

MapGo: Model-Assisted Policy Optimization for Goal-Oriented Tasks

Menghui Zhu^{1,2*}, Minghuan Liu^{1*}, Jian Shen¹, Zhicheng Zhang¹, Sheng Chen²,
Weinan Zhang^{1†}, Deheng Ye², Yong Yu¹, Qiang Fu² and Wei Yang²

¹Shanghai Jiao Tong University, Shanghai, China

²Tencent AI Lab, Shenzhen, China

{zerozmi7,minghuanliu,wnzhang}@sjtu.edu.cn, {victchen,dericye,leonfu,willyang}@tencent.com

Abstract

In Goal-oriented Reinforcement learning, relabeling the raw goals in past experience to provide agents with hindsight ability is a major solution to the reward sparsity problem. In this paper, to enhance the diversity of relabeled goals, we develop FGI (Foresight Goal Inference), a new relabeling strategy that relabels the goals by looking into the future with a learned dynamics model. Besides, to improve sample efficiency, we propose to use the dynamics model to generate simulated trajectories for policy training. By integrating these two improvements, we introduce the MapGo framework (Model-Assisted Policy Optimization for Goal-oriented tasks). In our experiments, we first show the effectiveness of the FGI strategy compared with the hindsight one, and then show that the MapGo framework achieves higher sample efficiency when compared to model-free baselines on a set of complicated tasks.

1 Introduction

In this paper, we consider the problem of Goal-oriented Reinforcement Learning (GoRL), where agents are expected to complete a series of similar tasks and learn to reach multiple goals. Typically, the reward is designed as binaries to indicate whether the agent reaches the goal, which, however, is too sparse for the agent to learn a good policy efficiently. To that end, in the literature, several works have been proposed to relabel or re-generalize different goals along past trajectories [Andrychowicz *et al.*, 2017; Fang *et al.*, 2018; Fang *et al.*, 2019], trying to reuse the past failed experiences by relabeling them with reachable goals. An early-stage example, Hindsight Experience Replay (HER) [Andrychowicz *et al.*, 2017], is inspired by human’s hindsight ability that one can always learn from an undesired outcome as from a desired one. Besides, other works [Florensa *et al.*, 2018; Ren *et al.*, 2019] pay attention to choosing suitable behavioral (training-time) goals to help explore the environment in a more efficient manner.

*Equal contribution. This work was done when the first author was an intern at Tencent. † marks the corresponding author. Our code is available at <https://github.com/apexrl/MapGo>.

However, existing works have two significant issues. First, most of them are model-free methods that are considered less sample-efficient at most times. In contrast, model-based reinforcement learning (MBRL) has shown better sample efficiency compared to the model-free counterparts [Langlois *et al.*, 2019]. This motivates us to investigate the feasibility of model-based techniques to GoRL. The second issue lies in the limited guidance ability of hindsight goals [Andrychowicz *et al.*, 2017]. Hindsight goals are usually sampled from trajectories of a past policy, which may not be appropriate for the current policy. Also, the diversity of the selected goals is limited since the relabeled goals and the updated tuples (s, a, s', g) are chosen from the same trajectory. Therefore, generating more diverse goals is expected.

Based on these considerations, we develop a model-based framework for GoRL, called MapGo (Model-Assisted Policy optimization for Goal-Oriented tasks). In MapGo, we learn a bootstrapped ensemble of dynamics models [Chua *et al.*, 2018] to help tackle the two issues above. First, we propose the Foresight Goal Inference (FGI) strategy, which uses the dynamics models to relabel the goals by looking into the future. Besides, the learned dynamics models are further used to generate branched short rollouts for policy optimization in a Dyna-style manner [Janner *et al.*, 2019].

We conduct extensive experiments on a set of continuous control tasks. Specifically, we first compare the relabeled goals of FGI with HER [Andrychowicz *et al.*, 2017] to verify the efficacy of FGI. Then we evaluate the MapGo framework on complicated goal-oriented learning benchmarks, indicating the higher sample efficiency compared with former model-free algorithms. Finally, we make a comprehensive ablation study to analyze the performance improvement of MapGo while also examining its limitation.

In a nutshell, our contributions are threefold.

- We design a novel goal relabeling mechanism called FGI that provides foresight goals with diversity.
- Based on FGI, we propose the MapGo framework, which incorporates model-based training techniques into GoRL to improve the sample efficiency.
- We provide an in-depth understanding of the components of MapGo for contributing to the final performance.

