

# One-Shot Visual Imitation Learning via Meta-Learning

Chelsea Finn<sup>\*1</sup>, Tianhe Yu<sup>\*1</sup>, Tianhao Zhang<sup>1</sup>, Pieter Abbeel<sup>1,2</sup>, Sergey Levine<sup>1</sup>

<sup>1</sup>University of California, Berkeley, <sup>2</sup>OpenAI  
cbfinn, tianhe.yu, tianhao.z, pabbeel, svlevine@berkeley.edu

**Abstract:** In order for a robot to be a generalist that can perform a wide range of jobs, it must be able to acquire a wide variety of skills quickly and efficiently in complex unstructured environments. High-capacity models such as deep neural networks can enable a robot to represent complex skills, but learning each skill from scratch then becomes infeasible. In this work, we present a meta-imitation learning method that enables a robot to learn how to learn more efficiently, allowing it to acquire new skills from just a single demonstration. Unlike prior methods for one-shot imitation, our method can scale to raw pixel inputs and requires data from significantly fewer prior tasks for effective learning of new skills. Our experiments on both simulated and real robot platforms demonstrate the ability to learn new tasks, end-to-end, from a single visual demonstration.

**Keywords:** imitation learning, deep learning, one-shot learning, meta-learning

## 1 Introduction

Enabling robots to be generalists, capable of performing a wide variety of tasks with many objects, presents a major challenge for current methods. Learning-based approaches offer the promise of a generic algorithm for acquiring a wide range of skills. However, learning methods typically require a fair amount of supervision or experience per task, especially for learning complex skills from raw sensory inputs using deep neural network models. Moreover, most methods provide no mechanism for using experience from previous tasks to more quickly solve new tasks. Thus, to learn many skills, training data would need to be collected independently for each and every task. By reusing data across skills, robots should be able to amortize their experience and significantly improve data efficiency, requiring minimal supervision for each new skill. In this paper, we consider the question: how can we leverage information from previous skills to quickly learn new behaviors?

We propose to combine meta-learning with imitation, enabling a robot to reuse past experience and, as a result, learn new skills from a single demonstration. Unlike prior methods that take the task identity [1, 2, 3, 4] or a demonstration [5] as the input into a contextual policy, our approach learns a parameterized policy that can be adapted to different tasks through gradient updates, effectively learning to imitation learn. As a result, the set of skills that can be learned is more flexible while using fewer overall parameters. For the first time, we demonstrate that vision-based policies can be fine-tuned end-to-end from one demonstration, using meta-learning as a pre-training procedure that uses demonstrations on a diverse range of other environments.

The primary contribution of this paper is to demonstrate an approach for one-shot imitation learning from raw pixels. We evaluate our approach on two simulated planar reaching domains, on simulated pushing tasks, and on visual placing tasks on a real robot (See Figure 1). Our approach is able to learn visuomotor policies that can adapt to new task variants using only one visual demonstration, including settings where only a raw video of the




Figure 1: The robot learns to place a new object into a new container from a single demonstration.

---

<sup>\*</sup>denotes equal contribution

demonstration is available without access to the controls applied by the demonstrator. By employing a parameter-efficient meta-learning method, our approach requires a relatively modest number of demonstrations for meta-learning and scales to raw pixel inputs. As a result, our method can successfully be applied to real robotic systems.

## 2 Related Work

We present a method that combines imitation learning [6] with meta-learning [7] for one-shot learning from visual demonstrations. Efficient imitation from a small number of demonstrations has been successful in scenarios where the state of the environment, such as the poses of objects, is known [8, 9, 10, 11]. In this work, we focus on settings where the state of the environment is unknown, where we must instead learn from raw sensory inputs. This removes the need for pre-defined vision systems while also making the method applicable to vision-based non-prehensile manipulation in unknown, dynamic environments. Imitation learning from raw pixels has been widely studied in the context of mobile robotics [12, 13, 14, 15]. However, learning from demonstrations has two primary challenges when applied to real-world settings. The first is the widely-studied issue of compounding errors [16, 17, 15], which we do not address in this paper. The second is the need for a large number of demonstrations for each task. This latter limitation is a major roadblock for developing generalist robots that can learn a wide variety of tasks through imitation. Inverse reinforcement learning [18] can reduce the number of demonstrations needed by inferring the reward function underlying a set of demonstrations. However, this requires additional robot experience to optimize the reward [19, 20, 21]. This experience typically comes in the form of trial-and-error learning or data for learning a model.

In this work, we drastically reduce the number of demonstrations needed for an individual task by sharing data across tasks. In particular, our goal is to learn a new task from a single demonstration of that task by using a dataset of demonstrations of many other tasks for meta-learning. Sharing information across tasks is by no means a new idea, e.g. by using task-to-task mappings [22], gating [23], and shared features [24]. These multi-task robotic learning methods consider the problem of generalization to new tasks from some specification of that task. A common approach, often referred to as contextual policies, is to provide the task as an input to the policy or value function, where the task is represented as a goal or demonstration [1, 2, 4, 3, 5]. Another approach is to train controllers for a variety of tasks and learn a mapping from task representations to controller parameters [10, 25, 26]. In this work, we instead use meta-learning to enable the robot to quickly learn new tasks with gradient-based policy updates. In essence, we learn policy parameters that, when finetuned on just one demonstration of a new task, can immediately learn to perform that task. This enables the robot to learn new tasks end-to-end with extreme efficiency, using only one demonstration, without requiring any additional mechanisms such as contexts or learned update functions.

## 3 Meta-Imitation Learning Problem Formulation

In this section, we introduce the visual meta-imitation learning problem, where a vision-based policy must adapt to a new task from a single demonstration. We also summarize a prior meta-learning method that we will extend into a meta-imitation learning algorithm in Section 4.

