

Scalable Multi-Task Imitation Learning with Autonomous Improvement

Avi Singh^{*1}, Eric Jang², Alexander Irpan², Daniel Kappler³, Murtaza Dalal¹,
 Sergey Levine^{1,2}, Mohi Khansari³, Chelsea Finn^{2,4}

Abstract—While robot learning has demonstrated promising results for enabling robots to automatically acquire new skills, a critical challenge in deploying learning-based systems is scale: acquiring enough data for the robot to effectively generalize broadly. Imitation learning, in particular, has remained a stable and powerful approach for robot learning, but critically relies on expert operators for data collection. In this work, we target this challenge, aiming to build an imitation learning system that can continuously improve through autonomous data collection, while simultaneously avoiding the explicit use of reinforcement learning, to maintain the stability, simplicity, and scalability of supervised imitation. To accomplish this, we cast the problem of imitation with autonomous improvement into a multi-task setting. We utilize the insight that, in a multi-task setting, a failed attempt at one task might represent a successful attempt at another task. This allows us to leverage the robot’s own trials as demonstrations for tasks other than the one that the robot actually attempted. Using an initial dataset of multi-task demonstration data, the robot autonomously collects trials which are only sparsely labeled with a binary indication of whether the trial accomplished any useful task or not. We then embed the trials into a learned latent space of tasks, trained using only the initial demonstration dataset, to draw similarities between various trials, enabling the robot to achieve one-shot generalization to new tasks. In contrast to prior imitation learning approaches, our method can autonomously collect data with sparse supervision for continuous improvement, and in contrast to reinforcement learning algorithms, our method can effectively improve from sparse, task-agnostic reward signals.

I. INTRODUCTION & RELATED WORK

Robotic learning holds the potential to enable *generalist* robots: robotic systems that can autonomously perform a wide range of different behaviors. In order to enable robots to be generalists, equipped with large repertoires of skills, we need them to be able to acquire each skill from a relatively modest amount of experience, and to make it feasible to add new skills using only a little bit of data. Such generalist systems not only have tremendous promise for applications across a range of domains that cannot currently be automated, but would also represent a significant step forward in artificial intelligence research. Two of the most successful approaches to general-purpose robotic learning are imitation learning, where a robot uses human-provided demonstration data to learn skills [1], [2], and reinforcement learning (RL), where it uses a trial-and-error learning process [3], [4].

When combined with high-capacity function approximators, such as deep neural networks, both techniques have been shown to enable complex robotic skills directly from low-level sensory observations, such as images [2], [5].

^{*}Work done while an AI Resident at X, Mountain View.

¹UC Berkeley, ²Robotics at Google, ³X, ⁴Stanford

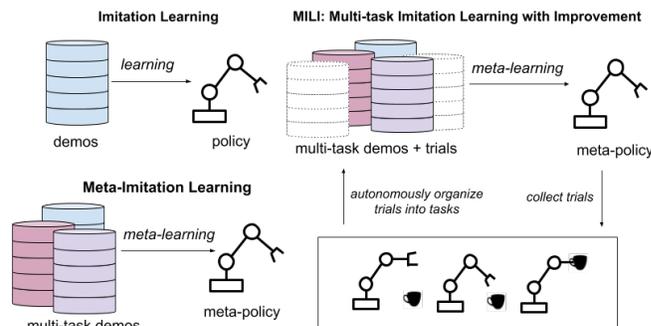


Fig. 1. In this figure, we describe our problem setting and differentiate it from the standard imitation learning and meta-imitation learning problem settings. In standard imitation learning, we are provided a dataset of demonstrations for a specific task, and we learn a policy to mimic the behavior in the demonstrations. In meta-imitation learning, we are provided with several such datasets drawn from a task distribution, and we learn a single meta-policy that is capable of performing can generalize to new tasks in this distribution from a single demonstration. We also consider a one-shot imitation learning setting in this paper; however, unlike standard meta-imitation, our method can also utilize autonomously collected experience. After filtering out trials that do not succeed on any task, our method adds the autonomously collected trials to the meta-training set, matching them with existing tasks using an automated pairing method. We use this newly collected and organized data to update the policy, establishing a learning loop that can perpetually improve from autonomously collected data.

However, both techniques also have substantial limitations, some of which we will aim to mitigate in this work. First, the combination of deep networks with imitation learning and reinforcement learning results in very high data complexity. In imitation learning, this means that human operators must provide a large number of demonstrations for each task, while in the case of reinforcement learning, this translates into very large autonomous data collection requirements particularly to overcome the challenges of exploration [6]. With the data requirements in both cases, it is thus difficult to scale these techniques to enable truly generalist robots equipped with large repertoires of skills. The two approaches also have some distinct strengths and weaknesses. Reinforcement learning enables a robot to improve, perhaps perpetually, from its own experience [5], [7], [8], but RL-based policies are generally more limited in the complexity of the tasks they can accomplish, due to the difficulty of discovering task solutions autonomously and the complexity of the RL optimization problem [9], [10], [11]. By comparison, imitation learning methods present a much simpler *supervised* learning problem, making them suitable for more complex skills [12], [2], but they require a more manual process of collecting human demonstration data [2], and the robot cannot continue to improve autonomously from its own experience [13]. This last point in particular is a major

