

TEACH: Temporal Variance-Driven Curriculum for Reinforcement Learning

Gaurav Chaudhary
IIT Kanpur
Kanpur, India
gauravch@iitk.ac.in

Laxmidhar Behera
IIT Kanpur
Kanpur, India
lbehera@iitk.ac.in

ABSTRACT

Reinforcement Learning (RL) has achieved significant success in solving single-goal tasks. However, uniform goal selection often results in sample inefficiency in multi-goal settings where agents must learn a universal goal-conditioned policy. Inspired by the adaptive and structured learning processes observed in biological systems, we propose a novel Student-Teacher learning paradigm with a Temporal Variance-Driven Curriculum to accelerate Goal-Conditioned RL. In this framework, the teacher module dynamically prioritizes goals with the highest temporal variance in the policy’s confidence score, parameterized by the state-action value (Q) function. The teacher provides an adaptive and focused learning signal by targeting these high-uncertainty goals, fostering continual and efficient progress. We establish a theoretical connection between the temporal variance of Q -values and the evolution of the policy, providing insights into the method’s underlying principles. Our approach is algorithm-agnostic and integrates seamlessly with existing RL frameworks. We demonstrate this through evaluation across 11 diverse robotic manipulation and maze navigation tasks. The results show consistent and notable improvements over state-of-the-art curriculum learning and goal-selection methods.

KEYWORDS

Adaptive Learning; Curriculum Learning; Robot Learning

ACM Reference Format:

Gaurav Chaudhary and Laxmidhar Behera. 2026. TEACH: Temporal Variance-Driven Curriculum for Reinforcement Learning. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 10 pages. <https://doi.org/10.65109/ADKV2733>

1 INTRODUCTION

Deep Reinforcement Learning (DRL) [31] has been successfully applied in solving complex problems such as robotic manipulation [21], continuous control [7, 10], flight control [27], intelligent perception system [9], and real-time strategy gameplay [2]. Building upon DRL’s success, learning a generalized policy to solve multiple goal-oriented tasks is of huge interest. In the past, the multi-goal or multi-task [52] problems have been addressed by merging Automatic Curriculum Learning (ACL) [1, 5, 26, 44, 45, 51] with policy learning. Broadly, ACL resorts to presenting goal or task

in increasing order of difficulty [37, 56, 57], and has been successfully applied to Sim2Real transfer [1], continuously-parameterized environments [43], sequencing of components in multi-agent systems [53], and to improve DRL agent’s sample efficiency and performance [16, 22].

This paper focuses on ACL from the perspective of Goal Conditioned Reinforcement Learning (GCRL) [48]. Specifically, we address multi-goal sparse reward scenarios where the agent is trained to learn a universal policy to attain any goal in the goal space, and a binary reward is provided when the desired goal is achieved. In such multi-goal scenarios under ACL, learning is most effective when the agent engages with goals positioned at the skill frontier—neither too easy nor too difficult [38, 50]. Based on this observation, [35, 43] use a curriculum that samples tasks with high Learning Progress (LP), while Value Disagreement Sampling (VDS) [57] samples goals with high epistemic uncertainty. Although ACL has improved their performance, these approaches rely on noisy value estimates to design a curriculum, and these noisy estimates become inefficient in a high-dimensional goal space with sparse reward settings.

Concretely, we are motivated by theoretical analysis on ACL for linear regression models [55] and curriculum design for teachers via demonstration [56]. Our key insight is that relying on uniform goal sampling to update policy and value function neglects the intrinsic dynamics of learning, where certain goals require more focus due to their evolving contributions to policy improvement and value function learning. To address this, we propose a curriculum learning framework based on temporal variance in Q -values, highlighting regions where the learning dynamics are active. Although this curriculum is defined using Q -value estimates, we establish that it inherently captures the co-evolution of both the policy and value function, a connection we formalize mathematically to provide theoretical support.

Specifically, we introduce a Student-Teacher framework for GCRL that utilizes learning progress measured by temporal variance in policy confidence score to guide curriculum design. In this framework, the teacher module (goal proposer) evaluates the agent’s policy and assigns a policy confidence score to each goal in the goal space. This policy confidence score, derived from state-action value (Q) estimates, corresponds to the expected return under the current policy for a given goal. The teacher leverages these confidence scores to dynamically design the curriculum.

The teacher ensures that the curriculum continuously adapts to the agent’s current capabilities and promotes efficient progress. To measure progress, we introduce the concept of Learning Progress (LP), which quantifies the temporal change in policy confidence scores. The teacher evaluates learning progress by monitoring the variance in these scores over time. This variance indicates the goals



This work is licensed under a Creative Commons Attribution International 4.0 License.

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). <https://doi.org/10.65109/ADKV2733>

at the skill frontier, where the agent’s policy is evolving most rapidly. This allows the teacher to focus on regions of the goal space where the agent’s learning is most active, promoting continual and effective policy improvement. The proposed algorithm has no static design heuristic [13, 52] and is ensemble-free [57]. We name our approach **TEACH** (Temporal variance driven Automatic Curriculum teachHer). Further, to demonstrate the effectiveness of TEACH, we compare the proposed algorithm with Proximal Curriculum for Reinforcement Learning (ProCurl) [52] and Self-Paced Context Evaluation (SPaCE) [13] approaches, which use value estimates for curriculum learning in contextual multi-task scenarios. We adapt their strategies to a multi-goal setting. We also compare with the current state-of-the-art VDS [57] in multi-goal GCRL. Finally, we show that TEACH consistently improves performance on different challenging tasks, including robotic manipulation [40], dexterous in-hand manipulation [40], and Maze navigation [57]. We summarize our contributions as follows:

- **Theoretical insights:** We establish a formal theoretical connection between Q-values and policy evolution, demonstrating that changes in Q-values relate to policy divergence and help identify regions of significant policy evolution.
- **Curriculum design:** We propose a curriculum strategy based on Learning Progress (LP) that mitigates the impact of noisy value estimates, promoting efficient learning in multi-goal sparse reward scenarios.
- **Validation and robustness:** We validate our method through extensive experiments across diverse robotic manipulation tasks, demonstrating its effectiveness and robustness across different settings.

2 RELATED WORK

Curriculum learning has been recognized as a critical factor in addressing various machine learning challenges [5, 6, 14, 49], aiming to structure the presentation of samples during the learning process. Although the experts can manually create such curricula tailored to specific problems, the concept of Automatic Curricula Learning [19, 43] concentrates on developing algorithms capable of autonomously arranging the sequences of learning problems to optimize agent performance.

Curriculum strategies for DRL: In DRL, applications of ACL are widespread as researchers often struggle to make generalist agents and search for strategies [11, 47, 57] that can train agents beyond their initial success. Curriculum reinforcement learning techniques [12, 17, 18, 24, 25, 28, 30, 43, 43, 45, 50, 57] primarily concentrate on enhancing an agent’s learning efficiency or effectiveness across a set of tasks. Curriculum Reinforcement Learning has also been successfully extended to sim-to-real transfer by adapting domain randomization [1]. Exploiting world models for curriculum learning have also been explored by [23, 36, 39, 42]. Planning Exploratory Goals (PEG) [23] is a goal-conditioned policy that optimizes directly for goals that would result in high exploratory value trajectories. However, our focus lies on ACL design for model-free GCRL.