2 Background

2.1 Goal-oriented Reinforcement Learning

Compared with standard Reinforcement Learning (RL) that is usually modeled as Markov Decision Process (MDP), Goal-oriented Reinforcement Learning (GoRL) augments it with a goal state. Formally, the Goal-oriented MDP (GoMDP) is denoted as a tuple $\langle \mathcal{S}, \mathcal{A}, M^*, \mathcal{G}, r_g, \gamma, \phi, \rho_0, p_g \rangle$, where $\mathcal{S} \in \mathbb{R}^{d_s}$ and $\mathcal{A} \in \mathbb{R}^{d_a}$ are the state space and the action space respectively, $M^* : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$ represents the transition function, $\gamma \in [0, 1]$ is the discount factor, $\mathcal{G} \in \mathbb{R}^{d_g}$ denotes the goal space and $\phi : \mathcal{S} \rightarrow \mathcal{G}$ is a tractable mapping function that provides goal representation. ρ_0 is the initial state distribution and p_g is the distribution of desired goals. In GoRL, a *desired goal* is a task to solve which are normally given by the environments; and a *behavioral goal* indicates the target used for sampling. Without carefully designed goal selection mechanism, these two goals are generally the same.

In GoMDP, the reward function $r : \mathcal{S} \times \mathcal{A} \times \mathcal{G} \rightarrow \mathbb{R}$ is goal-conditioned determining whether the goal is reached:

$$r_g(s_t, a_t, g) = \begin{cases} 0, & \|\phi(s_{t+1}) - g\| \leq \epsilon \\ -1, & \text{otherwise} \end{cases}. \quad (1)$$

Therefore, agents can utilize a universal value function [Schaul *et al.*, 2015] that uses a single function approximator to represent a set of goal-based value functions. Given a state $s_t \in \mathcal{S}$ and a goal $g \in \mathcal{G}$, the value under the policy π is defined as:

$$V^\pi(s_t, g) = \mathbb{E}_{\substack{a_t \sim \pi \\ s_{t+1} \sim M^*}} \left[\sum_t \gamma^t r(s_t, a_t, g) \right]. \quad (2)$$

The objective of GoRL is to find a goal-based policy $\pi(a|s, g)$ to maximize the expectation of discounted cumulative rewards over the distribution of p_g :

$$\eta = \mathbb{E}_{\substack{a_t \sim \pi \\ s_{t+1} \sim M^* \\ g \sim p_g}} \left[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t, g) \right]. \quad (3)$$

2.2 Hindsight Relabeling

In GoRL, agents are provided with a specific goal from environments at the beginning of an episode. Specifically, the agent first receives a behavioral goal g (from the environment or self-generated) at an initial state s_0 , then it starts to interact with the environment for T timesteps and generates a trajectory $\{(s_0, a_0, s_1, g, r_1), \dots, (s_T, a_T, s_{T+1}, g, r_{T+1})\}$. Due to the sparse reward signal shown in Eq. (1), the agent will not get any valuable reward r_t until it reaches the goal state. Therefore, Andrychowicz *et al.* [2017] proposed Hindsight Experience Replay (HER) to relabel the goals for those failed trajectories so that the agent can learn from failures. The often-used *future* strategy randomly chooses a state s_{t+k} after the relabeled state s_t ; the *final* strategy of HER is relabeling with the final state s_T of the corresponding trajectory; and the *episode* strategy instead chooses several states in the whole corresponding trajectory. Unlike these prior works, we may not try to find reasonable goals from historical data but let the agent predict and relabel the goals that it can achieve by building an environment dynamics model.

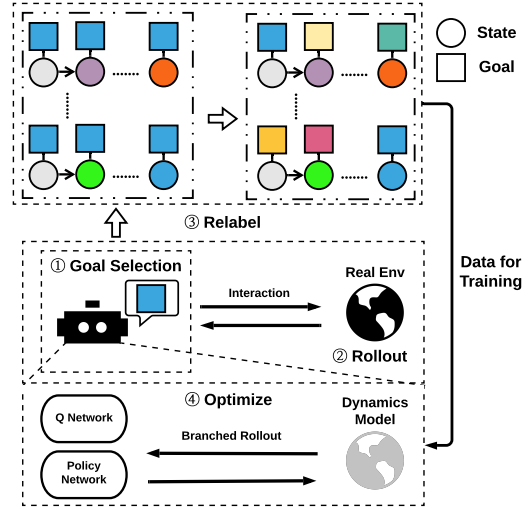


Figure 1: The overall framework of MapGo, containing four stages: i) *goal selection*: the agent selects a behavioral goal for sampling; ii) *rollout*: the agent performs rollout in the environment to generate trajectories; iii) *relabeling*: before training, trajectories samples from the environment are relabeled iv) *optimization*: the agent optimizes the policy and value networks. Different colors represent different states. Compared with other GoRL algorithms, MapGo focuses on enhancing the *relabeling* and *optimization* stages.

3 Methodology

In this section, we present our model-assisted policy optimization framework and explain how foresight goal inference works along with universal model-based policy optimization.

3.1 Overall Framework

We illustrate the overall framework of our Model-Assisted Policy Optimization for Goal-Oriented Tasks (MapGo) in Fig. 1. The agent starts with selecting the behavioral goals at an initial state and interacts with the environment to collect rollouts conditioned on the goal. Such experiences will then be relabeled by informative goal states, and the rewards will also be recomputed based on the relabeled goals, which will further be utilized for optimizing the policy and the value functions of the agent using off-policy RL algorithms. Our MapGo framework contributes mainly to the *relabeling* and *optimization* stage.

