

Teachable Robots: Understanding Human Teaching Behavior to Build More Effective Robot Learners

Andrea L. Thomaz and Cynthia Breazeal

Abstract

While Reinforcement Learning (RL) is not traditionally designed for interactive supervisory input from a human teacher, several works in both robot and software agents have adapted it for human input by letting a human trainer control the reward signal. In this work, we experimentally examine the assumption underlying these works, namely that the human-given reward is compatible with the traditional RL reward signal. We describe an experimental platform with a simulated RL robot and present an analysis of real-time human teaching behavior found in a study in which untrained subjects taught the robot to perform a new task. We report three main observations on how people administer feedback when teaching a robot a task through Reinforcement Learning: (a) they use the reward channel not only for feedback, but also for future-directed guidance; (b) they have a positive bias to their feedback — possibly using the signal as a motivational channel; and (c) they change their behavior as they develop a mental model of the robotic learner. Given this, we made specific modifications to the simulated RL robot, and analyzed and evaluated its learning behavior in four additional experiments with human trainers. We report significant improvements on several learning measures. This work demonstrates the importance of understanding the human-teacher/robot-learner partnership in order to design algorithms that support how people want to teach while simultaneously improving the robot's learning behavior.

Key words: Human-Robot Interaction, Reinforcement Learning, User Studies

1 Introduction

As robots enter the human environment to assist people in their daily lives, the ability for ordinary people to easily teach them new tasks will be key to their success. Various works have addressed some of the hard problems robots

face when learning in the real-world, e.g., real-time learning in environments that are partially observable, dynamic, continuous (Mataric, 1997; Thrun and Mitchell, 1993; Thrun, 2002). However, learning quickly from interactions with a human teacher poses additional challenges (e.g., limited human patience, ambiguous human input, etc.) as well as benefits for machine learning systems.

Several examples of agents that learn interactively with a human teacher are based on Reinforcement Learning (RL). RL has certain desirable qualities, such as the possibility to explore and learn from experience. However, many also question RL as a viable technique for learning in complex real-world environments because of practical problems, such as long training time requirements; non-scaling state representations; sparse rewards; and safe exploration strategies. Many of these considerations are particularly pertinent to robots using RL, prompting the use of human guidance. As a result, RL has been utilized for teaching robots and game characters; a popular approach incorporates real-time human feedback by having a person supply reward and/or punishment as an additional input to the reward function (Blumberg et al., 2002; Kaplan et al., 2002; Isbell et al., 2001; Evans, 2002; Stern et al., 1998).

Most of this work models the human input as indistinguishable from any other feedback coming from the environment, and implicitly assumes people will correctly communicate feedback as expected by the algorithm. We question these assumptions and argue that reinforcement-based learning approaches should be reformulated to more effectively incorporate a human teacher. To address this, we advocate a systems approach that integrates machine learning into a Human-Robot Interaction (HRI) framework. Our first goal is to understand the nature of the teacher’s input in order to adequately support a human teacher’s contribution in guiding a robot’s exploration. Specifically, we want to understand *how people want to teach* and *what they want to communicate to the robot learner*. Our second goal is to incorporate these insights into standard machine learning frameworks to improve a robot’s learning behavior.

This paper presents a series of five experiments analyzing the scenario of a human teaching a virtual robot to perform a novel task within a reinforcement-based learning framework. Our experimental system, *Sophie’s Kitchen*, is a computer game that allows a Q-Learning agent to be trained interactively to perform a task.¹

In the first experiment (Sec. 5) we use this game to study 18 people’s interactions with the agent and present an analysis of human teaching behavior. We have found several prominent characteristics for how human players approach

¹ Q-Learning is used as the instrument for this work because it is a standard and widely understood RL algorithm, thus affording the transfer of these lessons to other reinforcement-based approaches.

the task of explicitly teaching a RL agent with direct control of the reward signal. To our knowledge, this work is the first to explicitly address and report such results, relevant to any interactive learning algorithm.

- People want the ability to direct the agent’s attention to guide the exploration process.
- People have a positive bias in their rewarding behavior, suggesting that players convey both instrumental and motivational intents with their communication channel.
- People adapt their teaching strategy as they develop a mental model of the agent’s learning.

The second contribution of this work is to incorporate these findings into specific modifications of the agent’s graphical interface and its RL algorithm. We subsequently had over 200 people play the game in four follow-up experiments, showing that our modifications significantly improve the learning behavior of the agent and make the agent’s exploratory behavior more appropriately responsive to the human’s instruction.

- **Leveraging Human Guidance:** In the second experiment (Sec. 8), we show the positive effects of adding a guidance channel of communication. Human players are able to direct the agents attention to yield a faster and more efficient learning process.
- **Transparency to Guide a Human Teacher:** In the third experiment (Sec. 9), we show that transparency behaviors, such as gaze, that reveal the internal state of the agent can be utilized to improve the human’s teaching.
- **The Asymmetry of Human Feedback:** In the fourth and fifth experiments (Sec. 10), we show beneficial asymmetric interpretations of feedback from a human partner. The fourth experiment shows that giving human players a separate motivational communication channel decreases the positive rewards bias. The fifth experiment shows the benefits of treating negative feedback from the human as both feedback for the last action and a suggestion to reverse the action if possible.

We believe this work contributes to the design of real-time learning agents that are better matched to human teaching behavior. Not only do such agents learn more effectively, but they are also easier to teach. By understanding the coupled human-teacher/robot-learner system, we demonstrate that it is possible to design algorithms that support how people want to teach while simultaneously improving the machine’s ability to learn.

2 Background: Related Works in Human-Trainable Systems

A review of related works in machine learning yields several interesting dimensions upon which human-trainable systems can be characterized. One interesting dimension is implicit versus explicit training. For instance, personalization agents and adaptive user interfaces rely on the human as an implicit teacher to model human preferences or activities through passive observation of the user’s behavior (Lashkari et al., 1994; Horvitz et al., 1998; Mitchell et al., 2006). In contrast, this work addresses explicit training where the human teaches the learner through interaction.

For systems that learn via interaction, another salient dimension is whether the human or the machine leads the interaction. For instance, active learning or learning with queries is an approach that explicitly acknowledges an interactive supervisor (Cohn et al., 1995; Schohn and Cohn, 2000). Through queries, the algorithm is in control of the interaction without regard of what a human will be able to provide in a real scenario. In contrast, this work addresses the human-side of the interaction and specifically asks *how do humans want to teach machines?*

A third interesting dimension is the balance between having the machine rely on human guidance versus its own exploration to learn new tasks. A number of systems rely on a human guidance paradigm where the learning problem is essentially reduced to programming through natural interfaces — with little if any exploration on the part of the machine, yielding a dependence on having a human present to learn (e.g., learning by demonstration (Nicolescu and Matarić, 2003; Schaal, 1999; Voyles and Khosla, 1998; Lieberman, 2001), by tutelage (Lockerd and Breazeal, 2004), or straight communication (Lauria et al., 2002)). In contrast, modified reinforcement-based approaches (e.g., the human contributes to the reward function) are positioned strongly along the exploration dimension (Blumberg et al., 2002; Kaplan et al., 2002; Isbell et al., 2001; Evans, 2002; Stern et al., 1998).

In contrast, an important goal of this work is to create learning systems that can dynamically slide along the exploration-guidance spectrum, to leverage a human teacher when present as well as learn effectively on its own. While there are known practical issues with RL (training time requirements, representations of state and hidden state, practical and safe exploration strategies), we believe that an appropriate reformulation of RL-based approaches to include input from a human teacher could alleviate these current shortcomings. To do this properly, we must deeply understand the human teacher as a unique contribution that is distinct from other forms of feedback coming from the environment.

