TRANSFERRING HIERARCHICAL STRUCTURE WITH DUAL META IMITATION LEARNING

Anonymous authors
Paper under double-blind review

ABSTRACT

Hierarchical Imitation learning (HIL) is an effective way for robots to learn sub-skills from long-horizon unsegmented demonstrations. However, the learned hierarchical structure lacks the mechanism to transfer across multi-tasks or to new tasks, which makes them have to learn from scratch when facing a new situation. Transferring and reorganizing modular sub-skills require fast adaptation ability of the whole hierarchical structure. In this work, we propose Dual Meta Imitation Learning (DMIL), a hierarchical meta imitation learning method where the high-level network and sub-skills are iteratively meta-learned with model-agnostic meta-learning. DMIL uses the likelihood of state-action pairs from each sub-skill as the supervision for the high-level network adaptation, and use the adapted high-level network to determine different data set for each sub-skill adaptation. We theoretically prove the convergence of the iterative training process of DMIL and establish the connection between DMIL and the Expectation-Maximization algorithm. Empirically, we achieve state-of-the-art few-shot imitation learning performance on the meta-world (Yu et al., 2019b) benchmark.

1 INTRODUCTION

Imitation learning (IL) has shown promising results for intelligent robots to conveniently acquire skills from expert demonstrations (Zhu et al., 2018; Peng et al., 2018). Nevertheless, acquiring skills from long-horizon unsegmented demonstrations has been a challenge for IL algorithms, since the well-known issue of compounding errors (Ross et al., 2011). This is one of the crucial problems for the application of IL methods to robots since plenty of manipulation tasks are long-horizon. Hierarchical Imitation Learning (HIL) aims to tackle this problem by decomposing long-horizon tasks with a hierarchical model, in which a set of sub-skills are employed to accomplish specific parts of the long-horizon task, and a high-level network is responsible for determining the switching of sub-skills along with the task. Such a hierarchical structure is usually modeled with Options (Daniel et al., 2016; Krishnan et al., 2017; Jing et al., 2021) or goal-conditioned structure (Le et al., 2018). HIL expresses the nature of how humans solve complex tasks, and has been considered to be a valuable direction for IL algorithms (Osa et al., 2018).

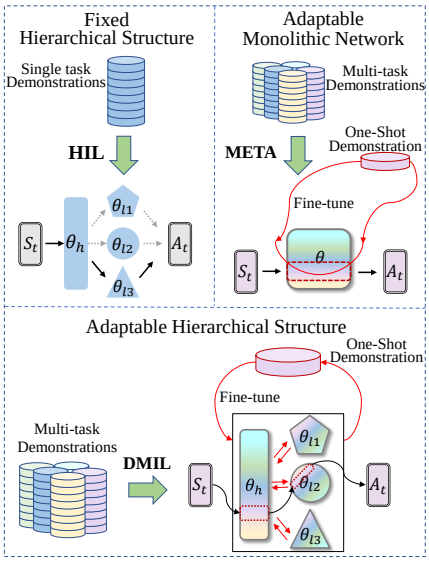

Figure 1: DMIL meta-learns the high-level network and sub-skills from multi-task demonstrations with mutual supervision, and can transfer the hierarchical structure with few-shot demonstration.

However, most current HIL methods have no mechanism to transfer and reuse previously learned skills in new tasks. This makes them have to learn the hierarchical structure from scratch, which makes a fair amount of new demonstrations have to be provided for each new task. Transferring learned skills can significantly improve data efficiency and reduce the supervision for each new task.

Transferring the learned hierarchical structure from HIL to new tasks is a two-fold process. On the one hand, sub-skills need to be transferred, since the same sub-skill may exhibit different specific forms in different tasks. For example, grasping a cube from the top and grasping a cup from the side both belong to the sub-skill *grasping*, and pushing a block directly forward and pushing it around the wall both need the sub-skill *pushing*. On the other hand, the high-level network also needs to be transferred, since directly using the previously learned high-level network in new tasks would lead to misclassification of sub-skills. This is because the high-level network can only employ *state* information at each time step to determine how to use sub-skills in new tasks. However, the expert may demonstrate totally different behaviors at the same state in different tasks. For example, when the robot arm is over a half-open drawer, the task can be either opening or closing the drawer, which may require different sub-skills. These facts highlight the necessity of an appropriate approach to quickly transfer the learned hierarchical structure to new tasks with few-shot demonstrations.

Recently, meta imitation learning (MIL) (Yu et al., 2018b; Finn et al., 2017b; Yu et al., 2018a) employ model-agnostic meta-learning (MAML) (Finn et al., 2017a) to imitation learning procedure to enable the learned policy to quickly adapt to new tasks with few shot demonstrations. MAML first fine-tunes the policy network in the inner loop, then evaluates the fine-tuned network to update initial parameters with end-to-end gradient descent at the outer step. This inspires us to integrate MAML into HIL to transfer the hierarchical structure. However, most MIL methods meta-learn a monolithic policy in an end-to-end fashion, which not conforms to the iterative training of the bi-level structure in HIL, where the training supervision of each level comes from the other level. This difference limits us to meta-learn each level of HIL seperately, but seperately meta-learning each level in an iterative fashion brings us two challenges. Firstly, we need to design appropriate forms of supervision for the meta-learning of each level. Secondly, we need to develop a method to schedule the inner loops and outer loops of the meta-learning of two levels in HIL to ensure convergence, since in different loops, the supervision from one level may be incorrect for the other level.

In this work, we propose a novel hierarchical meta imitation learning framework called dual meta imitation learning (DMIL) to successfully incorporate MAML into the iterative training process of HIL, as shown in figure 1. For the first challenge, we make the high-level network and sub-skills as mutual supervision during the bi-level MAML process. As we will depict in this paper, the likelihood of each state-action pair in few-shot demonstrations of each sub-skill can provide the supervisions for the meta-training of the high-level network, and the fine-tuned high-level network in turn determines the data sets for the meta-training of each sub-skill. For the second challenge, we design an elaborate training procedure for DMIL that first fine-tunes the high-level network and sub-skills in sequence at the inner step, then meta-updates them simultaneously at the outer loop. We theoretically prove its convergence by leveraging previous results from Amit & Meir (2018); Ravi & Beatson (2019); Finn et al. (2018); Zou & Lu (2020) to reframe both MAML and DMIL as hierarchical Bayes inference processes to incorporate MAML into HIL to get an integrated EM-like training process.

We test our method on the challenging meta-world benchmark environments (Yu et al., 2019b). In our experiments, we successfully acquire a set of sub-skills from a large scale of manipulation tasks and achieve state-of-the-art few-shot imitation learning abilities in the ML45 suites. In summary, the main contributions of this paper are as follows:

- We propose DMIL, a novel hierarchical meta imitation learning framework that meta-learns both the high-level network and sub-skills from unsegmented multi-task demonstrations.
- We propose an EM-like training algorithm to schedule the meta-learning and hierarchical imitation learning processes in DMIL and theoretically guarantee its convergence.
- We achieve state-of-the-art few-shot imitation learning performance with DMIL on the meta-world benchmark environments.

2 Related Work

2.1 Hierarchical Imitation Learning

Recovering inherent sub-skills contained in expert demonstrations and then reuse them with hierarchical structures has long been an interesting topic in the hierarchical imitation learning (HIL)

domain. According to whether there are pretraining tasks, we can divide HIL methods into two categories. The first kind aims to manually design a set of simple pretraining tasks that could encourage distinct skills or primitives, and then they learn a high-level network to master the switching of primitives to accomplish complex tasks (Florensa et al., 2017; Peng et al., 2019; Liu & Hodgins, 2017; Merel et al., 2019a; Le et al., 2018). However, for unsegmented multi-task demonstrations where no pretraining tasks are provided, which is the situation in our paper, these methods can not be applied.

The second kind of methods aim to learn sub-skills with unsupervised learning methods. Daniel et al. (2016); Krishnan et al. (2017) try to acquire Options (Sutton et al., 1999) from demonstrations with an Expectation-Maximization-like procedure and use the Baum-Welch algorithm to estimate the parameters of different options. Henderson et al. (2018); Jing et al. (2021) integrate the idea of generative adversarial networks into option discovery process to obtain a set of options. Li et al. (2017); Sharma et al. (2019); Lee & Seo (2020) incorporate generative-adversarial imitation learning (Ho & Ermon, 2016) framework and an information-theoretic metric (Chen et al., 2016) to simultaneously imitate the expert and maximize the mutual-information between latent sub-skill categories and corresponding trajectories to acquire decoupled sub-skills. There are some methods called mixture-of-expert (MoE) that compute the weighted sum of all primitives to get the action rather than only using one of them at each time step (Hausknecht & Stone, 2016; Neumann et al., 2009; Jacobs et al., 1991). There are also other methods that aim to seek an appropriate latent space that can map sub-skills into it and then condition a policy on the latent variable to reuse sub-skills (Merel et al., 2019b; Lynch et al., 2019; Haarnoja et al., 2018; Hausman et al., 2018).

For transferring learned sub-skills, some works aim to freeze sub-skills and relearn a new high-level network in new tasks (Peng et al., 2019). However, this is not appropriate in a few-shot imitation learning setting, since relearning a new high-level network with few-shot new task demonstrations would lead to over-fitting on few-shot data, and make the high-level network lost its capacity to extract useful information of states learned from previous experience.

2.2 Meta Imitation Learning

Meta imitation learning, or one-shot imitation learning, leverages various meta-learning methods and multi-task demonstrations to meta-learn a policy that can be quickly adapted to a new task with few-shot new task demonstrations. Duan et al. (2017); Cachet et al. (2020) employ self-attention modules to process the whole demonstration and the current observation to predict current action. Yu et al. (2018a); Finn et al. (2017b); Yu et al. (2018b) use model-agnostic meta-learning (MAML) (Finn et al., 2017a) to achieve one-shot imitation learning ability for various manipulation tasks with robot or human visual demonstrations. Recently, Xu et al. (2019); Yu et al. (2019a) propose to meta-learn a robust reward function that can be quickly adapted to new tasks and then use it to perform IRL in new tasks to transfer the policy. However, they need downstream inverse reinforcement learning after the adaptation of reward functions, thus conflicts with our goal of few-shot adaptation. In addition, they assume the reward function in new tasks is similar to the ones in previously learned tasks, which restricts their application to a large scale of diverse manipulation tasks.

Most meta imitation learning methods only learn one monolithic policy, lacking the ability to model multiple sub-skills in long-horizon tasks. Although there are some works that aim to tackle the multi-modal data problem in meta-learning by avoiding single parameters initialization across all tasks (Vuorio et al., 2019; Alet et al., 2018; Frans et al., 2018; Yao et al., 2019), they lack the mechanism to schedule the switching of different sub-skills over time. There are some works that also meta-learn a set of sub-skills in a hierarchical structure (Yu et al., 2018a; Frans et al., 2018). However, they either use pretraining tasks to manually learn a set of primitives or abandon the high-level network and relearn it in new tasks, which is not appropriate in the few-shot imitation learning setting.

