

MODEL-BASED META AUTOMATIC CURRICULUM LEARNING

Zifan Xu *

Department of Computer Science
The University of Texas at Austin
zfxu@utexas.edu

Yulin Zhang & Shahaf S. Shperberg

Department of Computer Science
The University of Texas at Austin
{yulin, shperbsh}@cs.utexas.edu

Reuth Mirsky

Department of Computer Science
Bar Ilan University.
mirskyr@cs.biu.ac.il

Yuqian Jiang, Bo Liu, Peter Stone

Department of Computer Science
The University of Texas at Austin
{jiangyuqian, bliu, pstone}@cs.utexas.edu

ABSTRACT

Curriculum learning (CL) has been widely explored to facilitate the learning of hard-exploration tasks in reinforcement learning (RL) by training a sequence of easier tasks, often called *a curriculum*. While most curricula are built either manually or automatically based on heuristics, e.g. choosing a training task which is barely beyond the current abilities of the learner, the fact that similar tasks might benefit from similar curricula motivates us to explore meta-learning as a technique for curriculum generation or *teaching* for a distribution of similar tasks. This paper formulates the meta CL problem that requires a *meta-teacher* to generate the curriculum which will assist the *student* to train toward any given *target task* from a task distribution based on the similarity of these tasks to one another. We propose a model-based meta automatic curriculum learning algorithm (MM-ACL) that learns to predict the performance improvement on one task when the student is trained on another, given the current status of the student. This predictor can then be used to generate the curricula for different target tasks. Our empirical results demonstrate that MM-ACL outperforms the state-of-the-art CL algorithms in a grid-world domain and a more complex visual-based navigation domain in terms of sample efficiency.

1 INTRODUCTION

Curriculum learning has emerged as a popular approach to address the challenges of sparse rewards and hard exploration in RL problems (Narvekar et al., 2020). Fig. 1 (a) shows a simple example of a student learning to navigate from a starting location (red) to a goal location (green), which might benefit from learning *source* tasks with closer starting locations (orange). While prior research has focused on developing curriculum learning algorithms that leverage some prior knowledge from experts, such as designing heuristics to select tasks (Florensa et al., 2018) or incorporating domain-specific knowledge to boost curriculum performance (Romac et al., 2021), such forms of prior knowledge are often limited. In this paper, we instead focus on *learning how to teach* for a given domain, assuming that a teacher can be learned offline from extensive experience teaching numerous students on a set of similar tasks in the same domain. As illustrated in Fig. 1 (b), a navigation domain can consist of different environments, but common domain knowledge might be learned to teach similar navigation tasks. Our goal is to train a generalized teacher that can identify effective curricula for new students learning an unseen, similar task within the domain. Since a good teacher can benefit any number of students, this setting places a greater focus on the eventual quality of the learned teacher, than on how long it takes to train the teacher.

While initial studies on CL in reinforcement learning used manually-designed curricula, based on humans’ knowledge of the domain (Sanger, 1994), the DRL community has developed a family of mechanisms called Automatic Curriculum Learning (ACL), which automatically adapts the distribution of source tasks according to the capabilities of learning agents (Portelas et al., 2020b). Such an adaptation of the source task distribution, however, is typically performed with heuristics or surrogate objectives. For example, GoalGAN (Florensa et al., 2018) trains a generator to generate source tasks with intermediate difficulty levels, which can avoid both too hard tasks that are infeasible to learn and too easy tasks that do not provide any novelty. These heuristics are often suboptimal and can even be

* Use footnote for providing further information about author (webpage, alternative address)—*not* for acknowledging funding agencies. Funding acknowledgements go at the end of the paper.

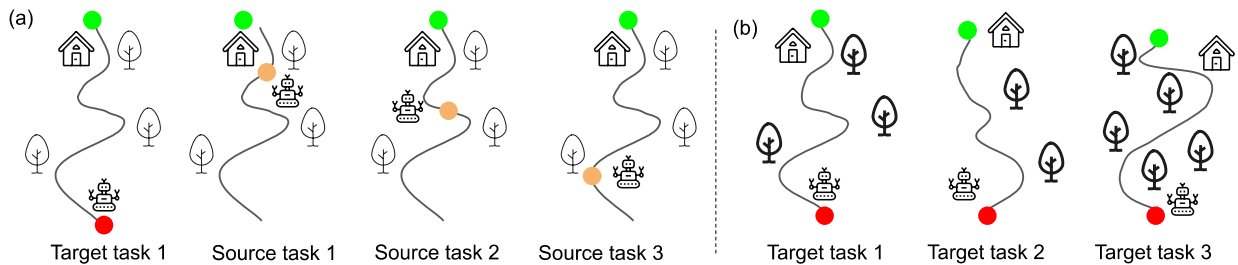


Figure 1: A navigation domain with a robot navigating from red to green markers, (a) a sample target task with a set of simpler source tasks (starting from the orange markers) that form the curriculum; (b) examples of four different target tasks in the same navigation domain.

arbitrarily bad. For example, the definition of intermediate difficulty level often relies on a manually-picked threshold of the evaluation metrics, which can lead to transitioning too early to hard tasks that are actually not learnable. In addition, estimating the heuristics can be costly due to random exploration. For example, `GoalGAN` has to try random tasks, measure their difficulty levels, and update the generator that generates the tasks of intermediate difficulty.

Alternatively, the CL problem can naturally be viewed as a *meta-learning* problem with an *inner loop* training the agent on each source task, and an *outer loop* updating the source task or the teacher module such that an *outer objective* of maximizing the agent’s performance on the target task will be improved. Such an explicitly defined outer objective forms the key to meta-learning (Hospedales et al., 2020). However, this meta-learning view of CL or meta-CL still remains mostly unexplored.

To solve such a bi-level problem, meta-learning usually trains the meta-representation (teacher) in an *offline* setting (Hospedales et al., 2020). Denoting the complete process of a student learning a target task *from scratch* under the supervision of a teacher as a *lifetime*, the teacher is trained in an offline manner over multiple lifetimes, where each lifetime corresponds to a new student learning a different target task. This process is usually called the *meta-training* phase. Then, the teacher is tested in new lifetimes with novel, but related target tasks, which is called the *meta-testing* phase. Looking back to the previous example in Fig. 1 (b), the meta-training phase will have the student taught by the teacher to learn a class of similar navigation tasks but with different landscapes (target tasks 1-3). Such teaching experiences can improve and prepare the teacher to teach a new target task 4 from this navigation domain. In this paper, we first formulate such a meta-CL problem that requires meta-training a teacher in an *offline* setting, which differs from an *online* setting where the teacher and the student co-evolve within one single lifetime.

To solve the meta-CL problem, we introduce a model-base meta automatic curriculum learning (MM-ACL) algorithm that is applicable with any policy gradient RL algorithm. MM-ACL selects the training tasks that will, to the greatest extent, improve the performance of the student on a given target task. To estimate the improvement, MM-ACL maintains a *learning dynamics model* which predicts the improvement on a given validation task after training on a given source task for some number of gradient updates. During the meta-training phase, MM-ACL performs “online validation” on a validation task both before and after the training of a selected source task. Then the improvement of the validation task given a source task can be measured and used as a data point for learning the learning dynamics model. Instead of learning a task selection model directly, learning such a learning dynamics model helps to improve sample efficiency and generalize more easily to different target tasks.

