

NoRML: No-Reward Meta Learning

Yuxiang Yang, Ken Caluwaerts, Atil Iscen, Jie Tan, Chelsea Finn

Robotics at Google

{yxyang,kencaluwaerts,atil,jietan,chelseaf}@google.com

ABSTRACT

Efficiently adapting to new environments and changes in dynamics is critical for agents to successfully operate in the real world. Reinforcement learning (RL) based approaches typically rely on external reward feedback for adaptation. However, in many scenarios this reward signal might not be readily available for the target task, or the difference between the environments can be implicit and only observable from the dynamics. To this end, we introduce a method that allows for self-adaptation of learned policies: No-Reward Meta Learning (NoRML). NoRML extends Model Agnostic Meta Learning (MAML) for RL and uses observable dynamics of the environment instead of an explicit reward function in MAML’s finetune step. Our method has a more expressive update step than MAML, while maintaining MAML’s gradient based foundation. Additionally, in order to allow more targeted exploration, we implement an extension to MAML that effectively disconnects the meta-policy parameters from the fine-tuned policies’ parameters. We first study our method on a number of synthetic control problems and then validate our method on common benchmark environments, showing that NoRML outperforms MAML when the dynamics change between tasks.

Videos and source-code are available at <https://sites.google.com/view/noreward-meta-rl/>.

KEYWORDS

Deep Learning; Reinforcement Learning; Meta-Learning

ACM Reference Format:

Yuxiang Yang, Ken Caluwaerts, Atil Iscen, Jie Tan, Chelsea Finn. 2019. NoRML: No-Reward Meta Learning. In *Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), Montreal, Canada, May 13–17, 2019*, IFAAMAS, 9 pages.

1 INTRODUCTION

Adapting to new environments is a crucial capability for autonomous robots to operate in the real world. For example, after a robot learns to walk, its dynamics may change due to hardware failures, its sensor measurements (e.g. IMU) may drift over time, and more importantly, the reward signals may no longer be available due to lack of corresponding sensors after the robot is deployed (e.g. a tracking system that measures walking distance). How can the robot learn to adapt even after these changes?

Model Agnostic Meta Learning (MAML) [10] tackles the above problem by training a *meta-policy* that is optimized for quick adaptation to new tasks. During adaptation, this meta-policy can be fine-tuned efficiently with a small amount of data that is collected in the new environment. While MAML is successful at adapting

policies to different tasks that are defined by reward changes (e.g. running forward vs. backward), it is less effective when adapting to other changes [9], such as dynamics changes, sensor drifts, or missing reward signals.

In this paper, we introduce No-Reward Meta Learning (NoRML) to address the above challenges. The key insight underlying NoRML is that we can simultaneously learn the meta-policy and the advantage function used for adapting the meta-policy, optimizing for the ability to effectively adapt to varying dynamics. The meta-learned advantage function serves two purposes. First, it allows the policy to be adapted without external rewards: even if the reward signal is not present during the adaptation stage, we can use the learned advantage function to evaluate the current policy and compute the gradient. Second, the learned advantage gives more expressive power to the meta-optimization over how the policy can be updated, since it can learn to implicitly shape its reward feedback. As a result, the meta-policy can better recognize and adapt to subtle dynamics changes.

Beyond learning to update the policy parameters, one critical aspect of the meta-RL problem is the sampling distribution induced by the meta-policy, i.e. how the meta-policy chooses to explore and collect trajectory samples that are maximally informative about the unknown task or environment [12, 27]. In the original MAML formulation, the same policy parameters are used for collecting experience as for the gradient computation, limiting the extent to which the meta-policy parameters can be used for each individual purpose. To further increase the expressive power of the meta-optimization, we propose to decouple these two roles by augmenting the adaptation process with a meta-learned parameter offset. The parameter offset is added to the initial MAML parameters after sampling and during gradient adaptation, allowing different parameter vectors to be used for each while defaulting to the case where they are the same, akin to how residual networks (ResNets) [13] default to the identity function. With the combination of the offset policy and the learned advantage function, NoRML can successfully adapt to changes in dynamics, sensor drifts, and missing reward signals.

We evaluate NoRML on three control domains with varying sources of dynamics changes: an illustrative point agent example with disoriented actions, a cartpole with sensor bias, and a half-cheetah with wiring errors. In comparison to MAML, NoRML enables adaptation from a single trial, does not require reward signal for adaptation, and in most cases even leads to improved asymptotic performance. We find that both the learned advantage function and the parameter offset are important for good performance.

2 RELATED WORK

Algorithms for learning to learn [1, 14, 25, 31], or meta learning, aim to acquire a procedure that can more efficiently and effectively learn to solve new tasks. We consider meta learning in the context of reinforcement learning, i.e. meta reinforcement learning [8, 33]. Prior

Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), N. Agmon, M. E. Taylor, E. Elkind, M. Veloso (eds.), May 13–17, 2019, Montreal, Canada. © 2019 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

model-free meta reinforcement learning algorithms can generally be categorized as being recurrence-based [8, 19, 27, 33], gradient-based [10, 12, 27], or a hybrid of the two [15, 28]. We build a gradient-based meta-RL method that extends the MAML algorithm [10]. Unlike these prior model-free meta-RL works, we focus on the problem of learning to adapt to different dynamics, rather than adapting to different rewards.

Prior model-based RL approaches have considered the problem of learning to adapt to different dynamics through meta-RL [4, 24] or through learned priors [11]. Using a learned model is suitable when sample efficiency is a concern, but achieves lower asymptotic performance than model-free meta-RL [4]. Recent work by Clavera et al. [5] used MAML to adapt to different learned dynamics models within an ensemble, to improve model-based RL. We also consider the problem of adapting to different dynamics, but in the context of adapting to different environments, rather than different estimated models of the same environment. Further, our method improves upon MAML by not requiring a reward function for adaptation.

Separate from the meta learning literature, many other works have considered the problem of adapting to varying dynamics through, e.g. adaptive inverse control [34], self-modeling [2], Bayesian optimization [6], or online system identification [18, 35]. Our approach leverages prior experience to learn to adapt a policy with as little as a single trial without observed reward, and has few assumptions about the nature of the dynamics changes. Other methods have leveraged prior experience to learn a single policy that is robust to many different dynamics [20–22, 30]. Our experiments illustrate several realistic scenarios where robustness is not sufficient and adaptation is critical to good performance.

3 PRELIMINARIES

In this section, we overview model-free reinforcement learning, describe gradient-based meta learning (MAML), and discuss the potential difficulties that the vanilla MAML algorithm can have in model-free RL scenarios. We also introduce notation.

