

HINDSIGHT FORESIGHT RELABELING FOR META-REINFORCEMENT LEARNING

Michael Wan, Jian Peng & Tanmay Gangwani

University of Illinois at Urbana-Champaign

{mw3, jianpeng, gangwan2}@illinois.edu

ABSTRACT

Meta-reinforcement learning (meta-RL) algorithms allow for agents to learn new behaviors from small amounts of experience, mitigating the sample inefficiency problem in RL. However, while meta-RL agents can adapt quickly to new tasks at test time after experiencing only a few trajectories, the meta-training process is still sample-inefficient. Prior works have found that in the multi-task RL setting, relabeling past transitions and thus sharing experience among tasks can improve sample efficiency and asymptotic performance. We apply this idea to the meta-RL setting and devise a new relabeling method called Hindsight Foresight Relabeling (HFR). We construct a relabeling distribution using the combination of *hindsight*, which is used to relabel trajectories using reward functions from the training task distribution, and *foresight*, which takes the relabeled trajectories and computes the utility of each trajectory for each task. HFR is easy to implement and readily compatible with existing meta-RL algorithms. We find that HFR improves performance when compared to other relabeling methods on a variety of meta-RL tasks.

1 INTRODUCTION

Deep Reinforcement Learning (RL) has achieved success on a wide variety of tasks, ranging from computer games to robotics. However, RL agents are typically trained on a single task and are extremely sample-inefficient, often requiring millions of samples to learn a good policy for just that one task. Ideally, RL agents should be able to utilize their prior knowledge and adapt to tasks quickly, just as humans do. Meta-learning, or learning to learn, has achieved promising results in this regard, allowing agents to exploit the shared structure between tasks in order to adapt to new tasks quickly during meta-test time.

However, although meta-learned policies can adapt quickly during meta-test time, training these meta-learned policies still requires a massive amount of experience. Since meta-RL is inherently on-policy due to the assumption that the meta-train time and the meta-test time input distributions should match, many state-of-the-art meta-RL methods (Duan et al., 2016; Finn et al., 2017; Mishra et al., 2017; Rothfuss et al., 2018; Wang et al., 2016) utilize strictly on-policy data and thus suffer from sample-inefficiency especially during meta-training. Recent works (Fakoor et al., 2019; Rakelly et al., 2019) have sought to incorporate off-policy RL (Fujimoto et al., 2018; Haarnoja et al., 2018) into meta-RL to improve sample efficiency.

The combination of off-policy RL and relabeling, in which experience is shared across tasks, has been utilized in the multi-task RL setting, in which an agent learns to achieve multiple different yet related tasks, for both goal-reaching tasks (Andrychowicz et al., 2017) and more general multi-task settings (Eysenbach et al., 2020). Experience collected for one task may be completely useless for training a policy to learn that task, but could be extremely informative in training a policy to learn a different task. For example, an agent trying to shoot a hockey puck into a net might miss to the right. This experience could easily be used to train an agent to shoot a puck into a net positioned further to the right (Andrychowicz et al., 2017).

Both meta-RL and multi-task RL involve training on a distribution of tasks, so it follows that we can also combine relabeling techniques with meta-RL algorithms in order to boost both sample

Code: <https://www.github.com/michaelwan11/hfr>

efficiency and asymptotic performance. In meta-RL, an agent learns to explore sufficiently to identify the task it is supposed to be solving, and then uses that knowledge to achieve high task returns. The agent collects exploratory pre-adaptation data, then undergoes some adaptation process using that pre-adaptation data. Finally, after adaptation, the agent attempts to solve the task. Meta-RL algorithms typically have a meta-training phase followed by a meta-test phase. The goal during meta-training is to train the meta-parameters such that they could be quickly adapted to solve any task from the meta-train task distribution, given a small amount of data from that task. At meta-test time, given a new unseen task, the goal is to rapidly adapt the *learned* meta-parameters for this task, using a small amount of task-specific data. The focus in this paper is to improve the sample efficiency of the meta-training phase via data sharing.

Using concepts from maximum entropy RL (MaxEnt RL), we introduce a relabeling scheme for the meta-RL setting. Prior relabeling methods for multi-task RL have used the total reward of the trajectory under different tasks to guide the relabeling (Eysenbach et al., 2020; Li et al., 2020). Direct application of this type of relabeling to the meta-RL setting is potentially sub-optimal since the multi-task RL and meta-RL objectives are distinct (learning to perform many tasks vs. learning to learn a new task). Towards developing an approach more suited to meta-RL, we define the notion of the *utility* of a trajectory under the different tasks, where the utility captures the usefulness of the trajectory for efficient adaptation under those tasks. We call our method Hindsight Foresight Relabeling (HFR) – we use *hindsight* in replaying the experience using reward functions from different tasks, and we use *foresight* in computing the utility of trajectories under different tasks and constructing a relabeling distribution over tasks using these utilities. We demonstrate the efficacy of our method on a variety of robotic manipulation and locomotion tasks. Notably, we show that our method, as the first meta-RL relabeling technique (applied during meta-training) that we are aware of, leads to improved performance compared to prior relabeling schemes designed for multi-task RL.

