, 60, 208]$, $\theta_{\text{TR}} = [-128, 13, 95, 111, 42, 243]$. We define task 1 to 6 as reaching from θ_{start} to those angles defined above.

C. Open-Computer Environment

The *Open-Computer* environment is for simulating a computer assembly, picking up a connector, and inserting it into a socket mounted on a mother board as illustrated in Fig. 2b. The picking part and insertion part is out of our focus, so the simulation starts from just above the connector place with an angle of $\theta_{\text{start}} = [-47, -8, 113, 0, 75, -138]$, and the goal is near the socket, denoted by $\theta_{\text{goal}} = [-90, -1, 138, -180, 46, 88]$. The real setup shown in Fig. 2c is the same as the above environment, except that the robot is grasping a connector with a harness. A video of the implementation of the algorithm on the real manipulator is provided in the supplementary materials.

VII. EXPERIMENTAL RESULTS

This section presents results from experiments designed to answer the questions described in VI. The baseline that we compared our proposed method with is a combination of a reference path and a PID-based trajectory tracking controller implemented in our simulator, as described in V-B. Note that the reference path is generated using RRT and is smoothed out by short-cutting, as described in IV-A.

A. Accelerating RL by Using Reference Paths

First, we evaluate the effectiveness of using reference paths to train an RL agent. We compare three learning methods. Firstly, we train an RL agent without a reference path by setting $w_7 = w_8 = 0$ in Eq.(9). Secondly, we train with reference paths, and finally, we train with a reference path while it is being updated in every episode if it satisfies the conditions described in IV-A. The evaluation metric is the cumulative episodic reward that an RL agent gets during an episode. For fair comparison between methods with and without reference paths, we omit reward terms that come from the reference path by setting $w_7 = w_8 = 0$ in Eq.(9).

Figure 4 shows the resulting episodic returns. It suggests that the use of reference paths improves convergence performance with respect to the training without reference path. Also, updating the reference path improves the performance more, because the initial reference path is jerky, and that may result in converging to a non-optimal trajectory. Thus, we see that the use of a reference trajectory for training of the RL agent helps in speeding up policy learning.

B. Generating Smoother and Shorter Trajectories using RL

Next, we compare the quality of the trajectories obtained by the proposed algorithm against the baseline method. We use two metrics to quantify the quality of the trajectories obtained: the time needed to reach the goal, and the magnitude of acceleration. Recall that part of the initial motivation to training the agent this way was to minimize control jumps, and thus generate trajectories with limited acceleration. Table I shows the time required by the proposed algorithm to reach the goal using the proposed algorithm and compared against the baseline. This clearly demonstrates that the proposed method generates high quality trajectories. Figure 5 shows the angular acceleration during a rollout of the proposed method, compared against the baseline method. It shows that the proposed method generates trajectories with much lower acceleration profiles than those generated by the baseline method in all joints, while also minimizing the time taken to reach the goal.

TABLE I: Time [sec] to reach goal.

Task	<i>Open-Computer</i>	<i>Book-Shelf</i>					
		1	2	3	4	5	6
Baseline	0.82	0.56	0.65	0.62	0.80	0.75	1.28
Ours	0.22	0.34	0.26	0.25	0.48	0.55	0.50

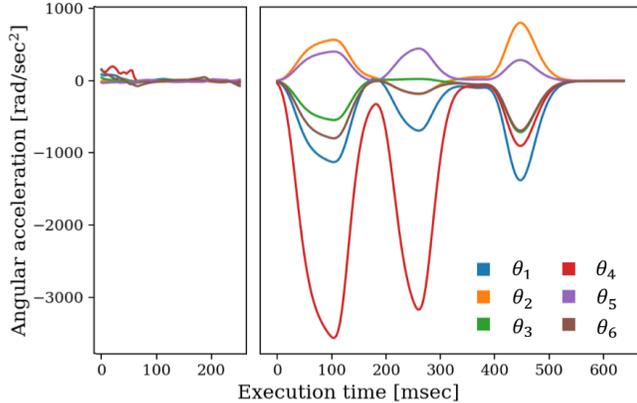


Fig. 5: Accelerations for *Book-Shelf* environment task 1. Left figure is the result of our proposed method and right figure is generated by RRT and PID controller. Note lower is better.

C. Curriculum Learning

Next, we investigate how our curriculum learning helps training our RL agent. We compare our full model with the one without curriculum learning and self imitation for task 1 in the *Book-Shelf* environment. Note that in TD3, the actor learns to maximize the Q function (critic) parameterized by a neural network. Therefore, if the estimation of Q function is insufficient, it gives undesirable gradients to the actor and that would result in lower episode rewards.

Fig. 6 shows a comparison of convergence rates of the agent using different methods. Without curriculum learning, the agent achieves slower convergence, because the training of the critic is harder due to a huge collision penalty, and it is harder to get positive reward which the agent receives only upon reaching the goal. Also, training without self imitation results in unstable training, because the critic needs to fit the noisy reward because of changing collision penalty w_3 in Eq. (8). Our full model is both stable and converges faster, because curriculum learning makes it easier for the critic to fit the Q function, and self imitation mitigates the noisy reward problem by imitating past good experiences.

