Learning from Guided Play: A Scheduled Hierarchical Approach for Improving Exploration in Adversarial Imitation Learning

Anonymous Author(s)
Affiliation
Address
email

Abstract

Effective exploration continues to be a significant challenge that prevents the deployment of reinforcement learning for many physical systems. This is particularly true for systems with continuous and high-dimensional state and action spaces, such as robotic manipulators. The challenge is accentuated in the sparse rewards setting, where the low-level state information required for the design of dense rewards is unavailable. Adversarial imitation learning (AIL) can partially overcome this barrier by leveraging expert-generated demonstrations of optimal behaviour and providing, essentially, a replacement for dense reward information. Unfortunately, the availability of expert demonstrations does not necessarily improve an agent's capability to explore effectively and, as we empirically show, can lead to inefficient or stagnated learning. We present Learning from Guided Play (LfGP), a framework in which we leverage expert demonstrations of, in addition to a main task, multiple auxiliary tasks. Subsequently, a hierarchical model is used to learn each task reward and policy through a modified AIL procedure, in which exploration of all tasks is enforced via a scheduler composing different tasks together. This affords many benefits: improved learning efficiency for main tasks with challenging bottleneck transitions, learning can be completed without true reward functions, reusable expert data between tasks, and transfer learning through the reuse of learned auxiliary task models becomes possible. Our experimental results in a challenging multitask robotic manipulation domain indicate that our method compares favourably to supervised imitation learning and to a state-of-the-art AIL method.

1 Introduction

Exploration is a crucial part of effective reinforcement learning (RL). A variety of methods have been developed for attempting to optimize the exploration-exploitation trade-off of RL agents [1, 2, 3], but the development of a method that generalizes across domains is still very much an open research problem. A simple, well-known approach for improving learning efficiency is to provide a dense, or "shaped" reward to learn from, but this can be very challenging to design appropriately [4]. Furthermore, the environment may not directly provide the low-level state information required for such a reward. An alternative to providing a dense reward is to learn a reward function from expert demonstrations of a task, in a process known as inverse RL (IRL). Most current, state-of-the-art approaches to IRL are in the class known as adversarial imitation learning (AIL). In AIL, rather than learning a true reward function, the policy and the discriminator form a two-player mini-max optimization problem: the policy aims to confuse the discriminator by minimizing the objective while the discriminator aims identify expert-like data by maximizing the objective [5, 6, 7].

Submitted to 35th Conference on Neural Information Processing Systems (NeurIPS 2021). Do not distribute.

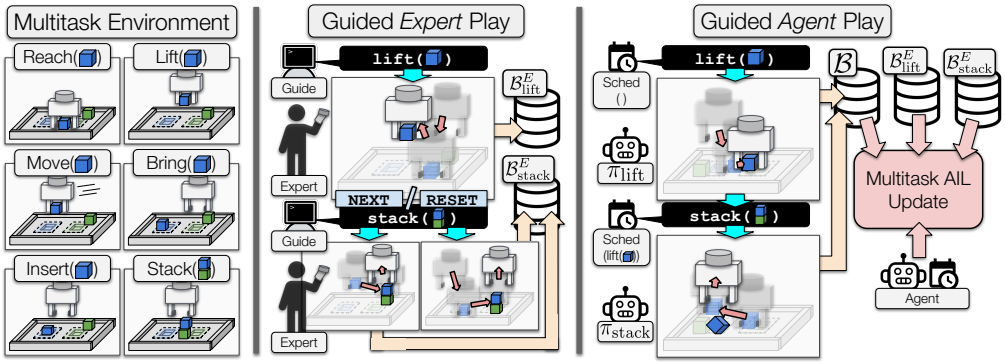

Figure 1: The main components of our system for learning from guided play. In a multitask environment, a guide prompts an expert for a mix of multitask demonstrations, after which we learn a multitask policy through scheduled hierarchical AIL.

AIL has been shown to be effective in many continuous-control environments [5, 7], but its application to long-horizon robotic manipulation tasks with a wide distribution of possible initial configurations remains challenging [7, 8, 9]. In this work, we investigate the use of AIL in a multitask robotic manipulation domain (see the left side of Figure 1), specifically investigating stacking, unstacking and then stacking, reaching followed by bringing to a target, and reaching followed by inserting. We find that a state-of-the-art AIL method, in which off-policy learning is used to maximize sample efficiency [7], performs poorly on these complex manipulation tasks. In all of our tasks, AIL even performs worse than the simpler alternative of supervised learning (often referred to as behavioural cloning (BC). We hypothesize that this is because AIL, and more generally IRL, can provide something akin to a dense reward, but because this learned reward is not necessarily optimal for teaching, AIL alone may not necessarily enforce sufficiently diverse exploration to allow effective learning.

To overcome this limitation of AIL, we present Learning from Guided Play (LfGP), in which we combine AIL with a scheduled approach to hierarchical RL (HRL), allowing an agent to "play" in the environment using the expert "guide." Using expert demonstrations of multiple auxiliary tasks, along with a main task, we demonstrate that a hierarchical agent is able to learn much more effectively than with AIL alone. This approach also affords several other benefits: learning a challenging main task becomes more efficient than supervised learning in terms of required expert data, no hand-crafted reward functions are needed as in [10], expert data on auxiliary tasks can be reused from one main task to another, and learning time can be reduced by reusing previously-learned auxiliary models on new main tasks through simple transfer learning. We choose this name in reference to the work upon which we build [10], which is in turn inspired by the "playful" phase of learning experienced by children, as opposed to goal-directed learning. In our case, *guided* represents two separate but related ideas: first, the expert guides this play, as opposed to requiring hand-crafted sparse rewards as in [10] (right side of Figure 1), and second, the expert gathering of multitask, semi-structured demonstrations is *guided* by uniform-random task selection (middle of Figure 1), rather than requiring the expert to craft their own high-level play goals as in [11]. Our specific contributions are:

1. A novel approach to AIL that learns a reward and policy for a challenging main task by simultaneously learning rewards and policies for auxiliary tasks.
2. A variety of manipulation experiments illustrating the viability of our approach, including comparisons with competing approaches to AIL and imitation learning.
3. An extension of our method to the domain of transfer learning.
4. An open source implementation of our algorithm and environments.[1]

2 Related Work

Imitation learning (IL) is often divided into two main categories: behavioural cloning (BC) [12] and inverse reinforcement learning (IRL) [13, 14]. BC recovers the expert policy via supervised learning, but it suffers from compounding errors due to covariate shift [15]. Alternatively, IRL

[1] To maintain double blind standards, for now we only include code as part of our supplementary material.

partially alleviates the covariate shift problem by estimating the reward function and then applying RL using the learned reward. The state-of-the-art approach to IRL is adversarial imitation learning (AIL) [5, 7, 16, 17], in which the expert policy is recovered by matching the occupancy measure between the generated data and the demonstration data. Our proposed method extends existing AIL algorithms by enabling more effective exploration via the use of a multitask model.

Agents learned via hierarchical reinforcement learning (HRL) methods, which aim to learn and act over multiple levels of temporal abstractions in long planning horizon tasks, are shown to provide more effective exploration than agents operating with only a single level [10, 18, 19, 20, 21]. Our work most closely resembles hierarchical AIL methods that attempt to combine AIL with HRL [17, 22, 23, 24]. Existing works [22, 23, 24] often formulate the hierarchical agent using the Option framework [18] and recover the reward function using AIL [5]. Both [22] and [24] leverage task-specific expert demonstrations to learn options using mixture-of-experts and expectation-maximization strategies respectively. Our work focuses on expert demonstrations that include multiple tasks, each of which encodes clear semantic meaning. In the multitask setting, [17] and [23] leverage unsegmented, multitask expert demonstrations to learn low-level policies via a latent variable model. In contrast, comparable to [25], we learn our models using sets of segmented demonstrations. Compared with learning from unsegmented demonstrations, the use of segmented demonstrations ensures that we know which auxiliary tasks our model will be learning, and opens up the possibility for transfer learning without further assumptions. Finally, we deviate from the Options framework and build upon Scheduled Auxiliary Control (SAC-X) to train our hierarchical agent as it is shown to work well in complex manipulation tasks [10, 26].

3 Problem Formulation

We define a Markov Decision Process (MDP) as $\mathcal{M} = \langle \mathcal{S}, \mathcal{A}, R, \mathcal{P}, \rho_0, \gamma \rangle$. The sets $\mathcal{S}$ and $\mathcal{A}$ are respectively the state and action space, $R : S \times A \to \mathbb{R}$ is a reward function, $\mathcal{P} : S \times A \times S \to \mathbb{R}_{\geq 0}$ is the state-transition environment dynamics distribution, $\rho_0 : S \to \mathbb{R}_{\geq 0}$ is the initial state distribution, and γ is the discount factor. Actions are sampled from a stochastic policy $\pi(a|s)$. The policy π interacts with the environment to yield experience (s_t, a_t, r_t, s_{t+1}) for $t = 0, \dots, \infty$, where $s_0 \sim \rho_0(\cdot), a_t \sim \pi(\cdot|s_t), s_{t+1} \sim \mathcal{P}(\cdot|s_t, a_t), r_t = R(s_t, a_t)$.

We assume infinite-horizon, non-terminating environments where t is unbounded for notational convenience, but our method also easily applies to the finite-horizon case. We aim to learn a policy π that maximizes the expected return $J(\pi) = \mathbb{E}_\pi \left[G(\tau_{0:\infty})\right] = \mathbb{E}_\pi \left[\sum_{t=0}^{\infty} \gamma^t R(s_t, a_t)\right]$, where $\tau_{t:\infty} = \{(s_t, a_t), \dots\}$ is the trajectory starting in state s_t, and $G(\tau_{t:\infty})$ is the return of the trajectory.

In this work, we focus on imitation learning (IL) in which R is unknown and instead, we are given a finite set of expert demonstration (s, a) pairs $\mathcal{B}^E = \left\{(s, a)^{E,(1)}, \dots\right\}$. In AIL, we attempt to simultaneously learn π and a discriminator $D : S \times A \to \mathbb{R}_{\geq 0}$ that differentiates between expert samples $(s, a)^E$ and policy samples $(s, a)^\pi$, and subsequently define R using D [5, 7]. To accommodate hierarchical learning, we augment $\mathcal{M}$ to contain auxiliary tasks, where $\mathcal{T}_{\text{aux}} = \{\mathcal{T}_1, \dots, \mathcal{T}_K\}$ are separate MDPs that share $\mathcal{S}, \mathcal{A}, \mathcal{P}, \rho_0$ and γ with the main task $\mathcal{T}_{\text{main}}$ but have their own reward functions R_k. With this modification, we refer to entities in our model that are specific to task $\mathcal{T} \in \mathcal{T}_{\text{all}}$, $\mathcal{T}_{\text{all}} = \mathcal{T}_{\text{aux}} \cup \{\mathcal{T}_{\text{main}}\}$, as $(\cdot)_\mathcal{T}$. We assume to have a set of expert data $\mathcal{B}^E_\mathcal{T}$ for each task.

4 Learning from Guided Play (LfGP)

We now introduce Learning from Guided Play (LfGP). We focus on complex manipulation tasks that require synthesizing a composition of skills. Recent work showed that, to successfully learn an agent to solve complex tasks, it is important to have effective exploration. The benefit of HRL is mainly attributed to exploration [21]. To improve exploration, we propose a method that combines the benefits of AIL and HRL. Our primary goal is to learn a policy $\pi_{\mathcal{T}_{\text{main}}}$ that can solve the main task $\mathcal{T}_{\text{main}}$, with a secondary goal of also learning auxiliary task policies $\pi_{\mathcal{T}_1}, \dots, \pi_{\mathcal{T}_K}$ for the purpose of transfer learning. More specifically, we derive a hierarchical learning objective that is decomposed into three parts: i) recovering the reward function of each task with expert demonstrations, ii) training all policies to achieve their respective goals, and iii) using all policies for effective exploration in $\mathcal{T}_{\text{main}}$. For a summary of the algorithm, see our supplementary material.