2 RELATED WORK

Meta-learning, or learning to learn (Baxter, 1998; Naik et al., 1992; Schmidhuber, 1987; Thrun & Pratt, 1998), has been a topic of interest since the 1980s. Various approaches have been developed in recent years. Prior works have attempted to represent the RL process using a recurrent neural network (RNN) (Duan et al., 2016; Miconi et al., 2018; Wang et al., 2016) - the hidden state is maintained across episode boundaries and informs the policy as to what task it is currently solving. Similarly, Mishra et al. (2017) also maintain the internal state across episode boundaries while incorporating temporal convolution and attention into a recursive architecture. Gradient-based meta-learning methods have also been explored (Finn et al., 2017; Nichol & Schulman, 2018; Xu et al., 2018; Zheng et al., 2020). MAML (Finn et al., 2017) seek to learn a good policy initialization so that only a few gradient steps are needed to achieve good performance on unseen meta-test tasks. Stadie et al. (2018) build on this gradient-based approach but explicitly consider the effect of the original sampling distribution on final performance. Similar to these works, our relabeling method also considers the impact of pre-adaptation data on the post-adaptation performance. Another body of work focuses on designing strategies for structured exploration in meta-RL such that task-relevant information could be efficiently recovered (Liu et al., 2020; Rakelly et al., 2019; Zintgraf et al., 2019). Rakelly et al. (2019) devise an off-policy meta-RL method called PEARL that trains an encoder to generate a latent context vector on which the meta-RL agent is conditioned. Although we use PEARL as our base algorithm in this work, our relabeling scheme is general enough to be integrated into any off-policy meta-RL algorithm.

Experience Relabeling in Meta-RL. Recent work has studied the scope of sharing experience among tasks in the meta-RL paradigm. Mendonca et al. (2020) propose to tackle meta-RL via a model identification process, where context-dependent neural networks parameterize the transition dynamics and the rewards function. Their method performs experience relabeling only at the meta-test time, with the purpose of consistent adaptation to the out-of-distribution tasks. Crucially, there is no relabeling or sharing of data amongst tasks during the meta-train time. In contrast, the goal of the relabeling in HFR is to improve the sample efficiency of the meta-training phase. Dorfman et al. (2020) study the offline meta-RL problem. They propose reward relabeling as a mechanism to mitigate the “MDP ambiguity” issue, which the authors note is specific to the offline meta-RL setting. Their relabeling is based on random task selection. HFR, on the other hand, operates in the online meta-RL setting and provides a principled approach to compute a relabeling distribution that suggests

tasks for relabeling. We compare with the random relabeling method used in [Dorfman et al. \(2020\)](#) in our experiments. A more detailed comparison to these two prior works is included in Appendix A.7.

In the context of multi-task RL ([Caruana, 1997](#); [Kaelbling, 1993](#); [Schaul et al., 2015](#)), recent methods have proposed relabeling to improve the sample-efficiency ([Andrychowicz et al., 2017](#); [Eysenbach et al., 2020](#); [Li et al., 2020](#)). HER ([Andrychowicz et al., 2017](#)) relabels transitions using goals that the agent actually achieves. Doing so allows for learning even with a sparse binary reward signal. However, HER is only applicable to goal-reaching tasks and cannot be incorporated into meta-RL algorithms because the meta-RL agent is trained on a batch of tasks sampled from a fixed task distribution. Similar to our work, [Eysenbach et al. \(2020\)](#) use MaxEnt RL to construct an optimal relabeling distribution for multi-task RL and apply this to both goal-reaching tasks and tasks with arbitrary reward functions.

3 BACKGROUND

3.1 REINFORCEMENT LEARNING

In reinforcement learning (RL), the environment is modeled as a Markov Decision Process (MDP) $\mathcal{M} = (\mathcal{S}, \mathcal{A}, r, p, \gamma, p_1)$, where \mathcal{S} is the state-space, \mathcal{A} is the action-space, r is the reward function, p is the transition dynamics, $\gamma \in [0, 1)$ is the discount factor, and p_1 is the initial state distribution. At timestep t , the agent π_θ , parameterized by parameters θ , observes the state $s_t \in \mathcal{S}$, takes an action $a_t \sim \pi_\theta(a_t | s_t)$, and observes the next state $s_{t+1} \sim p(s_{t+1} | s_t, a_t)$ and the reward $r(s_t, a_t)$. The goal is to maximize the expected cumulative discounted rewards: $\max_\theta \mathbb{E}_{s_t, a_t \sim \pi_\theta} [\sum_{t=1}^{\infty} \gamma^{t-1} r(s_t, a_t)]$.

3.2 META-REINFORCEMENT LEARNING

In the general meta-reinforcement learning (meta-RL) setting, there is a family of tasks that is characterized by a distribution $p(\psi)$, where each task ψ is represented by an MDP $\mathcal{M}_\psi = (\mathcal{S}, \mathcal{A}, r_\psi, p_\psi, \gamma, p_1)$. The tasks share the components $(\mathcal{S}, \mathcal{A}, \gamma, p_1)$, but can differ in the reward function r_ψ (e.g. navigating to different goal locations) and/or the transition dynamics p_ψ (e.g. locomotion on different terrains). In this work, we consider the setting where the tasks share the same transition dynamics (i.e., $p_\psi = p$), but differ in the reward function. The goal in meta-learning is to learn a set of meta-parameters such that given a new task from $p(\psi)$ and small amount of data for the new task, the meta-parameters can be efficiently adapted to solve the new task. In the context of meta-RL, given new task ψ , the agent collects some initial trajectories $\{\tau_{\text{pre}}\}$, each being a sequence $\{s_1, a_1, s_2, a_2, \dots\}$, and then undergoes some adaptation procedure $f_\phi(\pi_\theta, \tau_{\text{pre}}, r_\psi)$ (e.g., a gradient update ([Finn et al., 2017](#)) or a forward pass through an RNN ([Duan et al., 2016](#))). The adaptation procedure returns a new policy π' . Using this post-adaptation policy, the agent should seek to maximize the cumulative discounted rewards it achieves. Overall, the meta-RL objective is:

