



Diversity-Based Trajectory and Goal Selection with Hindsight Experience Replay

Tianhong Dai¹(✉), Hengyan Liu¹, Kai Arulkumaran^{1,2}, Guangyu Ren¹,
and Anil Anthony Bharath¹

¹ Imperial College London, London SW7 2AZ, UK

{tianhong.dai15,hengyan.liu15,g.ren19,a.bharath}@imperial.ac.uk

² Araya Inc., Tokyo 107-6024, Japan

kai_arulkumaran@araya.org

Abstract. Hindsight experience replay (HER) is a goal relabelling technique typically used with off-policy deep reinforcement learning algorithms to solve goal-oriented tasks; it is well suited to robotic manipulation tasks that deliver only sparse rewards. In HER, both trajectories and transitions are sampled uniformly for training. However, not all of the agent's experiences contribute equally to training, and so naive uniform sampling may lead to inefficient learning. In this paper, we propose diversity-based trajectory and goal selection with HER (DTGSH). Firstly, trajectories are sampled according to the diversity of the goal states as modelled by determinantal point processes (DPPs). Secondly, transitions with diverse goal states are selected from the trajectories by using k -DPPs. We evaluate DTGSH on five challenging robotic manipulation tasks in simulated robot environments, where we show that our method can learn more quickly and reach higher performance than other state-of-the-art approaches on all tasks.

Keywords: Deep reinforcement learning · Determinantal point processes · Hindsight experience replay

1 Introduction

Deep reinforcement learning (DRL) [3], in which neural networks are used as function approximators for reinforcement learning (RL), has been shown to be capable of solving complex control problems in several environments, including board games [27, 28], video games [4, 19, 30], simulated and real robotic manipulation [2, 9, 15] and simulated autonomous driving [12].

However, learning from a *sparse* reward signal, where the only reward is provided upon the completion of a task, still remains difficult. An agent may rarely or never encounter positive examples from which to learn in a sparse-reward environment. Many domains therefore provide dense reward signals [5], or practitioners may turn to reward shaping [20]. Designing dense reward functions

typically requires prior domain knowledge, making this approach difficult to generalise across different environments.

Fortunately, a common scenario is goal-oriented RL, where the RL agent is tasked with solving different goals within the same environment [11, 25]. Even if each task has a sparse reward, the agent ideally *generalises* across goals, making the learning process easier. For example, in a robotic manipulation task, the goal during a single episode would be to achieve a specific position of a target object.

Hindsight experience replay (HER) [1] was proposed to improve the learning efficiency of goal-oriented RL agents in sparse reward settings: when past experience is replayed to train the agent, the desired goal is replaced (in “hindsight”) with the achieved goal, generating many positive experiences. In the above example, the desired target position would be overwritten with the achieved target position, with the achieved reward also being overwritten correspondingly.

We note that HER, whilst it enabled solutions to previously unsolved tasks, can be somewhat inefficient in its use of uniformly sampling transitions during training. In the same way that prioritised experience replay [26] has significantly improved over the standard experience replay in RL, several approaches have improved upon HER by using data-dependent sampling [8, 32]. HER with energy-based prioritisation (HEBP) [32] assumes semantic knowledge about the goal-space and uses the energy of the target objects to sample trajectories with high energies, and then samples transitions uniformly. Curriculum-guided HER (CHER) [8] samples trajectories uniformly, and then samples transitions based on a mixture of proximity to the desired goal and the diversity of the samples; CHER adapts the weighting of these factors over time. In this work, we introduce diversity-based trajectory and goal selection with HER (DTGSH; See Fig. 1), which samples trajectories based on the diversity of the goals achieved within the trajectory, and then samples transitions based on the diversity of the set of samples. In this paper, DTGSH is evaluated on five challenging robotic manipulation tasks. From extensive experiments, our proposed method converges faster and reaches higher rewards than prior work, without requiring domain knowledge [32] or tuning a curriculum [8], and is based on a single concept—determinantal point processes (DPPs) [14].

