MHER: Model-based Hindsight Experience Replay

Rui Yang[*1] **, Meng Fang**[2,3]**, Lei Han**[2]**, Yali Du**[4]**, Feng Luo**[1]**, Xiu Li**[1]
[1]Tsinghua University, [2]Tencent Robotics X
[3]Eindhoven University of Technology, [4]University College London
yangrui19@mails.tsinghua.edu.cn, m.fang@tue.nl, lxhan@tencent.com
yali.du86@gmail.com, {luof19@mails, li.xiu@sz}.tsinghua.edu.cn

Abstract

Solving multi-goal reinforcement learning (RL) problems with sparse rewards is generally challenging. Existing approaches have utilized goal relabeling on collected experiences to alleviate issues raised from sparse rewards. However, these methods are still limited in efficiency and cannot make full use of experiences. In this paper, we propose *Model-based Hindsight Experience Replay* (MHER), which exploits experiences more efficiently by leveraging environmental dynamics to generate virtual achieved goals. Replacing original goals with virtual goals generated from interaction with a trained dynamics model leads to a novel relabeling method, *model-based relabeling* (MBR). Based on MBR, MHER performs both reinforcement learning and supervised learning for efficient policy improvement. Theoretically, we also prove the supervised part in MHER, i.e., goal-conditioned supervised learning with MBR data, optimizes a lower bound on the multi-goal RL objective. Experimental results in several point-based tasks and simulated robotics environments show that MHER achieves significantly higher sample efficiency than previous model-free and model-based multi-goal methods.

1 Introduction

Although reinforcement learning (RL) has been shown to be effective in a range of reward-driven problems [20, 18, 13, 19], current RL algorithms require massive amounts of training data [33] and lack sample efficiency in sparse reward settings [1]. In multi-goal RL, the problem of efficiency becomes more prominent as agents are required to accomplish multiple goals simultaneously. One of the most essential factors affecting sample efficiency in multi-goal RL is the sparse reward, in which case informative learning signals are very limited. Previous works have proposed many solutions such as reward shaping [22], curriculum learning [3], exploration [23, 34], and hindsight relabeling [16, 1]. Among those solutions, learning from mistakes is a useful strategy to handle sparse rewards in multi-goal RL settings. Hindsight Experience Replay (HER) [1] remarkably improves sample efficiency through goal-relabeling that relabels failed experiences with actually achieved goals. Following HER, a few works are put forward to improve goal sampling methods [9, 24, 17, 8], or utilize hindsight knowledge for supervised policy learning [31, 12] and adversarial imitation learning [7].

Goal relabeling provides a practical paradigm for generating pseudo demonstrations to train control policies, and deep reinforcement learning algorithms further improve upon the efficiency of relabeling strategies [16, 1]. Most goal relabeling methods depend on trajectories and goals that an agent collects from environments. However, intelligent agents can achieve goals in complex environments even though they never encounter the exact same situation [26]. This ability requires building representations of the dynamics from past experience that enable generalization to novel situations [27, 30]. Modeling dynamics offers an explicit way to represent an agent's knowledge about the task.

[*]Work was done during the internship with Tencent Robotics X.

Deep RL Workshop at Neural Information Processing Systems, 2021.

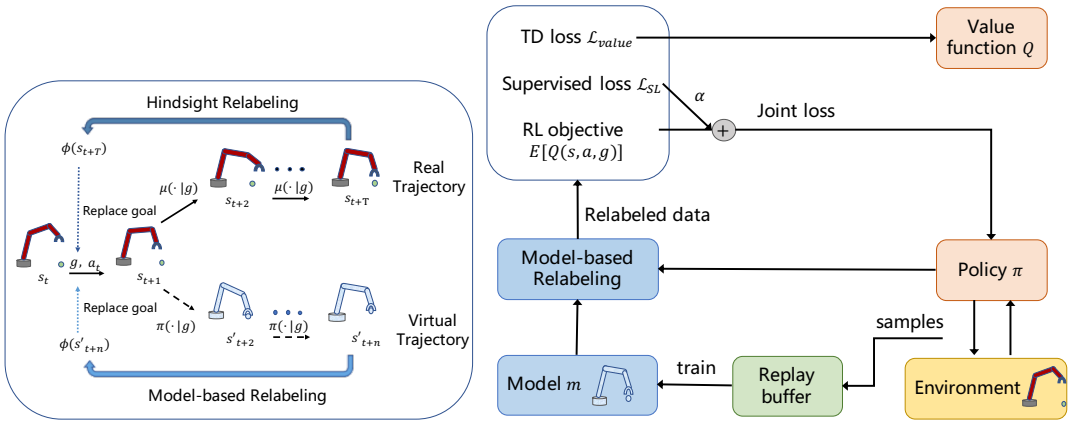

Figure 1: (a) Diagram of model-based relabeling. The real trajectory is collected by a past behavior policy μ and the virtual trajectory is generated from the interaction of the current policy π and a trained dynamics model. Hindsight relabeling such as HER uses real achieved goals (e.g., $\phi(s_{t+T})$, ϕ is a state-to-goal mapping) to relabel, while model-based relabeling utilizes virtual achieved goals (e.g., $\phi(s'_{t+n})$). (b) Overall framework of MHER. Detailed descriptions can be found in Section 4.4.

Existing methods [21, 15, 14, 28] have apply neural models to greatly facilitate predicting physical dynamics and the consequences of actions, and provide a strong inductive bias for generalization to novel environment situations. With the dynamics model and the current policy, we can predict future states along with achieved goals. Can we exploit goals generated from model-based interaction for sample efficient multi-goal reinforcement learning?

In this paper, we propose a novel framework, *Model-based Hindsight Experience Replay* (MHER), utilizing environmental dynamics to handle sparse rewards. In MHER, we introduce a new relabeling method, *model-based relabeling*, and then minimize a joint loss based on the model-based relabeled data for efficient policy learning. Unlike previous hindsight relabeling methods that relabel transitions with goals achieved at a later point during the same trajectory, model-based relabeling leverages dynamics models to generate pseudo goals for guiding the learning of the policy, as shown in Figure 1. The pseudo goals are collected in an efficient way without interaction with environments. Then goal-relabeling through imagined trajectories allows an agent to re-interpret its actions using a different goal from the perspective of the latest policy, leading to an implicit curriculum of goal relabeling guided by the policy. With the model-based relabeled data, we apply both supervised learning and reinforcement learning to update the policy. We theoretically prove that the policy can be improved by minimizing the goal-conditioned supervised learning loss [12] with the model-based relabeled data. To evaluate the performance of MHER, we conduct experiments on both point-based and Mujoco environments. Experimental results [2] show that MHER achieves significantly higher sample efficiency than previous works such as HER, Curriculum-guided HER [9], Maximum Entropy-based Prioritization [34], and Goal-Conditioned Supervised Learning [12].

The main contributions of this paper can be summarized as follows:

- We present a new goal relabeling method, model-based relabeling, leveraging dynamics models to handle sparse rewards in multi-goal RL;
- We apply supervised learning on the relabeled data and introduce a joint loss for RL training. We also prove that minimizing the supervised loss using the model-based relabeled data is equivalent to optimizing a lower bound on the original multi-goal RL objective;
- Empirical results on several benchmark environments demonstrate that the proposed method, MHER, exceeds previous multi-goal RL algorithms in sample efficiency.

[2] https://github.com/YangRui2015/Model-basedHER.

2 Related Work

Our work concentrates on sample efficiency in multi-goal reinforcement learning (RL) with sparse rewards. Hindsight Experience Replay (HER) [1] introduces hindsight relabeling for multi-goal RL, opening up a new way to learn from failed experiences with sparse and binary rewards. Based on HER, a few studies have been investigated to find more efficient goal sampling ways. Curriculum-guided HER (CHER) [9] selects goals in a heuristic way to balance the diversity of selected goals and the proximity to original goals. Focusing on the long horizon problem in multi-goal RL, Maximum Entropy Goal Achievement [24] samples from the frontier of achieved goals and gradually increases the entropy of achieved goals. Recent woks [17, 8] view hindsight relabeling as the inverse RL and sample relabeling goals according to their cumulative return or value function. Considering the visual tasks, [25] leverages GAN to generate virtual goals for relabeling, but it requires collecting a special dataset of near-goal states. In contrast to previous works, relabeling goals of MHER are generated from the interaction between current policy and a learned dynamics model.

In addition to improving the goal sampling methods, other works introduce prioritization replay and supervised learning to address sample inefficiency in multi-goal RL. Energy-Based Prioritization proposes to more frequently replay trajectories with higher energy. Maximum Entropy-based Prioritization (MEP) [34] prioritizes experiences based on trajectory entropy. By incorporating a few expert demonstrations, *goalGAIL* [7] significantly speeds up the convergence of policy via adversarial imitation learning. Self-supervised methods provides another simple but effective way for multi-goal RL. Policy Continuation with Hindsight Inverse Dynamics (PCHID) [31] extends hindsight inverse dynamics to the multi-step situation following dynamic programming, thus enabling learning in a self-imitated scheme. Goal-Conditioned Supervised Learning (GCSL) [12] provides theoretical guarantees that supervised learning from hindsight relabeled experiences optimizes a lower bound on the goal-oriented RL objective. Unlike PCHID and GCSL which both use real achieved goals for supervision, we leverage virtual goals for supervision and provide theoretical guarantees that supervised learning with virtual goals can lead to policy improvement.

Model-based RL algorithms have been studied for a long history and generally obtain higher sample efficiency over model-free algorithms [2]. Dyna [32] generates virtual samples with a trained dynamics model to augment training data. Learned dynamics models can also be incorporated into model-free algorithms to accelerate learning of policies and value functions [21]. Model-based value expansion (MVE) [10] improves model-free value estimation with predictive transitions. To tackle with the model bias in learned models, STEVE [5] uses an ensemble of models for a robust prediction. Model-Based Policy Optimization (MBPO) [15] proves a monotonic improvement with limited use of a predictive model. However, most of the model-based methods are designed for tasks with dense rewards. PlanGAN [6] is the first algorithm to use GANs and models for planning in sparse-reward multi-goal tasks. But PlanGAN is a planning method and requires a huge amount of computation to simulate trajectories for selecting a single action, and it may also suffer from the accumulated model-based errors in the long planning horizon. Different from these works, our work only utilizes the virtual achieved goals to avoid training with full virtual states and alleviate the impact of model bias.

