arXiv:1907.10247v1 [cs.LG] 24 Jul 2019

Efficient Exploration with Self-Imitation Learning via Trajectory-Conditioned Policy

Yijie Guo[1] **Jongwook Choi**[1] **Marcin Moczulski**[2]
Samy Bengio[2] **Mohammad Norouzi**[2] **Honglak Lee**[2,1]
[1]University of Michigan [2]Google Brain
{guoyijie,jwook}@umich.edu
{moczulski,bengio,mnorouzi,honglak}@google.com

Abstract

This paper proposes a method for learning a trajectory-conditioned policy to imitate diverse demonstrations from the agent's own past experiences. We demonstrate that such self-imitation drives exploration in diverse directions and increases the chance of finding a globally optimal solution in reinforcement learning problems, especially when the reward is sparse and deceptive. Our method significantly outperforms existing self-imitation learning and count-based exploration methods on various sparse-reward reinforcement learning tasks with local optima. In particular, we report a state-of-the-art score of more than 25,000 points on Montezuma's Revenge without using expert demonstrations or resetting to arbitrary states.

1 Introduction

Efficient exploration to learn a (near-)optimal behavior in the long term is a challenging problem in reinforcement learning (RL). With sparse reward and no expert demonstrations available, the agent must carefully balance the exploration and exploitation when taking a long sequence of actions to receive infrequent non-zero rewards. Many existing methods provide a guidance for exploration based on visitation counts [42, 5, 28, 45, 11] or errors in predicting dynamics [48, 38, 40, 31, 8, 7] to encourage visiting novel states, where the trade-off between exploration and exploitation is usually controlled by the weight of the exploration signal.

Self-Imitation Learning (SIL) [27] is one of the recent methods to tackle this problem of the exploration-exploitation dilemma. This method exploits the agent's previous good trajectories to improve the efficiency of learning, demonstrating that exploitation could indirectly drive further exploration in certain environments. However, in environments with locally optimal solutions, exploiting the experience of accumulating deceptive rewards may mislead the agent and prevent it from reaching a higher return in the long term. For example, as illustrated in Figure 1, the agent starts in the bottom left corner where it easily collects the apple near its initial location by random exploration and achieves a small positive reward. The SIL agent exploits the trajectory following the orange path and learns to collect the nearby rewards quickly. However, it is less likely to collect the key, open the door, and get the treasure within a given time limit. Therefore, in order to find the optimal path (purple), it is better to exploit the past experience in diverse directions (gray paths), instead of focusing only on the trajectories with the myopic and sub-optimal rewards.

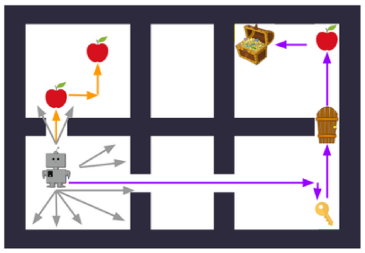

Figure 1: Map of Key-Door-Treasure domain, where the reward for getting an apple, picking up a key, opening a door and getting a treasure is 1, 1, 1, 5. The time limit for one episode is 35 steps.

This paper investigates the imitation of diverse past trajectories and how that leads further exploration and avoids getting stuck at a sub-optimal behavior. Specifically, we propose to use a buffer of the past trajectories to cover diverse possible directions. Then we learn a trajectory-conditioned policy to imitate any trajectory from the buffer, treating it as a demonstration. After completing the demonstration, the agent performs random exploration. In this way, the exploration frontier is

Preprint. Work in Progress.

indirectly pushed further in diverse directions. The explored region is gradually expanded and the chance of finding a globally optimal solution increases. After finding trajectories with near-optimal (or high) total rewards, we imitate them to learn the final policy that achieves good performance.

Our main contributions are summarized as follows: (1) We propose a novel architecture for a trajectory-conditioned policy to imitate diverse demonstrations. (2) We demonstrate the importance of imitating diverse past experiences to indirectly drive exploration to different regions of the environment, by comparing it with existing approaches on various sparse-reward reinforcement learning tasks with the discrete and continuous action space. (3) In particular, we achieve a performance competitive with the state-of-the-art on hard exploration Atari game of Montezuma's Revenge without using expert demonstrations or resetting to an arbitrary state.

2 Related Work

Self-Imitation Learning a good policy by imitating past experiences has been discussed in [27, 18, 20], where the agent is trained to imitate only the high-reward trajectories with the SIL [27] or GAIL objective [18]. In contrast, we store the past trajectories ending with diverse states in the buffer, because trajectories with low reward in the short term could lead to high reward in the long term, and thus following a diverse set of trajectories could be beneficial for discovering optimal solutions. Furthermore, our method focuses on explicit trajectory-level imitation while existing methods use sampled state-action pairs from the buffer to update the policy. Gangwani et al. [18] proposed to learn multiple diverse policies in a SIL framework using the Stein Variational Policy Gradient with the Jensen-Shannon kernel. Empirically, their exploration method can be limited by the number of policies learned simultaneously and the exploration performance of each single policy, as shown in Appendix F.

Exploration In a high level, exploration methods [48, 46, 47, 2, 9, 30, 42, 29] in RL tend to award a bonus (via intrinsic reward) to encourage an agent to visit novel states. Recently this idea was scaled up to large state spaces by utilizing approximation techniques [45], density models [5] and inverse dynamic models to localize the agent [11] or a random network to evaluate the novelty of a state [8]. We propose that instead of directly taking a quantification of novelty as an intrinsic reward signal, one can encourage exploration by rewarding the agent when it successfully follows demonstration trajectories that would lead to novel states. Go-Explore [16] also shows the benefit of exploration by returning to a promising state to solve hard-exploration Atari games, though its success relies on the assumption that the environment is deterministic and resettable. Here resetting to an arbitrary state can result in two to three orders of magnitude reduction in sample complexity, thus giving an unfair advantage over methods that do not make use of resetting; more importantly, such resetting is often infeasible in real environments. As discussed in Appendix H when using a perfect goal-conditioned policy, as opposed to a direct 'reset' function, Go-Explore could not explore as efficiently as our method. Previous works attempted reaching a goal state by learning a set of sub-policies [22] or a goal-conditioned policy in pixel observation space [13]. However, these policies do not perform well on sparse-reward environments such as Montezuma's Revenge. Our method provides an indirect 'reset' function in a stochastic environment by imitating a trajectory using a goal-conditioned policy. Several studies [19, 17, 34] seek a diversity of exploration by maximizing the entropy of mixture skill policies or generated goal states. However, these methods mainly focus on learning diverse skills or goal states as a continuous latent variable, and the experiments are performed mainly on a variety of simulated robotic tasks with a relatively simple observation space.

Goal-Conditioned Policy Many previous works [1, 26, 37, 32] studied learning a goal-conditioned policy. Similarly to hindsight experience replay [1], our approach samples goal states from past experiences. However, we use past experiences through both supervised learning and reinforcement learning objectives. Compared to a single goal state the state trajectory leads the agent to follow a demonstration in a soft order to reach the goal state even far away from the current state. Our method shares the same motivation as Duan et al. [14] which use an attention model over the demonstration and follow the idea of the sequence-to-sequence model [43, 10]. However, our architecture is simpler since it does not use an attention model over the current observation and it is evaluated on a variety of environments, while Duan et al. [14] mainly focuses on the block stacking task.

Imitation Learning The goal of imitation learning is to train a policy to mimic a given demonstration. For example, DQfD [21], Ape-X DQfD [33], TDC+CMC [3], and LfSD [36] achieve good results on the hard-exploration Atari games using human demonstrations. In contrast, our method does not rely on expert trajectories; instead, it treats the agent's own past trajectories as demonstrations.

