

HMRL: Hyper-Meta Learning for Sparse Reward Reinforcement Learning Problem

Yun Hua

East China Normal University
Shanghai, China
52194501002@stu.ecnu.edu.com

Xiangfeng Wang*

East China Normal University and
SRIAS
Shanghai, China
xfwang@sei.ecnu.edu.cn

Bo Jin*

East China Normal University and
SRIAS
Shanghai, China
bjin@cs.ecnu.edu.cn

Wenhao Li

East China Normal University
Shanghai, China
52194501026@stu.ecnu.edu.cn

Junchi Yan

Shanghai Jiao Tong University
Shanghai, China
yanjunchi@sjtu.edu.cn

Xiaofeng He

East China Normal University and
SRIAS
Shanghai, China
xfhe@sei.ecnu.edu.cn

Hongyuan Zha

School of Data Science and AIRS, The
Chinese University of Hong Kong
Shenzhen, China

ABSTRACT

In spite of the success of existing meta reinforcement learning methods, they still have difficulty in learning a meta policy effectively for RL problems with sparse reward. In this respect, we develop a novel meta reinforcement learning framework called Hyper-Meta RL (HMRL), for sparse reward RL problems. It is consisted with three modules including the cross-environment meta state embedding module which constructs a common meta state space to adapt to different environments; the meta state based environment-specific meta reward shaping which effectively extends the original sparse reward trajectory by cross-environmental knowledge complementarity and as a consequence the meta policy achieves better generalization and efficiency with the shaped meta reward. Experiments with sparse-reward environments show the superiority of HMRL on both transferability and policy learning efficiency.

CCS CONCEPTS

• **Computing methodologies** → **Reinforcement learning.**

KEYWORDS

Meta Learning; Reinforcement Learning; Sparse Reward

*Corresponding authors

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
KDD '21, August 14–18, 2021, Virtual Event, Singapore.

© 2021 Association for Computing Machinery.
ACM ISBN 978-1-4503-8332-5/21/08...\$15.00
<https://doi.org/10.1145/3447548.3467242>.

ACM Reference Format:

Yun Hua, Xiangfeng Wang, Bo Jin, Wenhao Li, Junchi Yan, Xiaofeng He, and Hongyuan Zha. 2021. HMRL: Hyper-Meta Learning for Sparse Reward Reinforcement Learning Problem. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '21), August 14–18, 2021, Virtual Event, Singapore*. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3447548.3467242>.

1 INTRODUCTION

Deep reinforcement learning (RL) have shown great success in games [18, 25, 28], and also practical areas like robotics [5] and traffic light management [30]. To achieve generalization and efficiency on multiple tasks, meta reinforcement learning (meta-RL), which combines RL and meta learning, has recently been studied by extracting meta knowledge to help efficient policy learning on new tasks. The works [6, 29] model task embedding as meta knowledge, to extract task features. Direct attention is applied to policy learning in [7], which is extended in [11, 12, 14, 32] by using basic components in RL as meta knowledge, e.g. hyper-parameters, return function, advantage function and etc.

These methods are still algorithmically limited, especially with sparse reward. Some meta-RL methods focus on the reward sequence: PEMIRL [34] introduces meta inverse RL to directly learn the reward signals. NoRML [33] learns meta advantage values as additional rewards. [37] employs meta-RL to design reward shaping function, to speed up training on new tasks. The sparse reward problem has been partially addressed by such methods, while it is still far from resolved. Existing meta-RL models mostly deal with tasks within one environment. The setting for transfer across different environments and the relevant techniques have not been well studied in literature, especially with sparse rewards. In particular, we make two observations: 1) The rewards of tasks in the same environment tend to be homogeneous which may weaken the capacity for effective cross-task reward utilization; 2) While

cross-environment meta knowledge can be complementary to enrich the reward over the tasks across environments, thus mitigates the sparse reward issue.

Motivated by these observations, this work aims to devise a cross-environment meta learning approach with sparse reward. We propose a hyper meta RL framework, called *Hyper-Meta RL*, which consists of three meta modules. We devise the cross-environment meta state embedding module and environment-specific meta reward shaping module to support the general meta policy learning module. The meta state is embedded across environments. Based on the general meta-state, the meta reward shaping module is dedicated to generate targeted reward shaping function. Further, the meta state and additional reward signals can improve the meta policy model to reduce the negative influence of missing reward signal. Moreover, the shared meta state space benefits meta reward shaping and provides a flexible mechanism to integrate domain knowledge from multiple environments. The Hyper-Meta RL framework is illustrated in Figure 1. The main contributions are:

- 1) For the sparse reward setting, a hyper-meta RL approach is proposed, to jointly utilize the information of tasks from multiple environments. This is in contrast to existing works for a single environment. As a natural way to improve generalization for meta policy learning, this proposes a novel meta-RL method on utilizing cross-environment information for sparse reward RL problems;

- 2) For each environment, an environment-specific meta reward shaping technique is developed. The shaping is fed with the meta state learned from domain knowledge in both the current and other environments. It can speed up new task learning either with or without meta policy, and meanwhile the policy consistency can be theoretically ensured;

- 3) We show that our methods can effectively improve both generalization capability and transfer efficiency, especially under environments with notable heterogeneity, as shown in our experiments.

Finally, we point out that our approach in fact can be interpreted as a hierarchical structure for task learning. The environment acts as the group of tasks, which can either be formed by prior knowledge or other means. In contrast, existing methods treat all tasks as in a single environment. Hence the improvement can be attributed to both our hyper-meta methods as well as the prior of the task set structure. At least, our method allows a direct and effective way to incorporate environment priors, which are often readily available or naturally arise. Furthermore, our method don't need too much computing resources because the only additive parts which need external computing resources are environment-specific meta reward shaping and cross-environment meta state embedding and they only need some simple neural networks.

