

Learning to Interactively Learn and Assist

Mark Woodward,* Chelsea Finn, Karol Hausman

Google Brain, Mountain View

{markwoodward, chelseaf, karolhausman}@google.com

Abstract

When deploying autonomous agents in the real world, we need effective ways of communicating objectives to them. Traditional skill learning has revolved around reinforcement and imitation learning, each with rigid constraints on the format of information exchanged between the human and the agent. While scalar rewards carry little information, demonstrations require significant effort to provide and may carry more information than is necessary. Furthermore, rewards and demonstrations are often defined and collected before training begins, when the human is most uncertain about what information would help the agent. In contrast, when humans communicate objectives with each other, they make use of a large vocabulary of informative behaviors, including non-verbal communication, and often communicate throughout learning, responding to observed behavior. In this way, humans communicate intent with minimal effort. In this paper, we propose such *interactive learning* as an alternative to reward or demonstration-driven learning. To accomplish this, we introduce a multi-agent training framework that enables an agent to learn from another agent who knows the current task. Through a series of experiments, we demonstrate the emergence of a variety of interactive learning behaviors, including information-sharing, information-seeking, and question-answering. Most importantly, we find that our approach produces an agent that is capable of learning interactively from a human user, without a set of explicit demonstrations or a reward function, and achieving significantly better performance cooperatively with a human than a human performing the task alone.

1 Introduction

Many tasks that we would like our agents to perform, such as unloading a dishwasher, straightening a room, or restocking shelves are inherently user-specific, requiring information from the user in order to fully learn all the intricacies of the task. The traditional paradigm for agents to learn such tasks is through rewards and demonstrations. However, iterative reward engineering with untrained human users is impractical in real-world settings, while demonstrations are often burdensome to provide. In contrast, humans learn from a variety of interactive communicative behaviors, including non-verbal gestures and partial demonstrations, each with their own information capacity and effort. Can we enable agents to learn tasks from humans through such unstructured interaction, requiring minimal effort from the human user?

*Work done as a part of the Google AI Residency
Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

The effort required by the human user is affected by many aspects of the learning problem, including restrictions on when the agent is allowed to act and restrictions on the behavior space of either human or agent, such as limiting the user feedback to rewards or demonstrations. We consider a setting where both the user and the agent are allowed to act throughout learning, which we refer to as *interactive learning*. Unlike collecting a set of demonstrations before training, interactive learning allows the user to selectively act only when it deems the information is necessary and useful, reducing the user’s effort. Examples of such interactions include allowing user interventions or agent requests, for demonstrations (Kelly et al. 2018), rewards (Warnell et al. 2018; Arumugam et al. 2019), or preferences (Christiano et al. 2017). While these methods allow the user to provide feedback throughout learning, the communication interface is restricted to structured forms of supervision, which may be inefficient for a given situation. For example, in a dishwasher unloading task, given the history of learning, it may be sufficient to point at the correct drawer rather than provide a full demonstration.

To this end, we propose to allow the agent and the user to exchange information through an unstructured interface. To do so, the agent and the user need a common prior understanding of the meaning of different unstructured interactions, along with the context of the space of tasks that the user cares about. Indeed, when humans communicate tasks to each other, they come in with rich prior knowledge and common sense about what the other person may want and how they may communicate that, enabling them to communicate concepts effectively and efficiently (Peloquin, Goodman, and Frank 2019).

In this paper, we propose to allow the agent to acquire this prior knowledge through joint pre-training with another agent who knows the task and serves as a human surrogate. The agents are jointly trained on a variety of tasks, where actions and observations are restricted to the physical environment. Since the first agent is available to assist, but only the second agent is aware of the task, interactive learning behaviors should emerge to accomplish the task efficiently. We hypothesize that, by restricting the action and observation spaces to the physical environment, the emerged behaviors can transfer to learning from a human user. An added benefit of our framework is that, by training on a variety of tasks from the target task domain, much of the non-user specific task prior knowledge is pre-trained into the agent, further reducing the effort required by the user.

We evaluate various aspects of agents trained with our




Figure 1: Episode traces after 100, 1k, and 40k pre-training steps for the cooperative fruit collection domain of Experiment 4. The principal agent “P” (pink) is told the fruit to be collected, lemons or plums, in its observations. Within an episode, the assistant agent “A” (blue) must infer the fruit to be collected from observations of the principal. Each agent observes an overhead image of itself and its nearby surroundings. By the end of training (c) the assistant is inferring the correct fruit and the agents are coordinating. This inference and coordination transfers to human principals (d). An interactive game and videos for all experiments are available at: <https://interactive-learning.github.io>

framework on several simulated object gathering task domains, including a domain with pixel observations, shown in Figure 1. We show that our trained agents exhibit emergent information-gathering behaviors in general and explicit question-asking behavior where appropriate. Further, we conduct a user study with trained agents, where the users score significantly higher with the agent than without the agent, which demonstrates that our approach can produce agents that can learn from and assist human users.

The key contribution of our work is a training framework that allows agents to quickly learn new tasks from humans through unstructured interactions, without an explicitly-provided reward function or demonstrations. Critically, our experiments demonstrate that agents trained with our framework generalize to learning test tasks from human users, demonstrating interactive learning with a human in the loop. In addition, we introduce a novel multi-agent model architecture for cooperative multi-agent training that exhibits improved training characteristics. Finally, our experiments on a series of object-gathering task domains illustrate a variety of emergent interactive learning behaviors and demonstrate that our method can scale to raw pixel observations.