Algorithm 1 Diverse Self-Imitation Learning with Trajectory-Conditioned Policy

Initialize parameter θ for the trajectory-conditioned policy $\pi_\theta(a_t|e_{\leq t}, o_t, g)$
Initialize the trajectory buffer $\mathcal{D} \leftarrow \emptyset$ ***# Store diverse past trajectories***
Initialize set of transitions in the current episode $\mathcal{E} \leftarrow \emptyset$ ***# Store current episode trajectory***
Initialize set of on-policy samples $\mathcal{F} \leftarrow \emptyset$ ***# Store data for on-policy PPO update***
Initialize demonstration trajectory $g \leftarrow \emptyset$
for each iteration **do**
  for each step t **do**
    Observe $s_t = \{o_t, e_t\}$ and choose an action $a_t \sim \pi_\theta(a_t|e_{\leq t}, o_t, g)$
    Execute action a_t in the environment to get r_t, o_{t+1}, e_{t+1}
    Store transition $\mathcal{E} \leftarrow \mathcal{E} \cup \{(o_t, e_t, a_t, r_t)\}$
    # Positive reward if agent follows demonstration g
    # No reward after agent completes g ***and then takes random exploration***
    Determine r_t^{DTSIL} by comparing $e_{\leq t+1}$ with g (Eq. 1)
    Store on-policy sample $\mathcal{F} \leftarrow \mathcal{F} \cup \{(o_t, e_t, a_t, g, r_t^{\text{DTSIL}})\}$
  end for
  if s_{t+1} is terminal **then**
    $\mathcal{D} \leftarrow$ UpdateBuffer($\mathcal{D}, \mathcal{E}$) (Alg. 2)
    Clear current episode trajectory $\mathcal{E} \leftarrow \emptyset$
    $g \leftarrow$ SampleDemo($\mathcal{D}$) (Alg. 3)
  end if
  # Perform PPO update using on-policy samples and supervised learning update
  $\theta \leftarrow \theta - \eta \nabla_\theta (\mathcal{L}^{\text{RL}} + \beta \mathcal{L}^{\text{SL}})$ (Eq. 2, Eq. 3)
  Clear on-policy samples $\mathcal{F} \leftarrow \emptyset$
end for

3 Method

The main idea of our method is to maintain a buffer of diverse trajectories collected during training and to train a trajectory-conditioned policy by leveraging reinforcement learning and supervised learning to roughly follow demonstration trajectories sampled from the trajectory buffer. Therefore, the agent is encouraged to explore beyond various visited states in the environment and gradually push its exploration frontier further. Ideally, we want to find trajectories with near-optimal total rewards. After this, we fine-tune the final policy to imitate the best trajectories found during training.[1] We name our method as Diverse Trajectory-conditioned Self-Imitation Learning (DTSIL).

3.1 Background and Notation

We first briefly describe the standard reinforcement learning setting that we build our approach upon. Specifically, at each time step t, an agent observes a state s_t and selects an action $a_t \in \mathcal{A}$, where it receives a reward $r_t = R(s_t, a_t)$ when transitioning from a state $s_t \in \mathcal{S}$ to a next state $s_{t+1} \in \mathcal{S}$, where $\mathcal{S}$ is a set of all states and $\mathcal{A}$ is a set of all actions. The goal of policy-based RL algorithms is to find a policy $\pi_\theta(a|s)$ parameterized by θ that maximizes the expected discounted return $\mathbb{E}_{\pi_\theta}[\sum_{t=0}^{T} \gamma^t r_t]$, where $\gamma \in (0, 1]$ is a discount factor.

In our work, we assume a state s_t includes the agent's observation o_t (e.g., raw pixel image) and a high-level abstract state embedding e_t (e.g., the agent's location in the abstract space). The embedding e_t may be learnable from o_t (or $o_{\leq t}$), but in this work, we assume that a high-level embedding is provided as a part of s_t. A *trajectory-conditioned policy* $\pi_\theta(a_t|e_{\leq t}, o_t, g)$ (which we refer to as $\pi_\theta(\cdot|g)$ in shorthand notation) takes a sequence of state embeddings $g = \{e_1^g, e_2^g, \cdots, e_{|g|}^g\}$ as input for a demonstration, where $|g|$ is the length of the trajectory g. A sequence of the agent's past state embeddings $e_{\leq t} = \{e_1, e_2, \cdots, e_t\}$ is provided to determine which part of the demonstration has been followed. Together with the current observation o_t, it helps to determine the correct action a_t to accurately imitate the demonstration. Our goal here is to find a set of optimal state embedding sequence(s) g^* and the policy $\pi_\theta^*(\cdot|g)$ to maximize the return: $g^*, \theta^* \triangleq \arg\max_{g,\theta} \mathbb{E}_{\pi_\theta(\cdot|g)}[\sum_{t=0}^{T} \gamma^t r_t]$. For robustness we may want to find multiple near-optimal embedding sequences with similar returns and a trajectory-conditioned policy for executing them.

[1] In our implementation, we train a trajectory-conditioned policy to imitate the best trajectories. Alternatively, an unconditional stochastic policy could also be trained to imitate them.

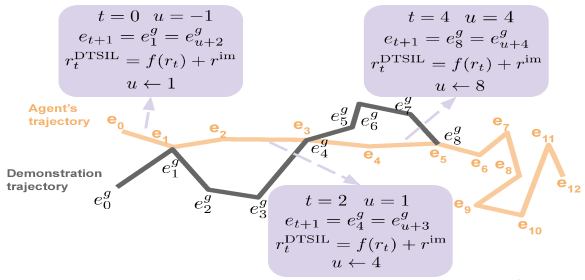

Figure 2: An example showing the calculation of r_t^{DTSIL} based on u and $\Delta t = 4$. At each step t, we check the state embedding e_{t+1} to determine the reward according to Equation 1. After completing the demonstration, the agent performs the random exploration with reward 0 ($e_5 \sim e_{12}$).

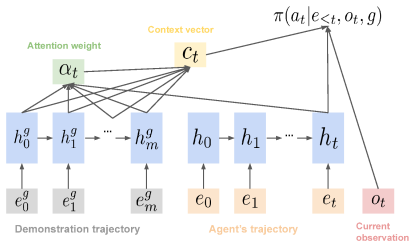

Figure 3: Architecture of the trajectory-conditioned policy.

3.2 Organizing Trajectory Buffer

We maintain a *trajectory buffer* $\mathcal{D} = \{(e^{(1)}, \tau^{(1)}, n^{(1)}), (e^{(2)}, \tau^{(2)}, n^{(2)}), \cdots\}$ of diverse past trajectories. For each embedding-trajectory-count tuple $(e^{(i)}, \tau^{(i)}, n^{(i)})$, $\tau^{(i)}$ is the best trajectory ending with a state with the high-level representation $e^{(i)}$, and $n^{(i)}$ is the number of times this state representation has been visited during training. In order to maintain the compact buffer a high-level discrete state representation $e^{(i)}$ is used (e.g., the agent's location in the discrete grid, discretized accumulated reward, etc.) and the existing entry is replaced if an improved trajectory $\tau^{(i)}$ is found.

When given a new episode $\mathcal{E} = \{(o_0, e_0, a_0, r_0), \cdots, (o_T, e_T, a_T, r_T)\}$, all the state representations $e_t (1 \leq t \leq T)$ in this episode $\mathcal{E}$ are considered because the buffer maintains all of the possible paths available for future exploration to avoid missing any possibility to find an optimal solution. If e_t is not yet stored in the buffer, $(e_t, \tau_{\leq t}, 1)$ is directly pushed into the buffer, where $\tau_{\leq t} = \{(o_0, e_0, a_0, r_0), \cdots, (o_t, e_t, a_t, r_t)\}$ is the agent's current partial episode ending with e_t. If the partial episode $\tau_{\leq t}$ is better (i.e., higher return or shorter trajectory) than the stored trajectory $\tau^{(k)}$ when reaching $e_t = e^{(k)}$, $\tau^{(k)}$ is replaced by the current trajectory $\tau_{\leq t}$. The algorithm is described in Appendix A.