We test MM-ACL on two simple MiniGrid domains and one more complex AI Habitat domain. The method is compared with a basic multi-task training baseline, an RL-based meta-curriculum learning algorithm, a handcrafted curriculum, and two state-of-the-art ACL algorithms. MM-ACL outperforms the baselines in terms of sample efficiency, and requires fewer meta-training lifetimes compared to the RL-based meta-curriculum learning algorithm.

2 RELATED WORK

This section provides an overview of prior research in three sub-fields, namely ACL for RL (Sec. 2.1), meta-RL (Sec. 2.2), and meta-ACL (Sec. 2.3).

2.1 ACL FOR RL

MM-ACL solves a meta-CL problem that reuses previous teaching experiences for future target tasks, which differs from existing ACL algorithms that focus on online adaptation of the teacher for a specific target task or no target task (Dennis et al., 2020; Wang et al., 2020). The implications of this difference are two-fold: First, MM-ACL is an offline algorithm that trains a teacher module during a meta-training phase, then deploys the teacher without further updating on a specific set of target tasks. Existing ACL algorithms, on the other hand, are online algorithms which co-evolve teachers during the training of their students (Portelas et al., 2020a; Florensa et al., 2018). Second, existing ACL algorithms do not maintain an explicit outer-loop objective. The tasks are selected by heuristics or surrogate objectives which might not align with the objective of improving the learner’s performance on the set of target tasks.

2.2 META-RL

One important line of research on meta-RL is meta-policy gradients (MPG), which utilizes differentiability of the update function with respect to some hyper-parameters in the policy-gradient-based RL algorithms. MPG performs online cross validation (Sutton, 1992) to collect hold-out trajectories to estimate outer-loop loss after every few gradient updates, then computes a meta-gradient of the hyper-parameters based on the outer-loop loss. Those hyper-parameters can be thought of as different tasks from a curriculum learning perspective, but are limited to differentiable parameters, for example, discount factor (Xu et al., 2018), intrinsic reward (Zheng et al., 2018), and auxiliary tasks (Veeriah et al., 2019). However, the parameters that facilitate learning most are sometimes directly related to the tasks. For example, transition dynamics and state space are not differentiable without having access to an analytic function of the transition dynamics. MM-ACL, however, can select tasks that are controlled by any form of parameters including non-differentiable ones. It deserves noting that, even though most MPG algorithms are online to make use of the efficiency of the meta-gradient, which differs from the offline setting employed by MM-ACL, MPG can also be performed offline by introducing episodic memory (Zheng et al., 2020).

Some existing meta-RL approaches do not use meta-gradients, and instead control specific aspects of RL that are generally non-differentiable, for example, the exploration-exploitation tradeoff (Ishii et al., 2002). They can also be thought of as a form of curriculum learning, but only limited to one specific aspect of the task, and therefore do not address the fully general CL problem.

2.3 META-ACL

There is also prior work that formulates the meta-CL problem in a similar fashion to the formalization presented in this paper. In CMDP’s, introduced by Narvekar & Stone (2019), the teacher is represented as a curriculum policy and is optimized by an RL algorithm in the outer-loop optimization. Such a curriculum policy has also been demonstrated to be generalizable to different target tasks (Narvekar & Stone, 2020). However, using RL as the outer-loop optimizer is extremely costly. Their approach uses hundreds of lifetimes to train an optimal curriculum policy with a tubular algorithm, where the low-level RL algorithm for learning was deployed in a simple grid-world domain. Modern deep RL algorithms may suffer from higher variance and longer training steps. Another meta-curriculum learning framework AGAIN (Portelas et al., 2020c) reuses previously generated curricula based on curriculum distillation. However, AGAIN has two limitations: i) it is only defined for one particular base ACL algorithm, ALP-GMM (Portelas et al., 2020a), and cannot be adapted to different ACL algorithms; ii) It aims to optimize a heuristic of *absolute learning progress* (used by ALP-GMM) instead of tackling a meta-CL problem directly. Intuitively, based on the previous curricula, AGAIN exploits the useful source tasks that fulfill the heuristic and prunes the bad tasks from random exploration to estimate the heuristic. This algorithm aims mostly to reduce the random exploration needed to estimate the heuristic when teaching a new student, which means that the algorithm is still limited to one specific heuristic.

Meta-CL has also been explored for supervised learning (Fan et al., 2018; Jiang et al., 2018). However, curriculum learning for RL encounters special challenges, such as high variance and exploration, that requires specially designed algorithms to tackle.

Lastly, in contrast to our work, which uses meta-learning to improve curriculum learning, some prior work (Stergiadis et al., 2021; Wu et al., 2021) focuses on curriculum learning for meta-learning. More specifically, this line of work aims at automatically adapting the training task distribution for generic meta-learning (not necessary CL) problems. These types of algorithms are similar to prior ACL algorithms without any meta-learned component.

3 PRELIMINARIES

This section formally introduces the meta-CL domain, as well as the student and teacher in Sec. 3.1, 3.2, and 3.3, respectively. In Sec. 3.4, we provide a formal definition of the meta-CL problem.

3.1 META-CL DOMAIN

Meta-curriculum learning aims to leverage a set of target tasks during the meta-training phase to learn *how to teach*, i.e. gain knowledge that can improve the learning of the agent on new target tasks that were not seen during meta-training. We use $\mathcal{M} := \{m\}$ to denote a meta-CL domain defined by a set of all possible tasks m , from which a target task m_{target} can be chosen. For example, in the navigation domain shown in Fig. 1, a target task m is uniquely defined by a combination of a navigation environment and start-goal locations. All the possible combinations constitute a meta-CL domain \mathcal{M} . However, it is often the case that the teacher can only control some aspects of the task when solving a given target task. For example, when the navigation robot in Fig. 1 is situated in the environment specified by the target task, the teacher can change the start-goal locations, but may not be able to move the student to a completely different environment (e.g. moving the trees or the house in the navigation domain). For this reason, we factorize a meta-CL domain into two orthogonal spaces: a controllable space \mathcal{W} and an uncontrollable space \mathcal{V} . A task m can, therefore, be represented as a tuple (w, v) with $w \in \mathcal{W}$ and $v \in \mathcal{V}$. Considering the same example of the navigation domain, w and v are the variables that specify the start-goal locations and the navigation environment respectively. For a given target task $m_{target} = (w_{target}, v_{target})$ in a lifetime, a teacher can control the variable w to generate a sequence of source tasks as a curriculum that can facilitate the learning of m_{target} . For the rest of the paper, we assume that the controllable and uncontrollable spaces are given as a part of the problem.