3.1 Model-free Reinforcement Learning

We study reinforcement learning problems where the agent makes a sequence of actions in a stochastic environment in order to maximize the cumulative reward. Formally, we define the problem as a Markov decision process (MDP) which consists of: a state space \mathcal{S} , an action space \mathcal{A} , a transition probability distribution $p(s_{t+1}|s_t, \mathbf{a}_t)$, a reward function $R: \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ and an initial state distribution $p(s_1)$.

In model-free RL, we aim to directly optimize a policy $\pi_\theta: \mathcal{S} \rightarrow \mathcal{P}(\mathcal{A})$, where $\mathcal{P}(\mathcal{A})$ is the set of probability distributions on the action space, and $\theta \in \mathbb{R}^n$ is a vector to parameterize the policy. The agent interacts with the environment over a finite horizon of length H and collects a trajectory $D = \{s_1, \mathbf{a}_1, r_1, \dots, s_{H+1}\}$ over $\mathcal{S} \times \mathcal{A} \times \mathcal{R}$. The agent tries to find policy parameters θ that maximize the expected return. Equivalently, we can instead minimize the expected loss, which can be written as:

$$\mathcal{L}(\theta, D) = -\mathbb{E}_{(s_t, \mathbf{a}_t) \sim \pi_\theta} \left[\sum_{t=1}^H R(s_t, \mathbf{a}_t) \right] = -\mathbb{E}_{(s_t, \mathbf{a}_t) \sim \pi_\theta} \left[\sum_{t=1}^H r_t \right]. \quad (1)$$

Policy gradients is a popular model-free algorithm to optimize a policy. The algorithm approximates the gradient $\nabla_\theta \mathcal{L}(\theta, D)$ using the *Policy Gradient Theorem* [29]:

$$\nabla_\theta \mathcal{L}(\theta, D) = -\mathbb{E}_{(s_t, \mathbf{a}_t) \sim \pi_\theta} \left[A^\pi(s_t, \mathbf{a}_t) \nabla_\theta \log \pi_\theta(\mathbf{a}_t | s_t) \right], \quad (2)$$

The expectation in Eq. 2 is computed by Monte-Carlo estimation, where we replaced the reward signal r_t with the advantage function $A^\pi(s_t, \mathbf{a}_t)$. To reduce the variance of the gradient estimations, an advantage function of the following form can be used:

$$A^\pi(s_t, \mathbf{a}_t) = \sum_{t'=t}^H \gamma^{t'-t} r_{t'} - V^\pi(s_t), \quad (3)$$

where $V^\pi(s)$ is a fitted value function estimator (also called *critic*) and γ is a discount factor.

3.2 Gradient-based Meta Learning

Meta learning algorithms optimize for a learning procedure that can quickly adapt to a particular task. Assuming that the training and testing tasks share some commonalities and are sampled from the same distribution, meta learning algorithms aim to learn the structure underlying the tasks and use this knowledge for fast learning. The meta-training process usually involves drawing data from different tasks and optimizing for performance after learning with a small amount of data.

Model-agnostic meta learning (MAML) [10] takes a gradient-based approach to the above problem. Formally, given a distribution over tasks $p(\mathcal{T})$, where each task defines a specific loss function $\mathcal{L}_{\mathcal{T}_i}$, MAML aims to find a good meta parameters θ that, with one step of gradient descent, can adapt to specific tasks with small amounts of data. The objective can be described as follows:

$$\arg \min_{\theta} \mathbb{E}_{\mathcal{T}_i \sim p(\mathcal{T})} \left[\mathcal{L}_{\mathcal{T}_i}(\theta_i) \right] \text{ s.t. } \theta_i = \theta - \alpha \nabla_{\theta} \mathcal{L}_{\mathcal{T}_i}(\theta). \quad (4)$$

To find a good set of meta parameters, θ , MAML uses gradient descent on the meta objective in Eq. 4. This requires second-order derivatives w.r.t. θ . When presented with data for a new test task \mathcal{T}_j , MAML adapts by simply performing one step of gradient descent starting from θ . Since the formulation of MAML is quite general, it can be applied to a range of problems, including model-free RL, which we describe next.

3.3 MAML on Model-free RL

Applying MAML in the context of model-free reinforcement learning (MAML-RL), θ parameterizes the *meta-policy* π_θ . To perform task-specific fine-tuning for a test task, one collects trajectories (*meta rollouts*) using the meta-policy π_θ on a sampled task \mathcal{T}_i and then uses the policy gradient equations (Eq. 2) to obtain the *fine-tuned* (also called *adapted*) policy π_{θ_i} . We can write the adaptation step for task \mathcal{T}_i explicitly by applying the gradient update rule in Eq. 4 to model-free RL and replacing the expectation with Monte-Carlo estimation. More precisely, we collect K trajectories $D_i^{\text{train}} = \{(s_1, \mathbf{a}_1, r_1, s_2, \mathbf{a}_2, r_2, \dots, s_{H+1})_k : k = 1 \dots K\}$ using the meta-policy π_θ on task \mathcal{T}_i and approximate the policy gradient:

$$\theta_i = \theta - \alpha \nabla_{\theta} \mathcal{L}_{\mathcal{T}_i}(\theta, D_i^{\text{train}}) \quad (5)$$

$$= \theta + \alpha \sum_{(s_t, \mathbf{a}_t, r_t) \in D_i^{\text{train}}} A^\pi(s_t, \mathbf{a}_t) \nabla_{\theta} \log \pi_\theta(\mathbf{a}_t | s_t). \quad (6)$$

During meta training, one also collects rollouts of the fine-tuned policies θ_i for \mathcal{T}_i to compute the meta objective in Eq. 4. As stated in Eq. 4, MAML optimizes the meta-policy parameters θ such that the expected loss over all tasks *after* adaptation is minimized. To this end, we use K trajectories D_i^{test} from the adapted policies π_{θ_i} to approximate the gradient of the meta-objective:

$$\theta \leftarrow \theta - \beta \sum_{\mathcal{T}_i \sim \mathcal{T}} \nabla_{\theta} \mathcal{L}_{\mathcal{T}_i}(\theta_i, D_i^{\text{test}}), \quad (7)$$

where

$$\mathcal{L}_{\mathcal{T}_i}(\theta_i, D_i^{\text{test}}) = \mathcal{L}_{\mathcal{T}_i}(\theta - \alpha \nabla_{\theta} \mathcal{L}_{\mathcal{T}_i}(\theta, D_i^{\text{train}}), D_i^{\text{test}}). \quad (8)$$

The gradient of $\mathcal{L}_{\mathcal{T}_i}(\theta_i, D_i^{\text{test}})$ can be approximated using the policy gradient algorithm (c.f. Eq. 6).

The complete MAML algorithm for reinforcement learning is listed in Algorithm 1. Algorithm 2 lists how to perform fine-tuning on a specific task, based on an optimized meta-policy. Unlike the original MAML-RL study [10] where different tasks corresponded to different rewards, our goal is to handle the setting where a different task entails a change of the environment, including both dynamics changes and sensor drifts, and where rewards are not present during adaptation, a scenario that MAML-RL cannot handle effectively.