3 Preliminaries

3.1 Multi-goal Reinforcement Learning

Multi-goal reinforcement learning (RL) can be characterized by the tuple $\langle S, A, G, r, p, \gamma \rangle$, where S, A, G, p, γ refer to state space, action space, goal space, transition function, and discount factor, respectively, and $r : S \times A \times G \to \mathbb{R}$ is the goal-conditioned reward function [26]. A commonly used sparse reward function in multi-goal setting is defined as:

$$r(s_t, a_t, g) = \begin{cases} 0, & ||\phi(s_t) - g||_2^2 < \text{threshold} \\ -1, & \text{otherwise} \end{cases}, \tag{1}$$

where $\phi : S \to G$ is a mapping function from states to achieved goals and is assumed to be available [1]. Agents are required to learn a policy $\pi : S \times G \to A$ to maximize returns of reaching goals from

the desired distribution $p(g)$:

$$J(\pi) = E_{g\sim p(g), a_t\sim\pi, s_{t+1}\sim p(\cdot|s_t,a_t)}\big[\sum_{t=0}^{\infty}\gamma^t r(s_t, a_t, g)\big] \tag{2}$$

3.2 Hindsight Experience Replay

Hindsight Experience Replay (HER) learns from failed experiences and tackles with sparse rewards in multi-goal RL. Given a trajectory $\tau = \{(s_t, a_t, g, r_t, s_{t+1})\}_{t=1}^{T}$ of length T, HER alternates g and r_t in the t-th transition $(s_t, a_t, g, r_t, s_{t+1})$ with a future achieved goal in the same trajectory $g' = \phi(s_{t+k}), 1 \leq k \leq T-t$ and $r'_t = r(s_t, a_t, g')$ computed by Eq. (1). After relabeling, transitions in failed trajectories can be assigned nonnegative rewards, therefore HER addresses the core issue of sparse rewards. HER can be combined with any off-policy algorithms such as DQN [20], DDPG [18], TD3 [11], and SAC [13].

In our paper, we adopt the DDPG+HER framework following [1, 9]. DDPG is an off-policy actor-critic algorithm consisting of a deterministic policy π and a value function $Q : S\times A\times G \rightarrow \mathbb{R}$. When collecting data with the policy π, Gaussian noise with zero mean and constant standard deviation is applied for exploration, as implemented in [11]. We denote the data distribution after hindsight relabeling as B_h. The value function Q is updated to minimize the TD error:

$$\mathcal{L}_{critic} = E_{(s_t,a_t,g',r'_t,s_{t+1})\sim B_h}\big[(y_t - Q(s_t, a_t, g'))^2\big], \tag{3}$$

where

$$y_t = r'_t + \gamma Q(s_{t+1}, \pi(s_{t+1}, g'), g').$$

The policy π is trained with policy gradient on the following loss:

$$\mathcal{L}_{actor} = -E_{(s_t,g')\sim B_h}\big[Q(s_t, \pi(s_t, g'), g')\big].$$

4 Methodology

In this section we will describe how our method, MHER, integrates environment dynamics into Hindsight Experience Replay [1] for training goal-conditioned policies. First we introduce the way we train the dynamics model that maps current state and action to the next state [32, 10]. Then we propose a novel model-based relabeling technique based on the learned model and introduce a joint loss combining policy gradient [29, 18] and goal-conditioned supervised learning [12] on the model-based relabeled data. Finally, we describe the overall framework of MHER.

4.1 Dynamics Models

When considering deterministic dynamics, the most direct way of training dynamics models is to learn to predict the next state given current state and action. Let $m(s_t, a_t)$ denote a learned dynamics function that takes the current state s_t and action a_t as input and outputs an estimation of the next state s_{t+1}. However, this model can be difficult to learn when the states s_t and s_{t+1} are very similar and the action has seemingly little effect on the changes in continuous control environments [21]. Therefore we follow [21] to learn a dynamics function that predicts the change between states by minimizing the following loss:

$$\mathcal{L}_{model} = E_{(s_t,a_t,s_{t+1})\sim B}\big[\|(s_{t+1} - s_t) - m(s_t, a_t)\|_2^2\big], \tag{4}$$

where the data (s_t, a_t, s_{t+1}) is sampled from the replay buffer B. Note that the dynamics model is independent of goals and rewards, therefore the data to train the model m can also be sampled from the relabeled data distribution such as B_h. With the trained dynamics model, current state s_t and action a_t, we can predict the next state as $s_{t+1} = s_t + m(s_t, a_t)$.

4.2 Model-based Relabeling

The insight of model-based relabeling (MBR) is that states in the virtual trajectory generated by the dynamics model can also be viewed as achieved goals for the starting state, so it can be used for relabeling the original transition. As shown in Figure 1, given a transition $(s_t, a_t, s_{t+1}, r_t, g)$ collected

Algorithm 1: MHER Framework

Given model-based interaction steps n, and α in the joint policy loss $\mathcal{L}_{joint}$;
Initialize policy π and value function Q;
Warmup the dynamics model m by minimizing $\mathcal{L}_{model}$ in Eq. (4) with random samples;
for *episode* = $1, 2, \ldots, M$ **do**
  Sample a desired goal g;
  Collect a trajectory with the policy π and save to the replay buffer B;
  Sample a minibatch b from the replay buffer : $\{(s_t, a_t, s_{t+1}, r_t, g)_i\}_{i=1}^{N} \sim B$;
  Update the dynamics model m with b;
  for $i = 1, 2, \ldots, N$ **do**
    $(s_t, a_t, s_{t+1}, r_t, g) \leftarrow b_i$;
    `// model-based relabeling`
    $s'_{t+1} = s_{t+1}$, $v = \{s'_{t+1}\}$;
    for $j = 1, 2, \ldots, n$ **do**
      $a'_{t+j} = \pi(s'_{t+j}, g)$;
      $s'_{t+j+1} = s'_{t+j} + m(s'_{t+j}, a'_{t+j})$;
      Append s'_{t+j+1} to v ;
    end for
    Sample random future state $s \sim v$;
    Get virtual achieved goal $g' = \phi(s)$;
    Recompute $r' = r(s_t, a_t, g')$ using reward function in Eq. (1);
    $b_i \leftarrow (s_t, a_t, r', s_{t+1}, g')$;
  end for
  Update value function Q with b to minimize $\mathcal{L}_{value}$ in Eq. (7);
  Update policy π with b to minimize $\mathcal{L}_{joint}$ in Eq. (6) ;
end for

by a behavior policy, MBR starts at s_{t+1} and interacts with the dynamics model m using current policy π for n steps. After interaction with the model, we have a virtual trajectory of $\{s'_{t+i}, a'_{t+i}, s'_{t+i+1}\}_{i=0}^{n}$, where $i = 0$ is the original transition and $a'_{t+i} = \pi(s'_{t+i}, g), s'_{t+i+1} = s'_{t+i} + m(s'_{t+i}, a'_{t+i}), 1 \leq i \leq n$. Then, any state in the virtual trajectory implies an achieved goal guided by current policy. MBR randomly samples from the virtual achieved goals $g' = \phi(s'_{t+j}), 1 \leq j \leq n$ to relabel the original transition: $(s_t, a_t, s_{t+1}, r'_t, g')$, where ϕ is the state-to-goal mapping and $r'_t = r(s_t, a_t, g')$ is computed according to Eq. (1). In our relabeled transitions, only the goal g' is virtually generated and other items are real experiences, avoiding the usage of full virtual states for training the value function. More detailed descriptions are provided in Appendix C. Note that the model-based interaction is driven by the current policy π under original goals g, thus MBR gradually pushes the relabeling goals towards the desired goal as the policy π improves. Theoretically, optimizing $J(\pi)$ is equivalent to minimizing an upper bound on the expected distance between virtual achieved goals and original desired goals. The proof can be found in Appendix A.3. We will also verify this property empirically in Section 5.4.

Two major advantages of model-based relabeling over previous relabeling strategies are as follows:

(1) MBR takes advantages of the current policy and environmental dynamics to generate more diverse goals for accelerating policy learning.

(2) As the policy improves, relabeling goals will gradually approach the assigned targets, therefore an implicit curriculum of automatic relabeling is introduced.

4.3 Learning Policy with Model-based Relabeled Data

After model-based relabeling (MBR), we further perform policy update based on the relabeled data. Similar to [12], the policy can be optimized through supervised learning on the model-based relabeled data. However, the difference is that we use virtually generated goals g' for supervision and minimize a mean square error rather than optimize the maximum likelihood, which are substantially the same under certain assumptions. We denote the model-based relabeled transition distribution as B_m, then

the supervised loss $\mathcal{L}_{SL}$ is defined as:

$$\mathcal{L}_{SL} = E_{(s_t,a_t,g')\sim B_m}\left[\|a_t - \pi(s_t, g')\|_2^2\right], \tag{5}$$

where g' is relabeled using MBR. Theoretically, we prove that minimizing $\mathcal{L}_{SL}$ in Eq. (5) is equivalent to maximizing a lower bound of $J(\pi)$, which is formally presented below.

Theorem 1. *Given a Diagonal Gaussian policy with a mean vector $\pi(s, g')$ and a non-zero positive constant variance σ^2, the discount factor γ, the real environmental dynamics $p(\cdot|s, a)$, a learned dynamics model $p_m(\cdot|s, a)$, the model-based relabeled data distribution B_m, and n-step model-based interactions, minimizing the supervised loss $\mathcal{L}_{SL}$ in Eq. (5) is equivalent to maximizing a lower bound on the multi-goal RL objective.*

$$J(\pi) \geq -\frac{n\gamma^n}{2\sigma^2}\mathcal{L}_{SL} + C_1\epsilon_m + C_2,$$

where C_1, C_2 are two constants independent of policy π. ϵ_m is the model error bounded at each timestep t: $\epsilon_m = max_t E_{s\sim\pi_{D,t}}[D_{TV}(p(s'|s, a)\|p_m(s'|s, a))]$, and π_D is the data collecting policy of B_m.