4.1 Learning the Reward Function

We first describe how to recover the reward functions from expert demonstrations. For each task $\mathcal{T} \in \mathcal{T}_{\text{all}}$, we learn a discriminator $D_{\mathcal{T}}(s, a)$ that is used to define the reward function for policy optimization. The discriminator is trained in an off-policy manner [16]:

$$\max_{D_{\mathcal{T}}} \mathbb{E}_{\mathcal{B}} \left[\frac{\rho_{\pi_{\mathcal{T}}}(s,a)}{\rho_{\mathcal{B}}(s,a)} \log\left(1 - D_{\mathcal{T}}(s,a)\right) \right] + \mathbb{E}_{\mathcal{B}_{\mathcal{T}}^{E}} \left[\log\left(D_{\mathcal{T}}(s,a)\right) \right], \tag{1}$$

where $\mathcal{B}$ is the replay buffer storing transitions collected while interacting with the environment, $\rho_{\pi_{\mathcal{T}}}(s, a)$ the intention occupancy measure, and $\rho_{\mathcal{B}}$ the distribution induced by the replay buffer $\mathcal{B}$. In practice, following [16], we sample the transitions from $\mathcal{B}$ without applying importance sampling:

$$\max_{D_{\mathcal{T}}} \mathbb{E}_{\mathcal{B}} \left[\log\left(1 - D_{\mathcal{T}}(s,a)\right) \right] + \mathbb{E}_{\mathcal{B}_{\mathcal{T}}^{E}} \left[\log\left(D_{\mathcal{T}}(s,a)\right) \right]. \tag{2}$$

Each discriminator tries to differentiate the occupancy measure between the distributions induced by $\mathcal{B}_{\mathcal{T}}^{E}$ and $\mathcal{B}$. As a result, the policy $\pi_{\mathcal{T}}$ attempts to confuse the discriminator by minimizing Equation 2 through changing the distribution induced by the replay buffer to be similar to the expert demonstrations. In our implementation, the discriminators minimize the joint discriminator loss

$$\mathcal{L}(D) = -\sum_{\mathcal{T}}^{\mathcal{T}_{\text{all}}} \mathbb{E}_{\mathcal{B}} \left[\log\left(1 - D_{\mathcal{T}}(s,a)\right) \right] + \mathbb{E}_{\mathcal{B}_{\mathcal{T}}^{E}} \left[\log\left(D_{\mathcal{T}}(s,a)\right) \right]. \tag{3}$$

The resulting $D_{\mathcal{T}}$ can be used to define various reward functions [16]—we define the reward function for each task $\mathcal{T}$ to be $D_{\mathcal{T}}$. Thus, we define the expected return for $\pi_{\mathcal{T}}$ as

$$J(\pi_{\mathcal{T}}) = \mathbb{E}_{\pi_{\mathcal{T}}} \left[G_{\mathcal{T}}(\tau_{t:\infty}) \right] = \mathbb{E}_{\pi_{\mathcal{T}}} \left[\sum_{t=0}^{\infty} \gamma^t D_{\mathcal{T}}(s_t, a_t) \right]. \tag{4}$$

4.2 Learning the Hierarchical Agent

We adapt SAC-X [10] to learn the hierarchical agent. The agent includes two levels of policies: intention policies (low-level) and a scheduler (high-level), as well as the Q-functions and the discriminators. The intentions aim to solve their corresponding tasks (i.e. intentions are $\pi_{\mathcal{T}}$ aiming to maximize task return $J(\pi_{\mathcal{T}})$), whereas the scheduler aims to maximize the expected return for $\mathcal{T}_{\text{main}}$ by selecting a sequence of intentions to interact with the environment. For the remainder of this paper, when we refer to a policy, we are referring to an intention policy, as opposed to the scheduler, unless otherwise specified.

4.2.1 Learning the Intentions

We learn each intention using Soft Actor Critic (SAC) [27], an actor-critic algorithm that optimizes for the entropy-regularized objective, though any off-policy RL algorithm would suffice. We augment Equation 4 with entropy regularization as

$$J(\pi_{\mathcal{T}}) = \mathbb{E}_{\pi_{\mathcal{T}}} \left[\sum_{t=0}^{\infty} \gamma^t \left(D_{\mathcal{T}}(s_t, a_t) + \alpha \mathcal{H}(\pi_{\mathcal{T}}(\cdot|s_t)) \right) \right], \tag{5}$$

where the learned temperature α determines the importance of the entropy term against the reward, and $\mathcal{H}(\pi_{\mathcal{T}}(\cdot|s_t))$ is the entropy of the intention $\pi_{\mathcal{T}}$ at state s_t. The corresponding soft Q-function $Q_{\mathcal{T}}(s_t, a_t)$ for task $\mathcal{T}$ is defined as

$$Q_{\mathcal{T}}(s_t, a_t) = D_{\mathcal{T}}(s_t, a_t) + \mathbb{E}_{\pi_{\mathcal{T}}} \left[\sum_{t=0}^{\infty} \gamma^t \left(D_{\mathcal{T}}(s_{t+1}, a_{t+1}) + \alpha \mathcal{H}(\pi_{\mathcal{T}}(\cdot|s_{t+1})) \right) \right]. \tag{6}$$

The intentions maximize the joint policy improvement objective:

$$\mathcal{L}(\pi) = \sum_{\mathcal{T}}^{\mathcal{T}_{\text{all}}} \mathbb{E}_{s \sim \mathcal{B}, a \sim \pi_{\mathcal{T}}(\cdot|s)} \left[Q_{\mathcal{T}}(s,a) - \alpha \log \pi_{\mathcal{T}}(a|s) \right]. \tag{7}$$

Notice that the replay buffer $\mathcal{B}$ contains transitions generated using different intentions. This allows each intention $\pi_{\mathcal{T}}$ to learn optimal actions for its respective task starting at any state visited by any intention. In other words, other intentions $\pi_{\mathcal{T}' \neq \mathcal{T}}$ interacting with the environment can provide exploration samples for $\pi_{\mathcal{T}}$.