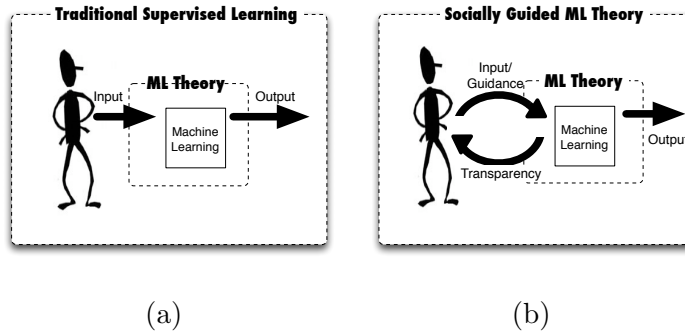


Fig. 1. 1(a) is a standard view of supervised learning, analyze input and output a model or classifier, etc. Our approach includes the human teacher, 1(b), emphasizing that teaching/learning is a two-way process. We add transparency, where the machine learner provides feedback to the human teacher during learning; and we augment the human input with guidance. We aim to enhance the performance of the tightly coupled partnership of a machine learner with a human teacher.

3 Approach: A HRI Framework for Machine Learning

Our approach is based on a *Social Learner Hypothesis*, namely that humans will naturally want to teach robots as social learners. As such, our work draws inspiration from Situated Learning Theory — a field of study that looks at the social world of children and how it contributes to their development. A key concept is *scaffolding*, where a teacher provides support such that a learner can achieve something they would not be able to accomplish independently (L. S. Vygotsky, 1978; Greenfield, 1984).

In a situated learning interaction, the teaching and learning processes are intimately coupled. A good instructor maintains a mental model of the learner’s state (e.g., what is understood so far, what remains confusing or unknown, etc.) in order to appropriately support the learner’s current needs. In particular, attention direction is one of the essential mechanisms that contribute to structuring the learning process (Wertsch et al., 1984). Other scaffolding acts include providing feedback, structuring successive experiences, regulating the complexity of information, and otherwise guiding the learner’s exploration. In general, this is a complex process where the teacher dynamically adjusts their support based on the learner’s demonstrated skill level and success.

The learner, in turn, helps the instructor by making their learning process *transparent* to the teacher through communicative acts (such as facial expressions, gestures, gaze, or vocalizations that reveal understanding, confusion, attention), and by demonstrating their current knowledge and mastery of the task (Krauss et al., 1996; Argyle et al., 1973). Through this reciprocal and tightly coupled interaction, the learner and instructor cooperate to simplify the task for the other — making each a more effective partner.

This situated learning process stands in dramatic contrast to typical machine learning scenarios that have traditionally ignored “teachability issues” such as how to make the teaching-learning process interactive or intuitive for a human partner. We advocate a new perspective that reframes the machine learning problem as an interaction between the human and the machine. This allows us to take advantage of human teaching behavior to construct a machine learning process that is more amenable to the human partner.

Figure 1(a) is a high level view of a supervised machine learning process. A human provides input to the learning mechanism, which performs its task and provides the output. Alternatively, an HRI perspective of machine learning models the complete human-machine system, characterized in Figure 1(b). This simple diagram highlights the key aspects of a social learning system, an interaction approach to machine learning forces the research community to consider many new questions. We need a principled theory of the content and dynamics of this tightly coupled process in order to design systems that can learn effectively from ordinary users.

Input Channels: A social interaction approach begins with the question: “How do humans want to teach?” In addition to designing the interaction based on what the machine needs to succeed in learning, we need to also understand what kinds of intentions people will try to communicate in their everyday teaching behavior. We can then change the input portion of the machine learning training process to better accommodate a human partner. It is important to understand the many ways that natural human social cues (e.g. referencing, attention direction, etc.) can frame the learning problem for a standard machine learning process. This paper explicitly examines the effect of allowing the human to guide the attention of a learner as well as provide feedback during its exploration process.

Output Channels: A social interaction approach asks: “How can the output provided by the learning agent improve the performance of the teaching-learning system?” In a tightly coupled interaction, a ‘black box’ learning process does nothing to improve the quality and relevance of the instructional guidance. However, transparency of the internal state of the machine could greatly improve the learning experience. By communicating its internal state, revealing what is known and what is unclear, the robot can guide the teaching process. To be most effective, the robot should reveal its internal state in a manner that is intuitive for the human partner (Breazeal, 2002; Arkin et al., 2003). For instance, facial expression, eye gaze, and behavior choices are a significant part of this output channel.

Input/Output Dynamics: Combining the previous two topics, we recognize that these input and output channels interact over time. The dynamics of the interaction can change the nature of the input from the human. In partic-



Fig. 2. *Sophie's Kitchen*. The agent is in the center, with a shelf on the right, oven on the left, a table in between, and five cake baking objects. The vertical bar is the interactive reward and is controlled by the human.

ular, the temporal structure of teaching versus performing may significantly influence the behavior of the human. An incremental, on-line learning system creates a very different experience for the human than a system that must receive a full set of training examples before its performance can be evaluated. Iterative feedback allows for on-line refinement; the human can provide another example or correct mistakes right away instead of waiting to evaluate the results at the end of the training process. Moreover, the sense that progress is being made may keep the human engaged with the training process for a longer period of time, which in turn benefits the learning system.

4 Experimental Platform: *Sophie's Kitchen*

To investigate the ways in which social interaction can impact machine learning for robots, we have implemented a Java-based simulation platform, “*Sophie's Kitchen*”, to experiment with learning algorithms and enhancements. *Sophie's Kitchen* is an object-based state-action MDP space for a single agent, Sophie, using a fixed set of actions on a fixed set of stateful objects.

4.1 *Sophie's Kitchen* MDP

The task scenario used is a kitchen world (see Fig. 2), where the agent, Sophie, learns to bake a cake. This system is defined by (L, O, Σ, T, A) .

- There are a finite set of k locations $L = \{l_1, \dots, l_k\}$. In our kitchen task, $k = 4$; $L = \{\text{Shelf}, \text{Table}, \text{Oven}, \text{Agent}\}$. As shown in Fig. 2, the agent is surrounded by a shelf, table and oven; and the location *Agent* is available to objects (i.e., when the agent picks up an object, then it has location *Agent*).
- There is a finite set of n objects $O = \{o_1, \dots, o_n\}$. Each object can be in one of an object-specific number of mutually exclusive object states. Thus, Ω_i is the set of states for object o_i , and $O^* = (\Omega_1 \times \dots \times \Omega_n)$ is the entire object

configuration space. In the kitchen task scenario $n = 5$: the objects **Flour**, **Eggs**, and **Spoon** each have only one object state; the object **Bowl** has five object states: **empty**, **flour**, **eggs**, **both**, **mixed**; and the object **Tray** has three object states: **empty**, **batter**, **baked**.

- Let L^A be the possible agent locations: $L^A = \{\text{Shelf, Table, Oven}\}$; and let L^O be the possible object locations: $L^O = \{\text{Shelf, Table, Oven, Agent}\}$. Then the legal set of states is $\Sigma \subset (L^A \times L^O \times O^*)$, and a specific state is defined by $(l_a, l_{o_1} \dots l_{o_n}, \omega)$: the agent's location, $l_a \in L^A$, and each object's location, $l_{o_i} \in L^O$, and the object configuration, $\omega \in O^*$.
- T is a transition function: $\Sigma \times A \mapsto \Sigma$. The action space A is expanded from four atomic actions (**GO**<x>, **PUT-DOWN**<x>, **PICK-UP**<x>, **USE**<x><y>): Assuming the locations L^A are arranged in a ring, the agent can always **GO left** or **right** to change location; she can **PICK-UP** any object in her current location; she can **PUT-DOWN** any object in her possession; and she can **USE** any object in her possession on any object in her current location. The agent can hold only one object at a time. Thus the set of actions available at a particular time is dependent on the particular state, and is a subset of the entire action space, A . Executing an action advances the world state in a deterministic way defined by T . For example, executing **PICK-UP** <Flour> advances the state of the world such that the **Flour** has location **Agent**. **USE**ing an ingredient on the **Bowl** puts that ingredient in it; using the **Spoon** on the **both**-**Bowl** transitions its state to the **mixed**-**Bowl**, etc.

In the initial state, s_0 , all objects and the agent are at location **Shelf**. A successful completion of the task will include putting flour and eggs in the bowl, stirring the ingredients using the spoon, then transferring the batter into the tray, and finally putting the tray in the oven. Some end states are so-called *disaster* states (for example—putting the eggs in the oven), which result in a negative reward ($r = -1$), the termination of the current trial, and a transition to state s_0 . In order to encourage short sequences, an inherent negative reward of $r = -.04$ is placed in any non-goal state.