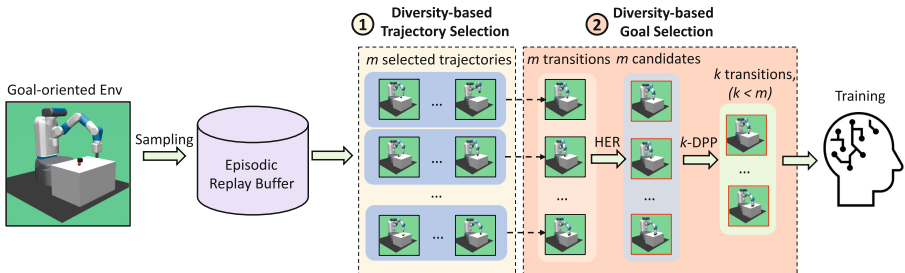


Fig. 1. Overview of DTGSH. Every time a new episode is completed, its diversity is calculated, and it is stored in the episodic replay buffer. During training, m episodes are sampled according to their diversity-based priority, and then k diverse, hindsight-relabelled transitions are sampled using a k -DPP [13].

2 Background

2.1 Reinforcement Learning

RL is the study of agents interacting with their environment in order to maximise their reward, formalised using the framework of Markov decision processes (MDPs) [29]. At each timestep t , an agent receives a state s_t from the environment, and then samples an action a_t from its policy $\pi(a_t|s_t)$. Next, the action a_t is executed in the environment to get the next state s_{t+1} , and a reward r_t . In the episodic RL setting, the objective of the agent is to maximise its expected return $\mathbb{E}[R]$ over a finite trajectory with length T :

$$\mathbb{E}[R] = \mathbb{E} \left[\sum_{t=1}^T \gamma^{t-1} r_t \right], \quad (1)$$

where $\gamma \in [0, 1]$ is a discount factor that exponentially downplays the influence of future rewards, reducing the variance of the return.

2.2 Goal-Oriented Reinforcement Learning

RL can be expanded to the multi-goal setting, where the agent’s policy and the environment’s reward function $\mathcal{R}(s_t, a_t)$ are also conditioned on a goal g [11, 25]. In this work, we focus on the goal-oriented setting and environments proposed by OpenAI [23].

In this setting, every episode comes with a desired goal g , which specifies the desired configuration of a target object in the environment (which could include the agent itself). At every timestep t , the agent is also provided with the currently achieved goal g_{t+1}^{ac} . A transition in the environment is thus denoted as: $(s_t, a_t, r_t, s_{t+1}, g, g_{t+1}^{ac})$. The environment provides a sparse reward function, where a negative reward is given unless the achieved goal is within a small distance ϵ of the desired goal:

$$\mathcal{R}(g, g_{t+1}^{ac}) := \begin{cases} 0 & \text{if } \|g_{t+1}^{ac} - g\| \leq \epsilon \\ -1 & \text{otherwise.} \end{cases} \quad (2)$$

However, in this setting, the agent is unlikely to achieve a non-negative reward through random exploration. To overcome this, HER provides successful experiences for the agent to learn from by relabelling transitions during training: the agent trains on a hindsight desired goal g' , which is set to the achieved goal g_{t+1}^{ac} , with r_t recomputed using the environment reward function (Eq. 2).

2.3 Deep Deterministic Policy Gradient

Deep deterministic policy gradient (DDPG) [16] is an off-policy actor-critic DRL algorithm for continuous control tasks, and is used as the baseline algorithm for HER [1, 8, 32]. The actor $\pi_\theta(s_t)$ is a policy network parameterised

by θ , and outputs the agent’s actions. The critic $Q_\eta(s_t, a_t)$ is a state-action-value function approximator parameterised by η , and estimates the expected return following a given state-action pair. The critic is trained by minimising $\mathcal{L}_c = \mathbb{E}[(Q_\eta(s_t, a_t) - y_t)^2]$ where $y_t = r_t + \gamma Q_\eta(s_{t+1}, \pi_\theta(s_{t+1}))$. The actor is trained by maximising $\mathcal{L}_a = \mathbb{E}[Q_\eta(s_t, \pi_\theta(s_t))]$, backpropagating through the critic. Further implementation details can be found in prior work [1, 16].

2.4 Determinantal Point Processes

A DPP [14] is a stochastic process that characterises a probability distribution over sets of points using the determinant of some function. In machine learning it is often used to quantify the diversity of a subset, with applications such as video [18] and document summarisation [10].