D. Generalization to Goal Perturbation

Next, we evaluate the generalization of our method with respect to goal perturbation. As shown in IV-D, we add the goal state to both the actor and critic networks, in the expectation that the method can generalize over goal states. The target task is task 2 of *Book-Shelf* environment, moving from start angle θ_{start} to θ_{ML} defined in VI-B. As for targets, we fix a plane in a cube of the *Book-Shelf* and change goal

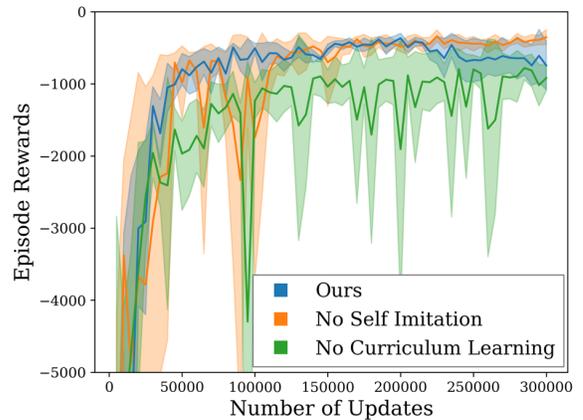


Fig. 6: A comparison between our proposed model, removing curriculum learning, and removing self imitation. The experimental conditions are same with Fig. 4

state within a $W \times H$ [mm²] rectangle. For training, we randomly sample N different positions in the rectangle. After training, we test the generalization by randomly sampling 50 positions in the same rectangle used during training, recording whether the robot reached the goal or not. To exploit the past good experiences, we prepared top- K buffer for each goal, and tried to make the training more stable and improve sample efficiency.

Table II shows the result of the experiments. It demonstrates that the RL agent successfully generalizes to goal perturbations over a reasonable area even in the presence of state constraints.

TABLE II: Performance on generalization task in simulation

(W, H, N)	(100, 100, 10)			(200, 200, 50)		
	Train	Test	Overall	Train	Test	Overall
Number of successes	10/10	49/50	59/60	49/50	45/50	94/100
Success rate	1	0.98	0.98	0.98	0.90	0.94

VIII. CONCLUSION

The research reported in this paper is based on the idea to combine RL with trajectory optimization for unknown systems in the presence of constraints on state, control and control-jumps. This kind of problems is common in robotics, where a manipulator has to be used for tasks in an environment cluttered with obstacles, in a position-control mode. We proposed a method based on RL, for the case when the dynamics are unknown, that generates optimal trajectories in the presence of obstacles and other constraints. For faster learning, we use an off-the-shelf sampling-based algorithm to first generate a reference trajectory which is then used by the RL agent to converge to an optimal solution faster. The proposed method was demonstrated on several simulated environments using a high-fidelity simulator for an industrial-grade manipulator. We compared the learned policy against a

baseline controller designed to track the trajectory obtained by smoothing the initial reference trajectory. The proposed algorithm was also tested for generalization to multiple new target states.

In future research, we would like to investigate the proposed algorithm by parameterizing it with the reference trajectory. We expect that as long as we do not change the environment, the agent would learn to produce a better solution respecting all the constraints.

