University of Groningen

Accelerating Reinforcement Learning for Reaching using Continuous Curriculum Learning

Luo, Sha; Kasaei, Hamidreza; Schomaker, Lambert

Published in:

Proceedings of 2020 International Joint Conference on Neural Networks (IJCNN)

DOI:

[10.1109/IJCNN48605.2020.9207427](https://doi.org/10.1109/IJCNN48605.2020.9207427)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version

Publisher's PDF, also known as Version of record

Publication date:

2020

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):

Luo, S., Kasaei, H., & Schomaker, L. (2020). Accelerating Reinforcement Learning for Reaching using Continuous Curriculum Learning. In *Proceedings of 2020 International Joint Conference on Neural Networks (IJCNN)* [9207427] IEEE. <https://doi.org/10.1109/IJCNN48605.2020.9207427>

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

Accelerating Reinforcement Learning for Reaching Using Continuous Curriculum Learning

Sha Luo, Hamidreza Kasaei, Lambert Schomaker

Bernoulli Institute

University of Groningen, The Netherlands

{s.luo, hamidreza.kasaei, l.r.b.schomaker}@rug.nl

Abstract—Reinforcement learning has shown great promise in the training of robot behavior due to the sequential decision making characteristics. However, the required enormous amount of interactive and informative training data provides the major stumbling block for progress. In this study, we focus on accelerating reinforcement learning (RL) training and improving the performance of multi-goal reaching tasks. Specifically, we propose a precision-based continuous curriculum learning (PCCL) method in which the requirements are gradually adjusted during the training process, instead of fixing the parameter in a static schedule. To this end, we explore various continuous curriculum strategies for controlling a training process. This approach is tested using a Universal Robot 5e in both simulation and real-world multi-goal reach experiments. Experimental results support the hypothesis that a static training schedule is suboptimal, and using an appropriate decay function for curriculum learning provides superior results in a faster way.

I. INTRODUCTION

In recent years, reinforcement learning (RL) has attracted growing interest in decision making domains, and achieved notable success in robotic control policy learning tasks [1] [2] [3] [4]. Reasons for this include RL requiring minimal engineering expertise to build complex robot models, and its framework for learning from interaction with the environment is suitable for training on sequential tasks. However, a low training efficiency is the primary challenge preventing reinforcement learning from being widely adopted in real robots. RL algorithms rely on experience that contains informative rewards, but random exploration methods usually only provide sparse rewards. Thus, the agent needs to acquire more experience to extract sufficient amounts of useful data, which results in long training time for a target policy. It is not practicable to train a robot for a long time due to safety factors. Then, in this study, we focus on improving the training efficiency of reinforcement learning for multi-goal reaching tasks¹ in which the target is randomly generated in the workspace; an example of the reaching process is illustrated in Fig. 1.

We are motivated by the concept of successive approximation, introduced by Skinner [5] [6]. He proposed this theory based on training a pigeon to bowl. This appeared to be a very tedious task because the desired behavior is far from the natural behavior of pigeons. Instead of rewarding the target

¹The reaching task in this study is defined as controlling the end effector to reach a pose: the positions in Cartesian space and the orientations in Euler angles.

Fig. 1: Three steps of the reaching task for UR5e and qbhand. The yellow ball with an attached coordinate frame represents the target pose, while the reference frame on the palm represents the pose of the hand. (a) The initial pose of the arm. (b-c) Robot postures after the first and second steps bringing the robot's hand closer to the target pose. (d) The final success step: when the distance of the positions and orientations between the hand and target are within the required precision, the step is regarded as a successful move. The target pose is computed using forward kinematics based on randomly selected joint angles within stipulated joint limits. The number of steps required for a pose reach may vary in different target poses.

behavior directly, they found that by positively reinforcing any behavior that is approximately in the direction of the desired final goal, the pigeon became a champion squash player within few minutes. From this example, we see that learning in animals is facilitated by appropriate guidance. They do not learn only from trial and error. The learning process is much faster through a suitable successive reward schedule with diverse guidance. The learned skills in easy early tasks can help them perform better in the later, more difficult situations. This concept has resurfaced in machine learning: curriculum learning (CL) [7], where the main idea is to train an agent to perform a defined sequence of basic tasks referred to as curriculum, and to increase the performance of agents and speed up the learning process for the final task.

For reaching tasks in robots, different levels of precision requirements can be seen as a curriculum, making the problem more or less difficult as designed. Thus, the scientific research question is how to design the curriculum to enable efficiency, optimally. A poorly designed curriculum can also drive the learning process in the wrong direction, which should be avoided. In this study, we consider the required precision as a continuous parameter coupled with the learning process and propose precision-based continuous curriculum learning (PCCL) to improve the training efficiency and performance of multi-goal reach by a UR5e robot arm. At the beginning of the learning process, the reach accuracy requirement is loose to allow the robot to gain more rewards for acquiring basic skills required to realize the *ultimate target*. To this end, the precision-based process is formulated with a cooling function that automatically selects the accuracy (ϵ) based on the current stage of training.