Formally, for a discrete set of points $\mathcal{Y} = \{x_1, x_2, \dots, x_N\}$, a point process \mathcal{P} is a probability measure over all $2^{|\mathcal{Y}|}$ subsets. \mathcal{P} is a DPP if a random subset \mathbf{Y} is sampled with probability:

$$\mathcal{P}_L(\mathbf{Y} = Y) = \frac{\det(L_Y)}{\sum_{Y' \subseteq \mathcal{Y}} \det(L_{Y'})} = \frac{\det(L_Y)}{\det(L + I)}, \quad (3)$$

where $Y \subseteq \mathcal{Y}$, I is the identity matrix, $L \in \mathbb{R}^{N \times N}$ is the positive semi-definite DPP kernel matrix, and L_Y is the sub-matrix with rows and columns indexed by the elements of the subset Y .

The kernel matrix L can be represented as the Gram matrix $L = X^T X$, where each column of X is the feature vector of an item in \mathcal{Y} . The determinant, $\det(L_Y)$, represents the (squared) volume spanned by vectors $x_i \in Y$. From a geometric perspective, feature vectors that are closer to being orthogonal to each other will have a larger determinant, and vectors in the spanned subspace are more likely to be sampled: $\mathcal{P}_L(\mathbf{Y} = Y) \propto \det(L_Y)$. Using orthogonality as a measure of diversity, we leverage DPPs to sample diverse trajectories and goals.

3 Related Work

The proposed work is built on HER [1] as a way to effectively augment goal-oriented transitions from a replay buffer: to address the problem of sparse rewards, transitions from unsuccessful trajectories are turned into successful ones. HER uses an episodic replay buffer, with uniform sampling over trajectories, and uniform sampling over transitions. However, these samples may be redundant, and many may contribute little to the successful training of the agent.

In the literature, some efforts have been made to increase the efficiency of HER by prioritising more valuable episodes/transitions. Motivated by the work-energy principle in physics, HEBP [32] assigns higher probability to trajectories in which the target object has higher energy; once the episodes are sampled, the transitions are then sampled uniformly. However, HEBP requires knowing the semantics of the goal space in order to calculate the probability, which is proportional to the sum of the target’s potential, kinetic and rotational energies.

CHER [8] dynamically controls the sampling of transitions during training based on a mixture of goal proximity and diversity. Firstly, m episodes are uniformly sampled from the episodic replay buffer, and then a minibatch of $k < m$ is sampled according to the current state of the curriculum. The curriculum initially biases sampling to achieved goals that are close to the desired goal (requiring a distance function), and later biases sampling towards diverse goals, using a k -nearest neighbour graph and a submodular function to more efficiently sample a diverse subset of goals (using the same distance function).

Other work has expanded HER in orthogonal directions. Hindsight policy gradient [24] and episodic self-imitation learning [6] apply HER to improve the efficiency of goal-based on-policy algorithms. Dynamic HER [7] and competitive ER [17] expand HER to the dynamic goal and multi-agent settings, respectively.

The use of DPPs in RL has been more limited, with applications towards modelling value functions of sets of agents in multiagent RL [21,31], and most closely related to us, finding diverse policies [22].

4 Methodology

We now formally describe the two main components of our method, DTGSH: 1) a diversity-based trajectory selection module to sample valuable trajectories for the further goal selection; 2) a diversity-based goal selection module to select transitions with diverse goal states from the previously selected trajectories. Together, these select informative transitions from a large area of the goal space, improving the agent’s ability to learn and generalise.

4.1 Diversity-Based Trajectory Selection

We propose a diversity-based prioritization method to select valuable trajectories for efficient training. Related to HEBP’s prioritisation of high-energy trajectories [32], we hypothesise that trajectories that achieve diverse goal states g_t^{ac} are more valuable for training; however, unlike HEBP, we do not require knowledge of the goal space semantics.

In a robotic manipulation task, the agent needs to move a target object from its initial position, g_1^{ac} , to the target position, g . If the agent never moves the object, despite hindsight relabelling it will not be learning information that would directly help in task completion. On the other hand, if the object moves a lot, hindsight relabelling will help the agent learn about meaningful interactions.