For policy evaluation, the soft Q-functions $Q_{\mathcal{T}}$ for each $\pi_{\mathcal{T}}$ minimize the joint soft Bellman residual

$$\mathcal{L}(Q) = \sum_{\mathcal{T}}^{\mathcal{T}_{\text{all}}} \mathbb{E}_{(s,a,r,s') \sim \mathcal{B}, a' \sim \pi_{\mathcal{T}}(\cdot|s')} \left[\frac{1}{2} (Q_{\mathcal{T}}(s,a) - \delta_{\mathcal{T}})^2 \right], \tag{8}$$

$$\delta_{\mathcal{T}} = r + \gamma Q_{\mathcal{T}}(s', a') - \alpha \log \pi_{\mathcal{T}}(a'|s'). \tag{9}$$

4.2.2 Learning the Scheduler

The scheduler composes a sequence of intentions to interact with the environment. The scheduler can switch between intentions every ξ timesteps. With $\gamma < 1$, we define H as the total number of possible intention switches within an episode.[2] The H intention choices made within the episode are defined as $\mathcal{T}^{0:H-1} = \left\{ \mathcal{T}^{(0)}, \ldots, \mathcal{T}^{(H-1)} \right\}$, where $\mathcal{T}^{(h)} \in \mathcal{T}_{\text{all}}$. The main task's return given the intention choices is then defined as

$$G_{\mathcal{T}_{\text{main}}}(\mathcal{T}^{0:H-1}) = \sum_{h=0}^{H-1} \sum_{t=h\xi}^{(h+1)\xi-1} \gamma^t D_{\mathcal{T}_{\text{main}}}(s_t, a_t), \tag{10}$$

where $a_t \sim \pi_{\mathcal{T}^{(h)}}(\cdot|s_t)$ is the action taken at timestep t, sampled from the chosen intention $\mathcal{T}^{(h)}$ in the hth scheduler period. We define the Q-function for the scheduler as $Q_S(\mathcal{T}^{0:h-1}, \mathcal{T}^{(h)}) = \mathbb{E}_{P_S} \left[G_{\mathcal{T}_{\text{main}}}(\mathcal{T}^{h:H-1}) | \mathcal{T}^{0:h-1} \right]$ and represent the scheduler as a Boltzmann distribution P_S over the Q-function. The scheduler maximizes the expected return of the main task following the scheduler:

$$\mathcal{L}(S) = \mathbb{E}_{P_S} \left[Q_S(\emptyset, \mathcal{T}^{(0)}) \right]. \tag{11}$$

We use Monte Carlo returns to estimate Q_S, estimating the expected return using the exponential moving average

$$Q_S(\mathcal{T}^{0:h-1}, \mathcal{T}^{(h)}) = (1-\phi) Q_S(\mathcal{T}^{0:h-1}, \mathcal{T}^{(h)}) + \phi G_{\mathcal{T}_{\text{main}}}(\mathcal{T}^{h:H}), \tag{12}$$

where $\phi \in [0, 1]$ represents the amount of discounting on older returns and $G_{\mathcal{T}_{\text{main}}}(\mathcal{T}^{h:H})$ is the cumulative discounted return of the trajectory τ starting at timestep $h\xi$.

4.3 Expert Data Collection

We assume that each $\mathcal{T} \in \mathcal{T}_{\text{all}}$ has, for evaluation purposes only, a binary indicator of success. In single-task imitation learning (IL) where this assumption is valid, expert data is typically collected by allowing the expert to control the agent until success conditions are met. At that point, the environment is reset following ρ_0 and collection is repeated for a fixed number of episodes or (s, a) pairs. In LfGP, and more generally for multitask IL, expert data can also be collected this way for each $\mathcal{T}$ separately: we refer to this strategy, which we use for most of our experiments, as *reset-based* expert data collection.

A limitation of reset-based expert data collection is that, when training LfGP, the initial state of each individual policy can be any state that $\mathcal{M}$ has been left in by a previous policy, which may include states not in the distribution of ρ_0. This "transition" initial state distribution, which we call $\rho_0(s|\mathcal{T}')$, where $\mathcal{T}' \in \mathcal{T}_{\text{all}}$ corresponds to the previously running policy $\pi_{\mathcal{T}'}$, would be challenging to sample from—it relies on the policies, and it may include states which are impractical to manually reset to (e.g. objects may start off as grasped or in mid-air). Consequently, we propose an alternative data collection strategy: *play-based* expert data collection, where we alternate between uniformly sampling the next task for an expert to complete and having the expert execute that task until success. In our implementation, we also reset the environment following ρ_0 periodically. See middle section of Figure 1 for a comparison of the two methods.

[2] We can choose H in the infinite horizon setting using the effective horizon $\frac{1}{1-\gamma}$.

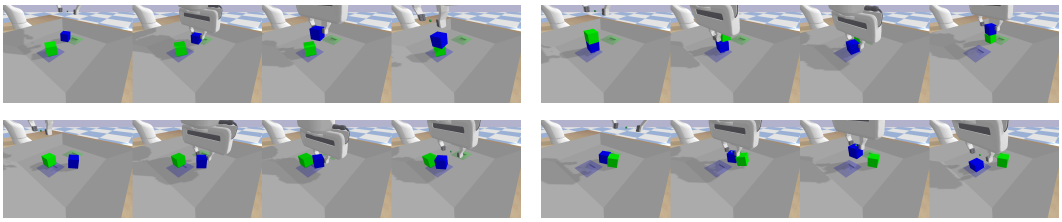

Figure 2: Example successful runs of our four main tasks. Left to right, top to bottom: Stack, Unstack-Stack, Bring, Insert.