MapGo involves two functional components: i) Foresight Goal Inference (FGI), which is used to infer an appropriate goal for the agent in the *relabeling* stage; and ii) Universal Model-based Policy Optimization (UMPO), which is used to train the goal-based policy via generated short rollouts together with the experience sampled from the environment in the *optimization* stage. We detail the whole algorithm in Algo. 1, where FGI and UMPO¹ represent the above two modules respectively. Our framework leaves flexible choices for *goal selection* stage and we denote it as GoalGen.

Below we move on to the details of FGI and UMPO.

¹The detailed algorithm of UMPO can be found at Appendix A.1. The full version of this paper with a detailed appendix can be found at <https://arxiv.org/abs/2105.06350>.

Algorithm 1 Model-Assisted Policy Optimization (MapGo)

Input: Policy parameter θ , Q value parameter ω , and dynamics model parameter ψ

$\mathcal{D}_{\text{env}} \leftarrow \emptyset;$

for $i = 1$ to K **do**

$(s_0, g) \leftarrow \text{GoalGen}(\mathcal{D}_{\text{env}}, \mathcal{T}^*, \pi_\theta);$

for $t = 0$ to $h - 1$ **do**

$a_t \leftarrow \pi_\theta(s_t, g);$

$s_{t+1} \sim M^*(\cdot | s_t, a_t);$

$r_t \leftarrow r_g(s_t, a_t, s_{t+1});$

end for

$\tau \leftarrow \{s_0, a_0, r_0, s_1, a_1, r_1 \dots\};$

$\mathcal{D}_{\text{env}} \leftarrow \mathcal{D}_{\text{env}} \cup \{\tau\};$

 Update M_ψ according to Eq. (5);

$\mathcal{D}_{\text{real}} \leftarrow \text{FGI}(M_\psi, \mathcal{D}_{\text{env}}, \pi, \phi, r);$

$\pi_\theta, Q_\omega \leftarrow \text{UMPO}(\pi_\theta, Q_\omega, \mathcal{D}_{\text{env}}, \mathcal{D}_{\text{real}}, M_\psi);$

end for

3.2 Foresight Goal Inference

Foresight Goal Inference (FGI) serves for the *relabeling* stage, which is designed for providing future reached goals w.r.t. the current policy based on the historical experience.

Dynamics Modeling

We implement such a foresight ability with simulated rollouts by utilizing an environment dynamics model. Generally, the environment dynamics, i.e., transition function M^* is assumed to be unknown in RL. However, we can build an approximated dynamics model M during interactions with the environment by executing a policy $\pi(a|s)$:

$$M(s'|s, a) = \mathbb{E}_{(s', a, s') \sim \mathcal{D}} [P(s'|s, a)], \quad (4)$$

where \mathcal{D} is a replay buffer in which the experiences are sampled by the policy π . Specifically, we use a bootstrapped probabilistic dynamics model [Chua *et al.*, 2018], which can capture both aleatoric uncertainty and epistemic uncertainty.

Formally, we denote M_ψ as a parameterized dynamics model and $\{M_\psi^i\}_N$ as its ensemble. For each M_ψ^i , it approximates the transition probability $p(s'|s, a)$ as a Gaussian distribution with diagonal covariances $\mathcal{N}(\mu_\psi^i(s, a), \Sigma_\psi^i(s, a))$ given a state s and an action a . The models are trained by minimizing the negative log likelihood:

$$\mathcal{L}(\psi) = \sum_{n=1}^N [\mu_\psi(s_n, a_n) - s_{n+1}]^\top \Sigma_\psi^{-1}(s_n, a_n) [\mu_\psi(s_n, a_n) - s_{n+1}] + \log \det \Sigma_\psi(s_n, a_n). \quad (5)$$

Goal Inference

With a well-trained dynamics model, the agent is capable of foreseeing a considerable future goal. Particularly, under certain state s_t , we replace the original desired goal state g_t with a foresight goal \hat{g}_t by reasoning forward K steps using an ensemble transition model M_ψ and a sampling policy $\hat{\pi}_\theta$:

$$\hat{g}_t = \mathbb{E}_{\hat{a}_{t+i} \sim \hat{\pi}_\theta(\hat{s}_{t+i}), \hat{s}_{t+i+1} \sim M_\psi} [\phi(\hat{s}_{t+K})], \quad 0 \leq i \leq K - 1, \quad (6)$$

Algorithm 2 Foresight Goal Inference (FGI)

Input: dynamics model f_θ , Dataset \mathcal{D}_{env} , Current policy π , Mapping function ϕ , Reward function r , Rollout maximum length H

Output: Relabeled dataset $\mathcal{D}'_{\text{env}}$

for every tuple $(s_t^i, a_t^i, r_t^i, g_t^i) \in \mathcal{D}_{\text{env}}$ **do**