In our approach, DPPs are used to model the diversity of achieved goal states g_t^{ac} in an episode, or subsets thereof. For a single trajectory \mathcal{T} of length T , we divide it into several partial trajectories τ_j of length b , with achieved goal states $\{g_t^{ac}\}_{t=n:n+b-1}$. That is, with a sliding window of $b = 2$, a trajectory \mathcal{T} can be divided into N_p partial trajectories:

$$\mathcal{T}_i = \left\{ \underbrace{\{g_1^{ac}, g_2^{ac}\}}_{\tau_1}, \underbrace{\{g_2^{ac}, g_3^{ac}\}}_{\tau_2}, \underbrace{\{g_3^{ac}, g_4^{ac}\}}_{\tau_3}, \dots, \underbrace{\{g_{T-1}^{ac}, g_T^{ac}\}}_{\tau_{N_p}} \right\}. \quad (4)$$

The diversity d_{τ_j} of each partial trajectory τ_j can be computed as:

$$d_{\tau_j} = \det(L_{\tau_j}), \quad (5)$$

where L_{τ_j} is the kernel matrix of partial trajectory τ_j :

$$L_{\tau_j} = M^T M, \quad (6)$$

and $M = [\hat{g}_n^{ac}, \hat{g}_{n+1}^{ac}, \dots, \hat{g}_{n+b-1}^{ac}]$, where each \hat{g}^{ac} is the ℓ_2 -normalised version of the achieved goal g^{ac} [13]. Finally, the diversity $d_{\mathcal{T}}$ of trajectory \mathcal{T} is the sum of the diversity of its N_p constituent partial trajectories:

$$d_{\mathcal{T}} = \sum_{j=1}^{N_p} d_{\tau_j}. \quad (7)$$

Similarly to HEBP [32], we use a non-uniform episode sampling strategy. During training, we prioritise sampling episodes proportionally to their diversity; the probability $p(\mathcal{T}_i)$ of sampling trajectory \mathcal{T}_i from a replay buffer of size N_e is:


$$p(\mathcal{T}_i) = \frac{d_{\mathcal{T}_i}}{\sum_{n=1}^{N_e} d_{\mathcal{T}_n}}. \quad (8)$$

4.2 Diversity-Based Goal Selection


In prior work [1, 32], after selecting the trajectories from the replay buffer, one transition from each selected trajectory is sampled uniformly to construct a minibatch for training. However, the modified goals g' in the minibatch might be similar, resulting in redundant information. In order to form a minibatch with diverse goals for more efficient learning, we use k -DPPs [13] for sampling goals. Compared to the standard DPP, a k -DPP is a conditional DPP where the subset Y has a fixed size k , with the probability distribution function:

$$\mathcal{P}_L^k(\mathbf{Y} = Y) = \frac{\det(L_Y)}{\sum_{|Y'|=k} \det(L_{Y'})}. \quad (9)$$

k -DPPs are more appropriate for goal selection with a minibatch of fixed size k . Given $m > k$ trajectories sampled from the replay buffer, we first uniformly sample a transition from each of the m trajectories. Finally, a k -DPP is used to sample a diverse set of transitions based on the relabelled goals g' (which, in this context, we denote as ‘‘candidate goals’’). Figure 2a gives an example of uniform vs. k -DPP sampling, demonstrating the increased coverage of the latter. Figure 2b provides corresponding estimated density plots; note that the density of the k -DPP samples is actually more uniform over the *support* of the candidate goal distribution.



(a) Plot of candidate goals and selected goals. k -DPP sampling is more likely to sample points from the full span of the goal space.



(b) Kernel density estimation of the distributions of goals. k -DPP leads to a more uniform selection of goals over the support of the goal space.

Fig. 2. Visualisation of $k = 32$ goals selected from $m = 100$ candidate goals of the Push task using either uniform sampling or k -DPP sampling, respectively. The candidate goals are distributed over a 2D (xy) space. Note that k -DPP sampling (right hand plots) results in a broader span of selected goals in xy space compared to uniform sampling (left hand plots).