REFERENCES

- [1] S. M. LaValle, *Planning algorithms*. Cambridge university press, 2006.
- [2] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. MIT press, 2018.
- [3] D. Silver, J. Schrittwieser, K. Simonyan, I. Antonoglou, A. Huang, A. Guez, T. Hubert, L. Baker, M. Lai, A. Bolton *et al.*, “Mastering the game of go without human knowledge,” *Nature*, vol. 550, no. 7676, p. 354, 2017.
- [4] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski *et al.*, “Human-level control through deep reinforcement learning,” *Nature*, vol. 518, no. 7540, p. 529, 2015.
- [5] S. Levine, C. Finn, T. Darrell, and P. Abbeel, “End-to-end training of deep visuomotor policies,” *The Journal of Machine Learning Research*, vol. 17, no. 1, pp. 1334–1373, 2016.
- [6] J. Kober, J. A. Bagnell, and J. Peters, “Reinforcement learning in robotics: A survey,” *The International Journal of Robotics Research*, vol. 32, no. 11, pp. 1238–1274, 2013.
- [7] J. Schulman, S. Levine, P. Abbeel, M. I. Jordan, and P. Moritz, “Trust region policy optimization,” in *ICML*, vol. 37, 2015, pp. 1889–1897.
- [8] S. Levine and V. Koltun, “Guided policy search,” in *Proceedings of the 30th International Conference on International Conference on Machine Learning - Volume 28*, ser. ICML’13. JMLR.org, 2013, pp. III–1–III–9. [Online]. Available: <http://dl.acm.org/citation.cfm?id=3042817.3042937>
- [9] S. Levine and P. Abbeel, “Learning neural network policies with guided policy search under unknown dynamics,” in *NIPS*, 2014, pp. 1071–1079.
- [10] A. Nair, B. McGrew, M. Andrychowicz, W. Zaremba, and P. Abbeel, “Overcoming exploration in reinforcement learning with demonstrations,” in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 6292–6299.
- [11] P. Pastor, H. Hoffmann, T. Asfour, and S. Schaal, “Learning and generalization of motor skills by learning from demonstration,” in *Robotics and Automation, 2009. ICRA’09. IEEE International Conference on*. IEEE, 2009, pp. 763–768.
- [12] K. Mülling, J. Kober, O. Kroemer, and J. Peters, “Learning to select and generalize striking movements in robot table tennis,” vol. 32, no. 3. Sage Publications Sage UK: London, England, 2013, pp. 263–279.
- [13] G. Thomas, M. Chien, A. Tamar, J. A. Ojea, and P. Abbeel, “Learning robotic assembly from CAD,” *CoRR*, vol. abs/1803.07635, 2018. [Online]. Available: <http://arxiv.org/abs/1803.07635>
- [14] Y. Tassa, T. Erez, and E. Todorov, “Synthesis and stabilization of complex behaviors through online trajectory optimization,” in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2012, pp. 4906–4913.
- [15] A. Y. Ng, D. Harada, and S. Russell, “Policy invariance under reward transformations: Theory and application to reward shaping,” in *ICML*, vol. 99, 1999, pp. 278–287.
- [16] G. Konidaris and A. Barto, “Autonomous shaping: Knowledge transfer in reinforcement learning,” in *ICML*. ACM, 2006, pp. 489–496.
- [17] B. Marthi, “Automatic shaping and decomposition of reward functions,” in *ICML*. ACM, 2007, pp. 601–608.
- [18] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra, “Continuous control with deep reinforcement learning,” 2015.
- [19] T. Schaul, J. Quan, I. Antonoglou, and D. Silver, “Prioritized experience replay,” *CoRR*, vol. abs/1511.05952, 2015. [Online]. Available: <http://arxiv.org/abs/1511.05952>
- [20] S. Fujimoto, H. van Hoof, and D. Meger, “Addressing function approximation error in actor-critic methods,” in *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholm, Sweden, July 10-15, 2018*, 2018, pp. 1582–1591. [Online]. Available: <http://proceedings.mlr.press/v80/fujimoto18a.html>
- [21] K. Hauser and V. Ng-Thow-Hing, “Fast smoothing of manipulator trajectories using optimal bounded-acceleration shortcuts,” in *2010 IEEE International Conference on Robotics and Automation*, May 2010, pp. 2493–2498.
- [22] J. Oh, Y. Guo, S. Singh, and H. Lee, “Self-imitation learning,” *CoRR*, vol. abs/1806.05635, 2018. [Online]. Available: <http://arxiv.org/abs/1806.05635>
- [23] “Melfa rv-fr,” <http://www.mitsubishielectric.com/fa/products/rbt/robot/>, accessed: 2019-01-15.
- [24] “Rt toolbox3,” https://eu3a.mitsubishielectric.com/fa/en/products/rbt/robot/rt_toolbox3, accessed: 2019-01-15.
- [25] M. Večerík, T. Hester, J. Scholz, F. Wang, O. Pietquin, B. Piot, N. Heess, T. Rothörl, T. Lampe, and M. Riedmiller, “Leveraging demonstrations for deep reinforcement learning on robotics problems with sparse rewards,” 2017.
- [26] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *CoRR*, vol. abs/1412.6980, 2014. [Online]. Available: <http://arxiv.org/abs/1412.6980>

APPENDIX

A. Curriculum Learning Setting

As written in IV-C, we used curriculum setting to train our RL agents. Let t_{goal} be the number of goals reached in an experiment. When $0 \leq t_{\text{goal}} < 100$, we train RL agents without checking collision with obstacles, and linearly decrease r_{goal} from 0.52 [rad] to 0.26 [rad]. Then, we linearly increase the value of the collision penalty from 0 to -10 in $100 \leq t_{\text{goal}} < 200$.

For self imitation, we set $K = 5$ for the top- K replay buffer, and do self imitation only when the buffer is filled with K episodes. For goal generalization experiments in VII-D, we set $K = 2$ for all different goal settings, and start self imitation when more than 20% of the top- K replay buffer is stored.

B. Training Details

Both the actor and critic networks have two hidden layers with 128 and 64 units for each layer. The hidden layers use the ReLU activation function, and the output layer of the actor uses the tanh activation function, so that an action lies in the range of $[-1, 1]$. We define the maximum step for an episode to be 300, and the agent randomly resets to a reference path with a probability of $p_{\text{reset}} = 0.3$, as described in IV-B. We train our TD3 agent for at most one million steps. Both the actor and the critic perform updating every time an episode finishes, collecting 1,000 samples, with a minibatch of size 100 sampled from a prioritized replay buffer. The prioritized replay buffer consists of 10^5 transitions with fixed $\alpha = 0.6$ and $\beta = 0.4$. For the ADAM optimization algorithm [26], we use learning rates of 0.001 for both the actor and the critic, and the default values from the TensorFlow framework for the other hyperparameters. The target networks are also updated every cycle using a decay coefficient of 0.995. We use a discount factor of $\gamma = 0.9$.