 Let L_i be the length of the trajectory of the tuple;

 Randomly choose a rollout step h uniformly in $(1, L_i)$;

if $L_i > H$ **then**

 Relabeling g_t^i with $\phi(s_{t+H}^i)$ by *future* strategy HER;

else

 Choose a new behavioral goal g_{new}^i by *future* strategy HER;

 Adopt h steps rollouts on s_{t+1}^i using $\hat{a}_{t'}^i \sim \pi(\cdot | \hat{s}_{t'}^i, g_{\text{new}}^i)$ in M_ψ , $t + 1 \leq t' \leq t + h$ and get \hat{s}_{t+h+1}^i ;

 Relabeling g_t^i with $\phi(\hat{s}_{t+h+1}^i)$;

end if

 Calculate the new reward $r'_i = r_g(s_t^i, a_t^i, g_t^i)$;

end for




Figure 2: Illustration of the FGI and the HER relabeling strategies.

where $\hat{a}_{t+i} \sim \hat{\pi}_\theta(\cdot | \hat{s}_{t+i}, g_t)$ is the action sampled by $\hat{\pi}_\theta$ and $\hat{s}_{t+i+1} \sim M_\psi(\cdot | \hat{s}_{t+i}, \hat{a}_{t+i})$ is the state sampled from M_ψ , which is the probabilistic dynamics model ensemble mentioned above. Specifically, at each timestep, we randomly choose one of the dynamics models to get a 1-step rollout.

Besides, there are more details worth noting. First, FGI should take the first rollout step from s_{t+1} when relabeling s_t since it is the tuple (s_t, a_t, g_t) that be used for training. Second, we take the rollout H steps at maximum due to the model precision and leave the longer relabeling with the future strategy HER. We show the FGI algorithm step-by-step in Algo. 2.

We illustrate how FGI works in Fig. 2 compared with prior HER relabeling methods. FGI advances in several aspects: i) HER limits the relabeled goal in the sampled trajectory, yet FGI generates various goals determined by the dynamics model; ii) FGI regenerates goals using the current model and current policy even for the historical trajectories, while HER still depends on old ones, which may significantly differ from the current policy.

3.3 Universal Model-Based Policy Optimization

In this paper, we expand model-based policy optimization [Janner *et al.*, 2019] into the goal-based setting and propose Universal Model-based Policy Optimization (UMPO) for efficient goal-based policy learning. By extending the Q function $Q(s, a)$ to the universal Q function $Q(s, a, g)$ and applying DDPG [Lillicrap *et al.*, 2015] as the policy optimization algorithm, the updating of the parameterized policy and Q function in UMPO can be written as follows:

$$\nabla_{\theta} J(\pi_{\theta}) = \mathbb{E}_{\substack{(s,g) \sim \hat{\mathcal{D}} \\ a \sim \pi}} [\nabla_a Q_{\omega}(s, a, g) \nabla_{\theta} \pi_{\theta}(s, g)] \quad (7)$$

$$\nabla_{\omega} J(Q_{\omega}) = \mathbb{E}_{\substack{(s,a,s',g) \sim \hat{\mathcal{D}} \\ a' \sim \pi}} [\nabla_{\omega} (Q_{\omega}(s, a, g) - r - \gamma Q_{\omega'}(s', a', g))^2], \quad (8)$$

where $Q_{\omega'}$ is the target Q function and $\hat{\mathcal{D}} \triangleq \alpha \mathcal{D}_{\text{env}} + (1 - \alpha) \mathcal{D}_{\text{model}}$ is a mixture of the real experience \mathcal{D}_{env} in the environment and generated short branched rollouts $\mathcal{D}_{\text{model}}$ with the learned model.

Unlike in classic RL environments, in GoRL setting, branched rollouts also need a behavioral goal g . Thus, how to choose appropriate goals when performing branched rollouts is a crucial problem. In UMPO, instead of taking the goal g from the tuple (s_t, a_t, s_{t+1}, g_t) directly, we randomly select a state s_{t+n} from the same trajectory and replace g_t with $\phi(s_{t+n})$ as in HER. It is believed that such a branched rollout strategy could improve the diversity of training samples and thus increase the performance, as evaluated in the experiment.

4 Related Work

4.1 Goal-oriented Reinforcement Learning

As mentioned above, GoRL algorithms roughly contain four steps: goal selection, rollout, relabeling, and optimization. An extensive set of previous works pay attention to the goal selection and the relabeling stage.

Goal Selection. The goal selection stage guides the agent to replace the raw desired goal with a currently appropriate one to achieve. For example, Florensa *et al.* [2018] proposed to utilize the generative adversarial network to generate the goals at appropriate difficulty. Ren *et al.* [2019] optimized the desired goals along with learning the policy, which selects them according to the result of the formulated Wasserstein Barycenter problem. Hartikainen *et al.* [2019] chose maximum distance goals based on a learned distance function. And Pitis *et al.* [2020a] sampled the goals which could maximize the entropy of the historical achieved goals.