2 Related Work

The traditional means of passing task information to an agent include specifying a reward function (Barto and Sutton 1998) that can be hand-crafted for the task (Singh, Lewis, and Barto 2009; Levine et al. 2016; Chebotar et al. 2017) and providing demonstrations (Schaal 1999; Abbeel and Ng 2004) before the agent starts training. More recent works explore the concept of the human supervision being provided throughout training by either providing rewards during training (Isbell et al. 2001; Thomaz et al. 2005; Warnell et al. 2018; Perez-Dattari et al. 2018) or demonstrations during training; either continuously (Ross, Gordon, and Bagnell 2011b; Kelly et al. 2018) or at the agent’s discretion (Ross, Gordon, and Bagnell 2011a; Borsa et al. 2017; Xu et al. 2018; Hester et al. 2018; James, Bloesch, and Davison 2018; Yu et al. 2018a; Krening 2018; Brown, Cui, and Niekum 2018). In all of these cases, however, the reward and demonstrations are the sole means of interaction.

Another recent line of research involves the human ex-

pressing their preference between agent generated trajectories (Christiano et al. 2017; Mindermann et al. 2018; Ibarz et al. 2018). Here again, the interaction is restricted to a single modality.

Our work builds upon the idea of meta-learning, or learning-to-learn (Schmidhuber 1987; Bengio, Bengio, and Cloutier 1991; Thrun and Pratt 2012). Meta-learning for control has been considered in the context of reinforcement learning (Duan et al. 2016; Wang et al. 2016; Finn, Abbeel, and Levine 2017) and imitation learning (Duan et al. 2017; Yu et al. 2018b). Our problem setting differs from these, as the agent is learning by observing and interacting with another agent, as opposed to using reinforcement or imitation learning. In particular, our method builds upon recurrence-based meta-learning approaches (Santoro et al. 2016; Duan et al. 2016; Wang et al. 2016) in the context of the multi-agent task setting.

When a broader range of interactive behaviors is desired, prior works have introduced a multi-agent learning component (Potter and Jong 1994; Palmer et al. 2018). The following methods are closely related to ours in that, during training, they also maximize a joint reward function between the agents and emerge cooperative behavior (Gupta, Egorov, and Kochenderfer 2017; Foerster et al. 2018; 2016; Lazaridou, Peysakhovich, and Baroni 2016; Andreas, Dragan, and Klein 2017). Multiple works (Gupta, Egorov, and Kochenderfer 2017; Foerster et al. 2018) emerge cooperative behavior but in task domains that do not require knowledge transfer between the agents, while others (Foerster et al. 2016; Lazaridou, Peysakhovich, and Baroni 2016; Lowe et al. 2017; Andreas, Dragan, and Klein 2017; Mordatch and Abbeel 2018) all emerge communication over a communication channel. Such communication is known to be difficult to interpret (Lazaridou, Peysakhovich, and Baroni 2016), without post-inspection (Mordatch and Abbeel 2018) or a method for translation (Andreas, Dragan, and Klein 2017). Critically, none of these prior works conduct user experiments to evaluate transfer to humans.

Mordatch and Abbeel (2018) experiment with tasks similar to ours, in which information must be communicated between the agents, and communication is restricted to the physical environment. This work demonstrates the emer-

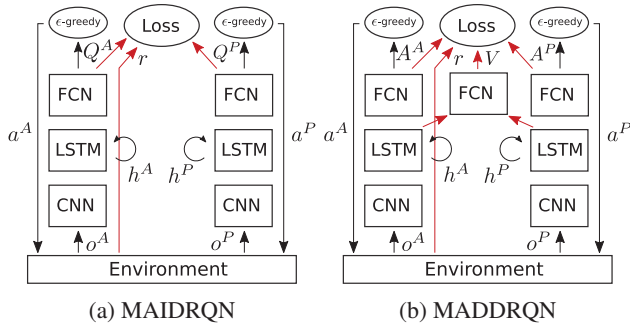


Figure 2: Information flow for the two models used in our experiments; red paths are only needed during training. The MADDRQN model (b) uses a centralized value-function with per-agent advantage functions. The centralized value function is only used during training. Superscripts A and P refer to the assistant and principal agents respectively. The MAIDRQN model (a) is used in experiments 1-3 and the MADDRQN model (b) is used in experiment 4 where it exhibits superior training characteristics for learning from pixels.

gence of pointing, demonstrations, and pushing behavior. Unlike this prior approach, however, our algorithm does not require a differentiable environment. We also demonstrate our method with pixel observations and conduct a user experiment to evaluate transfer to humans.

Laird et al. (2017) describe desiderata for interactive learning systems. Our method primarily addresses the desiderata of efficient interaction and accessible interaction.

3 Preliminaries

In this section, we review the cooperative partially observable Markov game (Littman 1994), which serves as the foundation for tasks in Section 4. A cooperative partially observable Markov game is defined by the tuple $\langle \mathcal{S}, \{\mathcal{A}^i\}, T, \mathcal{R}, \{\Omega^i\}, \{O^i\}, \gamma, H \rangle$, where $i \in \{1..N\}$ indexes the agent among N agents, \mathcal{S} , \mathcal{A}^i , and Ω^i are state, action, and observation spaces, $T : \mathcal{S} \times \{\mathcal{A}^i\} \times \mathcal{S}' \rightarrow \mathbb{R}$ is the transition function, $\mathcal{R} : \mathcal{S} \times \{\mathcal{A}^i\} \rightarrow \mathbb{R}$ is the reward function, $O^i : \mathcal{S} \times \Omega^i \rightarrow \mathbb{R}$ are the observation functions, γ is the discount factor, and H is the horizon.