The kitchen task has on the order of 10,000 states with between 2 and 7 actions available in each state. If the agent is left to learn on its own in this space (restarting in random states rather than s_0 after a disaster/goal is reached), it requires a few thousand actions before the goal state is reached for the first time. Due to the flexibility of the task, there are many action sequences that can lead to the desired goal. Here is one such sequence:

PICK-UP Bowl; GO right; PUT-DOWN Bowl; GO left; PICK-UP Flour; GO right; USE Flour, Bowl; PUT-DOWN Flour; GO left; PICK-UP Eggs; GO right; USE Eggs, Bowl; PUT-DOWN Eggs; GO left; PICK-UP Spoon; GO right; USE Spoon, Bowl; PUT-DOWN Spoon; GO left; PICK-UP Tray; GO right; PUT-DOWN Tray; PICK-UP Bowl; USE Bowl, Tray; PUT-DOWN Bowl; PICK-UP Tray; GO right; PUT-DOWN Tray.

Algorithm 1 Q-Learning with Interactive Rewards from a Human Partner

- 1: s =last state, s' =current state, a =last action, r =reward
- 2: **while** learning **do**
- 3: a = random select weighted by $Q[s, a]$ values
- 4: execute a , and transition to s'
 (small delay to allow for human reward)
- 5: sense reward, r
- 6: update Q-value:

$$Q[s, a] \leftarrow Q[s, a] + \alpha(r + \gamma(\max_{a'} Q[s', a']) - Q[s, a])$$

- 7: **end while**
-

4.2 Learning Algorithm

The algorithm implemented for the experiments presented in this paper is a standard Q-Learning algorithm (learning rate $\alpha = .3$ and discount factor $\gamma = .75$) (Watkins and Dayan, 1992). This is shown above in Algorithm 1. A slight delay happens in line 4 as the agent’s action is animated and also to allow the human time to issue interactive rewards. Q-Learning is used as the instrument for this work because it is a widely understood RL algorithm, thus affording the transfer of these lessons to other reinforcement-based approaches.

4.3 Interactive Rewards Interface

A central feature of *Sophie’s Kitchen* is the interactive reward interface. Using the mouse, a human trainer can—at any point in the operation of the agent—award a scalar reward signal $r \in [-1, 1]$. The user receives visual feedback enabling them to tune the reward signal before sending it to the agent. Choosing and sending the reward does not halt the progress of the agent, which runs asynchronously to the interactive human reward.

The interface also lets the user make a distinction between rewarding the whole state of the world or the state of a particular object (object specific rewards). An object specific reward is administered by doing a feedback message on a particular object (objects are highlighted when the mouse is over them to indicate that any subsequent reward will be object specific). This distinction exists to test a hypothesis that people will prefer to communicate feedback about particular aspects of a state rather than the entire state. However, object specific rewards are used only to learn about the human trainer’s behavior and communicative intent; the learning algorithm treats all rewards in the traditional sense of pertaining to a whole state and action pair.

5 Experiment: How People Teach RL Agents

Some may claim it is fairly obvious that a standard Reinforcement Learning agent is not suitable for learning an assistive task using interactive reward training as described above—if only due to the vast number of trials necessary to form a reasonable policy. However, the details of what exactly needs adjustment, and what human factors are dominant in such an interaction, are largely unexplored. It is these components that we wish to uncover and enumerate. The purpose of this initial experiment with *Sophie’s Kitchen* is to understand, when given a single reward channel (as in prior works), how do people use it to teach the agent?

5.1 Experiment Design

In the experiment, 18 participants were solicited from the campus community and came to our research lab. Each volunteer played a computer game, in which their goal was to get the virtual robot, Sophie, to learn how to bake a cake on her own. Participants were asked to rate their expertise with machine learning software and systems on a scale of 1 to 7, (1=no experience, 7=very experienced), and we found it was an above average but reasonably diverse population (mean=3.7; standard deviation=2.3).²

Participants were told they could not tell Sophie what to do, nor could they do actions directly, but they could send Sophie the following messages via a mouse to help her learn the task:

- Click and drag the mouse up to make a green box, a positive message; and down for red/negative.
- By lifting the mouse button, the message is sent to Sophie, she sees the color and size of the message.
- Clicking on an object, this tells Sophie your message is about that object. As in, “Hey Sophie, this is what I’m talking about...” If you click anywhere else, Sophie assumes your feedback pertains to everything in general.

The system maintains an activity log and records time step and real time of each of the following: state transitions, actions, human rewards, reward aboutness (if object specific), disasters, and goals. Additionally, there was an informal interview after subjects completed the task.

² We had both male and female participants, but did not keep gender statistics of the population.

5.2 Results of the Teaching Study

Of the 18 participants only one person did not succeed in teaching Sophie the task. During the first day of testing, four participants had to interrupt their trial due to a software error. As a result, some of the analysis below includes only the 13 individuals that finished the complete task. However, since participants who experienced this error still spent a significant amount of time training the agent, their data is included in those parts of the analysis that relate to overall reward behavior. In this section we present three main findings about how people approach the task of teaching an RL agent with an interactive reward signal. 1) They assume the ability to guide the agent. 2) Their teaching behavior changes as they develop a mental model for the learning agent. 3) There is a positive bias in user rewards.

5.2.1 Guidance Intentions

Even though the instructions clearly stated that communication of both general and object specific rewards were *feedback* messages, many people assumed that object specific rewards were future directed messages or guidance for the agent. Several people mentioned this in the interview, and this is also suggested through behavioral evidence in the game logs.

An object reward used in a standard RL sense, should pertain to the last object the agent used. Figure 3 has a mark for each player, indicating the percentage of object specific rewards that were about the last object the agent used: 100% would indicate that the player always used object rewards in a feedback connotation, and 0% would mean they never used object rewards as feedback. We can see that several players had object rewards that were rarely correlated to the last object (i.e., for 8 people less than 50% of their object rewards were about the last object).

Interview responses suggested these people’s rewards actually pertain to the future, indicating what they want (or do not want) the agent to use next. A single test case is used to show how many people used object rewards as a guidance mechanism: When the agent is facing the shelf, a guidance reward could be administered (i.e., what to pick up). Further, a positive reward given to either the empty bowl or empty tray on the shelf could *only* be interpreted as guidance since this state would not be part of any desired sequence of the task (only the initial state). Thus, rewards to empty bowls and trays in this configuration serve to measure the prevalence of guidance behavior.

Figure 4 indicates how many people tried giving rewards to the bowl or tray when they were empty on the shelf. Nearly all of the participants, 15 of 18, gave rewards to the bowl or tray objects sitting empty on the shelf. This leads




Fig. 3. There is one mark for each player, indicating their percentage of object rewards that were about the last object of attention. This graph shows that many people had object rewards that were rarely about the last object, thus rarely used in a feedback orientation.




Fig. 4. A reward to the empty bowl or tray on the shelf is assumed to be meant as guidance instead of feedback. This graph shows that 15 of the 18 players gave rewards to the bowl/tray empty on the shelf.

to the conclusion that many participants tried using the reward channel to guide the agent’s behavior to particular objects, giving rewards for actions the agent was *about to do* in addition to the traditional rewards for what the agent had just done.

These *anticipatory* rewards observed from everyday human trainers will require new attention in learning systems and algorithms in order for agents to correctly interpret their human partners. Section 8 covers the design, implementation, and evaluation of algorithm and interface modifications for utilizing guidance.