3.3 Sampling Demonstrations

When learning trajectory-conditioned policy π, demonstration trajectories are sampled from the trajectory buffer. We record the count $n^{(i)}$ of how many times this state embedding is visited, and set the sampling probability as $\frac{1}{\sqrt{n^{(i)}}}$. This is inspired by count-based exploration bonus [42, 5]: we sample a trajectory that ends with a less frequently visited state because this leads the agent to reach rarely visited regions in the state space and is more promising for discovering novel states. When trajectories with satisfactory accumulated rewards have been discovered, we sample the best trajectories stored in the buffer for imitation learning.[2] These trajectories are used to train the policy to converge to a high-reward behavior. The algorithm is described in Appendix B.

3.4 Learning Trajectory-Conditioned Policy

Imitation Reward A given demonstration trajectory $g = \{e_0^g, e_1^g, \cdots, e_{|g|}^g\}$ is used to provide rewards for imitation g, similarly to the imitation learning method introduced by [3]. At the beginning of an episode, the index u of the last visited state embedding in the demonstration is initialized as $u = -1$. At each step t, if the agent's new state s_{t+1} has an embedding e_{t+1} and it is the same as any of the next Δt state embeddings starting from the last visited state embedding e_u^g in the demonstration (i.e., $e_{t+1} = e_{u+j}^g$ where $1 \leq j \leq \Delta t$), then it receives a positive imitation reward r^{im}, and the index of the last visited state embedding in the demonstration is updated as $u \leftarrow u + j$. This encourages the agent to visit the state embeddings in the demonstration in a soft-order. When the last state embedding $e_{|g|}^g$ in the demonstration has been visited by the agent (i.e., $u = |g|$), then there is no further imitation reward r^{im} and the agent performs a random exploration until the episode terminates.

[2] We assume that the optimal episode reward or the threshold of an ideal episode reward is known in advance. Alternatively, we could switch between exploration and imitation by adjusting the probability of imitation in an episode as training goes on, as discussed in Appendix I.

To summarize, the agent receives a reward r_t^{DTSIL} defined as

$$r_t^{\text{DTSIL}} = \begin{cases} f(r_t) + r^{\text{im}} & \text{if } \exists j, 1 \leq j \leq \Delta t, \text{ such that } e^g_{u+j} = e_{t+1} \\ 0 & \text{otherwise,} \end{cases} \tag{1}$$

where $f(\cdot)$ is a monotonically increasing function (e.g., clipping [25]). Figure 2 illustrates the calculation of r_t^{DTSIL} and the update of u during an episode when the agent visits a state whose embedding appears in the demonstration g.

Policy Architecture For imitation learning with diverse demonstrations, we design a trajectory-conditioned policy $\pi_\theta(a_t|e_{\leq t}, o_t, g)$ that should imitate any given trajectory g. Inspired by neural machine translation methods [43, 10, 4, 23], the demonstration trajectory is the source sequence and the incomplete trajectory of the agent's state representations is the target sequence. We apply a recurrent neural network and an attention mechanism to the sequence data to predict actions that would make the agent to follow the demonstration trajectory. As illustrated in Figure 3, RNN computes the hidden features h_i^g for each state embedding e_i^g $(0 \leq i \leq m)$ in the demonstration and derives the hidden features h_t for the agent's state representation e_t. Then the attention weights α_t is computed by comparing the current agent's hidden features h_t with the demonstration's hidden features h_i^g $(0 \leq i \leq m)$. The context vector c_t is computed as an attention-weighted summation of the demonstration's hidden states to capture the relevant information in the demonstration trajectory and to predict the action a_t.

Reinforcement Learning Objective With the reward defined as r_t^{DTSIL} (Equation 1), the trajectory-conditioned policy π_θ can be trained with a policy gradient algorithm [44, 39]:

$$\mathcal{L}^{\text{RL}} = \mathbb{E}_{\pi_\theta}[-\log \pi_\theta(a_t|e_{\leq t}, o_t, g)\widehat{A}_t],$$
$$\text{where } \widehat{A}_t = \sum_{d=0}^{n-1} \gamma^d r_{t+d}^{\text{DTSIL}} + \gamma^n V_\theta(e_{\leq t+n}, o_{t+n}, g) - V_\theta(e_{\leq t}, o_t, g), \tag{2}$$

where the expectation $\mathbb{E}_{\pi_\theta}$ indicates the empirical average over a finite batch of on-policy samples and n denotes the number of rollout steps taken in each iteration. We use Proximal Policy Optimization (PPO) [39] as an actor-critic policy gradient algorithm for our experiments.

Supervised Learning Objective To improve trajectory-conditioned imitation learning and to better leverage the past trajectories, we propose a supervised learning objective. By sampling a trajectory τ from the buffer $\mathcal{D}$, the demonstration trajectory (i.e., source sequence) is formulated as $g = \{e_0, e_1, \cdots, e_{|g|}\}$ and assumed that the agent's incomplete trajectory (i.e., target sequence) is the partial trajectory $g_{\leq t} = e_{\leq t} = \{e_0, e_1, \cdots, e_t\}$ for any $1 \leq t \leq |g|$. Then a_t is the 'correct' action at step t for the agent to imitate the demonstration. Our supervised learning objective is to maximize the log probability of taking such actions:

$$\mathcal{L}^{\text{SL}} = -\log \pi_\theta(a_t|e_{\leq t}, o_t, g), \text{ where } g = \{e_0, e_1, \cdots, e_{|g|}\}. \tag{3}$$

4 Experiments

In the experiments, we aim to answer the following questions: (1) How well does the trajectory-conditioned policy imitate the diverse demonstration trajectories? (2) Does imitation of the past diverse experience enable the agent to further explore more diverse directions and guide the exploration to find the trajectory with a near-optimal total reward? (3) Can our proposed method aid in avoiding myopic behaviors and converge to near-optimal solutions?

4.1 Implementation Details

Our algorithm begins with an empty buffer $\mathcal{D}$ and we initialize the demonstration as a list of zero vectors. With such an input demonstration, the agent would perform random exploration to collect trajectories to fill the buffer $\mathcal{D}$. In practice, when $\mathcal{D} \neq \emptyset$, the sampled demonstration trajectory $g = \{e_0^g, e_1^g, \cdots, e_{|g|}^g\}$ could be lengthy. We present a part of the demonstration as the input into the policy, similar to translating a paragraph sentence by sentence. More specifically, we first input $\{e_0^g, e_1^g, \cdots, e_m^g\}$ $(m \leq |g|)$ into the policy. When the index of the agent's last visited state embedding in the demonstration u belongs to $\{m - \Delta t, \cdots, m\}$, we consider that the agent has accomplished this part of the demonstration, and therefore switch to the next part $\{e_u^g, e_{u+1}^g, \cdots, e_{u+m}^g\}$. We repeat this process until the last part of the demonstration. If the last part $\{e_u^g, e_{u+1}^g, \cdots, e_{|g|}^g\}$ is less than $m+1$ steps long, we pad the sequence with zero vectors.

A reward function $f(r_t) = r_t$ is used on the Mujoco domain, and $f(r_t) = 2 \cdot \text{clip}(r_t, -1, 1)$ is used on other environments. We use $r^{\text{im}} = 0.1$ as a reward to encourage the imitation. Further details

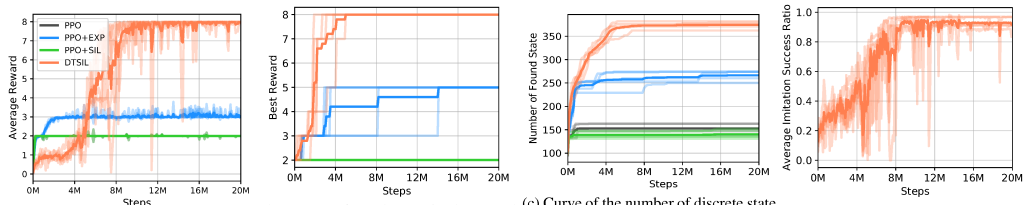

(a) Curve of the average episode reward (b) Curve of the best episode reward found during training (c) Curve of the number of discrete state embeddings found during training (d) Curve of the average imitation success ratio

Figure 4: Learning curves on Key-Door-Treasure domain averaged over 5 runs, where the curves in dark colors are average over the 5 curves in light colors. The x-axis and y-axis correspond to the number of steps and statistics about the performance, respectively. The average reward and average imitation success ratio are the mean values over 40 recent episodes.