Algorithm 1. Diversity-based Goal Selection using k -DPP

Require: set of m candidate goal states $\mathcal{G} := \{g_i\}_{i=1:m}$, minibatch size k

- 1: $J \leftarrow \emptyset, M \leftarrow [g_1, g_2, \dots, g_m]$
 - 2: Calculate the DPP kernel matrix L_M
 - 3: $\{v_n, \lambda_n\} \leftarrow \text{EigenDecomposition}(L_M)$
 - 4: $e_k(\lambda_1, \lambda_2, \dots, \lambda_m) := \sum_{\substack{J' \subseteq \{1, 2, \dots, m\} \\ |J'|=k}} \prod_{n \in J'} \lambda_n$ \triangleright elementary symmetric polynomial: e_k^m
 - 5: **for** $n = m, m - 1, \dots, 1$ **do**
 - 6: **if** $u \sim \text{Uniform}[0, 1] < \lambda_n \frac{e_{k-1}^{n-1}}{e_k^n}$ **then**
 - 7: $J \leftarrow J \cup \{n\}, k \leftarrow k - 1$
 - 8: **if** $k = 0$ **then**
 - 9: **break**
 - 10: **end if**
 - 11: **end if**
 - 12: **end for**
 - 13: $V \leftarrow \{v_n\}_{n \in J}, B \leftarrow \emptyset$
 - 14: **while** $|V| > 0$ **do**
 - 15: Select g_i from \mathcal{G} with $p(g_i) := \frac{1}{|V|} \sum_{v \in V} (v^T \mathbf{b}_i)^2$ $\triangleright \mathbf{b}_i$ is the i^{th} standard basis
 - 16: $B \leftarrow B \cup \{g_i\}$
 - 17: $V \leftarrow V_{\perp}$ \triangleright an orthonormal basis for the subspace of V orthogonal to \mathbf{b}_i
 - 18: **end while**
 - 19: **return** minibatch B with size k
-

Algorithm 2. Diversity-based Trajectory and Goal Selection with HER

Require: RL environment with episodes of length T , number of episodes N , off-policy RL algorithm \mathbb{A} , episodic replay buffer \mathcal{B} , number of algorithm updates U , candidate transitions size m , minibatch size k

- 1: Initialize the parameters θ of all models in \mathbb{A}
- 2: $\mathcal{B} \leftarrow \emptyset$
- 3: **for** $i = 1, 2, \dots, N$ **do**
- 4: Sample a desired goal g and an initial state s_0 ▷ start a new episode
- 5: **for** $t = 1, 2, \dots, T$ **do**
- 6: Sample an action a_t using the policy $\pi(s_t, g; \theta)$
- 7: Execute action a_t and get the next state s_{t+1} and achieved goal state g_{t+1}^{ac}
- 8: Calculate r_t according to Eq. (2)
- 9: Store transition $(s_t, a_t, r_t, s_{t+1}, g, g_{t+1}^{ac})$ in \mathcal{B}
- 10: **end for**
- 11: Calculate the diversity score of current episode $d_{\mathcal{T}_i}$ using Eq. (5) and Eq. (7)
- 12: Calculate the diversity-based priority $p(\mathcal{T})$ of each episode in \mathcal{B} using Eq. (8)
- 13: **for** iteration = $1, 2, \dots, U$ **do**
- 14: Sample m trajectories from \mathcal{B} according to priority $p(\mathcal{T})$
- 15: Uniformly sample one transition from each of the m trajectories
- 16: Relabel goals in each transition and recompute the reward to get m candidate transitions $\{(s_t, a_t, r'_t, s_{t+1}, g'_t)\}_{n=1:m}$
- 17: Sample minibatch B with size k from the m candidates using Algorithm 1
- 18: Optimise θ with minibatch B
- 19: **end for**
- 20: **end for**

Algorithm 1 shows the details of the goal selection subroutine, and Algorithm 2 gives the overall algorithm for our method, DTGSH.

5 Experiments

We evaluate our proposed method, and compare it with current state-of-the-art HER-based algorithms [1, 8, 32] on challenging robotic manipulation tasks [23], pictured in Fig. 3. Furthermore, we perform ablation studies on our diversity-based trajectory and goal selection modules. Our code is based on OpenAI Baselines¹, and is available at: <https://github.com/TianhongDai/div-hindsight>.

5.1 Environments

The robotic manipulation environments used for training and evaluation include five different tasks. Two tasks use the 7-DoF Fetch robotic arm with two-fingers parallel gripper: Push, and Pick&Place, which both require the agent to move a cube to the target position. The remaining three tasks use a 24-DoF Shadow Dexterous Hand to manipulate an egg, a block and a pen, respectively. The sparse reward function is given by Eq. (2).

¹ <https://github.com/openai/baselines>.




Fig. 3. Robotic manipulation environments. (a–b) use the Fetch robot, and (c–e) use the Shadow Dexterous Hand.