3 Method

3.1 Preliminaries

We denote a discrete-time finite-horizon Markov decision process (MDP) as a tuple $(\mathcal{S}, \mathcal{A}, T, P, r, \rho_0)$, where $\mathcal{S}$ is the state space, $\mathcal{A}$ is the action space, T is the time horizon,

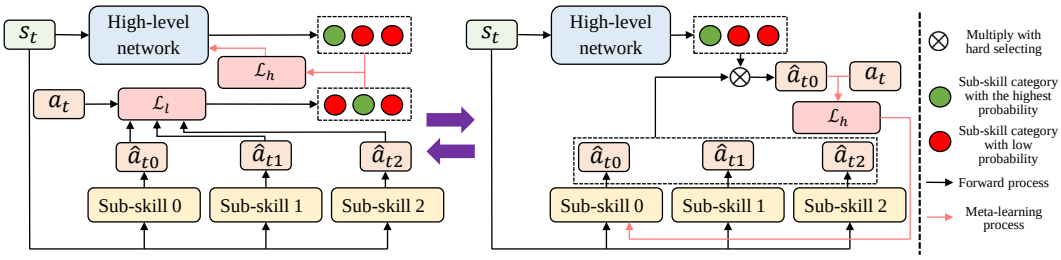

Figure 2: The iterative meta-learning process of DMIL at each time step. The fine-tuning and meta-updating step share the same framework, with different stages of the bi-level network.

$P : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \to [0, 1]$ is the transition probability distribution, $r : \mathcal{S} \times \mathcal{A} \to \mathbb{R}$ is the reward function, and ρ_0 is the distribution of the initial state s_0.

3.2 Formulation of Meta Imitation Learning Problem

We firstly introduce the general setting of the meta imitation learning problem. The goal of meta imitation learning is to extract some common knowledge from a set of robot manipulation tasks $\{\mathcal{T}_i\}$ that come from the same task distribution $p(\mathcal{T})$, and adapt it to new tasks quickly with few shot new task demonstrations. As in model-agnostic meta-learning algorithm (MAML) (Finn et al., 2017a), we formalize the common knowledge as the initial parameter θ of the policy network π_θ that can be efficiently adapted with new task gradients.

For each task $\mathcal{T}_i \sim p(\mathcal{T})$, a set of demonstrations $\mathcal{D}_i$ is provided, where $\mathcal{D}_i$ consists of N demonstration trajectories: $\mathcal{D}_i = \{\tau_{ij}\}_{j=1}^{N}$, and τ_{ij} consists of a sequence of state-action pairs: $\tau_{ij} = \{(s_t, a_t)\}_{t=1}^{T_{ij}}$, where T_{ij} is the length of τ_{ij}. Each $\mathcal{D}_i$ is randomly split into support set $\mathcal{D}_i^{tr}$ and query set $\mathcal{D}_i^{val}$ for meta-training and meta-testing respectively. During the training phase, we sample m tasks from $p(\mathcal{T})$, and in each task $\mathcal{T}_i$, we use $\mathcal{D}_i^{tr}$ to fine-tune π_θ to get the adapted task-specific parameters λ_i with gradient descent, and then evaluate it with $\mathcal{D}_i^{val}$ to get the meta-gradient of $\mathcal{T}_i$, and we optimize the initial parameters θ with the average of meta-gradients from all m tasks. As in Finn et al. (2017b); Duan et al. (2017), we use behavior cloning (Atkeson & Schaal, 1997) loss as our metrics for meta-training and meta-testing. It aims to train a policy $\pi_\theta : \mathcal{S} \to \mathcal{A}$ that maximizes the likelihood such that $\theta^* = \arg\max_\theta \sum_{i=1}^{N} \log \pi_\theta(a_i|s_i)$, where N is the number of provided state-action pairs. We denote the loss function of this optimization problem as $\mathcal{L}_{BC}(\theta, \mathcal{D})$, and the general objective of meta imitation learning problem is:

$$\min_\theta \sum_{i=1}^{m} \mathcal{L}_{BC}\left(\lambda_i, \mathcal{D}_i^{\text{val}}\right), \tag{1}$$

where $\lambda_i = \theta - \alpha \nabla_\theta \mathcal{L}_{BC}(\theta, \mathcal{D}_i^{tr})$, and α is a hyper-parameter which represents the inner-update learning rate.

3.3 Dual Meta Imitation Learning (DMIL)

In this work we assume at each time step t, the robot may switch to different sub-skills to accomplish the task. We define the sub-skill category at each time step t as $z_t = 1, \cdots, K$, where K is the maximum number of sub-skills. We assume a successful trajectory τ_{ij} of a task $\mathcal{T}_i$ is generated from several (at least one) sub-skill policies, i.e., $\tau_{ij} = \sum_{t=1}^{T_{ij}} \{(s_t, \pi_E(s_t|z_t))\}$, where π_E represents the expert policy. Our goal is to learn a hierarchical structure from multi-task demonstrations $\{\mathcal{D}_1, \cdots, \mathcal{D}_m\}$ in an unsupervised fashion. In our model, a high-level network π_{θ_h} that parameterized by θ_h determines the sub-skill category $\hat{z}_t$ at each time step t, and the *z-th* sub-skill among K different sub-skills $\pi_{\theta_{l1}}, \cdots, \pi_{\theta_{lK}}$ will be called to predict the corresponding action $\hat{a}_t$ of state s_t, where the hat symbol denotes the predicted result. We use λ_h and $\lambda_{l1}, \cdots, \lambda_{lk}$ to represent the adapted parameters of θ_h and $\theta_{l1}, \cdots, \theta_{lK}$ respectively. In order to achieve few-shot learning ability in new tasks, we condition the high-level network only on states, i.e., $\hat{z}_t = \pi_{\theta_h}(s_t)$, since at the testing phase we only have access to states and have no access to action information.

DMIL aims to first fine-tune both π_{θ_h} and $\pi_{\theta_{l0}}, \cdots, \pi_{\theta_{lK}}$ and then meta-update them. In a new task, π_{θ_h} may not provide correct sub-skill categories as stated in the introduction. However, sub-skills still retain the ability to give out supervision for the high-level network with knowledge learned from previously learned tasks. This is a reasonable assumption for robot manipulation tasks since most of them are made up of a set of basis skills like *reach, push* and *pick-place*. As shown in the left side of 2, the sub-skill that gives out the closet $\hat{a}_t$ to a_t can be seen as supervision for the high-level network to classify s_t into this sub-skill. On the other hand, the adapted high-level network can classify each datapoint in provided demonstrations to different sub-skills for them to perform fine-tuning, as shown in the right side of figure 2. In summary, DMIL contains four steps for one training iteration. We call them **High-Inner-Update (HI)**, **Low-Inner-Update (LI)**, **High-Outer-Update (HO)**, and **Low-Outer-Update (LO)**, which represents the fine-tuning and meta-updating process of the bi-level networks respectively. Since our framework simultaneously meta-learns both θ_h and $\theta_{l1}, \cdots, \theta_{lk}$ in an iterative fashion, we call our method dual meta-learning (DMIL). The whole procedure is summarized in algorithm 1.

HI: For each sampled task $\mathcal{T}_i$, we sample the first batch of trajectories $\{\tau_{i1}\}$ from $\mathcal{D}_i^{tr}$. The principle of this step is to use sub-skill that can predict the closest action to the expert action to provide self-supervised category ground truths for the training of high-level network, which is a classifier in form. We make *every* state-action pair passed directly to each sub-skill and compute $\mathcal{L}_{BC}(\theta_{lk}, \tau_{i1}), k = 1, \cdots, K$, and choose the ground truth at each time step as the sub-skill category k that minimizes $\mathcal{L}_{BC}(\theta_{lk}, (s_t, a_t))$:

$$p(z_{i1t} = k) = \begin{cases} 1, \text{ if } k = \arg\min_k \mathcal{L}_{BC}(\theta_{lk}, (s_t, a_t)) \\ 0, \text{ else} \end{cases} . \tag{2}$$

Then we get predicted sub-skill categories from the high-level network: $\hat{z_{i1t}} = \pi_{\theta_h}(s_t)$, and use a cross-entropy loss to train the high-level network:

$$\mathcal{L}_h(\theta_h, \tau_{i1}) = -\frac{1}{T_{i1}} \sum_{t=1}^{T_{i1}} \sum_{k=1}^{K} p(z_{i1t} = k) \log p(\hat{z_{i1t}} = k). \tag{3}$$

Finally we perform gradient descent on the high-level network and get $\lambda_h = \theta_h - \alpha \nabla_{\theta_h} \mathcal{L}_h(\theta_h, \tau_{i1})$. Note $\theta_{l1}, \cdots, \theta_{lk}$ are freezed here.

LI: We sample the second batch of trajectories $\{\tau_{i2}\}$ from $\mathcal{D}_i^{tr}$. The adapted high level network π_{λ_h} will process each state in τ_{i2} to get sub-skill category $\hat{z_{i2t}} = \pi_{\lambda_h}(s_t)$ at each time step, thus we get K separate data sets for different sub-skills: $\mathcal{D}_{2k} = \{(s_{i2t}, a_{i2t}) | \hat{z_{i2t}} = k\}, k = 1, \cdots, K$. Then we compute the adaptation loss for each sub-skill with the corresponding dataset. In case of continuous action space, we assume that actions belong to Gaussian distributions, so we have:

$$\mathcal{L}_{BC}(\theta_{lk}, \mathcal{D}_{2k}) = -\frac{1}{N_k} \sum_{t=1}^{N_k} (a_t - \pi_{\theta_{lk}}(s_t))^2, \tag{4}$$

where N_k is the number of state-action pairs in $\mathcal{D}_{2k}$. Finally we perform gradient descent on sub-skills and get $\lambda_{lk} = \theta_{lk} - \alpha \nabla_{\theta_{lk}} \mathcal{L}_{BC}(\theta_{lk}, \mathcal{D}_{2k})$. Note π_{λ_h} is frozen in this process.

HO: We sample the third batch of trajectories $\{\tau_{i3}\}$ from $\mathcal{D}_i^{tr}$ and get $\mathcal{L}(\lambda_h, \tau_{i3})$ as in the HI process. Then we compute meta-gradient for θ_h with $\mathcal{L}(\lambda_h, \tau_{i3})$ as follows:

$$\nabla_{\theta_h} \mathcal{L}(\lambda_h, \tau_{i3}) = \nabla_{\lambda_h} \mathcal{L}(\lambda_h, \tau_{i3})|_{\lambda_h = \theta_h - \alpha \nabla_{\theta_h} \mathcal{L}(\theta_h, \tau_{i2})} * \nabla_{\theta_h} \lambda_h. \tag{5}$$

LO: we sample τ_{i4} and get $\mathcal{L}(\lambda_{lk}, \mathcal{D}_{4k})$, $k = 1, \cdots, K$ as in the LI process, then we compute meta-gradient for θ_{lk} with $\mathcal{L}(\lambda_{lk}, \mathcal{D}_{4k})$ as follows:

$$\nabla_{\theta_{lk}} \mathcal{L}(\lambda_{lk}, \mathcal{D}_{4k}) = \nabla_{\lambda_{lk}} \mathcal{L}(\lambda_{lk}, \mathcal{D}_{4k})|_{\lambda_{lk} = \theta_{lk} - \alpha \nabla_{\theta_{lk}} \mathcal{L}(\theta_{lk}, \mathcal{D}_{1k})} * \nabla_{\theta_{lk}} \lambda_{lk}. \tag{6}$$

Note after the training of m tasks, we average all meta-gradients from m tasks and perform gradient descents on the initial parameters *together* to get new high-level parameters $\theta'_h = \theta_h - \beta \sum_{i=1}^{m} \nabla_{\theta_h} \mathcal{L}(\lambda_h, \tau_{i3})$ and new sub-skill policies parameters $\theta'_{lk} = \theta_{lk} - \beta \sum_{i=1}^{m} \nabla_{\theta_{lk}} \mathcal{L}(\lambda_{lk}, \tau_{i4})$, $k = 1, \cdots, K$, i.e., we do not update parameters at 5 and 6.