### 3.1 Problem Statement

Our goal is to learn a policy that can quickly adapt to new tasks from a single demonstration of that task. To remove the need for a large amount of task-specific demonstration data, we propose to reuse demonstration data from a number of other tasks to enable efficient learning of new tasks. By training for adaptation across tasks, meta-learning effectively treats entire tasks as datapoints. The amount of data available for each individual task is relatively small. In the context of robotics, this is precisely what we want for developing generalist robots – the ability to provide a small amount of supervision for each new task that the robot should perform. In this section, we will formally define the one-shot imitation learning problem statement and introduce notation.

We consider a policy  $\pi$  that maps observations  $\mathbf{o}$  to predicted actions  $\hat{\mathbf{a}}$ . During meta-learning, the policy is trained to adapt to a large number of tasks. Formally, each imitation task  $\mathcal{T}_i = \{\tau = \{\mathbf{o}_1, \mathbf{a}_1, \dots, \mathbf{o}_T, \mathbf{a}_T\} \sim \pi_i^*, \mathcal{L}(\mathbf{a}_{1:T}, \hat{\mathbf{a}}_{1:T}), T\}$  consists of demonstration data  $\tau$  generated by an

---

**Algorithm 1** Meta-Imitation Learning with MAML

---

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

**Require:**  $\alpha, \beta$ : step size hyperparameters

- 1: randomly initialize  $\theta$
  - 2: **while** not done **do**
  - 3:   Sample batch of tasks  $\mathcal{T}_i \sim p(\mathcal{T})$
  - 4:   **for all**  $\mathcal{T}_i$  **do**
  - 5:     Sample demonstration  $\tau = \{\mathbf{o}_1, \mathbf{a}_1, \dots, \mathbf{o}_T, \mathbf{a}_T\}$  from  $\mathcal{T}_i$
  - 6:     Evaluate  $\nabla_{\theta} \mathcal{L}_{\mathcal{T}_i}(f_{\theta})$  using  $\tau$  and  $\mathcal{L}_{\mathcal{T}_i}$  in Equation (2)
  - 7:     Compute adapted parameters with gradient descent:  $\theta'_i = \theta - \alpha \nabla_{\theta} \mathcal{L}_{\mathcal{T}_i}(f_{\theta})$
  - 8:     Sample demonstration  $\tau'_i = \{\mathbf{o}'_1, \mathbf{a}'_1, \dots, \mathbf{o}'_T, \mathbf{a}'_T\}$  from  $\mathcal{T}_i$  for the meta-update
  - 9:   **end for**
  - 10:   Update  $\theta \leftarrow \theta - \beta \nabla_{\theta} \sum_{\mathcal{T}_i \sim p(\mathcal{T})} \mathcal{L}_{\mathcal{T}_i}(f_{\theta'_i})$  using each  $\tau'_i$  and  $\mathcal{L}_{\mathcal{T}_i}$  in Equation 2
  - 11: **end while**
  - 12: **return** parameters  $\theta$  that can be quickly adapted to new tasks through imitation.
- 

expert policy  $\pi_i^*$  and a loss function  $\mathcal{L}$  used for imitation. Feedback is provided by the loss function  $\mathcal{L}(\mathbf{a}_1, \dots, \mathbf{a}_T, \hat{\mathbf{a}}_1, \dots, \hat{\mathbf{a}}_T) \rightarrow \mathbb{R}$ , which might be mean squared error for continuous actions, or a cross-entropy loss for discrete actions.

In our meta-learning scenario, we consider a distribution over tasks  $p(\mathcal{T})$ . In the one-shot learning setting, the policy is trained to learn a new task  $\mathcal{T}_i$  drawn from  $p(\mathcal{T})$  from only one demonstration generated by  $\mathcal{T}_i$ . During meta-training, a task  $\mathcal{T}_i$  is sampled from  $p(\mathcal{T})$ , the policy is trained using one demonstration from an expert  $\pi_i^*$  on  $\mathcal{T}_i$ , and then tested on a new demonstration from  $\pi_i^*$  to determine its training and test error according to the loss  $\mathcal{L}$ . The policy  $\pi$  is then improved by considering how the *test* error on the new demonstration changes with respect to the parameters. Thus, the test error on sampled demonstration from  $\pi_i^*$  serves as the training error of the meta-learning process. At the end of meta-training, new tasks are sampled from  $p(\mathcal{T})$ , and meta-performance is measured by the policy’s performance after learning from one demonstration. Tasks used for meta-testing are held out during meta-training.

### 3.2 Background: Model-Agnostic Meta-Learning

In our approach to visual meta-imitation learning, we will use meta-learning to train for fast adaptation across a number of tasks by extending model-agnostic meta-learning (MAML) [27] to meta-imitation learning from visual inputs. Previously, MAML has been applied to few-shot image recognition and reinforcement learning. The MAML algorithm aims to learn the weights  $\theta$  of a model  $f_{\theta}$  such that standard gradient descent can make rapid progress on new tasks  $\mathcal{T}$  drawn from  $p(\mathcal{T})$ , without overfitting to a small number of examples. Because the method uses gradient descent as the optimizer, it does not introduce any additional parameters, making it more parameter-efficient than other meta-learning methods. When adapting to a new task  $\mathcal{T}_i$ , the model’s parameters  $\theta$  become  $\theta'_i$ . In MAML, the updated parameter vector  $\theta'_i$  is computed using one or more gradient descent updates on task  $\mathcal{T}_i$ , i.e.  $\theta'_i = \theta - \alpha \nabla_{\theta} \mathcal{L}_{\mathcal{T}_i}(f_{\theta})$ . For simplicity of notation, we will consider one gradient update for the rest of this section, but using multiple gradient updates is a straightforward extension.