Relabeling. The works for relabeling aims to increase sample efficiency and can be traced back to Hindsight Experience Replay (HER) [Andrychowicz *et al.*, 2017], which shows that off-policy algorithms can reuse the data in the replay buffer by relabeling goals in the previous trajectory. As an improvement, Fang *et al.* [2019] proposed CHER that adaptively selects the failed experiences for replay according to the proximity to true goals and the curiosity of exploration. Other

works combine HER with other goals like generated goals by variational auto-encoder [Nair *et al.*, 2018], past achieved goals, and actual desired goals [Pitis *et al.*, 2020a] given by the environments.

Besides, Trott *et al.* [2019] proposed a sibling rivalry (SR) mechanism to avoid getting to local optima and get a better exploration while sampling in the environment, which can not be classified into the former two categories.

4.2 Model-Based Reinforcement Learning

Our work adopts the dynamics model for foresight goal inference and further develops the usage of model-based techniques in GoRL. Model-Based Reinforcement Learning (MBRL) has been active in the community for years due to its high sample efficiency, which needs fewer interactions with the environment to derive a good policy. A common way for using the model is the Dyna-style algorithm [Sutton, 1991; Feinberg *et al.*, 2018; Buckman *et al.*, 2018], which improves the policy through sampling in the transition model; besides, Model Predictive Control (MPC) [Camacho and Alba, 2013; Wang and Ba, 2019; Boney *et al.*, 2020] does not maintain a policy but uses the model to look forward finite steps with several action sequences and chooses the first action with maximum cumulative rewards.

Recently, Luo *et al.* [2018] built a lower bound of the expected reward and maximizes the lower bound jointly over the policy and the model. Chua *et al.* [2018] adopted an ensemble of neural network models while applying MPC to sample good actions. Another important work comes from Janner *et al.* [2019], which analyzes the compounding errors in model rollouts and builds a lower bound of the value discrepancy between the model and the environment. The resulting algorithm named Model-Based Policy Optimization (MBPO) generates short rollouts from the real states and optimizes the policy with both the real and the generated samples.

Also, some methods absorbing the idea of MBRL are proposed in the field of GoRL. Nair *et al.* [2020] built a goal-aware dynamics model to capture the dynamics in some multi-goal environments with only image inputs. Pitis *et al.* [2020b] learned local casual models to generate counterfactual experiences, which significantly improves the performances of RL agents in locally factored tasks.

5 Experiments

In this section, we conduct several experiments from different dimensions to answer the following questions:

1. Does our FGI relabeling trick work as expected and is more advanced than HER?
2. Can MapGo show advantage over previous methods on complicated goal-oriented tasks?
3. Which component of MapGo is the most effective?

To answer the first question, we demonstrate the effectiveness of FGI in a simple 2D locomotion task. We show the relabeled goals by FGI during different training procedures and the learning curve compared to HER. To answer the second question, we evaluate the performance of the proposed MapGo architecture on four complicated goal-oriented

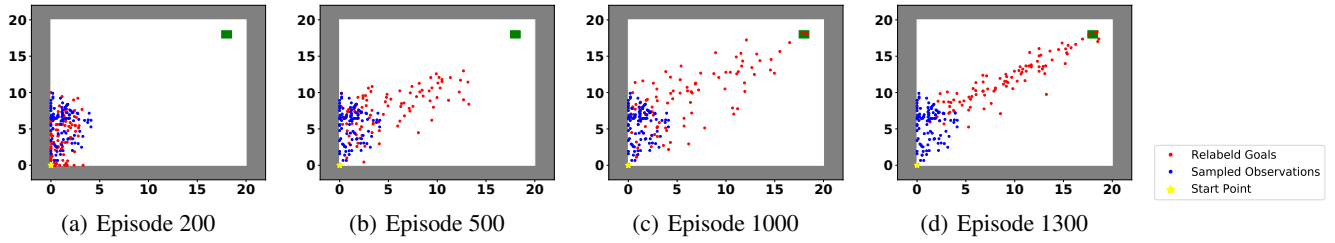


Figure 3: Relabeled goal points (red) and on the same trajectories points (blue) on 2D-World during training. Agents tries to learn to go to the goal (green square) from one start point (yellow star). Note that the trajectories are sampled by an initialized policy and are saved in the replay buffer for training. It is obvious that FGI relabel different goals according to the current policy even for the same trajectories, which provides much more information for the learning policy. However, HER is only constrained in the sampled trajectories (blue) at every training stage.

benchmarks. To answer the third question, we conduct an ablation study where we compare the functionality of FGI and UMPO separately. We also make further analysis on model rollout strategies and the limitations of MapGo.