The detailed proof is provided in Appendix A.1.

Previous works update policy via either policy gradient [1, 9] or supervised learning [12] with the hindsight relabeled data introduced in Sec 3.2. In contrast to these works, we propose a joint loss combining policy gradient and the supervised loss $\mathcal{L}_{SL}$ for more efficient policy learning. Specifically, at each iteration during training, the policy is trained to minimize the following joint loss on the model-based relabeled data B_m:

$$\begin{aligned}\mathcal{L}_{joint} = &-\mathbb{E}_{(s_t,g')\sim B_m}\left[Q(s_t, \pi(s_t, g'), g')\right] \\ &+ \alpha\mathbb{E}_{(s_t,a_t,g')\sim B_m}\left[\|a_t - \pi(s_t, g')\|_2^2\right],\end{aligned} \tag{6}$$

where $\alpha > 0$ is the weight balancing the expected return and the supervised loss $\mathcal{L}_{SL}$. The Q-function is updated to minimize the TD error on the model-based relabeled data B_m:

$$\mathcal{L}_{value} = E_{(s_t,a_t,g',r'_t,s_{t+1})\sim B_m}\left[(y_t - Q(s_t, a_t, g'))^2\right], \tag{7}$$

where $y_t = r'_t + \gamma Q(s_{t+1}, \pi(s_{t+1}, g'), g')$.

4.4 Algorithm

The overall framework of model-based hindsight experience replay (MHER) is presented in Algorithm 1 and Figure 1 (b). First, the dynamics model m is trained to minimize the loss $\mathcal{L}_{model}$ using Eq. (4) in both warm-up and training periods. For every episode, we sample a goal g from the desired goal distribution and collect a trajectory to the replay buffer B using current policy. After collecting data, we sample a minibatch b from the replay buffer B. For each transition in the minibatch, we leverage the current policy π to rollout a n-step trajectory with the learned dynamics model m and perform model-based relabeling as explained in Section 4.2. After MBR, the minibatch b belongs to the model-based relabeled distribution B_m and is sent for training the Q network and the policy network. The Q network is updated according to Eq. (7) and the policy is trained to minimize the joint loss $\mathcal{L}_{joint}$ in Eq. (6).

5 Experiments

We conduct experiments on both continuous Point2D and Mujoco environments and compare the performance of MHER against a number of leading multi-goal RL algorithms for sparse reward environments. We also demonstrate the effectiveness of our goal-relabeling method by visualizing the distribution of relabeling goals.

5.1 Experimental Settings

Environments Our experimental environments consist of two Point2D environments, one Sawyer robots, and two Mujoco robots modified from OpenAI Gym [4]. All of the five environments' states, actions, and goals are continuous. In the first two point-based environments the blue point aims to reach the green circle. The other four environments (FetchReach-v1, SawyerReachXYZEnv-v1, Reacher-v2) control a robot to reach a target position in the goal space. More detailed task description can be found in Appendix F.

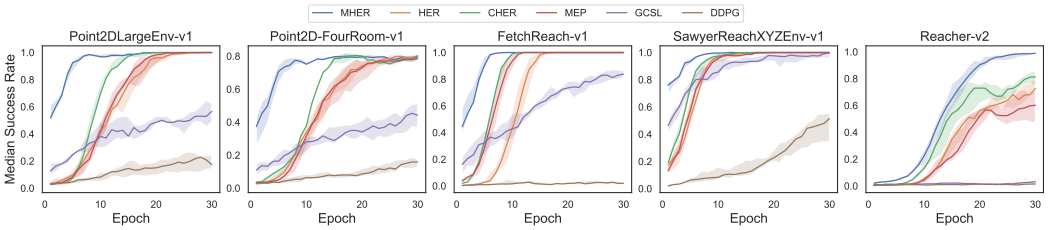

Figure 2: Median test success rate (line) with interquartile range (shaded area) on Point2D and Mujoco environments acorss 10 random seeds.

Baseline Implementation All the implementations of baseline algorithms are taken from their open-source code except GCSL, which we implement a deterministic version to fairly compare with other algorithms. For GCSL [12], we only maintain a policy network and train it to minimize the loss $\mathcal{L}_{GCSL} = E_{(s_t,a_t,g')\sim B_h}\left[\|a_t - \pi(s_t, g')\|_2^2\right]$, where g' is relabeled using future achieved goals similar to HER and B_h refers to the data distribution after hindsight relabeling.

Implementation of MHER As for the implementation of MHER, actor and critic networks are both 3-layer fully connected networks with 256 units each layer. All the networks are updated with Adam optimizer, learning rate 1×10^{-3}, and Polyak averaging coefficient of 0.9. To encourage exploration, the probability of random action is set as 0.3, and the scale of Gaussian noisy is 0.2. Following [1, 9], we perform model-based relabeling with a probability of 0.8. Note that the portion of data without MBR is not sent for goal-conditioned supervised learning, which can be implemented with a mask.

Training Settings We train all the algorithms for 30 epochs with one rollout worker. Each epoch contains 1 (Point2DLargeEnv-v1, Point2D-FourRoom-v1), 5 (FetchReach-v1, SawyerReachXYZEnv-v1), or 15 (Reacher-v2) episodes according to the difficulty of the environments. Every episode contains 100 interaction steps with the environment. At the end of each episode, we train all the algorithms for 5 batches with a batch size of 64.

Evaluation Settings After each epoch, we evaluate every algorithm for 100 episodes, and report the median test success rate and the interquartile range across 5 random seeds. To train the dynamics model m for MHER, we use a fully connected network with 4 hidden layers and 256 neurons each layer. In the warmup period, we train m for 100 updates with a batch size of 512 and a learning rate of 0.001. When training with MHER, we update m with the sampled minibatch for 2 times. More detailed hyperparameters are provided in Appendix D.

5.2 Benchmark Results

We compare MHER with baselines such as DDPG, HER [1], CHER [9], MEP [34], and GCSL [12] in five benchmark environments. For MHER, the model-based interaction step is set as 5 and the parameter α in the joint loss is set as 3. The median test success rates are reported in Figure 2. It is apparent that MHER achieves significantly higher performance than those baselines with a much faster learning speed. The results also show that DDPG learns slowly in all the environments, while HER, CHER, MEP are three effective baselines. Besides, GCSL's performance in different environments is not consistent, e.g., in Reacher-v2 it is very close to DDPG but in SawyerReachXYZEnv-v1 it outperforms DDPG by a large margin.

In Figure 3 (a) and 3 (b), we study how the parameters α and model-based interaction steps n impact MHER's performance. The parameter α impacts the weight of the supervised policy loss in $\mathcal{L}_{joint}$ when optimizing the policy, and $\alpha = 0$ indicates learning without the supervised loss. The results in Figure 3 (a) suggest that as α increases, performance increases when $\alpha \leq 3$ and decreases when $\alpha > 3$. As for the model-based interaction steps, more steps may contain more long-distance information but too many steps can lead to irrelevant goals. The results in Figure 3 (b) verify our analysis and the horizon of 5 achieves the best performance in the candidate set $\{0, 1, 3, 5, 7\}$. More experimental results are provided in Appendix B.

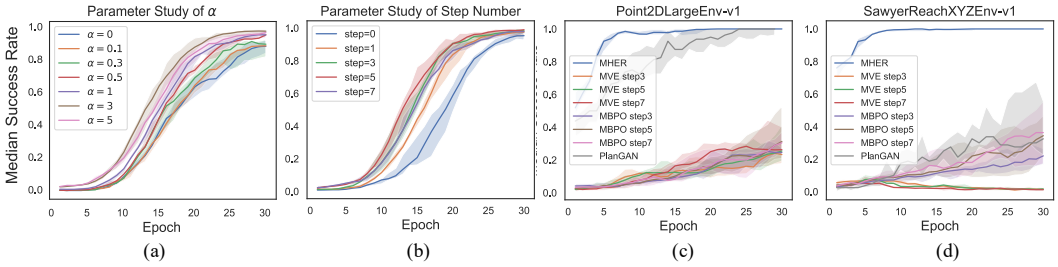

Figure 3: (a)(b) Parameter study of α and step number in Reacher-v2 environment. (c)(d) Comparison results with model-based based reinforcement learning baselines, MVE and MBPO, in Point2DLargeEnv-v1 and SawyerReachXYZEnv-v1 environments.

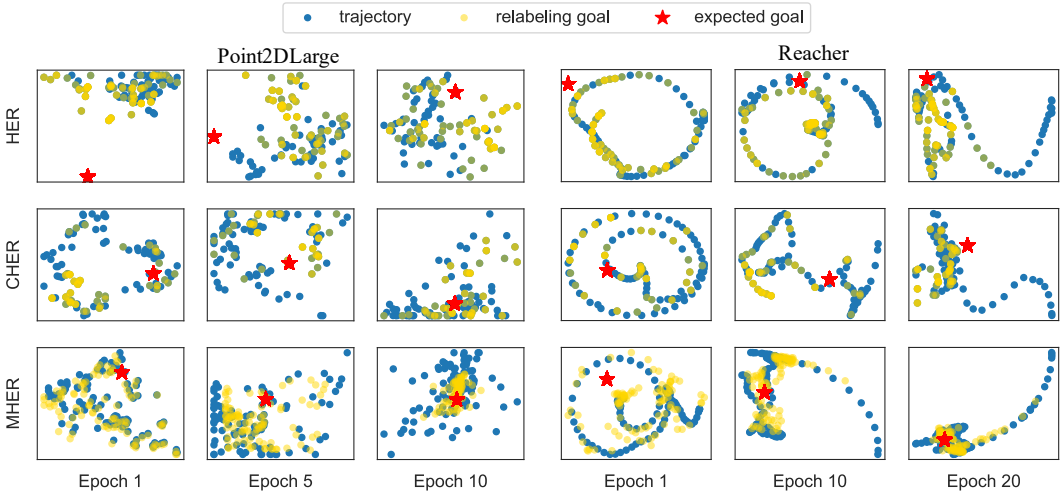

Figure 4: Relabeling goal distributions of HER, CHER, and MHER in Point2DLargeEnv-v1 and Reacher-v2. Blue points are real states in a trajectory, and red stars are their expected goals. Yellow points are relabeling goals for each transition in the trajectory.