2 RELATED WORK

Sparse Reward Reinforcement Learning. As the reward signal is the only supervised information in reinforcement learning, sparse reward problems will cause many negative influence on the training of reinforcement learning. Sparse reward reinforcement learning methods contain these classes: 1) Reward shaping or reward design [13, 20, 21] methods; 2) Curriculum learning methods [8, 19, 27, 31]; 3) Bayesian reinforcement learning [10, 22]. The reward shaping or reward design methods directly learns additional reward

to make the sparse reward dense while curriculum learning transfer policy from easier task as a prior and Bayesian reinforcement learning learns prior information about tasks.

Meta Reinforcement Learning. It aims to learn a common structure between different tasks to solve new and relevant tasks efficiently. Meta RL methods can also be divided into two categories: model-based meta RL and model-free meta RL. The model-based methods have the advantage to directly model dynamic environments, specifically by the environment parameters [1, 3, 23] or priors [9].

Model-free meta RL methods can be divided into these classes: 1) recurrence based meta-RL methods, which use recurrent neural network (RNN) to establish the meta task embedding [6, 17, 29]; 2) gradient based meta-RL methods, which are mostly extended from MAML [7] aim to learn a meta policy or some well-defined meta parameters simultaneously. The meta information can be set as the exploration strategy [11, 12, 35], advantage function [33], reward shaping [37], return function [32] and etc; 3) hybrid meta-RL methods [14, 26], considered as the hybrid of the above two techniques.

[11, 12] set the exploration policy as the meta knowledge, while [33] meta-learns the advantage function instead.

These meta-RL methods mostly focus on RL with dense reward, and only a few pay attention to sparse reward case. NoRML [33] meta-learns the advantage function for sparse reward tasks; [37] introduces MAML to enhance the reward shaping technique, to efficiently handle sparse reward; VariBAD [36] uses the Bayesian formulation to learn prior information to solve sparse reward problem in meta RL. Above meta-RL methods mostly address learning of different tasks from the same environment; DREAM [16] build a new exploration structure to solve no reward problem only in adapting but not training time. For a more general meta-RL problem with sparse reward, utilizing limited reward signals from relevant environments to design the meta-RL algorithm can be a **natural but currently under-investigated way** to increase the generalization and transfer abilities.

3 HYPER-META RL WITH SPARSE REWARD

Our key idea is to combine model-based meta RL and model-free meta RL. In our setting, tasks are sampled from different distributions based on environments, and environments are sampled from a fixed distribution. Instead of modeling the environment directly, our approach indirectly utilizes the common knowledge to model the cross-environment meta knowledge, which is denoted as the *Cross-Environment Meta State Embedding*.

As detailed in Figure 1, our Hyper-Meta RL method consists of three modules: 1) cross-environment meta state embedding module; 2) environment-specific meta reward shaping module; 3) meta policy learning module. In the following part of this section, we will first introduce the preliminaries and problem formulation, then we will introduce the three modules. Further we will prove the policy consistency to maintain on new tasks thus the method can be extended to different tasks in a stable fashion.

heuristic knowledge for more effective optimal policy search. It is worth noting that modifying the original reward will eventually change the task and may cause the policy inconsistency. To solve this problem, potential-based reward shaping is proposed in [20]:

$$\mathcal{F}(s, a, s') = \gamma\phi(s') - \phi(s), \quad (2)$$

where s, s' denote two consecutive states and a denotes the action; ϕ is a potential function; the function value \mathcal{F} denotes the reward shaping signal as the difference between the potential of s' and s ; γ denotes a discount rate. The dynamic potential-based reward shaping in [4] extends the original reward shaping to allow dynamic shaping and maintain the policy consistency:

$$\mathcal{F}(s, t, s', t') = \gamma\phi(s', t') - \phi(s, t), \quad (3)$$

where t and t' are two consecutive time stamps.

3.2 Cross-Environment Meta State Embedding

The cross-environment meta state embedding module focuses on embedding both the domain knowledge of the environment and the trajectory of the environment-specific tasks into a unified meta state space. The prior knowledge among different environments should be properly determined to define the shared meta state space \mathcal{S}_m . For example, the meta state can be defined as the agent coordination in both 1st person game and 3rd person game because the agent coordination is in a same space so that it can be shared. The setting of the meta state space is flexible, while the embedding function can further be pertinently learned in the closed loop of the whole algorithm. We can easily use some shared information in different environments' state spaces, and extract them into the shared meta state space. For the Hallway problem, the concatenation of the location information is established as the meta state embedding function h without learning. More generally, the function h can directly embed the original agent state together with some global features of the related task into the common meta state space (or a shared latent space). For the Maze problem, we employ a linear function parameterized by each environment's scenario size (domain knowledge z) as the state embedding function h , which makes a common mini-map as the common meta state.

3.3 Environment-specific Meta Reward Shaping

The environment-specific meta reward shaping module focuses on constructing the shaped reward, i.e., $\hat{\mathcal{R}} = \mathcal{R} + \mathcal{F}$. This module is based on the the potential-based reward shaping, while the cross-environment meta states are set to be the input of the potential function ϕ . The formulation is as:

$$\mathcal{F}(s_m, s'_m) = \gamma\phi(s'_m) - \phi(s_m), \quad (4)$$

where $s_m \in \mathcal{S}_m$ and $s'_m \in \mathcal{S}_m$. The potential function ϕ plays the key role in the above reward shaping, which is meta-learned. The parameter γ denotes the discount factor of the RL problem. The cross-environment meta state contains not only the state information but also the environment-specific task information. The potential value $\phi(s_m)$ will also be environment driven. Eventually, based on the cross-environment meta state embedding $h(s; z)$ introduced in 3.2, the cross-environment meta states based potential function can be further expressed as:

$$\mathcal{F}(s_m, s'_m) = \gamma\phi(h(s'; z)) - \phi(h(s; z)). \quad (5)$$

Algorithm 1 HMRL: Hyper-Meta $\{state, reward, policy\}$ for Reinforcement Learning with Sparse Reward