Algorithm 1 MAML Training [10]

Require: $p(\mathcal{T})$: task distribution

Require: α : adaptation learning rate

Require: β : meta learning rate

Randomly initialize θ

while not done **do**

 Sample a batch of tasks $\mathcal{T}_i \sim p(\mathcal{T})$

for all \mathcal{T}_i in batch **do**

 Sample K trajectories D_i^{train} using π_{θ} on task \mathcal{T}_i .

 Compute adapted parameters using D_i^{train} :

$\theta_i = \theta - \alpha \nabla_{\theta} \mathcal{L}_{\mathcal{T}_i}(\theta, D_i^{\text{train}})$.

 Sample K trajectories D_i^{test} using π_{θ_i} on task \mathcal{T}_i .

end for

 Update θ using all D_i^{train} , and D_i^{test} :

$\theta \leftarrow \theta - \beta \sum_{\mathcal{T}_i \sim \mathcal{T}} \nabla_{\theta} \mathcal{L}_{\mathcal{T}_i}(\theta_i, D_i^{\text{test}})$

 with $\mathcal{L}_{\mathcal{T}_i}(\theta_i, D_i^{\text{test}}) = \mathcal{L}_{\mathcal{T}_i}(\theta - \alpha \nabla_{\theta} \mathcal{L}_{\mathcal{T}_i}(\theta, D_i^{\text{train}}), D_i^{\text{test}})$.

end while

Algorithm 2 MAML Fine-tuning [10]

Require: \mathcal{T}_i : a new test task

Require: θ, α : from MAML training

 Sample K trajectories D using π_{θ} on task \mathcal{T}_i .

 Compute fine-tuned policy:

$\theta \leftarrow \theta - \alpha \nabla_{\theta} \mathcal{L}_{\mathcal{T}_i}(\theta, D)$.

4 NO-REWARD META LEARNING

Consider an agent with a single task (i.e. a fixed reward function) such as running forward. Intuitively, if the dynamics change such as calibration errors or motor malfunctions, an agent should not need reward supervision in order to adapt its behavior: the dynamics change is recognizable from the state-action transitions alone.

Similarly, a human can adapt to a new terrain without external reward feedback. However, model-free RL requires such rewards. Our goal is to develop a model-free meta-RL algorithm that can learn to quickly adapt a policy to dynamics changes and sensor drifts *without external rewards*. To do so, the meta learning algorithm needs to develop its own internal notion of reward, to learn to explore in a way that is maximally informative of the current conditions, and to be able to learn to recognize changes in dynamics and adapt appropriately.

In this section, we introduce our meta reinforcement learning algorithm, termed No-Reward Meta Learning (NoRML), that aims to address these challenges. NoRML consists of two additional components to the original MAML-RL formulation: a learned advantage function that internalizes the reward in a way that allows for reward-free adaptation (which we discuss next), and a learned parameter offset that enables better exploration (which we discuss in Section 4.2). The entire meta-training algorithm and meta-test procedure of NoRML are summarized in Algorithms 3 and 4. Note that we use the term “change of the environment” and “different tasks” interchangeably in the following discussion.

4.1 Learned Advantage Function

We introduce a *learned advantage function* $A_{\psi}(s_t, \mathbf{a}_t, s_{t+1})$ to replace the estimated advantage $A^{\pi}(s_t, \mathbf{a}_t)$ in Eq. 6. The reason is twofold. First, the learned advantage function can be used to evaluate a trajectory even if the reward signal is not present during adaptation. Thus, it solves the problem of missing reward signals. Second, it is a generalized function form for the advantage function, which considerably increases the expressiveness of the policy gradient for fine-tuning, giving the meta-optimization more control over how it can update the policy. A_{ψ} is a feed-forward neural network that takes in consecutive states and action $(s_t, \mathbf{a}_t, s_{t+1})$. We initialize the weights ψ of the advantage network randomly and train it end-to-end. More specifically, during the meta-training process, we adjust its weights according to the gradient of the meta objective (See Alg. 3 for more details). Note that the learned advantage A_{ψ} is only used during fine-tuning, while the reward-based advantage A^{π} is still used to compute the outer gradient during meta training.

Since A_{ψ} takes in $(s_t, \mathbf{a}_t, s_{t+1})$ as input, this allows A_{ψ} to detect changes in the dynamics $p(s_{t+1}|s_t, \mathbf{a}_t)$, and provide a more informed “evaluation” of the actions, compared to using only (s_t, \mathbf{a}_t) . An additional benefit of using A_{ψ} is that we eliminate the need to estimate the value function to calculate the observed advantage. In MAML, it is difficult to estimate a value function from only a few roll-outs, limiting the effectiveness of the resulting policy gradient. A_{ψ} directly transforms the policy gradient and can provide accurate information to the fine-tune step even when the sample size is small.

Given the learned advantage function A_{ψ} , The MAML adaptation step for the task \mathcal{T}_i is modified to be the following:

$$\theta_i = \theta - \alpha \nabla_{\theta} \mathcal{L}_{\mathcal{T}_i}^{\text{NoRML}}(\theta, D_i^{\text{train}}) \quad (9)$$

$$= \theta + \alpha \sum_{D_i^{\text{train}}} A_{\psi}(s_t, \mathbf{a}_t, s_{t+1}) \nabla_{\theta} \log \pi_{\theta}(\mathbf{a}_t | s_t), \quad (10)$$

where $D_i^{\text{train}} = \{(s_t, \mathbf{a}_t, s_{t+1})_k : k = 1 \dots K, t = 1 \dots H\}$ is generated based on K trajectories using the meta-policy π_θ on task \mathcal{T}_i (i.e. D_i^{train} contains KH state transitions).

We would like to emphasize that this learned advantage function differs fundamentally from approximating $A^\pi(s_t, \mathbf{a}_t)$ using a fitted value function. In the latter case, the learned advantage function is trained to predict the actual observed advantage values $A^\pi(s_t, \mathbf{a}_t)$ in Eq. 6. In contrast, our learned advantage function A_ψ is optimized to *transform* or *reshape* the policy gradient $\nabla_\theta \log \pi_\theta(\mathbf{a}_t | s_t)$ in a way that achieves more effective adaptation in a single fine-tune step. As a result, the output of our advantage network can be significantly different from the estimated advantage values (in a sense, A_ψ is not a true advantage function).

Algorithm 3 NoRML Training

Differences between MAML and NoRML are [highlighted](#).