$$\max_{\theta, \phi} \mathbb{E}_{\psi \sim p(\psi), (s_t, a_t) \sim \pi'(\theta, \phi)} \left[\sum_{t=1}^{\infty} \gamma^{t-1} r_\psi(s_t, a_t) \right]; \quad \pi'(\theta, \phi) = f_\phi(\pi_\theta, \tau_{\text{pre}}, r_\psi) \quad (1)$$

where θ, ϕ are the meta-parameters that are learned in the meta-training phase. A meta-RL agent must learn a good adaptation procedure f_ϕ that is proficient in extracting salient information about the task at hand, using few pre-adaptation trajectories τ_{pre} . At the same time, it should learn the policy meta-parameters θ such that it can achieve high returns after the adaptation process, i.e., while following the policy $\pi' = f_\phi(\pi_\theta, \tau_{\text{pre}}, r_\psi)$.

3.3 PEARL

In this work, we use PEARL ([Rakelly et al., 2019](#)) as our base meta-RL algorithm since it uses off-policy RL and provides structured exploration via posterior sampling. PEARL is built on top of Soft Actor-Critic ([Haamoja et al., 2018](#)) and trains an encoder network $q_\phi(z|c)$ that takes in the ‘‘context’’ c , which consists of a batch of (s_t, a_t, r_t, s_{t+1}) transitions, and produces the latent embedding z . The intent is to learn the encoder such that embedding z encodes some salient information about the task. The adaptation step f_ϕ in PEARL corresponds to generating this latent z and then conditioning the policy and the value function networks on it. The policy $\pi_\theta(a|s, z)$ is trained using loss L_{actor} :

$$L_{\text{actor}} = \mathbb{E}_{s \sim B, a \sim \pi_\theta, z \sim q_\phi(z|c)} \left[D_{\text{KL}} \left(\pi_\theta(a|s, z) \parallel \frac{\exp(Q_\theta(s, a, z))}{Z_\theta(s)} \right) \right] \quad (2)$$

where B is the replay buffer. The critic $Q_\theta(s, a, z)$ and the encoder $q_\phi(z|c)$ are trained with temporal difference learning:

$$L_{\text{critic}} = \mathbb{E}_{(s,a,r,s') \sim B, z \sim q_\phi(z|c)} \left[\left(Q_\theta(s, a, z) - (r + \bar{V}(s', \bar{z})) \right)^2 \right] \quad (3)$$

where \bar{V} is the target state value and \bar{z} denotes that the gradient does not flow back through the latent.

4 HINDSIGHT FORESIGHT RELABELING

The objective in this section is to derive a formalism for data-sharing amongst the tasks during the meta-training phase. This is achieved via trajectory-relabeling, wherein a trajectory collected for a training task ψ^i is reused or re-purposed for training a different task ψ^j . Reward-based trajectory-relabeling has received a lot of attention in recent works on multi-task RL and goal-conditioned RL (Andrychowicz et al., 2017; Eysenbach et al., 2020; Li et al., 2020). The intuition is that if a trajectory τ collected while solving for the task ψ^i achieves high returns under the reward definition for another task ψ^j (i.e., $\sum_t r_{\psi^j}(s_t, a_t)$ is large), then τ can be readily used for policy-optimization for the task ψ^j as well. The meta-RL setting presents the following subtlety – for any given task, the meta-RL agent generates trajectories with the aim of utilizing them in the adaptation procedure and subsequently seeks to maximize the post-adaptation returns (cf. §3.2). To improve the efficiency of the meta-training stage, we would like to share these pre-adaptation trajectories amongst the different tasks, accounting for the fact that the metric of interest with these trajectories is their usefulness for task-identification, rather than the returns (as in multi-task RL). This difference is illustrated in Figure 3. Hence, when deciding if a trajectory τ collected for task ψ^i is appropriate to be reused for task ψ^j , it is sub-optimal to consider the return value of this trajectory under ψ^i . Instead, we argue that this reuse compatibility should be determined based on the performance on the task ψ^j , after the agent has undergone adaptation using τ . Concretely, we define a function to measure the utility of the trajectory τ for a task ψ^j :

$$U_{\psi^j}(\tau) = \mathbb{E}_{s_t, a_t \sim \pi'} \left[\sum_{t=1}^{\infty} \gamma^{t-1} r_{\psi^j}(s_t, a_t) \right] \quad (4)$$

where $\pi' = f_\phi(\pi_\theta, \tau, r_{\psi^j})$ denotes the policy after using τ for adaptation. The trajectory-relabeling mechanism during meta-training now incorporates this function U_{ψ^j} , which we refer to as the *utility function*, rather than the return R_{ψ^j} . Broadly, a trajectory τ collected for task ψ^i can be relabeled for use in another task ψ^j if $U_{\psi^j}(\tau)$ is high. Subsection §4.1 makes this more precise by deriving a relabeling distribution $q(\psi|\tau)$ that informs us of the tasks for which τ should be reused. Figure 4, and the caption therein, describe a high-level overview of our approach, HFR.

Comparison to HIPI (Eysenbach et al., 2020) with a didactic example. We consider a toy environment to further motivate that return-value based data sharing and trajectory relabeling (as proposed by HIPI) is potentially sub-optimal for meta-RL. The Four-Corners environment consists of a point robot placed at the center of a square where each corner of the square represents a goal location, as shown in Figure 1. For each goal (task), there is a section of the space in the corresponding quadrant in which the robot receives a large negative reward. Consider a trajectory τ that hovers over the blue square in top-right quadrant. Note that τ could have been generated by the agent while collecting data for any of the four tasks. We examine if τ can be reused for the blue task. Since $R_{\text{blue}}(\tau)$ is highly negative, the relabeling strategy in HIPI does not reuse τ for meta-training on the blue task. It is clearly evident, however, that τ carries a significant amount of signal pertaining to task-identification on the blue task, making it a useful pre-adaptation trajectory. HFR reuses τ for the blue task since the utility $U_{\text{blue}}(\tau)$ is high.