5.1 Relabeling Goals Using FGI

We first show that our FGI relabeling trick offers fresher information for policy training than HER does. We conduct the comparison experiments on 2D-World, a simple 2D navigation task where the state space is the position of the agent $(x, y) \in [0, 20]^2$, and the action space is the movement $(\delta_x, \delta_y) \in [-1, 1]^2$. The initial state $s_0 = (0, 0)$ is fixed in all episodes. This task aims to reach a target that is randomly sampled in $[18.5, 19.5]^2$, and the agent receives a nonnegative reward only when the agent gets close enough to the target. We choose the maximum rollout length H as 20 and do not use extra relabeling methods from HER.

We illustrate the relabeled goals by FGI at different training stages in Fig. 3, where the blue points represent trajectories sampled by an initialized policy for off-policy training and the reds are the relabeled goals. As the training goes on, the relabeled goals are moving closer to the goal area, even for the same trajectories. It indicates that FGI relabels goals adapting the current policy and can provide much more useful information. On the other hand, HER can only relabel goals that coincide with the data in previous trajectories. As shown in Fig. 3, for a previous trajectory (blue points), HER still relabels goals from them without any additional information as the policy changes. The training curves are shown in Fig. 4, verifying the advantage and efficacy of FGI.

5.2 Complicated Goal-oriented Tasks

In this section, we aim to show the advantage of MapGo compared with previous methods on four more challenging goal-oriented tasks: Reacher, HalfCheetah, Fixed Ant Locomotion, and Diverse Ant Locomotion. In Reacher, we control a 2D robot to reach a randomly located target [Charlesworth and Montana, 2020]. Fixed Ant Locomotion and Diverse Ant Locomotion are similar to the environment in [Florensa *et al.*, 2018], where the agent is asked to move to a target position.

The main difference between them is whether the target is around a fixed far-away area or in the whole state space. It is




Figure 4: Training curves on 2D-World, where the solid line and the shade represent the mean and the standard deviation of the success rate separately, which is computed over 5 different random seeds. This shows the advantage of FGI which helps agent learn faster compared with HER.

worth noting that a fixed area does not mean that the environment is easier. A diverse distribution in another way provides the agent with the chance to learn from similar goals, while the fixed point desired goal makes the rewards even more sparse. HalfCheetah resembles the environment in [Finn *et al.*, 2017] and requires the agent to run while keeping a targeted speed until the end of the episodes.

In detail, we compare MapGo with HER, Hindsight Goal Generation (HGG) [Hartikainen *et al.*, 2019], GoalGAN [Florensa *et al.*, 2018] and OMEGA [Pitis *et al.*, 2020a]. We also test the performance of MapGo with HGG selecting the behavioral goal at each initial state (denoted as MapGo w. HGG). In Reacher, the rollout maximum length is 50 and in the others, we set it as 30. We utilize DDPG [Lillicrap *et al.*, 2015] as the learning algorithm for all the methods.

We plot the learning curve in Fig. 5, where MapGo outperforms all the other baselines and achieves higher sample efficiency in the most environments. Specifically, in Diverse Ant Locomotion, our algorithm takes about 2.0×10^5 timesteps to get an averaged success rate of 0.95, while the best baselines OMEGA is still at about 0.4 at the same time. Moreover, in HalfCheetah, MapGo learns much faster and only takes about 40 thousand steps to get a good performance. The only excep-

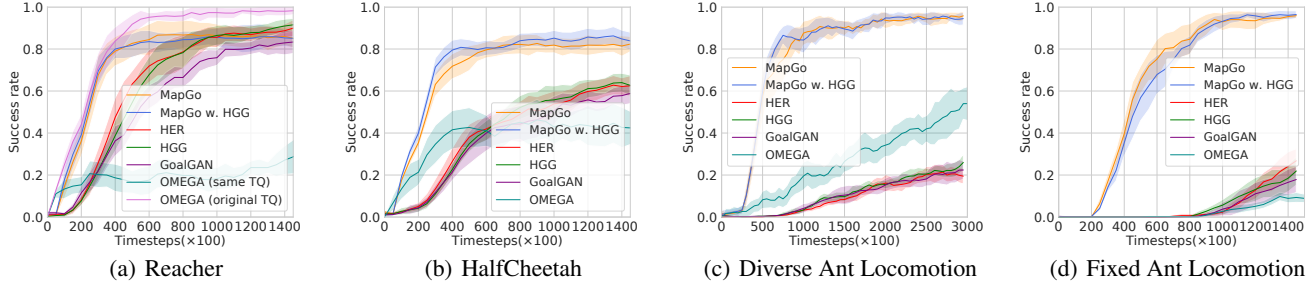


Figure 5: Learning curves of MapGo (ours) compared with baselines on three goal-oriented tasks. MapGo w. HGG denotes MapGo using HGG for behavioral goal selection. The solid line indicates the mean and the shaded area is the standard error over 5 different random seeds. Each trial is evaluated every 5000 environment steps, where each evaluation reports the average return over 100 episodes.

tion is that in Reacher, OMEGA performs better than MapGo because of a 50 times training frequency. For fairness we test OMEGA with different training frequencies (denoted as *original TQ* and *same TQ*), and show that MapGo actually still outperforms it with the same training frequency. We also find that the goal selection technique makes limited improvements on MapGo. One possible reason is that the model rollouts could provide enough exploration ability as HGG does.