- 1: **Input:** Task trajectory set: $\tau(\mathcal{M})$; Distribution over tasks: $p(\mathcal{T})$; Learning rates: α, β ;
 - 2: Initialization: $\theta, \{\theta_j^i\}, h, \phi$;
 - 3: **while** not done **do**
 - 4: Sample batch of environments $\{\tilde{\mathcal{E}} \sim p(\mathcal{E})\}$
 - 5: **for** each environments **do**
 - 6: Sample batch of tasks $\{\tilde{\mathcal{T}} \sim p(\mathcal{T}|\tilde{\mathcal{E}})\}$;
 - 7: **for** each $\tilde{\mathcal{T}}$ **do**
 - 8: Sample m extended trajectories $\mathcal{D} := \{\tau'(\tilde{\mathcal{T}})\}$ with meta states h , reward shaping ϕ , policy θ ;
 - 9: Evaluate $\nabla_{\theta} \mathcal{L}_{\tilde{\mathcal{T}}}(\theta)$ using \mathcal{D} and Eq. (6);
 - 10: Update task-specific policy $\theta_{\tilde{\mathcal{T}}}$ by SGD:

$$\theta_{\tilde{\mathcal{T}}} = \theta - \alpha \nabla_{\theta} \mathcal{L}_{\tilde{\mathcal{T}}}(\theta);$$
 - 11: Sample ℓ extended trajectories $\tilde{\mathcal{D}}' := \{\tilde{\tau}'(\tilde{\mathcal{T}})\}$ with meta state embedding h , reward shaping ϕ and task-specific policy $\theta_{\tilde{\mathcal{T}}}$;
 - 12: **end for**
 - 13: Update meta policy $\theta_{\tilde{\mathcal{E}}}$ with Eq. (6) and extended trajectories $\tilde{\mathcal{D}}'$:

$$\theta_{\tilde{\mathcal{E}}} = \theta_{\tilde{\mathcal{E}}} - \beta \sum_{\tilde{\mathcal{T}}} \nabla_{\theta} \mathcal{L}_{\tilde{\mathcal{T}}}(\theta_{\tilde{\mathcal{T}}});$$
 - 14: **end for**
 - 15: Update cross-environment meta state embedding h and environment-specific meta reward shaping ϕ with gradient descent by Eq. (11);
 - 16: **end while**
-

So that if we regard the cross environment meta state embedding $h(\cdot; z)$ as part of the potential function ϕ . Then it is likely to claim that the policy consistency for reward shaping will still holds. And we have proved it in the 3.5. The policy consistency will strength the adaption to new tasks.

3.4 Meta Policy Learning

As shown in Fig. 1, the meta policy is learned based on the extended trajectories of different tasks, while the meta state embedding and meta reward shaping are employed to support meta policy implicitly. We introduce the MAML framework to compute the meta policy iteratively. For the j -th task \mathcal{T}_j^i of environment \mathcal{E}_i , the agent aims to find the task-specific policy $\theta_{i,j}$ which can maximize the accumulated shaped reward $\hat{\mathcal{R}}$, and the loss for task \mathcal{T}_j^i is:

$$\mathcal{L}_{\mathcal{T}_j^i}(\theta_{i,j}) := -\mathbb{E}_{s_t^{i,j}, a_t^{i,j} \sim \theta_{i,j}} \left[\sum_{t=0}^T \gamma^t \hat{\mathcal{R}}(s_t^{i,j}, a_t^{i,j}) \right], \quad (6)$$

$$\hat{\mathcal{R}}(s_t^{i,j}, a_t^{i,j}) = \mathcal{R}(s_t^{i,j}, a_t^{i,j}) + \mathcal{F}(s_m^{i,j,t}, s_m^{i,j,t+1}),$$

where $s_t^{i,j} \in \mathcal{S}_j^i$ and $a_t^{i,j} \in \mathcal{A}_j^i$; $s_m^{i,j,t+1} \in \mathcal{S}_m$ and $s_m^{i,j,t} \in \mathcal{S}_m$ are both meta state vectors corresponding to $s_{t+1}^{i,j}$ and $s_t^{i,j}$ respectively. The meta state based meta reward shaping is introduced to modify original reward in order to increase the training efficiency. Actually

each meta state $s_m^{i,j,t}$ in \mathcal{T}_j^i is determined by the original state $s_t^{i,j}$ and some fixed task features, so that the shaping function value $\mathcal{F}(s_m^{i,j,t}, s_m^{i,j,t+1})$ is fixed with respect to $\theta_{i,j}$. The gradient of $\mathcal{L}_{\mathcal{T}_j^i}$ corresponding to $\theta_{i,j}$ is the same as follows:

$$\tilde{\mathcal{L}}_{\mathcal{T}_j^i}(\theta_{i,j}) := -\mathbb{E}_{s_t^{i,j}, a_t^{i,j} \sim \theta_{i,j}} \left[\sum_{t=0}^T \gamma^t \mathcal{R}(s_t^{i,j}, a_t^{i,j}) \right], \quad (7)$$

which also denotes the original loss function without meta reward shaping. The meta policy θ and task-specific policies $\{\theta_{i,j}\}$ can be iteratively learned by (6), following MAML.