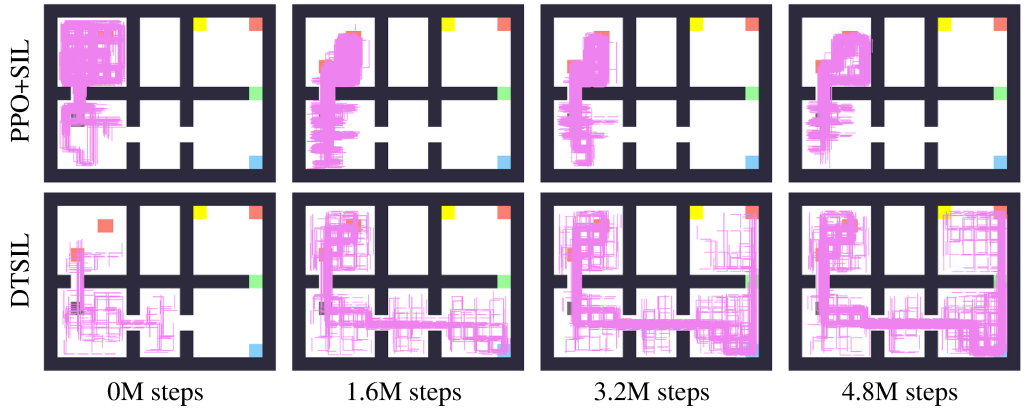

Figure 5: Visualization of the trajectories stored in the buffer for PPO+SIL and our method as training continues. The agent (gray), apple (red), key (blue), door (green), and treasure (yellow) are shown as squares for simplicity.

about the hyperparameters and the environment setup are described in the Appendix. Our PPO is based on OpenAI's implementation [12].

We compare our method with the following baselines: (1) PPO: Proximal Policy Optimization [39]; (2) PPO+EXP: PPO with count-based exploration bonus $r_t^+ = 1/\sqrt{N(e_t)}$, where $N(e)$ is the number of times the discrete state representation e was visited during training; (3) PPO+SIL: PPO with Self-Imitation Learning [27].

4.2 Key-Door-Treasure Domain

The Key-Door-Treasure domain (shown in Figure 1) is a simple grid-world environment with deceptive rewards that can lead the agent to a local optima. An observation consists of the agent's location (x_t, y_t) and binary variables showing whether the agent has gotten the key, opened the door, or collected the treasure. A state is represented as the agent's location and the cumulative reward: $e_t = (x_t, y_t, \sum_{i=1}^{t} r_i)$, indicating the location of the agent and identifying the collected objects.

As shown in Figure 4a, both PPO and PPO+SIL agents are stuck with the suboptimal policy of collecting the first two apples (Figure 1). The PPO+EXP agent explores further and gathers the two apples and one key. Our method learns to collect objects on the right side of the maze and achieves the highest total reward of 8 within the time limit. Figure 4b and Figure 4c show the highest episode reward and the number of different state embeddings found during training. The PPO+EXP agent occasionally scores the episode reward of 5 because its initial location is not fixed and it is possible to luckily collect three apples, pick up the key and open the door within the time limit if it started from a certain location. However, when the agent explores to find the reward signals in the right site of the maze, it never finds a path to the treasure.

In Figure 4d, we show the average success ratio of the imitation during training. It is defined as follows: for a given demonstration $g = \{e_0^g, e_1^g, \cdots, e_{|g|}^g\}$, let u be the index of the last visited state embedding in g, then the success ratio of imitating g is $\frac{u}{|g|}$ (i.e., the portion of trajectory imitated). Ideally, we want the success ratio to be 1.0, which indicates that the trajectory-conditioned policy could successfully follow any given demonstration from the buffer. At 3M steps, when the trajectories with the optimal total reward 8 are found, we start to sample the best trajectories stored in the buffer

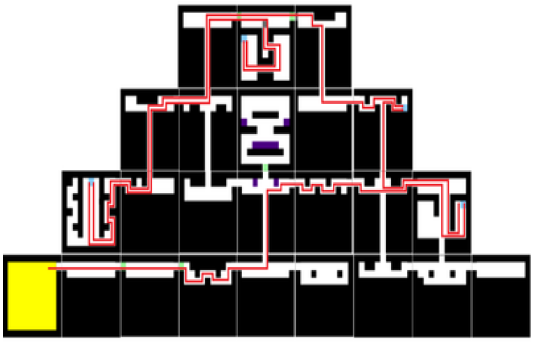

Figure 6: Map of Toy Montezuma's Revenge, where we show the agent (gray), key(blue), door(green), and treasure (yellow) as squares. The rewards are 100, 300, and 10000, respectively. An optimal path with the highest total reward of 11,600 is shown as a red line.

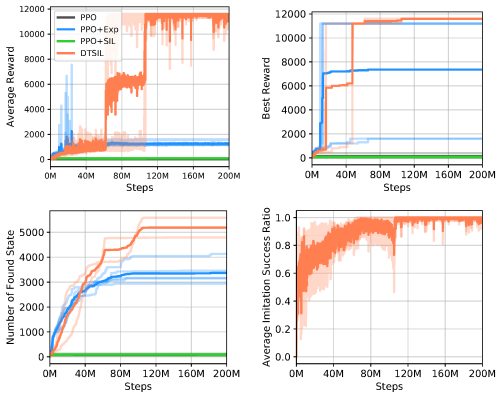

Figure 7: Learning curves on Toy Montezuma's Revenge averaged over 5 runs.

as demonstrations. Our trajectory-conditioned policy can imitate them well with a success ratio around 1.0.

Figure 5 visualizes a learning process. PPO+SIL fails on this task because the agent quickly exploits a good experience of collecting the apples but the buffer is filled with the trajectories exploring the nearby region. On the contrary, our method maintains a buffer of diverse trajectories which are used as demonstrations to guide the agent to explore different regions and discover an optimal behavior.

4.3 Toy Montezuma's Revenge Domain

We evaluate our method on a more challenging domain, Toy Montezuma's Revenge [35], which requires a more sophisticated strategy to explore the environment. As shown in Figure 6, there are 24 rooms similar to the layout of the first level of Atari Montezuma's Revenge, with a discrete grid for each room. The agent should navigate the labyrinth to locate the keys, unlock the doors and reach the goal (the treasure room). The observation is represented by the agent's location and cumulative episode reward. The state representation $e_t = (\mathrm{room}_t, x_t, y_t, \sum_{i=1}^{t} r_i)$ is the same as the observation.

The learning curve of the averaged episode reward in Figure 7 shows that PPO, PPO+SIL, and PPO+EXP could not learn a policy to reach the goal. The PPO+EXP agent occasionally finds a trajectory with the total reward of 11,200 reaching the treasure room, but fails to exploit this experience. On the other hand, our method learns a good behavior of not only reaching the goal room, but also collecting all of the keys to achieve an optimal total reward of 11,600.

4.4 Atari Montezuma's Revenge

We evaluate our method on the hard-exploration game Montezuma's Revenge of the Arcade Learning Environment (ALE) [6, 24], which consists of multiple levels and 24 rooms in each level (map shown in Figure 11). The observation is a frame of raw pixel images, and we use the state representation $e_t = (\mathrm{level}_t, \mathrm{room}_t, x_t, y_t, k_t)$ consisting of the agent's ground truth location (obtained from RAM) and the number of keys it holds. We use the number of keys k_t to reduce the size of the state embedding space. With such a discrete state representation, we can also add a count-based exploration bonus $r_t^+ = 1/\sqrt{N(e_t)}$ to r_t^{DTSIL} in Equation 1, in order to expedite exploration. This variant is denoted as 'DTSIL+EXP'.