Automatic curriculum learning for goal-conditioned RL is an approach for presenting goals to learning agents in a meaningful

order [34, 48]. Hindsight Experience Replay (HER) [3] is an implicit curriculum strategy that relabels unsuccessful trajectory rollout as successful. A more robust extension of it, Curriculum-guided HER (CHER) [15] uses an adaptive relabelling strategy based on diversity and goal-proximity. Combined with HER, [57] proposes a Value Disagreement Sampling (VDS) that uses a goal proposer module that prioritizes goals that maximize the epistemic uncertainty of the value function. A learning-based strategy using a Generative Adversarial Network (GAN) was used by [17], generating goals with intermediate success probability.

Learning progress-based curriculum strategies have also been studied, which sample goals more aggressively toward which agent shows the most progress. A student-teacher framework for a discrete task space that samples tasks with high LP was presented by [35]. To extend this to continuous task space [43] uses a Gaussian Mixture Model (GMM) on a tuple of tasks and absolute learning progress. The tasks are then sampled from a Gaussian chosen proportionally to its mean absolute learning progress. These methods use a dense reward structure and compute LP using either the nearest neighbour [43] or the change in episode return for a given task. This limits their direct applicability to GCRL with binary episodic returns with random initial states. SPaCE [13] uses the values function to design a curriculum learning based on learning progress. However, their work targets contextual RL for discrete task settings. In a similar context, ProCurl [52], inspired by the pedagogical concept of *Zone of Proximal Development* [54], uses value estimates to design a curriculum strategy that sample task that is neither too easy nor too hard.

Existing curriculum RL methods often struggle with adaptively selecting goals that maximize learning progress, relying on static heuristics [13, 52] or computationally expensive uncertainty [57] measures derived from the value function. To address this, we propose a temporal variance-driven curriculum strategy that prioritizes goals based on the variability of the policy learning progress over recent time steps. By capturing the temporal divergence in the policy learning signal, our approach adaptively identifies goals at the edge of the agent’s capability without relying on static heuristics or large ensembles, ensuring efficient and continual progress.

3 FORMAL PROBLEM SETUP

Multi-Goal MDP. We consider a multi-goal RL setting where the objective is to learn a universal policy capable of achieving any goal in a specified goal space. The environment is defined as $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathcal{G}, \mathcal{T}, \gamma, H, \mathcal{R})$, where \mathcal{S} and \mathcal{A} are the state and action spaces, $\mathcal{G} \subseteq \mathcal{S}$ is the goal space, $\mathcal{T} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$ defines the transition dynamics, $\gamma \in [0, 1]$ is the discount factor, H is the episode length, and $\mathcal{R} : \mathcal{S} \times \mathcal{A} \times \mathcal{G} \times \mathcal{S}' \rightarrow \mathbb{R}$ is the goal-conditioned reward function, where $\mathcal{R}(s, a, g, s') = 0$ if $\|s' - g\| < \epsilon$ and -1 otherwise. We model multi-goal RL as an RL problem with an extended state space $\mathcal{S} \times \mathcal{G}$, where the policy $\pi : \mathcal{S} \times \mathcal{G} \rightarrow \mathcal{A}$ selects actions conditioned on state and goal, and the Q-function $Q : \mathcal{S} \times \mathcal{G} \times \mathcal{A} \rightarrow \mathbb{R}$ represents the goal conditioned expected return. The agent optimizes the expected discounted return over goals sampled from \mathcal{G} :

$$J(\theta) = \mathbb{E}_{g \sim \mathcal{G}} \left[\mathbb{E}_{\tau \sim \pi_{\theta_t}(\cdot | s, g)} \sum_{h=0}^{H-1} \gamma^h \mathcal{R}(s_h, a_h, g, s'_h) \right]. \quad (1)$$

Algorithm 1 Interplay of Teacher-Student components in RL agent training

- 1: **Initialization:** Initialize RL policy parametrized by θ , replay buffer R , goal space \mathcal{G} .
 - 2: **for** each episode $e = 1, 2, \dots$ **do**
 - 3: Teacher (goal-proposer) picks a goal g from goal space \mathcal{G} .
 - 4: Student (RL agent) generates episodic rollout trajectory τ_e aiming to reach the proposed goal g .
 - 5: Update the parameterized RL agent using off-policy updates with replay buffer R .
 - 6: **end for**
-

Where $\tau = \{(s_h, a_h)\}_{h=0}^{H-1}$ is a trajectory generated by π_{θ_t} conditioned on g .

Curriculum Learning. Further, we extended this multi-goal RL setup to a student-teacher paradigm. The agent’s performance over goal g , which is uniformly sampled from goal space \mathcal{G} can be defined as a policy confidence score $C^{\pi_{\theta_t}}(g) = Q^{\pi_{\theta_t}}(s, g, a)$. In the context of our problem, we treat the policy as the student component and the goal proposer unit as the teacher component. The student component is updated using the goal-conditioned transitions sampled from the replay buffer R . On the other hand, the teacher component defines a curriculum over the goal space to improve the student’s performance. This work focuses on the design of a teacher (curriculum) that can guide the student to learn in a sample-efficient manner. The student-teacher interplay happens at the start of an episode, i.e., the teacher samples a target goal from the goal space for that particular episodic rollout. We summarize this student-teacher interplay in Algorithm 1.

4 TEMPORAL VARIANCE DRIVEN CURRICULUM DESIGN

This section introduces a novel curriculum learning approach that leverages policy confidence scores to evaluate learning progress. Existing methods primarily rely on value estimate-based metrics [13, 52, 57] to design curricula, which have shown promise in multi-task RL. However, value estimates are often noisy [32, 46] because they are trained using small, random batches of data collected from rollouts generated by the agent, which is still learning. One potential way to mitigate this issue is to use Polyak averaging [41] to obtain smoother value functions (refer to Section 5.4). Nevertheless, we argue that the effect of noisy values becomes more pronounced when employing static heuristic-based strategies [13, 52].

To address this, we propose leveraging the temporal evolution of the Q-function. We posit that even when value estimates are noisy, their temporal evolution can reveal a clear direction of improvement and suppress noise in curriculum learning. Further, the temporal variance in Q-values inherently captures the co-evolution of both the policy and value function, making it more informative and less susceptible to noisy value estimates. By focusing on this temporal signal, our strategy provides a more accurate representation of the agent’s learning progress, ultimately enabling the generation of a more effective learning curriculum. To this end, we first establish the theoretical connection between Q-value and policy evolution. Next, we introduce the policy confidence score, a formal metric

to quantify the performance of the policy, followed by curriculum design.