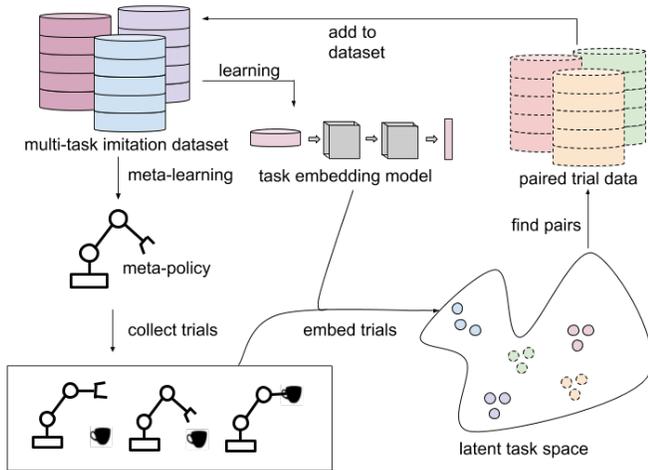


Fig. 2. The MILI algorithm. We bootstrap a one-shot imitation policy using a multi-task imitation learning dataset. We then use this policy to collect trials in new environments. A latent task space, learned using the same initial dataset, is then used to find similarities in the collected trials, and generate new tasks for meta-imitation learning. We update our meta-policy using the newly collected data, and repeat this process until convergence.

shortcoming of current imitation learning approaches: the capacity for continuous and autonomous self-improvement is an exceptionally powerful feature of RL.

Multi-task [14] and meta-imitation learning [15], [16], [17], [18], [19], [20], [21] alleviate the data collection problem to some extent by reusing data across many different tasks. These methods typically learn a parameterized policy that is a function of both the current observation and the task, thus obtaining a single policy that is capable of performing several different tasks and generalizing to new tasks. For multi-task imitation learning, the task is typically specified in the form of a task index, while in meta-imitation, the task is specified using a single demonstration for that task. Thus, meta-imitation learning systems learn a one-shot imitation policy, capable of performing a new task from just a single demonstration for it. However, these approaches still require manually collecting data for hundreds of different tasks, and cannot improve from autonomously collected data. In this paper, we will also consider a one-shot imitation learning problem. However, unlike these prior works, our goal is to develop an algorithm that can improve its one-shot imitation performance using additional *autonomously-collected* data.

To this end, we utilize the following insight: if a robot is trained in a multi-task setting consisting of multiple distinct tasks with hundreds of different objects, it is likely to perform some useful behavior when it attempts a task in a new setting, even if it is not the behavior that was actually intended. This data can then be leveraged to learn the skill that was actually performed. This insight can be viewed as a generalization of prior works on goal-conditioned learning via hindsight goal relabeling [22], [23], [24], [25], but in the space of *tasks* rather than the space of goals to be reached. It is this insight, and how we can leverage it in a meta-imitation learning framework, that provides the novel contribution of this work: a framework for autonomous data collection and improvement *without* explicit reinforcement

learning, that instead leverages a supervised meta-imitation learning method to enable a robot to be bootstrapped off of demonstration data, and then improve its repertoire of skills via its own autonomously collected experience.

Our approach, which we call Multi-task Imitation Learning with Improvement (MILI), is summarized in Figure 2. We bootstrap a meta imitation policy by performing one-shot imitation learning on a human-collected dataset of paired demonstrations, which consists of pairs of optimal trajectories corresponding to the same task (following the standard meta-imitation learning set-up in prior work [26], [16]). We then use this policy to collect trials autonomously. For each such trial, a user needs to provide an answer to the following yes/no question: did the robot succeed at *any* task of interest? If the answer is yes, this trajectory is added to an augmented meta-training set. This binary label can be provided via an automated reward function or direct human feedback. We use the initial demonstration dataset of paired demonstrations to also learn a latent space of tasks, where demonstrations corresponding to the same tasks are pushed close to each other, and demonstrations corresponding to different tasks are pushed apart. We then embed all of the autonomously collected trajectories into this learned task space, and generate a dataset of paired trajectories by pairing trajectories that are close to each other in the latent space.

The core contribution of this paper is a framework that allows imitation learning systems to continuously improve through autonomous data collection and learning. We empirically study a challenging, vision-based manipulation setting consisting of hundreds of tasks from four distinct task families: button pressing, sliding, grasping, and pick-and-place. Our framework allows for substantial improvements over standard imitation learning and meta-imitation learning approaches and, critically, has the potential to continuously improve itself without human-provided demonstrations.

II. PRELIMINARIES

In this section, we summarize the one-shot imitation learning problem setting considered in prior work [15], [16], [21], [19], [18]. We first formalize the single-task imitation learning setting, and then extend it to the one-shot setting.