5 Experiments

In our experiments, we attempt to answer the following questions:

1. How does the performance of LfGP compare with other existing state-of-the-art IL methods in challenging manipulation tasks, in terms of success rate, environment sample efficiency, and expert sample efficiency?
2. Are models trained using LfGP able to transfer from one main task to another? Does this transfer afford any benefits over learning the new main task from scratch?
3. How does performance using reset-based expert data compare with play-based expert data (see Section 4.3)?

5.1 Experimental Setup

We begin by describing our experimental setup. We complete experiments in a simulation environment containing a 7-DOF Franka Emika Panda manipulator, one green and one blue block, fixed zones corresponding to the green and blue blocks, and one slot in each zone with $<$ 1mm tolerance for fitting the blocks (see bottom right of Figure 2). The environment is designed such that several different challenging tasks can be completed with a common observation and action space. The main tasks that we investigate are Stack, Unstack-Stack, Bring, and Insert (see Figure 2 for examples). When learning Unstack-Stack, the green block is initialized on top of the blue block. For more details on our environment, see our supplementary material. We also define a set of auxiliary tasks: Open-Gripper, Close-Gripper, Reach, Lift, Move-Object, and Bring (Bring is both a main task and an auxiliary task for Insert), all of which have some degree of re-usability between main tasks. For a breakdown of how auxiliary tasks are used for each main task, see Table 1. Our supplementary material contains more detail on how we define success in each of the tasks.

We implemented several multitask and single-task baselines to compare our method to. The multitask algorithms, like our method, simultaneously attempt to learn to complete the main task as well as the auxiliary tasks. The single-task algorithms only attempt to learn to complete the main task. In general, we consider that a multitask algorithm is more useful than a single-task algorithm, given the potential to reuse trained models for learning new tasks, as we show in Section 5.3.

The first multitask algorithm that we compare to is a "No Schedule" version of our method (LfGP-NS), where the models are trained similarly to LfGP, but the scheduler only chooses the main task; this is meant to be a similar baseline to the Intentional Unintentional Agent shown in [10, 28]. The expected benefit is that the policies, the Q-functions, and the discriminators share their initial layers between tasks, so the auxiliary tasks can accelerate learning the representation for the main-task models [29]. We also compare to a multitask variant of behavioural cloning (BC), in which we simultaneously train every task by minimizing a multitask mean squared error objective

$$\mathcal{L}(\pi) = \sum_{\mathcal{T}}^{\mathcal{T}_{\text{all}}} \sum_{(s,a)\in\mathcal{B}_{\mathcal{T}}^{E}} \left(\pi_{\mathcal{T}}(s) - a\right)^2, \tag{13}$$

taking only the mean action outputs from each $\pi_{\mathcal{T}}$.

The single-task algorithms we compare to are BC and Discriminator-Actor-Critic (DAC) [7]. To make a fairer comparison and evaluate our method's expert sample efficiency, we provide BC and DAC with an equivalent amount of *total* expert data as our multitask methods as shown in Table 1.

Table 1: A summary of the reusability of the expert data used to generate the performance results described in Section 5.2. Each letter under "Dataset Sizes" is the first letter of a single (auxiliary) task, and bolded letters indicate that a dataset was reused for more than one main task (e.g., **O**pen-Gripper was used for all four main tasks). The numbers in this table refer to (s, a) pairs.

Algorithm Class	Task	Dataset Sizes	Reusable	Single	Total
Multitask	Stack	**OCSRLM**: 9k/task	**45k**	9k	54k
	U-Stack	**OCURLM**: 9k/task	**45k**	9k	54k
	Bring	**OCBRLM**: 9k/task	**54k**	0	54k
	Insert	**OCIBRLM**: 9k/task	**54k**	9k	63k
Single-task	Stack	S: 54k (less: 9k)	0	54k (9k)	54k (9k)
	U-Stack	U: 54k (less: 9k)	0	54k (9k)	54k (9k)
	Bring	B: 54k (less: 9k)	0	54k (9k)	54k (9k)
	Insert	I: 63k (less: 9k)	0	63k (9k)	63k (9k)

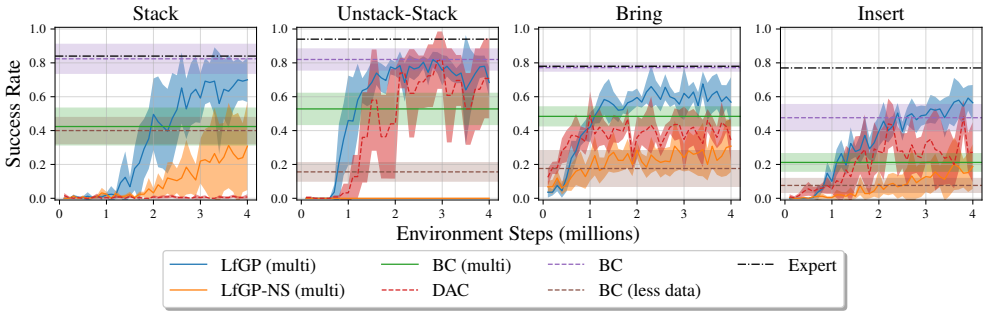

Figure 3: Performance results for our method (LfGP) and the baselines, as described in Section 5.1. Solid lines are multitask methods, dotted lines are single-task methods, and shaded area corresponds to standard deviation across five seeds.

To show the possible benefits or detriments of learning only a single main task with an equivalent amount of main task data as the multitask algorithm, we also show results for single-task BC with an equivalent amount of main task data to the multitask algorithms, which we call BC (less data).

We gathered expert data by first training an expert policy using SAC-X [10]. See further details of this procedure in our supplementary material. We then ran the expert policies to collect various amounts of expert data as described in Table 1. In these experiments, in the interest of testing the benefits of having reusable expert data, we use reset-based expert data (see Section 4.3).

5.2 Performance Results

For the AIL algorithms, we freeze the policies every 100k interactions and evaluate those policies for 50 randomized episodes, only using the mean action outputs. We trained each agent to 4 million environment steps. For the BC algorithms, we train the policies by dividing each expert buffer into a 70%/30% train/validation split, taking the policy after validation error has not improved for 100 epochs and evaluating it on 50 randomized episodes. For all algorithms, we test across five seeds and report the mean and standard deviation of all seeds. For more details of our training algorithms, including hyperparameters, see our supplementary material. The performance results for all methods with respect to the main tasks we tested are shown in Figure 3.