4.1 Theoretical Connection between Q-value and Policy Evolution

The Q-function, $Q^{\pi_{\theta_t}}(s, g, a) = \mathbb{E} \left[\sum_{h=0}^{H-1} \gamma^h \mathcal{R}(s_h, a_h, g, s'_h) \mid s_0 = s, a_0 = a, \pi_{\theta_t} \right]$, represents the expected return from a given state-action pair under policy π_{θ_t} for goal g . Changes in the policy π_{θ_t} directly influence the Q-values through the Bellman operator [4]:

$$Q^{\pi_{\theta_t}}(s, g, a) = \mathcal{R}(s, a, g, s') + \gamma \mathbb{E}_{s' \sim \mathcal{T}(\cdot | s, a)} \left[\mathbb{E}_{a' \sim \pi_{\theta_t}(\cdot | s', g)} Q^{\pi_{\theta_t}}(s', g, a') \right]. \quad (2)$$

As the policy evolves and improves, the state-action visitation distribution shifts, causing the Q-function to update and converge toward a new fixed point. In practical RL, Q-learning serves as a proxy for policy iteration, where temporal variance in Q-values reflects underlying changes in the policy. To formalize this, we consider a soft policy update mechanism [20]. While our training uses a deterministic policy via DDPG [33], we employ a soft policy update model to theoretically analyze the relationship between Q-value variance and policy evolution, as it provides a tractable approximation for policy divergence.

$$\pi_{\theta_t}(a | s, g) \propto \exp \left(\frac{Q^{\pi_{\theta_t}}(s, g, a)}{\alpha} \right), \quad (3)$$

where $\alpha > 0$ is a temperature parameter controlling exploration. The Kullback-Leibler (KL) divergence [29] between consecutive policies is:

$$\begin{aligned} \text{KL}(\pi_{\theta_{t+1}} \parallel \pi_{\theta_t}) &= \mathbb{E}_{s \sim \mathcal{D}, g \sim \mathcal{G}} \left[\sum_a \pi_{\theta_{t+1}}(a | s, g) \log \frac{\pi_{\theta_{t+1}}(a | s, g)}{\pi_{\theta_t}(a | s, g)} \right]. \end{aligned} \quad (4)$$

Expanding this using Q-values (see extended preprint [8] for complete proof) yields:

$$\begin{aligned} \text{KL}(\pi_{\theta_{t+1}} \parallel \pi_{\theta_t}) &\approx \frac{1}{2\alpha^2} \mathbb{E}_{s \sim \mathcal{D}, g \sim \mathcal{G}} \left[\text{Var}_{a \sim \pi_{\theta_t}(\cdot | s, g)} (\Delta Q^{\pi_{\theta_t}}(s, g, a)) \right], \end{aligned} \quad (5)$$

where $\Delta Q^{\pi_{\theta_t}}(s, g, a) = Q^{\pi_{\theta_{t+1}}}(s, g, a) - Q^{\pi_{\theta_t}}(s, g, a)$. This approximation shows that the KL divergence between successive policies is proportional to the variance of Q-value changes. This indicates that high Q-value variance signals significant policy evolution. By tracking this variance, we gain a dual perspective on value and policy dynamics, informing our curriculum design to focus on goals where learning is most active.

In Section 4.2, we introduce the policy confidence score, a Q-value-derived metric to quantify policy performance and guide curriculum updates.

Algorithm 2 Temporal Variance-Driven Curriculum Design

```

1: Initialization: Initialize RL policy parametrized by  $\theta$ , replay
   buffer  $R$ , goal space  $\mathcal{G}$ , temporal window  $n$  (for variance calcu-
   lation), interplay frequency  $\Delta$  (measured in episodes), and
   episode length  $H$ , training timesteps  $T$ .
2: Sample  $N$  goals from goal space  $\mathcal{G}$ 
3: for  $t = 0, 1, 2 \dots T$  do
4:   if  $t \bmod (\Delta \cdot H) = 0$  then
5:     Compute  $LP^\pi(g, t)$  over past  $n$  timesteps using Equa-
     tion (10) for each goal and sample target goal
6:   else
7:     Use the previous  $LP^\pi(g, t)$  to sample target goal
8:   end if
9:   Rollout goal-conditioned transition  $(s_t, a_t, r_t, s_{t+1}, g)$  and
     store in  $R$ 
10:  Update policy parameter  $\theta$  using stored experience
11: end for

```

4.2 Policy Confidence Score

Section 4.1 establishes that Q-value variance governs policy evolution, Equation (5). We use these insights to design a metric to assess policy performance in TEACH’s goal-conditioned RL. To translate this insight into practice, we adopt DDPG [33], an actor-critic method that updates the policy $\pi_{\theta_t}(s, g)$ that produces actions, augmented with exploration noise (e.g., Ornstein-Uhlenbeck) to mimic the Q-value-driven updates of a soft policy.

In DDPG, the policy parameters θ are updated to maximize $J(\theta)$ via the deterministic policy gradient:

$$\nabla_{\theta} J(\theta) \approx \mathbb{E}_{s \sim \mathcal{D}, g \sim \mathcal{G}} \left[\nabla_{\theta} \pi_{\theta_t}(s, g) \nabla_a Q^{\pi_{\theta_t}}(s, g, a) \Big|_{a=\pi_{\theta_t}(s, g)} \right], \quad (6)$$

where \mathcal{D} is the replay buffer state distribution. This gradient demonstrates that policy updates are guided by Q-value gradients, aligning with Section 4.1’s insight. Leveraging the Q-function’s significance, we define the *policy confidence score* for a goal g as:

$$C^{\pi_{\theta_t}}(g) = \mathbb{E}_{s \sim \mathcal{D}} \left[Q^{\pi_{\theta_t}}(s, g, \pi_{\theta_t}(s, g)) \right]. \quad (7)$$

This metric aggregates Q-values across states to estimate the expected return for an actor-critic policy $\pi_{\theta_t}(s, g)$, the term $Q^{\pi_{\theta_t}}(s, g, \pi_{\theta_t}(s, g))$ represents the expected return when starting from state s and pursuing goal g under π_{θ_t} . By taking the expectation over states, $C^{\pi_{\theta_t}}(g)$ yields a scalar metric that reflects the policy’s average performance for goal g , independent of specific state-action pairs.

The policy confidence score $C^{\pi_{\theta_t}}(g)$ provides a measure to evaluate the policy’s capability for each goal. This enables a curriculum that targets goals with significant learning potential, as detailed in Section 4.3. By leveraging Q-values updated through DDPG’s gradient, $C^{\pi_{\theta_t}}(g)$ captures the variance-driven progress central to TEACH’s design.

4.3 Curriculum Design

Extending Sections 4.1 and 4.2, we propose a curriculum that prioritizes goals $g \in \mathcal{G}$ to accelerate policy improvement in goal-conditioned RL. Building on the Q-value variance insight Equation (5) and policy confidence score Equation (7), we define a curriculum distribution $\mathcal{K}^\pi(g) \propto \delta_{C^\pi}(g, t)$, where temporal variance encodes learning progress for goal g at timestep t . This approach refines prior work like VDS [57] and SPaCE [13], which use epistemic uncertainty or value differences.

Following Section 4.2, we use the policy confidence score $C^{\pi_{\theta_t}}(g) = \mathbb{E}_{s \sim \mathcal{D}} \left[Q^{\pi_{\theta_t}}(s, g, \pi_{\theta_t}(s, g)) \right]$ to compute temporal variance for curriculum design. This score, driven by Q-value gradients in TEACH’s DDPG framework [33], reflects the policy’s performance for goal g .