3.2 STUDENT

We consider a student in a meta-CL problem as an RL algorithm that can improve its policy based on the trajectories collected from a given training task. We use ψ to denote the current status of a student, which contains all the internal states of the RL algorithm, such as all the weights of the actor and critic neural networks and the states of the optimizer if the algorithm is an actor-critic deep RL algorithm based on deep neural networks. Such a student can be updated from status ψ_t to ψ_{t+1} during a training iteration t based on a given training task m_t , which we formulate as follows:

$$\tau_t = \text{Rollout}(\psi_t, m_t), \quad \psi_{t+1} = \text{Train}(\psi_t, \tau_t) \quad (1)$$

Here, *Train* represents a training process that updates the student ψ , and *Rollout* generates a rollout trajectory τ_t based on the current student ψ_t on the task m_t . For example, the training process of REINFORCE (Sutton & Barto, 2018) with a linearly-decayed learning rate is a simple policy gradient update combined with a learning rate update:

$$\alpha_{t+1} = \alpha_t - 0.0001, \quad \theta_{t+1} = \theta_t - \alpha_t \nabla_{\theta_t} \mathcal{G}(\tau_t) \quad (2)$$

with $\psi_t = (\theta_t, \alpha_t)$ being the student status, θ_t the parameters of the policy, α_t the learning rate, and $\mathcal{G}(\tau_t)$ the returns of τ_t . Then, during a lifetime lasting N training iterations with a curriculum (m_1, m_2, \dots, m_N) , the student, starting from an initial status ψ_0 , is updated recursively following Eq. 1, which generates a sequence of student statuses $(\psi_0, \psi_1, \dots, \psi_N)$.

3.3 META-TEACHER

A meta-teacher will dynamically decide on which source tasks to train based on the information of (1) *which student is being trained* and (2) *what is the target task*. As defined in Sec. 3.1 and Sec. 3.2, the student and the target task can be represented as ψ and (w_{target}, v_{target}) respectively. However, for complex domains solved by deep RL algorithms, it might be prohibitive to take ψ as input since it contains all the neural network weights. Instead, in the ACL literature, a more commonly used representation of a student status is *contextual information* that comprises the training history of the student. In this paper, a context $c = (w_j, \bar{\mathcal{G}}_j)_{j=1}^J \in \mathcal{C}$ is a sequence of J training task-return pairs experienced so far in a lifetime, where w_j is the set of controllable parameters that define a source task (w_j, v_{target}) . As the teacher cannot control the v parameters, it is set to v_{target} throughout the lifetime. $\bar{\mathcal{G}}_j$ is the average return of the rollout trajectories, and \mathcal{C} is a set of all possible contextual information.

It is worth noting that the contextual information in our approach consists of the returns rather than the entire trajectories. This allows for the estimation of the student’s status over a longer horizon that spans hundreds of training iterations. Trajectory-based histories, on the other hand, are typically limited to only a few iterations, as longer histories would generate lengthy sequences that even state-of-the-art transformer-based encoders find difficult to encode (Dance

et al., 2021). The limited horizon fails to capture the full status of the student. For instance, consider the REINFORCE student with a linearly-decayed learning rate (Eq. 2). With only a few trajectories, the teacher can, at best, determine the current policy θ_t but not the learning rate α_t the student is using to learn its policy. This highlights the need to observe the changes in the student’s behavior over an extended period rather than just a limited horizon.

Formally, we define a meta-teacher f as a mapping from a context in \mathcal{C} and a target task in \mathcal{M} to a probability distribution over controllable parameter space \mathcal{W} . With the teacher being defined, the training task m_t in Eq. 1 can be selected by sampling a controllable parameter w_t from the distribution as follows:

$$w_t \sim f(c_t, m_{target}) \quad (3)$$

3.4 THE META-CL PROBLEM

To evaluate a meta-teacher function f ’s performance when teaching a given target task m_{target} over a student’s lifetime, we measure the return of the student ψ_N on the target task trained after an entire lifetime of N training iterations, where ψ_t denotes the student’s status at the t -th iteration. We use $\mathcal{G}(\psi_N, m_{target})$ to denote such an evaluated return. Then, a meta-CL problem can formally be defined as finding an optimal teacher such that the expectation of students’ evaluated returns over a distribution of target tasks $p(m)$ is maximized:

$$\begin{aligned} \max_f \quad & \mathbb{E}_{m_{target} \sim p(m)} [\mathcal{G}(\psi_N, m_{target})] \\ \text{s.t.} \quad & w_t \sim f(c_t, m_{target}) \quad \forall t = 0, \dots, N - 1 \\ & \tau_t = \text{Rollout}(\psi_t, m_t), \quad \psi_{t+1} = \text{Train}(\psi_t, \tau_t) \quad \forall t = 0, \dots, N - 1 \end{aligned} \quad (4)$$

An algorithm that solves the meta-CL problem is allowed to freely sample target tasks from $p(m)$ and teaches the students on the target tasks during a meta-training phase. These teaching experiences are used to update the teacher so that the objective defined in Eq. 4 will be improved. During the meta-testing phase, the teacher is evaluated by having it train new student lifetimes on unseen target tasks sampled from the same distribution $p(m)$. In summary, the meta-CL problem involves learning an optimal teacher that facilitates the training of any number of students from the same domain. This objective places a stronger emphasis on improving sample efficiency during the meta-testing phase, than how long it takes to train such a teacher during the meta training phase.

4 APPROACH

This section describes our proposed method MM-ACL in detail.

4.1 LOCAL LEARNING PROGRESS

The meta-CL problem requires optimizing the teacher f based on an objective that can be evaluated only after an entire lifetime training of N iterations (Eq. 4). Such a long-term objective, usually referred to as maximizing *global learning progress* (Portelas et al., 2020b), is typically considered intractable in hard-exploration domains that may require a few thousand iterations to solve, even with a good curriculum.

Instead, we consider a short-term objective that maximizes *local learning progress*, which is defined as the improvement of the target task performance after training on a source task for a shorter k iterations. The local learning progress is represented as a functional form $\mathcal{P}^k(\psi_i, m_i, m_{target})$ conditioned on the student status ψ_i , the source task m_i , and the target task m_{target} , which we formally define as follows:

$$\mathcal{P}^k(\psi_i, m_i, m_{target}) = \mathcal{G}(\psi_{i+k}, m_{target}) - \mathcal{G}(\psi_i, m_{target})$$

Then, at every i -th iteration with $i = 0, k, 2k, \dots, N$, the approximated meta-CL problem is to solve a bi-level optimization problem as follows:

$$\begin{aligned} \max_f \quad & \mathbb{E}_{m_{target} \sim p(m)} [\mathcal{P}^k(\psi_i, m_i, m_{target})] \\ \text{s.t.} \quad & w_i \sim f(c_i, m_{target}) \\ & \tau_{i+t} = \text{Rollout}(\psi_{i+t}, m_i), \quad \psi_{i+i+1} = \text{Train}(\psi_{i+t}, \tau_{i+t}) \quad \forall t = 0, \dots, k - 1 \end{aligned} \quad (5)$$

4.2 LEARNING DYNAMICS MODEL

The short-term objective of maximizing the local learning progress is analogous to some MPG algorithms that optimize the meta-representations based on online validations with a short interval of k -step iterations (Xu et al., 2018). Such

algorithms only consider meta-representations that are differentiable w.r.t. the *local learning progress*. Following the formulation defined in this paper requires that a derivative $\nabla_w \mathcal{P}^k(\psi, m, m_{target})$ of the local learning progress w.r.t. some controllable parameter w can be computed.

However, MPG is not applicable in CL, as \mathcal{P}^k is non-differentiable with respect to the controllable parameters w_t that control the source tasks unless an analytic form of the transition model is given. An alternative approach is to approximate the local learning progress $\mathcal{P}^k(\psi, m, m_{target})$ with a black box function, and then choose the best controllable parameters w based on the approximated learning progress.