5.3 Results of Model-based RL baselines

We also make fair comparison with model-based RL baselines in Figure 3 (c) and (d), including Model-based Value Expansion (MVE) [10] and Model-based Policy Optimization (MBPO) [15]. Implementation details of MVE and MBPO are provided in Appendix D. From the experimental results, we can conclude that previous model-based RL methods contribute little to multi-goal RL tasks with sparse rewards. In contrast to these algorithms, MHER can exploit environmental dynamics to learn policies more efficiently in the sparse reward setting. We also compare with PlanGAN [6], a planning method rather than a RL method, in Appendix B. The results show that MHER can achieve comparable or higher performance compared with PlanGAN, although PlanGAN requires extremely large amount of computations (detailed analysis is provided in Appendix B.1) to simulate trajectories for selecting actions.

5.4 Relationship with Curriculum Goal Relabeling

Prior work, Curriculum-guided HER (CHER) [9] introduces an explicit curriculum for automatically selecting relabeling goals. Nevertheless, the curriculum is hand-designed and needs to be adjusted through hyperparameters.

In Figure 4, we compare the relabeling goals of HER, CHER, and MHER in Point2DLargeEnv-v1 and Reacher-v2 environments. We can observe that HER selects goals randomly along the future achieved goals, leading to more goals in the end of trajectories. Besides, CHER selects goals to

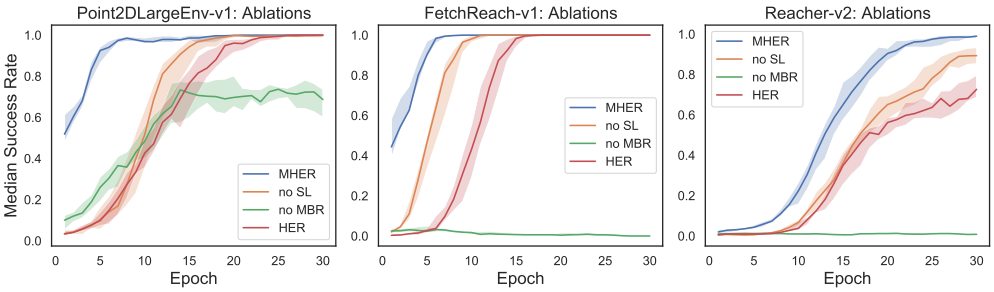

Figure 5: Ablation studies in Point2DLargeEnv-v1, FetchReach-v1, and Reacher-v2.

balance the diversity of selected goals and the proximity to the expected goal. Moreover, HER and CHER only select real achieved goals, thus the set of sampling goals is limited. On the contrary, MHER generates virtual relabeling goals according to the dynamics model and the current policy under the original desired goal. In the early stage, relabeling goals of MHER are near achieved goals. As the policy improves, they gradually approach their expected goals. The learning process of MHER is automatically adjusted by the policy, therefore MHER can substantially achieve more efficient curriculum learning.

5.5 Ablation Studies

To analyze the importance of *model-based relabeling* (MBR) and *supervised learning on the MBR data* (denoted as SL) in the MHER framework, we design ablation experiments to compare MHER variants with HER. For MBR and SL, the number of model-based interaction steps is 5 and $\alpha = 3$ by default. We experiment with the following settings:

- *MHER*: DDPG+MBR+SL;
- *no SL*: DDPG+MBR, which is equivalent to $\alpha = 0$;
- *no MBR*: DDPG+SL, adding an auxiliary task to DDPG that only uses model-based relabeled transitions to minimize the supervised policy loss.

Empirical results in Figure 5 demonstrate that MBR is more important than SL in the MHER framework. DDPG+MBR learns faster than HER, but DDPG+SL learns very slowly in FetchReach-v1 and Reacher-v2, and cannot converge to 100% success rate in Point2DLargeEnv-v1, which indicates that SL is less effective compared to MBR. By combining MBR and SL, the performance is significantly improved. The intuition behind the results is that MBR and SL achieve mutual improvement under the MHER framework, because MBR provides goals following an efficient curriculum to train the policy, and the supervised policy loss further improves the policy that guides the curriculum of MBR.

6 Conclusion

In this paper, we propose the framework of *Model-based Hindsight Experience Replay* (MHER) to improve sample efficiency in multi-goal RL with sparse rewards. In MHER, we introduce virtual goals for goal relabeling and policy improvement, corresponding to model-based relabeling (MBR) and a joint loss combining policy gradient and supervised learning with MBR data. Experiments in a range of continuous multi-goal tasks demonstrate that MHER achieves significantly higher sample efficiency than previous goal-conditioned works such as HER, CHER, MEP and GCSL. Moreover, we show that MHER is efficient in the following aspects: (1) virtual goals generated from interaction with the trained model are not limited to real experiences; (2) generated goals follow an efficient curriculum guided by the policy; (3) policy learning takes advantage of both reinforcement learning and supervised improvement. In the future, we would like to work on more efficient works inspired by MHER.

References

[1] Marcin Andrychowicz, Filip Wolski, Alex Ray, Jonas Schneider, Rachel Fong, Peter Welinder, Bob McGrew, Josh Tobin, OpenAI Pieter Abbeel, and Wojciech Zaremba. Hindsight experience replay. In *Advances in Neural Information Processing Systems*, pages 5048–5058, 2017.

[2] Christopher G Atkeson and Juan Carlos Santamaria. A comparison of direct and model-based reinforcement learning. In *IEEE International Conference on Robotics and Automation*, volume 4, pages 3557–3564, 1997.

[3] Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. Curriculum learning. In *International Conference on Machine Learning*, pages 41–48, 2009.

[4] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym. *arXiv preprint arXiv:1606.01540*, 2016.

[5] Jacob Buckman, Danijar Hafner, George Tucker, Eugene Brevdo, and Honglak Lee. Sample-efficient reinforcement learning with stochastic ensemble value expansion. In *NeurIPS*, 2018.

[6] Henry Charlesworth and Giovanni Montana. Plangan: Model-based planning with sparse rewards and multiple goals. In *Advances in Neural Information Processing Systems*, 2020.

[7] Yiming Ding, Carlos Florensa, Pieter Abbeel, and Mariano Phielipp. Goal-conditioned imitation learning. *Advances in Neural Information Processing Systems*, 32:15324–15335, 2019.

[8] Ben Eysenbach, Xinyang Geng, Sergey Levine, and Russ R. Salakhutdinov. Rewriting history with inverse RL: hindsight inference for policy improvement. In *Advances in Neural Information Processing Systems*, 2020.

[9] Meng Fang, Tianyi Zhou, Yali Du, Lei Han, and Zhengyou Zhang. Curriculum-guided hindsight experience replay. In *Advances in Neural Information Processing Systems*, pages 12623–12634, 2019.

[10] Vladimir Feinberg, Alvin Wan, Ion Stoica, Michael I Jordan, Joseph E Gonzalez, and Sergey Levine. Model-based value expansion for efficient model-free reinforcement learning. In *Proceedings of the 35th International Conference on Machine Learning (ICML 2018)*, 2018.

[11] Scott Fujimoto, Herke Hoof, and David Meger. Addressing function approximation error in actor-critic methods. In *International Conference on Machine Learning*, pages 1587–1596, 2018.

[12] Dibya Ghosh, Abhishek Gupta, Ashwin Reddy, Justin Fu, Coline Manon Devin, Benjamin Eysenbach, and Sergey Levine. Learning to reach goals via iterated supervised learning. In *International Conference on Learning Representations*, 2021.

[13] Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International Conference on Machine Learning*, pages 1861–1870, 2018.

[14] Danijar Hafner, Timothy Lillicrap, Jimmy Ba, and Mohammad Norouzi. Dream to control: Learning behaviors by latent imagination. In *International Conference on Learning Representations*, 2019.

[15] Michael Janner, Justin Fu, Marvin Zhang, and Sergey Levine. When to trust your model: Model-based policy optimization. In *Advances in Neural Information Processing Systems*, pages 12498–12509, 2019.

[16] Leslie Pack Kaelbling. Learning to achieve goals. In *International Joint Conference on Artificial Intelligence*, pages 1094–1099, 1993.

[17] Alexander C. Li, Lerrel Pinto, and Pieter Abbeel. Generalized hindsight for reinforcement learning. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.

[18] Timothy P Lillicrap, Jonathan J Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*, 2015.

[19] Jiafei Lyu, Xiaoteng Ma, Jiangpeng Yan, and Xiu Li. Efficient continuous control with double actors and regularized critics. *CoRR*, abs/2106.03050, 2021.

[20] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.

[21] Anusha Nagabandi, Gregory Kahn, Ronald S Fearing, and Sergey Levine. Neural network dynamics for model-based deep reinforcement learning with model-free fine-tuning. In *IEEE International Conference on Robotics and Automation*, pages 7559–7566, 2018.

[22] Andrew Y Ng, Daishi Harada, and Stuart Russell. Policy invariance under reward transformations: Theory and application to reward shaping. In *International Conference on Machine Learning*, volume 99, pages 278–287, 1999.

[23] Deepak Pathak, Pulkit Agrawal, Alexei A Efros, and Trevor Darrell. Curiosity-driven exploration by self-supervised prediction. In *International Conference on Machine Learning*, pages 2778–2787. PMLR, 2017.

[24] Silviu Pitis, Harris Chan, Stephen Zhao, Bradly Stadie, and Jimmy Ba. Maximum entropy gain exploration for long horizon multi-goal reinforcement learning. In *International Conference on Machine Learning*, pages 7750–7761, 2020.

[25] Himanshu Sahni, Toby Buckley, Pieter Abbeel, and Ilya Kuzovkin. Addressing sample complexity in visual tasks using HER and hallucinatory gans. In *Advances in Neural Information Processing Systems*, pages 5823–5833, 2019.

[26] Tom Schaul, Daniel Horgan, Karol Gregor, and David Silver. Universal value function approximators. In *International Conference on Machine Learning*, pages 1312–1320, 2015.

[27] Tom Schaul and Mark B Ring. Better generalization with forecasts. In *IJCAI*, pages 1656–1662, 2013.