Require: $p(\mathcal{T})$: distribution over tasks

Require: α, β : step size hyperparameters

Require: θ_σ : initial log standard deviation of meta-policy

Randomly initialize θ_μ and set $\theta = \begin{bmatrix} \theta_\mu^T & \theta_\sigma^T \end{bmatrix}^T$

Randomly initialize ψ

Initialize θ_{offset} to $[0 \dots 0]^T$

while not done **do**

 Sample a batch of tasks $\mathcal{T}_i \sim p(\mathcal{T})$

for all \mathcal{T}_i **do**

 Sample K trajectories [without rewards](#) using π_θ on task \mathcal{T}_i and store all state transitions as a set $D_i^{\text{train}} = \{(s_t, \mathbf{a}_t, s_{t+1}) : \forall k < K, \forall t \leq H\}$.

 Compute adapted parameters using D_i^{train} :

$$\theta_i = \theta + \theta_{\text{offset}} - \alpha \sum_{D_i^{\text{train}}} A_\psi(s_t, \mathbf{a}_t, s_{t+1}) \nabla_\theta \log \pi_\theta(\mathbf{a}_t | s_t).$$

 Sample K trajectories D_i^{test} using π_{θ_i} on task \mathcal{T}_i .

end for

 Update $\theta, \theta_{\text{offset}}$, and ψ using all D_i^{train} , and D_i^{test} :

$$\begin{bmatrix} \theta \\ \theta_{\text{offset}} \\ \psi \end{bmatrix} \leftarrow \begin{bmatrix} \theta \\ \theta_{\text{offset}} \\ \psi \end{bmatrix} - \beta \sum_{\mathcal{T}_i \sim \mathcal{T}} \nabla_{[\theta^T \theta_{\text{offset}}^T \psi^T]^T} \mathcal{L}_{\mathcal{T}_i}(\theta_i, D_i^{\text{test}})$$

with $\nabla \mathcal{L}_{\mathcal{T}_i}(\theta_i, D_i^{\text{test}}) = \sum_{D_i^{\text{test}}} A^\pi(s_t, \mathbf{a}_t) \nabla \log \pi_{\theta_i}(\mathbf{a}_t | s_t)$.

end while

Algorithm 4 NoRML Fine-tuning

Require: \mathcal{T}_i : a task

Require: $\theta, \theta_{\text{offset}}, \psi, \alpha$: from NoRML training

 Sample K trajectories without rewards using π_θ and store all state transitions as a set $D = \{(s_t, \mathbf{a}_t, s_{t+1}) : \forall k < K, \forall t < H\}$.

 Compute fine-tuned policy:

$$\theta \leftarrow \theta + \theta_{\text{offset}} - \alpha \sum_D A_\psi(s_t, \mathbf{a}_t, s_{t+1}) \nabla_\theta \log \pi_\theta(\mathbf{a}_t | s_t).$$

4.2 Offset Learning

Since one policy gradient step may be insufficient to adapt an exploratory meta-policy into a policy for the new task, we introduce a simple, yet effective, technique to decouple the meta-policy from

the adapted policies: a learned offset θ_{offset} that is added to the policy parameters θ when calculating an adapted policy:

$$\theta_i = \theta + \theta_{\text{offset}} - \alpha \sum_{D_i^{\text{train}}} A_\psi(s_t, \mathbf{a}_t, s_{t+1}) \nabla_\theta \log \pi_\theta(\mathbf{a}_t | s_t). \quad (11)$$

Note that the policy offset θ_{offset} is shared for *all tasks*. Hence, it does not influence task-specific adaptation. The adaptation step is still based on trajectories sampled from the meta-policy π_θ , and the policy gradient is still computed with respect to the meta policy parameters θ . Similar to the learned advantage function, the parameter offset is optimized end-to-end together with the meta-policy, as shown in Algorithm 3. Fig. 1 shows a geometric interpretation of MAML and NoRML.

5 EXPERIMENTAL SETUP

In this section, we describe the comparisons and implementation details of our experiments.

5.1 Comparisons

We compare NoRML to two existing approaches: vanilla MAML [10] and Domain Randomization [21, 23, 32] (DR). Domain randomization aims to learn a single robust policy by varying the environment for each rollout. We implement domain randomization by setting the adaptation learning rate α and the policy offset θ_{offset} to zero (and disabling meta learning of these parameters) in our MAML implementation. Hence the meta-policy is directly used to compute the average loss across tasks/randomizations. This eliminates other factors that could influence experimental results and ensures that we are doing a fair comparison.

In addition, we also perform an ablation study for different components of NoRML. We refer to them as *NoRML w/o offset*, which uses a learned advantage function but does not include offset as a trainable parameter, and *NoRML w/o LAF*, which, like MAML, uses ground-truth external reward but also includes the offset.

For a fair comparison, all algorithms are trained for the same number of iterations and with the same number of timesteps collected per iteration. For all experiments, we randomly sweep the following three hyperparameters: the outer learning rate β , the adaptation learning rate α , and the initial value of the policy standard deviation θ_σ . We then plot the learning curves using the best hyperparameters found.

5.2 Implementation Details

We represent our policy as a multivariate diagonal Gaussian distribution and use a fully-connected, feed-forward network to map states s_t to a distribution over action $\pi_\theta(\mathbf{a}_t | s_t)$. The neural network outputs the mean of the Gaussian policy, and we used standalone variables to represent the standard deviations of each dimension:

$\pi_\theta(\mathbf{a}_t | s_t) = \mathcal{N}(f(s_t | \theta_\mu), \text{diag}(e^{\theta_\sigma^2}))$ ($\theta = \begin{bmatrix} \theta_\mu^T & \theta_\sigma^T \end{bmatrix}^T$). We found this to greatly improve training stability, compared to having the network output both the mean and log standard deviation. We use a two-layer fully connected network with tanh activation function for the policy network (50 neurons per layer). Similarly, the learned advantage function A_ψ uses a fully-connected, two-layer neural network with rectifying linear units (50 neurons per layer).