The main contributions of this study are as follows: First, we propose a PCCL method for accelerating the RL training process. The experimental results showing that the proposed approach accelerates the training process and improves performance in both sparse reward and dense reward scenarios of multi-goal reach tasks. Secondly, the proposed PCCL is easily implemented in off-policy reinforcement learning algorithms. A similar average success rate in real world experiments demonstrate the feasibility of our method for physical applications without requiring fine-tuning when going from simulation to real-world applications. Finally, the robot's reaching movements in grasp and push tasks show that the learned reach policy is beneficial for advanced real robotic manipulations, which can pave the way for complex RL-based robotic manipulations.

II. RELATED WORK

The main goal of this study is to improve the training efficiency of reinforcement learning using curriculum learning. Although an extensive survey of efficient approaches for reinforcement learning in robotic control is beyond the scope of this study, we will review a few recent efforts.

Recently, various research groups have made substantial progress towards the development of continuous reinforcement learning for robot control. The long training durations for reinforcement learning algorithms demonstrate that it is difficult to train an agent to behave as expected using pure trial and error mechanisms [8] [9] [10] [11]. Much effort has been invested on the exploration and exploitation issue. One of the conventional methods for accelerating the learning of a control policy is learning from demonstration: Ashvin et al. [12] coupled deep deterministic policy gradients (DDPG) with demonstrations and showcased a method that can reduce the required training time for convergence by giving nonzero rewards in the early training stages. Another popular technique is prioritized experience replay [13]: shaping the experience by sampling more useful data. In this way, the exploitation ability to get more informative data for training is improved. Nachum et al. [14] showed how hierarchical reinforcement

learning (HRL) can train the robot to push and reach using a few million samples and get the same performance as a robot trained for several days.

For methods based on curriculum learning, hindsight experience replay (HER) [9] is one of the popular methods and can be considered a form of implicit curriculum that improves learning efficiency through learning from fatal experiences. It allows for skipping complicated reward engineering. Instead of judging a behavior to be a fatal miss, it is classified as a successful attempt to reach another target. This approach of learning from a fatal experience can be viewed as a curriculum that trains the easier-met goals first. Kerzel et al. [15] accelerated reach-for-grasp learning by dynamically adjusting the target size. When the robot has little knowledge of the task, the Kerzel et al. approach increases the target size to make it easier to succeed and get positive rewards for learning the policy. Fournier et al. [16] proposed a competence progress based curriculum learning method for automatically choosing the accuracy for a two-joint arm reach task.

The approaches used in the Kerzel et al. [15] and Fournier et al. [16] studies have similarities with the approach proposed in this study, including employing a *start-from-simple* strategy to accelerate the learning process, which can be categorized as curriculum learning. However, they use two-joint planar arm for a position reach environment, which is difficult to adjust into a more complicated environment. Furthermore, Kerzel's task simplification method deals only with discrete reach positions. Fournier's approach requires massive computation resources for multiple accuracy competence progress evaluations. Furthermore, it uses a discrete CL, and switching *curriculums* between pre-defined accuracies poses the risk of instability when training for a large dimensional task.

The PCCL method we propose is based on the continuous curriculum learning. By changing the precision requirement (ϵ) up to every epoch, we add a new task to the *curriculum*, which can smooth the training process and is easy to implement. Compared to the similar methods show in the Kerzel et al. [15] and Fournier et al. [16] studies, our continuous curriculum strategy obviates intensive computation for learning status evaluation and can be utilized in a real robotic arm.

III. BACKGROUND AND METHODOLOGY

Our PCCL algorithm is based on the continuous reinforcement learning mechanism, using the DDPG as the basic framework. Hence, we first introduce the background of RL and DDPG, followed by a description of reaching tasks. Then, the overall PCCL approach is illustrated.

A. Reinforcement Learning

We formalize our RL-based reaching problem as a Markov decision process (MDP). This process can be modeled as a tuple (S, A, P, R, γ) , where S and A denote the finite state and action space respectively, P is the transition probability matrix $P := P_{ss'}^a = \mathbb{P}[s_{t+1} = s' | s_t = s, a_t = a]$ representing the probability distribution to reach a state, given the current state and taking a specific action, R is the state-action associated

reward space and can be represented as r_t , given time t , and $\gamma \in [0, 1]$ is the discounted factor. The goal of the robot is to learn a policy π by interacting with the environment. The deterministic policy can be represented as $\pi(a_t|s_t)$, which guides the agent to choose action a_t under state s_t at a given time step t to maximize the future γ -discounted expected return $R_t = \mathbb{E}[\sum_{i=0}^{\infty} \gamma^i r_{t+i+1}]$.