More specifically, we define a *learning dynamics model* $g : \mathcal{C} \times \mathcal{W} \times \mathcal{W} \times \mathcal{V} \rightarrow \mathbb{R}$ that approximates the local learning progress $\mathcal{P}^k(\psi, m, m_{target})$ with a function $g(c, w, w_{target}, v_{target})$. Here, the dependence of the student status ψ is replaced with the contextual information c as discussed in Sec. 3.3, and the task m is replaced with controllable and uncontrollable parameters (w, v) . Once a learning dynamics model $g(c, w, w_{target}, v_{target})$ is learned, a teacher picks source tasks by sampling the controllable parameters w with the probabilities proportional to the local learning progress on the target task.

4.3 MM-ACL




Figure 2: (a) the neural network architecture of the learning dynamics model with FCs representing the fully connected layers; (b) The meta-training pipeline of MM-ACL. More details of the pipeline are presented in Sec. 4.3

With the concept of learning dynamics model, we present **Model-based Meta Automatic Curriculum Learning (MM-ACL)**, an offline meta-CL algorithm that uses a learning dynamics model with contextual information to generate curricula that facilitate student learning.

In the meta-learning literature, a task-agnostic setting is typically used, where the meta-learner is required to generalize over all possible in-distribution meta-tasks without being given the meta-task as an input. We consider a similar setting for MM-ACL, where the learning dynamics model g must generalize to all possible target tasks without explicitly knowing the uncontrollable parameters v_{target} . However, to effectively teach the student, it is still necessary to provide the controllable parameters w_{target} . Consequently, MM-ACL removes the dependence on v_{target} from the learning dynamics model. Despite this, contextual information can still encode some of the uncontrollable parameter information from the historical returns, allowing for generalization.

MM-ACL uses a neural network parameterized by ϕ to approximate the learning dynamics model, which we denote as $g_\phi(c, w, w_{target})$. The neural network architecture is shown in Fig. 2 (a). Each task-return pair in the contextual information is embedded by one fully-connected layer and is fed into a LSTM (Gers et al., 2000) as a sequence of task-return embeddings. Then, the output from the LSTM is concatenated with the embedding of w and w_{target} followed by a MLP to compute the prediction of local learning progress $\mathcal{P}^k(\psi, w, w_{target}, v_{target})$ on the target task w_{target} . More details of the learning dynamics model are shown in Appendix A.1.

Algorithm 1 shows the whole pipeline for training g_ϕ . The training data is stored in a replay buffer D collected from learners trained from scratch in multiple i.i.d. lifetimes. At every i -th iteration of each lifetime with $i = 0, k, 2k, \dots, N$, a training task w_i^{train} is selected based on the `SelectTask` function (Algorithm 2) that has an exploration probability p to select the task based on the current learning dynamics model and a probability $1 - p$ to select a randomly sampled exploratory task (lines 6-7). Once a training task w_i^{train} is selected, the student is updated for k iterations and the history of tasks and returns is appended to the contextual information (lines 8-12). Then, the policies before and after the update are evaluated on the same randomly sampled validation task w_i^{val} , resulting in the returns G_i^{val} and G_{i+k}^{val} , respectively (lines 13-16). The contextual information, training task, validation task, and learning progress constitute

Algorithm 1 MM-ACL

```

1: Input: Horizon  $k$ , RL algorithm  $Train$ , meta-CL domain  $\mathcal{M} = \mathcal{W} \times \mathcal{V}$ , target task distribution  $p(m)$ .
2: Output:  $\phi$ 
3: Initialize  $D \leftarrow \{\}, \phi$ 
4: for  $l = 1$  to  $NumLifetime$  do
5:   Initialize  $\psi_0, c_0 \leftarrow \emptyset, (w_{target}, v_{target}) \sim p(m), t = 0$ .
6:   while  $i < NumIteration$  do
7:      $w_i^{train} \leftarrow SelectTask(c_t, w_{target}, \phi)$ 
8:     for  $t = 0$  to  $k$  do
9:        $\tau_{i+t} \leftarrow Rollout(\psi_{i+t}, (w_i^{train}, v_{target}))$ 
10:       $\psi_{i+t+1} \leftarrow Train(\psi_{i+t}, \tau_{i+t})$ 
11:       $\mathcal{G}_{i+t}^{train} \leftarrow \mathcal{G}(\tau_{i+t})$ 
12:       $c_{i+t+1} \leftarrow c_{i+t} \cup (w_i^{train}, \mathcal{G}_{i+t}^{train})$ 
13:     end for
14:     for  $j = 1$  to  $NumValidationTask$  do
15:        $w_i^{val} \leftarrow GetRandomTask(\mathcal{W})$ 
16:        $\mathcal{G}_{i+k}^{val} \leftarrow \mathcal{G}(\psi_{i+k}, w_i^{val}, v_{target})$ 
17:        $\mathcal{G}_i^{val} \leftarrow \mathcal{G}(\psi_i, w_i^{val}, v_{target})$ 
18:        $D \leftarrow D \cup (c_t, w_i^{train}, w_i^{val}, \mathcal{G}_{i+k}^{val} - \mathcal{G}_i^{val})$ 
19:     end for
20:     if  $(i + l * NumIteration) \bmod UpdateInterval == 0$  then
21:        $\phi \leftarrow UpdateModel(D, \phi, NumEpoch)$ 
22:     end if
23:      $i \leftarrow i + k$ 
24:   end while
25: end for

```

Algorithm 2 SelectTask

```

1: Input: Contextual information  $c$ ,  $w_{target}$ , learning dynamics model  $\phi$ 
2: Output: Selected task  $w$ 
3: Initialize:  $p \leftarrow Uniform(0, 1)$ 
4: if  $p < ExplorationProbability$  then
5:    $w \leftarrow GetRandomTask(\mathcal{W})$ 
6: else
7:   Sample  $NumRandomTask$  tasks uniformly and randomly, call it  $\mathcal{W}^\dagger$ .
8:    $\forall i \in \mathcal{W}^\dagger, P[i] \leftarrow Softmax(g_\phi^k(c_t, w, w_{target}))[i]$ .
9:   Sample  $w \sim P$ .
10: end if

```

a training data point $(c_i, w_i^{train}, w_i^{val}, \mathcal{G}_{i+k}^{val} - \mathcal{G}_i^{val})$, which is added into the replay buffer D (line 17). At every fixed interval of iterations, the parameter ϕ is updated based on the training data in the replay buffer D (lines 19-21). Once the learning dynamics model is learned, the model can be deployed online following the `SelectTask` function with exploration probability $p = 0$. During the deployment, the contextual information c_i is built in the same way as line 12 in Algorithm 1. Notice that during deployment, lines 13 to 21 in Algorithm 1 that perform online validations are not executed; this omission guarantees that the algorithm does not add extra rollouts and gradient updates. A diagram of the meta-training pipeline is shown in Fig. 2 (b).