For testing, although our method needs totally two batches of trajectories for one round of adaptation, in practice we find only using one trajectory to perform HI and LI also works well in new tasks, thus DMIL can satisfy the one-shot imitation learning requirement. Besides the above process, we also add an auxiliary loss to better drive out meaningful sub-skills to avoid the excessively frequently switching between different sub-skills along with time. Detailed information can be found in B.

4 THEORETICAL ANALYSIS

DMIL is a brand new iterative hierarchical meta-learning procedure. Like other iterative learning methods, the convergence of DMIL is important to ensure feasibility. In this section, we prove the convergence of DMIL by rewriting both MAML and DMIL as hierarchical variational Bayes problems to establish the equivalence between them, since the convergence of MAML can be proved in Fallah et al. (2020). Proofs of all theorems are in Appendix C.

4.1 HIERARCHICAL VARIATIONAL BAYES FORMULATION OF MAML

According to (Ravi & Beatson, 2019), MAML is a hierarchical variational Bayes inference process. The general meta-learning objective 1 can be formulated as follows:

$$\begin{aligned}\mathcal{L}(\theta, \lambda_1, \cdots, \lambda_m) = \log \left[\prod_{i=1}^{m} p\left(\mathcal{D}_i|\theta\right)\right] \geq \sum_{i=1}^{m} \{&\mathrm{KL}(q(\phi_i; \lambda_i) \parallel p(\phi_i|\mathcal{D}_i, \theta)) \\ &+ E_{q(\phi_i;\lambda_i)}[\log p(\mathcal{D}_i, \phi_i|\theta) - \log q(\phi_i; \lambda_i)]\},\end{aligned} \tag{7}$$

where $\phi_i, i = 1, \cdots, m$ represent the local latent variables for task $\mathcal{T}_i$, and $\lambda_1, \cdots, \lambda_M$ are the variational parameters of the approximate posteriors over $\phi_1, \cdots, \phi_M$. We denote λ_i as $\lambda_i(\mathcal{D}_i, \theta)$ and $p(\phi_i|\mathcal{D}_i, \theta)$ as $p(\phi_i|\mathcal{D}_i^{tr}, \theta)$ to mean that λ_i and ϕ_i are determined with prior parameters θ and support data $\mathcal{D}_i^{tr}$. First we need to minimize $\mathrm{KL}(q(\phi_i; \lambda_i)\|p(\phi_i|\mathcal{D}_i^{tr}, \theta))$ w.r.t. λ_i. According to C.2, we have:

$$\lambda_i(\mathcal{D}_i^{tr}, \theta) = \arg\max_{\lambda_i} E_{q(\phi_i;\lambda_i)}[\log p(\mathcal{D}_i^{tr}|\phi_i)] - \mathrm{KL}(q(\phi_i; \lambda_i)\|p(\phi_i|\theta)), \tag{8}$$

and we can establish the connection between 8 and the fine-tuning process in MAML by the following Lemma:

Lemma 1 In case $q(\phi_i; \lambda_i)$ is a Dirac-delta function and choosing Gaussian prior for $p(\phi_i|\theta)$, equation 8 equals to the inner-update step of MAML, that is, maximizing $\log p(\mathcal{D}_i^{tr})$ w.r.t. λ_i by early-stopping gradient-ascent with choosing μ_θ as initial point:

$$\lambda_i(\mathcal{D}_i^{tr}; \theta) = \mu_\theta + \alpha \nabla_\theta \log p\left(\mathcal{D}_i^{tr}|\theta\right)|_{\theta=\mu_\theta}. \tag{9}$$

Then we need to optimize $\mathcal{L}(\theta, \lambda_1, \cdots, \lambda_M)$ w.r.t. θ. Since we evaluate $p(\mathcal{D}_i|\lambda_i(\mathcal{D}_i^{tr}, \theta))$ with only $\mathcal{D}_i^{val}$, we assume $p(\mathcal{D}_i|\lambda_i(\mathcal{D}_i^{tr}, \theta)) = p(\mathcal{D}_i^{val}|\lambda_i(\mathcal{D}_i^{tr}, \theta))$. We give out the following theorem to establish the connection between the meta-update process and the optimization of $\mathcal{L}(\theta, \lambda_1, \cdots, \lambda_M)$:

Theorem 1 In case that $\Sigma_\theta \to 0^+$, i.e., the uncertainty in the global latent variables θ is small, the following equation holds:

$$\nabla_\theta \mathcal{L}(\theta, \lambda_1, \cdots, \lambda_M) = \sum_{i=1}^{M} \nabla_{\lambda_i} \log p(\mathcal{D}_i^{val}|\lambda_i) \cdot \nabla_\theta \lambda_i(\mathcal{D}_i^{tr}, \theta). \tag{10}$$

A general EM algorithm will first compute the distribution of latent variables (E-step), then optimize the joint distribution of latent variable and trainable parameters (M-step), and the likelihood of data can be proved to be monotone increasing to guarantee the convergence since the evidence lower bound of likelihood is monotone increasing. Here $\phi_i, i = 1, \cdots, M$ are the latent variables, and θ corresponds to the trainable parameters. Lemma 1 and Theorem 1 correspond to the E-step and M-step respectively. In the following part we establish the equivalence between 9 with 3 and 4, and between 10 with 5 and 6 to prove the equivalence between DMIL and MAML.

4.2 MODELING DMIL WITH HIERARCHICAL VARIATIONAL BAYES FRAMEWORK

For simplicity, here we only derive in one specific task $\mathcal{T}_i$, since derivatives of parameters from multi-task can directly add up. We first establish the connection between the maximization of $\log p(\mathcal{D}_i^{tr}|\theta_h, \theta_{l1}, \cdots, \theta_{lK})$ with the particular loss functions in DMIL:

Theorem 2 In case of $p(a_t|s_t, \theta_{lk}) \sim \mathcal{N}(\mu_{\theta_{lk}(s_t)}, \sigma^2)$, we have:

$$\nabla_{\theta_h} \log p(\mathcal{D}_i^{tr}|\theta_h, \theta_{l1}, \cdots, \theta_{lK}) = \nabla_{\theta_h} \mathcal{L}_h(\theta_h, \mathcal{D}_i^{tr}), \tag{11}$$

and

$$\nabla_{\theta_{lk}} \log p(\mathcal{D}_i^{tr}|\theta_h, \theta_{l1}, \cdots, \theta_{lK}) = \nabla_{\theta_{lk}} \mathcal{L}_{BC}(\theta_{lk}, \mathcal{D}_{2k}), k = 1, \cdots, K. \tag{12}$$

Note in 12, $\mathcal{D}_{2k}$ coresponds to data sets determined by the adapted high level network λ_h, and this connects with 3 and 4 in DMIL. According to 8, finding λ_i equals to maximize $\log p\left(\mathcal{D}_i^{tr}|\theta\right)$ in specific conditions, and here in Theorem 2, we prove that maximize $\log p(\mathcal{D}_i^{tr}|\theta_h, \theta_{l1}, \cdots, \theta_{lK})$ corresponds to 3 and 4 in DMIL. Thus theorem 2 coresponds to the E-step of DMIL, where we take τ_{i1} and τ_{i2} as $\mathcal{D}_i^{tr}$, and optimize $\arg\max_{\lambda_i} E_{q(\phi_i;\lambda_i)}[\log p(\mathcal{D}_i^{tr}|\phi_i)] - \mathrm{KL}(q(\phi_i;\lambda_i)\|p(\phi_i|\theta))$ with coordinate descent method, which can be proved to be equal to 9 in C.5.

For the M-step, we take τ_{i3} and τ_{i4} as $\mathcal{D}_i^{val}$. According to Theorem 1, we can take the derivative of $\lambda_{ih}, \lambda_{il1}, \cdots, \lambda_{ilK}$ to maximize the joint distribution of latent variables and trainable parameters to maximize the likelihood of dataset, so we have:

$$\begin{aligned}\nabla_{\theta_h,\theta_l} \log p(\mathcal{D}_i^{val}|\lambda_{ih}, \lambda_{il}) = [&\nabla_{\lambda_{ih}} \log p(\mathcal{D}_i^{val}|\lambda_{ih}) * \nabla_{\theta_h} \lambda_{ih}(\mathcal{D}_i^{tr}, \theta_h), \\ &\nabla_{\lambda_{il}} \log p(\mathcal{D}_i^{val}|\lambda_{il}) * \nabla_{\theta_l} \lambda_{il}(\mathcal{D}_i^{tr}, \theta_l)]^T\end{aligned} \tag{13}$$

where $\theta_{il} = [\theta_{i1}, \cdots, \theta_{iK}]^T$ and $\lambda_{il} = [\lambda_{i1}, \cdots, \lambda_{iK}]^T$. This is exactly the gradients computed in HO and LO steps. Note this computation process can be automatically accomplished with standard deep learning libraries such as PyTorch (Paszke et al., 2019). To this end, we establish the equivalence between DMIL and MAML, and the convergence of DMIL can be proved.

For a clearer comparision, MAML is an iterative process of $\theta \rightarrow \lambda \rightarrow \theta'$, and DMIL is an iterative process of $\theta_h, \theta_l \rightarrow \lambda_h, \theta_l \rightarrow \lambda_h, \lambda_l \rightarrow \theta'_h, \theta'_l$, where the posterior estimation stages $\theta_h, \theta_l \rightarrow \lambda_h, \theta_l \rightarrow \lambda_h, \lambda_l$ has no effect on parameters θ_h, θ_l, thus can be divided to two steps as in DMIL. This decoupled fine-tuning fashion is exactly what we need to first adapt the high-level network and then adapt sub-skills. If we end-to-end fine-tune parameters like $\theta_h, \theta_l \rightarrow \lambda_h, \lambda_l$, sub-skills will receive supervisions from an unadapted high-level network, which may provide incorrect classifications. Different to this, the meta-updating process $\lambda_h, \lambda_l \rightarrow \theta'_h, \theta'_l$ must be done at the same time, since if we update θ_h and θ_l successively, the later one will receive different derivative (for example, $\nabla_{\theta_l} \log p(\mathcal{D}_i^{val}|\theta'_{ih}, \lambda_{il})$) from derivatives in MAML ($\nabla_{\theta_l} \log p(\mathcal{D}_i^{val}|\lambda_{ih}, \lambda_{il})$), and the equivalence would not be proved.