B. Deep Deterministic Policy Gradients (DDPG)

Because of the consideration of continuous state and action values, we restrict our interest to continuous control. Popular continuous control RL algorithms are: Deep Q-Network (DQN) [17], Continuous Actor Critic Learning Automation (CACLAL) algorithm [18], and the Trust Region Policy Optimization (TRPO) [19], Proximal Policy Optimization (PPO) [20] and Deep Deterministic Deep Policy Gradient (DDPG) algorithms [21]. We selected DDPG as our algorithm platform, which is an RL algorithm based on actor-critic architecture. The actor is used to learn the action policy and a critic is applied to evaluate how good the policy is. The learning process involves improving the evaluation performance of the critic and updating the policy following the direction of getting higher state-action values. This algorithm can be seen as a combination of Q-learning and actor-critic learning because it uses neural networks as the function approximator for both actor and critic (θ^Q, θ^π). To solve the instability problems of two networks and the data correlation in experiences, it uses a replay buffer and target networks ($\theta^{Q'}, \theta^{\pi'}$).

The differentiable loss for the critic is based on the state-action value function:

$$Loss(critic) = Q(s_t, a_t|\theta^Q) - y_t, \quad (1)$$

where Q is the output of the critic network that satisfies the Bellman equation, yielding the expected state-action value; and y_t is the real value:

$$y_t = r_t + \gamma * Q(s_{t+1}, a_{t+1}|\theta^{\pi'}). \quad (2)$$

During the training process, the parameters of these two networks can be updated as follows:

$$\theta^Q \leftarrow \theta^Q - \mu_Q \cdot \nabla_{\theta^Q} Loss(critic), \quad (3)$$

$$\theta^\pi \leftarrow \theta^\pi - \mu_\pi \cdot \nabla_a Q(s_t, \pi(s_t|\theta^\pi)|\theta^Q) \cdot \nabla_{\theta^\pi} \pi(s_t|\theta^\pi), \quad (4)$$

where the symbol ∇ denotes gradients, μ_Q and μ_π represent the learning rate of these two networks, respectively. After the update of the actor and critic, the target networks' weights are merely the copy of the actor and critic networks' weights.

There are several reasons behind the choice of this algorithm: First, it works in both the continuous state and the action space, which eliminates the need for the action discretization in the use of DQN (the previously dominant method for continuous state space). Secondly, it is a deterministic algorithm, which is beneficial for robotics domains as the learned policy can be easily verified with reproduction once the policy has been converged, compared with stochastic algorithms such as TRPO and PPO. Thirdly, DDPG uses the

Fig. 2: DDPG in a UVFAs framework

experience replay buffer to improve stability by decorrelating transitions. This technique can be modified to use additional sources of experience, such as expert demonstrations [10] [12], which has been shown to be effective for improving the training process.

C. Reach Task

The reach task is a fundamental component in robot manipulation that requires the end effector to reach to the target pose while satisfying Cartesian space constraints for various manipulation tasks. We aim to use reinforcement learning to train a UR5e on a multi-goal reach task, in which a value function (policy) is trained to learn how to map the current states into the robot's next action to reach the target goal. An example of a trained reach process is presented in Fig. 1, where the UR5e starts from a fixed initial pose and follows the policy in taking actions to reach the goal pose.

A large body of recent studies has focused on position-only reaching [22] or scenarios where the end effector is forced to reach the targets from one direction, usually top vertical [23]. Such simplification can greatly simplify the tasks but is difficult to generalize to interact with restricted objects. The skills learned for reaching are not easily used for further manipulations which normally require orientation constraints. We move one step forward by considering the reach orientation. For the sake of simplicity, we couple the 4th joint angles with the 2nd and 3rd joints by always setting the 4th joint angles vertical to the ground. This is supported by the feature of the UR5e in that the 2nd, 3rd, and 4th joints are in the same plane. In this way, we also decrease the orientation dimension of the end effector from 3 to 2 as the *Roll* axis is fixed. This joint coupling technique simplifies the task significantly while still being useful for right hand