The model parameters are trained by optimizing for the performance of  $f_{\theta'_i}$  with respect to  $\theta$  across tasks sampled from  $p(\mathcal{T})$ , corresponding to the following problem:

$$\min_{\theta} \sum_{\mathcal{T}_i \sim p(\mathcal{T})} \mathcal{L}_{\mathcal{T}_i}(f_{\theta'_i}) = \sum_{\mathcal{T}_i \sim p(\mathcal{T})} \mathcal{L}_{\mathcal{T}_i}(f_{\theta - \alpha \nabla_{\theta} \mathcal{L}_{\mathcal{T}_i}(f_{\theta})}) \quad (1)$$

Note that the meta-optimization is performed over the parameters  $\theta$ , whereas the objective is computed using the updated parameters  $\theta'_i$ . In effect, MAML aims to optimize the model parameters such that one or a small number of gradient steps on a new task will produce maximally effective behavior on that task. The meta-optimization across tasks uses stochastic gradient descent (SGD).

## 4 Meta-Imitation Learning with MAML

In this section, we describe how we can extend the model-agnostic meta-learning algorithm (MAML) to the imitation learning setting. The model’s input,  $\mathbf{o}_t$ , is the agent’s observation at

time  $t$ , e.g. an image, whereas the output  $\mathbf{a}_t$  is the action taken at time  $t$ , e.g. torques applied to the robot’s joints. We will denote a demonstration trajectory as  $\tau := \{\mathbf{o}_1, \mathbf{a}_1, \dots, \mathbf{o}_T, \mathbf{a}_T\}$  and use a mean squared error loss as a function of policy parameters  $\phi$  as follows:

$$\mathcal{L}_{\mathcal{T}_i}(f_\phi) = \sum_{\tau^{(j)} \sim \mathcal{T}_i} \sum_t \|f_\phi(\mathbf{o}_t^{(j)}) - \mathbf{a}_t^{(j)}\|_2^2. \quad (2)$$

We will primarily consider the one-shot case, where only a single demonstration  $\tau^{(j)}$  is used for the gradient update. However, we can also use multiple demonstrations to resolve ambiguity.

For meta-learning, we assume a dataset of demonstrations with at least two demonstrations per task. This data is only used during meta-training; meta-test time assumes only one demonstration for each new task. During meta-training, each meta-optimization step entails the following: A batch of tasks is sampled and two demonstrations are sampled per task. Using one of the demonstrations,  $\theta'_i$  is computed for each task  $\mathcal{T}_i$  using gradient descent with Equation 2. Then, the second demonstration of each task is used to compute the gradient of the meta-objective by using Equation 1 with the loss in Equation 2. Finally,  $\theta$  is updated according to the gradient of the meta-objective. In effect, the pair of demonstrations serves as a training-validation pair. The algorithm is summarized in Algorithm 1.

The result of meta-training is a policy that can be adapted to new tasks using a single demonstration. Thus, at meta-test time, a new task  $\mathcal{T}$  is sampled, one demonstration for that task is provided, and the model is updated to acquire a policy for that task. During meta-test time, a new task might involve new goals or manipulating new, previously unseen objects.

#### 4.1 Two-Head Architecture: Meta-Learning a Loss for Fast Adaptation

In the standard MAML setup, outlined previously, the policy is consistent across the pre- and post-gradient update stages. However, we can make a modification such that the parameters of the final layers of the network are not shared, forming two “heads,” as shown in Figure 2. The parameters of the pre-update head are not used for the final, post-update policy, and the parameters of the post-update head are not updated using the demonstration. But, both sets of parameters are meta-learned for effective performance after adaptation. Interestingly, this two head architecture amounts to using a different inner objective in the meta-optimization, while keeping the same outer objective. To see this, let us denote  $\mathbf{y}_t^{(j)}$  as the set of post-synaptic activations of the last hidden layer, and  $W$  and  $b$  as the weight matrix and bias of the final layer. The inner loss function is then given by:

$$\mathcal{L}_{\mathcal{T}_i}^*(f_\phi) = \sum_{\tau^{(j)} \sim \mathcal{T}_i} \sum_t \|W\mathbf{y}_t^{(j)} + b - \mathbf{a}_t^{(j)}\|_2^2, \quad (3)$$

where  $W$  and  $b$ , the weights and bias of the last layer, effectively form the parameters of the meta-learned loss function. We use the meta-learned loss function  $\mathcal{L}_{\mathcal{T}_i}^*$  to compute the adapted parameter  $\theta'_i$  of each task  $\mathcal{T}_i$ , via gradient descent. Then, the meta-objective becomes:

$$\min_{\theta, W, b} \sum_{\mathcal{T}_i \sim p(\mathcal{T})} \mathcal{L}_{\mathcal{T}_i}(f_{\theta'_i}) = \sum_{\mathcal{T}_i \sim p(\mathcal{T})} \mathcal{L}_{\mathcal{T}_i}(f_{\theta - \alpha \nabla_{\theta} \mathcal{L}_{\mathcal{T}_i}^*(f_{\theta})}). \quad (4)$$

This provides the algorithm more flexibility in how it adapts the policy parameters to the expert demonstration, which we found to increase performance in a few experiments (see Appendix A.3). However, the more interesting implication of using a learned loss is that we can omit the actions during 1-shot adaptation, as we discuss next.