5 EXPERIMENTS

To evaluate the effectiveness of the hierarchical structure and meta-learning process in DMIL, we perform meta imitation learning experiments on the meta-world (Yu et al., 2019b) benchmark, which consists of 50 diverse robot manipulation tasks. Video results of DMIL are also provided (see section 7). Detailed information about the experiment setup, hyper-parameters setting, training process, more experiment results, and descriptions of different methods can be found in D.2.

5.1 EXPERIMENTAL SETUP

Meta-training and Meta-testing Tasks: We evaluate all methods on the ML10 suite (10 meta-training tasks and 5 meta-testing tasks, see figure 7) and ML45 suite (45 meta-training tasks and 5 meta-testing tasks, see figure 6) in the meta-world benchmark. We collect 2K expert demonstrations for each meta-training task with scripted policies. The observation and action space are the same within the original paper (Yu et al., 2019b).

Model Details: We choose to use 3 sub-skills ($K = 3$) to perform our experiments, which is inspired by the grouping operation in Cachet et al. (2020), where they group the environments in mete-world to 3 groups intuitively. We use Reptile (Nichol et al., 2018) to warm up the whole network architecture and then use the original MAML for later training.

Comparison Methods: We compare DMIL to a number of prior methods, including: 1) an off-policy variant of Option-GAIL (Jing et al., 2021) to evaluate the effect of meta-learning in DMIL; 2) a transformer-based meta imitation learning method (MIL for short) (Finn et al., 2017b) to evaluate the effect of hierarchical structure in DMIL; 3) a behavior cloning variant of meta-learning shared hierarchies method (MLSH for short) (Vuorio et al., 2019) to evaluate the effect of keeping high-level network during training process; 4) and an off-policy variant of the contextual meta inverse RL

Table 1: Success rates of different methods. The result of each column is the mean and standard variance of 5 data points. Each data point comes from the success rate of 20 episodes.

	ML10				ML45			
	Meta-training		Meta-testing		Meta-training		Meta-testing	
Methods	1-shot	3-shot	1-shot	3-shot	1-shot	3-shot	1-shot	3-shot
OptionGAIL	0.545±0.011	**0.840±0.016**	0.241±0.042	0.640±0.025	0.506±0.008	0.715±0.006	0.220±0.013	0.381±0.010
MIL	**0.619±0.025**	0.809±0.029	0.361±0.040	0.509±0.032	0.384±0.011	0.645±0.017	0.205±0.024	0.310±0.005
PEMIRL	0.398±0.023	0.510±0.007	0.162±0.003	0.256±0.009	0.109±0.051	0.126±0.024	0.105±0.005	0.122±0.008
MLSH	0.406±0.134	0.625±0.021	0.106±0.032	0.135±0.009	0.235±0.093	0.295±0.021	0.050±0.000	0.050±0.000
DMIL-High	0.434±0.001	0.769±0.011	0.268±0.012	0.670±0.015	0.455±0.007	0.635±0.006	0.280±0.016	0.451±0.011
DMIL-Low	0.466±0.011	0.713±0.006	0.251±0.024	0.440±0.021	0.511±0.005	0.705±0.009	0.200±0.010	0.392±0.014
DMIL	0.421±0.010	0.788±0.009	**0.389±0.016**	**0.697±0.021**	**0.584±0.010**	**0.751±0.008**	**0.312±0.004**	**0.531±0.009**

method which transfer the reward function in the new task (PEMIRL for short) (Yu et al., 2019a) to show DMIL can imitate behaviors that are discrepant from previously learned ones.

Ablation Study: We use two variants of DMIL to see the effectiveness of the bi-level meta-earning process, including a model that only meta-learns the high-level network (DMIL-High) and a model that only meta-learns sub-skills (DMIL-Low).

5.2 Quantitative Results

We report the success rates of meta-training and meta-testing tasks in ML10 and ML45 suite of all methods in table 1. We perform 1-shot and 3-shot experiments respectively to show the progressive few-shot performance of different methods. Across comparable methods, DMIL achieves the best results in ML10 testing suite and ML45 training and testing suites. This shows the superiority of our method for transferring across a large scale of manipulation tasks. OptionGAIL achieves the highest success rate in ML10 training suite and achieves results close to DMIL in ML45 training suite. MLSH also achieves similar results. These results come from their hierarchical structures that have adequate capacity to fit potential multi-modal behaviors in multi-task demonstrations. Nevertheless, the variance of MLSH is larger than other methods. This is because the high-level network of MLSH is relearned. Compared to them, MIL achieves better results for all meta-testing tasks but is still a little worse than DMIL. PEMIRL is mediocre among all suites. This conveys that the reward functions across different tasks are difficult to transfer with few shot demonstrations.

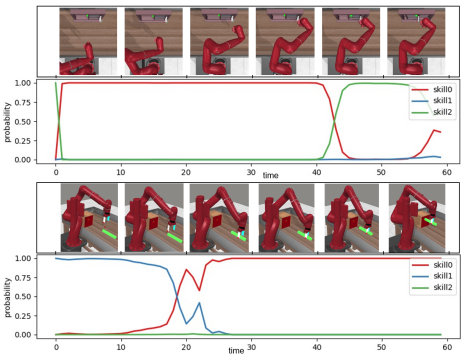

Figure 3: Sub-skills probability along two representative tasks *window-close* and *peg-insert-side*. The learned sub-skills develop distinct specializations.

For ablation studies, DMIL-High achieves close results with OptionGAIL in all meta-training suites and better results in all meta-testing suites. This comes from the meta-learned high-level network. DMIL-Low achieves poor results in all suites except in ML10 meta-training suites. We think this comes from two aspects. On the one hand, incorporating MAML brings training difficulty, thus the performance of DMIL-Low is similar to or no better than OptionGAIL in meta-training tasks. On the other hand, for a large scale of tasks such as ML45, and new tasks such as meta-testing tasks, DMIL-Low lacks the ability to transfer the high-level network to make a correct classification.

5.3 Qualitative Analysis

To illustrate the results of DMIL, we draw the curves of sub-skill probabilities along time of two tasks *window-close* and *peg-insert-side* in figure 3. We can see the activations of sub-skills show a strong correlation to different stages of tasks, with $\pi_{\theta_{l0}}$ being activated when the robot is closing to something, $\pi_{\theta_{l1}}$ being activated when the robot is picking up something, and $\pi_{\theta_{l2}}$ being activated when the robot is manipulating something.

For comparison of different methods, we illustrate the t-sne results of states of each sub-skill in an ML45 meta-testing task *hand-insert* in figure 4. We use 3 demonstrations for adaptation, and draw

the t-sne results on another 16 demonstrations. This task is a meta-testing task, so no method has ever been trained on this task before.

MLSH shows almost random clustering results no matter before and after adaptation, since its high-level network is relearned in new tasks. OptionGAIL clusters to three sub-skills after adaptation. Compared to them, DMIL clusters the data to only two sub-skills after adaptation. We believe fewer categories reflect more meaningful sub-skills are developed in DMIL.

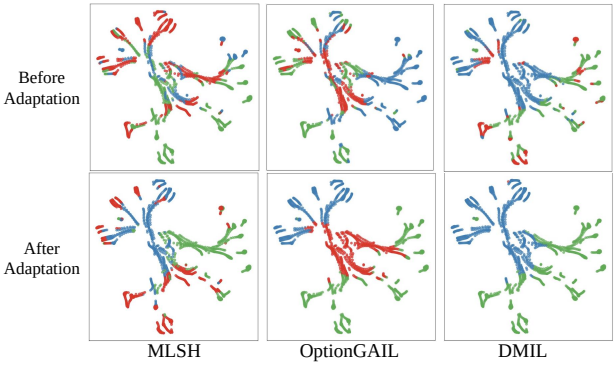

Figure 4: T-sne results of demonstrations and sub-skill categories of several hierarchical models for meta-testing task *hand-insert*.

Finally we perform ablation study for three DMIL variants in an manually-designed new task *push-around-wall* (figure D.1). In this task, the robot needs to grasp a cube and circle it around the wall. This is a brand new skill that is not in the meta-world suite. The accomplishment of this new task requires quickly adapting abilities of both the high-level network and sub-skills. We sample two demonstrations and use the first one as few-shot data, and illustrate the sub-skill categories of the second demonstration given by the high-level network at figure 5. Before adaptation, all variants give out approximately random results. However, after one-shot adaptation, DMIL classifies almost every state into sub-skill 0 and sub-skill 2, which indicates these two sub-skills in DMIL have been adapted to the new task, and the high-level network has also been adapted with the supervision from adapted sub-skills. Compared to DMIL, DMIL-Low still can not give out reasonable results after adaptation, since its high-level network lacks the ability to quickly transfer to new tasks. Instead, DMIL-High gives out plausible results after adaptation. This shows the high-level network has adapted to the new task according to the supervision from adapted sub-skills, but there is no enough ability in sub-skills to quickly adapt to new tasks, thus no certain sub-skill exceeds others to give out dominant supervision for the high-level network.

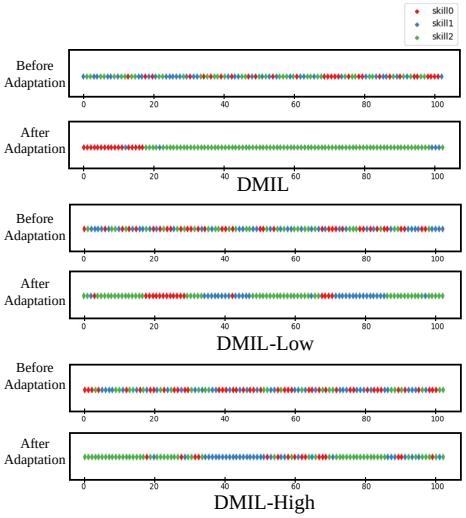

Figure 5: Sub-skill categories of task *push-around-wall* for ablation study.

6 Discussion and Future Works

In this work, we propose DMIL to meta-learn a hierarchical structure fashion from unsegmented multi-task demonstrations to endow it with fast adaptation ability to transfer to new tasks. We theoretically proved its convergence by reframing MAML and DMIL as hierarchical Bayes inference processes to get their equivalence. We successfully acquire a high-level network and a set of transferable sub-skills with DMIL from a large scale of robot manipulation tasks and exhibit its outstanding few-shot transfer ability compared to other baselines and ablation variants.

Future works can seek meaningful extensions of our method from several perspectives. For acquiring sub-skills from unsegmented multi-task demonstrations, it is worth studying the connection and (potential) exclusiveness between different metrics, such as likelihood-maximization metrics (EM-based methods like DMIL or Daniel et al. (2016); Krishnan et al. (2017)), information-theoretic metrics (like Eysenbach et al. (2019); Florensa et al. (2017); Sharma et al. (2020)) and occupancy measurement minimization metrics (like Li et al. (2017); Jing et al. (2021)). For the hierarchical architecture, a multi-level structure of skills is worth studying to better model manipulation skills in real-world, as shown in Shu et al. (2018). Self-attention modules like transformers (Vaswani et al., 2017) are also worth studying to be used as the high-level network to better capture the temporal relationships in demonstrations.