5.2.2 Inferring a Model of the Learner

Informed by related work (Isbell et al., 2001), it is reasonable to expect people would habituate to the activity and that feedback would decrease over the training session. However, just the opposite was found: the ratio of rewards to actions over the entire training session had a mean of .77 and standard deviation of .18. Additionally, there is an increasing trend in the rewards-to-

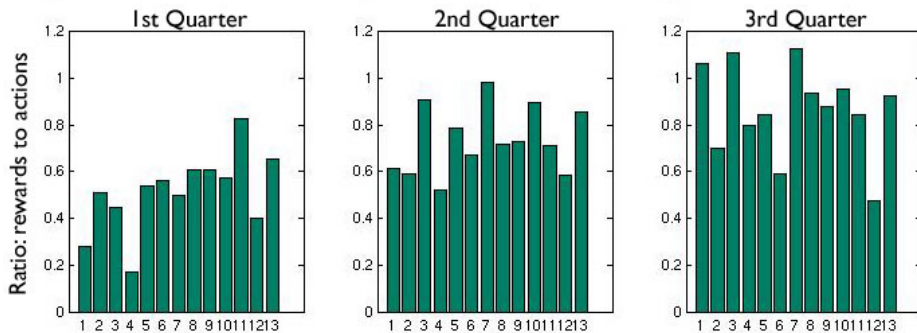


Fig. 5. Ratio of rewards to actions over the first three quarters of the training sessions shows an increasing trend.

actions ratio over the first three quarters of training. Fig. 5 shows data for the first three quarters for training, each graph has one bar for each individual indicating the ratio of rewards to actions. A 1:1 ratio in this case means that the human teacher gives a reward after every action taken by the agent. By the third graph more bars are approaching or surpassing a ratio of 1.

One explanation for this increasing trend is a shift in mental model; as people realize the impact of their feedback they adjusted their reward schedule to fit this model of the learner. This finds anecdotal support in the interview responses. Many users reported that at some point they came to the conclusion that their feedback was helping the agent learn and they subsequently gave more rewards. Many users described the agent as a “stage” learner, that it would seem to make large improvements all at once. This is precisely the behavior one sees with a Q-Learning agent: fairly random exploration initially, and the results of learning are not seen until the agent restarts after a failure. Without any particular understanding of the algorithm, participants were quickly able to develop a reasonable mental model of the agent through the interaction. They were encouraged by the learning progress, and subsequently gave more rewards.

A second expectation was that people would naturally use goal-oriented and intentional communication (measured by allowing people to specify object specific rewards, see Sec. 4.3). The difference between the first and last quarters of training shows that many people tried the object specific rewards at first but stopped using them over time (Fig. 6). In the interview, many users reported that the object rewards “did not seem to be working.” Thus, many participants tried the object specific rewards initially, but were able to detect over time that an object specific reward did not have a different effect on the learning process than a general reward (which is true), and therefore stopped using the object rewards.

These are concrete examples of the human trainer’s propensity to learn from

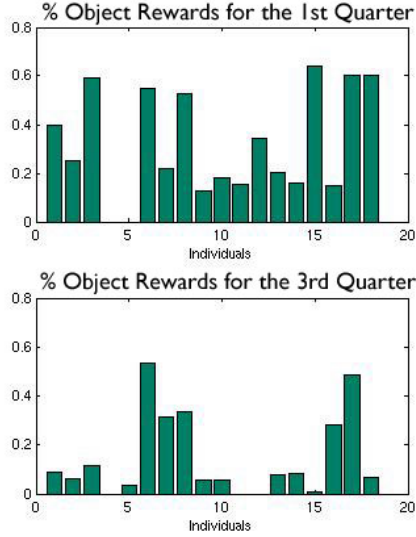


Fig. 6. Each bar represents an individual and the height is the percentage of object rewards. The difference in the first and last training quarters shows a drop off in usage over time.




Fig. 7. Histograms of rewards for each individual in the first quarter of their session. The left column is negative rewards and the right is positive rewards. Most people even in the first quarter of training have a much higher bar on the right.

the agent how to best impact the process. This presents a huge opportunity for an interactive learning agent to *improve its own learning environment* by communicating more internal state to the human teacher, making the learning process more transparent. Section 9 details the implementation and evaluation of a transparent gazing behavior to improve the learning environment.

5.2.3 An Asymmetric Use of Rewards

For many people, a large majority of rewards given were positive, the mean percentage of positive rewards for all players was 69.8%. This was thought at first to be due to the agent improving and exhibiting more correct behavior over time (soliciting more positive rewards); however, the data from the first quarter of training shows that well before the agent is behaving correctly, the majority of participants still show a positive bias. Fig. 7 shows reward histograms for each participant’s first quarter of training; the number of negative rewards on the left and positive rewards on the right, most participants have a much larger bar on the right. A plausible hypothesis is that people are falling into a natural teaching interaction with the agent, treating it as a social entity that needs encouragement. Some people specifically mentioned in the interview that they felt positive feedback would be better for learning. Section 10 details the implementation and evaluation of *Sophie’s Kitchen* with asymmetric human rewards.

6 Lessons Learned from the Teaching Study

The findings in this study offer empirical evidence to support our Social Learner Hypothesis and the concept of *partnership* when humans teach artificial agents. When untrained users are asked to interactively train a RL agent, we see them treat the agent in a social way, tending towards positive feedback, guiding the robot, and adjusting their training behavior as the interaction proceeds, reacting to the behavior of the learner. Importantly, we see this tendency even without specifically adding any behavior to the robot to elicit this attitude. This suggests that there is a human propensity to treat and understand other entities as intentional agents, and to adapt to them.

To date, RL does not account for the teacher’s commitment to adapt to the learner, presenting an opportunity for an interactive learning agent to improve its own learning environment by communicating more of its internal state.

Additionally, our findings indicate that the learning agent can take better advantage of the different kinds of messages a human teacher is trying to communicate. In common RL, a reward signal is stationary and is some function of the environment. It is usually a symmetrical scalar value indicating positive or negative feedback for being in the current state or for a particular state-action pair. Introducing human-derived real-time reward prompts us to reconsider these assumptions. We find that with a single communication channel people have various communicative intents—feedback, guidance, and motivation. Augmenting the human reward channel will likely be helpful to both the human teacher and the machine learning algorithm.

Finally, timing of rewards has been a topic in the RL community, particularly the credit assignment problem associated with delayed rewards. As opposed to delayed rewards, however, we saw that many human teachers administered anticipatory or guidance rewards to the agent. While delayed rewards have been discussed, the concept of rewarding the *action the agent is about to do* is novel and will require new tools and attention in the RL community.

7 Next Steps: Modifications and Follow-Up Experiments


The results from our first experiment suggest a few specific recommendations for interactive machine learning. One of these recommendations is that the communication from the human teaching partner cannot be merged into one single reward signal. We need to embellish the communication channel to account for the various intentions people wish to convey to the machine, particularly guidance intentions. Additionally, people tune their behavior to match the needs of the machine, and this process can be augmented with more transparency of the internal state of the learner.

These lessons motivate a set of specific extensions to the system in order to more deeply understand the impact of social guidance and transparency behaviors on a machine learning process. In follow-up versions of the Sophie game we examine the following extensions:

Guidance: Having found people try to communicate both guidance and feedback in their reward message, the next version of Sophie distinguishes between these two inputs. Users can still send a normal feedback message, but they can also communicate attention direction or guidance. The learning algorithm is biased to select an actions based on this attention direction signal when it is available.

Gaze as a Transparency Behavior: A second modification to *Sophie's Kitchen* explores the effect of gazing between the objects of attention for equally valuable candidate actions during the action selection phase. This communicates a level of uncertainty through the amount of gazing that precedes action. We expect this transparency behavior to improve the teacher's mental model of the learner, creating a more understandable interaction for the human and a better learning environment for the machine. Specifically, we expect more guidance to be administered when the agent seems uncertain.

Undo: A third modification has the Sophie agent respond to negative feedback with an UNDO behavior (natural correlate or opposite action) when possible. This is expected to increase the responsiveness and transparency of the agent and could balance the amount of positive and negative rewards seen. The



(a) Feedback message.

(b) Guidance message.

Fig. 8. The embellished communication channel includes the feedback messages as well as guidance messages. In 8(a), feedback is given by left-clicking and dragging the mouse up to make a green box (positive) and down for red (negative). In 8(b), guidance is given by right-clicking on an object of attention, selecting it with the yellow square.

algorithm changes such that in the step following negative feedback, the action selection mechanism chooses the action that ‘un-does’ the last action if possible.

Motivation: One hypothesis about the positive rewards bias is that people were using the reward channel for motivation. A fourth modification of the Sophie game allows explicit encouragement or discouragement by administering a reward *on* Sophie. This will allow people to distinguish specific feedback about the task (e.g., “That was good!”) from general motivational feedback (e.g., “Doing good Sophie!”).