#### 4.2 Learning to Imitate without Expert Actions

Conventionally, a demonstration trajectory consists of pairs of observations and actions, as we discussed in Section 4. However, in many scenarios, it is more practical to simply provide a video of the task being performed, e.g. by a human or another robot. One step towards this goal, which we consider in this paper, is to remove the need for the robot arm trajectory and actions at test time.<sup>1</sup> Though, to be clear, we will assume access to expert actions during meta-training. Without access to expert actions at test time, it is unclear what the loss function for 1-shot adaptation should be. Thus, we will meta-learn a loss function, as discussed in the previous section. We can simply modify the loss in Equation 3 by removing the expert actions:

$$\mathcal{L}_{\mathcal{T}_i}^*(f_\phi) = \sum_{\tau^{(j)} \sim \mathcal{T}_i} \sum_t \|W\mathbf{y}_t^{(j)} + b\|_2^2,$$

<sup>1</sup>We leave the problem of domain shift, i.e. between a video of a human and the robot’s view, to future work.




Figure 2: Diagrams of the policy architecture with a bias transformation (top and bottom) and two heads (bottom). The green arrows and boxes indicate weights that are part of the meta-learned policy parameters  $\theta$ .

This corresponds to a learned quadratic loss function on the final layer of activations, with parameters  $W$  and  $b$ . Though, in practice, the loss function could be more complex. With this loss function, we can learn to learn from the raw observations of a demonstration using the meta-optimization objective in Equation 4, as shown in our experiments in Sections 6.2 and 6.3.

## 5 Model Architectures for Meta-Imitation Learning

We use a convolutional neural network (CNN) to represent the policy, similar to prior vision-based imitation and meta-learning methods [13, 27]. The policy observation includes both the camera image and the robot’s configuration, e.g. the joint angles and end-effector pose. In this section, we overview the policy architecture, but leave the details to be discussed in Section 6. The policy consists of several strided convolutions, followed by ReLU nonlinearities. The final convolutional layer is transformed into spatial feature points using a spatial soft-argmax [28, 29] and concatenated with the robot’s configuration. The result is passed through a set of fully-connected layers with ReLU nonlinearities. Because the data within a demonstration trajectory is highly correlated across time, batch normalization was not effective. Instead, we used layer normalization after each layer [30].

Although meta-imitation learning can work well with standard policy architectures such as the one described above, the optimal architecture for meta-learning does not necessarily correspond to the optimal architecture for standard supervised imitation learning. One particular modification that we found improves meta-learning performance is to concatenate a vector of parameters to a hidden layer of post-synaptic activations, which leads to what we will refer to as a *bias transformation*. This parameter vector is treated the same as other parameters in the policy during both meta-learning and test-time adaptation. Formally, let us denote the parameter vector as  $\mathbf{z}$ , the post-synaptic activations as  $\mathbf{x}$ , and the pre-synaptic activations at the next layer as  $\mathbf{y}$ . A standard neural network architecture sets  $\mathbf{y} = W\mathbf{x} + b$ , for bias  $b$  and weight matrix  $W$ . The error gradient with respect to the standard bias  $\frac{d\mathcal{L}}{db}$  equals the error gradient with respect to  $\mathbf{y}$ ,  $\frac{d\mathcal{L}}{d\mathbf{y}}$ . Thus, a gradient update of the standard bias is directly coupled with the update to the weight matrix  $W$  and parameters in earlier layers of the network. The bias transformation, which we describe next, provides more control over the updated bias by eliminating this decoupling. With a bias transformation, we set  $\mathbf{y} = W_1\mathbf{x} + W_2\mathbf{z} + b$ , where  $W = [W_1, W_2]$  and  $b$  are the weight matrix and bias. First, note that including  $\mathbf{z}$  and  $W_2$  simply corresponds to a reparameterization of the bias,  $\tilde{b} = W_2\mathbf{z} + b$ , since neither  $W_2\mathbf{z}$  nor  $b$  depend on the input. The error gradient with respect to  $\mathbf{z}$  and  $W_2$  are:  $\frac{d\mathcal{L}}{dW_2} = \frac{d\mathcal{L}}{d\mathbf{y}}\mathbf{z}^T$  and  $\frac{d\mathcal{L}}{d\mathbf{z}} = W_2^T \frac{d\mathcal{L}}{d\mathbf{y}}$ . After one gradient step, the updated transformed bias is:  $\tilde{b}' = (W_2 - \alpha \frac{d\mathcal{L}}{d\mathbf{y}}\mathbf{z}^T)(\mathbf{z} - \alpha W_2^T \frac{d\mathcal{L}}{d\mathbf{y}}) + b - \alpha \frac{d\mathcal{L}}{d\mathbf{y}}$ . Thus, a gradient update to the transformed bias can be controlled more directly by the values of  $W_2$  and  $\mathbf{z}$ , whose values do not directly affect the gradients of other parameters in the network. To see one way in which the network might choose to control the bias, consider the setting where  $\mathbf{z}$  and  $b$  are zero. Then, the updated bias is:  $\tilde{b}' = -\alpha W_2 W_2^T \frac{d\mathcal{L}}{d\mathbf{y}} - \alpha \frac{d\mathcal{L}}{d\mathbf{y}}$ . In summary, the bias transformation increases the representational power of the *gradient*, without affecting the representation power of the network itself. In our experiments, we found this simple addition to the network made gradient-based meta-learning significantly more stable and effective. We include a diagram of the policy architecture with the bias transformation in Figure 2.




Figure 3: Example tasks from the policy’s perspective. In the top row, each pair of images shows the start and final scenes of the demonstration. The bottom row shows the corresponding scenes of the learned policy roll-out. Left: Given one demonstration of reaching a target of a particular color, the policy must learn to reach for the same color in a new setting. Center: The robot pushes the target object to the goal after seeing a demonstration of pushing the same object toward the goal in a different scene. Right: We provide a demonstration of placing an object on the target, then the robot must place the object on the same target in a new setting.