5 EXPERIMENTAL RESULTS

We test MM-ACL in a grid-world domain and a visual-based navigation domain. The domains are designed such that some form of curriculum is necessary, as training on the target task directly will always fail. We introduce the domains in detail in their corresponding sections. To evaluate the performance of ACL algorithms in these domains, we use a metric of *iterations to threshold*, which is formally defined as follows:

Definition 1 (Iterations to threshold) *Given a completed lifetime with a target task m_{target} and a sequence of student statuses $(\psi_i)_{i=1}^N$, the iterations to threshold (ItT) is defined as the minimum index i such that the student's past- K evaluated returns on the target task are all greater than or equal to a certain threshold δ with K being a history length. Denoting $\xi = \{i \mid i \in [K, N], \mathcal{G}(\psi_{i-j}, m_{target}) \geq \delta, \forall j \in [0, K]\}$,*

$$ItT = \begin{cases} \min_i \{i \in \xi\} & \xi \neq \emptyset \\ N & \text{otherwise.} \end{cases} \quad (6)$$

An algorithm with a low ItT is said to have a high sample efficiency, and vice versa.

The performance of MM-ACL is compared with four different baselines: (1) ALP-GMM; (2) a random curriculum; (3) a handcrafted curriculum; and, (4) a DQN-based method for solving CMDPs (as done by Narvekar & Stone (2020)), denoted as DQN-CMDP. Note that ALP-GMM is designed for domains of continuous control. To apply it in discrete domains, we modified the bandit algorithm from Matisen et al. (2019) but used the same heuristic as ALP-GMM. For

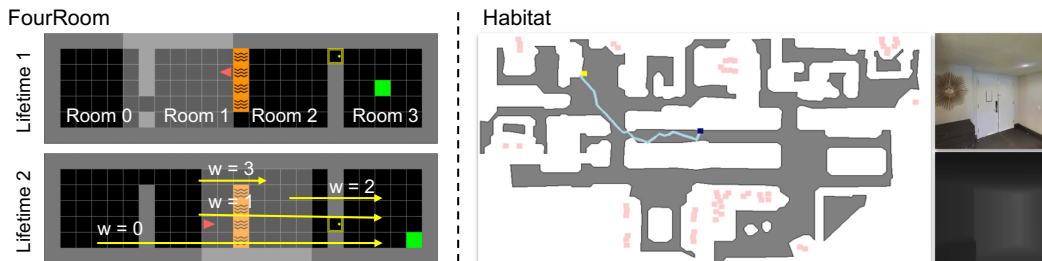


Figure 3: (Left) examples of **FourRoom** environments from two different lifetimes. The yellow arrows indicate the start and goal rooms for different controllable parameters w . (Right) a top-down-view image of **Habitat** environment with an agent’s trajectory of finding a chair in an environment. The blue dot representing the start location, the light blue curve tracing the path taken by the agent until reaching within 0.5m of a chair, and the pink dots marking the locations of chairs. The agent observes the RGB and depth images as shown at the right.

simplicity we refer to this customized bandit algorithm as ALP-GMM as well. In addition, due to the large amount of data required to train DQN-CMDP, we only compare with it in the grid-world domain. In all the experiments, the students are PPO agents (Schulman et al., 2017) with their policy neural networks and hyper-parameters tuned for each of the domains, as detailed in Appendix B.

5.1 FOURROOM DOMAIN

We test MM-ACL in the grid-world **FourRoom** domain based on the widely used MiniGrid environment by Chevalier-Boisvert et al. (2018). As shown in Figure 3, the target tasks in this domain require an agent to navigate through four rooms (labeled as room 0, 1, 2, and 3 from left to right) and reach a goal location (green tiles) in the right-most room (room 3) by executing one of the four discrete actions: move forward, turn left, turn right, and open the door. Three types of skills are required to pass the rooms: (1) navigate around the blocks (grey tiles) which the agent cannot overlap with; (2) avoid tiles covered with lava (orange tiles) that cause immediate failure of an episode once the agent visit them; (3) open the door that blocks the way to the goal location. The only positive reward of $1 - 0.0009 * T$ is assigned upon reaching the goal location, where T is the total number of time steps. In this domain, a threshold $\delta = 0.92$ and a history length $K = 4$ are used for measuring the ItT.

As shown in Fig. 3 on the left, different tasks in this domain may have different room configurations (e.g., openings on the walls) and different start-goal locations. These two types of variations correspond to the uncontrollable space \mathcal{V} and controllable space \mathcal{W} respectively. We assume that the start-goal locations can only be varied in a limited way: w can only specify the rooms from which the start and goal locations are randomly sampled. More specifically, the controllable space $\mathcal{W} = \{0, 1, 2, 3\}$ is a four-category discrete set with w equal to 0, 1, 2, 3 denoting the tasks of navigating from room 0 to 3, room 1 to 3, room 2 to 3, and room 1 to 2, respectively. The uncontrollable space \mathcal{V} is all the possible room configurations. Different lifetimes in this domain have the target tasks with the same $w_{target} = 0$, but different v_{target} uniformly sampled from \mathcal{V} .

We applied MM-ACL to the **FourRoom** domain. The *learning dynamics model* of MM-ACL is trained for 150 lifetimes during the meta-training phase. Then, the model is tested in five different lifetimes with previously unseen target tasks. The hyper-parameters of MM-ACL can be found in Appendix A.2. During meta-testing, each lifetime is repeated for five independent runs with different random seeds. While training a student, the return of the student is evaluated on the target task for 16 episodes after every training iteration. Fig. 4 (left) shows plots of return on the target task at different training iterations collected from two meta-testing lifetime. As shown in the figure, MM-ACL achieves the best sample efficiency with the lowest ItTs in both lifetimes, followed by ALP-GMM whose ItT is 29 more. On the other hand, both the random and handcrafted curricula fail to reach the threshold for one or a few runs with a 400-iteration budget. The average ItTs over all five lifetimes are reported in Table 1.

5.2 VISUAL-BASED NAVIGATION DOMAIN: HABITAT

We also test MM-ACL in a more complex visual-based navigation **Habitat** domain realized by AI Habitat (Savva et al., 2019), illustrated in Fig. 3 on the right. AI Habitat can simulate navigation in a variety of household environments, constructed from real-world houses (Chang et al., 2017), while providing photo-realistic camera images

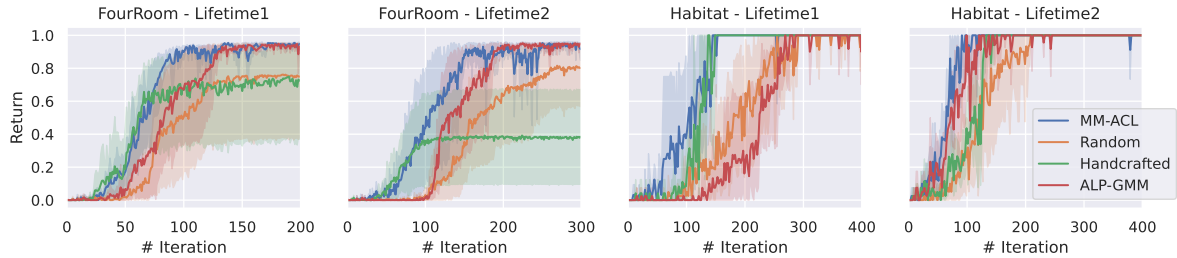


Figure 4: Average returns in the **FourRoom** and **Habitat** domains; the shaded areas show the 95% confidence interval.

as observations. We focus on an object-oriented navigation task that requires the agent to navigate to a given type of object in the environment. At each time step, based on an observation of RGB-D image and target object index, the agent can take one of three actions: move forward by 0.25 meter; turn left by 30° , or turn right by 30° . The episode is considered successful and a reward of +1 is obtained if the agent’s distance to the closet target object is less than 0.5 meter. In this domain, a threshold $\delta = 1$ and a history length $K = 4$ are used to measure the ItT.