The functions T , \mathcal{R} , and O^i are not accessible to the agents. At time t , the environment accepts actions $\{a_t^i\} \in \{\mathcal{A}^i\}$, samples $s_{t+1} \sim T(s_t, \{a_t^i\})$, and returns reward $r_t \sim \mathcal{R}(s_t, \{a_t^i\})$ and observations $\{o_{t+1}^i\} \sim \{O^i(s_{t+1})\}$. The objective of the game is to choose actions to maximize the expected discounted sum of future rewards:

$$\operatorname{argmax}_{\{a_{t_0}^i | o_{t_0}^i\}} \mathbb{E}_{s, o^i, r} \left[\sum_{t=t_0}^H \gamma^{t-t_0} r_t \right]. \quad (1)$$

Note that, while the action and observation spaces vary for the agents, they share a common reward which leads to a cooperative task.

4 The LILA Training Framework

We now describe our training framework for producing an assisting agent that can learn a task interactively from a human user. We define a task to be an instance of a cooperative partially observable Markov game as described in Section 3, with $N = 2$. To enable the agent to solve such tasks, we train the agent, whom we call the ‘‘assistant’’ (superscript A), jointly with another agent, whom we call the ‘‘principal’’ (superscript P) on a variety of tasks. Critically, the principal’s observation function informs it of the task.¹ The principal agent acts as a human surrogate which allows us to replace it with a human once the training is finished. By informing the principal of the current task and withholding rewards and gradient updates until the end of each task, the agents are encouraged to emerge interactive learning behaviors in order to inform the assistant of the task and allow them to contribute to the joint reward. We limit actions and observations to the physical environment, with the hope of emerging human-compatible behaviors.

In order to train the agents, we consider two different models. We first introduce a simple model that we find works well in tabular environments. Then, in order to scale our approach to pixel observations, we introduce a modification to the first model that we found was important in increasing the stability of learning.

Multi-Agent Independent DRQN (MAIDRQN):

The first model uses two deep recurrent Q -networks (DRQN) (Hausknecht and Stone 2015) that are each trained with Q -learning (Watkins 1989). Let $Q_{\theta^i}(o_t^i, a_t^i, h_t^i)$ be the action-value function for agent i , which maps from the current action, observation, and history, h_t^i , to the expected discounted sum of future rewards. The MAIDRQN method optimizes the following loss:²

$$\mathcal{L}^{\text{MAIDRQN}} := \frac{1}{N} \sum_{i,t} [y_t - Q_{\theta^i}(o_t^i, a_t^i, h_t^i)]^2 \quad (2)$$

$$y_t := r_t + \gamma \max_{a_{t+1}^i} Q_{\theta^i}(o_{t+1}^i, a_{t+1}^i, h_{t+1}^i)$$

The networks are trained simultaneously. The model architecture is a recurrent neural network, depicted in Figure 2a. We use this model for experiments 1-3.

Multi-Agent Dueling DRQN (MADDRQN): With independent Q -Learning, as in MAIDRQN, the other agent’s changing behavior and unknown actions make it difficult to estimate the Bellman target y_t in Equation 2, which leads to instability in training. This model addresses part of the instability that is caused by unknown actions.

If $Q^*(o, a, h)$ is the optimal action-value function, then the optimal value function is $V^*(o, h) = \operatorname{argmax}_a Q^*(o, a, h)$, and the optimal advantage function is defined as $A^*(o, a, h) = Q^*(o, a, h) - V^*(o, h)$ (Wang et

¹Our tasks are similar to tasks in Hadfield-Menell et al. (2016), but with partially observable state and without access to the other agent’s actions, which should better generalize to learning from humans in natural environment.

²In our experiments we do not use a lagged ‘‘target’’ Q -network (Mnih et al. 2013), but we do stop gradients through the Q -network in y_t .

Table 1: Experimental configurations for our 4 experiments. Experiment 1 has two sub experiments, 1A and 1B. In 1B, the agents incur a penalty whenever the principals moves. The Observation Window column lists the radius of cells visible to each agent.

EXP.	MODEL	PRINCIPAL MOTION PENALTY	GRID SHAPE	NUM. OBJECTS	OBSERVATIONS	OBSERVATION WINDOW
1A	MAIDRQN	0.0	5x5	10	BINARY VECTORS	FULL
1B	MAIDRQN	-0.4	5x5	10	BINARY VECTORS	FULL
2	MAIDRQN	0.0	5x5	10	BINARY VECTORS	1-CELL
3	MAIDRQN	-0.1	3x1 "L"	1	BINARY VECTORS	1-CELL
4	MADDRQN	0.0	5x5	10	$64 \times 64 \times 3$ PIXELS	~ 2 -CELLS

al. 2015). The advantage function captures how inferior an action is to the optimal action in terms of the expected sum of discounted future rewards. This allows us to express Q in a new form, $Q^*(o, a, h) = V^*(o, h) + A^*(o, a, h)$. We note that the value function is not needed when selecting actions: $\operatorname{argmax}_a Q^*(o, a, h) = \operatorname{argmax}_a (V^*(o, h) + A^*(o, a, h)) = \operatorname{argmax}_a A^*(o, a, h)$. We leverage this idea by making the following approximation to an optimal, centralized action-value function for multiple agents:

$$Q^*({o^i, a^i, h^i}) = V^*({o^i, h^i}) + A^*({o^i, a^i, h^i}) \quad (3)$$

$$\approx V^*({o^i, h^i}) + \sum_i A^{i*}({o^i, a^i, h^i}),$$

where $A^{i*}({o^i, a^i, h^i})$ is an advantage function for agent i and $V^*({o^i, h^i})$ is a joint value function.³