## 6 Experiments

The goals of our experimental evaluation are to answer the following questions: (1) can our method learn to learn a policy that maps from image pixels to actions using a single demonstration of a task? (2) how does our meta-imitation learning method compare to prior one-shot imitation learning methods with varying dataset sizes? (3) can we learn to learn without expert actions? (4) how well does our approach scale to real-world robotic tasks with real images?

We evaluate our method on one-shot imitation in three experimental domains. In each setting, we compare our proposed method to a subset of the following methods:

- **random policy:** A policy that outputs random actions from a standard Normal distribution.
- **contextual policy:** A feedforward policy, which takes as input the final image of the demonstration, to indicate the goal of the task, and the current image, and outputs the current action.
- **LSTM:** A recurrent neural network which ingests the provided demonstration and the current observation, and outputs the current action, as proposed by Duan et al. [5].
- **LSTM+attention:** A recurrent neural network using the attention architecture proposed by Duan et al. [5]. This method is only applicable to non-vision tasks.

The contextual policy, the LSTM policies, and the proposed approach are all trained using the same datasets, with the same supervision. All policies, including the proposed approach, were meta-trained via a behavioral cloning objective (mean squared error) with supervision from the expert actions, using the Adam optimizer with default hyperparameters [31].

### 6.1 Simulated Reaching

The first experimental domain is a family of planar reaching tasks, as illustrated in Figure 3, where the goal of a particular task is to reach a target of a particular color, amid two distractors with different colors. This simulated domain allows us to rigorously evaluate our method and compare with prior approaches and baselines. We consider both vision and non-vision variants of this task, so that we can directly compare to prior methods that are not applicable to vision-based policies. See Appendix A.1 for more details about the experimental setup and choices of hyperparameters.

We evaluate each method on a range of meta-training dataset sizes and show the one-shot imitation success rate in Figure 4. Using vision, we find that meta-imitation learning is able to handle raw pixel inputs, while the LSTM and contextual policies struggle to learn new tasks using modestly-sized meta-learning datasets. In the non-vision case, which involves far fewer parameters, the LSTM policies fare much better, particularly when using attention, but still perform worse than MIL. Prior work demonstrated these approaches using 10,000 or more demonstrations [5]. Therefore, the mediocre performance of these methods on much smaller datasets is not surprising. We also provide a comparison with and without the bias transformation discussed in Section 5. The results demonstrate that MIL with the bias transformation (bt) can perform more consistently across dataset sizes.




Figure 4: One-shot success rate on test tasks as a function of the meta-learning dataset size in the simulated domains. Our meta-imitation learning approach (MIL) can perform well across a range of dataset sizes, and can more effectively learn new tasks than prior approaches that feed in the goal image (contextual) or demonstration (LSTM) as input. A random policy achieves 25.7% reaching success and 0.45% pushing success. For videos of the policies, see the supplementary video<sup>3</sup>.

## 6.2 Simulated Pushing

The goal of our second set of experiments is to evaluate our approach on a challenging domain, involving 7-DoF torque control, a 3D environment, and substantially more physical and visual diversity across tasks. The experiment consists of a family of simulated table-top pushing tasks, illustrated in Figure 3, where the goal is to push a particular object with random starting position to the red target amid one distractor. We designed the pushing environment starting from the OpenAI Gym PusherEnv, using the MuJoCo physics engine [32, 33]. We modified the environment to include two objects, vision policy input, and, across tasks, a wide range of object shapes, sizes, textures, frictions, and masses. We selected 116 mesh shapes from [thingiverse.com](http://thingiverse.com), 105 meshes for meta-training and 11 for evaluation. The meshes include models of chess pieces, models of animals like teddy bears and pufferfish, and other miscellaneous shapes. We randomly sampled textures from a set of over 5,000 images and used held-out textures for meta-testing. A selection of the objects and textures are shown in Figure 5. For more experimental details, hyperparameters, and ablations, see Appendix A.2.

The performance of one-shot pushing with held-out objects, shown in Figure 4, indicates that MIL effectively learned to learn to push new objects, with 85.8% one-shot learning success using the largest dataset size. Furthermore, MIL achieves, on average, 6.5% higher success than the LSTM-based approach across dataset sizes. The contextual policy struggles, likely because the full demonstration trajectory information is informative for inferring the friction and mass of the target object.

In Table 1, we provide two additional evaluations, using the largest dataset size. The first evaluates how each approach handles input demonstrations with less information, e.g. without actions and/or the robot arm state. For this, we trained each method to be able to handle such demonstrations, as discussed in Section 4.2. We see that the LSTM approach has difficulty learning without the expert actions. MIL also sees a drop in performance, but one that is less dramatic. The second evaluation shows that all approaches can improve performance if five demonstrations are available, rather than one, despite all policies being trained for 1-shot learning. In this case, we averaged the predicted action over the 5 input demonstrations for the contextual and LSTM approaches, and averaged the gradient for MIL.