The meta state embedding h and meta reward shaping ϕ will be alternatively updated together with the meta policy. Before discussing details of learning the meta state embedding h and the meta reward shaping potential function ϕ , we can assume that for each meta state $s_m^* \in \mathcal{S}_m$, no less than one original state corresponds to it. Without loss of generality, the state set corresponding to s_m^* is denoted as $\mathcal{S}(s_m^*)$. The potential function ϕ can be expressed as:

$$\phi(s_m^*) = \mathbb{E}_{\{\mathcal{T}_j^i\}} \left[\mathbb{E}_{s \in \mathcal{S}(s_m^*)} \left(V_{\mathcal{T}_j^i}^*(s) \right) \right], \quad (8)$$

where function $V^*(\cdot)$ denotes the classical optimal value function. However, the optimal value function V^* is extremely hard to guarantee, we employ the return function to approximately estimate the optimal value function, i.e.,

$$\phi(s_m^*) = \mathbb{E}_{\{\mathcal{T}_j^i\}} \left[\mathbb{E}_{s \in \mathcal{S}(s_m^*)} \left(\text{Return}_{\mathcal{T}_j^i}(s) \right) \right], \quad (9)$$

where the return function is calculated as follows:

$$\text{Return}(s_t) = \begin{cases} \hat{\mathcal{R}}(s_t) + \gamma \text{Return}(s_{t+1}), & t < T; \\ \hat{\mathcal{R}}(s_t), & t = T. \end{cases}$$

The loss function with respect to the meta state embedding and meta reward shaping can be established based on Eq. 9. For each $s_m \in \mathcal{S}_m$, we define the loss function as follows:

$$\mathcal{L}_{s_m} := \left(\mathbb{E}_{\{\mathcal{T}_j^i\}} \left[\mathbb{E}_{s \in \mathcal{S}(s_m)} \left(\text{Return}_{\mathcal{T}_j^i}(s) \right) \right] - \phi(s_m) \right)^2, \quad (10)$$

Further by fixing the last $\phi(s_m)$ with h and ϕ from former iteration, and considering all significant meta state $s_m \in \mathcal{S}_m$ (depending on specific test problems), we obtain the overall loss for meta state embedding and meta reward shaping:

$$\mathcal{L}(\phi, h) := \sum_{s_m} \mathcal{L}_{s_m}. \quad (11)$$

The meta state embedding function h and meta reward shaping potential function ϕ are both updated by gradient descent via (11). The overall meta-RL scheme is shown in Alg. 1. With the obtained meta policy θ , meta state embedding h and meta reward shaping ϕ , we can do fine-tune on new tasks and the detailed algorithm is given in Alg. 2. To emphasize, new tasks can utilize only the obtained meta state embedding and meta reward shaping modules, without the meta policy. These two meta modules can independently be applied to new tasks with new policy structure, thus their policy structure can be different from that of the meta policy.

Algorithm 2 New Task Fine-tuning with Hyper-Meta RL

- 1: **Input:** New task \mathcal{T}_{new} ; meta state embedding h , reward shaping ϕ and policy θ ;
- 2: Sample extended trajectories \mathcal{D}_{new} with h and ϕ ;
- 3: Evaluate $\nabla_{\theta} \mathcal{L}_{\mathcal{T}_{new}}(\theta)$ using trajectories \mathcal{D}_{new} ;
- 4: Update task-specific policy $\theta_{\mathcal{T}_{new}}$ by gradient descent:

$$\theta_{\mathcal{T}_{new}} = \theta - \alpha \nabla_{\theta} \mathcal{L}_{\mathcal{T}_{new}}(\theta).$$

- 5: **if not fixed then**
 - 6: Update cross-environment meta state embedding h and environment-specific meta reward shaping ϕ with gradient descent by Eq. (11)
 - 7: **end if**
-

3.5 Theoretical Study and Discussions

Policy Consistency. For classical reward shaping, any potential-based reward shaping can guarantee the consistency of optimal policy [4, 20]. However HMRL aims to learn meta modules to fulfill policy adaptation on different tasks rather than getting the optimal policy on a specific task. We are concerned with the question if the meta reward shaping scheme is applied on new task learning, whether the optimal policy keeps still. Our theory shows, using meta reward shaping can maintain policy consistency on new tasks.

Theorem 1. *Assume a new task $\mathcal{T} = \{\mathcal{S}, \mathcal{A}, \mathcal{P}, \gamma, \mathcal{R}\}$ has the same meta state space \mathcal{S}_m with the obtained meta RL model. If the environment-specific meta reward shaping $\mathcal{F}(s_m, s_m')$ is utilized, then it will have the consistency of optimal policy between the extended new task $\mathcal{T}' = \{\mathcal{S}, \mathcal{S}_m, \mathcal{A}, \mathcal{P}, \gamma, \mathcal{R} + \mathcal{F}\}$ and original task \mathcal{T} .*

Proof. Consider the accumulated reward for task \mathcal{T} , i.e. $G_{\mathcal{T}} = \sum_{i=0}^T \gamma^i \mathcal{R}(s_i)$, where $s_i \in \mathcal{S}$. Further the accumulated reward for extended task \mathcal{T}' can be calculated as

$$\begin{aligned} G_{\mathcal{T}'} &= \sum_{i=1}^T \gamma^i [\mathcal{R}(s_i) + \mathcal{F}(s_m^i, s_m^{i+1})] \\ &= \sum_{i=1}^T \gamma^i [\mathcal{R}(s_i) + (\gamma \phi(s_m^{i+1}) - \phi(s_m^i))] \\ &= \sum_{i=1}^T \gamma^i \mathcal{R}(s_i) + \sum_{i=1}^T (\gamma^{i+1} \phi(s_m^{i+1}) - \gamma^i \phi(s_m^i)) \\ &= G_{\mathcal{T}} - \phi(s_m^0), \end{aligned} \quad (12)$$

which indicates that the difference between the accumulated reward $G_{\mathcal{T}'}$ and $G_{\mathcal{T}}$ is a constant because $\phi(s_m^0)$ is independent with the policy. Further the goal to calculate the optimal policy for two tasks is to maximize the accumulated rewards $G_{\mathcal{T}'}$ and $G_{\mathcal{T}}$ respectively. The optimal policies are equivalent based on (12). \square

4 EXPERIMENTS

4.1 Evaluation Protocol

In this section, two sparse reward RL applications are chosen from the popular *gym-miniworld* [2] to test the efficiency of the proposed HMRL, i.e., *Hallway* and *Maze*. The reward can be obtained only after each task is finished and in all the two problems, we set the reward as $-used-steps$. Both 1st and 3rd scenarios are studied, which also are considered as different environments¹. The agent in the 1st person environment can only obtain the observed information

¹In experiment section of this paper, scenario is equivalent to environment




Figure 2: The Hallway Problem. 2(c) shows the scenario/map of the evaluation task, and the task in both 1st person environment and 3rd person environment are evaluated.

from its own moving direction, while the 3rd environment can obtain the whole scenario observation, e.g., Figure 2 and Figure 6. Eventually, 1st person environment is much more difficult than 3rd person environment because of partially observation.