The training loss for this model is:

$$\mathcal{L}^{\text{MADDRQN}} := \sum_t [y_t - Q_{\{\theta^i, \phi\}}({o_t^i, a_t^i, h_t^i})]^2 \quad (4)$$

$$y_t := r_t + \gamma \max_{\{a_{t+1}^i\}} Q_{\{\theta^i, \phi\}}({o_{t+1}^i, a_{t+1}^i, h_{t+1}^i})$$

where

$$Q_{\{\theta^i, \phi\}}({o_t^i, a_t^i, h_t^i}) := V_\phi({o_t^i, h_t^i}) + \sum_i A_{\theta^i}({o_t^i, a_t^i, h_t^i}). \quad (5)$$

Once trained, each agent selects their actions according to their advantage function A_{θ^i} ,

$$a_t^i = \operatorname{argmax}_a A_{\theta^i}({o_t^i, a, h_t^i}), \quad (6)$$

as opposed to the Q-functions Q_{θ^i} in the case of MAIDRQN.

In the loss for the MAIDRQN model, Equation 2, there is a squared error term for each Q_{θ^i} which depends on the joint reward r . This means that, in addition to estimating the immediate reward due their own actions, each Q_{θ^i} must

³The approximation is due to the substitution of $\sum_i A^{i*}({o^i, a^i, h^i})$ for $A^*({o^i, a^i, h^i})$ in Equation 3, which implies that the agents' current actions have independent effects on expected future rewards, and is not true in general. Nevertheless, it is a useful approximation.

estimate the immediate reward due to the actions of the other agent, without access to their actions or observations. By using a joint action value function and decomposing it into advantage functions and a value function, each A^i can ignore the immediate reward due to the other agent, simplifying the optimization.

We refer to this model as a multi-agent dueling deep recurrent Q-network (MADDRQN), in reference to the single agent dueling network of Wang et al. (2015). The MAIDRQN model, which adds a fully connected network for the shared value function, is depicted in Figure 2b; The MAIDRQN model is used in experiment 4.

Training Procedure: We use a standard episodic training procedure, with the task changing on each episode. The training procedure for the MADDRQN and MAIDRQN models differ only in the loss function. Here, we describe the training procedure with reference to the MADDRQN model. We assume access to a subset of tasks, $\mathcal{D}^{\text{Train}}$, from a task domain, $\mathcal{D} = \{\dots, \mathcal{T}_j, \dots\}$. First, we initialize the parameters θ^P , θ^A , and ϕ . Then, the following procedure is repeated until convergence. A batch of tasks are uniformly sampled from $\mathcal{D}^{\text{Train}}$. For each task \mathcal{T}_b in the batch, a trajectory, $\tau_b = (o_0^P, o_0^A, a_0^P, a_0^A, r_0, \dots, o_H^P, o_H^A, a_H^P, a_H^A, r_H)$, is collected by playing out an episode in an environment configured to \mathcal{T}_b , with actions chosen ϵ -greedy according to A_{θ^P} and A_{θ^A} . The hidden states for the recurrent LSTM cells are reset to $\mathbf{0}$ at the start of each episode. The loss for each trajectory is calculated using Equations 4 and 5. Finally, a gradient step is taken with respect to θ^P , θ^A , and ϕ on the sum of the episode losses.

5 Experimental Results

We designed a series of experiments in order to study how different interactive learning behaviors may emerge, to test whether our method can scale to pixel observations, and to evaluate the ability for the agents to transfer to a setting with a human user.

We conducted four experiments on grid-world environments, where the goal was to cooperatively collect all objects from one of two object classes. Two agents, the principal and the assistant, act simultaneously and may move in one of the four cardinal directions or may choose not to move, giving five possible actions per agent.

Table 2: Results for Experiments 1A and 1B. Experiment 1B includes a motion penalty for the principal’s motion. In both experiments, MAIDRQN outperforms the principal acting alone, demonstrating that the assistant learns from and assists the principal. All performance increases are significant (confidence > 99%), except for FeedFwd-A and Solo-P in Experiment 1A, which are statistically equivalent.