7 REPRODUCIBILITY STATEMENT

We provide trained models of DMIL, some basic codes to run few-shot testing in meta-world environments, and video results in an anonymous repository [1]. The complete training code will be provided later. Details of experiments such as environments, hyper-parameter setting, model setup, training procedure and data preprocessing can be found at D.

REFERENCES

Ferran Alet, Tomás Lozano-Pérez, and Leslie Pack Kaelbling. Modular meta-learning. In *2nd Annual Conference on Robot Learning*, 2018.

Ron Amit and Ron Meir. Meta-learning by adjusting priors based on extended pac-bayes theory. In *Proceedings of the 35th International Conference on Machine Learning*, 2018.

Christopher G. Atkeson and Stefan Schaal. In *Proceedings of the Fourteenth International Conference on Machine Learning*, 1997.

T Cachet, J Perez, and S Kim. Transformer-based meta-imitation learning for robotic manipulation. In *Neural Information Processing Systems, Workshop on Robot Learning*, 2020.

Xi Chen, Yan Duan, Rein Houthooft, John Schulman, Ilya Sutskever, and Pieter Abbeel. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. In *Advances in Neural Information Processing Systems*, 2016.

Christian Daniel, Herke van Hoof, Jan Peters, and Gerhard Neumann. Probabilistic inference for determining options in reinforcement learning. *Mach. Learn.*, 2016.

Yan Duan, Marcin Andrychowicz, Bradly C. Stadie, Jonathan Ho, Jonas Schneider, Ilya Sutskever, Pieter Abbeel, and Wojciech Zaremba. One-shot imitation learning. In *Advances in Neural Information Processing Systems*, 2017.

David Duvenaud, Dougal Maclaurin, and Ryan Adams. Early stopping as nonparametric variational inference. In *Artificial Intelligence and Statistics*. PMLR, 2016.

Benjamin Eysenbach, Abhishek Gupta, Julian Ibarz, and Sergey Levine. Diversity is all you need: Learning skills without a reward function. In *7th International Conference on Learning Representations*, 2019.

Alireza Fallah, Aryan Mokhtari, and Asuman E. Ozdaglar. On the convergence theory of gradient-based model-agnostic meta-learning algorithms. In *The 23rd International Conference on Artificial Intelligence and Statistics*, 2020.

Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of the 34th International Conference on Machine Learning*, 2017a.

Chelsea Finn, Tianhe Yu, Tianhao Zhang, Pieter Abbeel, and Sergey Levine. One-shot visual imitation learning via meta-learning. In *1st Annual Conference on Robot Learning*, 2017b.

Chelsea Finn, Kelvin Xu, and Sergey Levine. Probabilistic model-agnostic meta-learning. In *Advances in Neural Information Processing Systems*, 2018.

Carlos Florensa, Yan Duan, and Pieter Abbeel. Stochastic neural networks for hierarchical reinforcement learning. In *5th International Conference on Learning Representations*, 2017.

Kevin Frans, Jonathan Ho, Xi Chen, Pieter Abbeel, and John Schulman. Meta learning shared hierarchies. In *6th International Conference on Learning Representations*, 2018.

Seyed Kamyar Seyed Ghasemipour, Shixiang Gu, and Richard S. Zemel. Smile: Scalable meta inverse reinforcement learning through context-conditional policies. In *Advances in Neural Information Processing Systems 32:*, 2019.

[1] https://anonymous.4open.science/r/DMIL-6FCC

Erin Grant, Chelsea Finn, Sergey Levine, Trevor Darrell, and Thomas L. Griffiths. Recasting gradient-based meta-learning as hierarchical bayes. In *6th International Conference on Learning Representations*, 2018.

Tuomas Haarnoja, Kristian Hartikainen, Pieter Abbeel, and Sergey Levine. Latent space policies for hierarchical reinforcement learning. In Jennifer G. Dy and Andreas Krause (eds.), *Proceedings of the 35th International Conference on Machine Learning*, 2018.

Matthew J. Hausknecht and Peter Stone. Deep reinforcement learning in parameterized action space. In *4th International Conference on Learning Representations*, 2016.

Karol Hausman, Jost Tobias Springenberg, Ziyu Wang, Nicolas Heess, and Martin A. Riedmiller. Learning an embedding space for transferable robot skills. In *6th International Conference on Learning Representations*, 2018.

Peter Henderson, Wei-Di Chang, Pierre-Luc Bacon, David Meger, Joelle Pineau, and Doina Precup. Optiongan: Learning joint reward-policy options using generative adversarial inverse reinforcement learning. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.

Jonathan Ho and Stefano Ermon. Generative adversarial imitation learning. In *Advances in Neural Information Processing Systems*, 2016.

Robert A. Jacobs, Michael I. Jordan, Steven J. Nowlan, and Geoffrey E. Hinton. Adaptive mixtures of local experts. *Neural Comput.*, 1991.

Mingxuan Jing, Wenbing Huang, Fuchun Sun, Xiaojian Ma, Tao Kong, Chuang Gan, and Lei Li. Adversarial option-aware hierarchical imitation learning. In *Proceedings of the 38th International Conference on Machine Learning*, 2021.

Ilya Kostrikov, Kumar Krishna Agrawal, Sergey Levine, and Jonathan Tompson. Addressing sample inefficiency and reward bias in inverse reinforcement learning. *arXiv preprint arXiv:1809.02925*, 2018.

Sanjay Krishnan, Roy Fox, Ion Stoica, and Ken Goldberg. DDCO: discovery of deep continuous options for robot learning from demonstrations. In *1st Annual Conference on Robot Learning*, 2017.

Hoang Minh Le, Nan Jiang, Alekh Agarwal, Miroslav Dudík, Yisong Yue, and Hal Daumé III. In *Proceedings of the 35th International Conference on Machine Learning*, 2018.

Sang-Hyun Lee and Seung-Woo Seo. Learning compound tasks without task-specific knowledge via imitation and self-supervised learning. In *Proceedings of the 37th International Conference on Machine Learning*, 2020.

Yunzhu Li, Jiaming Song, and Stefano Ermon. Infogail: Interpretable imitation learning from visual demonstrations. In *Advances in Neural Information Processing Systems*, 2017.

Libin Liu and Jessica K. Hodgins. Learning to schedule control fragments for physics-based characters using deep q-learning. *ACM Trans. Graph.*, 2017.

Corey Lynch, Mohi Khansari, Ted Xiao, Vikash Kumar, Jonathan Tompson, Sergey Levine, and Pierre Sermanet. Learning latent plans from play. In *3rd Annual Conference on Robot Learning*, 2019.

Josh Merel, Arun Ahuja, Vu Pham, Saran Tunyasuvunakool, Siqi Liu, Dhruva Tirumala, Nicolas Heess, and Greg Wayne. Hierarchical visuomotor control of humanoids. In *7th International Conference on Learning Representations*, 2019a.

Josh Merel, Leonard Hasenclever, Alexandre Galashov, Arun Ahuja, Vu Pham, Greg Wayne, Yee Whye Teh, and Nicolas Heess. Neural probabilistic motor primitives for humanoid control. In *7th International Conference on Learning Representations*, 2019b.

Gerhard Neumann, Wolfgang Maass, and Jan Peters. Learning complex motions by sequencing simpler motion templates. In Andrea Pohoreckyj Danyluk, Léon Bottou, and Michael L. Littman (eds.), *Proceedings of the 26th Annual International Conference on Machine Learning*, 2009.

Alex Nichol, Joshua Achiam, and John Schulman. On first-order meta-learning algorithms. *arXiv preprint arXiv:1803.02999*, 2018.

Takayuki Osa, Joni Pajarinen, Gerhard Neumann, J. Andrew Bagnell, Pieter Abbeel, and Jan Peters. An algorithmic perspective on imitation learning. 2018.

Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 2019.

Xue Bin Peng, Pieter Abbeel, Sergey Levine, and Michiel van de Panne. Deepmimic: example-guided deep reinforcement learning of physics-based character skills. *ACM Trans. Graph.*, 2018.

Xue Bin Peng, Michael Chang, Grace Zhang, Pieter Abbeel, and Sergey Levine. MCP: learning composable hierarchical control with multiplicative compositional policies. In *Advances in Neural Information Processing Systems*, 2019.

Sachin Ravi and Alex Beatson. Amortized bayesian meta-learning. In *7th International Conference on Learning Representations*, 2019.

Stéphane Ross, Geoffrey J. Gordon, and Drew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, 2011.

Reginaldo J Santos. Equivalence of regularization and truncated iteration for general ill-posed problems. *Linear algebra and its applications*, 1996.

Archit Sharma, Shixiang Gu, Sergey Levine, Vikash Kumar, and Karol Hausman. Dynamics-aware unsupervised discovery of skills. In *8th International Conference on Learning Representations*, 2020.

Mohit Sharma, Arjun Sharma, Nicholas Rhinehart, and Kris M. Kitani. Directed-info GAIL: learning hierarchical policies from unsegmented demonstrations using directed information. In *7th International Conference on Learning Representations*, 2019.

Tianmin Shu, Caiming Xiong, and Richard Socher. Hierarchical and interpretable skill acquisition in multi-task reinforcement learning. In *6th International Conference on Learning Representations*, 2018.

Richard S. Sutton, Doina Precup, and Satinder P. Singh. Between mdps and semi-mdps: A framework for temporal abstraction in reinforcement learning. *Artif. Intell.*, 1999.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, 2017.

Risto Vuorio, Shao-Hua Sun, Hexiang Hu, and Joseph J. Lim. Multimodal model-agnostic meta-learning via task-aware modulation. In *Advances in Neural Information Processing Systems*, 2019.

Kelvin Xu, Ellis Ratner, Anca D. Dragan, Sergey Levine, and Chelsea Finn. In *Proceedings of the 36th International Conference on Machine Learning*, 2019.

Huaxiu Yao, Ying Wei, Junzhou Huang, and Zhenhui Li. Hierarchically structured meta-learning. In *Proceedings of the 36th International Conference on Machine Learning*, 2019.

Lantao Yu, Tianhe Yu, Chelsea Finn, and Stefano Ermon. Meta-inverse reinforcement learning with probabilistic context variables. In *Advances in Neural Information Processing Systems*, 2019a.

Tianhe Yu, Pieter Abbeel, Sergey Levine, and Chelsea Finn. One-shot hierarchical imitation learning of compound visuomotor tasks. *arXiv preprint arXiv:1810.11043*, 2018a.

Tianhe Yu, Chelsea Finn, Sudeep Dasari, Annie Xie, Tianhao Zhang, Pieter Abbeel, and Sergey Levine. One-shot imitation from observing humans via domain-adaptive meta-learning. In *Robotics: Science and Systems XIV*, 2018b.

Tianhe Yu, Deirdre Quillen, Zhanpeng He, Ryan Julian, Karol Hausman, Chelsea Finn, and Sergey Levine. Meta-world: A benchmark and evaluation for multi-task and meta reinforcement learning. In *3rd Annual Conference on Robot Learning*, 2019b.

Yuke Zhu, Ziyu Wang, Josh Merel, Andrei A. Rusu, Tom Erez, Serkan Cabi, Saran Tunyasuvunakool, János Kramár, Raia Hadsell, Nando de Freitas, and Nicolas Heess. Reinforcement and imitation learning for diverse visuomotor skills. In *Robotics: Science and Systems XIV*, 2018.