Four baseline meta RL algorithms are compared with the proposed HMRL: MAML [7] (the groundbreaking meta RL algorithm), NoRML [33] (the state-of-the-art meta RL algorithm for sparse reward setting), variBAD [36] (the state-of-the-art Bayesian meta RL algorithm for sparse reward setting) variBAD [36] and HMRL-w/o-ms (HMRL without meta state embedding, while the potential based reward shaping directly conditioned with different environment states, which is from [37]). To emphasize, MAML is enhanced with the flexible meta-SGD extension technique in [15].

All baselines except HMRL-w/o-ms are trained on each scenario independently, because they are not adaptive to the cross environment setting. While HMRL-w/o-ms is trained on cross environment setting as the ablation experiments.

For fair comparison, all methods are trained with the same inner-loop and outer-loop iteration numbers.

In the experiment, our policy network contains two layers Convolutional Neural Networks with the kernel size of [5, 5], then we use two fully connected layer with the size of 128. We use the ReLU activation for all models. All the baselines use the same policy network. MAML and NoRML’s implementations are from <https://github.com/google-research/google-research/tree/master/norml>. While variBAD’s implementation is from <https://github.com/lmzintgraf/variBAD>. In the variBAD, we use a 128 unit RNN as the trajectories embedding.

4.2 Problem I: Hallway

The agent aims to reach the red box (or target point) at the end of a straight hallway, while avoiding bumping into walls. Different tasks in each environment (e.g., 1st person or 3rd person) are defined by the random initialization of the start location. For both 1st and 3rd person environments, the observation of the agent in each step contains a $80 \times 60 \times 12$ tensor (including recent four frames with RGB channels of its view towards its moving direction, while Figure 2(a) and Figure 2(b) denotes each frame of two environments respectively) and the coordinates of the agent and the red box. The action space contains four moving directions

$$\{\text{up, down, left, right}\},$$

with 0.2 unit. The task ends if the agent reaches within 1.0 distance to the red box and each rollout has a maximum horizon of 80 steps.

The common meta state s_m is set to be the coordinates of the agent in two environments and the red box, which is denoted as a 3×2 matrix. As a result, the meta state embedding function h can be considered as simple concatenation of the location information without learning. The meta reward shaping potential function ϕ is set as a fully-connected network with two 32 units hidden layers. The policy network π is set as a Convolutional Neural Networks and the setting details are proposed in 4.1.

Additionally for evaluation, we introduce a new task with both 1st person and 3rd person environments within a different map, i.e. Figure 2(c). The evaluation task is set with the start point of the agent at [1.0, 6.0] and the red box at [3.0, 9.0]. In order to show the transfer-ability of our algorithm within sparse reward and cross-environment tasks, we further modify a new action space setting:

$$\{\text{turn left, turn right, move forward, move back}\},$$

where “turn left” and “turn right” aims to change the agent’s direction left and right respectively in a counterclockwise/clockwise direction with 30 degree, while “move forward” and “move back” aims to move the agent on its current direction with 0.2 unit.

From Figure 3(a) and Figure 3(b), our proposed HMRL makes an improvement compared with all other baselines in both 1st person and 3rd person environments.

Compared with HMRL-w/o-ms which directly use the reward shaping without meta state embedding, HMRL can obviously benefit from the cross-environment knowledge. While directly using reward shaping will introduce more noise from heterogeneous state of different environments.

Figure 3(c) presents the step numbers for the agent arriving at the target red box from 10 randomly chosen tasks, and our proposed HMRL uses the minimum steps in the worst task on 1st environment and achieves a decent holistic performance. HMRL can promote agent in relatively harder case (e.g., 1st person environment compared with 3rd person environment) exploring more efficiently by the utilization of cross-environment meta information.

Figure 3(d) shows the heat map of the averaged potential value in the cross environment meta reward shaping on a specific training task. The targets (red boxes) are located on the right of the map, and the agent will obtain a positive reward shaping if it moves from the smaller potential value to the larger one. The heat map proves that the learned cross-environment meta reward shaping encourages the agent moving to the target (red box).

If we proceed further, we can find the learning curves of most baselines shows that training in 1st person environment(Fig.3(a)) is more difficult than 3rd person environment(Fig.3(b)) at the early time because of the low sample efficiency caused by sparse reward and partial observation. But HMRL can bring the effective prior from 3rd person environment to 1st person environment through the cross-environment part. So it represented the training process with a smaller gap between the two environments.

All the evidence above confirmed that HMRL can reach our main motivation of using cross environment information to solve sparse reward problem in meta reinforcement learning because it assuredly learned the cross environment information and the learned cross environment information positively help the meta reinforcement learning tasks with sparse reward problem.