[28] Julian Schrittwieser, Ioannis Antonoglou, Thomas Hubert, Karen Simonyan, Laurent Sifre, Simon Schmitt, Arthur Guez, Edward Lockhart, Demis Hassabis, Thore Graepel, et al. Mastering atari, go, chess and shogi by planning with a learned model. *Nature*, 588(7839):604–609, 2020.

[29] David Silver, Guy Lever, Nicolas Heess, Thomas Degris, Daan Wierstra, and Martin Riedmiller. Deterministic policy gradient algorithms. In *International conference on machine learning*, pages 387–395. PMLR, 2014.

[30] Satinder P Singh, Michael L Littman, Nicholas K Jong, David Pardoe, and Peter Stone. Learning predictive state representations. In *Proceedings of the 20th International Conference on Machine Learning (ICML-03)*, pages 712–719, 2003.

[31] Hao Sun, Zhizhong Li, Xiaotong Liu, Bolei Zhou, and Dahua Lin. Policy continuation with hindsight inverse dynamics. In *Advances in Neural Information Processing Systems*, pages 10265–10275, 2019.

[32] Richard S Sutton. Dyna, an integrated architecture for learning, planning, and reacting. *ACM Sigart Bulletin*, 2(4):160–163, 1991.

[33] Jane X Wang, Zeb Kurth-Nelson, Dhruva Tirumala, Hubert Soyer, Joel Z Leibo, Remi Munos, Charles Blundell, Dharshan Kumaran, and Matt Botvinick. Learning to reinforcement learn. *arXiv preprint arXiv:1611.05763*, 2016.

[34] Rui Zhao, Xudong Sun, and Volker Tresp. Maximum entropy-regularized multi-goal reinforcement learning. In *International Conference on Machine Learning*, pages 7553–7562, 2019.

[35] Menghui Zhu, Minghuan Liu, Jian Shen, Zhicheng Zhang, Sheng Chen, Weinan Zhang, Deheng Ye, Yong Yu, Qiang Fu, and Wei Yang. Mapgo: Model-assisted policy optimization for goal-oriented tasks. *arXiv preprint arXiv:2105.06350*, 2021.

A Theoretical Analysis

A.1 Proof of Theorem 1

Proof. Following GCSL [12], we denote a real trajectory as $\tau = \{s_0, a_0, \ldots, s_T, a_T\}$. The expected goal is g, and the real dynamics is $p(s_{t+1}|s_t, a_t)$. We assume the initial state s_0 is uniformly distributed in the state space. To simplify the analysis, the relabeling goal is defined as the achieved goal of the last state $\mathcal{G}(\tau) \triangleq \phi(s_T)$, where ϕ is the state-to-goal mapping and $\mathcal{G}$ is the mapping from trajectories to relabeling goals. This simplification is also adopted in [12]. Then we denote a virtual trajectory generated by the trained dynamics model $p_m(s'_{t+1}|s'_t, a'_t)$ as $\tau_m = \{s'_0, a'_0, \ldots, s'_n, a'_n\}$, $s'_0 = s_0$, and suppose $T \geq n$. The model-based relabeling goal is defined as $\mathcal{G}(\tau_m) \triangleq \phi(s'_n)$. We utilize a sparse reward function: $r(s_t, a_t, g) = 1[\phi(s_t) = g]$, therefore the reward at each timestep is bounded $r_t \in [0, 1]$. Other variables are defined in Section 3.

In our setting, the original multi-goal RL objective is:

$$J(\pi) = E_{g\sim p(g), a_t\sim\pi, s_{t+1}\sim p(\cdot|s_t,a_t)}\big[\sum_{t=0}^{T} \gamma^t r(s_t, a_t, g)\big].$$

Similarly, the expected return of the virtual trajectory τ_m is:

$$J_m(\pi) = E_{g\sim p(g), a_t\sim\pi, s_{t+1}\sim p_m(\cdot|s_t,a_t)}[\sum_{t=0}^{n} \gamma^t r(s_t, a_t, g)].$$

Following MBPO [15], we first bound model error using ϵ_m. As our model-based approach collects data with current policy π, the policy shift $\epsilon_\pi = 0$. Next, we apply Lemma 2 to bound $J(\pi)$ with $J_m(\pi)$ and absorb constants into C_1 as:

$$J(\pi) \geq J_m(\pi) + C_1\epsilon_m,$$

where $C_1 = -\frac{2\gamma r_{max}}{(1-\gamma)^2} = -\frac{2\gamma}{(1-\gamma)^2}$ is derived from Lemma 2.

In GCSL, agents optimize the goal-reaching objective, i.e., maximizing the probability of reaching g at the last step:

$$\hat{J}(\pi) = E_{g\sim p(g), a_t\sim\pi, s_{t+1}\sim p(\cdot|s_t,a_t)}[1|\mathcal{G}(\tau) = g].$$

With Lemma 3, we can also obtain the goal-reaching return bound for model-based return $J_m(\pi)$:

$$J_m(\pi) \geq \gamma^n E_{g\sim p(g), a_t\sim\pi, s_{t+1}\sim p_m(\cdot|s_t,a_t)}\big[1[\mathcal{G}(\tau_m) = g]\big],$$

and denote the inequality as $J_m(\pi) \geq \gamma^n \hat{J}_m(\pi)$. For $\hat{J}_m(\pi)$, we follow Theorem 4.1 in [12] and have:

$$\hat{J}_m(\pi) \geq E_{a_t\sim\pi, s_{t+1}\sim p_m(\cdot|s_t,a_t)}\big[\sum_{t=0}^{n} log\pi(a_t|s_t, \mathcal{G}(\tau_m))\big] + C_2$$

As we collect data using the current policy π, we don't have the policy variation term. The right side of this inequality means using relabeled data for maximum likelihood estimation, therefore the policy π is required to be stochastic and select actions with non-zero probability.

Combining the intermediate results, we can conclude that:

$$\begin{aligned} J(\pi) &\geq J_m(\pi) + C_1\epsilon_m \geq \gamma^n \hat{J}_m(\pi) + C_1\epsilon_m \\ &\geq \gamma^n E_{a_t\sim\pi, s_{t+1}\sim p_m(\cdot|s_t,a_t)}\big[\sum_{t=0}^{n} log\pi(a_t|s_t, \mathcal{G}(\tau_m))\big] + C_1\epsilon_m + C_2 \end{aligned}$$

Then, we take the Diagonal Gaussian policy with mean vector $\pi(s_t, g)$ of dimension $|a|$ and non-zero positive constant variance σ^2:

$$P(a_t|s_t, g) = \frac{1}{(\sigma\sqrt{2Pi})^{|a|}} e^{\frac{-\|a_t - \pi(s_t,g)\|_2^2}{2\sigma^2}}.$$

Note that we use Pi to represent 'pi' in the Gaussian distribution and to distinguish from the policy π. Pi is a constant and we will absorb it in C_2 in the following derivation. Denoting the model-based relabeled data distribution as B_m and relabeled goal as g', we have:

$$\begin{aligned} J(\pi) &\geq -\frac{\gamma^n}{2\sigma^2} E_{(a_t,s_t,g')_{t=0}^n \sim B_m} \big[\sum_{t=0}^{n} \|a_t - \pi(s_t, g')\|_2^2 \big] + C_1\epsilon_m + C_2 \\ &= -\frac{n\gamma^n}{2\sigma^2} E_{(a_t,s_t,g') \sim B_m} \|a_t - \pi(s_t, g')\|_2^2 + C_1\epsilon_m + C_2 \end{aligned}$$

This completes the proof. □

A.2 Useful Lemmas

Lemma 1 (Finite Horizon Return Bound)**.** *Suppose the model error is bounded as* $max_t E_{s\sim p_1^t(s)} D_{TV}(p_1(s'|s,a)\|p_2(s'|s,a)) \leq \epsilon_m$*, and* $max_s D_{TV}(\pi_1(a|s)\|\pi_2(a|s)) \leq \epsilon_\pi$*. And We denote* $J^{(T)}(\pi)$ *as expected return of finite horizon* T *under policy* π*, and* $r_{max} > 0$ *is the maximum reward. Then the difference of finite horizon returns are bounded as:*

$$|J^{(T)}(\pi_1) - J^{(T)}(\pi_2)| \leq 2r_{max} \sum_{t=0}^{T} \gamma^t [t(\epsilon_m + \epsilon_\pi) + \epsilon_\pi]$$

Proof. This lemma can be derived by Lemma B.3 in [15], we also provide a sketch proof.

$$\begin{aligned} |J^{(T)}(\pi_1) - J^{(T)}(\pi_2)| &= |\sum_{s,a}\sum_{t} \gamma^t (p_1^t(s,a) - p_2^t(s,a)) r(s,a)| \\ &\leq r_{max} \sum_{t}\sum_{s,a} \gamma^t |p_1^t(s,a) - p_2^t(s,a)| \\ &= 2r_{max} \sum_{t} \gamma^t D_{TV}(p_1^t(s,a)\|p_2^t(s,a)) \\ &\leq 2r_{max} \sum_{t} \gamma^t [D_{TV}(p_1^t(s)\|p_2^t(s)) + max_s D_{TV}(\pi_1(a|s)\|\pi_2(a|s))] \\ &\leq 2r_{max} \sum_{t=0}^{T} \gamma^t [t(\epsilon_m + \epsilon_\pi) + \epsilon_\pi] \end{aligned}$$

where $D_{TV}(p_1^t(s)\|p_2^t(s)) \leq t(\epsilon_m + \epsilon_\pi)$ is also prove in Lemma B.3 in [15]. The branched returns bound (Lemma B.3 in [15]) is a special case of our Lemma 1 when horizon $T \to \infty$. □

Lemma 2. *Let* $J^{(\infty)}(\pi)$ *denote* $J(\pi)$ *with infinite horizon, where the superscript refers to the horizon.* $J(\pi) = J^{(T)}(\pi), J_m(\pi) = J_m^{(n)}(\pi), T \geq n$*. Given model error* ϵ_m *and policy shift* ϵ_π *defined in [15], non-negative reward function* $r(s_t, a_t) \geq 0$*, and the maximum reward* r_{max}*,* $J(\pi)$ *can be bounded as:*

$$J(\pi) \geq J_m(\pi) - \frac{2r_{max}}{(1-\gamma)^2}(\gamma\epsilon_m + 2\epsilon_\pi),$$