In the Fetch environments, the state s_t contains the position and velocity of the joints, and the position and rotation of the cube. Each action a_t is a 4-dimensional vector, with three dimensions specifying the relative position of the gripper, and the final dimension specifying the state of the gripper (i.e., open or closed). The desired goal g is the target position, and the achieved goal g_t^{ac} is the position of the cube. Each episode is of length $T = 50$.

In the Shadow Dexterous Hand environments, the state s_t contains the position and velocity of the joints. Each action a_t is a 20-dimensional vector which specifies the absolute position of 20 non-coupled joints in the hand. The desired goal g and achieved goal g_t^{ac} specify the rotation of the object for the block and pen tasks, and the position + rotation of the object for the egg task. Each episode is of length $T = 200$.

5.2 Training Settings

We base our training setup on CHER [8]. We train all agents on minibatches of size $k = 64$ for 50 epochs using MPI for parallelisation over 16 CPU cores; each epoch consists of 1600 (16×100) episodes, with evaluation over 160 (16×10) episodes at the end of each epoch. Remaining hyperparameters for the baselines are taken from the original work [1, 8, 32]. Our method, DTGSH, uses partial trajectories of length $b = 2$ and $m = 100$ as the number of candidate goals.

5.3 Benchmark Results

We compare DTGSH to DDPG [16], DDPG+HER [1], DDPG+HEBP [32] and DDPG+CHER [8]. Evaluation results are given based on repeated runs with 5 different seeds; we plot the median success rate with upper and lower bounds given by the 75th and 25th percentiles, respectively.

Figure 4 and Table 1 show the performance of DDPG+DTGSH and baseline approaches on all five tasks. In the Fetch tasks, DDPG+DTGSH and DDPG+HEBP both learn significantly faster than the other methods, while in the Shadow Dexterous Hand tasks DDPG+DTGSH learns the fastest and achieves higher success rates than all other methods. In particular, DDPG cannot solve any tasks without using HER, and CHER performs worse in the Fetch




Fig. 4. Success rate of DTGSH and baseline approaches.

tasks. We believe the results highlight the importance of sampling both diverse trajectories *and* goals, as in our proposed method, DTGSH.

Table 1. Final mean success rate \pm standard deviation, with best results in **bold**.

	Push	Pick&Place	EggFull	BlockRotate	PenRotate
DDPG [16]	0.09 \pm 0.01	0.04 \pm 0.00	0.00 \pm 0.00	0.01 \pm 0.00	0.00 \pm 0.00
DDPG+HER [1]	1.00 \pm 0.00	0.89 \pm 0.03	0.11 \pm 0.01	0.55 \pm 0.04	0.15 \pm 0.02
DDPG+HEBP [32]	1.00 \pm 0.00	0.91 \pm 0.03	0.14 \pm 0.02	0.59 \pm 0.02	0.20 \pm 0.03
DDPG+CHER [8]	1.00 \pm 0.00	0.91 \pm 0.04	0.15 \pm 0.01	0.54 \pm 0.04	0.17 \pm 0.03
DDPG+DTGSH	1.00 \pm 0.00	0.94 \pm 0.01	0.17 \pm 0.03	0.62 \pm 0.02	0.21 \pm 0.02

5.4 Ablation Studies

In this section, we perform the following experiments to investigate the effectiveness of each component in DTGSH: 1) diversity-based trajectory selection with HER (DTSH) and diversity-based goal selection with HER (DGSH) are evaluated independently to assess the contribution of each stage; 2) the performance using different partial trajectory lengths b ; 3) the performance of using different candidate goal set sizes m .

Figure 5 shows the performance of using DTSH and DGSH independently. DDPG+DTSH outperforms DDPG+HER substantially in all tasks, which supports the view that sampling trajectories with diverse achieved goals can substantially improve performance. Furthermore, unlike DDPG+HEBP, DTSH does not require knowing the structure of the goal space in order to calculate the energy of the target object; DDPG+DGSH achieves better performance than




Fig. 5. Success rate of HER, DTGSH, and ablations DTSH and DGS.

DDPG+HER in three environments, and is only worse in one environment. DGS performs better in environments where it is easier to solve the task (e.g., Fetch tasks), and hence the trajectories selected are more likely to contain useful transitions. However, DTGSH, which is the combination of both modules, performs the best overall.