## 6.3 Real-World Placing

The goal of our final experiment is to evaluate how well a real robot can learn to learn to interact with new, unknown objects from a single visual demonstration. Handling unseen objects is a challenge for both learning-based and non-learning-based manipulation methods, but is a necessity for robots to be capable of performing diverse tasks in unstructured real-world environments. In practice, most robot learning approaches have focused on much more narrow notions of generalization, such as a

method		video+state +action	video +state	video
LSTM	1-shot	78.38%	37.61%	34.23%
contextual		n/a	58.11%	56.98%
MIL (ours)		<b>85.81%</b>	<b>72.52%</b>	<b>66.44%</b>
LSTM	5-shot	83.11%	39.64%	31.98%
contextual		n/a	64.64%	59.01%
MIL (ours)		<b>88.75%</b>	<b>78.15%</b>	<b>70.50%</b>

Table 1: One-shot and 5-shot simulating pushing success rate with varying demonstration information provided at test-time. MIL can more successfully learn from a demonstration without actions and without robot state and actions than LSTM and contextual policies.

<sup>3</sup>For video results and code, see <https://sites.google.com/view/one-shot-imitation>



Figure 5: Training and test objects used in our simulated pushing (left) and real-world placing (right) experiments. Note that we only show a subset of the  $\sim 100$  training objects used for the pushing and placing experiments, and a subset of the textures and object scales used for training and testing robot pushing.

varied target object location [28] or block stacking order [5]. With this goal in mind, we designed a robotic placing experiment using a 7-DoF PR2 robot arm and RGB camera, where the goal is to place a held item into a target container, such as a cup, plate, or bowl, while ignoring two distractors. We collected roughly 1300 demonstrations for meta-training using a diverse range of objects, and evaluated one-shot learning using held-out, unseen objects (see Figure 5). The policy is provided a single visual demonstration of placing the held item onto the target, but with varied positions of the target and distractors, as illustrated in Figure 3. Demonstrations were collected using human teleoperation through a motion controller and virtual reality headset [34], and each demo included the camera video, the sequence of end-effector poses, and the sequence of actions – the end-effector linear and angular velocities. See Appendix A.3 for more explanation of data collection, evaluation, and hyperparameters.

The results, in Table 2, show that the MIL policy can learn to localize the previously-unseen target object and successfully place the held item onto the target with 90% success, using only a single visual demonstration with those objects. We found that the LSTM and contextual policies were unable to localize the correct target object, likely due to the modestly-sized meta-training dataset, and instead placed onto an arbitrary object, achieving 25% success. Using the two-head approach described in 4.2, we also experimented with only providing the video of the demonstration, omitting the robot end-effector trajectory and controls. MIL can also learn to handle this setting, although with less success, suggesting the need for more data and/or further research. We include videos of all placing policies in the supplementary video<sup>4</sup>.

method	test performance
LSTM	25%
contextual	25%
MIL	<b>90%</b>
MIL, video only	<b>68.33%</b>

Table 2: One-shot success rate of placing a held item into the correct container, with a real PR2 robot, using 29 held-out test objects. Meta-training used a dataset with  $\sim 100$  objects. MIL, using video only receives the only video part of the demonstration and not the arm trajectory or actions.

## 7 Discussion and Future Work

We proposed a method for one-shot visual imitation learning that can learn to perform tasks using visual inputs from just a single demonstration. Our approach extends gradient-based meta-learning to the imitation learning setting, and our experimental evaluation demonstrates that it substantially outperforms a prior one-shot imitation learning method based on recurrent neural networks. The use of gradient-based meta-learning makes our approach more efficient in terms of the number of demonstrations needed during meta-training, and this efficiency makes it possible for us to also evaluate the method using raw pixel inputs and on a real robotic system.

The use of meta-imitation learning can substantially improve the efficiency of robotic learning methods without sacrificing the flexibility and generality of end-to-end training, which is especially valuable for learning skills with complex sensory inputs such as images. While our experimental evaluation uses tasks with limited diversity (other than object diversity), we expect the capabilities of our method to increase substantially as it is provided with increasingly more diverse demonstrations for meta-training. Since meta-learning algorithms can incorporate demonstration data from all available tasks, they provide a natural avenue for utilizing large datasets in a robotic learning context, making it possible for robots to not only learn more skills as they acquire more demonstrations, but to actually become *faster* and *more effective* at learning new skills through the process.

<sup>4</sup>For video results and code, see <https://sites.google.com/view/one-shot-imitation>



## Acknowledgments

This work was supported by the National Science Foundation through IIS-1651843, IIS-1614653, IIS-1637443, and a graduate research fellowship, by Berkeley DeepDrive, by the ONR PECASE award N000141612723, as well as an ONR Young Investigator Program award. The authors would like to thank Yan Duan for providing a reference implementation of [5] and the anonymous reviewers for providing feedback.