Learning Progress via Temporal Variance. Per Section 4.1, high Q-value variance signals policy evolution. We extend this to define learning progress via the temporal variance of the policy confidence score. A simple difference like $C^{\pi_{\theta_t}}(g) - C^{\pi_{\theta_{t-\Delta t}}}(g)$ over a fixed interval Δt could be noisy due to unreliable Q-value estimates in reinforcement learning [32, 46]. Instead, we propose a robust metric: the temporal variance of $C^{\pi_{\theta_t}}(g)$ over a window of n timesteps, defined as:

$$\delta_{C^\pi}(g, t) = \frac{1}{n} \sum_{k=t-n+1}^t (C^{\pi_{\theta_k}}(g) - \bar{C}^\pi(g, t))^2, \quad (8)$$

where the mean confidence score over the window is defined as:

$$\bar{C}^\pi(g, t) = \frac{1}{n} \sum_{k=t-n+1}^t C^{\pi_{\theta_k}}(g).$$

This variance mitigates noisy Q-value estimates, as averaging over n timesteps filters fluctuations, enhancing robustness in TEACH’s curriculum design strategy. The learning progress for goal g at timestep t is then:

$$LP^\pi(g, t) = \frac{\delta_{C^\pi}(g, t)}{Z_t}, \quad (9)$$

where $Z_t = \int_{\mathcal{G}} \delta_{C^\pi}(g, t) dg$ normalizes the distribution over the goal space \mathcal{G} .

Since \mathcal{G} is continuous, computing Z_t directly is intractable. We address this by uniformly sampling N goals $\{g_1, g_2, \dots, g_N\} \subset \mathcal{G}$ at the start of training. For each sampled goal g_i , we compute $\delta_{C^\pi}(g_i, t)$ and approximate the learning progress as:

$$LP^\pi(g_i, t) = \frac{\delta_{C^\pi}(g_i, t)}{\sum_{j=1}^N \delta_{C^\pi}(g_j, t)}. \quad (10)$$

The curriculum distribution is:

$$\mathcal{K}^\pi(g_i) = \frac{\delta_{C^\pi}(g_i, t)}{\sum_{j=1}^N \delta_{C^\pi}(g_j, t)}, \quad (11)$$

and goals are sampled proportionally to these probabilities during training. This process is repeated periodically to update the curriculum based on the latest policy and Q-value estimates.

Rationale and Benefits. The temporal variance $\delta_{C^\pi}(g, t)$, rooted in Sections 4.1 and 4.2, leverages Q-value dynamics to focus on goals where the policy evolves rapidly. By prioritizing high-variance goals, TEACH’s curriculum enhances sample efficiency, as shown in Section 5. Unlike simple differences, averaging over n timesteps




Figure 1: We evaluate the performance of TEACH on 4 FetchArm, 4 HandManipulate OpenAI Gymnasium [40] environments, and 3 Maze navigation environments taken from [57].

reduces sensitivity to Q-value noise, improving robustness in training. This approach unifies policy and value dynamics, tailored to the continuous, noisy nature of goal-conditioned RL.

Finally, we summarize the proposed method in Algorithm 2. Our approach employs the Deep Deterministic Policy Gradient (DDPG) [33] algorithm as the base RL policy. To enhance the efficiency of universal policy learning, we incorporate the Hindsight Experience Replay (HER) [3] strategy into the replay buffer manipulation. HER allows the off-policy RL algorithm to learn more effectively by reinterpreting unsuccessful trajectories as successful.

A notable advantage of our method is that the teacher component, responsible for curriculum learning, does not introduce any new hyperparameters except the temporal window (n), which captures the evolution of the policy confidence score over the past n time steps. This design choice ensures that the training routine remains consistent with the base RL policy, simplifying the implementation and reducing the potential for hyperparameter tuning complexity. As a result, our method maintains the stability and reliability of the base RL algorithm while enhancing its learning efficiency through strategic replay buffer manipulation and curriculum learning.

5 EXPERIMENTS

We test our method on 11 multi-goal binary reward tasks. The complexity details of the task are shown in Table 1, context refers to the goal’s dimension, and goal space size (N) refers to the number of goals sampled from the continuous goal space to make the problem tractable. While our main focus relies on the robotics task environment [40], we also include three Maze navigation tasks [57] to evaluate performance in low-context problems.

5.1 Baselines

To establish the contribution and effectiveness of the proposed approach, we compare our method with the following baselines. The VDS and ProCurl are currently state-of-the-art in multi-goal and multi-task settings, respectively. All the baselines use value estimates to design ACL, which makes them a strong choice for baselines to highlight the advantage of our approach.

HER-IID [3] In HER-IID, the RL agent uses a hindsight experience replay buffer which independently and identically samples goals from task space. We use the official code-based implementation to reproduce the results.

VDS [57] Value Disagreement Sampling samples goals from the goal space based on value disagreement. Their strategy prioritizes goals that maximize the epistemic uncertainty of the Q-function of the policy. We use the official implementation to reproduce results. **SPaCE** [13] Self-Paced Context Evaluation provides a curriculum learning explicitly using the agent’s performance as an ordering criterion. They use the agent’s state value predictions to generate a curriculum. Which uses the temporal difference to design a curriculum that heuristically selects a new goal if it observes significant learning in the value function. We extended their strategy to our setting using DDPG+HER [3] implementation.

ProCuRL [52] Proximal Curriculum for Reinforcement Learning Agents is inspired by the pedagogical concept of the zone of proximal development. They capture the idea of proximal using state value estimates with respect to the learner’s current policy. Their strategy essentially serves as a geometric mean of the value function over goals and the sample goal with the highest geometric mean. We extended their strategy to our setting using DDPG+HER [3].

5.2 Implementation Details

Our curriculum design comprises two primary components: the teacher (goal proposer) and the student (RL policy). The teacher is a non-learning module that evaluates the student’s capabilities (the RL policy) to propose goals that promote efficient learning. The RL policy, parameterized by θ , is deterministic by design, and we add noise to its actions to enable better exploration. We combine DDPG with HER. Incorporating the HER replay buffer relabeling strategy enhances learning efficiency.

At the beginning of each episode, the student queries the teacher for a target goal. The teacher evaluates the change in the student’s learning progress for each goal within the goal space, which consists of N goals. The teacher selects and assigns the goal with maximum temporal uncertainty in the student’s confidence score as the target goal. The RL policy then focuses on achieving this assigned goal. Transitions generated through agent-environment interactions are