Algorithm 1: PCCL with DDPG based on UVFAs

```
1 Initialize critic network  $\hat{Q}(s, a, g|\theta^{\hat{Q}})$  and actor
  network  $\pi(s, g|\theta^{\pi})$  with weights  $\theta^{\hat{Q}}$  and  $\theta^{\pi}$ 
  randomly;
2 Initialize target critic network  $\hat{Q}'$  and actor network  $\pi'$ 
  weights  $\theta^{\hat{Q}'} = \theta^{\hat{Q}}$  and  $\theta^{\pi'} = \theta^{\pi}$ ;
3 Initialize experience replay buffer  $R$ ;
4 for  $e$  in  $max\_epochs$  do
5   Compute curriculum  $\epsilon$  from equation (11);
6   for  $i$  in  $max\_episodes$  do
7      $s, g = Env.reset()$  ;
8      $g = g.append(\epsilon)$ ;
9     for  $j$  in  $max\_steps$  do
10       $a = \epsilonpsilon\_greedy$  exploration;
11       $next\_s, r, done = Env.step(a, g)$ ;
12      Add( $s, a, next\_s, r, g$ ) into  $R$ ;
13       $s = next\_s$ ;
14      if  $done$  then
15        break;
16   for  $k$  in  $training\_steps$  do
17     Randomly sample  $N$  experience from  $R$ ;
18     Critic and actor network weights update:
      following equations (1)(2)(3)(4) ;
19     Update the target networks;;
20      $\theta^{\hat{Q}'} = \tau\theta^{\hat{Q}} + (1 - \tau)\theta^{\hat{Q}'}$ ;
21      $\theta^{\pi'} = \tau\theta^{\pi} + (1 - \tau)\theta^{\pi'}$ ;
```

oriented reach tasks as it covers almost all of the reaching directions from the right side and the top.

In this study, we formulate the reach task as an MDP process with definitions of the *state*, *action*, *goal* and *reward*:

State: The current pose of the end effector and joint angles, as defined in Eq. (5). The pose is composed of positions and orientations of the end effector in which we set the position with meters as the unit and the orientation with the *roll-pitch-yaw* order of the Euler angles in radian units. The angles of the joints $j = [j_1, j_2, j_3, j_4, j_5, j_6]$, which are also in radian units.

$$s = [ee_x, ee_y, ee_z, ee_{Rx}, ee_{Ry}, ee_{Rz}, j], \quad (5)$$

Action: We consider the joints' increments as the action (defined as a in Eq. (6)), in which the a_i represents the normalized joint angle change for the i th joint. The angles of the 4th joint is related to a_2 and a_3 to ensure its always vertical to the ground.

$$a = [a_i]; \quad i \in \{1, \dots, 6\} \quad a_i \in [-1, 1],$$
$$a_4 = \begin{cases} -\pi - a_2 & \text{if } \left| \frac{\pi}{2} + a_2 \right| + a_3 \geq \pi, \\ -a_3 - a_2 & \text{otherwise.} \end{cases} \quad (6)$$

For stable training, we bound the action value within the range of -1 and 1. The maximum amount of change for the

joint value is $\pi/6$. Thus, the action of -1 means decreasing the joint value by $\pi/6$, while 1 means increasing the joint value by $\pi/6$. Furthermore, the final computed actions are bounded by the joints limitations.

Goal: The goal is derived by uniformly sampling joint angles within the joint's limitations. Based on these joint angles, setting the forward kinematics result (end effector's pose) in Cartesian space as the goal can ensure a reachable target pose.

$$goal = [ee_x, ee_y, ee_z, ee_{Rx}, ee_{Ry}, ee_{Rz}]. \quad (7)$$

Reward: Designing the reward function is a key to the success of RL. We compare our algorithm based on two types of reward: dense and sparse reward functions, as they are commonly used for reaching tasks. Thus, we give both definitions as follows:

For the dense reward, the reward is given based on the Euclidean distance between the end effector and the goal pose. If both the position distance $\text{dist}(p)$ and the orientation distance $\text{dist}(o)$ are smaller than the required precision (ϵ), we consider it a successful action and set the reward as 1. Otherwise, the reward is set to penalize the distance, which is a weighted summation of the position distance and the orientation distance.

$$\text{dist}(ee, goal) = \alpha \cdot \text{dist}(p) + \beta \cdot \text{dist}(o), \quad (8)$$

where $\alpha, \beta \geq 0$ and $\alpha + \beta = 1$. They are representing the weight factors of the position distance and orientation distance, respectively. In this study, these two distances are within the near range, and we set both weight factors to 0.5.

$$r = \begin{cases} 1 & \text{if } \epsilon < \text{dist}(p) \text{ and } \epsilon < \text{dist}(o), \\ -\text{dist}(ee, goal) & \text{otherwise.} \end{cases} \quad (9)$$

For the sparse reward condition, a successful move is rewarded by 1, otherwise, 0.02 punishment is given for the energy consumption.

$$r = \begin{cases} -0.02 & \epsilon < \text{dist}(ee, goal), \\ 1 & \epsilon \geq \text{dist}(ee, goal). \end{cases} \quad (10)$$

D. Precision-based Continuous Curriculum Learning (PCCL)

Curriculum learning is a technique that focuses on improving the training efficiency on difficult tasks by learning simple tasks first [7]. Most of the previous studies in this field are focused on defining discrete sub-tasks as the *curriculum* for agents to learn. For our multi-degree-of-freedom (DoF) reach environment, the shift between discrete precision can easily result in an unstable performance as the skills learned in the previous tasks also require time to adjust to the new, stricter tasks. However, the continuous CL can smooth the learning by continuously shrinking the precision.