8 Leveraging Human Guidance

We modified the *Sophie’s Kitchen* game to add a guidance channel of communication. Based on prior work with supervised RL, the expectation is that a guidance signal should improve learning (Clouse and Utgoff, 1992; Maclin et al., 2005; Smart and Kaelbling, 2002) (which we confirm in an experiment with an expert trainer). The contribution of our work is the focus on non-expert human trainers. In an experiment we show that everyday human teachers can use attention direction as a form of guidance, to improve the learning behavior of an RL agent.

8.1 Modification to Game Interface

The guidance intentions identified in our teaching experiment suggests that people want to speak directly to the action selection part of the algorithm to

influence the exploration strategy. To accomplish this, we added a guidance channel of communication to distinguish this intention from feedback. Clicking the right mouse button draws an outline of a yellow square. When the yellow square is administered on top of an object, this communicates a guidance message to the learning agent and the content of the message is the object. Figure 8(b) shows the player guiding Sophie to pay attention to the bowl. Note, the left mouse button still allows the player to give feedback as described previously, but there are no longer object rewards.

8.2 Modification to Learning Algorithm

Algorithm 1 describes the standard Q-Learning algorithm used for the initial interactive training sessions with *Sophie's Kitchen*. Conceptually, our modified version gives the algorithm a pre-action and post-action phase in order to incorporate the new guidance input. In the pre-action phase the agent registers guidance communication to bias action selection, and in the post-action phase the agent uses the reward channel in the standard way to evaluate that action and update a policy. The modified learning process is shown in Algorithm 2.

The agent begins each iteration of the learning loop by pausing to allow the teacher time to administer guidance (1.5 seconds). The agent saves the object of the human's guidance messages as g . During the action selection step, the default behavior chooses randomly between the set of actions with the highest Q-values, within a bound β . However, if any guidance messages were received, the agent will *instead* choose randomly between the set of actions that have to do with the object g . In this way the human's guidance messages bias the action selection mechanism, narrowing the set of actions the agent considers.

8.3 Evaluation: Guidance Improves Learning

The first experiment we ran with the guidance modification evaluates the effects of guidance from an expert trainer. This is analogous to prior works, and serves to confirm that supervision is beneficial to the agent in *Sophie's Kitchen*. We collected data from expert³ training sessions, in two conditions:

- (1) **No guidance:** has feedback only and the trainer gives one positive or negative reward after every action.
- (2) **Guidance:** has both guidance and feedback available; the trainer uses the same feedback behavior and additionally guides to the desired object at every opportunity.

³ one of the authors

Algorithm 2 Interactive Q-Learning modified to incorporate interactive human guidance in addition to feedback.

```
1: while learning do
2:   while waiting for guidance do
3:     if receive human guidance message then
4:        $g = \text{guide-object}$ 
5:     end if
6:   end while
7:   if received guidance then
8:      $a = \text{random selection of actions containing } g$ 
9:   else
10:     $a = \text{random selection weighted by } Q[s, a] \text{ values}$ 
11:   end if
12:   execute  $a$ , and transition to  $s'$ 
   (small delay to allow for human reward)
13:   sense reward,  $r$ 
14:   update Q-value:
```

$$Q[s, a] \leftarrow Q[s, a] + \alpha(r + \gamma(\max_{a'} Q[s', a']) - Q[s, a])$$

```
15: end while
```

For the user’s benefit, we limited the task for this testing (e.g., taking out the spoon/stirring step, among other things). We had one user follow the above expert protocol for 10 training sessions in each condition. The results of this experiment are summarized in Table 1, showing that guidance improves several learning metrics. The number of training trials needed to learn the task was significantly less, 30%; as was the number actions needed to learn the task, 39% less. In the **guidance** condition the number of unique states visited was significantly less, 40%; thus the task was learned more efficiently. And finally the **guidance** condition was more successful, the number of trials ending in failure was 48% less, and the number of failed trials before the first successful trial was 45% less.

Having found guidance has the potential to drastically improve several metrics of the agent’s learning behavior, our final evaluation of the guidance modification looks at how the agent performs with ordinary human trainers. We solicited 11 people to come to our research lab to play the *Sophie’s Kitchen* game using both feedback and guidance messages. We added the following instructions about the guidance messages to the instructions from the previous experiment (and took out object specific rewards):

You can direct Sophie’s attention to particular objects with guidance messages. Click the right mouse button to make a yellow square, and use it to help guide Sophie to objects, as in ‘Pay attention to this!’

Table 1

An **expert** user trained 20 agents, with and without guidance, following a strict best-case protocol in each condition; this yields theoretical best-case effects of guidance on learning. (F = failed trials, G = first success). Results from 1-tailed t-tests.

Measure	Mean no guide	Mean guide	chg	t(18)	p
# trials	6.4	4.5	30%	2.48	.01
# actions	151.5	92.6	39%	4.9	<.01
# F	4.4	2.3	48%	2.65	<.01
# F before G	4.2	2.3	45%	2.37	.01
# states	43.5	25.9	40%	6.27	<.01

We compare the game logs of these players (the **guidance** condition), to 17 who played without the guidance signal (the **no guidance** condition). This comparison is summarized in Table 2.

Guidance players were faster than **no guidance** players. The number of training trials needed to learn the task was 48.8% less, and the number actions needed was 54.9% less. Thus, the ability for the human teacher to guide the agent’s attention to appropriate objects at appropriate times creates a significantly faster learning interaction.

The **guidance** condition provided a significantly more successful training experience. The number of trials ending in failure was 37.5% less, and the number of failed trials before the first successful trial was 41.2% less. A more successful training experience is particularly desirable when the learning agent is a robot that may not be able to withstand very many failure conditions. Additionally, a successful interaction, especially reaching the first successful attempt sooner, may help the human teacher feel that progress is being made and prolong their engagement in the process.

Finally, agents in the **guidance** condition learned the task by visiting a significantly smaller number of unique states, 49.6% less than the **no guidance** condition. Additionally, we analyze the percentage of time spent in a good portion of the state space, defined as $G = \{\text{every unique state in } X\}$, where $X = \{\text{all non-cyclic sequences, } s_0, \dots, s_n, \text{ such that } n \leq 1.25(\text{min_sequence_length}), \text{ and } s_n = \text{a goal state}\}$. The average percentage of time that **guidance** agents spent in G was 72.4%, and is significantly higher than the 60.3% average of **no guidance** agents. Thus, attention direction helps the human teacher keep the exploration of the agent within a smaller and more positive (useful) portion of the state space. This is a particularly important result since that the ability to deal with large state spaces has long been a criticism of RL. A human partner may help the algorithm overcome this challenge.

Table 2

Non-expert human players trained Sophie with and without guidance communication available and also show positive effects of guidance on the learning. (F = failed trials, G = first success). Results from 1-tailed t-tests.

Measure	Mean no guide	Mean guide	chg	t(26)	p
# trials	28.52	14.6	49%	2.68	<.01
# actions	816.44	368	55%	2.91	<.01
# F	18.89	11.8	38%	2.61	<.01
# F before G	18.7	11	41%	2.82	<.01
# states	124.44	62.7	50%	5.64	<.001
% good states	60.3	72.4		-5.02	<.001

9 Transparency to Guide a Human Teacher

In the previous section, we saw that the ability for the human teacher to direct the Sophie agent’s attention has significant positive effects on several learning metrics. This section reports a related result – that the ability of the agent to use gaze as a transparency behavior results in measurably better human guidance instruction.

Gaze requires that the learning agent have a physical/graphical embodiment that can be understood by the human as having a forward heading. In general, gaze precedes an action and communicates something about the action that is going to follow. In this way gaze serves as a transparency device, allowing an onlooker to make inferences about what the agent is likely to do next, their level of confidence and certainty about the environment, and perhaps whether or not guidance is necessary. A gaze behavior was added to the *Sophie’s Kitchen* game. The modified game was deployed on the World Wide Web, and data was collected from over 50 people playing the game, allowing for a concrete analysis of the effect Sophie’s gaze had on a human teacher’s behavior.

9.1 Modification to Game Interface

Recall the interactive Q-Learning algorithm modified for guidance (Algorithm 2). The gaze behavior modification makes one alteration to the stage at which the agent is waiting for guidance, shown in Algorithm 3. When the agent is waiting for guidance, it finds the set of actions, A^+ , with the highest Q-values, within a bound β . $\forall a \in A^+$, the learning agent gazes for 1 second at

the human. This also communicates overall task certainty or confidence as the agent will speed up when every set, A^+ , has a single action. The hypothesis is that this transparency will improve the teacher’s model of the learner, creating a more understandable interaction for the human and a better learning environment for the agent.