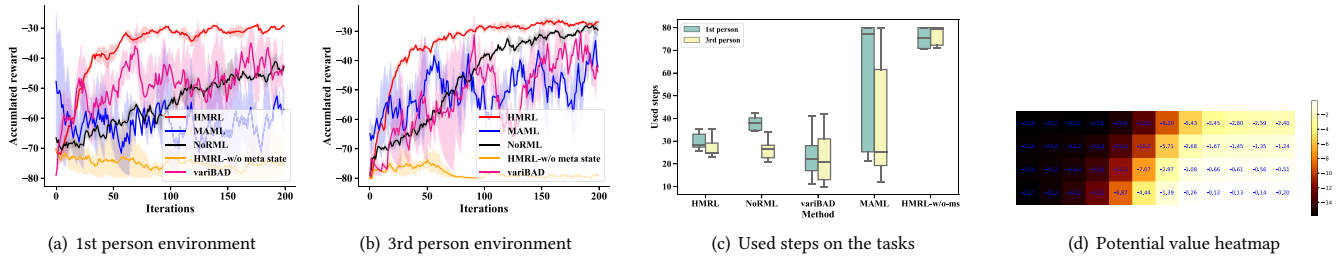


Figure 3: Results of the 1st person and 3rd person environments in Hallway problem. 3(a) and 3(b) shows the learning curve of both environments in training process. 3(c) shows the step numbers for the agent finishes the 10 randomly chosen tasks with the trained policy. 3(d) shows the heat map of the averaged potential value in the cross environment meta reward shaping.

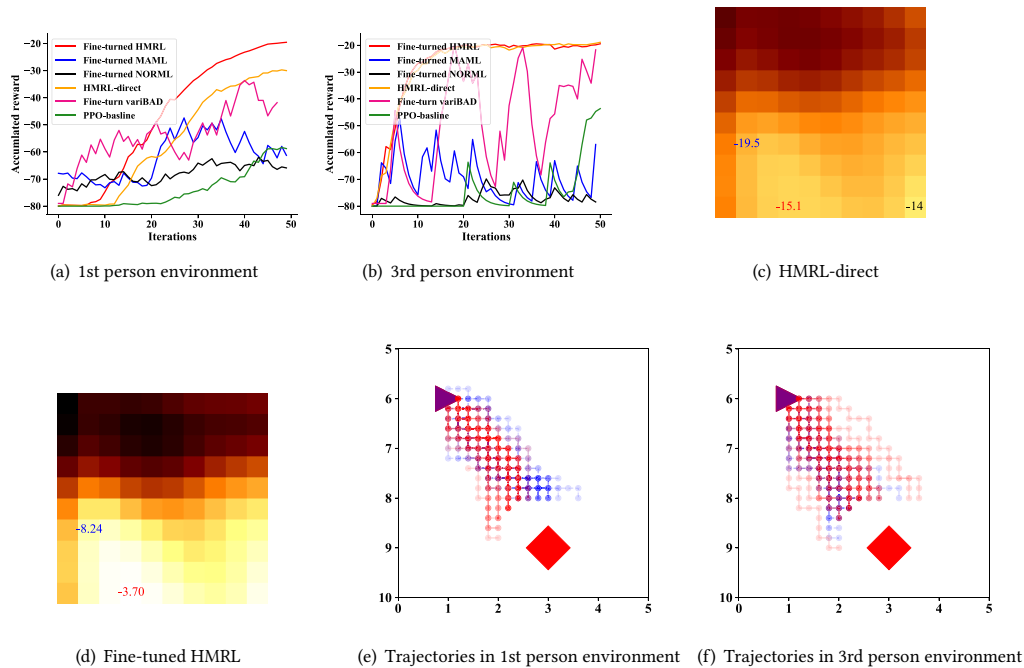


Figure 4: The results on the new evaluation tasks shown in Figure 2(c) within both 1st person environment and 3rd person environment. The blue, red and black colors denote the “start” point, “target” point and “misleading” point (defined in the main context) respectively. The triangle is the start point and the red box is the target. And the blue trajectories are from the policy trained by HMRL-direct, the red trajectories are from the policy trained by Fine-tuned HMRL. All the trajectories are collected from the policies ran on the corresponding evaluation task for 20 times.

Figure 4 shows the result comparison on the evaluation task in Figure 2(c) in order to show the transfer-ability of different algorithms. The policy obtained by our proposed HMRL can be directly applied to the new task, which is one of the advantages of HMRL and denoted as HMRL-direct. All compared algorithms are fine-tuned for better adapting to the evaluation task including the HMRL. Moreover, the RL method PPO² [24] is compared and trained from the beginning on this evaluation task as a bottom line.

In Figure 4(a), we can find that all meta RL algorithms are more efficient than PPO.

Most important, the HMRL outperforms all other baselines on both direct transfer and fine-tuned transfer cases in both environments, while the fine-tuned transfer version performs much better in the 1st person environment. Figure 4(c) and 4(d) present the heatmap of the potential value obtained by HMRL-direct and fine-tuned HMRL respectively. In Figure 4(c), the potential value of the target point is not the largest one, as a result we call the point with the “misleading” point because of the larger potential value

²All meta learning methods in this paper use PPO as the inner training methods




Figure 5: The learning curve of different environments in Maze problem.




Figure 6: Visualization of Maze Problem and corresponding potential value heatmap. In the chosen task within 2RS4, the red box was located in the upper right.

usually guiding the agent to move towards. Conversely, there is no “misleading” point for the fine-tuned HMRL heatmap, so that the agent have a higher possibility to go on a closer way. This guess can also be supported by the average trajectories of 1st environment in Fig 4(e), where the policy trained by HMRL-direct guides the agent go to the right-side where the misleading point located. But at the same time, the average trajectories of 3rd person environment in Fig 4(f) shows the policy trained by HMRL-direct doesn’t have this tendency, and also both HMRL cases have the similar final result in the learning curve in Fig 3(b). In this point, it seems that Fine-tuned HMRL can get a better performance when treating with new partial observed tasks suffer from sparse reward problem.

Overall, in the experiment, HMRL can jointly utilize the information of tasks from different environments and using the learned cross environment information to speed up and improve both new task learning as well as meta policy training under sparse reward problem. All these figures prove the transferability and policy learning efficiency of HMRL.