In this section, we focus on a precision-based continuous curriculum learning (PCCL) method, which uses a continuous function to change the curriculum (required precision). The scientific challenge arises how to design a function that works

as the *generator* of the curriculum, just like the job of a teacher to choose the practical level of curriculum for the students.

Framework Definition: A detailed mathematical description of the discrete CL framework has been formulated in the study by Narvekar et al. [24]. Bassich et al. [25] extended it with decay functions as the curriculum generator for continuous CL. We formulate our continuous curriculum learning model based on their work. The task domain is described as D , which is associated with a vector of parameters $F^t = [F_0, \dots, F_n]$. The curriculum *generator* is described as τ , which is used to change the curriculum $M_t = \tau(D, F^t)$ [24]. An important rule when designing the *generator* is that the difficulty should increase monotonically. We refer to the difficulty of the curriculum in the environment as O . Thus, it should satisfy the constraint: $O(M_{t-1}) \leq O(M_t)$.

Compared with discrete CL, the distinguishing feature of continuous CL is the parameter t . It can be referred to as an episode, which means changing the curriculum up to every episode. In this study, we refer to the parameter t as the epoch number for updating the policy.

For the decay function, Bassich et al. [25] categorized it into two classes: fixed decay and adaptive decay functions, based on if it depends on the performance of the agent when deriving the curriculum. The fixed decay functions have the benefit of being easily implemented with only a few parameters to be considered, with the downside being that it lacks flexibility. The adaptive decay functions may generalize to different configurations, but they take up much time evaluating the performance to get feedback for computing the suitable curriculum, especially for complex high-dimensional environments. Furthermore, designing the adaptive decay function requires prior expert knowledge. Thus, we choose the fixed decay function for generating the training curriculum. Once the parameters of the fixed decay function are well designed at the outset, the policy can enjoy the advantages continuously. We formulate the decay function based on a cooling schedule [26] to generate the curriculum. It contains four parameters: start and end precision (e_0 and e_m respectively), the number of epochs that a decay function should experience (s), and the monotonic reduction slope of the decay (α). The decay function is formalized as a power function (11):

$$\epsilon = e_m + \left(\frac{s - k}{s} \right)^\alpha (e_m - e_0), \quad (11)$$

where $k \in [0, s]$ represents the current epoch number, s is the total number of epochs for the entire precision reduction process, and $\alpha \in (0, \infty)$ is the slope of the decay function. The precision (ϵ) of the training process decreases gradually following the equation (11). If α is smaller than 1, the initial decaying process is slower with a smaller slope α . The decay function is a linear function when α is equal to 1. If α is larger than 1, then the larger of the slope, the faster the decay of the initial part of the precision behaves.

The implicit problem with using CL in the precision (ϵ) for a reaching task is that it can violate the MDP structure of the reach problem by changing the requirements for a

successful reach. In this way, the state can be terminated or under different precision requirements, leading to fake updates in the RL algorithm. One of the methods to solve this problem is using universal value function approximators (UVFAs), as proposed by Schaul et al. [27]. It formalizes a value function to generalize both states and goals. In our multi-goals reach environment, the goals are generated randomly in the entire workspace, which means the training policy not only need to consider the state but also the goals. Furthermore, in different stages of the training process of our proposed PCCL method, the precision level for goals are different, which makes the model perfect to be used with the architecture of the UVFAs. Under the architecture of the UVFAs, the goal of the environment is composed of the pose of the targets, and the precision of a successful reach. Therefore, the inputs for the actor and critic in DDPG are being extended with the renewed goal, updating those fake experience to the right one that meets the MDP feature. The framework of the UVFAs based DDPG algorithm can be found in Fig. 2. We also illustrated the pseudo-code of our proposed algorithm in Algorithm. 1, in which the actor π , critic Q and their target networks (π' , critic Q') are embedded into the UVFAs framework, including the goal as input.

IV. EXPERIMENTS

In this section, we first introduce the environmental setup and baseline. Because there are no standard environments for multi-goal reach, we create our environment using the Gazebo simulator and setup a Python interface between the RL algorithm and the agent environment. We then compare the performance and efficiency of our proposed PCCL approach with that of the vanilla DDPG in the context of various reach tasks performed by a UR5e robot arm in both simulation and real world environments.

A. Environmental Setup

All of the experiments are conducted in the Ubuntu 16.04 system and ROS 'Melodic' package. Gazebo 9.0 is used for simulation. The system is encapsulated into a Singularity container², which can be easily reproduced with a bootstrap file in any environment. For the parameter search process, the university's HPC facility with CPU nodes was used.