Yayi Zou and Xiaoqi Lu. Gradient-em bayesian meta-learning. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems*, 2020.

A ALGORITHM

Algorithm 1 Dual Meta Imitation Learning

Require: task distribution $p(\mathcal{T})$, multi-task demonstrations $\{\mathcal{D}_i\}, i = 1, \cdots, m$, initial parameters of high-level network θ_h and sub-skill policies $\theta_{l1}, \cdots, \theta_{lK}$, inner and outer learning rate α, β.
while not done **do**
  Sample batch of tasks $\mathcal{T}_i \sim p(\mathcal{T})$
  for all $\mathcal{T}_i$ **do**
    Sample $\{\tau_{i1}\}, \{\tau_{i2}\}, \{\tau_{i3}\}, \{\tau_{i4}\}$ from $\{\mathcal{D}_i\}$
    Evaluate $\nabla_{\theta_h} \mathcal{L}_h(\theta_h, \tau_{i1})$ according to 3 and τ_{i1}
    Compute adapted parameters of high-level network: $\lambda_h = \theta_h - \alpha \nabla_{\theta_h} \mathcal{L}_h(\theta_h, \tau_{i1})$
    Evaluate $\nabla_{\theta_{lk}} \mathcal{L}_{BC}(\theta_{lk}, \mathcal{D}_{2k})$ according to 4 and τ_{i2}, $k = 1, \cdots, K$
    Compute adapted parameters of sub-skills: $\lambda_{lk} = \theta_{lk} - \alpha \nabla_{\theta_{lk}} \mathcal{L}_{BC}(\theta_{lk}, \mathcal{D}_{2k})$, $k = 1, \cdots, K$
    Evaluate $\nabla_{\theta_h} \mathcal{L}_{\mathcal{T}_i}(\lambda_h, \tau_{i3})$ and $\nabla_{\theta_{lk}} \mathcal{L}_{\mathcal{T}_i}(\lambda_{lk}, \mathcal{D}_{4k})$, $k = 1, \cdots, K$
  end for
  Update $\theta_h \leftarrow \theta_h - \beta \nabla_{\theta_h} \sum_{\mathcal{T}_i \sim p(\mathcal{T})} \mathcal{L}_{\mathcal{T}_i}(\lambda_h, \tau_{i3})$
  Update $\theta_{lk} \leftarrow \theta_{lk} - \beta \nabla_{\theta_{lk}} \sum_{\mathcal{T}_i \sim p(\mathcal{T})} \mathcal{L}_{\mathcal{T}_i}(\lambda_{lk}, \mathcal{D}_{4k})$, $k = 1, \cdots, K$
end while

B AUXILIARY LOSS

We adopt an auxiliary loss for DMIL to better drive out meaningful sub-skills by punishing excessive switching of sub-skills along the trajectory. This comes from an intuitive idea: each sub-skill should be a temporal-extended macro-action, and the high-level policy only needs to switch to different skills few times along a task, as the same idea of *macro-action* in MLSH Frans et al. (2018). We denote:

$$sign(x) = \begin{cases} 1, \text{ if } x = True \\ 0, \text{ if} x = False \end{cases}, \tag{14}$$

and the auxiliary loss is:

$$\mathcal{L}_{aux}(\tau) = \sum_{t=0}^{T-1} sign(\hat{z}_{t+1} \neq \hat{z}_t) \ / \ len(\tau). \tag{15}$$

Although this operation seems discrete, in practice we can use the operations in modern deep learning framework such as PyTorch (Paszke et al., 2019) to make it differentiable. We add this loss function to the $\mathcal{L}_h(\theta_h, \tau_{i1})$ and $\nabla_{\theta_h} \mathcal{L}(\lambda_h, \tau_{i3})$ with a coefficient $\lambda = 0.1$ in a warm-up period from beginning, i.e. λ will be 0 after this period. In the ML10 suite, we set this period as 1000 epochs, and set it as 2000 epochs in ML45 suite.

C PROOFS

C.1 PROOF OF LEMMA 1

Lemma 1 In case $q(\phi_i; \lambda_i)$ is a Dirac-delta function and choosing Gaussian prior for $p(\phi_i|\theta)$, equation 8 equals to the inner-update step of MAML, that is, maximizing $\log p(\mathcal{D}_i^{tr})$ w.r.t. λ_i by early-stopping gradient-ascent with choosing μ_θ as initial point:

$$\lambda_i(\mathcal{D}_i^{tr}; \theta) = \mu_\theta + \alpha \nabla_\theta \log p\left(\mathcal{D}_i^{tr}|\theta\right)|_{\theta=\mu_\theta}. \tag{16}$$

Proof: in case of the conditions of Lemma 1, we have:

$$\lambda_i(\mathcal{D}_i^{tr}; \theta) = \arg\max_{\lambda_i} [\log p(\mathcal{D}_i^{tr}|\mu_{\lambda_i}) - \left\|\mu_{\lambda_i} - \mu_\theta\right\|^2 / 2\Sigma_\theta^2], \tag{17}$$

As stated in Grant et al. (2018), firstly in the case of linear models, early stopping of an iterative gradient descent process of λ equals to the maximum posterior estimation (MAP) (Santos, 1996).

In our case the posterior distribution refers to $q(\phi_i|\lambda_i)$, and MAML is a Bayes process to find the MAP estimate as the point estimate of $\lambda_i(\mathcal{D}_i^{tr};\theta)$. In the nonlinear case, this point estimate is not necessarily the global mode of the posterior, and we can refer to Duvenaud et al. (2016) for another implicit posterior distribution over ϕ_i and making the early stopping procedure of MAML acting as priors to get the similar result.

C.2 Proof of Equation 8

Equation 8 can be written as:

$$\begin{aligned}\lambda_i(\mathcal{D}_i^{tr},\theta) &= \arg\min_{\lambda_i} \mathrm{KL}(q(\phi_i;\lambda_i)\|p(\phi_i|\mathcal{D}_i^{tr},\theta)) \\ &= \arg\max_{\lambda_i} E_{q(\phi_i;\lambda_i)}[\log p(\phi_i|\mathcal{D}_i^{tr},\theta) - \log q(\phi_i;\lambda_i)] \\ &= \arg\max_{\lambda_i} E_{q(\phi_i;\lambda_i)}[\log p(\mathcal{D}_i^{tr}|\phi_i)] - \mathrm{KL}(q(\phi_i;\lambda_i)\|p(\phi_i|\theta)),\end{aligned} \tag{18}$$

where in MAML we assume $p(\mathcal{D}_i^{tr}|\phi_i) = p(\mathcal{D}_i^{tr}|\phi_i,\theta)$, and use the joint distribution $p(\mathcal{D}_i^{tr},\phi_i|\theta)$ to replace $p(\phi_i|\mathcal{D}_i^{tr},\theta)$ since we assume that $p(\mathcal{D}_i^{tr})$ subjects to uniform distribution. Thus 8 can be proved.

C.3 Proof of Theorem 1

Theorem 1 In case that $\Sigma_\theta \to 0^+$, i.e., the uncertainty in the global latent variables θ is small, the following equation holds:

$$\nabla_\theta \mathcal{L}(\theta,\lambda_1,\cdots,\lambda_M) = \sum_{i=1}^{M} \nabla_{\lambda_i} \log p(\mathcal{D}_i^{val}|\lambda_i) * \nabla_\theta \lambda_i(\mathcal{D}_i^{tr},\theta). \tag{19}$$

Proof:

$$\begin{aligned}\nabla_\theta \mathcal{L}(\theta,\lambda_1,\cdots,\lambda_M) &\approx \sum_{i=1}^{M}\{\nabla_\theta E_{q(\phi_i;\lambda_i)}[\log p(\mathcal{D}_i,\phi_i|\theta) - \log q(\phi_i;\lambda_i)]\} \\ &= \sum_{i=1}^{M} \nabla_\theta[\log p(\mathcal{D}_i^{val}|\lambda_i(\mathcal{D}_i^{tr},\theta)) - \log p(\lambda_i(\mathcal{D}_i^{tr},\theta)|\theta)] \\ &\approx \sum_{i=1}^{M} \nabla_\theta \log p(\mathcal{D}_i^{val}|\lambda_i(\mathcal{D}_i^{tr},\theta)) \\ &= \sum_{i=1}^{M} \nabla_{\lambda_i} \log p(\mathcal{D}_i^{val}|\lambda_i) * \nabla_\theta \lambda_i(\mathcal{D}_i^{tr},\theta)\end{aligned} \tag{20}$$

where the first approximate equal holds because the VI approximation error is small enough, and the second approximate equal holds because that in case $\Sigma_\theta \to 0^+$ and assuming λ_i be a neuron network, $\log p(\lambda_i(\mathcal{D}_i^{tr},\theta)|\theta) \approx 0$ holds almost everywhere, so $\nabla_\theta \log p(\lambda_i(\mathcal{D}_i^{tr},\theta)|\theta) \approx 0$. Note the condition of theorem 1 is usually satisfied since we are using MAML, and the initial parameters θ are assumed to be deterministic.

From another perspective, the right side of above equation is the widely used meta-gradient in MAML, and it is equal to $\frac{1}{m}\sum_{i=1}^{m}\left(I - \alpha\nabla_\theta^2 \mathcal{L}_{BC}\left(\theta,\mathcal{D}_i^{tr}\right)\right) * \nabla_{\lambda_i}\mathcal{L}_{BC}\left(\theta - \alpha\nabla_\theta \mathcal{L}_{BC}\left(\theta,\mathcal{D}_i^{tr}\right),\mathcal{D}_i^{val}\right)$, which is proved to be converged by Fallah et al. (2020).

C.4 Proof of Theorem 2

Theorem 2 In case of $p(a_t|s_t,\theta_{lk}) \sim \mathcal{N}(\mu_{\theta_{lk}(s_t)},\sigma^2)$, we have:

$$\nabla_{\theta_h} \log p(\mathcal{D}_i^{tr}|\theta_h,\theta_{l1},\cdots,\theta_{lK}) = \nabla_{\theta_h}\mathcal{L}_h(\theta_h,\mathcal{D}_i^{tr}), \tag{21}$$

and

$$\nabla_{\theta_{lk}} \log p(\mathcal{D}_i^{tr}|\theta_h, \theta_{l1}, \cdots, \theta_{lK}) = \nabla_{\theta_{lk}} \mathcal{L}_{BC}(\theta_{lk}, \mathcal{D}_{2k}), k = 1, \cdots, K. \tag{22}$$

Proof: since $p(\mathcal{D}_i^{tr}|\theta_h, \theta_{l1}, \cdots, \theta_{lK}) = \prod_{t=1}^{N} p(a_t|s_t, \theta_h, \theta_{l1}, \cdots, \theta_{lK}) p(s_t|\theta_h, \theta_{l1}, \cdots, \theta_{lK})$ and the second term is independent of θ, we consider the first conditional probability:

$$p(a_t|s_t, \theta_h, \theta_{l1}, \cdots, \theta_{lK}) = \sum_{k=1}^{K} p(z_k|s_t, \theta_h) p(a_t|s_t, \theta_{lk}). \tag{23}$$

In the HI step, $p(a_t|s_t, \theta_{lk})$ is fixed, thus 23 becomes a convex optimization problem:

$$\max_{\theta_h} \sum_{k=1}^{K} p(z_k|s_t, \theta_h) p(a_t|s_t, \theta \tag{24}$$

The solution of this problem is λ_h^* which satisfies $p(z_k|s_t, a_t, \lambda_h^*) = 1$, $k = \arg\max_k p(a_t|s_t, \theta_{lk})$. This means that π_{θ_h} needs to predict the sub-skill category at time step t as k, in which case $\pi_{\theta_{lk}}$ can maximize $p(a_t|s_t, \theta_{lk})$. In case we choose π_{θ_h} to be a classifier that employs a Softmax layer at the end, minimizing the cross entropy loss 3 equals to maximize 24, thus 21 can be proved.