To quantify this effect, we include the numerical data on the returns and the utility values for a sampled trajectory that hovers over the blue square in top-right quadrant. The values for the returns $\{R_{\text{purple}}(\tau), R_{\text{blue}}(\tau), R_{\text{orange}}(\tau), R_{\text{green}}(\tau)\}$ are $\{-20, -58, -20, -20\}$, while the utility values $\{U_{\text{purple}}(\tau), U_{\text{blue}}(\tau), U_{\text{orange}}(\tau), U_{\text{green}}(\tau)\}$ are $\{-1015, -756, -935, -931\}$. These (unnormalized) numbers show that the probability of relabeling this trajectory with the blue task is low under HIPI, but high under HFR. Further analysis is included in Appendix A.6. Figure 2 compares HFR and HIPI in terms of the success-rate in the Four-Corners environment, and shows the performance benefit of using the utility function for trajectory relabeling.




Figure 1: The Four-Corners environment




Figure 2: Success-rate on Four-Corners




Figure 3: An illustration of the differences between multi-task RL and meta-RL. In multi-task RL (blue) the agent simply maximizes its returns given a task ψ , while in meta-RL (orange) the agent must first quickly identify the task with a limited number of exploratory trajectories (first two orange stacks in the figure), before adapting to the task and maximizing returns. Because of these differences, existing multi-task relabeling methods may be sub-optimal for meta-RL.

4.1 DERIVING A META-RL RELABELING DISTRIBUTION

Our derivation in this subsection largely follows HIPI (Eysenbach et al., 2020), but differs in that we adapt it to the meta-RL setting to promote sharing of pre-adaptation trajectories amongst tasks, using the concept of trajectory utility. Assume a dataset \mathcal{D} of trajectories gathered by the meta-RL agent when solving the different tasks in the meta-train task distribution. We wish to learn a trajectory relabeling distribution $q(\psi|\tau)$ such that, given any trajectory $\tau \sim \mathcal{D}$, we could reuse τ for tasks with high density under this posterior distribution. To that end, we start by defining a variational distribution $q(\tau|\psi)$ to designate the trajectories used for the adaptation process f_ϕ , for a given task ψ . Using the definition of the utility function (Eq. 4), the meta-RL objective from Eq. 1 could be written as: $\max_{\theta, \phi} \mathbb{E}_{\psi \sim p(\psi)} \mathbb{E}_{\tau \sim q(\tau|\psi)} [U_\psi(\tau)]$. For fixed meta-parameters (θ, ϕ) , a natural approach to optimize the variational distribution $q(\tau|\psi)$ is to use this same objective since it facilitates alignment with the goals of the meta-learner. Thus, the combined objective for the variational distributions for *all* the tasks, augmented with entropy regularization, is:

$$\max_q \mathbb{E}_{\psi \sim p(\psi)} \left[\mathbb{E}_{\tau \sim q(\tau|\psi)} [U_\psi(\tau)] + \mathcal{H}_{q(\tau|\psi)} \right] \quad (5)$$

where $\mathcal{H}_{q(\tau|\psi)}$ denotes the causal entropy of the policy associated with $q(\tau|\psi)$. Now, it is easy to show that the above optimization is equivalent to a reverse-KL divergence minimization objective: $\min_{q(\tau, \psi)} D_{\text{KL}} [q(\tau, \psi) || p(\tau, \psi)]$. In this objective, the joint distributions over the tasks and the trajectories are defined as $q(\tau, \psi) = q(\tau|\psi)p(\psi)$ and $p(\tau, \psi) = p(\tau|\psi)p(\psi)$, where

$$p(\tau|\psi) \triangleq \frac{1}{Z(\psi)} p_1(s_1) e^{U_\psi(\tau)} \prod_{t=1}^T p(s_{t+1}|s_t, a_t) \quad (6)$$

Our goal is to formulate the trajectory relabeling distribution $q(\psi|\tau)$. To make this explicit in our objective, we use the trick proposed in HIPI (Eysenbach et al., 2020) and factor $q(\tau, \psi)$ as $q(\psi|\tau)q(\tau)$, thereby rewriting the reverse-KL divergence minimization objective as:

$$\min_{q(\tau, \psi)} \mathbb{E}_{\substack{\tau \sim q(\tau) \\ \psi \sim q(\psi|\tau)}} \left[\log q(\psi|\tau) + \log q(\tau) - \log p(\psi) + \log Z(\psi) - U_\psi(\tau) - \log p_1(s_1) - \sum_t \log p(s_{t+1}|s_t, a_t) \right] \quad (7)$$

Ignoring the terms independent of ψ , we can analytically solve (by differentiating and setting to zero) for the *optimal* trajectory relabeling distribution for the meta-RL setting:

$$q(\psi|\tau) \propto p(\psi) e^{U_\psi(\tau) - \log Z(\psi)} \quad (8)$$

Given a trajectory $\tau \sim \mathcal{D}$, the HFR algorithm uses this relabeling distribution to sample tasks for which τ should be reused. Concretely, we compute the utility function under τ for all the tasks, construct the distribution $q(\psi|\tau)$ using these utilities (Eq. 8), and sample tasks from it. Please see Figure 4 for details. We assume a uniform prior $p(\psi)$ over the tasks in our experiments.