It is clear that our method outperforms multitask BC and LfGP-NS, demonstrating the necessity of exploring each of the policies to maximize performance. Intriguingly, even with substantially more main-task expert data, DAC alone is unable to match the performance of our method, and completely fails in the Stacking task. Single-task BC with more main-task data outperforms our method in Stack and Bring, but we emphasize that this baseline has two major limitations compared with ours: none of the collected expert data for this policy is reusable for other tasks, and the policy itself does not have any learned auxiliary policies that could be reused.

The performance results for all multitask methods and all auxiliary tasks are shown in Figure 4. One notable finding is that multitask BC tends to perform quite well on the auxiliary tasks, and, in fact,

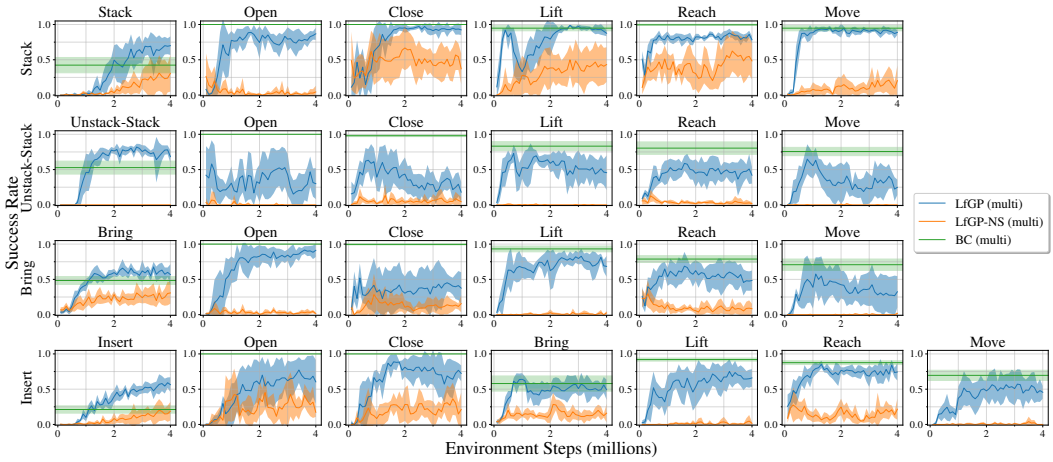

Figure 4: Performance for LfGP and the multitask baselines across all tasks, shaded area corresponds to standard deviation.

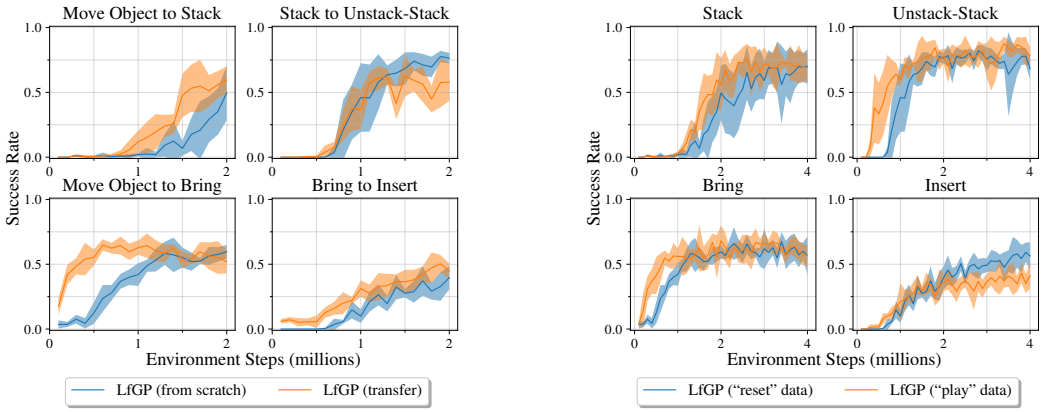

Figure 5: **Left:** The results of our transfer experiments. **Right:** The results of our play-based expert data experiments. Shaded area corresponds to standard deviation.

for all auxiliary tasks, outperforms LfGP. We suspect this is because, compared to the main task, the auxiliary tasks are shorter horizon and simpler than the main task, making them very good candidates for BC. Furthermore, since LfGP maximizes the expected return of the main task, LfGP does not necessarily perform as well on the auxiliary tasks compared to multitask BC, which simultaneously matches the expert for all tasks. A natural followup question is whether we could combine the benefits of quick learning using multitask BC on the simpler auxiliary tasks with the improved performance of LfGP on main tasks—our preliminary results were negative so we leave this as future work.

5.3 Transfer Learning Results

In this section, we show the results of experiments in which we attempt to reuse policies that learned to complete one main task to complete another main task. Specifically, we are interested in reducing the learning time for a main task compared with learning from scratch. To complete these experiments, we chose a relatively simple formulation as a proof-of-concept. For each main task, we take a high-performing existing model that uses a subset of auxiliary tasks required for the new main task and save the parameters and the replay buffer of the existing model. We then train a new model on the new task by loading the saved parameters and replay buffer, adding new model outputs for the new main task. For Stack and Bring, we transfer from a model with $\mathcal{T}_{\text{main}}$ = Move-Object and $\mathcal{T}_{\text{aux}}$ = {Open-Gripper, Close-Gripper, Reach, Lift}. For Unstack-Stack and Insert, we use $\mathcal{T}_{\text{main}}$ = Stack and $\mathcal{T}_{\text{main}}$ = Bring models respectively, with the same $\mathcal{T}_{\text{aux}}$ as in Table 1.

The results of these experiments are shown on the left side of Figure 5. It is clear that in all cases except for Unstack-Stack, transferring trained models from one main task to another is more sample

efficient than from scratch. This result is particularly clear in the case of Bring. For Unstack-Stack, training speed remains roughly consistent, but final performance is actually considerably lower with the transfer model. Given that this experiment was only meant to show that transfer learning would be possible and could provide benefits in some cases, we leave investigating this issue for future work.