In the LI step, $p(z_k|s_t, \lambda_h)$ is fixed, and the data sets for optimizing $\theta_{l1}, \cdots, \theta_{lK}$ are also fixed as $\mathcal{D}_{2k} = \{(s_{ijt}, a_{ijt})|\hat{z_{i2t}} = k\}_{t=1}^{N_k}$. Thus we need to maximize each $p(a_t|s_t, \theta_{lk})$ with $\mathcal{D}_{2k}$. In case of $p(a_t|s_t, \theta_{lk}) \sim \mathcal{N}(\mu_{\theta_{lk}(s_t)}, \sigma^2) \propto \exp[-\frac{(a_t - \pi_{\theta_{lk}}(s_t))^2}{2\sigma^2}]$, we have $\max_{\theta_{lk}} p(a_t|s_t, \theta_{lk}) \Leftrightarrow \min_{\theta_{lk}} (a_t - \pi_{\theta_{lk}}(s_t))^2$, which leads to the loss function 4, thus 22 can be proved, which finishes the prove of Theorem 2.

C.5 Proof of the E-step of DMIL

According to Theorem 1, we aim to maximize 17 w.r.t λ_i from the initial point θ_i with coordinate gradient ascent. We here need to prove that in DMIL, we could also achieve the global maximum point of λ_i as in MAML. We first give out the following Lemma:

Lemma 2 Let x be the solution found by coordinate gradient descent of $f(x)$. Let $x_i, i = 1, \cdots, n$ be the n coordinate directions used in the optimization process. If $f(x)$ can be decomposed as:

$$f(x) = g(x) + \sum_{i=1}^{n} h_i(x), \tag{25}$$

where $g(x)$ is a differentiable convex function, and each $h_i(x)$ is a convex function of the coordinate direction x_i, then x is the global minimum of $f(x)$.

Proof: Let y be another arbitrary point, we have:

$$\begin{aligned} f(y) - f(x) &= g(y) + h(y) - (g(x) + h(x)) \\ &\geq \nabla_x g(x)^T (y - x) + \sum_{i=1}^{n} h_i(y_i) - h_i(x_i) \\ &= \sum_{i=1}^{n} (\nabla_i g(x)(y_i - x_i) + h_i(y_i) - h_i(x_i)) \\ &\geq 0. \end{aligned} \tag{26}$$

Now let's consider our problem. Consider

$$
\begin{aligned}
\log p(\mathcal{D}_i^{tr}|\theta_{ih},\theta_{il}) &= \log \sum_{t=1}^{T} p(a_t|s_t,\theta_{ih},\theta_{il})p(s_t) \\
&= \log \sum_{t=1}^{T} p(s_t) \sum_{k=1}^{K} p(z_k|s_t,a_t,\theta_{ih})p(a_t|s_t,\theta_{il}) \\
&\geq \sum_{t=1}^{T} [\log p(s_t) + \log \sum_{k=1}^{K} p(z_k|s_t,a_t,\theta_{ih})p(a_t|s_t,\theta_{il})] \\
&\geq \sum_{t=1}^{T} [\log p(s_t) + \sum_{k=1}^{K} \log p(z_k|s_t,a_t,\theta_{ih})p(a_t|s_t,\theta_{il})] \\
&= \sum_{t=1}^{T} [\log p(s_t) + \sum_{k=1}^{K} \log p(z_k|s_t,a_t,\theta_{ih}) + \sum_{k=1}^{K} \log p(a_t|s_t,\theta_{il})].
\end{aligned}
\tag{27}
$$

In our case, two coordinate directions are θ_{ih} and θ_{il}. Let's consider the terms inside the brackets. According to Lemma 2, we can think $\log p(s_t)$ as $g(x)$(here it equals to constant), $\sum_{k=1}^{K} \log p(z_k|s_t,a_t,\theta_{ih})$ as $h_1(x)$ and $\sum_{k=1}^{K} \log p(a_t|s_t,\theta_{il})$ as $h_2(x)$. Thus the optimum can be proved.

D Experiment Details

D.1 Environments

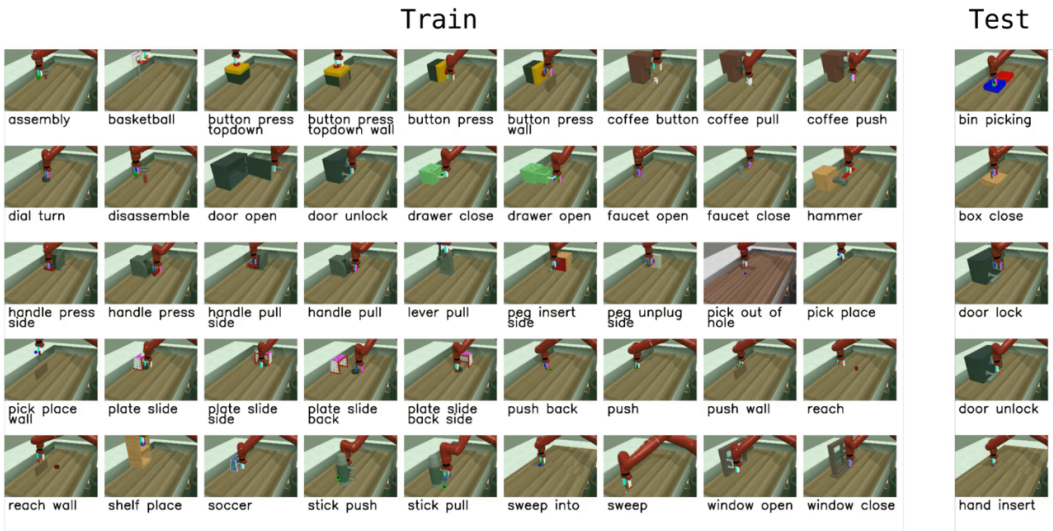

Figure 6: The ML45 environment.

D.2 Model Setup

DMIL: The high-level network and each sub-skill is modeled with a 4-layer fully-connected neuron network, with 512 ReLU units in each layer. We use Adam as the meta-optimizer. **DMIL-High** and **DMIL-Low** use the same architecture with DMIL. Hyper-parameter settings are available in table 2.

MIL: We use a transformer (Vaswani et al., 2017) as the policy to perform MAML. The input of the encoder is the whole one-shot demonstration or 3-shot demonstrations with concatenated state and

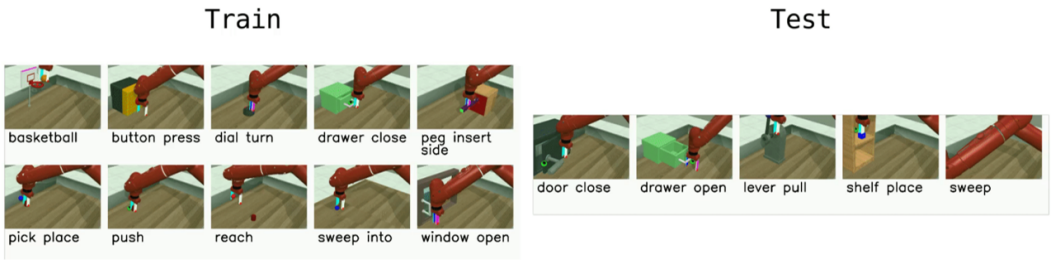

Figure 7: The ML10 environment.

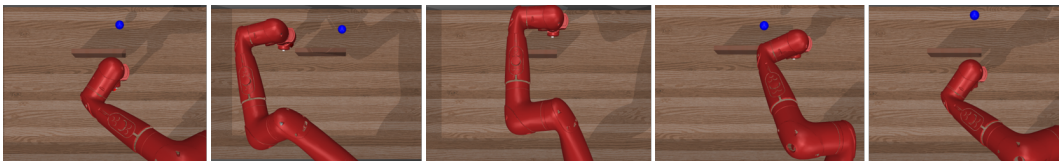

Figure 8: Task *push-around-wall*

Table 2: DMIL hyper-parameters.

Parameter	**Value**
α	5e-4
β	1e-4
fine-tune iterations	3
batch size (in trajectory)	16
λ	0.1

Table 3: MIL hyper-parameters.

Parameter	**Value**
n_{head}	8
n_{layer}	3
d_{model}	512
d_k	64
d_v	64
$n_{position}$	250
dropout	0.1
batch size (in state-action pair)	512

action. The input of the decoder is current state. The output is the predicted action. Hyper-parameter settings are available in table 3.

MLSH: We use the same settings of network with DMIL here. The macro step of high-level network is 3. Since our problem is not a reinforcement learning process, we use the *behavior cloning* variant of MLSH. The pseudo reward is defined by the negative mean square loss of the predicted action and the ground truth, and we perform pseudo reinforcement learning process with off-policy demonstration data. We use PPO as our reinforcement learning algorithm. Note in this way we ignore the importance sampling weights that required by replacing the sampling process in the environments with the demonstrations in the replay buffer, which has been shown to be effective in practice in Kostrikov et al. (2018); Ghasemipour et al. (2019). Hyper-parameter settings are available in table 4.

PEMIRL: We use the same setting of the high-level network as the policy π_ω and the inference model q_ψ in PEMIRL. We use PPO as our reinforcement learning algorithm. We use a 3-layer fully-connected neuron network as the context-dependent disentangled reward estimator $r_\theta(s, m)$ and the

Table 4: MLSH hyper-parameters.

Parameter	Value
high-level learning rate	1e-3
sub-skill learning rate	1e-4
PPO clip threshold	0.02
high-level warmup step	500
joint update step	1000
batch size (in state-action pair)	900

Table 5: PEMIRL hyper-parameters.

Parameter	Value
learning rate of all models	1e-4
PPO clip threshold	0.02
coefficient (γ) of h_ϕ	1
β	0.1
batch size (in trajectory)	16

context-dependent potential function $h_\phi(s, m)$. Here we also use the *behavior cloning* variant of PEMIRL to only train models on the off-policy data. Hyper-parameter settings are available in table 5.

D.3 Training Details

For fine-tuning, OptionGAIL, DMIL-Low and DMIL-High have no meta-learning mechanism for (some parts of) the trained model. In the few-shot adaptation process, we have different fine-tune method for these baselines:

OptionGAIL: We train OptionGAIL models on the provided few-shot demonstrations for a few epochs.

DMIL-Low: We fix the high-level network and only fine-tune sub-skills with few-shot demonstrations.