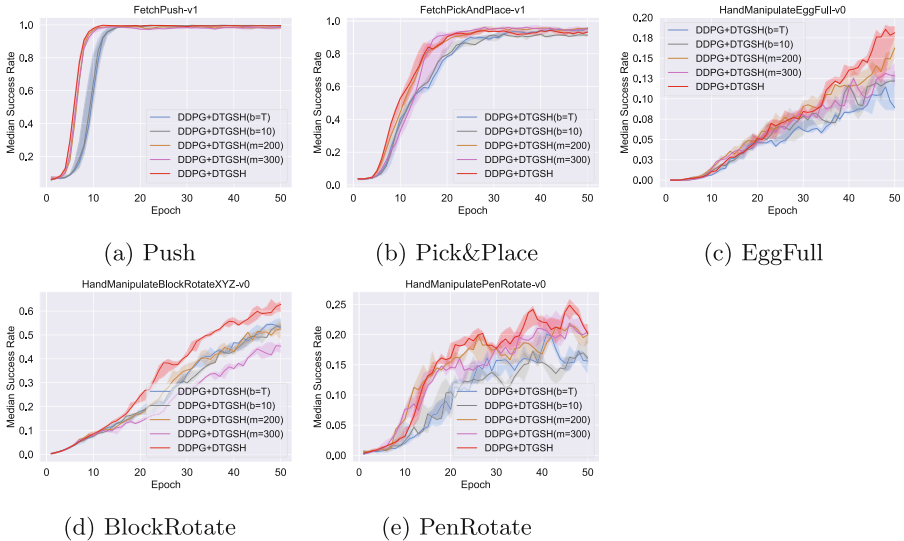


Fig. 6. Success rate of DTGSH with different partial trajectory lengths b and different candidate goal set sizes m .

Figure 6 shows the performance of DDPG+DTGSH with different partial trajectory lengths b and different candidate goal set sizes m . In this work, we use $b = 2$ and $m = 100$ as the defaults. Performance degrades with $b \gg 2$, indicating that pairwise diversity is best for learning in our method. $m \gg 100$ does not affect performance in the Fetch environments, but degrades performance in the Shadow Dexterous Hand environments.

5.5 Time Complexity

Table 2 gives example training times of all of the HER-based algorithms. DTGSH requires an additional calculation of the diversity score of $\mathcal{O}(N_p b^3)$ at the end of every training episode, and sampling of $\mathcal{O}(mk^2)$ for each minibatch.

Table 2. Training time (hours:minutes:seconds) of DTGSH and baseline approaches on the Push task for 50 epochs.

	DDPG+HER [1]	DDPG+HEBP [32]	DDPG+CHER [8]	DDPG+DTGSH
Time	00:55:08	00:56:32	03:02:18	01:52:30

6 Conclusion

In this paper, we introduced diversity-based trajectory and goal selection with hindsight experience replay (DTGSH) to improve the learning efficiency of goal-orientated RL agents in the sparse reward setting. Our method can be divided into two stages: 1) valuable trajectories are selected according to diversity-based priority, as modelled by determinantal point processes (DPPs) [14]; 2) k -DPPs [13] are leveraged to sample transitions with diverse goal states from previously selected trajectories for training. Our experiments empirically show that DTGSH achieves faster learning and higher final performance in five challenging robotic manipulation tasks, compared to previous state-of-the-art approaches [1, 8, 32]. Furthermore, unlike prior extensions of hindsight experience replay, DTGSH does not require semantic knowledge of the goal space [32], and does not require tuning a curriculum [8].

Acknowledgements. This work was supported by JST, Moonshot R&D Grant Number JPMJMS2012.

References

1. Andrychowicz, M., et al.: Hindsight experience replay. In: Neural Information Processing Systems (2017)
2. Andrychowicz, O.M., et al.: Learning dexterous in-hand manipulation. Int. J. Robot. Res. **39**(1), 3–20 (2020)