As shown in Figure 8, in the early stage, the average episode reward of DTSIL+EXP is worse than PPO+EXP because our policy is trained to imitate diverse demonstrations instead of directly maximizing the environment reward. Compared to PPO+EXP, our agent is not eager to blindly follow high-reward paths since the path with a relatively low score in the short term might lead to higher rewards in the long term. As training continues, our method successfully discovers trajectories to accomplish the first level with a satisfactory total reward of more than 20,000. After switching to imitating the best trajectories in the buffer by sampling them as demonstrations, the average episode reward dramatically increases to over 25,000[3]. Table 1 compares our proposed method with previous

[3]Demo videos of the learned policies for both PPO+EXP and DTSIL+EXP are available at: https://sites.google.com/view/diverse-sil/home.

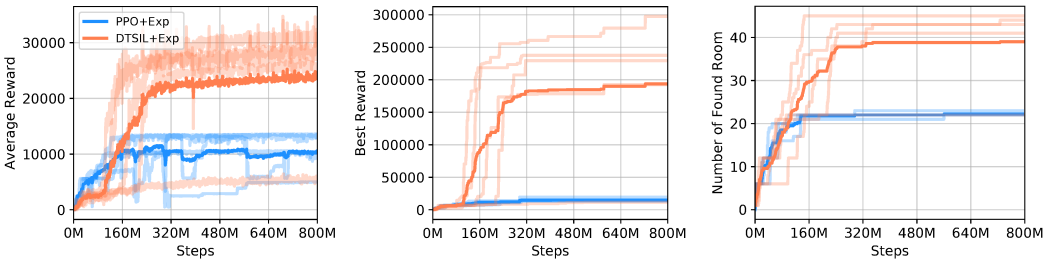

Figure 8: Learning curves of the average episode reward, the best episode reward, and the number of different rooms found on Atari Montezuma's Revenge, averaged over 5 runs. For DTSIL+EXP, 29,278 is the maximum mean score (averaged over 40 recent episodes) achieved over 800M environment timesteps (3,200M frames) during training, averaged over 5 runs. DTSIL+EXP discovers around 40 rooms on average while PPO+EXP never finds a path to pass through all 24 rooms at the first level and then proceed to the next level.

Method	DTSIL	Abstract-HRL	A2C+CoEX+RAM	SmartHash	DeepCS	A2C+SIL	PPO+CoEX	RND
Score	**29,278**	11,500	6,600	5,661	3,500	2,500	11,618	10,070

Table 1: Comparison with the state-of-the-art results on Montezuma's Revenge without using expert demonstrations or resetting to an arbitrary state. DTSIL, Abstract-HRL [22], A2C+CoEX+RAM [11], SmartHash [45], and DeepCS [41] make use of information from RAM, while A2C+SIL [27], PPO+CoEX [11], and RND [8] do not use RAM information. The score is averaged over multiple runs, gathered from each paper.

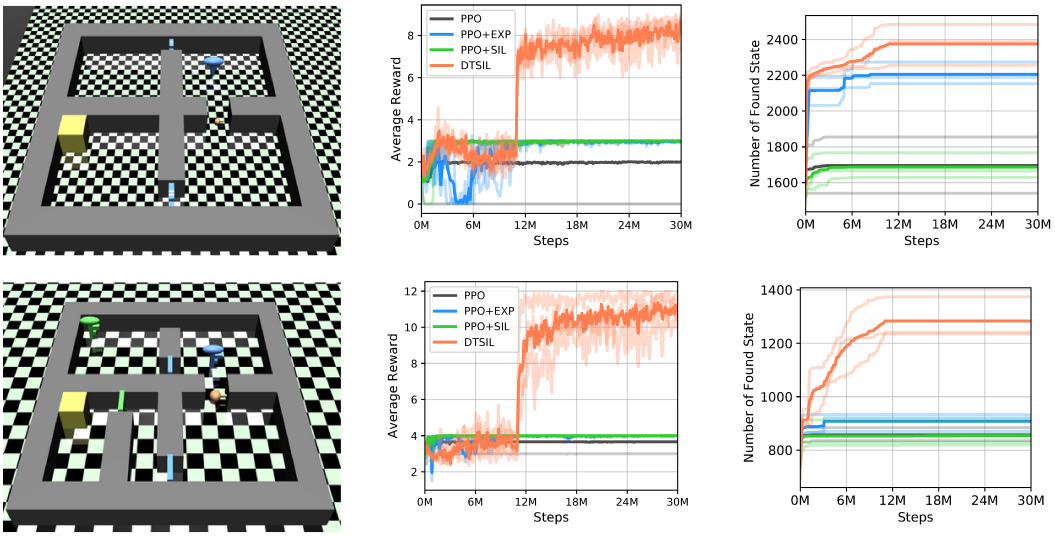

Figure 9: Point Maze in Mujoco domain. The reward for getting the key, opening the door, and collecting the treasure (yellow block) is 1, 2, and 6 respectively. Once the agent uses a key to open a door with the matched color, the key cannot be used again. The learning curve of the episode reward and the number of found state representations are averaged over 3 independent runs.

works without using any expert demonstration or resetting to an arbitrary state, where our approach significantly outperforms the other approaches.

4.5 Mujoco

We evaluate our proposed method on continuous control tasks. We adapt the maze environment introduced in [15] to construct a set of challenging tasks which require the point mass agent to collect the key, open the door with the same color and finally reach the treasure to get a high score. One key cannot be used any more once it was used before to open a door with the same color, which makes the agent to be easily trapped. A visualization of these environments is shown in Figure 9.

The observation is the agent's location and range sensor reading about nearby objects. The state representation is $e_t = (x_t, y_t, \sum_{i=1}^{t} r_i)$. We discretize the continuous variables x_t, y_t as integers $\lfloor x_t \rfloor, \lfloor y_t \rfloor$ when determining whether a state embedding has a corresponding trajectory stored in the buffer and whether a state embedding in the demonstration has been visited by the agent.

In the first maze of Figure 9, the agent can easily get the blue key near its initial location and open the blue door in the upper part. However, the optimal path is to bring the key to open the blue door in the bottom and obtain the treasure, reaching an episode reward of 9. In the second maze, the agent should bring the blue key and pick up the green key while avoiding opening the blue door in the upper part. Then, the green and blue key can open the two doors at the bottom of the maze, which results in the total reward of 12. The learning curves in Figure 9 show that PPO, PPO+EXP, and PPO+SIL may get stuck at a sub-optimal behavior, whereas our policy eventually converges to the behavior achieving the optimal episode reward.

5 Conclusion

This paper proposes to learn diverse policies by imitating diverse trajectory-level demonstrations through count-based exploration over these trajectories. Imitation of diverse past trajectories can guide the agent to rarely visited states and encourages further exploration of novel states. We demonstrate that on a variety of environments with local optima, our method significantly improves self-imitation learning (SIL). It avoids prematurely converging to a sub-optimal solution and learns a near-optimal behavior in order to achieve a high total reward.

References

[1] 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*, pages 5048–5058, 2017.

[2] P. Auer. Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine Learning Research*, 3(Nov):397–422, 2002.

[3] Y. Aytar, T. Pfaff, D. Budden, T. Paine, Z. Wang, and N. de Freitas. Playing hard exploration games by watching youtube. In *Advances in Neural Information Processing Systems*, pages 2930–2941, 2018.

[4] D. Bahdanau, K. Cho, and Y. Bengio. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*, 2014.

[5] M. Bellemare, S. Srinivasan, G. Ostrovski, T. Schaul, D. Saxton, and R. Munos. Unifying count-based exploration and intrinsic motivation. In *Advances in Neural Information Processing Systems*, pages 1471–1479, 2016.

[6] M. G. Bellemare, Y. Naddaf, J. Veness, and M. Bowling. The arcade learning environment: An evaluation platform for general agents. *Journal of Artificial Intelligence Research*, 47:253–279, 2013.

[7] Y. Burda, H. Edwards, D. Pathak, A. Storkey, T. Darrell, and A. A. Efros. Large-scale study of curiosity-driven learning. *arXiv preprint arXiv:1808.04355*, 2018.

[8] Y. Burda, H. Edwards, A. Storkey, and O. Klimov. Exploration by random network distillation. *arXiv preprint arXiv:1810.12894*, 2018.