4.2 ALGORITHM AND IMPLEMENTATION DETAILS

Our relabeling algorithm is summarized in Algorithm 1 and fits seamlessly into the meta-training process of any of the base meta-RL algorithms. Once the meta-RL agent generates a trajectory τ for a training task, τ is fed as input to HFR, and it returns another task that can reuse this experience τ .

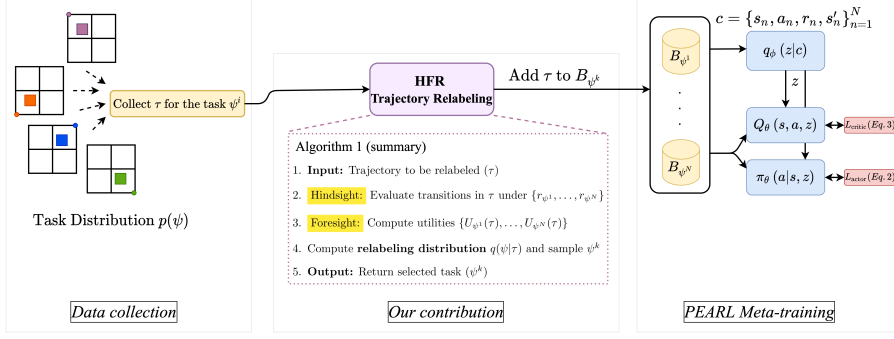


Figure 4: During meta-training, after a trajectory τ is collected for task ψ^i , HFR uses *hindsight* to relabel this trajectory using reward functions for different tasks, and then uses *foresight* to compute the utility of the relabeled trajectory for the different tasks. A distribution over tasks is constructed using the utilities, and a task ψ^k is sampled from the distribution, with tasks for which the trajectory has higher (normalized) utility having higher probability mass. The trajectory is then relabeled using the reward function r_{ψ^k} and added to the task-specific replay buffer B_{ψ^k} . Finally, the meta-training update rules are applied. This process repeats throughout the entirety of meta-training. HFR uses PEARL as the base meta-RL algorithm and does not alter its data collection or meta-gradient computation rules. Please see the Algorithm 1 box for details.

Algorithm 1: Hindsight Foresight Relabeling (HFR)

Input : Trajectory to be relabeled (τ)

Output : Task to relabel the trajectory with (ψ)

for each training task ψ^i do

$U_{\psi^i}(\tau) \leftarrow \text{ComputeUtility}(\tau, \psi^i)$
 $\log Z(\psi^i) \leftarrow \text{GetLogPartition}(\psi^i)$

end

Return $\psi \sim \text{softmax}\{U_{\psi^i}(\tau) - \log Z(\psi^i)\}$ (Eq. 8)

Function $\text{GetLogPartition}(\psi)$:

Sample batch of trajectories $\{\tau^i\}_{i=1}^N \sim B_\psi$

for each trajectory τ^i do

$U_\psi(\tau^i) \leftarrow \text{ComputeUtility}(\tau^i, \psi)$

end

Return $\log Z(\psi) \approx \log\left(\frac{1}{N} \sum_{i=1}^N e^{U_\psi(\tau^i)}\right)$

Function $\text{ComputeUtility}(\tau, \psi)$:

for each $(s_t, a_t, r_t) \in \tau$ do

Replace r_t with $r_\psi(s_t, a_t)$

end

Sample embedding using encoder $z \sim q_\phi(z|\tau)$

Sample a batch of initial states $\{s_1^i\}_{i=1}^N \sim B_\psi$

Sample actions for these states using the post-adaptation policy $\pi_\theta(\cdot|s, z)$:

$\{a_1 \sim \pi_\theta(a_1|s_1^i, z)\}_{i=1}^N$

Return $U_\psi = \frac{1}{N} \sum_{i=1}^N Q_\theta(s_1^i, a_1^i, z)$ (Eq. 9)

We compute the utility of the input trajectory for every training task, along with an empirical estimate of the log-partition function of the tasks. The task to relabel the trajectory with is then sampled from a categorical distribution. For our experiments, we build on top of the PEARL algorithm (Rakelly et al., 2019), which is a data-efficient off-policy meta-RL method. PEARL maintains task-specific replay buffers B_ψ . If HFR returns the task ψ^k , then τ is relabeled using the reward function r_{ψ^k} and added to B_{ψ^k} for meta-training on the task ψ^k .

The adaptation procedure $\pi' = f_\phi(\pi_\theta, \tau, r_\psi)$ for a task ψ corresponds to a sequence of steps: 1.) augment τ by marking each transition with a reward value computed using $r_\psi(s_t, a_t)$, 2.) condition the encoder on τ to sample an embedding, $z \sim q_\phi(z|\tau)$; and 3.) condition the policy on z to obtain the post-adaptation policy, $\pi' = \pi_\theta(\cdot|s, z)$. The calculation of the utility function (Eq. 4) requires generation of post-adaptation trajectories, which could be computationally inefficient, especially if the number of tasks is large. To avoid this cost, for each task, we sample a batch of initial states $s_1 \sim p_1(s_1)$ and the corresponding actions from the post-adaptation policy, and compute the utility based on an estimate of the state-action value function $Q_\psi^{\pi'}(s_1, a_1)$ as:

$$U_\psi(\tau) = \mathbb{E}_{s_1 \sim p_1, a_1 \sim \pi'(\cdot|s_1)} \left[Q_\psi^{\pi'}(s_1, a_1) \right] \quad (9)$$

Since we use PEARL, we can avoid training separate task-specific value functions Q_ψ , and instead get the required estimates from the task-conditioned critic $Q(s, a, z)$ already used by PEARL (Eq. 3). We




Figure 5: MuJoCo environments we evaluate on – sparse reward manipulation tasks (a, b, c), as well as sparse and dense reward locomotion tasks (c, d).