Imitation Learning via Behavior Cloning. In the standard single task imitation learning setting [1], [27], we are provided a dataset of expert demonstrations $\mathcal{D} = \{d_1, d_2, \dots, d_N\}$. Each demonstration d_n consists of a trajectory of observations and actions denoting optimal behavior for that task $d_n = [(o_1^n, a_1^n), \dots, (o_T^n, a_T^n)]$, and we need to learn a policy π that can mimic this expert behavior. While there are several ways to perform imitation learning from expert demonstrations such as inverse reinforcement learning [28], [29], [30], [31], [32] and occupancy matching [33], we consider behavior cloning [1] in this paper due its stability and ease of use. We train a policy π_θ , parameterized using a neural network with parameters θ that takes the observation o_t as input and outputs a distribution over actions. The parameters θ of the policy are trained with stochastic gradient descent to minimize the following loss function:

$$\mathcal{L}_{bc}(\theta) = \sum_{n=1}^N \sum_{t=1}^T -\log \pi_{\theta}(a_t^n | o_t^n) \quad (1)$$

where $\pi_{\theta}(a_t | o_t)$ represents the distribution over actions for observation o_t .

One-Shot Imitation Learning. In the one-shot imitation learning setting [26], [16], the goal is to learn a meta-policy that can adapt to new, unseen tasks from just a single demonstration for that task. In order to achieve this for tasks with high-dimensional observations such as pixels, which typically require large datasets to learn an effective policy [2], [34], we need to transfer knowledge from demonstrations of previously seen tasks to the task at hand. Thus, instead of assuming access to expert demonstrations for a single task, one-shot imitation learning assumes an unknown distribution of tasks $p(\mathbb{T})$, and is provided with a set of tasks $\{\mathbb{T}_i\}$ from this distribution, which are called meta-training tasks. More concretely, for each training task \mathbb{T}_i , we have access to a set of demonstrations $\mathcal{D}_i = \{d_1^i, d_2^i, \dots, d_N^i\}$. Different tasks contain different objects, and different actions can be performed on those objects. For example, as shown in Figure 4, a mug could be picked up, a plate could be pushed across a table, a button could be pressed, and a glass could be placed on a tray. The combination of an action and an object constitutes a unique task.

One-shot imitation learning techniques learn a meta-policy π_{θ} , which takes as input both the current observation o_t and a demonstration d corresponding to the task which is to be performed, and outputs a distribution over actions. The demonstration d specifies to the meta-policy what task is to be performed, and conditioning on different demos can lead to different tasks being performed for the same observation. At training time, we first sample a task \mathbb{T}_i , and then sample two demonstrations d_m and d_n corresponding to this task for $m \neq n$. We condition the meta-policy on one of these two demonstrations, say demonstration d_n , and optimize the following loss on the expert observation-action pairs from the other demonstration, d_m :

$$\mathcal{L}_{bc}(\theta, d_m, d_n) = \sum_{t=1}^T -\log \pi_{\theta}(a_t^m | o_t^m, d_n)$$

We obtain the complete one-shot imitation learning loss by summing across all tasks and all possible demonstration pairs that can be drawn from the same task:

$$\mathcal{L}_{oil}(\theta, \{\mathcal{D}_i\}) = \sum_{i=1}^M \sum_{d_m, d_n \sim \mathcal{D}_i} \mathcal{L}_{bc}(\theta, d_m, d_n) \quad (2)$$

where M is the total number of training tasks.

III. PROBLEM STATEMENT

In this section, we describe our problem setting, and highlight its differences from the standard imitation and reinforcement learning problem settings. Similar to the one-shot imitation learning problem setting described in Section II, our goal is to learn a meta-policy that can adapt to new,

unseen tasks from just a single demonstration for that task, and we assume access to a set of expert trajectories \mathcal{D}_i for a subset of the meta-training tasks $\{\mathbb{T}_i\}$. However, unlike the one-shot imitation learning setting, which only assumes access to a static set of demonstrations, we also assume that we can attempt new trials τ on the meta-training tasks $\{\mathbb{T}_i\}$ using a meta-policy trained only on $\{\mathcal{D}_i\}$.

This is similar to the multi-task reinforcement learning and meta-reinforcement learning settings [35], [36], but we do not assume access to any task-specific reward functions for any of the tasks. We may also wish to constrain the space of learnable behaviors so that the robot avoids associating “knocking objects onto the floor” or “not touching any objects” as meaningful tasks. This requires access to a filtering function \mathcal{F} , which operates on the trial τ and returns TRUE if any useful behavior was performed during a trial. This could be automated via a learned classifier for “useful” behaviors, or implemented manually via a human annotator that annotates trajectories collected by the meta-policy. Note that the human annotator does not need to specify what task was achieved by the robot in a particular trial: it could be pushing a bowl, picking up a mug or pressing a button, but the human does not need to provide a task label, as long as it is “useful”. This makes it possible to scale to hundreds of tasks without increasing the difficulty of the annotation. If the filtering function is not provided, the robot can in principle learn to perform all possible behaviors, including undesirable behaviors such as throwing objects onto the floor. While this is not necessarily unreasonable, it could in practice crowd out the desirable behaviors and increase training time.

IV. MILI: MULTI-TASK IMITATION LEARNING WITH IMPROVEMENT

As shown in Figure 2, we extend a multi-task, meta-imitation learning pipeline with the ability to learn from the policy’s own experience, while preserving the stability and simplicity of supervised learning methods. In order to achieve this, we need to (a) learn a policy that can perform a diverse set of behaviors in a variety of different environments, (b) utilize this policy to autonomously collect data in new environments, and (c) learn an improved one-shot imitation policy that leverages both the initial dataset and the autonomously collected data. For obtaining a data collection policy, we perform meta-imitation learning on a human-provided dataset. We condition this meta-policy on random demos for collecting trials in new environments. We also learn a *latent space* of skills using the same initial dataset, which is then used for organizing the collected trials into new tasks. We now run meta-imitation learning on this expanded dataset, resulting in an improved policy that utilizes both human-provided and autonomously collected data.