On the algorithm side, a UVFAs based multi-goal DDPG is used as the baseline, in which the robot is told what to do by using an additional goal as input. The pipeline of the algorithm is illustrated in Fig. 2. The input layer of the actor is composed of state and goal. The output for the actor is the joint angles increments for all of the six joints. For the critic, the input layer is composed of current state, action, and goal, while the output is the Q value for the actions given the state and goal. Both of the actor and critic networks are equipped with three hidden layers with the size of (512, 256, 64), followed by the same ReLU [28] activation function. For the output layer, the critic network does not use an activation function to keep the real

²<https://sylabs.io/guides/3.4/user-guide/>

Q value (state value estimation) unbounded, while the policy network uses the tanh activation function. The DDPG hyperparameters used in the experiments are listed in Table I. For the hyper-parameters in PCCL, the start curriculum means the starting reach precision for the position and orientation. For the dense reward setting, we set the starting precision (ϵ) to 0.15 (units for the position: m, units for the orientation: radian). For the sparse reward setting, it is harder to get positive rewards. Through the experiment, we find out that 0.25 is the start precision at which something can be learned using the sparse reward setting. We, therefore, enlarge the start precision to 0.25 for the sparse reward setting. Furthermore, because the sparse reward environment is sensitive to the change of the precision (ϵ), we design the curriculum changing in a much slower way with the curriculum learning period (s) of 2.5K epochs compared to 1.0K with dense reward to reduce to the final precision. In both of the settings, a slope of 0.8 is selected for the decay function as it works well for both rewards settings.

TABLE I: DDPG Training Hyperparameters

Hyperparameters	Symbol	Value
Discount factor	γ	0.98
Target update ratio	τ	0.01
Actor learning rate	μ_Q	0.0001
Critic learning rate	μ_π	0.001
Gaussian action noise	σ	0.1
Replay buffer size	B	5e+6
Batch size	N	128
Epochs of training	E	3K
Episodes per epoch	M	10
Steps per episode	T	100
Training steps per epoch	K	64
Exploration method	*	ϵ_{greedy}

B. Simulation Experiments

We evaluate the training efficiency and performance of the proposed approach against the vanilla DDPG algorithm on both sparse reward and distance-based dense reward scenarios.

At the beginning of each episode in training, the robot is initialized to have the same joint angles: $[-\frac{\pi}{2}, -\frac{\pi}{2}, 0, \frac{\pi}{2}, \frac{\pi}{2}, -\frac{\pi}{2}]$, and then is trained to reach a target pose that is derived from a randomly selected joints configuration. If the pose distance is within the precision before running out of the maximum steps, this step is regarded as a successful reach and the episode is being terminated. We record the accumulated steps every 10 epochs and evaluate the averaged success rate every 100 epochs over 100 various target poses. The simulation experiment results can be seen in Fig. 4. Our PCCL method is able to improve the performance in both rewards settings, especially for the sparse situations, where PCCL allows the task to be learned at all compared to the multi-goal DDPG baseline (from 0.0% to 75%). We notice that there is a tendency of decreasing performance at the later stage of training; we attribute it to the increased goal difficulty that making the task harder to learn. Besides, from Fig. 4a, we can see that our proposed approach can reduce the number of necessary steps by nearly $1.0e + 05$ with less variance in the

dense-reward situation. For reinforcement learning, an efficient policy could learn the skills faster and result in an earlier termination in episodes, thus requires fewer steps during the training period, implying less time on training. Specifically, the total training time is decreased by 19.9% in the dense-reward condition, which is summarized in Table II. In the sparse-reward setting, the training time of our PCCL method is even less than the dense-reward setting with the normal multi-goal DDPG algorithm.

Once the policy has been learned in simulation, we evaluate the average success rate in both simulated and physical UR5e robots over 100 runs. The averaged results are shown in Table III. Based on the obtained results, it is obvious that the policy trained in simulation can be applied to the real robot without any further fine tuning. Furthermore, the reach policy learned in simulation can be tested in the real world with almost the same performance.

TABLE II: Training time comparison

Methods	Total time	
	Dense reward	Sparse reward
DDPG	22h29min6s	52h30min
PCCL-DDPG	18h01min33s	19h41min20s

TABLE III: Performance on trained policy

Methods	Success rate (%)			
	Dense reward		Sparse reward	
	Sim	Real	Sim	Real
DDPG	0.950	0.950	0.000	0.000
PCCL-DDPG	0.982	0.980	0.715	0.705

C. Real World Experiments

We test the manipulation performance on two tasks: pushing and grasping based on the reaching policy learned by our PCCL algorithm. In this experiment, a UR5e arm, fitted with a Qbhand, is used to test the performance of the proposed approach. Towards this goal, the proposed approach is integrated into the RACE robotic system [29] [30]. In particular, the RACE framework provides the pose of active objects as the goal to reach by the arm. The experimental settings for pushing and grasping tasks are shown in Fig. 3.