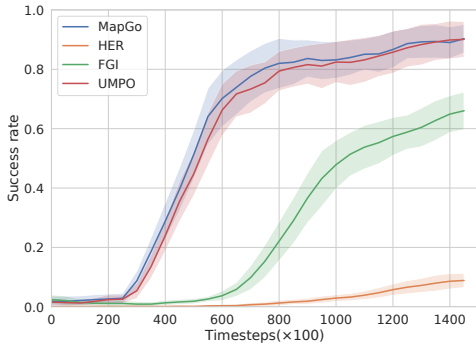


Figure 6: Ablation experiments on Diverse Ant Locomotion over 5 different random seeds. We analyse which component of MapGo works most by running UMPO and FGI respectively. It is obvious that UMPO contributes much more but by combining UMPO with FGI, MapGo achieves the best sample efficiency.

5.3 Ablation Study

To figure out how FGI and UMPO contribute to the overall performance of MapGo, we perform an ablation study on Diverse Ant Locomotion. Specifically, we compare the performance of the whole MapGo as conducted in the last section with running UMPO and FGI separately. Note that we do not use any relabel or goal selection techniques for UMPO.

As shown in Fig. 6, it is easy to conclude that UMPO is more effective, which can achieve close sample efficiency as MapGo. However, with combining FGI and UMPO, MapGo still improves and achieves the best sample efficiency. Besides, compared to HER, FGI still outperforms it substantially, which proves its effectiveness.

5.4 Different Branched Rollout Strategies

In MapGo, to encourage the exploration and increase the diversity in model rollouts, we utilize hindsight relabeling to set a behavioral goal before sampling with the model. In this section, we test it with other different rollout strategies on Diverse Ant Locomotion: 1) *NoRelabel*, taking the original desired goal from the environment for rollouts; 2) *NowDesired*, choosing the behavior goals using a current HGG goal selector. As shown in Fig. 7, the relabel strategy used in MapGo when performing branched rollouts is much more effective than the others. This implies that the goals should be diverse to achieve better performance.




Figure 7: Ablation experiments on Diverse Ant Locomotion over 3 different random seeds. We analyse the model rollout strategy of MapGo frameworks and we compare MapGo with *NoRelabel* and *NowDesired* strategy.

5.5 Limitations of MapGo

Although MapGo provides visible improvements in sample efficiency on complicated goal-oriented tasks, we surprisingly find it performs poorly for some robotic tasks. To understand such a situation, we compare the loss of the learned models on Ant locomotion and FetchPush and draw the learning curves in Appendix A.5. It shows that the current way of learning a dynamics model is unhelpful for such a robotic task, and the loss diverges.

This can be attributed to the fact that the transition function in such environments is somewhere discontinuous.

For example, small movements of robots could separate the arm and the object, thus significantly affecting the transition function. Since MapGo relies heavily on the learned model, which serves to both FGI and UMPO, this can lead to poor performances. Similar observations can be found in Pitis *et al.* [2020b], where the authors try to build a mask to represent the adjacency matrix of the sparsest local causal graph by dynamics model. However, in robotics tasks like *FetchPush*, all the masks are generated by domain knowledge rather than a dynamics model. It implies that the dynamics model limits the application of MapGo only if we can improve the accuracy of the learned model in these environments.

6 Conclusion

We tackle the problem of GoRL and aim to solve the sample efficiency and the diversity of relabeled goals. Specifically, we propose MapGo, which is integrated by a novel relabeling strategy named FGI that equips the agent with the foresight ability by looking into the future with the current policy, and UMPO that applies branch rollout to increase the sample efficiency. Qualitative and quantitative experiments show the effectiveness of MapGo and evaluate each component comprehensively. In the future, we will further study cases of discontinuous transitions when using dynamics models.

Acknowledgements

This work is supported by the “New Generation of AI 2030” Major Project (2018AAA0100900), National Natural Science Foundation of China (62076161, 61632017) and the Tencent Rhino-Bird Research Project.