3. Arulkumaran, K., Deisenroth, M.P., Brundage, M., Bharath, A.A.: Deep reinforcement learning: a brief survey. *IEEE Signal Process. Mag.* **34**(6), 26–38 (2017)
4. Berner, C., et al.: Dota 2 with large scale deep reinforcement learning. [arXiv:1912.06680](https://arxiv.org/abs/1912.06680) (2019)
5. Brockman, G., et al.: Openai gym. arXiv preprint [arXiv:1606.01540](https://arxiv.org/abs/1606.01540) (2016)
6. Dai, T., Liu, H., Bharath, A.A.: Episodic self-imitation learning with hindsight. *Electronics* **9**(10), 1742 (2020)
7. Fang, M., Zhou, C., Shi, B., Gong, B., Xu, J., Zhang, T.: DHER: Hindsight experience replay for dynamic goals. In: *International Conference on Learning Representations* (2018)
8. Fang, M., Zhou, T., Du, Y., Han, L., Zhang, Z.: Curriculum-guided hindsight experience replay. In: *Neural Information Processing Systems* (2019)
9. Gu, S., Holly, E., Lillicrap, T., Levine, S.: Deep reinforcement learning for robotic manipulation with asynchronous off-policy updates. In: *International Conference on Robotics and Automation* (2017)
10. Hong, K., Nenkova, A.: Improving the estimation of word importance for news multi-document summarization. In: *Conference of the European Chapter of the Association for Computational Linguistics* (2014)
11. Kaelbling, L.P.: Learning to achieve goals. In: *International Joint Conference on Artificial Intelligence* (1993)
12. Kiran, B.R., et al.: Deep reinforcement learning for autonomous driving: a survey. *IEEE Trans. Intell. Transp. Syst.* 1–18 (2021)
13. Kulesza, A., Taskar, B.: k-DPPs: fixed-size determinantal point processes. In: *International Conference on Machine Learning* (2011)
14. Kulesza, A., et al.: *Determinantal Point Processes for Machine Learning*. Foundations and Trends in Machine Learning (2012)
15. Levine, S., Finn, C., Darrell, T., Abbeel, P.: End-to-end training of deep visuomotor policies. *J. Mach. Learn. Res.* **17**, 1–40 (2016)
16. Lillicrap, T.P., et al.: Continuous control with deep reinforcement learning. In: *International Conference on Learning Representations* (2016)
17. Liu, H., Trott, A., Socher, R., Xiong, C.: Competitive experience replay. In: *International Conference on Learning Representations* (2019)
18. Mahasseni, B., Lam, M., Todorovic, S.: Unsupervised video summarization with adversarial lstm networks. In: *IEEE Conference on Computer Vision and Pattern Recognition* (2017)
19. Mnih, V., et al.: Human-level control through deep reinforcement learning. *Nature* **518**, 529–533 (2015)
20. Ng, A.Y., Harada, D., Russell, S.: Theory and application to reward shaping. In: *International Conference on Machine Learning* (1999)
21. Osogami, T., Raymond, R.: Determinantal reinforcement learning. In: *AAAI Conference on Artificial Intelligence* (2019)
22. Parker-Holder, J., Pacchiano, A., Choromanski, K.M., Roberts, S.J.: Effective diversity in population based reinforcement learning. In: *Neural Information Processing Systems* (2020)
23. Plappert, M., et al.: Multi-goal reinforcement learning: challenging robotics environments and request for research. [arXiv:1802.09464](https://arxiv.org/abs/1802.09464) (2018)
24. Rauber, P., Ummadisingu, A., Mutz, F., Schmidhuber, J.: Hindsight policy gradients. In: *International Conference on Learning Representations* (2019)
25. Schaul, T., Horgan, D., Gregor, K., Silver, D.: Universal value function approximators. In: *International Conference on Machine Learning* (2015)

26. Schaul, T., Quan, J., Antonoglou, I., Silver, D.: Prioritized experience replay. In: International Conference on Learning Representations (2016)
27. Schrittwieser, J., et al.: Mastering Atari, go, chess and shogi by planning with a learned model. *Nature* **588**, 604–609 (2020)
28. Silver, D., et al.: Mastering the game of go without human knowledge. *Nature* **550**, 354–359 (2017)
29. Sutton, R.S., Barto, A.G.: Reinforcement Learning: An Introduction. The MIT Press, Cambridge (2018)
30. Vinyals, O., et al.: Grandmaster level in Starcraft ii using multi-agent reinforcement learning. *Nature* **575**, 350–354 (2019)
31. Yang, Y., et al.: Multi-agent determinantal q-learning. In: International Conference on Machine Learning (2020)
32. Zhao, R., Tresp, V.: Energy-based hindsight experience prioritization. In: Conference on Robot Learning (2018)