During each lifetime in the **Habitat** domain, the agent has to search for one fixed type of object in the same environment. The uncontrollable space \mathcal{V} is, therefore, all the possible object-environment pairs. The controllable parameter $w \in \mathcal{W}$ is a variable that specifies the *geodesic distance* to the closest target object, i.e., given a parameter value w , the teacher samples a starting location for the agent such that the shortest path between the agent and the closest object is exactly w . Lastly, the target tasks from different lifetimes have a random v_{target} sampled uniformly from \mathcal{V} , and w_{target} is the largest possible *geodesic distance* to the target object in the environment specified by v_{target} .

Similarly to **FourRoom**, we apply MM-ACL to this domain with the *learning dynamics model* trained over 80 lifetimes. Then, the model is deployed in five lifetimes, each repeated for five independent runs. When training a student, the return of the student is evaluated on the target task for eight episodes after every two training iterations. Fig. 4 (right) shows two plots of the return on the target task at different training iterations collected from two meta-testing lifetimes. As seen in the figure, MM-ACL achieves the best sample efficiency in both lifetimes, while both ALP-GMM and the handcrafted curriculum perform sub-optimally in one of the lifetimes.

	FourRoom	Habitat
MM-ACL	129 ± 25	82 ± 21
DQN-CMDP	202 ± 39	-
ALP-GMM	158 ± 28	105 ± 32
Handcrafted	175 ± 84	102 ± 18
MM-ACL-single	175 ± 73	-
Random	213 ± 46	126 ± 54

Table 1: Average ItTs in **FourRoom** and **Habitat**.




Figure 5: The probability distributions used by MM-ACL (left) and ALP-GMM (right) to sample source tasks at different training iterations in the same meta-testing lifetime.

6 ANALYSIS

In this section, we conduct an analysis and ablation study. We focus our attention on the grid-world domains due to the lower computational cost compared to **Habitat**.

What curriculum does MM-ACL generate? To study why MM-ACL achieves lower sample efficiency, in Fig. 5, we plot the probability distributions used by MM-ACL and ALP-GMM to sample the tasks at different training iterations. As seen in the figure, MM-ACL focuses equally on $w = 3$ and $w = 2$ at the beginning, which correspond to the two basic skills of avoiding lava and opening the door. Then it switches to $w = 1$ that combines the two skills in the same tasks. Finally, MM-ACL turns to the target task $w = 0$ after about 50 iterations, when the target task is learnable. On

the other hand, ALP-GMM heavily exploits $w = 2$ at the beginning since it gives the highest absolute learning progress, which actually slows down the learning of the target task.




Figure 6: (left) the sample efficiencies of MM-ACL (blue) trained for 10, 50, 100, and 150 lifetimes, with the red solid line marking the sample efficiency of ALP-GMM, and the dashed lines marking its 95% confidence interval; (right) evaluation of the target task return of the students taught by MM-ACL and a fixed curriculum generated by MM-ACL.

How many lifetimes does MM-ACL need? To answer this question, we train MM-ACL with limited data collected from 10, 50, 100 and 150 lifetimes, and test it in the **FourRoom** domain, similarly to Sec. 5.1. The results are shown in Fig. 6 (left), in which the sample efficiency of ALP-GMM is marked as the red line. MM-ACL is outperformed by ALP-GMM with only ten lifetimes, but is comparable to ALP-GMM after 50 lifetimes and clearly outperforms it after 100 lifetimes.

Does MM-ACL generalize to different student statuses? As discussed in Sec. 3.3, a meta-teacher has to infer the student status from the contextual information if it does not have access to the full representation of the student status (e.g. the weights of the neural network). To study the ability of MM-ACL to infer and adapt to the current status of the student, we record a fixed history curriculum generated by MM-ACL in one of the runs, and deploy it on the same target task, but with five different random seeds that may lead to different student statuses. As shown in Fig. 6 (right), the curve of evaluated returns of MM-ACL is above the curve of the fixed curriculum throughout the training and shows much lower variance. On the other hand, the fixed curriculum shows higher variances which indicates a failure to adapt to different student statuses.

Does MM-ACL benefit from experiencing different target tasks during training? To verify that MM-ACL can improve its generalization to unseen target tasks from experiencing more target tasks during training, we conduct an ablation study that trains MM-ACL with the same number of 150 lifetimes, with all the lifetimes having the same *single* target task. We denote the resulting model as MM-ACL-single, which is then evaluated in the same way as MM-ACL in Sec. 5.1. The results in Table 1 (rows 1 and 6) show that MM-ACL-single achieves a lower sample efficiency of 175 ItT compared to 129 by MM-ACL. Notice that this sample efficiency is similar to that of a handcrafted curriculum (row 5 in Table 1) designed for a few target tasks, which further supports the necessity of generalizing teachers to different target tasks.

7 CONCLUSION

This paper first formulates a new meta-CL problem, in which a teacher can benefit from previous lifetimes' teaching experience to be able to teach new students and new target tasks. As an initial approach to solving this problem, a new meta-CL algorithm MM-ACL is proposed that learns a *learning dynamics model* that predicts the improvement on one task when trained on another, which can be utilized to select the source task that maximizes the improvement on the target task. Our empirical results show an improvement of sample efficiency compared to the state-of-the-art ACL algorithms and existing meta-CL algorithms.

However, MM-ACL suffers from myopia that might lead to failures in cases in which training on multiple source tasks is required before observing any improvement in the student's performance on the target task. In addition, we only tested one type of contextual information based on the history tasks and returns; more informative contextual information based on the trajectories could also be considered in future work.

REFERENCES

- Angel Chang, Angela Dai, Thomas Funkhouser, Maciej Halber, Matthias Niessner, Manolis Savva, Shuran Song, Andy Zeng, and Yinda Zhang. Matterport3d: Learning from rgb-d data in indoor environments. *International Conference on 3D Vision (3DV)*, 2017.
- Maxime Chevalier-Boisvert, Lucas Willems, and Suman Pal. Minimalistic gridworld environment for openai gym. <https://github.com/maximecb/gym-minigrid>, 2018.