References

- [Andrychowicz *et al.*, 2017] 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 *NeurIPS*, pages 5048–5058, 2017.
- [Boney *et al.*, 2020] Rinu Boney, Juho Kannala, and Alexander Ilin. Regularizing model-based planning with energy-based models. In *CoRL*, pages 182–191, 2020.
- [Buckman *et al.*, 2018] Jacob Buckman, Danijar Hafner, George Tucker, Eugene Brevdo, and Honglak Lee. Sample-efficient reinforcement learning with stochastic ensemble value expansion. In *NeurIPS*, pages 8224–8234, 2018.
- [Camacho and Alba, 2013] Eduardo F Camacho and Carlos Bordons Alba. *Model predictive control*. Springer Science & Business Media, 2013.
- [Charlesworth and Montana, 2020] Henry Charlesworth and Giovanni Montana. Plangan: Model-based planning with sparse rewards and multiple goals. *arXiv preprint arXiv:2006.00900*, 2020.
- [Chua *et al.*, 2018] Kurtland Chua, Roberto Calandra, Rowan McAllister, and Sergey Levine. Deep reinforcement learning in a handful of trials using probabilistic dynamics models. In *NeurIPS*, pages 4754–4765, 2018.
- [Fang *et al.*, 2018] Meng Fang, Cheng Zhou, Bei Shi, Boqing Gong, Jia Xu, and Tong Zhang. Dher: Hindsight experience replay for dynamic goals. In *International Conference on Learning Representations*, 2018.
- [Fang *et al.*, 2019] Meng Fang, Tianyi Zhou, Yali Du, Lei Han, and Zhengyou Zhang. Curriculum-guided hindsight experience replay. In *NeurIPS*, pages 12623–12634, 2019.
- [Feinberg *et al.*, 2018] Vladimir Feinberg, Alvin Wan, Ion Stoica, Michael I Jordan, Joseph E Gonzalez, and Sergey Levine. Model-based value estimation for efficient model-free reinforcement learning. *arXiv preprint arXiv:1803.00101*, 2018.
- [Finn *et al.*, 2017] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. *arXiv preprint arXiv:1703.03400*, 2017.
- [Florensa *et al.*, 2018] Carlos Florensa, David Held, Xinyang Geng, and Pieter Abbeel. Automatic goal generation for reinforcement learning agents. In *ICML*, pages 1515–1528. PMLR, 2018.
- [Hartikainen *et al.*, 2019] Kristian Hartikainen, Xinyang Geng, Tuomas Haarnoja, and Sergey Levine. Dynamical distance learning for semi-supervised and unsupervised skill discovery. In *International Conference on Learning Representations*, 2019.
- [Janner *et al.*, 2019] Michael Janner, Justin Fu, Marvin Zhang, and Sergey Levine. When to trust your model: Model-based policy optimization. In *NeurIPS*, pages 12498–12509, 2019.
- [Langlois *et al.*, 2019] Eric Langlois, Shunshi Zhang, Guodong Zhang, Pieter Abbeel, and Jimmy Ba. Benchmarking model-based reinforcement learning. *arXiv preprint arXiv:1907.02057*, 2019.
- [Lillicrap *et al.*, 2015] 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.
- [Luo *et al.*, 2018] Yuping Luo, Huazhe Xu, Yuanzhi Li, Yuandong Tian, Trevor Darrell, and Tengyu Ma. Algorithmic framework for model-based deep reinforcement learning with theoretical guarantees. *arXiv preprint arXiv:1807.03858*, 2018.
- [Nair *et al.*, 2018] Ashvin V Nair, Vitchyr Pong, Murtaza Dalal, Shikhar Bahl, Steven Lin, and Sergey Levine. Visual reinforcement learning with imagined goals. In *NeurIPS*, pages 9191–9200, 2018.
- [Nair *et al.*, 2020] Suraj Nair, Silvio Savarese, and Chelsea Finn. Goal-aware prediction: Learning to model what matters. In *International Conference on Machine Learning*, pages 7207–7219. PMLR, 2020.

- [Pitis *et al.*, 2020a] Silviu Pitis, Harris Chan, Stephen Zhao, Bradly Stadie, and Jimmy Ba. Maximum entropy gain exploration for long horizon multi-goal reinforcement learning. *arXiv preprint arXiv:2007.02832*, 2020.
- [Pitis *et al.*, 2020b] Silviu Pitis, Elliot Creager, and Animesh Garg. Counterfactual data augmentation using locally factored dynamics. *NeurIPS*, 33, 2020.
- [Ren *et al.*, 2019] Zhizhou Ren, Kefan Dong, Yuan Zhou, Qiang Liu, and Jian Peng. Exploration via hindsight goal generation. In *NeurIPS*, pages 13464–13474, 2019.
- [Schaul *et al.*, 2015] Tom Schaul, Daniel Horgan, Karol Gregor, and David Silver. Universal value function approximators. In *International conference on machine learning*, pages 1312–1320, 2015.
- [Sutton, 1991] Richard S Sutton. Dyna, an integrated architecture for learning, planning, and reacting. *ACM Sigart Bulletin*, 2(4):160–163, 1991.
- [Trott *et al.*, 2019] Alexander Trott, Stephan Zheng, Caiming Xiong, and Richard Socher. Keeping your distance: Solving sparse reward tasks using self-balancing shaped rewards. In *NeurIPS*, pages 10376–10386, 2019.
- [Wang and Ba, 2019] Tingwu Wang and Jimmy Ba. Exploring model-based planning with policy networks. *arXiv preprint arXiv:1906.08649*, 2019.