Proof. Let π_D denote the data collecting policy, we follow the proof of Theorem A.1 in [15] and have:

$$J^{(T)}(\pi) - J_m^{(T)}(\pi) = \underbrace{J^{(T)}(\pi) - J^{(T)}(\pi_D)}_{L1} + \underbrace{J^{(T)}(\pi_D) - J_m^{(T)}(\pi)}_{L2}$$

For L_1 and L_2, we apply Lemma 1 to derive the following bound (note that there should be no model error in L_1 term):

$$\begin{aligned} J^{(T)}(\pi) - J_m^{(T)}(\pi) &\geq \underbrace{-2r_{max} \sum_{t=0}^{T} \gamma^t \big[t\epsilon_\pi + \epsilon_\pi\big]}_{L_1} \underbrace{-2r_{max} \sum_{t=0}^{T} \gamma^t \big[t(\epsilon_m + \epsilon_\pi) + \epsilon_\pi\big]}_{L_2} \\ &= -2r_{max} \sum_{t=0}^{T} \gamma^t \big[t(\epsilon_m + 2\epsilon_\pi) + 2\epsilon_\pi\big], \end{aligned}$$

To verify the above inequality, let $T \to \infty$ and leverage the properties $\sum_{t=0}^{\infty} \gamma^t t \leq \frac{\gamma}{(1-\gamma)^2}$ and $\sum_{t=0}^{\infty} \gamma^t \leq \frac{1}{1-\gamma}$, then we can get MBPO performance bound (Theorem A.1 in [15]). Utilizing the assumptions $r_t \in [0,1], T \geq n$ and properties $\epsilon_m \geq 0, \epsilon_\pi \geq 0, J_m^{(T)} \geq J_m^{(n)}$, we can then prove the lower bound

$$\begin{aligned}
J^{(T)}(\pi) &\geq J_m^{(T)}(\pi) - 2r_{max} \sum_{t=0}^{T} \gamma^t \big[t(\epsilon_m + 2\epsilon_\pi) + 2\epsilon_\pi\big] \\
&\geq J_m^{(T)}(\pi) - 2r_{max} \sum_{t=0}^{\infty} \gamma^t \big[t(\epsilon_m + 2\epsilon_\pi) + 2\epsilon_\pi\big] \\
&\geq J_m^{(n)}(\pi) - \frac{2r_{max}[\gamma\epsilon_m + 2\gamma\epsilon_\pi + 2\epsilon_\pi(1-\gamma)]}{(1-\gamma)^2} \\
&\geq J_m^{(n)}(\pi) - \frac{2r_{max}}{(1-\gamma)^2}(\gamma\epsilon_m + 2\epsilon_\pi).
\end{aligned}$$

Alternating $J_m^{(n)}(\pi), J^{(T)}(\pi)$ with $J_m(\pi), J(\pi)$ completes the proof. □

Lemma 3 (Goal-reaching Return Bound)**.** *Given trajectories $\tau = \{s_0, a_0, \ldots, s_T, a_T\}$ and reward function $r(s_t, a_t, g) = 1[\phi(s_t) = g]$, the goal-reaching objective $\gamma^T \hat{J}(\pi)$ is a lower bound of the original multi-goal objective $J(\pi)$.*

Proof. Leveraging the definition of $J(\pi)$, $\hat{J}(\pi)$, and the reward function $r(s_t, a_t, g) = 1[\phi(s_t) = g]$, we have:

$$\begin{aligned}
J(\pi) &= E_{g\sim p(g), a_t\sim\pi, s_{t+1}\sim p(\cdot|s_t,a_t)}\big[\sum_{t=0}^{T} \gamma^t r(s_t, a_t, g)\big] \\
&\geq E_{g\sim p(g), a_t\sim\pi, s_{t+1}\sim p(\cdot|s_t,a_t)}\big[\gamma^T r(s_T, a_T, g)\big] \\
&= \gamma^T E_{g\sim p(g), a_t\sim\pi, s_{t+1}\sim p(\cdot|s_t,a_t)}\big[1[\mathcal{G}(\tau) = g]\big] \\
&= \gamma^T \hat{J}(\pi).
\end{aligned}$$

□

A.3 Theoretical Analysis of Virtual Achieved Goals

We provide formal analysis on the expected distance between virtual achieved goals and original desired goals. First, we define the *expected distance* ($ED_\pi(g)$) as the sum of the product of distance and the discounted visitation frequencies:

$$ED_\pi(g) = \sum_s \rho_\pi(s) \|\phi(s) - g\|_2^2$$

where g refer to original desired goals, ϕ is a state-to-goal mapping, and $\rho_\pi(s) = P(s_0 = s) + \gamma P(s_1 = s|\pi) + \gamma^2 P(s_2 = s|\pi) + ..., s_0 \sim p(s_0)$ and actions are sampled according to π. The expected distance $ED_\pi(g)$ measure the proximity of trajectories generated by the policy π and the desired goal g. The following theory demonstrates the expected distance in the model-based relabeling can be minimized by optimizing the policy.

Theorem 2 (Expected Distance Bound, Infinite Horizon)**.** *Suppose the state space is finite, the policy π is optimized to maximize $J(\pi)$, the reward function is sparse $r(s, a, g) = 1[\|\phi(s) - g\| \leq \epsilon]$, ϵ is a small positive threshold, the goal space is bounded, i.e., $\|g - g'\|_2^2 \leq C_g, \forall g, g' \in \mathcal{G}, C_g \geq 0$. Denote the environmental dynamics as $p(s'|s, a)$, the model error ϵ_m, the learned dynamics model as $p_m(s'|s, a)$, and the discounted visitation frequencies with the dynamics model as $\rho_{\pi,m}(s)$. Then,* ***optimizing*** *$J(\pi)$* ***is equivalent to minimizing the expected distance between virtual achieved goals and original desired goals****, i.e., $ED_{\pi,m}(g) = \sum_s \rho_{\pi,m}(s)\|\phi(s) - g\|_2^2$. Specifically, the following inequality holds:*

$$ED_{\pi,m}(g) \leq -C_g \cdot J(\pi) + \frac{C_g + \epsilon}{1-\gamma} + \frac{2\gamma\epsilon_m C_g}{(1-\gamma)^2}$$

Proof. We first bound the expected distance as:

$$
\begin{aligned}
ED_\pi(g) &= \sum_s \rho_\pi(s)\|\phi(s)-g\|_2^2 \cdot 1[\|\phi(s)-g\|_2^2 \le \epsilon] + \sum_s \rho_\pi(s)\|\phi(s)-g\|_2^2 \cdot 1[\|\phi(s)-g\|_2^2 > \epsilon] \\
&\le \frac{\epsilon}{1-\gamma} + \sum_s \rho_\pi(s)\|\phi(s)-g\|_2^2 \cdot 1[\|\phi(s)-g\|_2^2 > \epsilon] \\
&\le \frac{\epsilon}{1-\gamma} + C_g \sum_s \rho_\pi(s) \cdot 1[\|\phi(s)-g\|_2^2 > \epsilon]
\end{aligned} \tag{8}
$$

The above derivation leverages the property that:

$$
\sum_s \rho_\pi(s) = \sum_s P(s_0 = s) + \gamma \sum_s P(s_1 = s) + \gamma^2 \sum_s P(s_2 = s) + \cdots = \frac{1}{1-\gamma}
$$

As in our setting, rewards only depend on states, thus the RL objective can be written as:

$$
\begin{aligned}
J(\pi) &= E_{s_0 \sim p(s_0), a_t \sim \pi(a_t|s_t), s_{t+1} \sim p(s_{t+1}|s_t, a_t)} \big[\sum_{t=0}^{\infty} \gamma^t r(s_t) \big] \\
&= \sum_{t=0}^{\infty} \sum_s P(s_t = s|\pi) \cdot \gamma^t r(s) = \sum_s \sum_{t=0}^{\infty} \gamma^t P(s_t = s|\pi) \cdot r(s) \\
&= \sum_s \rho_\pi(s) \cdot 1[\|\phi(s)-g\|_2^2 \le \epsilon] \\
&= \sum_s \rho_\pi(s) \cdot \big[1 - 1[\|\phi(s)-g\|_2^2 > \epsilon]\big] \\
&= \frac{1}{1-\gamma} - \sum_s \rho_\pi(s) \cdot 1[\|\phi(s)-g\|_2^2 > \epsilon]
\end{aligned}
$$

Taking the upper bound of $ED_\pi(g)$ in Eq 8 into $J(\pi)$, we have:

$$
J(\pi) \le \frac{1}{1-\gamma} - \frac{1}{C_g}[ED_\pi(g) - \frac{\epsilon}{1-\gamma}]
$$

and

$$
ED_\pi(g) \le -C_g \cdot J(\pi) + \frac{C_g + \epsilon}{1-\gamma}.
$$

The virtual achieved goals are generated with the learned dynamics model, therefore the following inequality also holds for the learned dynamics model:

$$
ED_{\pi,m}(g) \le -C_g \cdot J_m(\pi) + \frac{C_g + \epsilon}{1-\gamma}.
$$

From the proof of Lemma 2, $J_m^{(\infty)}(\pi)$ can be bounded as: $J_m(\pi) \ge J(\pi) - \frac{2\gamma\epsilon_m}{(1-\gamma)^2}$, where ϵ_m is the model error defined before. Finally, we can conclude that:

$$
\begin{aligned}
ED_{\pi,m}(g) &\le -C_g \cdot J_m(\pi) + \frac{C_g + \epsilon}{1-\gamma} \\
&\le -C_g \cdot J(\pi) + \frac{C_g + \epsilon}{1-\gamma} + \frac{2\gamma\epsilon_m C_g}{(1-\gamma)^2}.
\end{aligned}
$$

This completes the proof. □

Theorem 3 (Expected Distance Bound, Finite Horizon)**.** *Assume the finite horizon for model-based rollout equals the horizon in real environment, i.e., $n = T$. Following Theorem 2 and [Ross et al. 2011], the expected distance bound for finite horizon is:*

$$
ED_{\pi,m}(g) \le -C_g \cdot J(\pi) + (C_g + \epsilon)T + 2\gamma\epsilon_m C_g T^2
$$

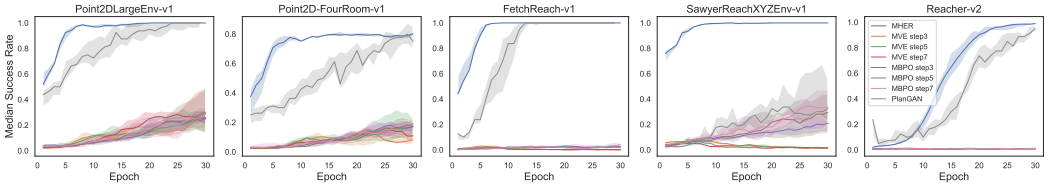

Figure 6: Comparison results with model-based baselines, including PlanGAN.