Learning a data collection policy Naïvely performing behavior cloning (as described in Section II) on the human-provided dataset is unlikely to result in diverse behavior: since a robot can perform many useful tasks in any given scene, cloning the actions from all human-provided trajectories without providing any context will lead to an averaging

Algorithm 1 MILI: Multi-Task Imitation Learning with Improvement

- 1: **Input:** Training tasks $\{\mathbb{T}_i\}$, dataset $\{\mathcal{D}_i\}$
 - 2: **Input:** Data collection batch size B
 - 3: **Input:** Pairing threshold α
 - 4: $\{\textit{One-shot imitation pre-training}\}$
 - 5: $\theta \leftarrow \arg \min_{\theta} \mathcal{L}(\theta, \{\mathcal{D}_i\})$ (see Equation 3)
 - 6: $\{\textit{Autonomous data collection and improvement}\}$
 - 7: Initialize empty trial dataset \mathcal{R}
 - 8: **for** iter= 1, . . . **do**
 - 9: **for** trial= 1, . . . , B **do**
 - 10: Sample training task \mathbb{T}_i
 - 11: Sample demo d from \mathcal{D}_i
 - 12: Collect trial τ_n from task \mathbb{T}_i with meta-policy $\pi_{\theta}(a|o, d)$
 - 13: **if** $\mathcal{F}(\tau_n)$ **then**
 - 14: Add τ_n to \mathcal{R}
 - 15: **end if**
 - 16: **end for**
 - 17: **for** τ_n, τ_m in \mathcal{R} **do**
 - 18: Compute $p = f_{\theta}(\tau_n)^T f_{\theta}(\tau_m)$
 - 19: **if** $p > \alpha$ **then**
 - 20: Create new task dataset $\mathcal{D} := \{\tau_n, \tau_m\}$
 - 21: $\{\mathcal{D}_i\} = \{\mathcal{D}_i\} \cup \mathcal{D}$
 - 22: **end if**
 - 23: **end for**
 - 24: $\theta \leftarrow \arg \min_{\theta} \mathcal{L}_{\text{oil}}(\theta, \{\mathcal{D}_i\})$ (see Equation 2)
 - 25: **end for**
 - 26: **return** π_{θ}
-

of the different possible actions, and the behavior generated by the policy will not be useful. Therefore, we instead train a demo-conditioned meta-policy using the loss function described in Equation 2 (see Section II). We can now use this meta-policy to collect trials in new scenes by conditioning it on different demonstrations from $\{\mathcal{D}_i\}$.

Utilizing autonomously collected data A straightforward way to improve our meta-policy’s performance using the collected trials would be to optimize the one-shot imitation learning loss defined in Equation 2 on the collected trajectories. However, in order to optimize this loss, we need at least two successful trials for any given task, as described in Section II. That is, we need at least two trials depicting optimal behavior for the same task for it to be useful for learning. Since our filtering function only labels trials as being useful for any task, we cannot use it for assigning trajectories to specific tasks. If we can find two trials that perform a similar behavior, we can add them to our dataset as corresponding to a new task.

However, finding similar trials can be non-trivial: we wish to learn policies from high-dimensional observations spaces such as visual inputs, which implies that our trials consist of videos, and finding the distance between two videos by computing a standard distance metric like the Euclidean distance is unlikely to result in useful pairings. This motivates the need to learn a *latent space of tasks*, in which we can embed any given trial, and compute meaningful distances between the trials. If two trials are found to be close to each other in the latent space, we can add them to $\{\mathcal{D}_i\}$ as a pair of demonstrations corresponding to a new

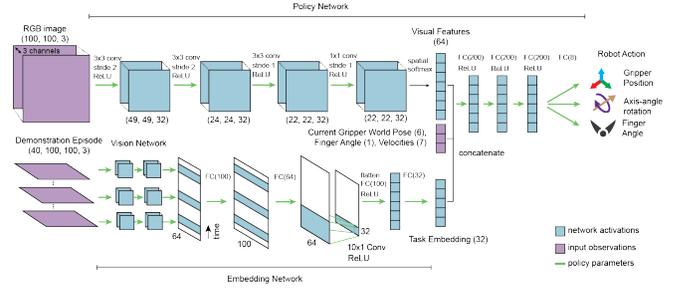


Fig. 3. We train visuomotor meta-imitation policies end-to-end, based on the setup described by [20]. Upper left and right: RGB observations are mapped to visual features, which are transformed into desired gripper positions. Lower: The task embedding network aggregates visual features over a sequence of image observations (obtained from a demonstration episode or policy rollout) into a “task embedding”, which is used to condition multi-task behavior from the policy network. Unlike [20], the task embedding network does not condition on observed gripper positions in the demo/rollouts.

task, and update the meta-policy to minimize the loss in Equation 2. In Section IV-A, we describe how we can learn a latent space of tasks jointly with the meta-policy, using only the human-provided one-shot imitation dataset and no extra supervision. In Section IV-B, we describe how we utilize a learned latent space of tasks and meta-policy to improve from autonomously collected data.