5.4 Expert Data Collection Schemes

As described in Section 4.3, there are (at least) two strategies for collecting expert data to be used with LfGP. In this section, we compare the training performance of LfGP between reset-based and play-based expert data.

We show the comparison between the results of training LfGP with reset-based and with play-based data on the right side of Figure 5. For Stack, Unstack-Stack, and Bring, play-based data appears to generally increase the learning speed of LfGP, implying that matching the transition distribution appears to be beneficial for learning, although there is no significant effect on final performance. Conversely, for Insert, play-based data appears to have only a marginal effect on learning speed, while having fairly significant negative impact on final performance. This could be because the Insert task is the least forgiving in terms of the required final state of the object, and reset-based data may contain more transitions between near-insertions and complete insertions than play-based data.

Compared with reset-based data, while play-based data assures that the expert distribution better matches the learning distribution, it also has the downside of making it harder to reuse. In the case of reset-based data, one can easily add a new dataset corresponding to a new task, while keeping existing datasets the same. In play-based, and in our experiments, each individual main task has its own dataset, given that the "transition" initial state distribution should contain data from $\mathcal{T}_{\text{all}}$, which changes depending on $\mathcal{T}_{\text{main}}$.

6 Conclusion

In this work, we identified and addressed an exploration problem in adversarial imitation learning (AIL) algorithms applied to manipulation tasks by introducing Learning from Guided Play (LfGP), a scheduled hierarchical AIL method that enables more effective exploration through the use of "playful" scheduled auxiliary tasks, "guided" by expert data, as opposed to hand-crafted rewards. We demonstrated that an existing state-of-the-art AIL algorithm is substantially outperformed by our method. As well, with significantly less main task data, our algorithm can perform comparably to behaviour cloning (BC) and significantly outperforms multitask BC given the same data. Furthermore, our method can leverage reusable expert data and its learned intentions can be applied to transfer learning. As future work, we intend to further investigate alternative approaches to transfer learning, such as a combination of multitask BC with our method, to see if the complementary advantages of supervised learning and AIL can reduce the number of required environment interactions.

Broader Impact

In this paper, we introduced a scheduled hierarchical approach to adversarial imitation learning (AIL), and demonstrated its successful application to a simulated robotic manipulation domain. Robotic manipulation systems are inherently dangerous, though significant strides have been made in the field of *collaborative* robotics in the past few decades [30], in which lighter arms with speed and force limitations are used in close proximity with humans. The simulated arm from this work, the Franka Emika Panda, is an example of such an arm. Like any other reinforcement learning (RL) method applied to a physical system, ours shares the safety concern that a system that exhibits any degree of random exploration can act in unexpected ways, potentially causing harm through damage to the agent itself, the environment, or to present humans. In our method, this could be exacerbated by the abrupt switching from one task to another; when using this approach on a real platform, care should be taken to constrain the robot's movements through soft and/or hard boundaries. This abruptness could also be reduced by using some kind of soft switching mechanism between hierarchical members [31, 32]. Compared with RL, our method could also learn unexpected behaviours due to the existence of suboptimal expert data, but this is a well-studied issue in the field of IL [33], so there is no shortage of approaches that could be combined with ours to address potentially imperfect demonstrations.

References

[1] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.

[2] Marc Bellemare, Sriram Srinivasan, Georg Ostrovski, Tom Schaul, David Saxton, and Remi Munos. Unifying Count-Based Exploration and Intrinsic Motivation. In *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016.

[3] Ashvin Nair, Bob McGrew, Marcin Andrychowicz, Wojciech Zaremba, and Pieter Abbeel. Overcoming Exploration in Reinforcement Learning with Demonstrations. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6292–6299, 2018.

[4] Andrew Y. Ng and Michael I. Jordan. *Shaping and Policy Search in Reinforcement Learning*. PhD thesis, University of California, Berkeley, 2003.

[5] Jonathan Ho and Stefano Ermon. Generative adversarial imitation learning. *Advances in neural information processing systems*, 29:4565–4573, 2016.

[6] Justin Fu, Katie Luo, and Sergey Levine. Learning robust rewards with adverserial inverse reinforcement learning. In *International Conference on Learning Representations*, 2018.

[7] Ilya Kostrikov, Ofir Nachum, and Jonathan Tompson. Imitation learning via off-policy distribution matching. In *International Conference on Learning Representations*, 2019.

[8] Manu Orsini, Anton Raichuk, Léonard Hussenot, Damien Vincent, Robert Dadashi, Sertan Girgin, Matthieu Geist, Olivier Bachem, Olivier Pietquin, and Marcin Andrychowicz. What matters for adversarial imitation learning? *arXiv preprint arXiv:2106.00672*, 2021.

[9] Rafael Rafailov, Tianhe Yu, Aravind Rajeswaran, and Chelsea Finn. Visual adversarial imitation learning using variational models. In *ICML 2021 Workshop on Unsupervised Reinforcement Learning*, 2021.

[10] Martin Riedmiller, Roland Hafner, Thomas Lampe, Michael Neunert, Jonas Degrave, Tom Wiele, Vlad Mnih, Nicolas Heess, and Jost Tobias Springenberg. Learning by playing solving sparse reward tasks from scratch. In *International Conference on Machine Learning*, pages 4344–4353. PMLR, 2018.

[11] Corey Lynch, Mohi Khansari, Ted Xiao, Vikash Kumar, Jonathan Tompson, Sergey Levine, and Pierre Sermanet. Learning latent plans from play. In *Conference on Robot Learning*, pages 1113–1132. PMLR, 2020.

[12] Dean A. Pomerleau. ALVINN: An Autonomous Land Vehicle in a Neural Network. In D. S. Touretzky, editor, *Proceedings of the Annual Conference on Neural Information Processing Systems (NIPS'89)*, pages 305–313, Denver, CO, USA, 1989. Morgan-Kaufmann.