[9] N. Chentanez, A. G. Barto, and S. P. Singh. Intrinsically motivated reinforcement learning. In *Advances in neural information processing systems*, pages 1281–1288, 2005.

[10] K. Cho, B. Van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*, 2014.

[11] J. Choi, Y. Guo, M. Moczulski, J. Oh, N. Wu, M. Norouzi, and H. Lee. Contingency-aware exploration in reinforcement learning. *arXiv preprint arXiv:1811.01483*, 2018.

[12] P. Dhariwal, C. Hesse, O. Klimov, A. Nichol, M. Plappert, A. Radford, J. Schulman, S. Sidor, Y. Wu, and P. Zhokhov. Openai baselines. `https://github.com/openai/baselines`, 2017.

[13] H. Dong, J. Mao, X. Cui, and L. Li. Explicit recall for efficient exploration, 2019.

[14] Y. Duan, M. Andrychowicz, B. Stadie, O. J. Ho, J. Schneider, I. Sutskever, P. Abbeel, and W. Zaremba. One-shot imitation learning. In *Advances in neural information processing systems*, pages 1087–1098, 2017.

[15] Y. Duan, X. Chen, R. Houthooft, J. Schulman, and P. Abbeel. Benchmarking deep reinforcement learning for continuous control. In *International Conference on Machine Learning*, pages 1329–1338, 2016.

[16] A. Ecoffet, J. Huizinga, J. Lehman, K. O. Stanley, and J. Clune. Go-explore: a new approach for hard-exploration problems. *arXiv preprint arXiv:1901.10995*, 2019.

[17] B. Eysenbach, A. Gupta, J. Ibarz, and S. Levine. Diversity is all you need: Learning skills without a reward function. *arXiv preprint arXiv:1802.06070*, 2018.

[18] T. Gangwani, Q. Liu, and J. Peng. Learning self-imitating diverse policies. *arXiv preprint arXiv:1805.10309*, 2018.

[19] K. Gregor, D. J. Rezende, and D. Wierstra. Variational intrinsic control. *arXiv preprint arXiv:1611.07507*, 2016.

[20] Y. Guo, J. Oh, S. Singh, and H. Lee. Generative adversarial self-imitation learning. *arXiv preprint arXiv:1812.00950*, 2018.

[21] T. Hester, M. Vecerik, O. Pietquin, M. Lanctot, T. Schaul, B. Piot, D. Horgan, J. Quan, A. Sendonaris, I. Osband, et al. Deep q-learning from demonstrations. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.

[22] E. Z. Liu, R. Keramati, S. Seshadri, K. Guu, P. Pasupat, E. Brunskill, and P. Liang. Learning abstract models for long-horizon exploration, 2019.

[23] M.-T. Luong, H. Pham, and C. D. Manning. Effective approaches to attention-based neural machine translation. *arXiv preprint arXiv:1508.04025*, 2015.

[24] M. C. Machado, M. G. Bellemare, E. Talvitie, J. Veness, M. Hausknecht, and M. Bowling. Revisiting the arcade learning environment: Evaluation protocols and open problems for general agents. *Journal of Artificial Intelligence Research*, 61:523–562, 2017.

[25] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, S. Petersen, C. Beattie, A. Sadik, I. Antonoglou, H. King, D. Kumaran, D. Wierstra, S. Legg, and D. Hassabis. Human-level control through deep reinforcement learning. *Nature*, 2015.

[26] A. Nair, D. Chen, P. Agrawal, P. Isola, P. Abbeel, J. Malik, and S. Levine. Combining self-supervised learning and imitation for vision-based rope manipulation. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2146–2153. IEEE, 2017.

[27] J. Oh, Y. Guo, S. Singh, and H. Lee. Self-imitation learning. *arXiv preprint arXiv:1806.05635*, 2018.

[28] G. Ostrovski, M. G. Bellemare, A. van den Oord, and R. Munos. Count-based exploration with neural density models. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 2721–2730. JMLR. org, 2017.

[29] P.-Y. Oudeyer and F. Kaplan. What is intrinsic motivation? a typology of computational approaches. *Frontiers in neurorobotics*, 1:6, 2009.

[30] P.-Y. Oudeyer, F. Kaplan, and V. V. Hafner. Intrinsic motivation systems for autonomous mental development. *IEEE transactions on evolutionary computation*, 11(2):265–286, 2007.

[31] D. Pathak, P. Agrawal, A. A. Efros, and T. Darrell. Curiosity-driven exploration by self-supervised prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 16–17, 2017.

[32] D. Pathak, P. Mahmoudieh, G. Luo, P. Agrawal, D. Chen, Y. Shentu, E. Shelhamer, J. Malik, A. A. Efros, and T. Darrell. Zero-shot visual imitation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 2050–2053, 2018.

[33] T. Pohlen, B. Piot, T. Hester, M. G. Azar, D. Horgan, D. Budden, G. Barth-Maron, H. van Hasselt, J. Quan, M. Večerík, et al. Observe and look further: Achieving consistent performance on atari. *arXiv preprint arXiv:1805.11593*, 2018.

[34] V. H. Pong, M. Dalal, S. Lin, A. Nair, S. Bahl, and S. Levine. Skew-fit: State-covering self-supervised reinforcement learning. *arXiv preprint arXiv:1903.03698*, 2019.

[35] M. Roderick, C. Grimm, and S. Tellex. Deep abstract q-networks. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, pages 131–138. International Foundation for Autonomous Agents and Multiagent Systems, 2018.

[36] T. Salimans and R. Chen. Learning montezuma's revenge from a single demonstration. *arXiv preprint arXiv:1812.03381*, 2018.

[37] T. Schaul, D. Horgan, K. Gregor, and D. Silver. Universal value function approximators. In *International conference on machine learning*, pages 1312–1320, 2015.

[38] J. Schmidhuber. Curious model-building control systems. In *[Proceedings] 1991 IEEE International Joint Conference on Neural Networks*, pages 1458–1463. IEEE, 1991.

[39] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.

[40] B. C. Stadie, S. Levine, and P. Abbeel. Incentivizing exploration in reinforcement learning with deep predictive models. *arXiv preprint arXiv:1507.00814*, 2015.

[41] C. Stanton and J. Clune. Deep curiosity search: Intra-life exploration improves performance on challenging deep reinforcement learning problems. *arXiv preprint arXiv:1806.00553*, 2018.

[42] A. L. Strehl and M. L. Littman. An analysis of model-based interval estimation for markov decision processes. *Journal of Computer and System Sciences*, 74(8):1309–1331, 2008.

[43] I. Sutskever, O. Vinyals, and Q. V. Le. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pages 3104–3112, 2014.

[44] R. S. Sutton, D. A. McAllester, S. P. Singh, and Y. Mansour. Policy gradient methods for reinforcement learning with function approximation. In *Advances in neural information processing systems*, pages 1057–1063, 2000.

[45] H. Tang, R. Houthooft, D. Foote, A. Stooke, O. X. Chen, Y. Duan, J. Schulman, F. DeTurck, and P. Abbeel. # exploration: A study of count-based exploration for deep reinforcement learning. In *Advances in neural information processing systems*, pages 2753–2762, 2017.

[46] S. B. Thrun. Efficient exploration in reinforcement learning. 1992.

[47] S. B. Thrun and K. Möller. Active exploration in dynamic environments. In *Advances in neural information processing systems*, pages 531–538, 1992.

[48] J. urgen Schmidhuber. Adaptive confidence and adaptive curiosity. Technical report, Citeseer, 1991.