A. Learning a Latent Task Space Jointly with the Policy

Following prior work in one-shot imitation learning [26], [18], [20], we learn a meta-policy $\pi_{\theta}(a|o, d)$ that consists of two neural networks: an embedding network f_{θ} and a policy network g_{θ} , where θ represents the parameters of the two neural networks (shown together in Figure 3). Below, we detail the roles performed by these networks, their architecture and loss functions that we use to train them.

Embedding network and contrastive loss functions. The embedding network, represented as f_{θ} , accepts as input a demo of the task to be performed. The embedding network consists of a convolutional neural network followed by 1-D temporal convolutions, and embeds the demo into a fixed-length vector, which we denote as $f_{\theta}(d)$, and refer to as the demo embedding. Intuitively, we would like the embeddings to satisfy the following property: we want two demo embeddings to be close to each other if the demos correspond to the same task, and we wish them to be further apart for demos that belong to different tasks. Formally, this can be accomplished by considering the distance function $H_{\theta}(d_m, d_n) = \|f_{\theta}(d_m) - f_{\theta}(d_n)\|_2^2$. H_{θ} should be low when d_m and d_n correspond to the same task, and it should be high when they correspond to different tasks. Contrastive loss functions [37] satisfy this property, and are given by the following expression:

$$\mathcal{L}_c(\theta, \{\mathcal{D}_i\}) = \sum_{j=1}^M \sum_{k=1}^M \sum_{\substack{d_m^j \sim \mathcal{D}_j \\ d_n^k \sim \mathcal{D}_k}} \mathbb{1}(j = k) H_{\theta}(d_m^j, d_n^k) + \mathbb{1}(j \neq k) \max(0, \beta - H_{\theta}(d_m^j, d_n^k)),$$

where β is a margin (which we set to 1.0).

Policy network. The remaining part of the meta-policy consists of a policy network. It takes as input an image of the current scene (along with other parts of the robot’s state such as end-effector pose), and outputs a distribution over actions. In order to learn to predict actions from this information, we use the one-shot imitation learning loss \mathcal{L}_{oil} in Equation 2 in Section II. Our complete loss function,

$$\mathcal{L}(\theta, \{\mathcal{D}_i\}) = \mathcal{L}_{\text{oil}}(\theta, \{\mathcal{D}_i\}) + \mathcal{L}_c(\theta, \{\mathcal{D}_i\}) \quad (3)$$

is minimized using the Adam [38] optimizer, and we train both the embedding and policy networks jointly, sharing the convolutional layers between them. The output of this procedure results in the meta-policy π_θ , where θ denotes the learned parameters.

B. Autonomous Improvement

Our method is summarized in Alg. 1. We utilize the learned meta-policy π_θ for collecting data in new scenes. Given a scene, we condition the learned meta-policy on a randomly sampled demo from the initial dataset $\{\mathcal{D}_i\}$ (see Lines 9-11 in Alg. 1). We then run the meta-policy in this new scene, and add it to the trial dataset \mathcal{R} if the filtering function \mathcal{F} deems the trial to have performed some useful task (but does not provide any label about what the performed task was). We then search the trial dataset \mathcal{R} for similar trials by computing the cosine distance (i.e. a normalized dot-product) between the embedding of the newly collected trial, and all trials collected in the past (see Lines 17-18 in Alg. 1). If the cosine distance for any pair of trials is found to be above a pre-specified threshold α (which we determine using cross-validation to be 0.9), we add the pair of trials to the dataset as corresponding to a new task (see Lines 19-21 in Alg. 1). We then update our meta-policy π_θ using the one-shot imitation learning loss (see Line 24 in Alg. 1).

V. EXPERIMENTS

Our experiments seek to answer the following questions:

- 1) Does our method enable autonomous improvement? That is, can the trials generated by the policy in new environments improve one-shot imitation performance in comparison to only meta-imitation learning or only behavior cloning on a static dataset?
- 2) How does the performance vary with the number of trials collected?
- 3) How does the performance of our learned latent space-based pairing model compare to an oracle that pairs the autonomously collected data optimally?

We conduct our experiments on a realistic 3D simulation created using the Bullet physics engine [39], shown in Figure 4. It consists of a 7-DoF robotic gripper controlled with continuous position control from visual observations at 10Hz to accomplish manipulation tasks from four distinct task families, and contains over a hundred different kitchenware objects [40]. The four task families that we consider in our experiments are: button pressing, grasping, pushing and pick and place. Example tasks from each of the four families are shown in Figure 4. Different tasks within these families

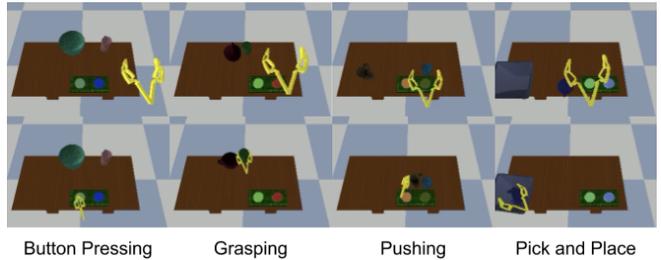


Fig. 4. Our dataset of tasks consists of four distinct task families: button pressing, grasping, pushing and pick and place. Within each task family, we have hundreds of tasks, and our dataset contains over a hundred distinct kitchenware objects. Our train set has 520 tasks, while our validation and test sets each have 40 tasks each. We tune hyperparameters on the validation tasks, and report final performance on the test tasks.