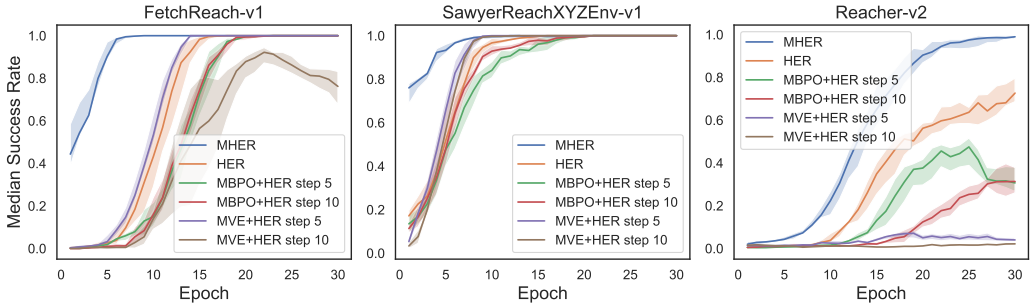

Figure 7: Comparison results with MBPO+HER and MVE+HER in FetchReach-v1, SawyerReachXYZEnv-v1, and Reacher-v2.

B Extra Experimental Results

In this section, we provide additional experimental results to further understand MHER and its effectiveness. We report the average success rate and standard deviation across 10 random seeds for all the experiments.

B.1 Comparison Results with Model-based Baselines

In our paper, we have compared with model-based RL baseines such as MVE [10] and MBPO [15]. Furthermore, we include PlanGAN as our baseline, which is not a RL method but a planning method requiring huge amount of computation. The procedure of selection one single action of planner in PlanGAN is: 1) first random sampling 20 initial actions, 2) further sampling 50 trajectories with model and GANs for each initial action, 3) computing average return of each initial action, and 4) selecting the best action in the initial set. Even if we ignore the additional computation to train GANs, PlanGAN needs at least 10^5 (20 initial action$\times$ 50 trajectories $\times$50 steps $\times$2 policy and model propagation) times forward propagations than MHER to select a single action. As PlanGAN is highly dependent on the learned model, we train the model and GANs in PlanGAN with a batch size of 128 and each epoch train 25 times. In MHER, we only update the model 2 times each epoch with a batch size of 64, therefore MHER is more computationally friendly compared to PlanGAN. We report the median test success rate and interquartile across 10 random seeds in Figure 6, which apparently shows MHER is more sample efficient than PlanGAN. Moreover, MHER is much more stable compared with PlanGAN and other model-based baselines considering the deviation in the results.

Moreover, we provide comparison results with MBPO+HER and MVE+HER in Figure 7. In the two methods, we perform hindsight relabeling before model-based interaction, and then the model-based interaction is also driven by hindsight goals. Through combining with HER, MBPO and MVE work better in sparse reward tasks. However, their performance still cannot surpasses MHER's.

B.2 Study of the Model Layers

In this section, we study the impact of model layers in MHER by varying the number of layers (each of which is of 256 neurons). The empirical results in Figure 8 show the number of layers does not affect the performance much and even one layer can perform comparably to 4 layers. The dynamics model with 4 layers achieves stable and good performance in different tasks.

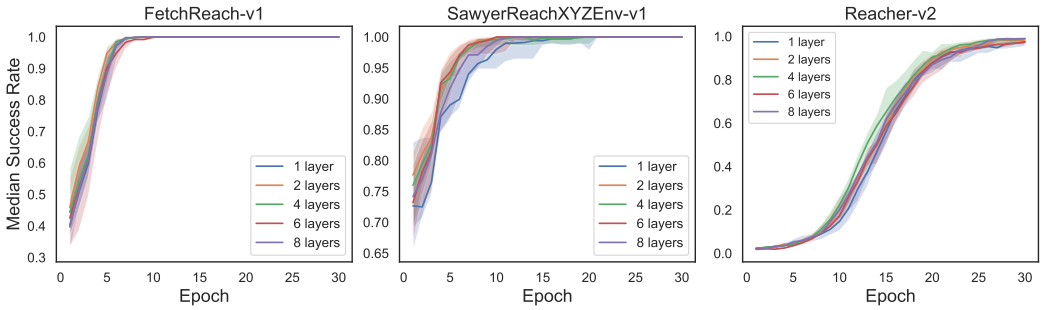

Figure 8: Varying the number of model layers in FetchReach-v1, SawyerReachXYZEnv-v1, and Reacher-v2.

B.3 Parameter Study

Additional parameter studies of α and model-based interaction steps are shown in Figure 9. In the results, we can conclude that the performance of MHER improves when parameter α increases from 0 to 3 in the two environments, but decrease slightly when α is beyond 3 in Reacher-v2. Regarding the model-based interaction steps, we can observe a performance increase as steps increase from 0 to 5, and a performance decrease as steps are more than 5.

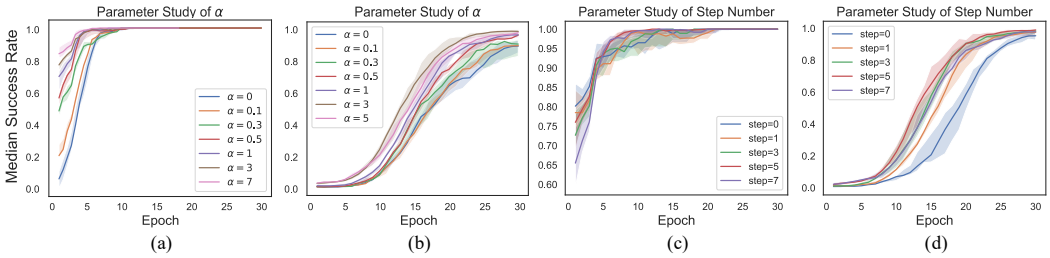

Figure 9: (a)(b) Additional parameter studies of α in SawyerReachXYZEnv-v1 and Reacher-v2. (c)(d) Additional parameter studies of model-based interaction steps in SawyerReachXYZEnv-v1 and Reacher-v2.

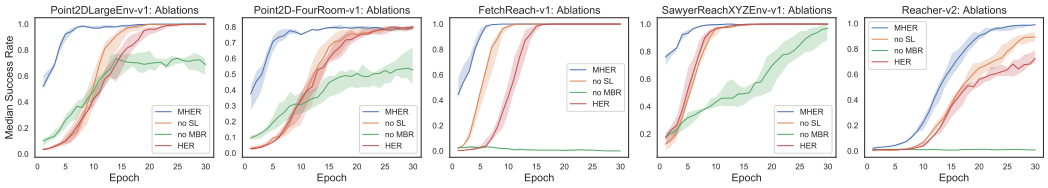

Figure 10: Additional ablation studies in benchmark environments.

B.4 Ablations

We also provide more ablation study results in Figure 10. The results are consistent with the conclusions in our paper: (1) MBR is more important than SL, (2) we can outperform HER using MBR alone, and (3) SL is not robust and fails to learn in FetchReach-v1 and Reacher-v2.

B.5 Additional Comparisons

We conducted additional experiments in Figure 11 to study how does MHER compare to *random relabeling* and HER with *goal noise*. For *random relabeling*, we relabel the transitions with random goals sampled from the goal space. Regarding *goal noise*, Gaussian noise with zero mean and constant (0.01) standard deviation is applied to hindsight goals of HER. We can observe from the results that *goal noise* provides slight improvement over HER in FetchReach-v1 and SawyerReachXYZEnv-v1,

but it worsens the result in Reacher-v2. In addition, it seems that *random relabeling* does not work, which might be because it hardly helps with the sparse reward problem.

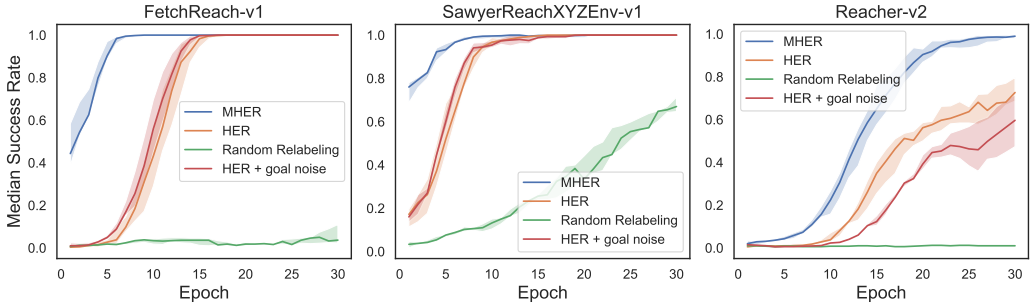

Figure 11: Comparison results between MHER, *random relabeling* and *goal noise*.

C Details about Virtual Achieved Goals

Previous work, HER [1], assumes there exists a state-to-goal mapping function: $\phi : S \to G$, i.e., given a state s, we can find a goal $g = \phi(s) \in G$ achieved by this state. In the case where each goal corresponds to a state we want to achieve (e.g., 2D point reaching task), goal space equals state space $G = S$ and the mapping ϕ is exactly an identity transformation.