- Christopher R Dance, Julien Perez, and Théo Cachet. Conditioned reinforcement learning for few-shot imitation. In *International Conference on Machine Learning*, pp. 2376–2387. PMLR, 2021.
- Michael Dennis, Natasha Jaques, Eugene Vinitzky, Alexandre Bayen, Stuart Russell, Andrew Critch, and Sergey Levine. Emergent complexity and zero-shot transfer via unsupervised environment design. *Advances in neural information processing systems*, 33:13049–13061, 2020.
- Yang Fan, Fei Tian, Tao Qin, Xiang-Yang Li, and Tie-Yan Liu. Learning to teach. *arXiv preprint arXiv:1805.03643*, 2018.
- Carlos Florensa, David Held, Xinyang Geng, and Pieter Abbeel. Automatic goal generation for reinforcement learning agents. In *International conference on machine learning*, pp. 1515–1528. PMLR, 2018.
- Felix A Gers, Jürgen Schmidhuber, and Fred Cummins. Learning to forget: Continual prediction with lstm. *Neural computation*, 12(10):2451–2471, 2000.
- Timothy Hospedales, Antreas Antoniou, Paul Micaelli, and Amos Storkey. Meta-learning in neural networks: A survey. *arXiv preprint arXiv:2004.05439*, 2020.
- Shin Ishii, Wako Yoshida, and Junichiro Yoshimoto. Control of exploitation–exploration meta-parameter in reinforcement learning. *Neural networks*, 15(4-6):665–687, 2002.
- Lu Jiang, Zhengyuan Zhou, Thomas Leung, Li-Jia Li, and Li Fei-Fei. Mentornet: Learning data-driven curriculum for very deep neural networks on corrupted labels. In *International conference on machine learning*, pp. 2304–2313. PMLR, 2018.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Dzmitry Bahdanau Maxime Chevalier-Boisvert Lucas Willems, Ye Yuan. RL starter files. <https://github.com/lcswillems/rl-starter-files>, 2013.
- Tambet Matiisen, Avital Oliver, Taco Cohen, and John Schulman. Teacher–student curriculum learning. *IEEE transactions on neural networks and learning systems*, 31(9):3732–3740, 2019.
- Sanmit Narvekar and Peter Stone. Learning curriculum policies for reinforcement learning. In *Proceedings of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, May 2019.
- Sanmit Narvekar and Peter Stone. Generalizing curricula for reinforcement learning. In *4th Lifelong Learning Workshop at the International Conference on Machine Learning (ICML 2020)*, July 2020.
- Sanmit Narvekar, Bei Peng, Matteo Leonetti, Jivko Sinapov, Matthew E. Taylor, and Peter Stone. Curriculum learning for reinforcement learning domains: A framework and survey. *Journal of Machine Learning Research*, 21(181): 1–50, 2020. URL <http://jmlr.org/papers/v21/20-212.html>.
- Rémy Portelas, Cédric Colas, Katja Hofmann, and Pierre-Yves Oudeyer. Teacher algorithms for curriculum learning of deep rl in continuously parameterized environments. In *Conference on Robot Learning*, pp. 835–853. PMLR, 2020a.
- Rémy Portelas, Cédric Colas, Lilian Weng, Katja Hofmann, and Pierre-Yves Oudeyer. Automatic curriculum learning for deep rl: A short survey. *arXiv preprint arXiv:2003.04664*, 2020b.
- Rémy Portelas, Clément Romac, Katja Hofmann, and Pierre-Yves Oudeyer. Meta automatic curriculum learning. *arXiv preprint arXiv:2011.08463*, 2020c.
- Clément Romac, Rémy Portelas, Katja Hofmann, and Pierre-Yves Oudeyer. Teachmyagent: a benchmark for automatic curriculum learning in deep rl. In *International Conference on Machine Learning*, pp. 9052–9063. PMLR, 2021.
- Terence D Sanger. Neural network learning control of robot manipulators using gradually increasing task difficulty. *IEEE transactions on Robotics and Automation*, 10(3):323–333, 1994.
- Manolis Savva, Abhishek Kadian, Oleksandr Maksymets, Yili Zhao, Erik Wijmans, Bhavana Jain, Julian Straub, Jia Liu, Vladlen Koltun, Jitendra Malik, Devi Parikh, and Dhruv Batra. Habitat: A Platform for Embodied AI Research. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019.

- John Schulman, Philipp Moritz, Sergey Levine, Michael Jordan, and Pieter Abbeel. High-dimensional continuous control using generalized advantage estimation. *arXiv preprint arXiv:1506.02438*, 2015.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- Emmanouil Stergiadis, Priyanka Agrawal, and Oliver Squire. Curriculum meta-learning for few-shot classification. *arXiv preprint arXiv:2112.02913*, 2021.
- Richard S Sutton. Adapting bias by gradient descent: An incremental version of delta-bar-delta. In *AAAI*, pp. 171–176. San Jose, CA, 1992.
- Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.
- Vivek Veeriah, Matteo Hessel, Zhongwen Xu, Janarthanan Rajendran, Richard L Lewis, Junhyuk Oh, Hado P van Hasselt, David Silver, and Satinder Singh. Discovery of useful questions as auxiliary tasks. *Advances in Neural Information Processing Systems*, 32, 2019.
- Rui Wang, Joel Lehman, Aditya Rawal, Jiale Zhi, Yulun Li, Jeffrey Clune, and Kenneth Stanley. Enhanced poet: Open-ended reinforcement learning through unbounded invention of learning challenges and their solutions. In *International Conference on Machine Learning*, pp. 9940–9951. PMLR, 2020.
- Tongtong Wu, Xuekai Li, Yuan-Fang Li, Gholamreza Haffari, Guilin Qi, Yujin Zhu, and Guoqiang Xu. Curriculum-meta learning for order-robust continual relation extraction. In *AAAI*, pp. 10363–10369. AAAI Press, 2021.
- Zhongwen Xu, Hado P van Hasselt, and David Silver. Meta-gradient reinforcement learning. *Advances in neural information processing systems*, 31, 2018.
- Zeyu Zheng, Junhyuk Oh, and Satinder Singh. On learning intrinsic rewards for policy gradient methods. *Advances in Neural Information Processing Systems*, 31, 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.

A MM-ACL

A.1 LEARNING DYNAMICS MODEL

A data point that trains the learning dynamics model is a 5-tuple $(c_i, w_i^{train}, w_i^{val}, \mathcal{G}_i^{val}, \mathcal{G}_{i+k}^{val})$, where c_i is the context at the t -th training iteration, \mathcal{G}_i^{val} and \mathcal{G}_{i+k}^{val} are the returns evaluated on the task w_i^{val} before and after training on the task w_i^{train} for k iterations. Then, the learning dynamics model $g_\phi(c_i, w_i^{train}, w_i^{val})$, parameterized by ϕ , is trained to predict the learning progress $\mathcal{G}_{i+k}^{val} - \mathcal{G}_i^{val}$ given an input of $(c_i, w_i^{train}, w_i^{val})$. More specifically, the parameters ϕ is optimized to minimize the mean square error of the prediction of learning progress evaluated over all the data points in a replay buffer D , or:

$$\min_{\phi} \mathbb{E}_{(c_i, w_i^{train}, w_i^{val}, \mathcal{G}_i^{val}, \mathcal{G}_{i+k}^{val}) \sim D} [(\mathcal{G}_{i+k}^{val} - \mathcal{G}_i^{val} - g_\phi(c_i, w_i^{train}, w_i^{val}))^2]$$

In addition, since the evaluated return \mathcal{G}_i^{val} is available in each data point, we use the the same parameter ϕ with an additional head of two fully-connected layers to perform an auxiliary task that predicts \mathcal{G}_i^{val} . The final loss function is the linear combination of the two losses, which can be written as follows:

$$\min_{\phi} \mathbb{E}_{(c_i, w_i^{train}, w_i^{val}, \mathcal{G}_i^{val}, \mathcal{G}_{i+k}^{val}) \sim D} [(\mathcal{G}_{i+k}^{val} - \mathcal{G}_i^{val} - g_\phi(c_i, w_i^{train}, w_i^{val}))^2 + \alpha * (\mathcal{G}_i^{val} - h_\phi(c_i, w_i^{train}, w_i^{val}))^2],$$

where α is a constant coefficient, and h_ϕ represents the function that predicts the evaluated return \mathcal{G}_i^{val} . We found that such an auxiliary task can reduce the over-fitting when the model is tested in new lifetimes.