Appendix to "Efficient Exploration with Self-Imitation Learning via Trajectory-Conditioned Policy"

A Algorithm of Trajectory Buffer Update

Algorithm 2 Update Trajectory Buffer

Input: the trajectory buffer $\mathcal{D} = \{(e^{(1)}, \tau^{(1)}, n^{(1)}), (e^{(2)}, \tau^{(2)}, n^{(2)}), \cdots\}$
Input: the current episode $\mathcal{E} = \{(o_0, e_0, a_0, r_0), (o_1, e_1, a_1, r_1), \cdots, (o_T, e_T, a_T, r_T)\}$
 # Consider all the states in $\mathcal{E}$
for each step t **do**
  # Consider state s_t ***and partial episode*** $\tau_{\leq t} = \{(o_0, e_0, a_0, r_0), \cdots, (o_t, e_t, a_t, r_t)\}$
  if there exists $(e^{(k)}, \tau^{(k)}, n^{(k)}) \in \mathcal{D}$ where $e^{(k)} = e_t$ **then**
    # Compare partial episode $\tau_{\leq t}$ ***with stored trajectory*** $\tau^{(k)}$
    if $\tau_{\leq t}$ has higher total reward or reaches the same total reward with less steps **then**
      $\tau^{(k)} \leftarrow \tau_{\leq t} = \{(o_0, e_0, a_0, r_0), (o_1, e_1, a_1, r_1), \cdots, (o_t, e_t, a_t, r_t)\}$
    end if
    $n^{(k)} \leftarrow n^{(k)} + 1$
  else
    $\mathcal{D} \leftarrow \mathcal{D} \cup (e_t, \tau_{\leq t}, 1)$ where $\tau_{\leq t} = \{(o_0, e_0, a_0, r_0), (o_1, e_1, a_1, r_1), \cdots, (o_t, e_t, a_t, r_t)\}$
  end if
end for
return $\mathcal{D}$

B Algorithm of Sampling Demonstrations

Algorithm 3 Sample Demonstration Trajectories

Input: the trajectory buffer $\mathcal{D} = \{e^{(1)}, \tau^{(1)}, n^{(1)}), (e^{(2)}, \tau^{(2)}, n^{(2)}), \cdots\}$
if trajectories with satisfactory reward have been found **then**
  # sample the trajectories reaching near-optimal score in the buffer
  $g \leftarrow \{e_0, e_1, \cdots, e_{|g|}\}$ for all $(o_t, e_t, a_t, r_t) \in \tau^{best}$
else
  Calculate probability distribution $p \leftarrow [\frac{1}{\sqrt{n^{(1)}}}, \frac{1}{\sqrt{n^{(2)}}}, \cdots]$
  $p \leftarrow \frac{p}{\sum_j p_j}$
  Sample $(e, \tau, n) \sim$ Categorical$(\mathcal{D}, p)$
  $g \leftarrow \{e_0, e_1, \cdots, e_{|g|}\}$ for all $(o_t, e_t, a_t, r_t) \in \tau$
end if
return g

C Hyperparameters

The hyper-parameters for our proposed method used in each experiment are listed in Table 2.

Environment	Key-Door-Treasure	Toy MontezumaRevenge	Atari MontezumaRevenge	Mujoco
Learning Rate η	2.5e-4	2.5e-4	2.5e-4	1e-4
Δt	2	4	8	8
Length of demonstration input part m	10	10	10	10
Wight of supervised learning loss β	10 decreases to 1 when action prediction accuracy >0.85	10 decreases to 1 when action prediction accuracy >0.85	10 decreases to 1 when action prediction accuracy >0.75	0
Satisfactory episode reward	8	11600	20000	9 or 12

Table 2: Hyper-parameters on various environments for our experiments.

D Environment Setting

On Atari Montezuma's Revenge, the number of keys the agent holds at step t is approximated by $\sum_{i=1}^{t} \mathbf{1}_{r_i=100} - \sum_{i=1}^{t} \mathbf{1}_{r_i=300}$ because the agent gets one key with reward 100 and loses one key opening a door with reward 300. More details of the environment setting are summarized in Table 3.

Environment	Key-Door-Treasure	Toy MontezumaRevenge	Atari MontezumaRevenge	Mujoco
Observation	agent's location (x, y) in 17x13 grid and binary variables indicating whether key, door or treasure is collected	agent's location (room, x, y) in 24x11x11 grid and accumulated reward	stacked most recent 4 gray observations with shape 84x84x4	agent's location (x,y) in 22x22 space and range sensor reading about nearby objects
State Representation	$(x_t, y_t, \sum_{i=1}^{t} r_i)$	$(\text{room}_t, x_t, y_t, \sum_{i=1}^{t} r_i)$	$(\text{level}_t, \text{room}_t, x_t, y_t, k_t)$ where (x_t, y_t)in 9x9 grid	$(x_t, y_t, \sum_{i=1}^{t} r_i)$
Action	5 discrete actions: up, down, left, right, noop	5 discrete actions: up, down, left, right, noop	18 discrete actions: noop, fire, left, ...	(dx, dy) in continuous action space
Reward	key +1 door +1 treasure +5 apple +1	key +100 door +300 treasure +10000	key +100 door +300 ...	key +1 door +2 treasure +6
Time limit	35 steps	1000 steps	4500 steps	1000 steps
Stochasticity	take 3 random steps before the episode starts	take 5 random steps before the episode starts	MontezumaRevenge--NoFrameskip-v4; take a random number between 0 and 30 of noop actions before the episode starts	take random normal noise from the agent's initial position

Table 3: The setting on various environments for our experiments.

E Details of Network Architecture and Training Process

In the trajectory-conditioned policy, we first embed the input state e_t (or e_i^g) with a fully-connected layer with 64 units. Next, RNN with gated recurrent units (GRU) computes the feature h_t (or h_i^g) with 128 units. The attention weight α_t is calculated based on the Bahdanau attention mechanism [4]. The concatenation of the context vector c_t, the hidden feature of agent's current state h_t, and convolutional features from the observation are used to predict $\pi(a_t|e_{\leq t}, o_t, g)$. For experiments on the Key-Door-Treasure domain, Toy Montezuma's Revenge, and Mujoco, the features from o_t are not required for the policy. However, on the Atari game of Montezuma's Revenge, it is necessary to take the raw observation o_t as input into policy because the location information in $e_{\leq t}$ could not guide the agent to use the temporal context (e.g. avoiding moving skulls and passing laser gates).

During the imitation phase in our implementation, we sample the top 1% best trajectories (with the highest total reward and least number of steps) in the buffer as demonstrations, and train our trajectory-conditioned policy to imitate them and achieve the good performance.

F Comparison with Learning Diverse Policies by SVPG

While the code for the Stein variational policy gradient (SVPG) in [18] has not yet been released, we replicate the method in [18] to learn diverse policies. Their experiments focus on continuous control tasks with relatively simple observation spaces with limited local optimal branches in the state space. We learn 8 diverse policies in parallel following their method on our Key-Door-Treasure domain with discrete action space. Figure 10 shows a visualization of the learning progress: the 8 policies learn to cover different regions of the environment. The method explores better than PPO+SIL, but the exploration of each individual agent is not strong enough to find the optimal path to achieve the highest episode reward.

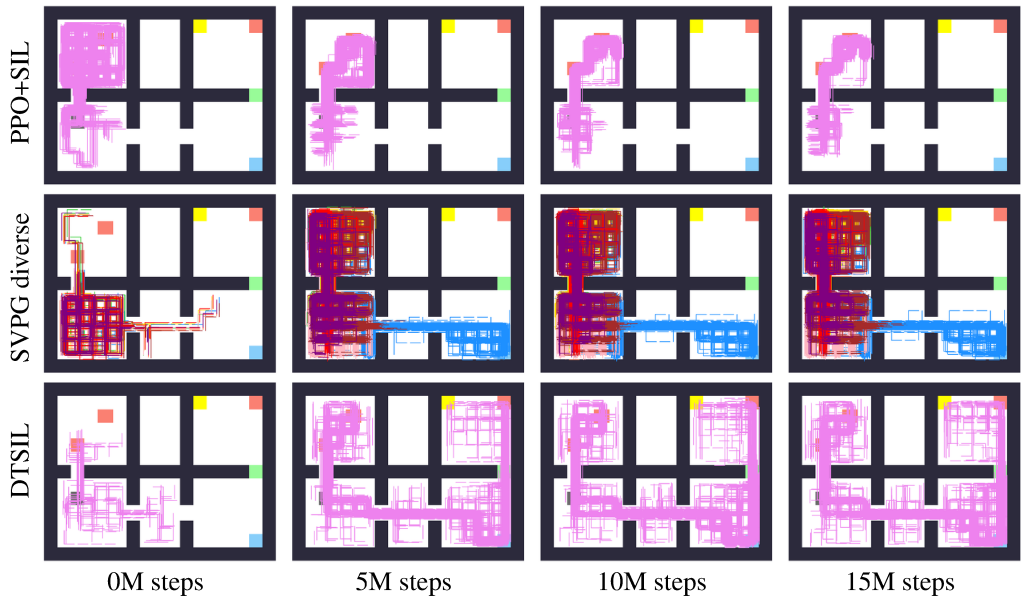

Figure 10: Visualization of the trajectories stored in the buffer for PPO+SIL, SVPG diverse [18] and our method as training continues. In the second row, we show the trajectories for a total of 8 policies learned simultaneously with the SVPG method proposed in [18], where each color corresponds to the trajectories collected by each policy.