Table 1: Complexity of environments

Environment	Reward	State	Context	Action	Goal space size	Episode length
FetchReach-v2	binary	\mathbb{R}^{10}	\mathbb{R}^3	\mathbb{R}^3	$1e^3$	50
FetchPickAndPlace-v2	binary	\mathbb{R}^{25}	\mathbb{R}^3	\mathbb{R}^4	$1e^3$	50
FetchSlide-v2	binary	\mathbb{R}^{25}	\mathbb{R}^3	\mathbb{R}^3	$1e^3$	50
FetchPush-v2	binary	\mathbb{R}^{25}	\mathbb{R}^3	\mathbb{R}^3	$1e^3$	50
HandManipulateBlock-v1	binary	\mathbb{R}^{61}	\mathbb{R}^7	\mathbb{R}^{20}	$1e^3$	100
HandManipulatePen-v1	binary	\mathbb{R}^{61}	\mathbb{R}^7	\mathbb{R}^{20}	$1e^3$	100
HandManipulateEgg-v1	binary	\mathbb{R}^{61}	\mathbb{R}^7	\mathbb{R}^{20}	$1e^3$	100
HandReach-v1	binary	\mathbb{R}^{63}	\mathbb{R}^{15}	\mathbb{R}^{20}	$1e^3$	50
MazeA-v0	binary	\mathbb{R}^2	\mathbb{R}^2	\mathbb{R}^2	$1e^3$	50
MazeB-v0	binary	\mathbb{R}^2	\mathbb{R}^2	\mathbb{R}^2	$1e^3$	50
MazeC-v0	binary	\mathbb{R}^2	\mathbb{R}^2	\mathbb{R}^2	$1e^3$	50

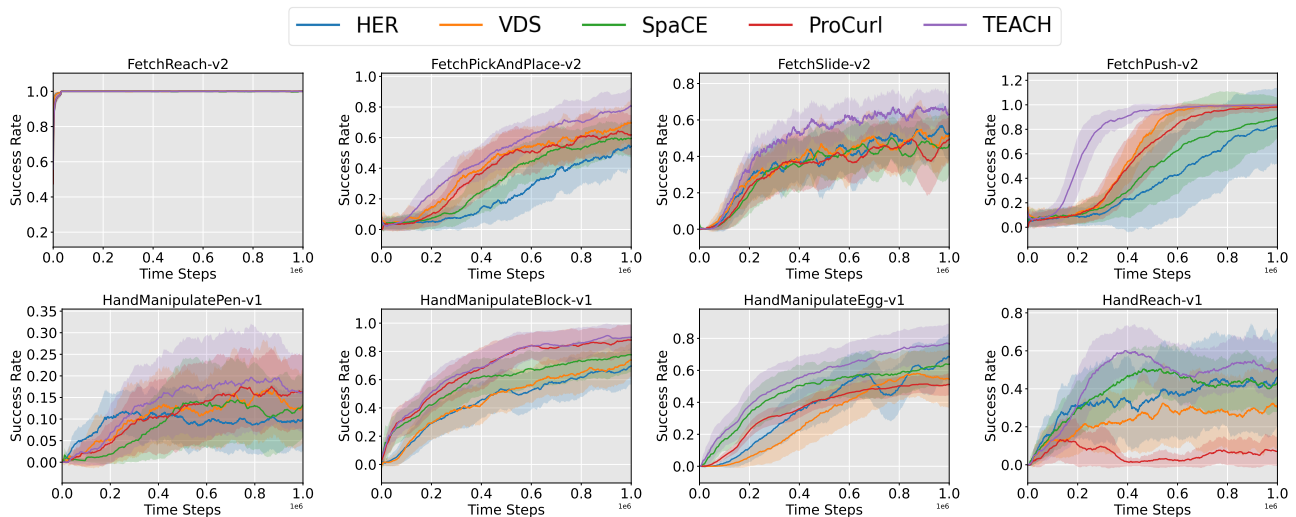


Figure 2: Results show the performance across 8 robotic manipulation tasks (refer extended preprint [8] for complete results). The plots show the success rate along the y-axis, evaluated through the current policy. The reported results are the mean across seven seeds, with shaded regions highlighting the standard deviation.

stored in the replay buffer and subsequently used to update the policy. The network architecture is an MLP with 2 layers for the Fetch and Maze navigation tasks and 3 layers for the HandManipulation task, respectively, for both the actor and critic networks. The Q network is trained using a learning rate of 0.001 with a batch size of 1024. All robotics tasks are trained for 1M time steps and maze navigation tasks for 400K time steps, respectively. The agent’s performance is evaluated by randomly sampling goals from the goal space. All the reported results highlight the agent’s success rate in achieving those randomly sampled goals. Specifically, the success rate is calculated as the average goal-achievement rate over 20 episodes at each evaluation step.

5.3 Improvement Through TEACH

The performance of all methods on robotics manipulation tasks is shown in Figure 2. These results demonstrate that our proposed

approach, which inherently combines dual exploitation of the information from the current policy and state-action value estimates through temporal evolution of Q-value, achieves superior sample efficiency. The performance improvements are evident across low-context tasks, such as *Maze navigation*, and high-context tasks, such as *HandManipulation*. The reported performance of TEACH is computed using a temporal window of $n = 10$, i.e., the curriculum is designed based on the temporal divergence of the Q-function over the past 10 time steps. However, we highlight in Section 5.5 that the performance of TEACH is mostly invariant with the size of the temporal history of the value function.

Our approach consistently yields better policies while maintaining high sample efficiency, as evidenced by extensive comparisons across tasks and baselines. The improvements stem from addressing two primary challenges: (1) the noisy and often inaccurate nature of value estimates and (2) the additional noise introduced by the

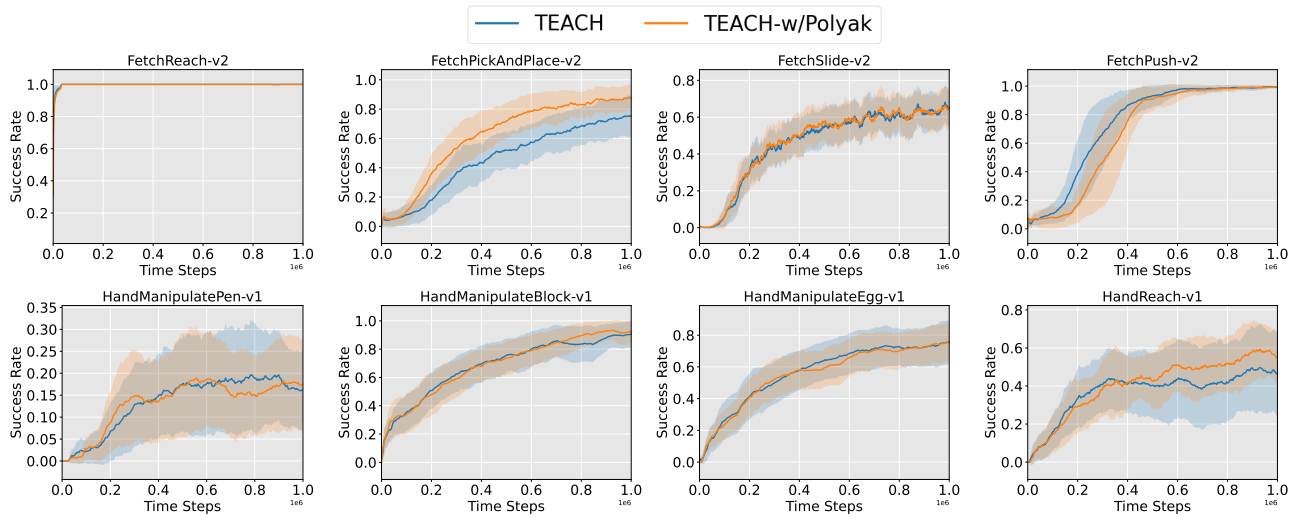


Figure 3: Results show the effect of smooth target confidence score for temporal window size $n = 3$. The results are across seven seeds, where the shaded region represents the standard deviation. We observe that the smooth target confidence score compared to the standard confidence score led to similar performances. This validates the hypothesis that the proposed method measures are robust to noisy value estimates.