Figure 6: Performance of our relabeling algorithm HFR (shown in blue) on sparse reward tasks. HFR consistently outperforms baselines on both sparse reward robotic manipulation and locomotion tasks. Visual-Reacher uses image observations, while the other environments use proprioceptive states.

highlight that HFR facilitates efficient data-sharing among the training tasks via trajectory-relabeling *without* altering the meta-train and test-time adaptation rules of the base meta-RL algorithm.

5 EXPERIMENTS

The goal in this section is to quantitatively evaluate the benefit of sharing experience among tasks using HFR, during the meta-train stage. We evaluate on a set of both sparse and dense reward MuJoCo environments (Todorov et al., 2012) modeled in OpenAI Gym (Brockman et al., 2016). Please refer to the Appendix for environment details. We compare HFR with two relabeling methods: **Random**, in which each trajectory is relabeled with a randomly chosen task, and **HPI** (Eysenbach et al., 2020), which utilizes MaxEnt RL to devise an algorithm that relabels each transition using a distribution over tasks involving the soft Q values for that transition. In contrast, HFR proposes the concept of the utility of a trajectory for a given task. Using the utility aligns the relabeling methodology with the objective of the meta-RL agent (*cf.* section 4). All methods are built on top of PEARL. Finally, we compare to PEARL with no relabeling at all, which we refer to as **None**.

5.1 RESULTS

Sparse Reward Environments. We first evaluate HFR on a set of sparse reward robotic manipulation and locomotion tasks. We use five environments: a goal-reaching task involving a quadruped Ant robot, a pushing task on the Sawyer robot, a reaching task on the Sawyer robot, a velocity-matching task involving a bipedal Cheetah robot, and a reaching task involving the MuJoCo Reacher where the agent learns directly from images. The environments are described in detail in Appendix A.4 and shown pictorially in Figure 5. Figure 6 plots the performance (average returns or success-rate) on the held-out meta-test tasks on the y -axis, with the total timesteps of environment interaction for meta-training on the x -axis. We note that HFR tends to be more sample-efficient than the baselines and achieves a higher asymptotic score. Meta-RL in sparse reward tasks is hard due to the challenges of task-identification and efficient exploration. HFR is especially useful for these tasks as the data-sharing afforded by the trajectory relabeling algorithm mitigates the need for an elaborate exploration strategy during meta-training. This leads to the sample-efficiency gains exhibited in Figure 6.

Dense Reward Environments. We next evaluate HFR on dense reward environments (Figure 7). We experiment with the Cheetah-Highdim environment, in which the bipedal robot is required to




Figure 7: Performance of our relabeling algorithm HFR (shown in blue) on dense reward tasks. With the exception of Cheetah-Highdim, relabeling in general offers no benefit in dense reward meta-RL tasks, likely due to the highly informative nature of a dense reward function.

best match its state vector to a set of predefined vectors, as well as the Cheetah-Vel and quadruped Ant-Vel environments, in which the robots are required to run at various velocities.

Note that the impact of HFR is much less pronounced for these environments, with the exception of Cheetah-Highdim. We believe this is because exploration is not as critical for these environments as it was for the sparse reward tasks. This hypothesis is supported by the fact that, in these environments, PEARL with no relabeling is competitive with the various relabeling methods, which all share similar performance, whereas in the sparse reward environments HFR is the relabeling method that performs best, with the two other relabeling methods also vastly outperforming vanilla PEARL. In the case of the Cheetah-Vel and Ant-Vel environments, agents are provided with an informative dense reward that immediately informs them as to which task they’re supposed to be solving. Although the agent in Cheetah-Highdim is also provided with an informative dense reward, the reward function in this task is a linear combination of an 18-dimensional state vector. Thus, reasonably good exploration is needed to determine optimal values for each of these 18 dimensions. HFR relabeling provides improvement over the baselines for this task.

5.2 ABLATION STUDIES

Batch Size. We investigate the impact of the batch size N used in Algorithm 1 for computing an empirical estimate of the state-action value: $\frac{1}{N} \sum_{i=1}^N Q_{\theta}(s_1^i, a_1^i, z)$. We compare the effect of batch size across batch sizes $\{16, 32, 64, 128, 256\}$. We expect lower values of N to lead to a higher variance estimate of the post-adaptation cumulative discounted rewards and thus potentially a worse approximation of the optimal meta-RL relabeling distribution. In Figure 8a, we see some evidence of this in the comparatively worse performance when using $N = 16$ and $N = 32$. However, we note that even with these small batch sizes, our method still performs well, and in general achieves high returns across all choices of N .

Partition Function. We investigate the impact of the log-partition function $\log Z(\psi)$ in the optimal meta-RL relabeling distribution (Eq. 8). Prior work has noted the importance of the partition function in the multi-task RL setting when tasks may have different reward scales (Eysenbach et al., 2020). We believe that in the meta-RL setting, the partition function may be crucial even if the tasks share the reward scale, since some tasks in the meta-training distribution may be easier to solve than others. We speculate that with the omission of the partition function from our relabeling distribution, trajectories would find high utility and thus be disproportionately labeled with the easily-solved tasks, causing a degradation in overall performance. In our experiments, we find that the partition function serves as an essential normalization factor; Figure 8b shows an example.