METHOD NAME	EXPERIMENT 1A			EXPERIMENT 1B		
	JOINT REWARD	REWARD DUE TO P	REWARD DUE TO A	JOINT REWARD	REWARD DUE TO P	REWARD DUE TO A
ORACLE-A	4.9 ± 0.2	2.5 ± 0.1	2.4 ± 0.1	4.0 ± 0.1	0.0 ± 0.0	4.0 ± 0.1
MAIDRQN	4.6 ± 0.2	3.3 ± 0.2	1.3 ± 0.3	3.6 ± 0.1	0.4 ± 0.1	3.2 ± 0.1
FEEDFWD-A	4.1 ± 0.1	4.1 ± 0.1	0.0 ± 0.0^4	2.0 ± 0.4	0.7 ± 0.3	1.3 ± 0.6
SOLO-P	4.0 ± 0.1	4.0 ± 0.1	N/A	1.2 ± 0.1	1.2 ± 0.1	N/A





Figure 3: Training curve for Experiment 1B. Error bars are 1 standard deviation. At the end of training, nearly all of the joint reward in an episode is due to the assistant’s actions, indicating that the trained assistant can learn the task and then complete it independently.

Within an experiment, tasks vary by the placement of objects, and by the class of objects to be collected, which we call the “target class”. The target class is supplied to the principal as a two dimensional, one-hot vector. If either agent enters a cell containing an object, the object disappears and both agents receive a reward: +1 for objects of the target class and -1 otherwise. Each episode consisted of a single task and lasted for 10 time-steps. Table 1 gives the setup for each experiment.

We collected 10 training runs per experiment, and we report the aggregated performance of the 10 trained agent pairs on 100 test tasks not seen during training. The training batch size was 100 episodes and the models were trained for 150,000 gradient steps (Experiments 1-3) or 40,000 gradient steps (Experiment 4). Videos for all experiments, as well as an interactive game, are available on the paper website.⁵

⁴The FeedForward assistant moves 80% of the time, but it never collects an object.

⁵<https://interactive-learning.github.io>



(a) The assistant learns from a single principal movement. (b) The assistant learns from a lack of principal movement.

Figure 4: Episode traces of trained agents on test tasks from Experiment 1B. Both agents begin in the center square and cooperatively collect as many instances of the target shape as possible. The target shape is shown in green. The principal agent P observes the target shape, but the assistant agent A does not and must learn from the principal’s movement or lack of movement. The assistant rapidly learns the target shape from the principal and collects all instances.

Experiment 1 A&B – Learning and Assisting: In this experiment we explore if the assistant can be trained to learn and assist the principal. Table 2 shows the experimental results without and with a penalty for motion of the principal (Experiments 1A and 1B respectively). Figures 3 and 4 show the learning curve and trajectory traces for trained agents in Experiment 1B.

The joint reward of our approach (MAIDRQN) exceeds that of a principal trained to act alone (Solo-P), and approaches the optimal setting where the assistant also observes the target class (Oracle-A). Further, we see that the reward due to the assistant is positive, and even exceeds the reward due to the principal when the motion penalty is present (Experiment 1B). This demonstrates that the assistant learns the task from the principal and assists the principal. Our approach also outperforms an ablation in which the assistant’s LSTM is replaced with a feed forward network (FeedFwd-A), highlighting the importance of memory.

Experiment 2 – Active Information Gathering: In this




Figure 5: Visualization of the 1-cell observation window used in experiments 2 and 3. Cell contents outside of an agent’s window are hidden from that agent.




Figure 6: Episode traces of trained agents on test tasks from Experiment 2. Both agents begin in the center square and cooperatively collect as many instances of the target shape as possible. The target shape is shown in green. The principal agent P observes the target shape, but the assistant agent A does not and must learn from the principal’s movement. With restricted observations, the assistant moves with the principal until it observes a disambiguating action, and then proceeds to collect the target shape on its own.

experiment we explore if, in the presence of additional partial observability, the assistant will take actions to actively seek out information. This experiment restricts the view of each agent to a 1-cell window and only places objects around the exterior of the grid, requiring the assistant to move with the principal and observe its behavior, see Figure 5. Figure 6 shows trajectory traces for two test tasks. The average joint reward, reward due to the principal, and reward due to the assistant are 4.7 ± 0.2 , 2.8 ± 0.2 , and 1.9 ± 0.1 respectively. This shows that our training framework can produce information seeking behaviors.

Experiment 3 – Interactive Questioning and Answering: In this experiment we explore if there is a setting where explicit questioning and answering can emerge. On 50% of the tasks, the assistant is allowed to observe the target class. This adds uncertainty for the principal, and discourages it from proactively informing the assistant. Figure 7 shows the first several states of tasks in which the assistant does not

Table 3: Results for Experiment 4. Trained assistants learned from human principals and significantly increased their scores (Human&Agent) over the humans acting alone (Human), demonstrating the potential for our training framework to produce agents that can learn from and assist humans.⁷

PLAYERS	JOINT REWARD	REWARD DUE TO P	REWARD DUE TO A
AGENT&AGENT	4.6 ± 0.2	2.6 ± 0.2	2.0 ± 0.2
HUMAN&AGENT	4.2 ± 0.4	2.9 ± 0.3	1.3 ± 0.5
AGENT	3.9 ± 0.1	3.9 ± 0.1	N/A
HUMAN	3.8 ± 0.3	3.8 ± 0.3	N/A

observe the target class.⁶

The emerged behavior is for the assistant to move into the visual field of the principal, effectively asking the question, then the principal moves until it sees the object, and finally answers the question by moving one step closer only if the object should be collected. The average joint reward, reward due to the principal, and reward due to the assistant are 0.4 ± 0.1 , -0.1 ± 0.1 , and 0.5 ± 0.1 respectively. This demonstrates that our framework can emerge question-answering, interactive behaviors.