4.3 Problem II: Maze

The agent in Maze problem aims to reach the red box (or target point) in a randomly generated maze. Three different environments are considered, i.e., 2RS3, 2RS4 and 3RS3 (Figure 6). The task for each environment is defined by the random initialization of scenario (both start point and target point). The brief introduction of these environments are proposed as follows: 1) 2RS3: a maze with 2 rooms and the size of each room is 3 (shown in 6(a)) which is called $2R(room)S(size)3$ with 2RS3 for short. The observation is obtained in the 3rd person view, and the action space is set to be moving to {up, down, left, right} with a length of 0.2; 2) 2RS4: a maze with 2 rooms and the size of each room is 4 (shown in 6(b)). The observation is obtained in the 1st person view, and the action space is set to be changing the agent’s direction left and

right ({turn left, turn right}) and moving to forward or backward ({move forward, move back}) with a length of 0.2; 3) 3RS3: a maze with 3 rooms and the size of each room is 3 (shown in 6(c)), we also set the step length as 0.4 to speed up the task. The other settings are the same as 2RS3.

For all three environments, the observation contains a $80 \times 60 \times 12$ tensor (including recent four frames with RGB channels and the coordinates of the agent and target red box). Additionally, the task will end if the agent reaches within 1.0 to the target red box and each rollout has a maximum horizon of 150 steps. Due to the difference among these environments, directly setting the common meta state s_m to be the coordinates of the agent and the red box is useless. For this reason, we set the state embedding function h to be a linear function based on each environment’s scenario size (denoted as the domain knowledge z), with a 3×2 mini-map as the common meta state s_m . The meta reward shaping potential function ϕ is set as a fully-connected neural network with two 32 units hidden layers. The policy network π is set as a Convolutional Neural Networks and the setting details are proposed in 4.1. The learning curves are shown in Fig. 5.

It is obvious that our HMRL performs better than other baselines in all environments. And in Figure 6(d), Figure 6(e) and Figure 6(f), the potential values on different environments’ meta states are proposed. The meta reward shaping can lead the agent to the true target red box, for the reason that the closer to the target leads to the larger potential value. So the evidence above indicates that HMRL can transfer meta knowledge among multiple environments with more efficient reward shaping. Furthermore, it shows that in this experiment with a different setting in cross environment meta reward shaping and scenarios than the experiment in 4.2, our HMRL can also assuredly learn the cross environment information from more different environments and using the cross environment

information to help training the meta reinforcement learning tasks with sparse reward problem.

In this experiment, HMRL still can utilize the cross environment information in a more general setting. So that we can say this experiment further prove the transferability and policy learning efficiency of HMRL.

5 CONCLUSION

A hyper-meta reinforcement learning algorithm (HMRL) has been proposed to handle the typical sparse reward problem. The cross-environment meta state embedding and environment-specific meta reward shaping modules are devised to cooperate with the general meta policy learning module, which increase the generalization and learning efficiency of meta-RL. Furthermore, a shared meta state space is designed with a specific meta reward shaping technique, which enables and improves the capability of extracting complementary meta knowledge, introducing domain knowledge, speeding up the policy learning for cross-environment tasks. Experimental results with sparse reward demonstrate the superiority of the proposed method on both transfer-ability and efficiency. Our work can also be considered as an early search of designing intrinsic reward in cross-environment problems.

In the future, we will focus more on searching optimal reward in more flexible cross-environment problems. As recent attempts on cross-environment situations are limited in action space and state space, we have taken a small but pioneering step forward. Also, we will explore more challenging environment changes (perhaps by collaborating with other party to access the real-world platforms) whereby our method may have much potential. The source code will be made public available.

6 ACKNOWLEDGE

This work was supported in part by National Key Research and Development Program of China (No. 2020AAA0107400), NSFC (No. 12071145), Shanghai Municipal Science and Technology Major Project (No. 2021SHZDZX0102), STCSM (No. 18DZ2270700 and 20511101100), the Open Research Projects of Zhejiang Lab (NO.2021KE0AB03) and a grant from Shenzhen Institute of Artificial Intelligence and Robotics for Society.