Reward Function. One assumption our method assumes is access to the true reward function $r_{\psi}(s, a)$, which we can query to get the reward for an individual transition under any task ψ . This availability has also been utilized in existing works on relabeling for multi-task RL (Eysenbach et al., 2020; Li et al., 2020). In many real-world applications of meta-RL, *e.g.* a distribution over robotic tasks, it is reasonable to assume that the task-designer outlines a rough template for the rewards corresponding to the different tasks in the distribution. Furthermore, several of the rewards used in our experiments are success/failure indicators, which are simple to specify. Nevertheless, we consider the scenario where we cannot query the true reward function for individual transitions. Figures 8c and 8d show good performance even when we use a learned reward function rather than the true reward function to relabel trajectories.

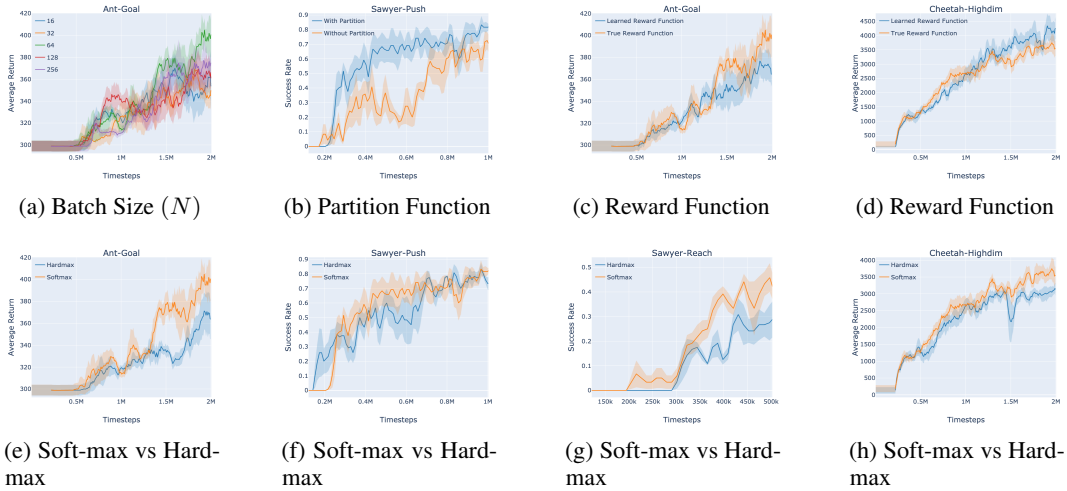


Figure 8: Ablation analyses. (a) HFR with different values for the batch size N used in approximating Eq. 9. HFR is relatively robust to choice of batch size, with some smaller choices of N leading to slightly worse performance. (b) HFR with and without the log-partition function ($\log Z(\psi)$). The partition function serves as a necessary normalization factor and prevents against simply relabeling every trajectory using the easiest task. (c), (d) HFR with true reward function (orange) and HFR with learned reward function (blue). (e), (f), (g), (h) HFR with soft-max (orange) vs hard-max (blue) relabeling distribution

Soft-max vs Hard-max Relabeling Distribution. Given a trajectory τ , the relabeling distribution derived in Section §4.1 samples tasks for which τ should be reused. Specifically, tasks are sampled as: $\psi \sim \text{softmax}\{U_{\psi^i}(\tau) - \log Z(\psi^i)\}$. This raises the following question: is it crucial to have stochasticity in the relabeling distribution, or could we deterministically select the task for which the normalized utility value is the highest, *i.e.*, $\psi = \text{argmax}\{U_{\psi^i}(\tau) - \log Z(\psi^i)\}$?

A minor modification to the equations in Section §4.1 yields the hard-max relabeling distribution. Concretely, we can add a term to the starting objective for the variational distribution q that explicitly *minimizes* the entropy of the relabeling distribution $q(\psi|\tau)$:

$$\max_q \mathbb{E}_{\psi \sim p(\psi)} \left[\mathbb{E}_{\tau \sim q(\tau|\psi)} [U_{\psi}(\tau)] + \mathcal{H}_{q(\tau|\psi)} \right] - (1 - \epsilon) \mathbb{E}_{\tau \sim q(\tau)} [\mathcal{H}_{q(\psi|\tau)}]$$

where ϵ is a value less than 1. Proceeding with the derivation in the exact same manner as in Section §4.1, we obtain the adjusted relabeling distribution:

$$q_{\epsilon}(\psi|\tau) \propto e^{\frac{U_{\psi}(\tau) - \log Z(\psi)}{\epsilon}}$$

In the limit when $\epsilon \rightarrow 0$, $q_{\epsilon}(\psi|\tau)$ is the hard-max relabeling distribution. Figures 8e, 8f, 8g, and 8h compare the performance of HFR when sampling tasks from a soft-max relabeling distribution (orange) vs a hard-max distribution (blue). The results indicate that stochasticity is an important factor.

6 CONCLUSION

In this paper, we introduced HFR, a trajectory relabeling method for meta-RL that enables data sharing between tasks during meta-train. We argue that unlike the multi-task RL setting, where the appropriateness of a trajectory for a task could be measured by the returns under the task, for meta-RL, it is preferable to consider the future (expected) task-returns of an agent adapted using that trajectory. We capture this notion by defining the *utility* function for a trajectory-task pair and incorporate these utilities in our relabeling mechanism. Inspired by prior work on multi-task RL, an optimal relabeling distribution is then derived that informs us of the tasks for which a generated trajectory should be reused. *Hindsight* is used to relabel trajectories with different reward functions, while *foresight* is used in computing the utility of each trajectory under different tasks and constructing a relabeling distribution. HFR is easy to implement, can be integrated into any existing meta-RL algorithm, and yields improvement on a variety of meta-RL tasks, especially those with sparse rewards. To the best of our knowledge, HFR is the first relabeling method designed explicitly for the meta-RL paradigm.