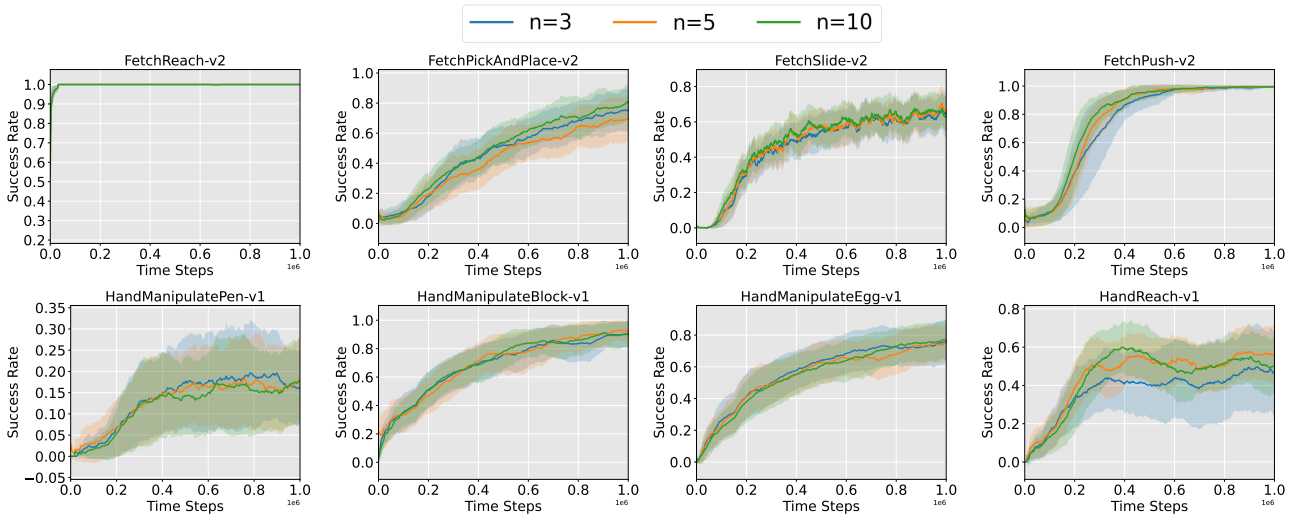


Figure 4: Ablation analysis of the effect of temporal window size n . The results are across seven seeds, where the shaded region represents the standard deviation. We observe that the presented approach is insensitive to the temporal window size.

HER relabeling strategy, which adds fictitious and biased data into the replay buffer used for training the value function.

We mitigate these challenges by designing a curriculum that better aligns with the current policy’s capabilities and is less prone to the detrimental effects of noisy value estimates. By leveraging dual exploitation—drawing information from both the current policy and state-action value estimates using temporal divergence to measure learning progress, our strategy reduces the impact of noise and enhances the robustness and effectiveness of policy learning.

5.4 Effect of Smooth Target Confidence Score on Performance

The issue of noisy value estimates and bias can be addressed by Polyak averaging [41] to obtain a smooth value function. The smoothed values function is often called the ‘target’ values function. The target function parameters are computed using Equation (12), where θ' are parameters of the target values function and θ are

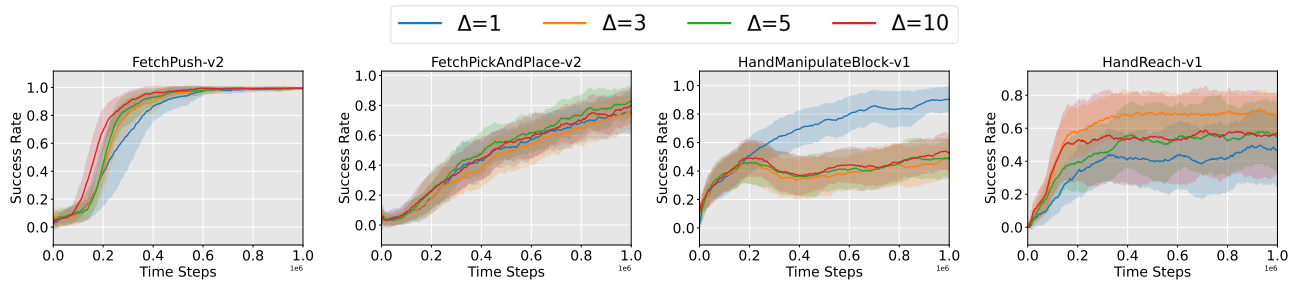


Figure 5: Ablation analysis of the effect of student-teacher interplay frequency Δ . The results are across seven seeds, where the shaded region represents the standard deviation. We observe that the performance can be improved using different interplay frequencies for different tasks.

parameters of the value function.

$$\theta' = \alpha\theta' + (1 - \alpha)\theta \quad (12)$$

We conduct an ablation study incorporating a smooth state-action value function to evaluate the impact of a smooth confidence score. The results, presented in Figure 3, indicate that a smooth target state-action value function does not improve performance compared to the noisy state-action value function used in the proposed algorithm. The smooth success score provides better results in *FetchPickAndPlace* task but performs poorly for *FetchPush* task while performing similarly across the remaining tasks. This underscores the ability of the TEACH to be robust to noisy value estimates, which further validates the arguments made in Section 4. The impact of the noisy value estimates can be reduced further to some extent by using a larger temporal window size.

5.5 Effect of Size of Temporal Window

In this analysis, we investigate the effect of the temporal window size n on our proposed automatic curriculum learning method, TEACH. As shown in Figure 4, the proposed approach demonstrates a high degree of invariance to the size of the temporal window n , which captures the length of past time steps used for computing the temporal divergence of the policy confidence score. This invariance highlights the robustness of TEACH in adapting to different temporal contexts, making it suitable for diverse task settings. However, while the overall performance remains consistent across different values of n , we observe that $n = 10$ produces slightly better results in some tasks. This may be because $n = 10$ strikes an optimal balance between capturing sufficient historical information and avoiding sensitivity to noise compared to smaller temporal windows. Therefore, we report results using $n = 10$ in Figure 1 for consistency and reproducibility.

5.6 Effect of Teacher-Student Interplay Frequency

In this ablation study, we investigate the effect of the teacher-student interplay frequency (Δ) on learning performance. The interplay frequency controls the temporal interval (number of episodes), after which changes in the policy confidence score are reevaluated. As illustrated in Figure 5, a less aggressive interplay frequency can lead to faster learning. These results suggest that a continuous

curriculum generation strategy may induce forgetting behaviors, as agents are exposed to diverse goals at varying time steps. Conversely, oversampling specific goals allows agents to focus on a consistent set of objectives, thereby facilitating improved learning. However, our findings reveal that achieving more stable and consistent performance requires an adaptive interplay approach, yielding superior results to fixed strategies, which may require extensive fine-tuning. While this work employs a fixed teacher-student interplay frequency ($\Delta = 1$) to underscore the benefits of a temporal divergence-driven curriculum strategy. Nevertheless, we emphasize that the design of adaptive curriculum strategies that dynamically adjust the interplay between the teacher and the student to optimize learning performance can result in superior performance. Notably, prior works have not addressed the importance of adaptability in curriculum generation, and we believe that future research should explore the design of adaptive curriculum strategies that dynamically adjust the interplay between the teacher and the student to optimize learning performance.