In the experiments for both the **FourRoom** and the **Habitat** domains, we use the same neural network architecture to represent the learning dynamics model, which we detail as follows. The contextual information $c_i = \langle (w_j, \mathcal{G}_j^{train}) \rangle_{j=1}^i$ is a sequence of task-return pairs. The model first feeds each task-return pair into a 16-dimensional linear embedding layer. Then, the output of the sequence of embeddings is fed into two 64-hidden-size layers of LSTMs. The resulting 64-dimensional contextual embedding is concatenated with two 16-dimensional embeddings that embeds w_i^{train} and w_i^{val} respectively. Lastly, the 96-dimensional feature embedding is fed into two fully-connected layers with 64 and 32 hidden units, which outputs the learning progress prediction. In addition, another two 64×32 layers are created to predict the evaluated return. We used the constant coefficient $\alpha = 0.1$, which is chosen from grid search that minimizes the loss on hold-out validation data points. The model is optimized by Adam optimizer Kingma & Ba (2014) with an learning rate of 0.0001 and a batch size of 512. We used a function `UpdateModel($D, \phi, NumEpoch$)` that will learn the model on replay buffer D for $NumEpoch$ epochs with the aforementioned hyper-parameters.

A.2 HYPER-PARAMETERS

The hyper-parameters that specify the learning dynamics model are described in Appendix A.1. In the table below, we list the hyper-parameters that specify the meta-training process.

	FourRoom	Habitat
<i>NumLifetime</i>	150	120
<i>NumIteration</i>	300	500
<i>NumValidationTask</i>	4	4
<i>Horizon k</i>	1	2
<i>UpdateInterval</i>	45000	20000
<i>NumEpoch</i>	100	100
<i>ExplorationProbability</i>	1	[0.9, 0.5, 0.1]

Note that in the **FourRoom** domain, *UpdateInterval* is set to 450000 that only updates the learning dynamics model once with the data collected purely from random curriculum. This works for the **FourRoom** domain with relatively simple controllable parameter space and can speed up the training by parallelizing the training of multiple lifetimes. In the **Habitat** domain, the *ExplorationProbability* is decreased by 0.4 at every 40 lifetimes when the model is updated for 100 epochs.

B BASELINES

B.1 STUDENT (PPO)

We used a customized PPO implementation [Schulman et al. \(2017\)](#) for both **FourRoom** domain and **Habitat** domain. We detail the neural network architectures and hyper-parameters as follows.

For **FourRoom** domain, the observation is composed of a three-channel seven-by-seven image with the channels encode color, object type and object status respectively. The image is fed into a CNN neural network of sequential layers: **Input** \rightarrow **Conv2D**[16, 2, 1] \rightarrow **MaxPool**[2] \rightarrow **Conv2D**[32, 2, 1] \rightarrow **Conv2D**[64, 1, 1] \rightarrow **LSTM**[128] \rightarrow **FC**[128], where **FC** is the fully connected layers with the number denoting the number of hidden units, **LSTM** is one layer of LSTM with the number denoting hidden size, and **Conv2D** is the 2D convolution neural network with the numbers denoting the number of channels, kernel size, and stride respectively. All the layers are followed by the ReLU activation. This design of the architecture is modified from an RL implementation [Lucas Willems \(2013\)](#) that is known to solve the MiniGrid environment.

For **Habitat** domain, the observation is composed of a four-channel 84-by-84 RGB-D image. The image is fed into a CNN neural network of sequential layers: **Input** \rightarrow **Conv2D**[32, 8, 4] \rightarrow **Conv2D**[64, 4, 2] \rightarrow **Conv2D**[64, 2, 2] \rightarrow **FC**[128] \rightarrow **FC**[128], where **FC** is the fully connected layers with the number denoting the number of hidden units, and **Conv2D** is the 2D convolution neural network with the numbers denoting the number of channels, kernel size, and stride respectively. All the layers are followed by the ReLU activation.

Other PPO related hyper-parameters are listed as follows:

PPO	FourRoom	Habitat
Discount factor	0.9	0.9
Ratio clip	0.2	0.2
Entropy coefficient	0.01	0.01
GAE Schulman et al. (2015) coefficient	0.97	0.97
Batch size	256	256
Learning rate	0.0001	0.0001
Number of epochs per update	4	4
Number of steps between updates	2000	8000

B.2 ALP-GMM

The core idea of ALP-GMM is to fit a Gaussian Mixture Model (GMM) every n episodes on recently sampled tasks *concatenated with their respective absolute learning progress (ALP) value*. The ALP is defined as the return difference of the current sampled task and the same task sampled in the history. Since it is impossible to sample exactly the same task twice for continuous task space, ALP-GMM computes a per-task ALP from the entire history of sampled tasks using a knn-based approach. The Gaussian from which to sample a new task is then chosen proportionally to its respective learning progress. The original ALP-GMM is applied to the **Habitat** domain with the hyper-parameters listed as follows:

ALP-GMM (continuous)	Habitat	ALP-GMM (discrete)	FourRoom
Fit rate	100	Buffer size	50
Buffer size	200	Random task ratio	0.2
Random task ratio	0.2		
Number of bootstrap episodes	50		
Number of mixture models	[2, 3, 4, 5]		

For the **FourRoom** domain with discrete task space, instead of fitting a GMM, ALP-GMM maintains one replay buffer of the returns for each of the tasks. The ALP is measured by the return difference between the current task and the last sampled task in the same replay buffer. Then, the tasks are sampled proportional to the their average ALPs in the replay buffers. The related hyper-parameters are as follows:

B.3 DQN-CMDP

DQN-CMDP models the curriculum learning problem as Markov Decision Process (MDP) and solves the MDP with a standard DQN algorithm. More specifically, the action space of such an MDP is the uncontrollable task space \mathcal{W} , and the state space is the space of contextual information \mathcal{C} . An episode of this MDP is an entire lifetime that terminates when the evaluated return reaches the threshold (the same criteria used to compute the *iterations to threshold*). The teacher receives a +10 reward only when the evaluated return reaches the threshold.

We only applied DQN-CMDP to the **FourRoom** domain that uses discrete task space. We used the same encoder as described in Appendix A.1 to encode the contextual information. Then, the extracted feature embedding is fed into a 64×32 fully-connected layers to output a value prediction. Such a DQN policy is trained for the same number of 150 lifetimes (episodes) for comparison.

B.4 HANDCRAFTED CURRICULUM

The handcrafted curricula are empirically created based on their performances on the target tasks that are different from the five hold-out target tasks used for the meta-testing. For the **FourRoom** domain, the handcrafted curriculum trains a student on controllable parameter $w = 2$, $w = 3$, $w = 1$, and $w = 0$ for 50k, 50k, 50k, and 100k time steps respectively. For the **Habitat** domain, the handcrafted curriculum trains a student on controllable parameter $w = 4.5$, $w = 5.75$, $w = 7$, $w = 8.25$, $w = 9.5$, $w = 10.75$, and $w = 12$ with each trained for 200k steps, and finally the target task trained until the end of the lifetime.