correspond to different sets of objects and manipulating those objects in different manners. For example, a pushing task in a scene containing both a mug and a bowl could be to push the bowl to the mug, or it could be to push the mug to the bowl. To succeed at a task, the robot must be able to perform the task from multiple different initial object arrangements such that it cannot simply memorize the motion from the one provided demonstration. The policy network, visualized in Figure 3, uses pixel inputs for all the tasks, i.e. an RGB image of size 100×100 , which gets passed in to the control network along with the 7-DoF gripper pose, three dimensions of which correspond to the its position in 3D space, three correspond to the Euler angles, and one corresponds to the finger angle. The output of the policy is the desired 7-DoF pose of the gripper, represented as the 3D position. We train on 520 tasks, use a validation set of 40 tasks for hyperparameter tuning, and show performance on a test set of 40 tasks, which contain held-out objects.

Demonstrations for the initial dataset $\{\mathcal{D}_i\}$ are collected by a human using an HTC Vive virtual reality system. Four demonstrations are collected for each task, and the object positions are randomized between demonstrations. At evaluation time, we are provided with one demonstration for a given test task, and the policy performance is evaluated with an object arrangement that is different from the demonstration.

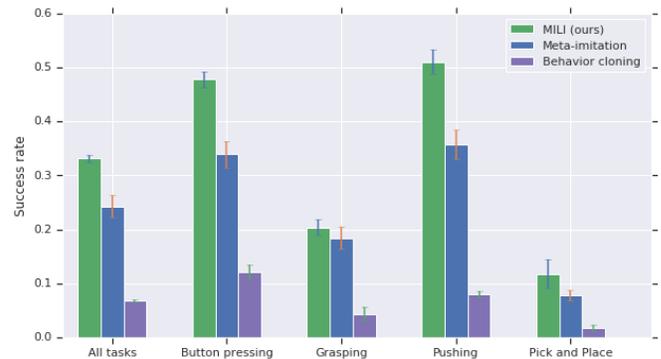


Fig. 5. We evaluate one-shot imitation performance of our policy on 40 unseen test tasks from four distinct task families. We ran our method, as well as both of the comparisons, with five random seeds, and report the average final performance and standard error across seeds. All methods use the same human demonstration data. We see that our method outperforms the meta-imitation learning baseline for all task families, indicating that the autonomously collected experience results in a better policy.

A. Autonomous Improvement

We first evaluate the central question of our paper: can the robot collect trials in new scenes and use them to improve its performance? We bootstrap the meta-policy from 800 human teleoperation demos on 200 different tasks, and learn a one-shot imitation meta-policy as well a latent task space from this data (see Line 5 in Algorithm 1). We collect 60K trials from 150 different scenes. We find that 8.2K of these trials pass the filtering function \mathcal{F} and get added to our trial dataset \mathcal{R} . Our method finds pairs of similar trials using a pairing threshold of $\alpha = 0.9$ in the trial dataset \mathcal{R} , adds them to $\{\mathcal{D}_i\}$, and re-runs meta-imitation learning (see Lines 17-24 in Algorithm 1). The performance of this policy is shown in Figure 5, and videos can be found on our project website¹. While it is possible to run several iterations of our method for continuous improvement, we only evaluate one round of our method in this section. We compare our method against: **Meta-imitation** This corresponds to optimizing the loss function described in Equation 2 on the human demonstrations dataset $\{\mathcal{D}_i\}$, and is representative of prior work [18]. **Behavior cloning** This corresponds to optimizing the loss function described in Equation 1 on all the demos from $\{\mathcal{D}_i\}$.

As shown in Figure 5, our method substantially outperforms the meta-imitation learning baseline, achieving a relative improvement of 36.8%. The meta-imitation baseline in turn outperforms the behavior cloning baseline, which is expected since the behavior cloning baseline does not incorporate task-specific information. This shows that our proposed method is capable of incorporating policy roll-outs in the learning process so as to improve the policy’s one-shot imitation performance on new, unseen tasks.

B. Varying the Number of Trials

A central claim in our paper is that a policy can continuously improve if it is able to collect more and more data and use it for learning. In this section, we aim to answer the following question: do we achieve monotonically improving performance as we collect more trials and add them to the dataset? To answer this question, we run our method with five random seeds while varying the number of trials collected: 10K, 20K, 40K and 60K. The results of this experiment are summarized in Figure 5, and show the following trend: as we keep collecting more trials, the one-shot imitation performance of our meta-policy for unseen tasks keeps improving. The relative improvements in performance are highest for the first 10K trials collected, but the performance starts levelling off as we collect more data.