Experiment 4 – Learning from and Assisting a Human Principal with Pixel Observations: In this final experiment we explore if our training framework can extend to pixel observations and whether the trained assistant can learn from a human principal. Figure 8 shows examples of the pixel observations. Ten participants, who were not familiar with this research, were paired with the 10 trained assistants, and each played 20 games with the assistant and 20 games without the assistant. Participants were randomly assigned which setting to play first. Figure 1 shows trajectory traces on test tasks at several points during training and with a human principal after training.

Unlike the previous experiments, stability was a challenge in this problem setting; most training runs of MAIDRQN became unstable and dropped below 0.1 joint reward before the end of training. Hence, we chose to use the MADDRQN model because we found it to be more stable than MAIDRQN. The failure rate was 64% vs 75% for each method respectively, and the mean failure time was 5.6 hours vs 9.7 hours (confidence > 99%), which saved training time and was a practical benefit.

Table 3 shows the experimental results. The participants scored significantly higher with the assistant than without (confidence > 99%). This demonstrates that our framework can produce agents that can learn from humans.

While inclusion of the assistant increases the human’s score, it is still less than the score when the assistant acts with the principal agent with which it was trained. What is

⁶The test and training sets are the same in Experiment 3, since there are only 8 possible tasks

⁷Significance is based on a t-test of the participants’ change in score, which is more significant than the table’s standard deviations would suggest (confidence > 99%).

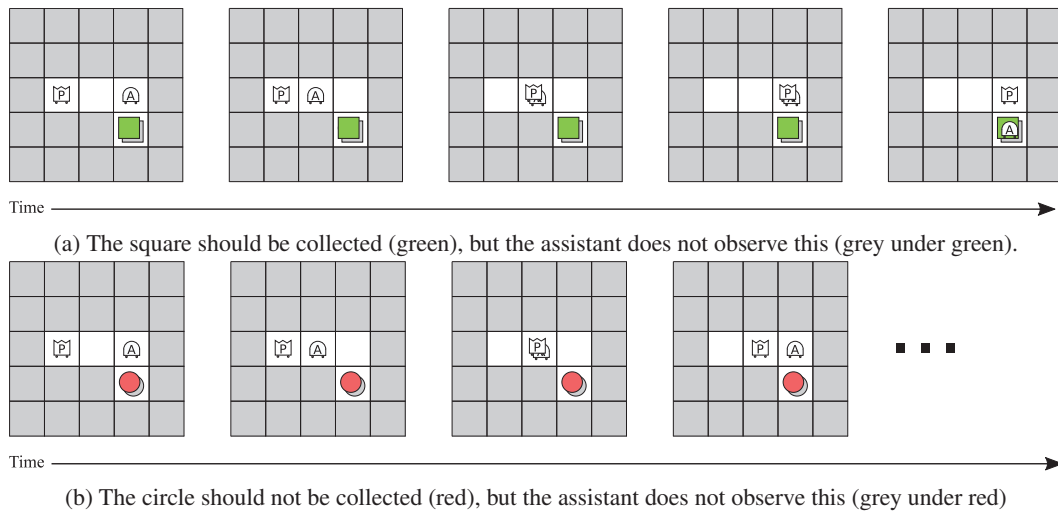


Figure 7: Episode roll-outs for trained agents from Experiment 3. When the assistant is uncertain of an object, it requests information from the principal by moving into its visual field and observing the response.




Figure 8: Example observations for experiment 4. The principal’s observation also includes a 2 dimensional one-hot vector indicating the fruit to collect, plums in this case. These are the 7th observations from the human-agent trajectory in Figure 1d.

the cause of this gap? To answer this question, we identified that 12% of the time the assistant incorrectly infers which object to collect (episodes where the assistant always collects the wrong object). If we exclude these episodes, we obtain the performance when the assistant has correctly inferred the task, but must still coordinate with the human principal. Humans with “correct” assistants achieve rewards (4.7 ± -0.3 , 2.8 ± 0.3 , 1.9 ± 0.2) that are statistically equivalent to Agent&Agent and statistically superior to Human&Agent in Table 3. This means that the assistants coordinate equivalently with human principals and artificial principals, but can experience problems inferring the task from human principals, resulting in the observed drop in score. This is an example of co-adaptation to communicating with the principal agent during training. The next section suggests an approach to address such co-adaptation.

6 Summary and Future Work

We introduced the LILA training framework, which trains an assistant to learn interactively from a knowledgeable principal through only physical actions and observations in the environment. LILA produces the assistant by jointly training it with a principal, who is made aware of the task through its observations, on a variety of tasks, and restricting the observation and action spaces to the physical environment. We further introduced the MADDRQN algorithm, in which the agents have individual advantage functions but share a value function during training. MADDRQN showed improved stability over MAIDRQN, which was a practical benefit in the experiments. The experiments demonstrate that, depending on the environment, LILA emerges behaviors such as demonstrations, partial demonstrations, information seeking, and question answering. Experiment 4 demonstrated that LILA scales to environments with pixel observations, and, crucially, that LILA is able to produce agents that can learn from and assist humans.

A possible future extension involves training with populations of agents. In our experiments, the agents sometimes emerged overly co-adapted behaviors. For example, in Experiment 2, the agents tend to always move in the same direction in the first time step, but the direction varies by the training run. This makes agents paired across runs less compatible, and less likely to generalize to human principals. We believe that training an assistant across populations of agents will reduce such co-adapted behaviors. Finally, LILA’s emergence of behaviors means that the trained assistant can only learn from behaviors that emerged during training. Further research should seek to minimize these limitations, perhaps through advances in online meta-learning (Finn et al. 2019).