## References

- [1] M. P. Deisenroth, P. Englert, J. Peters, and D. Fox. Multi-task policy search for robotics. In *International Conference on Robotics and Automation (ICRA)*, 2014.
- [2] A. G. Kupcsik, M. P. Deisenroth, J. Peters, G. Neumann, et al. Data-efficient generalization of robot skills with contextual policy search. In *AAAI Conference on Artificial Intelligence*, 2013.
- [3] T. Schaul, D. Horgan, K. Gregor, and D. Silver. Universal value function approximators. In *International Conference on Machine Learning (ICML)*, 2015.
- [4] F. Stulp, G. Raiola, A. Hoarau, S. Ivaldi, and O. Sigaud. Learning compact parameterized skills with a single regression. In *International Conference on Humanoid Robots*, 2013.
- [5] Y. Duan, M. Andrychowicz, B. Stadie, J. Ho, J. Schneider, I. Sutskever, P. Abbeel, and W. Zaremba. One-shot imitation learning. *arXiv preprint arXiv:1703.07326*, 2017.
- [6] S. Schaal, A. Ijspeert, and A. Billard. Computational approaches to motor learning by imitation. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 2003.
- [7] S. Thrun and L. Pratt. *Learning to learn*. Springer Science & Business Media, 1998.
- [8] A. Billard, Y. Epars, S. Calinon, S. Schaal, and G. Cheng. Discovering optimal imitation strategies. *Robotics and autonomous systems*, 2004.
- [9] S. Schaal, J. Peters, J. Nakanishi, and A. Ijspeert. Learning movement primitives. *Robotics Research*, 2005.
- [10] P. Pastor, H. Hoffmann, T. Asfour, and S. Schaal. Learning and generalization of motor skills by learning from demonstration. In *International Conference on Robotics and Automation (ICRA)*, 2009.
- [11] N. Ratliff, J. A. Bagnell, and S. S. Srinivasa. Imitation learning for locomotion and manipulation. In *International Conference on Humanoid Robots*, 2007.
- [12] D. Pomerleau. ALVINN: an autonomous land vehicle in a neural network. In *Advances in Neural Information Processing Systems (NIPS)*, 1989.
- [13] M. Bojarski, D. Del Testa, D. Dworakowski, B. Firner, B. Flepp, P. Goyal, L. Jackel, M. Monfort, U. Muller, J. Zhang, X. Zhang, J. Zhao, and K. Zieba. End to end learning for self-driving cars. *arXiv preprint arXiv:1604.07316*, 2016.
- [14] A. Giusti, J. Guzzi, D. C. Cireşan, F.-L. He, J. P. Rodríguez, F. Fontana, M. Faessler, C. Forster, J. Schmidhuber, G. Di Caro, et al. A machine learning approach to visual perception of forest trails for mobile robots. *IEEE Robotics and Automation Letters*, 2016.
- [15] J. Zhang and K. Cho. Query-efficient imitation learning for end-to-end simulated driving. In *AAAI Conference on Artificial Intelligence*, 2017.
- [16] S. Ross, G. J. Gordon, and D. Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *AISTATS*, 2011.
- [17] M. Laskey, S. Staszak, W. Y.-S. Hsieh, J. Mahler, F. T. Pokorny, A. D. Dragan, and K. Goldberg. Shiv: Reducing supervisor burden in dagger using support vectors for efficient learning from demonstrations in high dimensional state spaces. In *International Conference on Robotics and Automation (ICRA)*, 2016.

- [18] A. Y. Ng, S. J. Russell, et al. Algorithms for inverse reinforcement learning. In *International Conference on Machine Learning (ICML)*, 2000.
- [19] C. Finn, S. Levine, and P. Abbeel. Guided cost learning: Deep inverse optimal control via policy optimization. In *International Conference on Machine Learning (ICML)*, 2016.
- [20] A. Nair, P. Agarwal, D. Chen, P. Isola, P. Abbeel, and S. Levine. Combining self-supervised learning and imitation for vision-based rope manipulation. *International Conference on Robotics and Automation (ICRA)*, 2017.
- [21] P. Sermanet, K. Xu, and S. Levine. Unsupervised perceptual rewards for imitation learning. *Robotics: Science and Systems (RSS)*, 2017.
- [22] S. Barrett, M. E. Taylor, and P. Stone. Transfer learning for reinforcement learning on a physical robot. In *Ninth International Conference on Autonomous Agents and Multiagent Systems-Adaptive Learning Agents Workshop (AAMAS-ALA)*, 2010.
- [23] K. Mülling, J. Kober, O. Kroemer, and J. Peters. Learning to select and generalize striking movements in robot table tennis. *The International Journal of Robotics Research (IJRR)*, 2013.
- [24] A. Gupta, C. Devin, Y. Liu, P. Abbeel, and S. Levine. Learning invariant feature spaces to transfer skills with reinforcement learning. *International Conference on Learning Representations (ICLR)*, 2017.
- [25] J. Kober, A. Wilhelm, E. Oztop, and J. Peters. Reinforcement learning to adjust parametrized motor primitives to new situations. *Autonomous Robots*, 2012.
- [26] B. C. Da Silva, G. Konidaris, and A. G. Barto. Learning parameterized skills. In *International Conference of Machine Learning (ICML)*, 2012.
- [27] C. Finn, P. Abbeel, and S. Levine. Model-agnostic meta-learning for fast adaptation of deep networks. *International Conference on Machine Learning (ICML)*, 2017.
- [28] S. Levine, C. Finn, T. Darrell, and P. Abbeel. End-to-end learning of deep visuomotor policies. *Journal of Machine Learning Research (JMLR)*, 2016.
- [29] C. Finn, X. Y. Tan, Y. Duan, T. Darrell, S. Levine, and P. Abbeel. Deep spatial autoencoders for visuomotor learning. In *International Conference on Robotics and Automation (ICRA)*, 2016.
- [30] J. L. Ba, J. R. Kiros, and G. E. Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.
- [31] D. Kingma and J. Ba. Adam: A method for stochastic optimization. *International Conference on Learning Representations (ICLR)*, 2015.
- [32] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba. Openai gym. *arXiv preprint arXiv:1606.01540*, 2016.
- [33] E. Todorov, T. Erez, and Y. Tassa. Mujoco: A physics engine for model-based control. In *International Conference on Intelligent Robots and Systems (IROS)*, 2012.
- [34] T. Zhang, Z. McCarthy, O. Jow, D. Lee, K. Goldberg, and P. Abbeel. Deep imitation learning for complex manipulation tasks from virtual reality teleoperation. *arXiv preprint arXiv:1710.04615*, 2017.

## A Additional Experimental Details

In this section, we provide additional experimental details for all experiments, including information regarding data collection, evaluation, and training hyperparameters.