For the pushing task, a juice box is randomly placed at a reachable position in the workspace. The initial goal is to move the arm to the pre-push area and then push the juice box from point A to point B as shown in Fig. 3a.

For the grasping task, a cup and a basket are randomly put on the table that can be reached by the arm. The initial goal of this task is to move the arm to go to the pre-grasp pose of the cup (i.e., represented as point A in Fig. 3b) using the trained PCCL algorithm. Once it reaches the target pose, the RACE system will take over the responsibility of controlling the hand to grasp the cup. After grasping the object, the robot moves the cup to the top of the basket using the PCCL algorithm (i.e., reach to pose B, as shown in Fig. 3b). Finally, the robot releases the cup into the basket. A video of the experiment

(a) Push scenario

(b) Grasp scenario

Fig. 3: Push and grasp experiments: the proposed algorithm is mainly responsible for moving the robotic arm to the target pose. Object perception and object grasping are done by the RACE system [29] [30]. In the case of push scenario (a), the arm first goes to the pre-push pose, as shown by a red point A. Then, the robot pushes the juice box object from point A to the target pose B. In the case of object grasping scenario (b), the robotic arm first reaches the pre-grasp pose as shown by the red point A, then grasps the cup and manipulates it on top of the basket, as shown by point B.

(a) Average accumulated steps with dense reward.

(b) Average test success rate with dense reward.

(c) Average accumulated steps with sparse reward.

(d) Average test success rate with sparse reward.

Fig. 4: Simulation results on average accumulated steps, i.e., accumulated simulation steps during the learning process, and average success rate: The shaded regions represent the variance over ten runs. The gray dash lines in (b) and (d) indicate the epoch that arriving the final curriculum, after which the PCCL method has the same precision criteria as the baseline.

is available³. These demonstrations show that although the policy is trained on the initial *up* pose, it can be generalized to other different initial states (e.g. pre-push and pre-grasp A poses). Furthermore, the policy learned in simulation can be directly applied on the physical robot. Additionally, the learned skill can be further utilized in more complicated scenarios, as shown by the success of the manipulation tasks. Accelerating the training process of RL for reaching tasks may benefit not only the RL field but the robotics autonomy domains as well.

³see <https://youtu.be/WY-1EbYBSGo>

V. DISCUSSION AND CONCLUSION

We have proposed a precision-based continuous curriculum learning (PCCL) approach to accelerate reinforcement learning for multi-DOF reaching tasks. The technique utilizes the *start-from-simple* rule to design a continuous curriculum *generator*: decay function. The experimental results show that the proposed method can improve the RL algorithm training efficiency and performance, especially for a sparse and binary reward scenario, which fails to learn the task in the baseline environment.

Compared to other curriculum learning based methods on multi-goal reach tasks [15] [16], the proposed approach enables the training of a high-dimensional pose reaching task by deploying a continuous-function based curriculum strategy. Although the predefined decay function lacks general flexibility compared with adaptive functions, it is an efficient way to learn complex tasks and environments. Conversely, the adaptive functions would be required to evaluate the difficulty level of the environment and trained policy status frequently during the learning process, which normally means success rate evaluation or environment analysis. Those evaluation steps entail a heavy computation burden for the agents and contain delayed information for training. Our method can be generalized to different robotic arms and environment. The main limitation is that the difficulty of the situation should be related to the precision requirements.

Note that although the reach task that we consider in this study is simple, it mainly serves as a convenient benchmark to straightforwardly demonstrate the significant improvement of training speed with our proposed PCCL method in high-dimensional continuous environment. By improving the efficiency of training RL algorithm, we, to some extent, pave the way for further investigation of some RL-based robotics applications, because the trained reach policy is a preliminary step towards further real world complex robotics manipulation tasks. As a future work, we plan to take obstacle avoidance into consideration, and explore the possibility of curriculum learning on complex motion planning.

ACKNOWLEDGMENT

We are thankful to Dr. Marco Wiering for his thoughtful suggestions. We also thank Weijia Yao for the fruitful discussion about the experimental design and the useful comments for the paper. Sha Luo is funded by the China Scholarship Council.