9.2 Evaluation: Gaze Improves Guidance

To evaluate the use of transparency, we deployed the *Sophie’s Kitchen* game on the World Wide Web. Participants were asked to play the computer game and were given instructions on administering feedback and guidance. Each of the 52 participants played the game in one of the following test conditions:

- **Guidance:** Players had both the feedback and the guidance channels of communication.
- **Gaze-guide:** Players had the feedback and guidance channels. Additionally, the agent used the gaze behavior.

The system maintained an activity log and recorded time step and real time of each of the following: state transitions, actions, human rewards, guidance messages and objects, gaze actions, disasters, and goals. These logs were analyzed to test the transparency hypothesis: Learners can help shape their learning environment by communicating aspects of the internal process. In particular, the gaze behavior will improve a teacher’s guidance instruction.

This hypothesis is evaluated through the comparison of players that had the **guidance** condition versus those that had the **gaze-guide** condition. These results are summarized in Table 3. Note that the players that did not have the gaze behavior still had ample opportunity to administer guidance; however, the time that the agent waits is uniform throughout.

Looking at the timing of each player’s guidance instruction, their communication can be separated into two segments: the percentage of guidance that was given when the number of action choices was ≥ 3 (high uncertainty), and when choices were ≤ 3 (low uncertainty), note that these are overlapping classes. Three is chosen as the midpoint because the number of action choices available to the agent at any time in the web-based version of *Sophie’s Kitchen* is at most 5. Thus we describe a situation where the number of equally valued action choices is ≥ 3 as high uncertainty, and ≤ 3 as low uncertainty.

Players in the **gaze-guide** condition had a significantly lower percentage of guidance when the agent had low uncertainty compared to the players in the **guidance** condition, $t(51) = -2.22, p = .015$. And conversely the percentage of guidance when the agent had high uncertainty increased from the **guidance**

Table 3

1-tailed t-test showing the effect of gaze on guidance. Compared to the guidance distribution without gaze, the gaze condition caused a decrease when uncertainty was low and an increase when uncertainty was high. (uncertainty low = number of action choices ≤ 3 , high = number of choices ≥ 3).

Measure	Gaze-Guide	Guidance	t(51)	p
% Guidance when uncertainty low	79	85	-2.22	<.05
% Guidance when uncertainty high	48	36	1.96	<.05

to the gaze-guide condition, $t(51) = 1.96, p = .027$. Thus, when the agent uses the gaze behavior to indicate which actions it is considering, the human trainers do a better job matching their instruction to the needs of the agent throughout the training session. They give more guidance when it is needed and less when it is not.

10 The Asymmetry of Human Feedback

One of the main findings in our initial experiment concerned the biased nature of positive and negative feedback from a human partner. Clearly, people have different intentions they are communicating with their positive and negative feedback messages. In this section we present to modifications to the game interface that address the asymmetric meanings of human feedback.

One hypothesis is that people are falling into a natural teaching interaction with the agent, treating it as a social entity that needs motivation and encouragement. People may feel bad giving negative rewards to the agent, or feel that it is important to be both instrumental and motivational with their communication channel. In interviews a number of participants mentioned that they believed the agent would learn better from positive feedback.

Another hypothesis is that negative rewards did not produce the expected reaction from the robot. A typical RL agent does not have an instantaneous reaction to either positive or negative rewards, but in the case of negative rewards, this could be interpreted as the agent “ignoring” the human’s feedback. In that case, the user may stop using them when they feel the agent is not taking their input into account. One way to address this is to introduce an UNDO behavior. Many actions (PICK-UP, PUT-DOWN, TURN) have a natural correlate or opposite action that can be performed in response to negative feedback. This could add to the responsiveness and transparency of the agent and balance the amount of positive and negative rewards seen.




Fig. 10. A reward is considered motivational rather than instrumental if it is administered *on* Sophie, as pictured here. Instructions about this input channel indicate that it is for general feedback about the task (e.g. "Doing good Sophie!" or "Doing bad!") as opposed to feedback about a particular action.

We explore both of these hypotheses in this section. First, we look at adding a motivation channel of communication, to test if the positive bias was due to motivational intentions. Second, we add the **UNDO** behavior and show that this reaction to a person's negative feedback produces a significantly better learning behavior for the RL agent.

10.1 Motivation

In this experiment, we add a motivation communication channel. Our hypothesis is that we should see the positive bias decrease when the players have a separate channel for motivational versus instrumental communication.

10.1.1 Modification to the Game Interface

For this experiment we have the original feedback channel of communication, and a dedicated motivational input. This is done by considering a reward motivational if it is administered *on* Sophie. For visual feedback the agent is shaded yellow to let the user know that a subsequent reward will be motivational. Figure 10 shows a positive motivational message to Sophie. The game instructions given to players indicate that this input channel is for general feedback about the task (e.g. "Doing good Sophie!" or "Doing bad!") as opposed to feedback about a particular action.

10.1.2 Evaluation: Motivation Intentions Confirmed

To test our hypothesis about people's motivational intents, we deployed the *Sophie's Kitchen* game on the World Wide Web and had 98 volunteers play the game. Players that had the motivation signal had a significantly more balanced feedback valance than the players that did not have it. Players that did not have a motivational channel had a mean ratio ($\frac{\#positive}{\#negative}$) of 2.07; whereas

those with the motivational channel had a mean ratio of 1.688. This is a significant effect, $t(96) = -2.02, p = .022$. Thus, we conclude that motivation is a separate intention that was folded into the players' positive feedback in the initial study. Future work is to understand how an agent can utilize this signal in a different way to improve the learning interaction.

10.2 UNDO Behavior

The UNDO modification addresses a second aspect of the asymmetric meaning of human feedback. The intuition is that positive feedback tells a learner undeniably, "what you did was good." However, negative feedback has multiple meanings: 1) that the last action was bad, and 2) that the current state is bad and future actions should correct that. Thus, negative feedback is about both the past and about future intentions for action. In the final modification to *Sophie's Kitchen*, the algorithm assumes that a negatively reinforced action should be reversed if possible. This UNDO interpretation of negative feedback shows significant improvements in several metrics of learning.

10.2.1 Modification to the Algorithm

This baseline algorithm is modified to respond to negative feedback with an UNDO behavior (a natural correlate or opposite action) when possible. Thus a negative reward affects the policy in the normal fashion, but also alters the subsequent action selection if possible. The proper UNDO behavior is represented within each primitive action and is accessed with an *undo* function: GO [direction] returns GO [-direction]; PICK-UP [object] returns PUT-DOWN [object]; PUT-DOWN [object] returns PICK-UP [object]; USE actions are not reversible. Algorithm 4 shows how this is implemented with the changes in lines 2 – 6, as compared to the baseline Algorithm 1.

10.2.2 Evaluation of UNDO Behavior

We found the UNDO response to negative feedback from the human trainer significantly improves the learning performance of the agent in a number of ways. Data was collected from 97 human participants by deploying the *Sophie's Kitchen* game on the World Wide Web.

The *Sophie's Kitchen* platform offers a measurable comparison between two conditions of the learning algorithm. In the **baseline** case the algorithm handles both positive and negative feedback in a standard way, feedback is incorporated into the value function (Alg. 1). In the **undo** case the algorithm uses feedback to update the value function but then also uses negative feedback

Algorithm 4 Interactive Q-Learning with the addition of the UNDO behavior

```
1: while learning do
2:   if (reward last cycle < -0.25) and (can undo last action,  $a_{last}$ ) then
3:      $a = \text{undo}(a_{last})$ 
4:   else
5:      $a = \text{random select weighted by } Q[s, a] \text{ values}$ 
6:   end if
7:   execute  $a$ , and transition to  $s'$ 
   (small delay to allow for human reward)
8:   sense reward,  $r$ 
9:   update policy:
```

$$Q[s, a] \leftarrow Q[s, a] + \alpha(r + \gamma(\max_{a'} Q[s', a']) - Q[s, a])$$

```
10: end while
```

in the action selection stage as an indication that the best action to perform next is the reverse of the negatively reinforced action (Alg. 4). Statistically significant differences were found between the **baseline** and **undo** conditions on a number of learning metrics (summarized in Table 4).