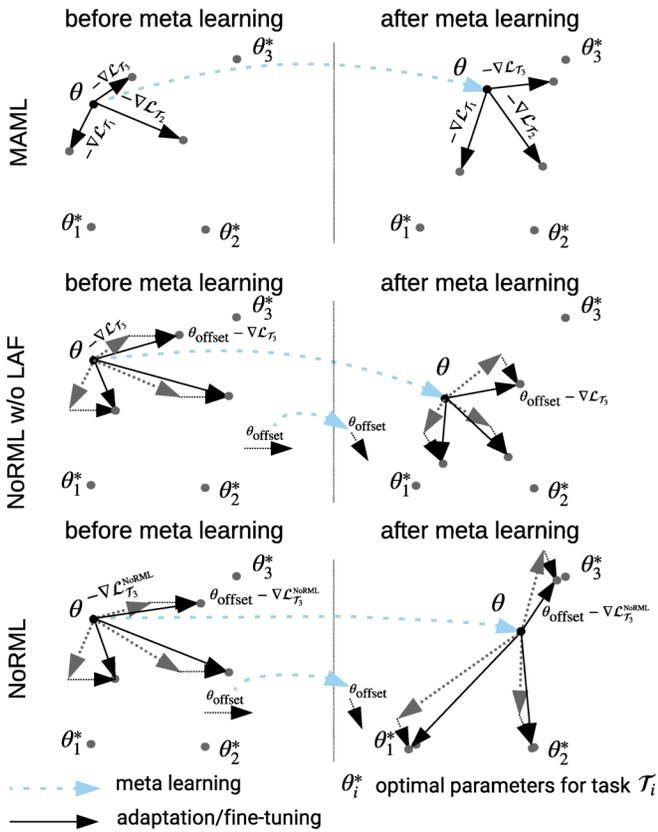


Figure 1: Geometric interpretations of MAML, NoRML w/o Learned Advantage Function (LAF), and NoRML. MAML (top) optimizes a meta-policy θ such that a single fine-tune step using the policy gradient $\nabla_{\theta} \mathcal{L}_{\mathcal{T}_i}$ is likely to significantly improve the performance on a specific task \mathcal{T}_i . NoRML without LAF (middle) learns an additional parameter vector θ_{offset} . This vector is added to the meta-policy parameters θ and decouples the meta-policy from the fine-tuned policies. NoRML (bottom) learns an advantage function during meta learning, which results in a modified, more expressive policy gradient for fine-tuning. We also use a policy offset vector θ_{offset} with NoRML to maximize performance.

For our MAML implementation, we also included the Meta-SGD extension by Li et al. [17], which replaces the fixed inner learning rate of MAML with a learned vector of the same dimension as the policy parameter. In practice, hyperparameter optimization of the initial value of α is required as tuning learning rates is a much slower process than optimizing the meta-policy θ . We used this extension in our experiments to improve the expressiveness of MAML’s adaptation step.

For meta training, we use Proximal Policy Optimization [26] (PPO) for MAML’s adaptation and meta objectives to improve performance. For NoRML PPO is only used for the meta objective as the learned advantage function A_{ψ} learns to transform the vanilla

policy gradient. As in the original MAML paper [10], we use polynomial regression per task to fit the value function [7] (in both the adaptation and meta learning steps for MAML and only in the meta learning step for NoRML). We use Adam [16] as our meta-policy optimizer and use vanilla policy gradient in the inner loop to avoid over-complicating the meta objective. To speed up computation, our MAML and NoRML implementations are parallelized across tasks and rollouts.

For all of our experiments, we use $K = 25$ rollouts for adaptation and meta learning. We sample 10 tasks \mathcal{T}_i during each meta-training iteration.

6 POINT AGENT CASE STUDY

We now introduce a simple, 2D point agent control problem to illustrate the advantage function learned by NoRML and study its effect with both shaped and sparse rewards. For this simple, synthetic control task, we show that NoRML learns an intuitive, realistic advantage function and achieves similar performance to MAML. As we increase the task’s difficulty by making the reward sparse, MAML struggles to learn, while NoRML still learns a similar advantage function that allows it to adapt to dynamics changes.

6.1 Task Setup

Consider an agent in a plane that is trying to move to the right from $(0, 0)$ to $(1, 0)$. The agent observes its current position $s_t = [x_t \ y_t]^T$, and each action specifies its movement $a_t = [dx_t \ dy_t]^T$, where $dx_t, dy_t \in [-1, 1]$. We introduce a dynamic change where the action is rotated by ϕ (one ϕ per task) that is unknown to the agent. The new dynamics can be expressed as:

$$\begin{bmatrix} x_{t+1} \\ y_{t+1} \end{bmatrix} = \begin{bmatrix} \cos \phi & -\sin \phi \\ \sin \phi & \cos \phi \end{bmatrix} \begin{bmatrix} dx_t \\ dy_t \end{bmatrix} + \begin{bmatrix} x_t \\ y_t \end{bmatrix}. \quad (12)$$

We also restrict the agent’s movement to a square region $-2 \leq x_{t+1}, y_{t+1} \leq 2$.

We consider two types of reward functions: shaped and sparse. In the former case, each rollout has a fixed horizon of 10 steps, and the reward at each step is the negative Euclidean distance to the destination $(1, 0)$. For the sparse reward case, the reward is -1 for each step taken, but an episode can terminate early when the agent successfully reaches within 0.1 radius to the goal. Each rollout has a maximum horizon of 100 steps.

As a meta learning problem, different tasks are defined by the unknown rotation ϕ , where the task distribution is uniform on $[0, 2\pi]$ radians. The agent needs to gather information about the rotation amount ϕ in the meta rollouts, and make corresponding changes to the policy during fine-tuning.

6.2 Impact of the Learned Advantage Function

To visualize the advantage function learned by NoRML, we transformed the original action space to polar coordinates and plotted the learned advantage functions, as seen in fig. 2b and 2d. In both cases, the agent is located at the origin and takes an action of length 1 and angle ω , which would move the agent from $[0, 0]^T$ to $[\cos \omega, \sin \omega]^T$. Therefore, an action with angle $\omega = 0$ would take the agent directly to the desired destination.