REFERENCES

- 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*, pp. 5048–5058, 2017.
- Jonathan Baxter. Theoretical models of learning to learn. In *Learning to learn*, pp. 71–94. Springer, 1998.
- Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym. *arXiv preprint arXiv:1606.01540*, 2016.
- Rich Caruana. Multitask learning. *Machine learning*, 28(1):41–75, 1997.
- Ron Dorfman, Idan Shenfeld, and Aviv Tamar. Offline meta learning of exploration. *arXiv preprint arXiv:2008.02598*, 2020.
- Yan Duan, John Schulman, Xi Chen, Peter L Bartlett, Ilya Sutskever, and Pieter Abbeel. $R1^2$: Fast reinforcement learning via slow reinforcement learning. *arXiv preprint arXiv:1611.02779*, 2016.
- Benjamin Eysenbach, Xinyang Geng, Sergey Levine, and Ruslan Salakhutdinov. Rewriting history with inverse rl: Hindsight inference for policy improvement. *arXiv preprint arXiv:2002.11089*, 2020.
- Rasool Fakoor, Pratik Chaudhari, Stefano Soatto, and Alexander J Smola. Meta-q-learning. *arXiv preprint arXiv:1910.00125*, 2019.
- Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. *arXiv preprint arXiv:1703.03400*, 2017.
- Scott Fujimoto, Herke Van Hoof, and David Meger. Addressing function approximation error in actor-critic methods. *arXiv preprint arXiv:1802.09477*, 2018.
- Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. *arXiv preprint arXiv:1801.01290*, 2018.
- Leslie Pack Kaelbling. Learning to achieve goals. In *IJCAI*, pp. 1094–1099. Citeseer, 1993.
- Alexander C Li, Lerrel Pinto, and Pieter Abbeel. Generalized hindsight for reinforcement learning. *arXiv preprint arXiv:2002.11708*, 2020.
- Evan Zheran Liu, Aditi Raghunathan, Percy Liang, and Chelsea Finn. Decoupling exploration and exploitation for meta-reinforcement learning without sacrifices. *arXiv preprint arXiv:2008.02790*, 2020.
- Russell Mendonca, Xinyang Geng, Chelsea Finn, and Sergey Levine. Meta-reinforcement learning robust to distributional shift via model identification and experience relabeling. *arXiv preprint arXiv:2006.07178*, 2020.
- Thomas Miconi, Kenneth Stanley, and Jeff Clune. Differentiable plasticity: training plastic neural networks with backpropagation. In *International Conference on Machine Learning*, pp. 3559–3568. PMLR, 2018.
- Nikhil Mishra, Mostafa Rohaninejad, Xi Chen, and Pieter Abbeel. A simple neural attentive meta-learner. *arXiv preprint arXiv:1707.03141*, 2017.
- DK Naik, Richard Mammone, and A Agarwal. Meta-neural network approach to learning by learning. In *Proceedings of the 1992 Artificial Neural Networks in Engineering, ANNIE'92*, 1992.
- Alex Nichol and John Schulman. Reptile: a scalable metalearning algorithm. *arXiv preprint arXiv:1803.02999*, 2(2):1, 2018.
- Kate Rakelly, Aurick Zhou, Chelsea Finn, Sergey Levine, and Deirdre Quillen. Efficient off-policy meta-reinforcement learning via probabilistic context variables. In *International conference on machine learning*, pp. 5331–5340, 2019.

- Jonas Rothfuss, Dennis Lee, Ignasi Clavera, Tamim Asfour, and Pieter Abbeel. Promp: Proximal meta-policy search. *arXiv preprint arXiv:1810.06784*, 2018.
- Tom Schaul, Daniel Horgan, Karol Gregor, and David Silver. Universal value function approximators. In *International conference on machine learning*, pp. 1312–1320. PMLR, 2015.
- Jurgen Schmidhuber. Evolutionary principles in self-referential learning. *On learning how to learn: The meta-meta-... hook.*) *Diploma thesis, Institut f. Informatik, Tech. Univ. Munich*, 1(2), 1987.
- Bradly C Stadie, Ge Yang, Rein Houthoofd, Xi Chen, Yan Duan, Yuhuai Wu, Pieter Abbeel, and Ilya Sutskever. Some considerations on learning to explore via meta-reinforcement learning. *arXiv preprint arXiv:1803.01118*, 2018.
- Sebastian Thrun and Lorien Pratt. Learning to learn: Introduction and overview. In *Learning to learn*, pp. 3–17. Springer, 1998.
- Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 5026–5033. IEEE, 2012.
- 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.
- Zhongwen Xu, Hado van Hasselt, and David Silver. Meta-gradient reinforcement learning. *arXiv preprint arXiv:1805.09801*, 2018.
- Zeyu Zheng, Junhyuk Oh, Matteo Hessel, Zhongwen Xu, Manuel Kroiss, Hado Van Hasselt, David Silver, and Satinder Singh. What can learned intrinsic rewards capture? In *International Conference on Machine Learning*, pp. 11436–11446. PMLR, 2020.
- Luisa Zintgraf, Kyriacos Shiarlis, Maximilian Igl, Sebastian Schulze, Yarin Gal, Katja Hofmann, and Shimon Whiteson. Varibad: A very good method for bayes-adaptive deep rl via meta-learning. *arXiv preprint arXiv:1910.08348*, 2019.