C. Using Oracle Task Labels

Evaluating meta-imitation learning on an optimally paired dataset of policy rollouts allows us to measure what success rates we might hope MILI to achieve if we learned a perfect latent space. This oracle achieves a one-shot imitation success rate of 33.7% as compared to MILI’s 32.7%, suggesting that MILI can drastically reduce the burden of labeling tasks without significant loss in performance.

¹Project page: <https://sites.google.com/view/scalable-mili/>

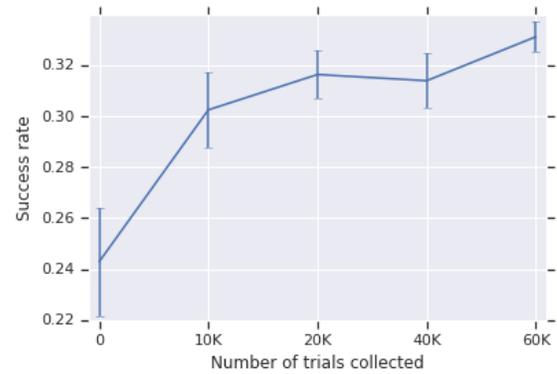


Fig. 6. Performance of MILI improves as we increase the number of trials.

VI. DISCUSSION AND FUTURE WORK

In this paper, we study how meta-imitation learning can be used to enable an agent to improve with autonomously collected data. Our aim is to preserve all of the desirable properties of imitation learning – simplicity and stability – while augmenting our method with one of the most powerful features of RL: the ability to improve from autonomously collected trials. Our approach is based on a simple observation: in a multi-task setting, the robot’s attempts to perform a new task, even if unsuccessful, may still serve as successful trials for *other* tasks. Therefore, if we can match them with corresponding prior trials and thereby form new task datasets for meta-learning, we can incorporate the robot’s own attempts as pseudo-demonstrations for other tasks. Unlike RL, our method does not require any task-specific reward function, only a general, task-agnostic filtering function that indicates whether a given trial is useful for *any* task. This makes it straightforward to scale our approach to large skill repertoires. After bootstrapping our meta-imitation policy from a few hundred human-provided demonstrations, we can collect tens of thousands of autonomous trials, and use them to meta-train a final policy that substantially outperforms standard meta-imitation learning on the initial dataset.

Limitations. Our work opens several possible directions for future work. Currently, our method still requires a filtering function, which means we need to obtain a binary label from a human oracle for each trial that we collect to decide if it should be included in the training set. While this signal is easy to obtain, future work might investigate learning-based methods that perform this filtering automatically, further boosting the scalability of this approach. Further, while our method improves with more data, the overall success rates are still relatively low due to the challenging nature of the environment. We expect that further advances in vision-based meta-learning will complement the approach in this paper to boost performance. Finally, another exciting direction for future work is to devise a complete lifelong learning approach based on our method, performing repeated iterations of data collection and meta-imitation on a real-world robotic system so as to continually improve the policy on both the original set of tasks and new tasks discovered during data collection.