The **UNDO** behavior helps the agent avoid failure. The total number of failures during the learning phase was 37% less in the **undo** case. This is particularly interesting for robotic agents that need to learn in the real world where learning from failure may not be a viable option. The **undo** case also had 40% less failures before the first success. This is especially important when the agent is learning with a human partner. The human partner will have limited patience and will need to see progress quickly in order to remain engaged in the task. Thus, the **undo** behavior seems to be a good technique for reaching the first success faster.

There was a nearly significant effect for the number of actions required to learn the task, with the **undo** condition requiring 12% less steps (the high degree of variance in the number of steps needed to learn the task leads to the higher p value). Another indication of the efficiency of the **undo** case compared to the **baseline** is in the state space needed to learn the task. The number of unique states visited is 13% less in the **undo** case. This indicates that when the algorithm interprets negative feedback as a directive for reversing the previous action, or returning to the previous state, the resulting behavior is more efficient in its use of the state space to learn the desired task.

Table 4

1-tailed t-test: Significant differences were found between the **baseline** and **undo** conditions, in training sessions with nearly 100 non-expert human subjects playing the *Sophie’s Kitchen* game online.

Measure	Mean baseline	Mean undo	chg	t(96)	p
# states	48.3	42	13%	-2.26	=.01
# F	6.94	4.37	37%	-3.76	<.001
# F before G	6.4	3.87	40%	-3.7	<.001
# actions to G	208.86	164.93	21%	-2.25	=.01
# actions	255.68	224.2	12%	-1.32	=.095

11 Discussion

Robotic and software agents that operate in human environments will need the ability to learn new skills and tasks ‘on the job’ from everyday people. It is important for designers of learning systems to recognize that while the average consumer is not familiar with machine learning techniques, they are intimately familiar with various forms of social learning (e.g., tutelage, imitation, etc.). This raises two important and related research questions for the machine learning community. 1) How do people want to teach machines? 2) How do we design machines that learn effectively from natural human interaction? In this paper we have demonstrated the utility of a *socially guided machine learning* approach, exploring the ways machines can be designed to more fully take advantage of a natural human teaching interaction.

Our work emphasizes the *interactive* elements in teaching. There are inherently two sides to an interaction, in this case the human teacher and the machine learner. Our approach aims to enhance standard machine learning algorithms from both perspectives of this interaction: modifying the algorithm to build a better learning agent, and modifying the interaction techniques to provide a better experience for the human teacher. Understanding how humans want to teach is an important part of this process.

The scenario of human input has received some attention in the machine learning community (see Sec. 2 for several examples). Many prior works have addressed how human input can theoretically impact a learning algorithm or interaction. In contrast, this work addresses the nature of *real* people as teachers; our ground truth evaluation is the performance of the machine learner with non-expert human teachers. Whereas prior works typically lend control either to the machine or the human, our contribution is the focus on how a machine learner can use transparency behaviors to steer the instruction it receives from a human, creating more reciprocal control of the interaction.

Several prior works that utilize a human teacher are inspired by animal or human learning. For instance, game characters that the human player can shape through interaction have been successfully incorporated into a few computer games (Evans, 2002; Stanley et al., 2005; Stern et al., 1998). Breazeal *et al.* have demonstrated aspects of collaboration and social learning on a humanoid robot, using social cues to guide instruction (Breazeal et al., 2004). Animal training techniques and human tutelage have been explored in several robotic agents (Kaplan et al., 2002; Saksida et al., 1998; Steels and Kaplan, 2001; Lockerd and Breazeal, 2004). As a software agent example, Blumberg’s virtual dog character can be taught via clicker training, and behavior can be shaped by a human teacher (Blumberg et al., 2002).

Many of these works agree with our situated learning paradigm for machines, and have emphasized that an artificial agent should use social techniques to create a better interface for a human partner. This work goes beyond gleaning inspiration from natural forms of social learning/teaching to formalize this inspiration and empirically ground it in observed human teaching behavior through extensive user studies. One of the main contributions of this work is empirical evidence that social interaction is not only a good idea because it creates a good interface for a human partner, but also because it create a better learning environment and significant learning benefits for an agent.

Our findings indicate that a learning agent can take better advantage of the different kinds of messages a human teacher is trying to communicate. We find that, given a single communication channel, people have various communicative intents.

In addition to common instrumental feedback, *people assume they can guide the agent*, even when we explicitly told people that only feedback messages were supported. In their guidance communication, people mean to bias the action selection mechanism of the RL algorithm. When we allow this, introducing a separate interaction channel for attention direction and modifying the action selection mechanism of the algorithm, we see a significant improvement in the agent’s learning. The agent is able to learn tasks using fewer actions over fewer trials. It has a more efficient exploration strategy that wasted less time in irrelevant states. We argue that a less random and more sensible exploration will lead to more understandable and teachable agents. Guidance also led to fewer failed trials and less time to the first successful trial. This is a particularly important improvement in that it implies a less frustrating teaching experience, which in turn creates a more engaging interaction for the human.

We also see that players treat the agent as a social entity and want a *motivational channel of communication* to encourage it. This is seen despite the fact that the learning agent in this work is very mechanistic and simplistic.

One can assume that this effect will only be more prominent with characters that are explicitly designed to be socially and emotionally appealing. We argue that to build successful agents that learn from people, attention of the research community should focus on understanding and supporting the psychology and social expectations of the human teacher. It remains future work to explore how this motivational or encouragement channel of communication should influence the learning algorithm in a different way than the ordinary positive and negative feedback. Our hypothesis is that this communication is intended to influence the internal motivations, drives and goals of the agent.

This work offers a concrete example that the *transparency of the agent's behavior to the human can improve its learning environment*. In a social learning interaction both learner and teacher influence the performance of the tutorial dyad. While this observation seems straightforward in the human literature, little attention has been paid to the communication between human teacher and artificial agent in the traditional machine learning literature. Particularly, we believe that the transparency of the learner's internal process is paramount to the success of the tutorial dialog. Specifically, this work has shown that when the learning agent uses gazing behaviors to reveal its uncertainties and potential next actions, people were significantly better at providing more guidance when it was needed and less when it was not. Thus the agent, through its own behavior, was able to shape the human's input to be more appropriate. Gaze is just one such transparency device, the exploration of various devices and their relation to the learning process is part of our future work.

Additionally these transparency behaviors serve to boost the realism and believability of the character overall, thereby making it more engaging for the human. The creation of believable characters that people find emotionally appealing and engaging has long been a challenge (Thomas and Johnson, 1981; Bates, 1997). Autonomy complicates this goal further, since the character has to continually make action choices that are reasonable and useful as well as believable and engaging. Blumberg *et al.* has some of the most extensive work in this domain (Blumberg, 1997; Tomlinson and Blumberg, 2002) within a dog learning context. Thus another challenge for teachable characters is to be appropriately responsive to the human's instruction.

In this work we have studied one aspect of such responsiveness, informed by our initial user study. Negative feedback from a human teacher can be treated as both feedback for the action and suggestion to perform an UNDO behavior and *reverse the action if possible*. When this is part of the agent's behavior, learning is improved in both speed and efficiency.

We chose to use the Q-Learning algorithm for this work because it is standard and widely understood. This affords the transfer of these lessons and modifications to other reinforcement-based approaches. We have shown significant

improvements in an RL domain showing that learning in a situated interaction with a human partner can help overcome some of the well recognized problems of RL. Furthermore, these improvements in learning will contribute to higher quality interactive learning agents that are better equipped to take advantage of the ways that people naturally approach teaching.

12 Conclusion

This work shows that designing for the complete human-machine learning system creates a more successful robot learner. Our initial experiment with an interactive computer game character lead to three main findings: people assume they can guide the agent, they dynamically adjust their behavior as they develop a model of the agent, and they have a positive bias in their rewards.