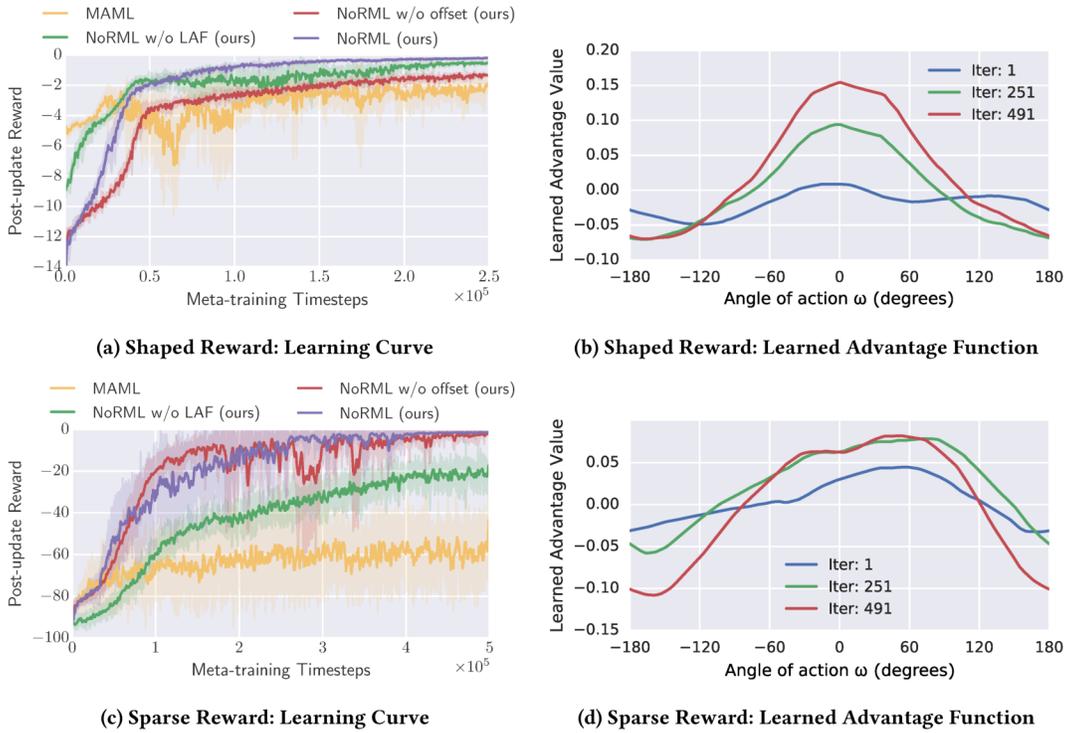


Figure 2: Point agent: Meta-training curves and the learned advantage function for held-out tasks. LAF means "learned advantage function" (see section 5.1 for details). The advantage values are plotted on actions of length 1 on origin (0, 0), evaluated on the task with no rotation ($\phi = 0$). Therefore, an action at angle 0° takes the point to the goal (1, 0), and an action at angle 180° takes the point to (-1, 0). NoRML is able to learn a shaped advantage function that leads to effective adaptation to dynamics, even when only sparse rewards are provided.

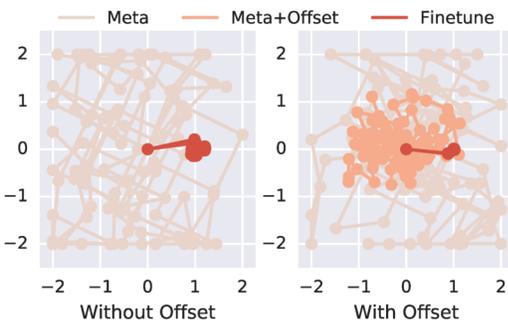


Figure 3: 10 rollout trajectories for NoRML policies trained with and without the offset. "Meta+Offset" means we only add the learned offset θ_{offset} to the meta-policy parameters θ without a gradient step and evaluate policy $\pi_{\theta+\theta_{\text{offset}}}$.

In the shaped-reward case (Fig. 2b), we see that the learned advantage function gradually converges to a bell-like shape peaked around 0° , which rewards actions that move the agent towards the destination. Note that the advantage network only takes in (s_t, a_t, s_{t+1}) as input and does not have access to the ground-truth

reward value. After meta training, it is able to learn a smooth advantage function that guides the meta-policy for proper fine-tuning.

As we make the task harder by introducing the sparse reward function, NoRML still learns a similarly-shaped reward function shown in Fig. 2d with peak value around angle 0° . In this more difficult task, MAML’s performance degrades dramatically (Fig. 2c). The sparse reward function is challenging for vanilla MAML to adapt to changes in dynamics, since it uses a single policy gradient step with a fitted value function. In contrast, the learned advantage function in NoRML provides a more informative reward signal that enables the agent to adapt in one policy gradient step.

6.3 Impact of the Offset

We find that the parameter offset θ_{offset} encourages exploration in NoRML. In Fig. 3, we train two NoRML policies, one with and one without a policy offset and we plot the trajectories sampled from these policies. The meta-policies in both cases tend to be exploratory and have larger variance, and the fine-tuned policies are more consistent and move directly to the destination. However, with offset enabled, the fine-tuning process in Fig. 3 learns to reduce variance even further, since the offset already reduces the explorative policy to a more conservative one.

The effect of the offset is further illustrated in Fig. 2a, where the offset allows the fine-tuned policy to converge to better final values. Without the offset, both MAML and NoRML could not achieve a total reward greater than -1 . When the policy offset is enabled, however, the learning curves converged to values as high as -0.4 , which is a significant improvement for this task.

7 CONTINUOUS CONTROL TASKS

To study how NoRML scales to more complex deep RL problems, we apply it to two continuous control problems in the OpenAI Gym [3], and compare it with vanilla MAML and domain randomization.

7.1 Cartpole with Sensor Bias

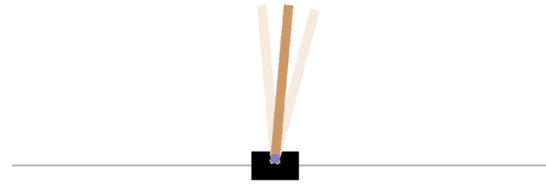
We introduce a variant of the Cartpole environment, in which the agent needs to move the cart to balance an inverted pendulum. For each time step, the agent observes the position and velocity of both the cart and the pole, and applies a force to the cart in order to balance the system. In our variant, the position sensor can drift: its reading can be offset by an unknown amount ranging from -10° to 10° (Fig. 4a). Hence, our meta-training task distribution corresponds to a uniform distribution over this range of sensor reading. We also make the task more difficult by increasing the required duration of balancing the pole from 4 seconds, as in the original OpenAI Gym environment, to 10 seconds.

As seen in Fig. 4b, although both vanilla MAML and NoRML converge to a reward of 500 in the end, NoRML converges faster despite having fewer assumptions—NoRML does not require an external reward signal for its adaptation. Domain randomization cannot solve the task in this case, demonstrating that adaptation is necessary to solve these tasks. The plot also shows an ablation study of different components: without the offset, NoRML could not converge to a high final reward, and without the learned advantage function, convergence is slower.

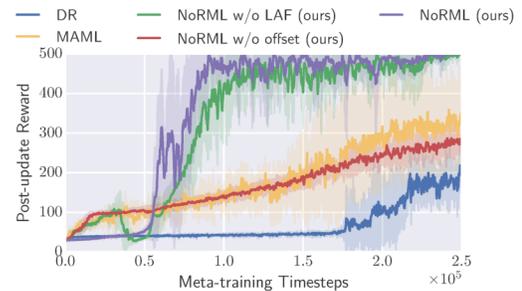
7.2 Half Cheetah with Swapped Actions

We next evaluate on the half cheetah environment in OpenAI Gym. To test NoRML’s adaptability to dynamic changes, we purposefully allow the torque outputs of the two hip joints to be swapped, leading to two different tasks. In real robotic problems, this change could occur due to wiring and signal transmission errors. Another challenge of robot locomotion in the real world is the unavailability of accurate, real-time on-board localization systems. To simulate this, we remove the position and linear velocity from half cheetah’s observation space. Note that the lack of localization information also makes it difficult to compute the distance-based reward function. Without reward function, MAML can no longer perform adaptation, while NoRML can still adapt using the learned advantage.