REFERENCES

- [1] D. A. Pomerleau, "Alvin: An autonomous land vehicle in a neural network," in *Neural Information Processing Systems (NIPS)*, 1989, pp. 305–313.
- [2] 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.
- [3] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. MIT press Cambridge, 1998, vol. 1, no. 1.
- [4] J. Kober and J. Peters, "Reinforcement learning in robotics: A survey," in *Reinforcement Learning*. Springer, 2012, pp. 579–610.
- [5] D. Kalashnikov, A. Irpan, P. Pastor, J. Ibarz, A. Herzog, E. Jang, D. Quillen, E. Holly, M. Kalakrishnan, V. Vanhoucke, and S. Levine, "Scalable deep reinforcement learning for vision-based robotic manipulation," in *Proceedings of The 2nd Conference on Robot Learning*.
- [6] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. Riedmiller, "Playing atari with deep reinforcement learning," *arXiv preprint arXiv:1312.5602*, 2013.
- [7] T. Haarnoja, A. Zhou, K. Hartikainen, G. Tucker, S. Ha, V. K. Jie Tan, H. Zhu, A. Gupta, P. Abbeel, and S. Levine, "Soft actor-critic algorithms and applications," Tech. Rep., 2018.
- [8] A. Singh, L. Yang, K. Hartikainen, C. Finn, and S. Levine, "End-to-end robotic reinforcement learning without reward engineering," *Robotics: Science and Systems*, 2019.
- [9] P. Henderson, R. Islam, P. Bachman, J. Pineau, D. Precup, and D. Meger, "Deep reinforcement learning that matters," in *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [10] A. Irpan, "Deep reinforcement learning doesn't work yet," <https://www.alexirpan.com/2018/02/14/rl-hard.html>, 2018.
- [11] G. Dulac-Arnold, D. Mankowitz, and T. Hester, "Challenges of real-world reinforcement learning," *arXiv preprint arXiv:1904.12901*, 2019.
- [12] W. Sun, A. Venkatraman, G. J. Gordon, B. Boots, and J. A. Bagnell, "Deeply aggregated: Differentiable imitation learning for sequential prediction," in *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org, 2017, pp. 3309–3318.
- [13] S. Ross and D. Bagnell, "Efficient reductions for imitation learning," in *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, 2010, pp. 661–668.
- [14] 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.
- [15] Y. Duan, M. Andrychowicz, B. Stadie, J. Ho, J. Schneider, I. Sutskever, P. Abbeel, and W. Zaremba, "One-shot imitation learning," *Neural Information Processing Systems (NIPS)*, 2017.
- [16] C. Finn, T. Yu, T. Zhang, P. Abbeel, and S. Levine, "One-shot visual imitation learning via meta-learning," *Conference on Robot Learning (CoRL)*, 2017.
- [17] T. Yu, C. Finn, A. Xie, S. Dasari, T. Zhang, P. Abbeel, and S. Levine, "One-shot imitation from observing humans via domain-adaptive meta-learning," *Robotics: Science and Systems (RSS)*, 2018.
- [18] S. James, M. Bloesch, and A. J. Davison, "Task-embedded control networks for few-shot imitation learning," *arXiv preprint arXiv:1810.03237*, 2018.
- [19] D. Huang, S. Nair, D. Xu, Y. Zhu, A. Garg, L. Fei-Fei, S. Savarese, and J. C. Niebles, "Neural task graphs: Generalizing to unseen tasks from a single video demonstration," 2018.
- [20] A. Zhou, E. Jang, D. Kappler, A. Herzog, M. Khansari, P. Wohlhart, Y. Bai, M. Kalakrishnan, S. Levine, and C. Finn, "Watch, try, learn: Meta-learning from demonstrations and reward."
- [21] T. L. Paine, S. G. Colmenarejo, Z. Wang, S. Reed, Y. Aytar, T. Pfaff, M. W. Hoffman, G. Barth-Maron, S. Cabi, D. Budden *et al.*, "One-shot high-fidelity imitation: Training large-scale deep nets with rl," *arXiv preprint arXiv:1810.05017*, 2018.
- [22] L. P. Kaelbling, "Learning to achieve goals," in *Proceedings of the 13th International Joint Conference on Artificial Intelligence 1993*, 1993.
- [23] M. Andrychowicz, D. Crow, A. Ray, J. Schneider, R. Fong, P. Welinder, B. McGrew, J. Tobin, P. Abbeel, and W. Zaremba, "Hindsight experience replay," in *Neural Information Processing Systems (NIPS)*, 2017, pp. 5055–5065.
- [24] T. Yu, G. Shevchuk, D. Sadigh, and C. Finn, "Unsupervised visuomotor control through distributional planning networks," *arXiv preprint arXiv:1902.05542*, 2019.
- [25] P. Rauber, A. Ummadisingu, F. Mutz, and J. Schmidhuber, "Hindsight policy gradients," *arXiv preprint arXiv:1711.06006*, 2017.
- [26] Y. Duan, M. Andrychowicz, B. Stadie, J. Ho, J. Schneider, I. Sutskever, P. Abbeel, and W. Zaremba, "One-shot imitation learning," *Neural Information Processing Systems (NIPS)*, 2017.
- [27] 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.
- [28] A. Y. Ng and S. J. Russell, "Algorithms for inverse reinforcement learning," in *International Conference on Machine Learning (ICML)*, 2000.
- [29] P. Abbeel and A. Y. Ng, "Apprenticeship learning via inverse reinforcement learning," in *International Conference on Machine Learning (ICML)*. ACM, 2004, p. 1.
- [30] N. D. Ratliff, J. A. Bagnell, and M. A. Zinkevich, "Maximum margin planning," in *International Conference on Machine Learning (ICML)*. ACM, 2006.
- [31] M. Wulfmeier, P. Ondruska, and I. Posner, "Maximum entropy deep inverse reinforcement learning," in *Neural Information Processing Systems Conference, Deep Reinforcement Learning Workshop*, vol. abs/1507.04888, 2015.
- [32] 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.
- [33] J. Ho and S. Ermon, "Generative adversarial imitation learning," in *Neural Information Processing Systems (NIPS)*, 2016.
- [34] R. Rahmatizadeh, P. Abolghasemi, L. Bölöni, and S. Levine, "Vision-based multi-task manipulation for inexpensive robots using end-to-end learning from demonstration," *CoRR*, vol. abs/1707.02920, 2017.
- [35] K. Hausman, J. T. Springenberg, Z. Wang, N. Heess, and M. A. Riedmiller, "Learning an embedding space for transferable robot skills," in *International Conference on Learning Representations (ICLR)*, 2018.
- [36] Y. Duan, J. Schulman, X. Chen, P. L. Bartlett, I. Sutskever, and P. Abbeel, "R²: Fast reinforcement learning via slow reinforcement learning," *arXiv preprint arXiv:1611.02779*, 2016.
- [37] R. Hadsell, S. Chopra, and Y. LeCun, "Dimensionality reduction by learning an invariant mapping," in *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2006)*, 17–22 June 2006, New York, NY, USA, 2006.
- [38] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," *International Conference on Learning Representations (ICLR)*, 2015.
- [39] E. Coumans and Y. Bai, "Pybullet, a python module for physics simulation for games, robotics and machine learning," *GitHub repository*, 2016.
- [40] X. Yan, J. Hsu, M. Khansari, Y. Bai, A. Pathak, A. Gupta, J. Davidson, and H. Lee, "Learning 6-dof grasping interaction via deep geometry-aware 3d representations," in *International Conference on Robotics and Automation (ICRA)*, 2018.