### A.1 Simulated Reaching

**Experimental Setup:** In both vision and no-vision cases of this experiment, the input to the policy includes the arm joint angles and the end-effector position. In the vision variant, the  $80 \times 64$  RGB image is also provided as input. In the non-vision version, the 2D positions of the objects are fed into the policy, but the index of the target object within the state vector is not known and must be inferred from the demonstration. The policy output corresponds to torques applied to the two joints of the arm. A policy roll-out is considered a success if it comes within 0.05 meters of the goal within the last 10 timesteps, where the size of the arena is  $0.6 \times 0.6$  meters.

To obtain the expert policies for this task, we use iLQG trajectory optimization to generate solutions for each task (using knowledge of the goal), and then collect several demonstrations per task from the resulting policy with injected Gaussian noise. At meta-test time, we evaluate the policy on 150 tasks and 10 different trials per task (1500 total trials) where each task corresponds to a held-out color. Note that the demonstration provided at meta-test time usually involves different target and distractor positions than its corresponding test trial. Thus, the one-shot learned policy must learn to localize the target using the demonstration and generalize to new positions, while meta-training must learn to handle different colors.

**Hyperparameters:** For all vision-based policies, we use a convolutional neural network policy with 3 convolution layers each with  $40 \ 3 \times 3$  filters, followed by 4 fully-connected layers with hidden dimension 200. For this domain only, we simply flattened the final convolutional map rather than transforming it into spatial feature points. The recurrent policies additionally use an LSTM with 512 units that takes as input the features from the final layer. For non-vision policies, we use the same architecture without the convolutional layers, replacing the output of the convolutional layers with the state input. All methods are trained using a meta batch-size of 5 tasks. The policy trained with meta-imitation learning uses 1 meta-gradient update with step size 0.001 and bias transformation with dimension 10. We also find it helpful to clip the meta-gradient to lie in the interval  $[-20, 20]$  before applying it. We use the normal single-head architecture for MIL as shown in Figure 2.

### A.2 Simulated Pushing

**Experimental Setup:** The policy input consists of a  $125 \times 125$  RGB image and the robot joint angles, joint velocities, and end-effector pose. A push is considered a success if the center of the target object lands on the red target circle for at least 10 timesteps within a 100-timestep episode. The reported pushing success rates are computed over 74 tasks with 6 trials per task (444 total trials).

We acquired a separate demonstration policy for each task using the trust-region policy optimization (TRPO) algorithm. The expert policy inputs included the target and distractor object poses rather than vision input. To encourage the expert policies to take similar strategies, we first trained a single policy on a single task, and then initialized the parameters of all of the other policies with those from the first policy. When initializing the policy parameters, we increased the variance of the converged policy to ensure appropriate exploration.

**Hyperparameters:** For all methods, we use a neural network policy with 4 strided convolution layers with  $16 \ 5 \times 5$  filters, followed by a spatial softmax and 3 fully-connected layers with hidden dimension 200. For optimization, each method use a meta-batch size of 15 tasks. MIL uses 1 inner gradient descent step with step size  $\alpha = 0.01$ , inner gradient clipping within the range  $[-10, 10]$ , and bias transformation with dimension 20. We also use an additional bias transformation that is concatenated to the image at each time step with the same size as the input image. The LSTM policy uses 512 hidden units.

Because this domain is significantly more challenging than the simulating reaching domain, we found it important to use the two-head architecture described in section 4.2. We include an ablation of the two-head architecture in Table 3, demonstrating the benefit of this choice.

method	1-head	2-head
MIL with 1-shot	80.63%	85.81%
MIL with 5-shot	82.63%	88.75%

Table 3: Ablation test on 1-head and 2-head architecture for simulated pushing as shown in Figure 2, using a dataset with 9200 demonstrations for meta-learning. Using two heads leads to significantly better performance in this domain.

### A.3 Real-World Placing

**Experimental Setup:** The videos in the demo are composed of a sequence of  $320 \times 240$  RGB images from the robot camera. We pre-process the demonstrations by downsampling the images by a factor of 2 and cropping them to be of size  $100 \times 90$ . Since the videos we collected are of variable length, we subsample the videos such that they all have fixed time horizon 30.

To collect demonstration data for one task, we randomly select one holding object and three placing containers from our training set of objects (see the third image of Figure 5), and place those three objects in front of the robot in 8 random positions. In this way, we collect 1293 demonstrations, where we use 96 of them as validation set and the rest as the training set.

During policy evaluation, we evaluate the policy with 18 tasks and 4 trials per task (72 total trials) where we use 1 placing target and 1 holding object from the test set for each task. In addition, we manually code an "open gripper" action at the end of the trajectory, which causes the robot to drop the holding object. We define success as whether or not the held object landed in or on the target container after the gripper is opened.

**Hyperparameters:** We use a neural network policy with 3 strided convolution layers and 2 non-strided convolutions layers with  $64 \ 3 \times 3$  filters, followed by a spatial softmax and 3 fully-connected layers with hidden dimension 100. We initialize the first convolution layer from VGG-19 and keep it fixed during meta-training. Following prior work [34], we change the objective to be a mixture of  $\ell_1$  and  $\ell_2$  loss, where  $\ell_2$  loss is scaled down by 100, and add an auxiliary loss that regresses from the learned features at the first time step to the 2D position of the target container. We determine the position of the target container from the end-effector position at the final timestep of the demonstration; this does not require additional supervision beyond the demonstration. Additionally, we also feed the predicted 2D position of the target into the fully-connected layers of the network. MIL uses a meta-batch size of 12 tasks, 5 inner gradient descent steps with step size 0.005, inner gradient clipping within the range  $[-30, 30]$ , and bias transformation with dimension 20. We also use the single-head architecture for MIL just as what we do for simulated reaching. The LSTM policy uses 512 hidden units.