With the learned advantage function, NoRML significantly outperforms vanilla MAML both in convergence speed and final return (Fig. 5a). Moreover, although domain randomization achieved similar performance in terms of final return, we found the gait learned by domain randomization to be less stable (Fig. 5b): the body of the half cheetah oscillates a lot during running (Fig. 5c), as domain randomization must learn a single policy that can handle all action swaps. On the other hand, with NoRML, the policy can adapt



(a) The cartpole’s uncalibrated angle sensor adds an unknown bias to the actual reading. When the pole is at the position shown in the figure, the sensor’s reading can indicate any reading between the two blurred ones.



(b) Meta-training curve on held-out tasks: since the reward is 1 for every time step that the pole does not fall down, the total reward reflects the length of the episode and the maximum reward is 500. Note that NoRML only uses the cart and pole’s position and velocity during adaptation, while MAML additionally uses external rewards.

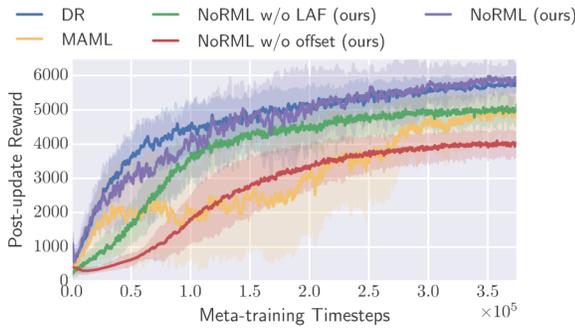
Figure 4: Illustration of the cartpole task and results. We find that adaptation is critical for this task, that our method converges faster, and that both the learned advantage and learned offset are helpful.

quickly and gracefully, without reward or tracking information, regardless of the actions being swapped.

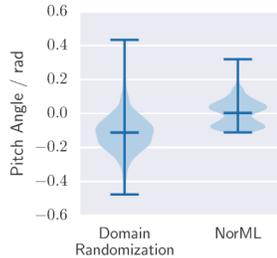
For the HalfCheetah policy, we also took the top-performing policy and tested its adaptation performance using a smaller number of rollouts. As shown in Fig. 6, even with as little as a single rollout, the fine-tuned policy still achieves a reasonable performance. This shows that the learned advantage function is noise-tolerant and can sense dynamics changes using a small amount of sample data, as little as a single trajectory. On the other hand, MAML needs at least 5 meta rollouts to achieve a similar post-adaptation performance, due to the noisy estimation of value function and advantage.

8 CONCLUSION

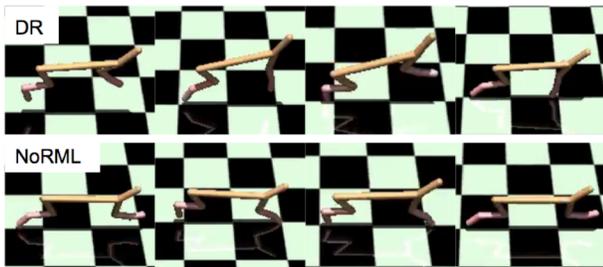
In this paper, we introduce NoRML, a new meta reinforcement learning algorithm that adapts to changes in dynamics and sensor drifts, without the need for external reward signals during adaptation. The key insight is to learn an advantage function, a parameter



(a) Meta-training curve for the half cheetah. In this case, MAML struggles to recognize and adapt to dynamics changes, while the learned advantage function enables effective adaptation. Domain randomization also leads to high reward, but cannot produce as stable of a gait (see below).



(b) Pitch distribution showing increased stability of the fine-tuned NoRML gait compared to the gait learned with Domain Randomization.



(c) Gait learned by Domain Randomization (top) and NoRML (bottom). The fine-tuned gait learned by NoRML is more stable and the body oscillates less.

Figure 5: Learning curve, IMU readings and snapshots of the running gait for the HalfCheetah environment.

offset between meta and adapted policies, and the meta policy simultaneously. The learned advantage function results in a more expressive adaptation step by generalizing MAML’s update rule. The parameter offset encourages exploration during meta rollouts by decoupling the fine-tuned policies from the meta-policy.

We evaluate NoRML on three control problems with varying types of dynamics changes: an illustrative point agent example with distorted actions, a cartpole with sensor bias, and a half-cheetah with wiring errors. Our experiments show that, by incorporating both a learned advantage and a learned offset, NoRML can adapt to

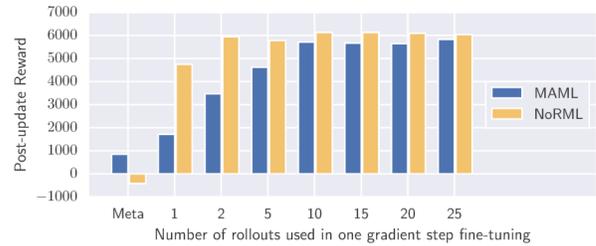


Figure 6: Post-update rewards for HalfCheetah using variable numbers of rollouts. NoRML can adapt effectively with as few as one roll-out, while MAML cannot adapt well using fewer than 5 rollouts.

all these types of changes, even when the reward signals are not present during adaptation.

A promising future research direction is to apply NoRML to transfer policies from simulation to real robots. Sim-to-real transfer is a challenging problem in robotics, which is caused by model errors between the simulation and the real-world physics. We can treat the model error as the dynamics change and apply NoRML to adapt to this change with a few shots of real robot data.