Acknowledgments The authors are especially grateful to Alonso Martinez for designing and iterating on the Unity environment used in Experiment 4.

References

- Abbeel, P., and Ng, A. Y. 2004. Apprenticeship learning via inverse reinforcement learning. In *Proceedings of the twenty-first international conference on Machine learning*, 1. ACM.
- Andreas, J.; Dragan, A.; and Klein, D. 2017. Translating neuralese. *arXiv preprint arXiv:1704.06960*.
- Arumugam, D.; Lee, J. K.; Saskin, S.; and Littman, M. L. 2019. Deep reinforcement learning from policy-dependent human feedback. *arXiv preprint arXiv:1902.04257*.
- Barto, A., and Sutton, R. S. 1998. *Reinforcement Learning: An Introduction*. MIT Press.
- Bengio, Y.; Bengio, S.; and Cloutier, J. 1991. Learning a synaptic learning rule. In *Proc. of the International Joint Conference on Neural Networks (IJCNN)*.
- Borsa, D.; Piot, B.; Munos, R.; and Pietquin, O. 2017. Observational learning by reinforcement learning.
- Brown, D. S.; Cui, Y.; and Niekum, S. 2018. Risk-aware active inverse reinforcement learning. In *Proc. of Conference on Robot Learning (CoRL)*.
- Chebotar, Y.; Hausman, K.; Zhang, M.; Sukhatme, G.; Schaal, S.; and Levine, S. 2017. Combining model-based and model-free updates for trajectory-centric reinforcement learning. In *International Conference on Machine Learning (ICML)*.
- Christiano, P.; Leike, J.; Brown, T. B.; Martic, M.; Legg, S.; and Amodei, D. 2017. Deep reinforcement learning from human preferences. In *Proc. of the Conference on Neural Information Processing Systems (NIPS)*.
- Duan, Y.; Schulman, J.; Chen, X.; Bartlett, P. L.; Sutskever, I.; and Abbeel, P. 2016. RL^2 : Fast reinforcement learning via slow reinforcement learning.
- Duan, Y.; Andrychowicz, M.; Stadie, B.; Ho, O. J.; Schneider, J.; Sutskever, I.; Abbeel, P.; and Zaremba, W. 2017. One-shot imitation learning. In *Advances in neural information processing systems*, 1087–1098.
- Finn, C.; Abbeel, P.; and Levine, S. 2017. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proc. of the International Conference on Machine Learning (ICML)*.
- Finn, C.; Rajeswaran, A.; Kakade, S.; and Levine, S. 2019. Online meta-learning. *arXiv preprint arXiv:1902.08438*.
- Foerster, J. N.; Assael, Y. M.; de Freitas, N.; and Whiteson, S. 2016. Learning to communicate with deep multi-agent reinforcement learning. In *Proc. of the Conference on Neural Information Processing Systems (NIPS)*.
- Foerster, J. N.; Farquhar, G.; Afouras, T.; Nardelli, N.; and Whiteson, S. 2018. Counterfactual multi-agent policy gradients. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Gupta, J. K.; Egorov, M.; and Kochenderfer, M. 2017. Cooperative multi-agent control using deep reinforcement learning. In *Proc. of Autonomous Agents and Multiagent Systems (AAMAS)*.
- Hadfield-Menell, D.; Dragan, A.; Abbeel, P.; and Russell, S. 2016. Cooperative inverse reinforcement learning. In *Proc. of the Conference on Neural Information Processing Systems (NIPS)*.
- Hausknecht, M., and Stone, P. 2015. Deep recurrent q-learning for partially observable mdps. In *AAAI Fall Symposium Series*.
- Hester, T.; Vecerik, M.; Pietquin, O.; Lanctot, M.; Schaul, T.; Piot, B.; Horgan, D.; Quan, J.; Sendonaris, A.; Dulac-Arnold, G.; Osband, I.; Agapiou, J.; Leibo, J. Z.; and Grusly, A. 2018. Deep q-learning from demonstrations. In *AAAI*.
- Ibarz, B.; Leike, J.; Pohlen, T.; Irving, G.; Legg, S.; and Amodei, D. 2018. Reward learning from human preferences and demonstrations in atari. In *Proc. of the Conference on Neural Information Processing Systems (NIPS)*.
- Isbell, C. L.; Shelton, C. R.; Kearns, M.; Singh, S. P.; and Stone, P. 2001. Cobot: A social reinforcement learning agent. In *Proc. of the Conference on Neural Information Processing Systems (NIPS)*.
- James, S.; Bloesch, M.; and Davison, A. J. 2018. Task-embedded control networks for few-shot imitation learning. In *Proc. of the Conference on Robot Learning (CoRL)*.
- Kelly, M.; Sidrane, C.; Driggs-Campbell, K.; and Kochenderfer, M. J. 2018. Hg-dagger: Interactive imitation learning with human experts.
- Krenging, S. 2018. Newtonian action advice: Integrating human verbal instruction with reinforcement learning.
- Laird, J. E.; Gluck, K.; Anderson, J.; Forbus, K. D.; Jenkins, O. C.; Lebiere, C.; Salvucci, D.; Scheutz, M.; Thomaz, A.; Trafton, G.; et al. 2017. Interactive task learning. *IEEE Intelligent Systems* 32(4):6–21.
- Lazaridou, A.; Peysakhovich, A.; and Baroni, M. 2016. Multi-agent cooperation and the emergence of (natural) language. *arXiv preprint arXiv:1612.07182*.
- Levine, S.; Finn, C.; Darrell, T.; and Abbeel, P. 2016. End-to-end training of deep visuomotor policies. *The Journal of Machine Learning Research* 17(1):1334–1373.
- Littman, M. L. 1994. Markov games as a framework for multi-agent reinforcement learning. In *Machine learning proceedings 1994*. Elsevier, 157–163.
- Lowe, R.; Wu, Y.; Tamar, A.; Harb, J.; Abbeel, O. P.; and Mordatch, I. 2017. Multi-agent actor-critic for mixed cooperative-competitive environments. In *Advances in Neural Information Processing Systems*, 6379–6390.
- Mindermann, S.; Shah, R.; Gleave, A.; and Hadfield-Menell, D. 2018. Active inverse reward design. In *Proc. of the ICML/IJCAI/AAMAS Workshop on Goals for Reinforcement Learning*.
- Mnih, V.; Kavukcuoglu, K.; Silver, D.; Graves, A.; Antonoglou, I.; Wierstra, D.; and Riedmiller, M. A. 2013. Playing atari with deep reinforcement learning. In *Proc. of the Conference on Neural Information Processing Systems (NIPS), Workshop on Deep Learning*.
- Mordatch, I., and Abbeel, P. 2018. Emergence of grounded compositional language in multi-agent populations. In *Proc. of the AAAI Conference on Artificial Intelligence (AAAI)*.
- Palmer, G.; Tuyls, K.; Bloembergen, D.; and Savani, R. 2018. Lenient multi-agent deep reinforcement learning. In *Proc. of Autonomous Agents and Multiagent Systems (AAMAS)*.
- Peloquin, B. N.; Goodman, N. D.; and Frank, M. C. 2019. The interactions of rational, pragmatic agents lead to efficient language structure and use. *PsyArXiv*.
- Perez-Dattari, R.; Celemin, C.; del Solar, J. R.; and Kober, J. 2018. Interactive learning with corrective feedback for policies based on deep neural networks. In *Proc. of the International Symposium on Experimental Robotics (ISER)*.
- Potter, M. A., and Jong, K. A. D. 1994. A cooperative coevolutionary approach to function optimization. *The Third Parallel Problem Solving From Nature* 249–257.
- Ross, S.; Gordon, G.; and Bagnell, D. 2011a. 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*, 627–635.