REFERENCES

- [1] Elena Arcari, Andrea Carron, and Melanie N Zeilinger. 2020. Meta Learning MPC using Finite-Dimensional Gaussian Process Approximations. *arXiv preprint arXiv:2008.05984* (2020).
- [2] Maxime Chevalier-Boisvert. 2018. gym-miniworld environment for OpenAI Gym. <https://github.com/maximecb/gym-miniworld>.
- [3] Ignasi Clavera, Anusha Nagabandi, Ronald S Fearing, Pieter Abbeel, Sergey Levine, and Chelsea Finn. 2018. Learning to Adapt: Meta-learning for model-based control. *arXiv preprint arXiv:1803.11347* 3 (2018).
- [4] Sam Michael Devlin and Daniel Kudenko. 2012. Dynamic potential-based reward shaping. In *AAMAS*. 433–440.
- [5] Yan Duan, Xi Chen, Rein Houthoofd, John Schulman, and Pieter Abbeel. 2016. Benchmarking deep reinforcement learning for continuous control. In *ICML*. 1329–1338.
- [6] Yan Duan, John Schulman, Xi Chen, Peter L Bartlett, Ilya Sutskever, and Pieter Abbeel. 2016. RL²: Fast reinforcement learning via slow reinforcement learning. *arXiv preprint arXiv:1611.02779* (2016).
- [7] Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017. Model-agnostic meta-learning for fast adaptation of deep networks. In *ICML*. 1126–1135.
- [8] Carlos Florensa, David Held, Markus Wulfmeier, Michael Zhang, and Pieter Abbeel. 2017. Reverse Curriculum Generation for Reinforcement Learning. In *CoRL*. 482–495.
- [9] Justin Fu, Sergey Levine, and Pieter Abbeel. 2016. One-shot learning of manipulation skills with online dynamics adaptation and neural network priors. In *IROS*. 4019–4026.
- [10] Mohammad Ghavamzadeh, Shie Mannor, Joelle Pineau, and Aviv Tamar. 2016. Bayesian reinforcement learning: A survey. *arXiv preprint arXiv:1609.04436* (2016).
- [11] Abhishek Gupta, Russell Mendonca, YuXuan Liu, Pieter Abbeel, and Sergey Levine. 2018. Meta-reinforcement learning of structured exploration strategies. In *NeurIPS*. 5302–5311.
- [12] Swaminathan Gurumurthy, Sumit Kumar, and Katia Sycara. 2019. MAME: Model-Agnostic Meta-Exploration. *arXiv preprint arXiv:1911.04024* (2019).
- [13] Joshua Hare. 2019. Dealing with Sparse Rewards in Reinforcement Learning. *arXiv preprint arXiv:1910.09281* (2019).
- [14] Lin Lan, Zhenguo Li, Xiaohong Guan, and Pinghui Wang. 2019. Meta Reinforcement Learning with Task Embedding and Shared Policy. In *IJCAI*. 2794–2800.
- [15] Zhenguo Li, Fengwei Zhou, Fei Chen, and Hang Li. 2017. Meta-SGD: Learning to learn quickly for few-shot learning. *arXiv preprint arXiv:1707.09835* (2017).
- [16] Evan Zheran Liu, Aditi Raghunathan, Percy Liang, and Chelsea Finn. 2020. Explore then Execute: Adapting without Rewards via Factorized Meta-Reinforcement Learning. *arXiv preprint arXiv:2008.02790* (2020).
- [17] Nikhil Mishra, Mostafa Rohaninejad, Xi Chen, and Pieter Abbeel. 2018. A simple neural attentive meta-learner. In *ICLR*.
- [18] 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. 2015. Human-level control through deep reinforcement learning. *Nature* 518, 7540 (2015), 529–533.
- [19] Sanmit Narvekar, Jivko Sinapov, Matteo Leonetti, and Peter Stone. 2016. Source task creation for curriculum learning. In *AAMAS*. 566–574.
- [20] Andrew Y Ng, Daishi Harada, and Stuart Russell. 1999. Policy invariance under reward transformations: Theory and application to reward shaping. In *ICML*. 278–287.
- [21] Martin Riedmiller, Roland Hafner, Thomas Lampe, Michael Neunert, Jonas Degrave, Tom Wiele, Vlad Mnih, Nicolas Heess, and Jost Tobias Springenberg. 2018. Learning by Playing Solving Sparse Reward Tasks from Scratch. In *ICML*. 4344–4353.
- [22] Stephane Ross, Brahim Chaib-draa, and Joelle Pineau. 2007. Bayes-Adaptive POMDPs. In *NeurIPS*. 1225–1232.
- [23] Steindór Sæmundsson, Katja Hofmann, and Marc Peter Deisenroth. 2018. Meta reinforcement learning with latent variable Gaussian processes. In *UAI*. 642–652.
- [24] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal Policy Optimization Algorithms. *arXiv preprint arXiv:1707.06347* (2017).
- [25] David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. 2016. Mastering the game of Go with deep neural networks and tree search. *Nature* 529, 7587 (2016), 484–489.
- [26] Flood Sung, Li Zhang, Tao Xiang, Timothy Hospedales, and Yongxin Yang. 2017. Learning to learn: Meta-Critic networks for sample efficient learning. *arXiv preprint arXiv:1706.09529* (2017).
- [27] Maxwell Svetlik, Matteo Leonetti, Jivko Sinapov, Rishi Shah, Nick Walker, and Peter Stone. 2017. Automatic curriculum graph generation for reinforcement learning agents. In *AAAI*. 2590–2596.
- [28] Chen Tessler, Shahar Givony, Tom Zahavy, Daniel J. Mankowitz, and Shie Mannor. 2016. A Deep Hierarchical Approach to Lifelong Learning in Minecraft. In *AAAI*. 1553–1561.
- [29] Jane X Wang, Zeb Kurth-Nelson, Dhruva Tirumala, Hubert Soyer, Joel Z Leibo, Remi Munos, Charles Blundell, Dharshan Kumaran, and Matt Botvinick. 2016. Learning to reinforcement learn. *arXiv preprint arXiv:1611.05763* (2016).
- [30] Hua Wei, Guanjie Zheng, Huaxiu Yao, and Zhenhui Li. 2018. IntelliLight: A Reinforcement Learning Approach for Intelligent Traffic Light Control. In *KDD*. 2496–2505.
- [31] Jan Wohlke, Felix Schmitt, and Herke van Hoof. 2020. A Performance-Based Start State Curriculum Framework for Reinforcement Learning. In *AAMAS*. 1503–1511.
- [32] Zhongwen Xu, Hado P van Hasselt, and David Silver. 2018. Meta-gradient reinforcement learning. In *NeurIPS*. 2396–2407.
- [33] Yuxiang Yang, Ken Caluwaerts, Atil Iscen, Jie Tan, and Chelsea Finn. 2019. NoRML: No-Reward Meta Learning. In *AAMAS*. 323–331.
- [34] Lantao Yu, Tianhe Yu, Chelsea Finn, and Stefano Ermon. 2019. Meta-Inverse Reinforcement Learning with Probabilistic Context Variables. In *NeurIPS*. 11749–11760.
- [35] Jin Zhang, Jianhao Wang, Hao Hu, Yingfeng Chen, Changjie Fan, and Chongjie Zhang. 2020. Learn to Effectively Explore in Context-Based Meta-RL. *arXiv preprint arXiv:2006.08170* (2020).
- [36] Luisa Zintgraf, Kyriacos Siarlis, Maximilian Igl, Sebastian Schulze, Yarin Gal, Katja Hofmann, and Shimon Whiteson. 2020. VariBAD: A Very Good Method for Bayes-Adaptive Deep RL via Meta-Learning. In *ICLR*.
- [37] Haosheng Zou, Tongzheng Ren, Dong Yan, Hang Su, and Jun Zhu. 2019. Reward shaping via meta-learning. *arXiv preprint arXiv:1901.09330* (2019).