REFERENCES

- [1] S. Gu, E. Holly, T. Lillicrap, and S. Levine, "Deep reinforcement learning for robotic manipulation with asynchronous off-policy updates," in *2017 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2017, pp. 3389–3396.
- [2] 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.
- [3] L. Tai, G. Paolo, and M. Liu, "Virtual-to-real deep reinforcement learning: Continuous control of mobile robots for mapless navigation," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017, pp. 31–36.
- [4] G. Lan, M. D. Carlo, F. van Diggelen, J. M. Tomczak, D. M. Roijers, and A. E. Eiben, "Learning directed locomotion in modular robots with evolvable morphologies," *ArXiv*, vol. abs/2001.07804, 2019.
- [5] B. F. Skinner, "Reinforcement today," *American Psychologist*, vol. 13, no. 3, p. 94, 1958.
- [6] G. B. Peterson, "The discovery of shaping: Bf skinner's big surprise," *The Clicker Journal: The Magazine for Animal Trainers*, vol. 43, pp. 6–13, 2000.
- [7] Y. Bengio, J. Louradour, R. Collobert, and J. Weston, "Curriculum learning," in *Proceedings of the 26th annual international conference on machine learning*. ACM, 2009, pp. 41–48.
- [8] M. Plappert, M. Andrychowicz, A. Ray, B. McGrew, B. Baker, G. Powell, J. Schneider, J. Tobin, M. Chociej, P. Welinder *et al.*, "Multi-goal reinforcement learning: Challenging robotics environments and request for research," *arXiv preprint arXiv:1802.09464*, 2018.
- [9] M. Andrychowicz, F. Wolski, A. Ray, J. Schneider, R. Fong, P. Welinder, B. McGrew, J. Tobin, O. P. Abbeel, and W. Zaremba, "Hindsight experience replay," in *Advances in Neural Information Processing Systems*, 2017, pp. 5048–5058.
- [10] T. Jurgenson and A. Tamar, "Harnessing reinforcement learning for neural motion planning," *arXiv preprint arXiv:1906.00214*, 2019.
- [11] A. R. Mahmood, D. Korenkevych, G. Vasan, W. Ma, and J. Bergstra, "Benchmarking reinforcement learning algorithms on real-world robots," *arXiv preprint arXiv:1809.07731*, 2018.
- [12] 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.
- [13] T. Schaul, J. Quan, I. Antonoglou, and D. Silver, "Prioritized experience replay," *arXiv preprint arXiv:1511.05952*, 2015.
- [14] O. Nachum, S. S. Gu, H. Lee, and S. Levine, "Data-efficient hierarchical reinforcement learning," in *Advances in Neural Information Processing Systems*, 2018, pp. 3303–3313.
- [15] M. Kerzel, H. B. Mohammadi, M. A. Zamani, and S. Wermter, "Accelerating deep continuous reinforcement learning through task simplification," in *2018 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 2018, pp. 1–6.
- [16] P. Fournier, O. Sigaud, M. Chetouani, and P.-Y. Oudeyer, "Accuracy-based curriculum learning in deep reinforcement learning," *arXiv preprint arXiv:1806.09614*, 2018.
- [17] 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.
- [18] H. Van Hasselt and M. A. Wiering, "Reinforcement learning in continuous action spaces," in *2007 IEEE International Symposium on Approximate Dynamic Programming and Reinforcement Learning*. IEEE, 2007, pp. 272–279.
- [19] J. Schulman, S. Levine, P. Abbeel, M. Jordan, and P. Moritz, "Trust region policy optimization," in *International conference on machine learning*, 2015, pp. 1889–1897.
- [20] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.
- [21] 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," *arXiv preprint arXiv:1509.02971*, 2015.
- [22] T.-H. Pham, G. De Magistris, and R. Tachibana, "Optlayer-practical constrained optimization for deep reinforcement learning in the real world," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 6236–6243.
- [23] I. Lenz, H. Lee, and A. Saxena, "Deep learning for detecting robotic grasps," *The International Journal of Robotics Research*, vol. 34, no. 4-5, pp. 705–724, 2015.
- [24] S. Narvekar, J. Sinapov, M. Leonetti, and P. Stone, "Source task creation for curriculum learning," in *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 2016, pp. 566–574.
- [25] A. Bassich and D. Kudenko, "Continuous Curriculum Learning for Reinforcement Learning," in *Proceedings of the 2nd Scaling-Up Reinforcement Learning (SURL) Workshop*. IJCAI, 2019.
- [26] L. Schomaker and M. Bulacu, "Automatic writer identification using connected-component contours and edge-based features of uppercase western script," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 26, no. 6, pp. 787–798, 2004.
- [27] T. Schaul, D. Horgan, K. Gregor, and D. Silver, "Universal value function approximators," in *International Conference on Machine Learning*, 2015, pp. 1312–1320.
- [28] A. F. Agarap, "Deep learning using rectified linear units (relu)," *arXiv preprint arXiv:1803.08375*, 2018.
- [29] S. H. Kasaei, M. Oliveira, G. H. Lim, L. S. Lopes, and A. M. Tomé, "Towards lifelong assistive robotics: A tight coupling between object perception and manipulation," *Neurocomputing*, vol. 291, pp. 151–166, 2018.
- [30] S. H. Kasaei, N. Shafii, L. S. Lopes, and A. M. Tomé, "Interactive open-ended object, affordance and grasp learning for robotic manipulation," in *2019 IEEE/RSJ International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 3747–3753.