[13] Andrew Ng and Stuart Russell. Algorithms for inverse reinforcement learning. *Proceedings of the Seventeenth International Conference on Machine Learning*, 0:663–670, 2000.

[14] Pieter Abbeel and Andrew Y Ng. Apprenticeship learning via inverse reinforcement learning. In *Proceedings of the twenty-first international conference on Machine learning*, page 1, 2004.

[15] Stéphane Ross, Geoffrey Gordon, and Drew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 627–635. JMLR Workshop and Conference Proceedings, 2011.

[16] Ilya Kostrikov, Kumar Krishna Agrawal, Sergey Levine, and Jonathan Tompson. Addressing sample inefficiency and reward bias in inverse reinforcement learning. *CoRR*, abs/1809.02925, 2018.

[17] Karol Hausman, Yevgen Chebotar, Stefan Schaal, Gaurav Sukhatme, and Joseph Lim. Multi-modal imitation learning from unstructured demonstrations using generative adversarial nets. In *31st Annual Conference on Neural Information Processing Systems (NIPS 2017)*, pages 1236–1246. Curran Associates, Inc., 2018.

[18] Richard S Sutton, Doina Precup, and Satinder Singh. Between mdps and semi-mdps: A framework for temporal abstraction in reinforcement learning. *Artificial intelligence*, 112(1-2):181–211, 1999.

[19] Tejas D Kulkarni, Karthik Narasimhan, Ardavan Saeedi, and Josh Tenenbaum. Hierarchical deep reinforcement learning: Integrating temporal abstraction and intrinsic motivation. *Advances in neural information processing systems*, 29:3675–3683, 2016.

[20] Ofir Nachum, Shixiang Shane Gu, Honglak Lee, and Sergey Levine. Data-efficient hierarchical reinforcement learning. *Advances in Neural Information Processing Systems*, 31:3303–3313, 2018.

[21] Ofir Nachum, Haoran Tang, Xingyu Lu, Shixiang Gu, Honglak Lee, and Sergey Levine. Why does hierarchy (sometimes) work so well in reinforcement learning? *arXiv preprint arXiv:1909.10618*, 2019.

[22] Peter Henderson, Wei-Di Chang, Pierre-Luc Bacon, David Meger, Joelle Pineau, and Doina Precup. Optiongan: Learning joint reward-policy options using generative adversarial inverse reinforcement learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32, 2018.

[23] Mohit Sharma, Arjun Sharma, Nicholas Rhinehart, and Kris M Kitani. Directed-info gail: Learning hierarchical policies from unsegmented demonstrations using directed information. In *International Conference on Learning Representations*, 2018.

[24] Mingxuan Jing, Wenbing Huang, Fuchun Sun, Xiaojian Ma, Tao Kong, Chuang Gan, and Lei Li. Adversarial option-aware hierarchical imitation learning. *arXiv preprint arXiv:2106.05530*, 2021.

[25] Felipe Codevilla, Matthias Müller, Antonio López, Vladlen Koltun, and Alexey Dosovitskiy. End-to-end driving via conditional imitation learning. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4693–4700. IEEE, 2018.

[26] Tim Hertweck, Martin Riedmiller, Michael Bloesch, Jost Tobias Springenberg, Noah Siegel, Markus Wulfmeier, Roland Hafner, and Nicolas Heess. Simple sensor intentions for exploration. *arXiv preprint arXiv:2005.07541*, 2020.

[27] Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *ICML*, 2018.

[28] Serkan Cabi, Sergio Gómez Colmenarejo, Matthew W Hoffman, Misha Denil, Ziyu Wang, and Nando Freitas. The intentional unintentional agent: Learning to solve many continuous control tasks simultaneously. In *Conference on Robot Learning*, pages 207–216. PMLR, 2017.

[29] Max Jaderberg, Volodymyr Mnih, Wojciech Marian Czarnecki, Tom Schaul, Joel Z Leibo, David Silver, and Koray Kavukcuoglu. Reinforcement learning with unsupervised auxiliary tasks. *arXiv preprint arXiv:1611.05397*, 2016.

[30] J. Edward Colgate, J. Edward, Michael A. Peshkin, and Witaya Wannasuphoprasit. Cobots: Robots For Collaboration With Human Operators, 1996.

[31] Thomas Kipf, Yujia Li, Hanjun Dai, Vinicius Zambaldi, Alvaro Sanchez-Gonzalez, Edward Grefenstette, Pushmeet Kohli, and Peter Battaglia. Compile: Compositional imitation learning and execution. In *International Conference on Machine Learning (ICML)*, 2019.

[32] Youngwoon Lee, Shao-Hua Sun, Sriram Somasundaram, Edward S. Hu, and Joseph J. Lim. Composing complex skills by learning transition policies. In *Proceedings of International Conference on Learning Representations*, 2019.

[33] Saurabh Arora and Prashant Doshi. A Survey of Inverse Reinforcement Learning: Challenges, Methods and Progress. *arXiv:1806.06877 [cs, stat]*, June 2018.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes] We specifically state our contributions in Section 1.
   (b) Did you describe the limitations of your work? [Yes] In Section 5, we describe several specific ways in which our method falls short of our baselines, or when things we tried didn't work out.
   (c) Did you discuss any potential negative societal impacts of your work? [Yes] See our broader impact statement following our conclusion.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [N/A]
   (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] See our supplementary material. We had to limit how much we could include due to size restrictions, so we only included a representative set of the data.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] Short details are in Section 5.1, with expanded details included in the supplementary material.
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] We showed standard deviation across five seeds in all of our plots.
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] See our supplementary material.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [N/A] We used our own environment and data for this work, though we do cite the creator of the underlying simulator (PyBullet).
   (b) Did you mention the license of the assets? [No] Since PyBullet is well-known in the community as being free and open source, we do not mention that it has a permissive zlib license.
   (c) Did you include any new assets either in the supplemental material or as a URL? [Yes] We include code for producing our environment, as well as the data and expert models for a single task. Due to space restrictions, we could not include expert data/models for all four tasks.
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]