G Map of Atari Montezuma's Revenge at the First Level

A map of Atari Montezuma's Revenge at the first level is shown in Figure 11. It is challenging to bring two keys to open the two doors in room 17 behind the treasure in room 15, where the agent can pass to the next level.

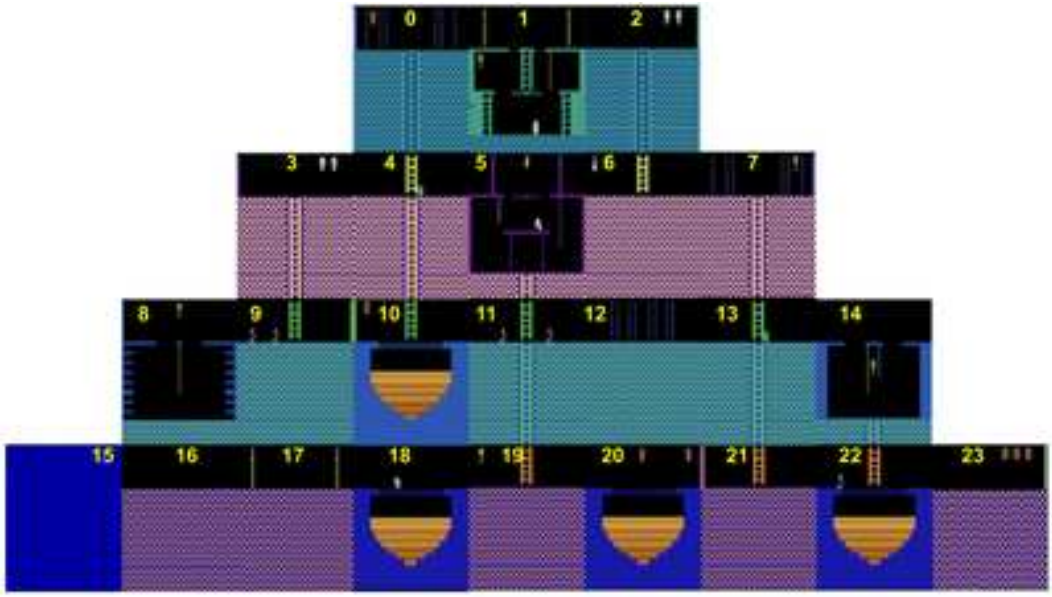

Figure 11: Map of Atari Montezuma's Revenge at the first level with 24 rooms.

H Comparison with Go-Explore

To evaluate the efficiency of exploration, we compare our method with the "exploration phase" in the Go-Explore algorithm [16]. The idea behind Go-Explore is to reset the agent to any interesting state sampled from the buffer of state embeddings, and then explore further using random actions. For fair comparison to our method, we modify the Go-Explore code such that we could not reset to any arbitrary states in the environment. We also use the same state representation $(\text{level}_t, \text{room}_t, x_t, y_t, k_t)$ and the same sampling weight $\frac{1}{\sqrt{n^{(i)}}}$ to sample goal states from the buffer.

In the Go-Explore method without using the direct 'reset' function and with a perfect goal-conditioned policy to visit any state sampled from the buffer, the agent could precisely advance to the goal state by following the stored trajectory. The total steps taken in the environment are counted by summing the number of steps taken to follow the stored trajectories and the number of steps taken to explore.

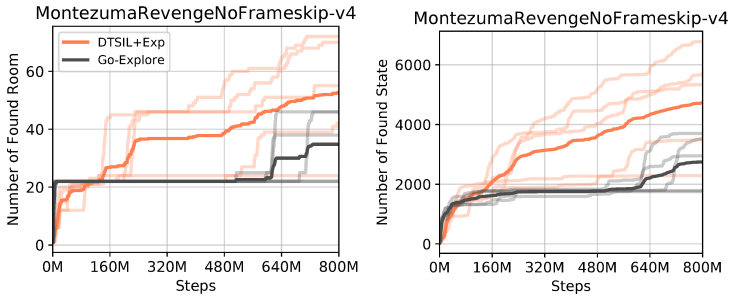

Figure 12: Learning curves of the number of rooms and the number of different state representations found on Atari Montezuma's Revenge, averaged over 5 runs. The curves in dark colors are the average of the 5 curves in light colors. During training, the state representation used is $(\text{level}_t, \text{room}_t, x_t, y_t, k_t)$.

In Figure 12, we show the average number of rooms found and the number of different state representations found during training. Even if we assume that there is a perfect goal-conditioned policy in Go-Explore to guide the agent to follow the stored trajectory exactly and visit the goal state, the learning curves demonstrate that our method is more efficient for exploring diverse state representations and consequently visits several rooms. This is because our method uses the count-based exploration bonus to encourages the exploration around and beyond the stored trajectories and the imitation reward allows the agent to follow the demonstrations in a soft-order.

I Automatic Switching between Exploration and Imitation

As stated in Section 3, when the agent has discovered trajectories with a satisfactory total reward the imitation is switched on. However, instead of setting the threshold of the satisfactory total reward in a handcrafted way we can automatically switch between exploration and imitation. In the early stage of training the agent prefers the exploration mode because there are few good trajectories to imitate. As the training progresses it is reasonable to train the policy to follow the best trajectories stored in the buffer. Assuming the agent has taken x_{current} steps in the environment and that the total number of steps is x_{total} (e.g., 800M steps for the experiments on Montezuma's Revenge) the probability of using the imitation mode is set to $p = \frac{x_{\text{current}}}{x_{\text{total}}}$. Therefore, the probability of performing imitation of good trajectories is increasing as training continues and the agent switches between exploration and imitation automatically without the need for setting the threshold of a satisfactory trajectory reward.

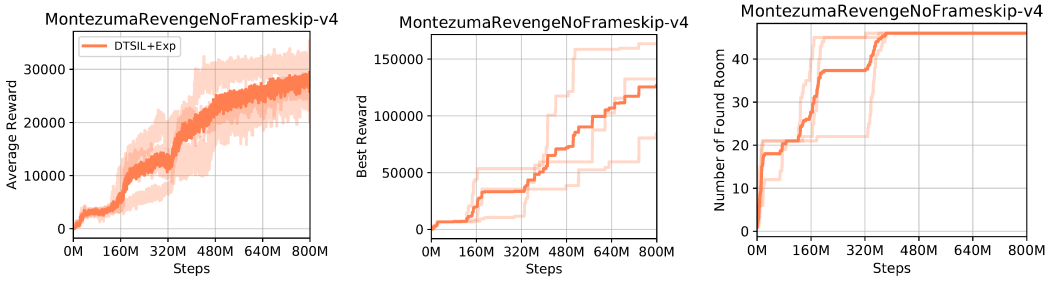

Figure 13: Learning curves of the average episode reward, the best episode reward, and the number of different rooms found on Atari Montezuma's Revenge, averaged over 3 runs. The probability of taking the imitation mode is automatically adjusted during training. 30,877 is the maximum mean score (averaged over 40 recent episodes) achieved over 800M environment timesteps (3,200M frames) during training, averaged over 3 runs.

As shown in Figure 13, by automatically switching between exploration and imitation our method can reach the average episode reward of around 30,000 over 3 runs within 800M environment steps.