6 CONCLUSION

In this work, we introduced a novel curriculum strategy for goal-conditioned reinforcement learning that leverages a temporal divergence based approach to address the challenges posed by noisy value estimates in curriculum design. We first demonstrated theoretically that temporal divergence provides an upper bound on policy divergence, offering a more reliable mechanism for evaluating both policy performance and the value function. Subsequently, we conducted extensive experiments across 11 robotics and navigation environments with binary rewards, showcasing the practical effectiveness of our strategy. In these experiments, we highlighted the robustness of our algorithm, demonstrating its insensitivity to smooth target success score estimates and its ability to mitigate the effects of noisy value estimates. Finally, we demonstrated that an adaptive curriculum learning strategy, built upon our temporal divergence-based framework, can outperform fixed strategies, yielding improved and more stable results.

REFERENCES

- [1] İlge Akkaya, Marcin Andrychowicz, Maciek Chociej, Mateusz Litwin, Bob McGrew, Arthur Petron, Alex Paino, Matthias Plappert, Glenn Powell, Raphael Ribas, et al. 2019. Solving rubik’s cube with a robot hand. *arXiv preprint arXiv:1910.07113* (2019).

- [2] Per-Arne Andersen, Morten Goodwin, and Ole-Christoffer Granmo. 2018. Deep RTS: a game environment for deep reinforcement learning in real-time strategy games. In *2018 IEEE conference on computational intelligence and games (CIG)*. IEEE, 1–8.
- [3] Marcin Andrychowicz, Filip Wolski, Alex Ray, Jonas Schneider, Rachel Fong, Peter Welinder, Bob McGrew, Josh Tobin, OpenAI Pieter Abbeel, and Wojciech Zaremba. 2017. Hindsight experience replay. *Advances in neural information processing systems* 30 (2017).
- [4] Richard Bellman. 1966. Dynamic programming. *science* 153, 3731 (1966), 34–37.
- [5] Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. 2009. Curriculum learning (icml). *Google Scholar Google Scholar Digital Library Digital Library* (2009).
- [6] Angelo Cangelosi and Matthew Schlesinger. 2015. *Developmental robotics: From babies to robots*. MIT press.
- [7] Gaurav Chaudhary and Laxmidhar Behera. 2025. From Novelty to Imitation: Self-Distilled Rewards for Offline Reinforcement Learning. arXiv:2507.12815 [cs.LG] <https://arxiv.org/abs/2507.12815>
- [8] Gaurav Chaudhary and Laxmidhar Behera. 2025. TEACH: Temporal Variance-Driven Curriculum for Reinforcement Learning. *arXiv preprint arXiv:2512.22824* (2025).
- [9] Gaurav Chaudhary, Laxmidhar Behera, and Tushar Sandhan. 2023. Active perception system for enhanced visual signal recovery using deep reinforcement learning. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 1–5.
- [10] Gaurav Chaudhary, Wassim Uddin Mondal, and Laxmidhar Behera. 2025. MOORL: A Framework for Integrating Offline-Online Reinforcement Learning. *arXiv preprint arXiv:2506.09574* (2025).
- [11] Karl Cobbe, Chris Hesse, Jacob Hilton, and John Schulman. 2020. Leveraging procedural generation to benchmark reinforcement learning. In *International conference on machine learning*. PMLR, 2048–2056.
- [12] Michael Dennis, Natasha Jaques, Eugene Vinitzky, Alexandre Bayen, Stuart Russell, Andrew Critch, and Sergey Levine. 2020. Emergent complexity and zero-shot transfer via unsupervised environment design. *Advances in neural information processing systems* 33 (2020), 13049–13061.
- [13] Theresa Eimer, André Biedenkapp, Frank Hutter, and Marius Lindauer. 2021. Self-paced context evaluation for contextual reinforcement learning. In *International Conference on Machine Learning*. PMLR, 2948–2958.
- [14] Jeffrey L Elman. 1993. Learning and development in neural networks: The importance of starting small. *Cognition* 48, 1 (1993), 71–99.
- [15] Meng Fang, Tianyi Zhou, Yali Du, Lei Han, and Zhengyou Zhang. 2019. Curriculum-guided hindsight experience replay. *Advances in neural information processing systems* 32 (2019).
- [16] Yannic Flet-Berliac and Philippe Preux. 2019. Samples Are Useful? Not Always: denoising policy gradient updates using variance explained. (2019).
- [17] Carlos Florensa, David Held, Xinyang Geng, and Pieter Abbeel. 2018. Automatic goal generation for reinforcement learning agents. In *International conference on machine learning*. PMLR, 1515–1528.
- [18] Carlos Florensa, David Held, Markus Wulfmeier, Michael Zhang, and Pieter Abbeel. 2017. Reverse curriculum generation for reinforcement learning. In *Conference on robot learning*. PMLR, 482–495.
- [19] Alex Graves, Marc G Bellemare, Jacob Menick, Remi Munos, and Koray Kavukcuoglu. 2017. Automated curriculum learning for neural networks. In *International conference on machine learning*. Pmlr, 1311–1320.
- [20] Tuomas Haarnoja, Haoran Tang, Pieter Abbeel, and Sergey Levine. 2017. Reinforcement learning with deep energy-based policies. In *International conference on machine learning*. PMLR, 1352–1361.
- [21] Dong Han, Beni Mulyana, Vladimir Stankovic, and Samuel Cheng. 2023. A survey on deep reinforcement learning algorithms for robotic manipulation. *Sensors* 23, 7 (2023), 3762.
- [22] Dan Horgan, John Quan, David Budden, Gabriel Barth-Maron, Matteo Hessel, Hado Van Hasselt, and David Silver. 2018. Distributed prioritized experience replay. *arXiv preprint arXiv:1803.00933* (2018).
- [23] Edward S Hu, Richard Chang, Oleh Rybkin, and Dinesh Jayaraman. 2023. Planning goals for exploration. *arXiv preprint arXiv:2303.13002* (2023).
- [24] Minqi Jiang, Michael Dennis, Jack Parker-Holder, Jakob Foerster, Edward Grefenstette, and Tim Rocktäschel. 2021. Replay-guided adversarial environment design. *Advances in Neural Information Processing Systems* 34 (2021), 1884–1897.
- [25] Minqi Jiang, Edward Grefenstette, and Tim Rocktäschel. 2021. Prioritized level replay. In *International Conference on Machine Learning*. PMLR, 4940–4950.
- [26] Avi Karni, Gundela Meyer, Christine Rey-Hipolito, Peter Jezard, Michelle M Adams, Robert Turner, and Leslie G Ungerleider. 1998. The acquisition of skilled motor performance: fast and slow experience-driven changes in primary motor cortex. *Proceedings of the National Academy of Sciences* 95, 3 (1998), 861–868.