DMIL-High: We fix sub-skills and only fine-tune the high-level network with few-shot demonstrations.

D.4 Detailed Results

We provided detailed success rates of each environments of different methods in figure 9 and 10.

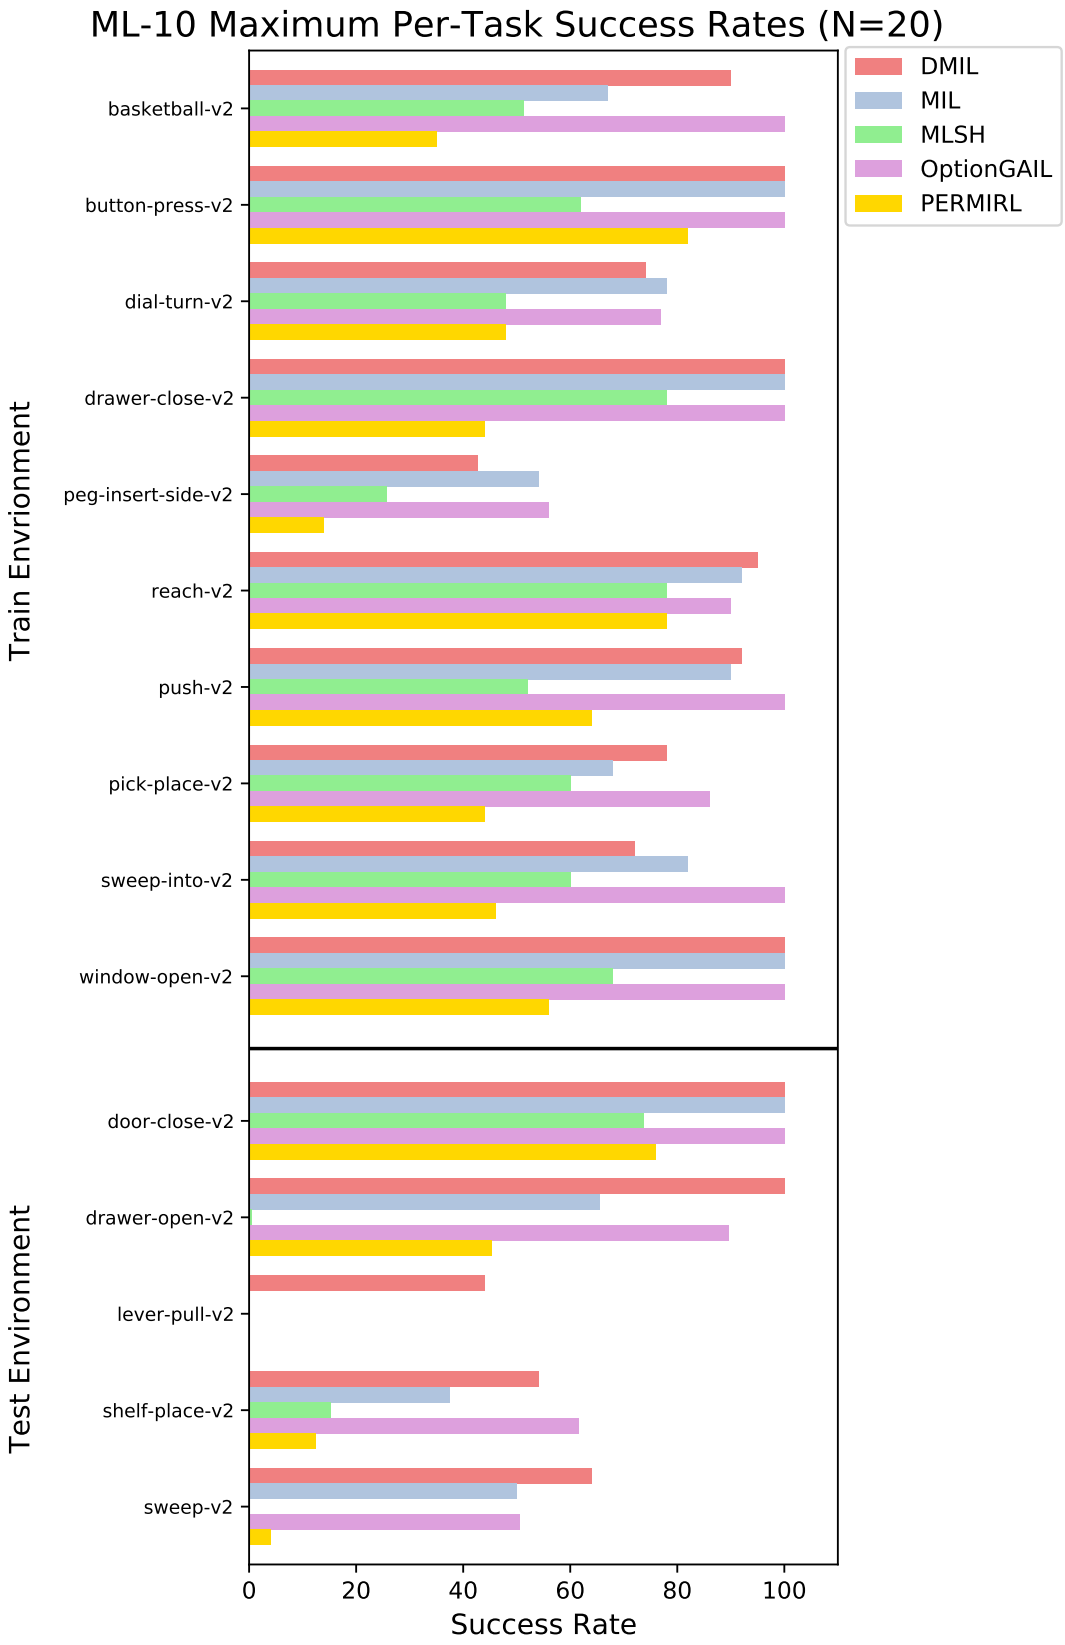

Figure 9: The ML10 results of different methods after 3-shot adaptation.

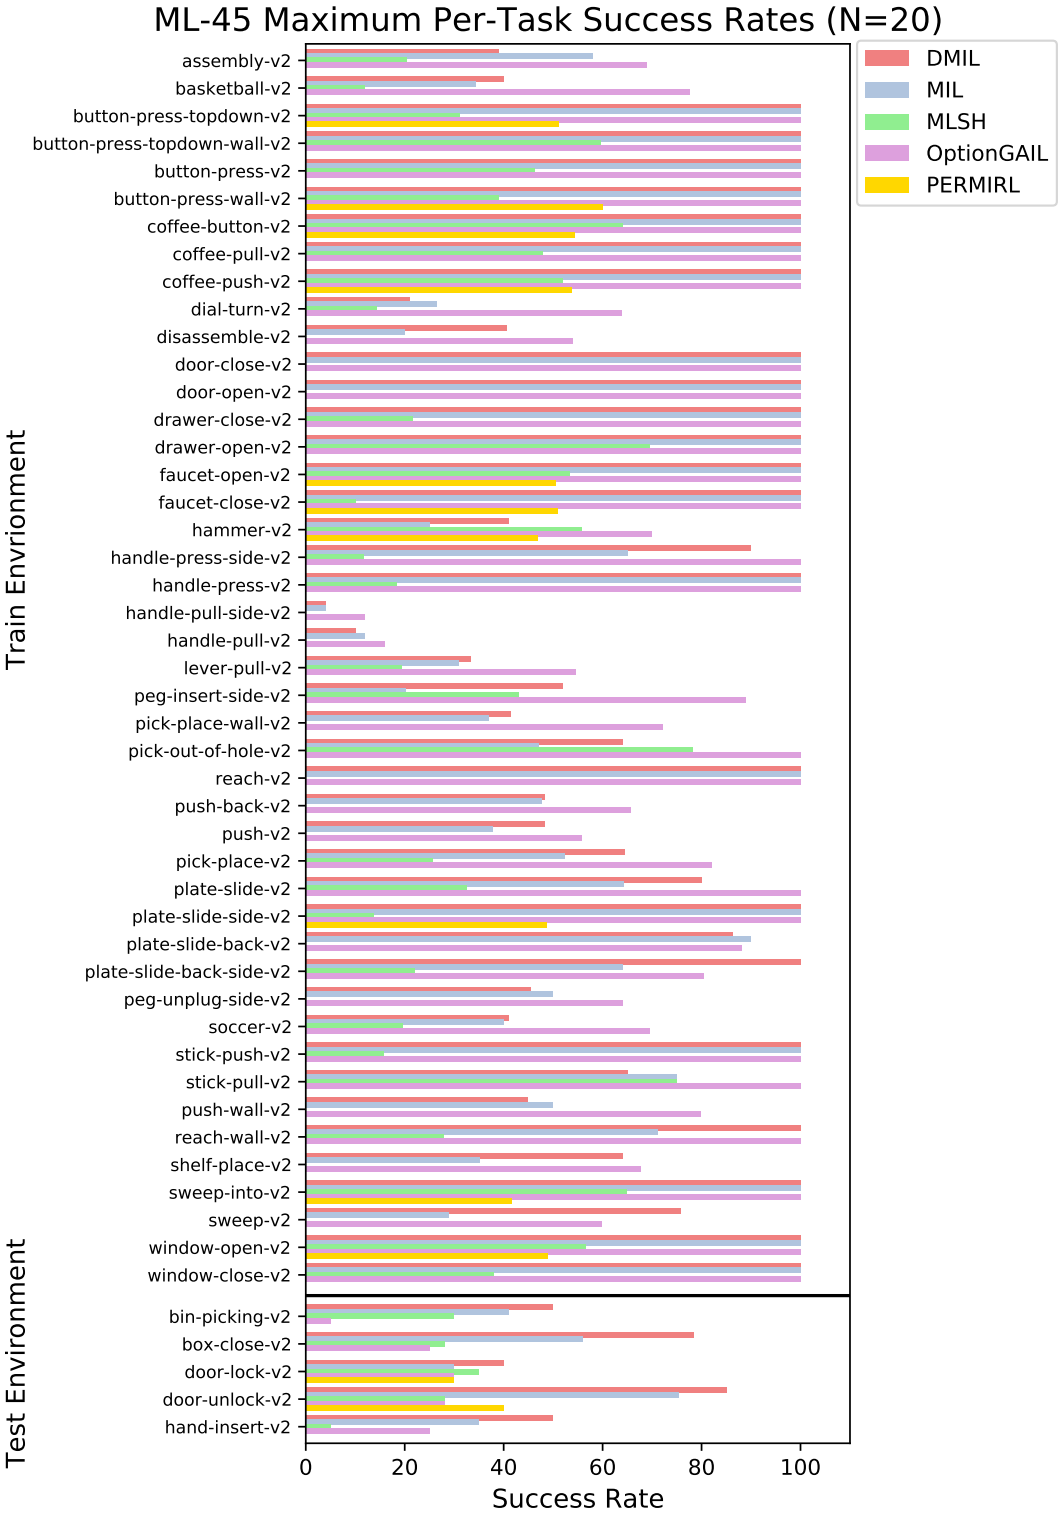

Figure 10: The ML45 results of different methods after 3-shot adaptation.