We addressed these findings in four follow-up experiments with modified versions of the *Sophie's Kitchen* game. Our modifications include: an embellished channel of communication that distinguishes between guidance, feedback, and motivational intents; endowing the character with transparency behaviors that reveal specific aspects of the agent's learning process; and providing a more natural reaction to negative feedback. A series of user studies show that these empirically informed modifications result in several learning improvements across several dimensions including the speed of task learning, the efficiency of state exploration, the understandability of the agent's learning process for the human, and a significant drop in the number of failed trials encountered during learning.

Importantly, in this work we acknowledge that the ground truth evaluation for systems meant to learn from people is performance with non-expert humans. This topic deserves more attention from the machine learning community, as it will be important for progress towards a social learning scenario for machines. This series of experiments with an interactive learning agent illustrates the effectiveness of this approach for building machines that can learn from ordinary people. The ability to utilize and leverage social skills is far more than a nice interface technique. It can positively impact the dynamics of underlying learning mechanisms to show significant improvements in a real-time interactive learning session with non-expert human teachers.

Acknowledgements

This work is funded in part by the Media Lab corporate sponsors of the Things that Think and Digital Life consortia.

References

- Argyle, M., Ingham, R., McCallin, M., 1973. The different functions of gaze. *Semiotica*, 19–32.
- Arkin, R., Fujita, M., Takagi, T., Hasegawa, R., 2003. An ethological and emotional basis for human-robot interaction. In: *Proceedings of the Conference on Robotics and Autonomous Systems*.
- Bates, J., 1997. The role of emotion in believable agents. *Communications of the ACM* 37 (7), 122–125.
URL citeseer.ist.psu.edu/bates94role.html
- Blumberg, B., 1997. Old tricks, new dogs: ethology and interactive creatures. Ph.D. thesis, Massachusetts Institute of Technology.
- Blumberg, B., Downie, M., Ivanov, Y., Berlin, M., Johnson, M., Tomlinson, B., 2002. Integrated learning for interactive synthetic characters. In: *Proceedings of the ACM SIGGRAPH*.
- Breazeal, C., 2002. *Designing Sociable Robots*. MIT Press, Cambridge, MA.
- Breazeal, C., Brooks, A., Gray, J., Hoffman, G., Lieberman, J., Lee, H., Lockerd, A., Mulanda, D., 2004. Tutelage and collaboration for humanoid robots. *International Journal of Humanoid Robotics* 1 (2).
- Clouse, J., Utgoff, P., 1992. A teaching method for reinforcement learning. In: *Proc. of the Ninth International Conf. on Machine Learning (ICML)*. pp. 92–101.
- Cohn, D., Ghahramani, Z., Jordan, M., 1995. Active learning with statistical models. In: *Tesauro, G., Touretzky, D., Alspector, J. (Eds.), Advances in Neural Information Processing. Vol. 7. Morgan Kaufmann*.
- Evans, R., 2002. Varieties of learning. In: *Rabin, S. (Ed.), AI Game Programming Wisdom*. Charles River Media, Hingham, MA, pp. 567–578.
- Greenfield, P. M., 1984. Theory of the teacher in learning activities of everyday life. In: *Rogoff, B., Lave, J. (Eds.), Everyday cognition: its development in social context*. Harvard University Press, Cambridge, MA.
- Horvitz, E., Breese, J., Heckerman, D., Hovel, D., Rommelse, K., July 1998. The lumiere project: Bayesian user modeling for inferring the goals and needs of software users. In: *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence*. Madison, WI, pp. 256–265.
- Isbell, C., Shelton, C., Kearns, M., Singh, S., Stone, P., 2001. Cobot: A social reinforcement learning agent. *5th Intern. Conf. on Autonomous Agents*.
- Kaplan, F., Oudeyer, P.-Y., Kubinyi, E., Miklosi, A., 2002. Robotic clicker training. *Robotics and Autonomous Systems* 38(3-4), 197–206.
- Krauss, R. M., Chen, Y., Chawla, P., 1996. Nonverbal behavior and nonverbal communication: What do conversational hand gestures tell us? In: *Zanna, M.*

- (Ed.), *Advances in experimental social psychology*. Tampa: Academic Press, pp. 389–450.
- L. S. Vygotsky, E. M. C., 1978. *Mind in society: the development of higher psychological processes*. Harvard University Press, Cambridge, MA.
- Lashkari, Y., Metral, M., Maes, P., 1994. Collaborative Interface Agents. In: *Proceedings of the Twelfth National Conference on Artificial Intelligence*. Vol. 1. AAAI Press, Seattle, WA.
 URL citeseer.ist.psu.edu/lashkari94collaborative.html
- Lauria, S., Bugmann, G., Kyriacou, T., Klein, E., 2002. Mobile robot programming using natural language. *Robotics and Autonomous Systems* 38(3-4), 171–181.
- Lieberman, H. (Ed.), 2001. *Your Wish is My Command: Programming by Example*. Morgan Kaufmann, San Francisco.
- Lockerd, A., Breazeal, C., 2004. Tutelage and socially guided robot learning. In: *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*.
- Maclin, R., Shavlik, J., Torrey, L., Walker, T., Wild, E., July 2005. Giving advice about preferred actions to reinforcement learners via knowledge-based kernel regression. In: *Proceedings of the The Twentieth National Conference on Artificial Intelligence (AAAI)*. Pittsburgh, PA.
- Mataric, M., 1997. Reinforcement learning in the multi-robot domain. *Autonomous Robots* 4 (1), 73–83.
- Mitchell, T. M., Wang, S., Huang, Y., 2006. Extracting knowledge about users' activities from raw workstation contents. In: *Proceedings of the 21st National Conference on Artificial Intelligence (AAAI)*.
- Nicolescu, M. N., Matarić, M. J., July 2003. Natural methods for robot task learning: Instructive demonstrations, generalization and practice. In: *Proceedings of the 2nd Intl. Conf. AAMAS*. Melbourne, Australia.
- Saksida, L. M., Raymond, S. M., Touretzky, D. S., 1998. Shaping robot behavior using principles from instrumental conditioning. *Robotics and Autonomous Systems* 22(3/4), 231.
- Schaal, S., 1999. Is imitation learning the route to humanoid robots? *Trends in Cognitive Sciences* 3, 233242.
- Schohn, G., Cohn, D., 2000. Less is more: Active learning with support vector machines. In: *Proc. 17th ICML*. Morgan Kaufmann, San Francisco, CA, pp. 839–846.
- Smart, W. D., Kaelbling, L. P., 2002. Effective reinforcement learning for mobile robots. In: *In Proceedings of the IEEE International Conference on Robotics and Automation*. pp. 3404–3410.
- Stanley, K. O., Bryant, B. D., Miikkulainen, R., 2005. Evolving neural network agents in the nero video game. In: *Proceedings of IEEE 2005 Symposium on Computational Intelligence and Games (CIG'05)*.
- Steels, L., Kaplan, F., 2001. Aibo's first words: The social learning of language and meaning. *Evolution of Communication* 4 (1), 3–32.
- Stern, A., Frank, A., Resner, B., 1998. Virtual petz (video session): a hybrid approach to creating autonomous, lifelike dogz and catz. In: *AGENTS '98: Proceedings of the second international conference on Autonomous agents*. ACM Press, New York, NY, USA, pp. 334–335.
- Thomas, F., Johnson, O., 1981. *Disney Animation: The Illusion of Life*. Abbeville

- Press, New York.
- Thrun, S., 2002. Robotics. In: Russell, S., Norvig, P. (Eds.), *Artificial Intelligence: A Modern Approach* (2nd edition). Prentice Hall.
- Thrun, S. B., Mitchell, T. M., 1, 1993. Lifelong robot learning. Tech. Rep. IAI-TR-93-7.
- Tomlinson, B., Blumberg, B., 2002. Social synthetic characters. *Computer Graphics* 26 (2).
- Voyles, R., Khosla, P., 1998. A multi-agent system for programming robotic agents by human demonstration. In: *Proceedings of AI and Manufacturing Research Planning Workshop*.
- Watkins, C., Dayan, P., 1992. Q-learning. *Machine Learning* 8 (3), 279–292.
- Wertsch, J. V., Minick, N., Arns, F. J., 1984. Creation of context in joint problem solving. In: Rogoff, B., Lave, J. (Eds.), *Everyday cognition: its development in social context*. Harvard University Press, Cambridge, MA.