Model-based relabeling is a novel goal relabeling method different from previous relabeling methods and enjoys many advantages by leveraging the dynamics model. MHER follows the same setting as HER. The difference is that MHER leverages the virtual states generated from model-based interaction rather than past collected states. As shown in Figure 12, the blue trajectory is collected by past policy, while the green trajectory is generated by interaction of current policy and the learned dynamics model.

Given a transition in a past collected trajectory $(s_t, a_t, r_t, s_{t+1}, g)$, model-based relabeling (MBR) aims to find a virtual achieved goal to replace the original goal g and reward r_t. MBR interacts with the dynamics model m for n steps starting from s_{t+1}, and collects a virtual trajectory $\{s'_{t+i}, a'_{t+i}, s'_{t+i+1}\}_{i=1}^{n}$, where $s'_{t+1} = s_{t+1}, a'_{t+i} = \pi(s'_{t+i}, g), s'_{t+i+1} = s'_{t+i} + m(s'_{t+i}, a'_{t+i}), i \in [1, n]$. Note that we start interaction from s_{t+1}, which is reasonable because starting from s_t means a_t, s_{t+1} also need to be replaced. After model-based interaction, MBR samples from virtual achieved goals $g' = \phi(s_{t+i}), i \in [1, n]$ and alternates the original transition as $(s_t, a_t, r'_t, s_{t+1}, g')$, where r'_t is the recomputed reward $r'_t = r(s_t, a_t, g')$ according to Eq. 1. The model-based relabeled data can be further used for off-policy reinforcement learning and goal-conditioned supervised learning.

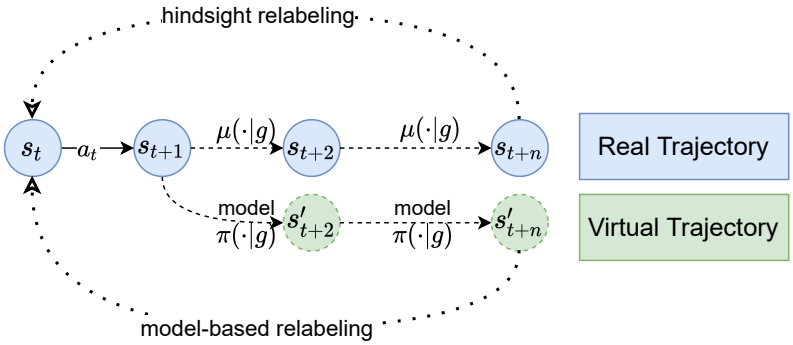

Figure 12: Simplified schematic of model-based relabeling.

D Implementation Details

D.1 Implementation of MHER

In this section, we will provide implementation details of MHER. The actor and critic networks of MHER are both 3-layer fully connected networks with 256 units each and ReLU non-linearities. The actor and critic are updated with Adam optimizer, learning rate 1×10^{-3}. Target nets of actor and critic with the same network structure are adopted for stabilizing training, the Polyak averaging coefficient of the target net is set as 0.9. The probability of random action is 0.3, and the scale of Gaussian noisy for exploration is 0.2. The model-based relabeling rate is 80%, and we keep 20% samples without relabeling. MHER uses one rollout worker, and the buffer size is 10^6. At the end of each episode (including 100 environment steps), we train the networks with 5 batches of size 64. We only use one single CPU and one GPU to train MHER.

Regarding the dynamics model in MHER, we use a fully connected network with 4 hidden layers and 256 neurons each layer. The optimizer is Adam and the learning rate is 0.001. The dynamics model is trained to minimize the loss $\mathcal{L}_{model}$ in Eq. 4. In the warmup period, we train the dynamics model for 100 updates with a batch size of 512. When training with MHER, every batch we update m with a batch size of 64 for 2 times.

D.2 Implementation of GCSL

We implement GCSL with a Diagonal Gaussian policy with a constant standard deviation of 0.2. GCSL only keeps a single policy network (3 layers, 256 neurons each) as the actor in MHER, without target network. The policy network is trained to minimize the loss

$$\mathcal{L}_{GCSL} = E_{(s_t,a_t,g')\sim B_h}\left[\|a_t - \pi(s_t, g')\|_2^2\right]$$

where g' is relabeled using future achieved goals like HER [1]. The optimizer is Adam and the learning rate is also $1 \cdot 10^{-3}$. GCSL shares other parameters with MHER for fair comparison, such as replay buffer size and batch size.

D.3 Implementation of MVE and MBPO

Model-based value expansion (MVE) [10] rollouts with a learned dynamics model and introduces multi-step value estimation based on the expanded transitions. Denote the replay buffer as B, a real transition $(s_t, a_t, r_t, s_{t+1}, g)$, generated transitions $\{\hat{s}_{t+i}, \hat{a}_{t+i}, \hat{r}_{t+i}\}, i \in [1, H], \hat{s}_{t+1} = s_{t+1}$, MVE updates the Q-function to minimize the following loss:

$$\mathcal{L}_{critic} = E_{(s_t,a_t,g,r_t,s_{t+1})\sim B}\Big[\Big(\sum_{i=0}^{H-1} \gamma^i \hat{r}_{t+i} + \gamma^H Q(\hat{s}_{t+H}, \pi(\hat{s}_{t+H}, g), g) - Q(s_t, a_t, g)\Big)^2\Big], \quad (9)$$

where $\hat{r}_t = r_t, \hat{r}_{t+i} = r(\hat{s}_{t+i}, \pi(\hat{s}_{t+i}, g), g)$, r is the reward function in Eq. 1. We also use DDPG to learn the policy, and the hyper-parameters of DDPG and dynamics model are the same as MHER. We provide empirical results with varying H in Figure 6.

Model-based policy optimization (MBPO) [15] stores short model-generated rollouts branched from real data to the model dataset D_{model}, and optimizes policy using the model dataset D_{model}. For fair comparison, we also utilize DDPG for policy optimization and the size of D_{model} is $1 \cdot 10^6$. Hyper-parameters of DDPG and dynamics model are the same as MHER. The results with model-based horizon $\{3, 5, 7\}$ are reported in Figure 6.

E Comparison with Concurrent Work

We noticed a concurrent work, MapGo [35], which shares similar idea of model-based relabeling (MBR) with us. However, MapGo doesn't exploit the MBR goals deeper. Instead, we further leverage MBR goals for supervised policy learning with theoretical guarantees. Moreover, MapGo utilizes a large number of virtual rollout for policy improvement, and thus introduces model error in both states and actions. In Section 4.2, we have emphasized that MHER avoids training with fully virtual states, and only the goals in the training data are generated by the model. Therefore, our method can be

more robust to model errors. Despite the differences, MapGo found the model cannot fit exactly for hard manipulation tasks, which we have also observed in our experiments and we believe it is worth further research.

F Task Descriptions

All of the five tasks have continuous state space, action space and goal space. In this section, we will introduce these tasks in detail.

F.1 Point2DLargeEnv-v1

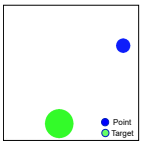

Point2DLarge

Point2DLargeEnv-v1 is taken from the open-source package *multi-world* and it requires the blue point to reach the green circle. The state space $[-5, 5] \times [-5, 5]$ has two dimensions representing Cartesian coordinates of the blue point, and the action space $[-1, 1] \times [-1, 1]$ also has two dimensions meaning the horizontal and vertical displacement. The goal space is the same as state space, which means $\phi(s) = s$. The bule point and the green circle are randomly initialized in the state space. The allowable error ϵ of reaching goal is the radius of the target circle and is set as 1. The reward function is defined as:

$$r(s_{XY}, a, g_{XY}) = -1(\|s_{XY} - g_{XY}\|_2^2 > \epsilon).$$

F.2 Point2D-FourRoom-v1

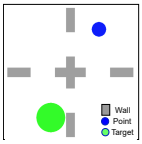

PointFourRoom

The Point2D-FourRoom-v1 environment is also built on *multi-world*. The state space, the action space, the goal space, and the reward function are same as Point2DLarge-v1. The difference is that there are four rooms separated by gray walls.

F.3 FetchReach-v1

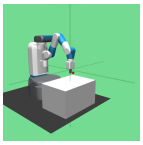
FetchReach

The FetchReach-v1 environment is taken from OpenAI Gym [4]. In this environment, a 7-DoF robotic arm is expected to touch a desired location with its two-finger gripper. The state space is 10-dimensional, including the gripper's position and linear velocities. The action space is 4-dimensional, which represents the gripper's movements and its status about the opening and closing. Moreover, the goals are 3-dimensional vectors representing the target place of the gripper. The state-to-goal mapping is $\phi(s) = s[0:3]$, because the first 3 dimensions of state describe the position of the gripper. The allowable error in FetchReach is $\epsilon = 0.05$. The reward function is defined as:

$$r(s_{XYZ}, a, g_{XYZ}) = -1(\|s_{XYZ} - g_{XYZ}\|_2^2 > \epsilon).$$

F.4 SawyerReachXYZEnv-v1

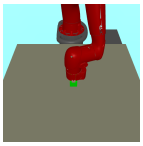
SawyerReach

The SawyerReachXYZEnv environment is taken from *multi-world*. The Sawyer robot aims to reach a target position with its end-effector. The observation space is 3-dimensional, representing the 3D Cartesian position of the end-effector. Correspondingly, the goal space is 3-dimensional and describes the expected position, and the state-to-goal mapping is $\phi(s) = s$. Besides, the action space has 3 dimensions describing the next position of the end-effector. The reward function is the same as FetchReach except the allowable error $\epsilon = 0.06$.

F.5 Reacher-v2

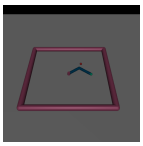
Reacher

The Reacher environment is revised from OpenAI Gym [4]. States are 11-dimensional, which indicates the angles, the positions, and the velocity of the joints. Actions are 2-dimensional and control the movement of two joints. The

goals are 2-dimension representing the expected XY position. And the state-to-goal mapping is $\phi(s) = s[-3:-1]$, where the last three dimensions are the XYZ position of the end-effect. The reward function is defined as:

$$r(s_{XY}, a, g_{XY}) = -1(\|s_{XY} - g_{XY}\|_2^2 > \epsilon),$$

where the allowable error ϵ is set as 0.02.