REFERENCES

- [1] Samy Bengio, Yoshua Bengio, Jocelyn Cloutier, and Jan Gecsei. 1992. On the optimization of a synaptic learning rule. In *Preprints Conf. Optimality in Artificial and Biological Neural Networks*. Univ. of Texas, 6–8.
- [2] Josh Bongard, Victor Zykov, and Hod Lipson. 2006. Resilient machines through continuous self-modeling. *Science* 314, 5802 (2006), 1118–1121.
- [3] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. 2016. OpenAI Gym. (2016). arXiv:arXiv:1606.01540
- [4] Ignasi Clavera, Anusha Nagabandi, Ronald S Fearing, Pieter Abbeel, Sergey Levine, and Chelsea Finn. 2018. Learning to Adapt: Meta-Learning for Model-Based Control. *arXiv preprint arXiv:1803.11347* (2018).
- [5] Ignasi Clavera, Jonas Rothfuss, John Schulman, Yasuhiro Fujita, Tamim Asfour, and Pieter Abbeel. 2018. Model-Based Reinforcement Learning via Meta-Policy Optimization. In *Proceedings of The 2nd Conference on Robot Learning (Proceedings of Machine Learning Research)*, Aude Billard, Anca Dragan, Jan Peters, and Jun Morimoto (Eds.), Vol. 87. PMLR, 617–629. <http://proceedings.mlr.press/v87/clavera18a.html>
- [6] Antoine Cully, Jeff Clune, Danesh Tarapore, and Jean-Baptiste Mouret. 2015. Robots that can adapt like animals. *Nature* 521, 7553 (2015), 503.
- [7] Yan Duan, Xi Chen, Rein Houthoofd, John Schulman, and Pieter Abbeel. 2016. Benchmarking deep reinforcement learning for continuous control. In *International Conference on Machine Learning*. 1329–1338.
- [8] Yan Duan, John Schulman, Xi Chen, Peter L Bartlett, Ilya Sutskever, and Pieter Abbeel. 2016. RL2: Fast Reinforcement Learning via Slow Reinforcement Learning. *arXiv preprint arXiv:1611.02779* (2016).
- [9] Chelsea Finn. 2018. *Learning to Learn with Gradients*. Ph.D. Dissertation. UC Berkeley.
- [10] Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017. Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks. *International Conference on Machine Learning (ICML)* (2017).
- [11] Justin Fu, Sergey Levine, and Pieter Abbeel. 2016. One-shot learning of manipulation skills with online dynamics adaptation and neural network priors. In *International Conference on Intelligent Robots and Systems (IROS)*. IEEE.
- [12] Abhishek Gupta, Russell Mendonca, YuXuan Liu, Pieter Abbeel, and Sergey Levine. 2018. Meta-Reinforcement Learning of Structured Exploration Strategies. *arXiv preprint arXiv:1802.07245* (2018).
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Identity mappings in deep residual networks. In *European conference on computer vision*. Springer, 630–645.
- [14] Sepp Hochreiter, A Steven Younger, and Peter R Conwell. 2001. Learning to learn using gradient descent. In *International Conference on Artificial Neural Networks*. Springer, 87–94.

- [15] Rein Houthoofd, Richard Y Chen, Phillip Isola, Bradley C Stadie, Filip Wolski, Jonathan Ho, and Pieter Abbeel. 2018. Evolved policy gradients. *arXiv preprint arXiv:1802.04821* (2018).
- [16] Diederik P. Kingma and Jimmy Ba. 2014. Adam: A Method for Stochastic Optimization. *CoRR* abs/1412.6980 (2014). arXiv:1412.6980 <http://arxiv.org/abs/1412.6980>
- [17] Zhenguo Li, Fengwei Zhou, Fei Chen, and Hang Li. 2017. Meta-SGD: Learning to Learn Quickly for Few Shot Learning. *CoRR* abs/1707.09835 (2017). arXiv:1707.09835 <http://arxiv.org/abs/1707.09835>
- [18] Lennart Ljung. 1998. System identification. In *Signal analysis and prediction*. Springer, 163–173.
- [19] Nikhil Mishra, Mostafa Rohaninejad, Xi Chen, and Pieter Abbeel. 2018. A simple neural attentive meta-learner. *International Conference on Learning Representations (ICLR)* (2018).
- [20] Igor Mordatch, Kendall Lowrey, and Emanuel Todorov. 2015. Ensemble-CIO: Full-body dynamic motion planning that transfers to physical humanoids. In *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on*. IEEE, 5307–5314.
- [21] Xue Bin Peng, Marcin Andrychowicz, Wojciech Zaremba, and Pieter Abbeel. 2018. Sim-to-real transfer of robotic control with dynamics randomization. In *International Conference on Robotics and Automation (ICRA)*.
- [22] Aravind Rajeswaran, Sarveer Ghotra, Balaraman Ravindran, and Sergey Levine. 2017. Epopt: Learning robust neural network policies using model ensembles. *International Conference on Learning Representations (ICLR)* (2017).
- [23] Fereshteh Sadeghi and Sergey Levine. 2017. CAD2RL: Real single-image flight without a single real image. *Robotics: Science and Systems (RSS)* (2017).
- [24] Steindór Sæmundsson, Katja Hofmann, and Marc Peter Deisenroth. 2018. Meta Reinforcement Learning with Latent Variable Gaussian Processes. *arXiv preprint arXiv:1803.07551* (2018).
- [25] Jürgen Schmidhuber. 1987. *Evolutionary principles in self-referential learning, or on learning how to learn: the meta-meta-... hook*. Ph.D. Dissertation. Technische Universität München.
- [26] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal Policy Optimization Algorithms. *CoRR* abs/1707.06347 (2017). arXiv:1707.06347 <http://arxiv.org/abs/1707.06347>
- [27] Bradley C Stadie, Ge Yang, Rein Houthoofd, Xi Chen, Yan Duan, Yuhuai Wu, Pieter Abbeel, and Ilya Sutskever. 2018. Some considerations on learning to explore via meta-reinforcement learning. *arXiv preprint arXiv:1803.01118* (2018).
- [28] Flood Sung, Li Zhang, Tao Xiang, Timothy Hospedales, and Yongxin Yang. 2017. Learning to learn: Meta-critic networks for sample efficient learning. *arXiv preprint arXiv:1706.09529* (2017).
- [29] Richard S Sutton, David A McAllester, Satinder P Singh, and Yishay Mansour. 2000. Policy gradient methods for reinforcement learning with function approximation. In *Advances in neural information processing systems*. 1057–1063.
- [30] Jie Tan, Tingnan Zhang, Erwin Coumans, Atıl İscen, Yunfei Bai, Danijar Hafner, Steven Bohez, and Vincent Vanhoucke. 2018. Sim-to-Real: Learning Agile Locomotion For Quadruped Robots. *Robotics: Science and Systems (RSS)* (2018).
- [31] Sebastian Thrun and Lorien Pratt. 1998. Learning to learn. In *Learning to learn*. Springer.
- [32] Joshua Tobin, Rachel Fong, Alex Ray, Jonas Schneider, Wojciech Zaremba, and Pieter Abbeel. 2017. Domain Randomization for Transferring Deep Neural Networks from Simulation to the Real World. *CoRR* abs/1703.06907 (2017). arXiv:1703.06907 <http://arxiv.org/abs/1703.06907>
- [33] Jane X Wang, Zeb Kurth-Nelson, Dhruva Tirumala, Hubert Soyer, Joel Z Leibo, Remi Munos, Charles Blundell, Dhharshan Kumaran, and Matt Botvinick. 2017. Learning to reinforcement learn. *CogSci* (2017).
- [34] Bernard Widrow. 1990. Adaptive inverse control. In *Applications of Artificial Neural Networks*, Vol. 1294. International Society for Optics and Photonics, 13–22.
- [35] Wenhao Yu, Jie Tan, C Karen Liu, and Greg Turk. 2017. Preparing for the unknown: Learning a universal policy with online system identification. *arXiv preprint arXiv:1702.02453* (2017).