- Ross, S.; Gordon, G. J.; and Bagnell, J. A. 2011b. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proc. of the International Conference on Artificial Intelligence and Statistics (AISTATS)*.
- Santoro, A.; Bartunov, S.; Botvinick, M.; Wierstra, D.; and Lillincrap, T. P. 2016. One-shot learning with memory-augmented neural networks. In *Proc. of the International Conference on Machine Learning (ICML)*.
- Schaal, S. 1999. Is imitation learning the route to humanoid robots? *Trends in Cognitive Sciences* 233–242.
- Schmidhuber, J. 1987. *Evolutionary Principles in Self-Referential Learning*. Ph.D. Dissertation, Institut f. Informatik, Tech. Univ. Munich.
- Singh, S.; Lewis, R. L.; and Barto, A. G. 2009. Where do rewards come from. In *Proceedings of the annual conference of the cognitive science society*, 2601–2606.
- Thomaz, A. L.; Hoffman, G.; ; and Breazeal, C. 2005. Real-time interactive reinforcement learning for robots. In *AAAI*.
- Thrun, S., and Pratt, L. 2012. *Learning to learn*. Springer Science & Business Media.
- Wang, Z.; Schaul, T.; Hessel, M.; Van Hasselt, H.; Lanctot, M.; and De Freitas, N. 2015. Dueling network architectures for deep reinforcement learning. *arXiv preprint arXiv:1511.06581*.
- Wang, J. X.; Kurth-Nelson, Z.; Tirumala, D.; Soyer, H.; Leibo, J. Z.; Munos, R.; Blundell, C.; Kumaran, D.; and Botvinick, M. 2016. Learning to reinforcement learn. *arXiv preprint arXiv:1611.05763*.
- Warnell, G.; Waytowich, N.; Lawhern, V.; and Stone, P. 2018. Deep tamer: Interactive agent shaping in high-dimensional state spaces. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Watkins, C. J. C. H. 1989. *Learning from delayed rewards*. Ph.D. Dissertation, King’s College, Cambridge.
- Xu, K.; Ratner, E.; Dragan, A.; Levine, S.; and Finn, C. 2018. Learning a prior over intent via meta-inverse reinforcement learning.
- Yu, T.; Finn, C.; Xie, A.; Dasari, S.; Zhang, T.; Abbeel, P.; and Levine, S. 2018a. One-shot imitation from observing humans via domain-adaptive meta-learning. In *Proc. of Robotics: Science and Systems (RSS)*.
- Yu, T.; Finn, C.; Xie, A.; Dasari, S.; Zhang, T.; Abbeel, P.; and Levine, S. 2018b. One-shot imitation from observing humans via domain-adaptive meta-learning. *arXiv preprint arXiv:1802.01557*.