- [27] Elia Kaufmann, Antonio Loquercio, Rene Ranftl, Alexey Dosovitskiy, Vladlen Koltun, and Davide Scaramuzza. 2018. Deep drone racing: Learning agile flight in dynamic environments. In *Conference on Robot Learning*. PMLR, 133–145.
- [28] Pascal Klink, Hany Abdulsamad, Boris Belousov, Carlo D’Eramo, Jan Peters, and Joni Pajarinen. 2021. A probabilistic interpretation of self-paced learning with applications to reinforcement learning. *Journal of Machine Learning Research* 22, 182 (2021), 1–52.
- [29] Solomon Kullback and Richard A Leibler. 1951. On information and sufficiency. *The annals of mathematical statistics* 22, 1 (1951), 79–86.
- [30] Qiyang Li, Yuexiang Zhai, Yi Ma, and Sergey Levine. 2023. Understanding the complexity gains of single-task rl with a curriculum. In *International Conference on Machine Learning*. PMLR, 20412–20451.
- [31] Yuxi Li. 2017. Deep reinforcement learning: An overview. *arXiv preprint arXiv:1701.07274* (2017).
- [32] Gabriele Libardi, Gianni De Fabritiis, and Sebastian Dittert. 2021. Guided exploration with proximal policy optimization using a single demonstration. In *International Conference on Machine Learning*. PMLR, 6611–6620.
- [33] Timothy P Lillicrap, Jonathan J Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. 2015. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971* (2015).
- [34] Minghuan Liu, Menghui Zhu, and Weinan Zhang. 2022. Goal-conditioned reinforcement learning: Problems and solutions. *arXiv preprint arXiv:2201.08299* (2022).
- [35] Tamber Matiisen, Avital Oliver, Taco Cohen, and John Schulman. 2019. Teacher-student curriculum learning. *IEEE transactions on neural networks and learning systems* 31, 9 (2019), 3732–3740.
- [36] Russell Mendonca, Oleh Rybkin, Kostas Daniilidis, Danijar Hafner, and Deepak Pathak. 2021. Discovering and achieving goals via world models. *Advances in Neural Information Processing Systems* 34 (2021), 24379–24391.
- [37] Michèle Molina and François Jouen. 1998. Modulation of the palmar grasp behavior in neonates according to texture property. *Infant Behavior and Development* 21, 4 (1998), 659–666.
- [38] Lrel Pinto, James Davidson, Rahul Sukthankar, and Abhinav Gupta. 2017. Robust adversarial reinforcement learning. In *International Conference on Machine Learning*. PMLR, 2817–2826.
- [39] Silviu Pitis, Harris Chan, Stephen Zhao, Bradly Stadie, and Jimmy Ba. 2020. Maximum entropy gain exploration for long horizon multi-goal reinforcement learning. In *International Conference on Machine Learning*. PMLR, 7750–7761.
- [40] Matthias Plappert, Marcin Andrychowicz, Alex Ray, Bob McGrew, Bowen Baker, Glenn Powell, Jonas Schneider, Josh Tobin, Maciek Chociej, Peter Welinder, et al. 2018. Multi-goal reinforcement learning: Challenging robotics environments and request for research. *arXiv preprint arXiv:1802.09464* (2018).
- [41] Boris T Polyak and Anatoli B Juditsky. 1992. Acceleration of stochastic approximation by averaging. *SIAM journal on control and optimization* 30, 4 (1992), 838–855.
- [42] Vitchyr H Pong, Murtaza Dalal, Steven Lin, Ashvin Nair, Shikhar Bahl, and Sergey Levine. 2019. Skew-fit: State-covering self-supervised reinforcement learning. *arXiv preprint arXiv:1903.03698* (2019).
- [43] Rémy Portelas, Cédric Colas, Katja Hofmann, and Pierre-Yves Oudeyer. 2020. Teacher algorithms for curriculum learning of deep rl in continuously parameterized environments. In *Conference on Robot Learning*. PMLR, 835–853.
- [44] Rémy Portelas, Cédric Colas, Lilian Weng, Katja Hofmann, and Pierre-Yves Oudeyer. 2020. Automatic curriculum learning for deep rl: A short survey. *arXiv preprint arXiv:2003.04664* (2020).
- [45] Sébastien Racanière, Andrew K Lampinen, Adam Santoro, David P Reichert, Vlad Firoiu, and Timothy P Lillicrap. 2019. Automated curricula through setter-solver interactions. *arXiv preprint arXiv:1909.12892* (2019).
- [46] Roberta Raileanu and Rob Fergus. 2021. Decoupling value and policy for generalization in reinforcement learning. In *International Conference on Machine Learning*. PMLR, 8787–8798.
- [47] Aravind Rajeswaran, Sarvjeet Ghotra, Balaraman Ravindran, and Sergey Levine. 2016. Epopt: Learning robust neural network policies using model ensembles. *arXiv preprint arXiv:1610.01283* (2016).
- [48] Tom Schaul, Daniel Horgan, Karol Gregor, and David Silver. 2015. Universal value function approximators. In *International conference on machine learning*. PMLR, 1312–1320.
- [49] Oliver G Selfridge, Richard S Sutton, and Andrew G Barto. 1985. Training and Tracking in Robotics.. In *Ijcai*. 670–672.
- [50] Sainbayar Sukhbaatar, Zeming Lin, Ilya Kostrikov, Gabriel Synnaeve, Arthur Szlam, and Rob Fergus. 2017. Intrinsic motivation and automatic curricula via asymmetric self-play. *arXiv preprint arXiv:1703.05407* (2017).
- [51] Open Ended Learning Team, Adam Stooke, Anuj Mahajan, Catarina Barros, Charlie Deck, Jakob Bauer, Jakub Sygnowski, Maja Trebacz, Max Jaderberg, Michael Mathieu, et al. 2021. Open-ended learning leads to generally capable agents. *arXiv preprint arXiv:2107.12808* (2021).
- [52] Georgios Tzannetos, Bárbara Gomes Ribeiro, Parameswaran Kamalaruban, and Adish Singla. 2023. Proximal Curriculum for Reinforcement Learning Agents. *arXiv preprint arXiv:2304.12877* (2023).
- [53] Oriol Vinyals, Igor Babuschkin, Wojciech M Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H Choi, Richard Powell, Timo Ewalds, Petko Georgiev, et al. 2019. Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature* 575, 7782 (2019), 350–354.

- [54] Lev Semenovich Vygotsky and Michael Cole. 1978. *Mind in society: Development of higher psychological processes*. Harvard university press.
- [55] Daphna Weinshall, Gad Cohen, and Dan Amir. 2018. Curriculum learning by transfer learning: Theory and experiments with deep networks. In *International Conference on Machine Learning*. PMLR, 5238–5246.
- [56] Gaurav Yengera, Rati Devidze, Parameswaran Kamalaruban, and Adish Singla. 2021. Curriculum Design for Teaching via Demonstrations: Theory and Applications. *Advances in Neural Information Processing Systems* 34 (2021), 10496–10509.
- [57] Yunzhi Zhang, Pieter Abbeel, and Lerrel Pinto. 2020. Automatic curriculum learning through value disagreement. *Advances in Neural Information Processing Systems* 33 (2020), 7648–7659.