

Face valuing: Training user interfaces with facial expressions and reinforcement learning

Vivek Veeriah, Patrick M. Pilarski, Richard S. Sutton

Reinforcement Learning & Artificial Intelligence Lab

Department of Computing Science, University of Alberta

{vivekveeriah, pilarski, rsutton}@ualberta.ca

Abstract

An important application of interactive machine learning is extending or amplifying the cognitive and physical capabilities of a human. To accomplish this, machines need to learn about their human users' intentions and adapt to their preferences. In most current research, a user has conveyed preferences to a machine using explicit corrective or instructive feedback; explicit feedback imposes a cognitive load on the user and is expensive in terms of human effort. The primary objective of the current work is to demonstrate that a learning agent can reduce the amount of explicit feedback required for adapting to the user's preferences pertaining to a task by learning to perceive a value of its behavior from the human user, particularly from the user's facial expressions—we call this *face valuing*. We empirically evaluate face valuing on a grip selection task. Our preliminary results suggest that an agent can quickly adapt to a user's changing preferences with minimal explicit feedback by learning a value function that maps facial features extracted from a camera image to expected future reward. We believe that an agent learning to perceive a value from the body language of its human user is complementary to existing interactive machine learning approaches and will help in creating successful human-machine interactive applications.

1 Introduction

One important objective of human-machine interaction is to augment existing human capabilities, which requires machines and their human users to closely collaborate and form a productive partnership. To achieve this, it is crucial for the machines to learn interactively from their users, specifically their intents and preferences. In current research trends, the user's preferences are conveyed via explicit instructions (Kuhlmann et al., 2004) or expensive corrective feedback (Knox & Stone, 2015)—which can be in the form of pre-defined words or sentences, push buttons, mouse clicks etc. In many real-world, ongoing scenarios, these methods of feedback impose a cognitive load on human users. Moreover, in complex domains like prosthetic limbs, it is demand-

ing for the user to provide these kinds of explicit feedback. It is important to have an alternative approach that is both scalable and would allow the machines to learn their human users' intents and preferences via ongoing interactions.

In this paper, we explore the idea that a reinforcement learning agent can learn a value function that relates a user's body language, specifically from the user's facial expressions, to expectations of future reward. The agent can use this value function to adapt its actions to a user's preferences quickly with minimal explicit feedback. This approach is analogous to an agent learning to understand the body language of its human user. It could also be imagined as building a form of communicative capital between a human user and a learning agent (c.f., Pilarski et al., 2015). Learning from interactions with a human user tend to be continual; reinforcement learning methods are therefore naturally suited for this purpose.

To the best of our knowledge, our system is the first to learn a value function in real-time for a user's body language, specifically a value function that relates future reward to the facial features of the user. Additionally, this work is the first example of how such a value function can be used to guide the learning process of an agent interacting with a human user. Importantly, our approach does not utilize explicit reward channels, for example those discussed by Thomaz & Breazeal (2012) and by Knox et al. (2009). As it operates in real time, we believe that our approach is well suited for realistic human-machine interaction tasks and complements existing interactive machine learning approaches. Learning a language between an agent and its user in the form of value functions represents a new and powerful class of human-machine interaction technologies, and we expect the approaches discussed in this work will have broad applicability in many different real-world domains.

2 Related Methods

Significant research effort has been directed toward creating successful human-robot partnerships (e.g., as summarized in Knox & Stone, 2015; Mead et al., 2013; Breazeal et al., 2012; Pilarski & Sutton, 2012; Edwards et al., 2015). A natural approach is for an agent to learn from ongoing interactions with a human user via *human-delivered rewards*. Research by Thomaz & Breazeal (2008), Knox & Stone (2009), Breazeal et al. (2012), Loftin et al. (2015), and Peng et al. (2016) adopts this perspective, and it has been extensively studied

in recent work by Knox & Stone (2015). In the TAMER approach of Knox and Stone (2015), a system was able to learn a predictive model of a human user’s shaping rewards, such that the model could be used to successfully train an agent even in the presence of human-related delays and inconsistencies. As a potential drawback of learning a reward model, when the user needs to modify the agent’s behavior, the model would have to be changed (e.g., via additional shaping rewards from the user). We desire a method for delivering feedback that does not require a large number of costly interactions from the human, and that transfers well to new or changed situations without the need for retraining.

Another interesting approach to the interactive instruction of a machine learner involves a human teaching a robot to perform a task through demonstrations, a process aptly named as *learning from demonstration*. This approach can also be called programming by demonstration. There are numerous works exploring learning from demonstration (e.g., Koenig & Mataric, 2012; Schulman et al., 2013; Alizadeh et al., 2014). One noted downside is that this form of learning is reported to be at times a tiring experience for a human user. Many approaches are also limited in their ability to scale up to a full range of real-world tasks (e.g., it is impossible to tractably provide demonstrations or instructions covering all possible situations).

A key difference between many existing methods and our approach is that we are concerned with designing a general, scalable approach that would allow an agent to adapt its behavior to a user’s changing preferences with minimal explicit human-generated feedback. This is in contrast to approaches that seek to use body language like facial features as an input channel to directly control a robot or other machine’s operation (e.g., Breazeal (1998) and Liu & Picard (2003)). As a significant contribution of the present work, we describe the use of facial features not as a channel of control but as a means of valuation. To this end, we propose to learn a value function that is grounded in the user’s body language, independent of the features of a task, and to use this value function to help influence an agent’s real-time decision making in a way that spans multiple tasks and settings of use.

3 Reinforcement Learning

In a reinforcement learning setting (Sutton & Barto, 1998), a learning agent interacts with an unknown environment in order to achieve a goal. In this setting, the goal is to maximize the cumulative reward accumulated by adapting its behavior within the environment.

Markov Decision Processes (MDPs) are the mathematical notations used for formalizing a reinforcement learning problem. An MDP consists of a tuple $\langle \mathcal{S}, \mathcal{A}, p, r, \gamma \rangle$, consisting of \mathcal{S} , set of all states; \mathcal{A} , set of all actions; $p(s'|s, a)$, a transition probability function, which gives the probability of transition to state $s' \in \mathcal{S}$ at the next time-step given for the current state $s \in \mathcal{S}$ and action $a \in \mathcal{A}$; $r(s, a, s')$, the reward function that gives the expected reward for a transition from state $s \in \mathcal{S}$ to $s' \in \mathcal{S}$ by taking action $a \in \mathcal{A}$; γ is the discount factor, that specifies the relative importance between immediate and long term rewards. In episodic prob-

lems, the MDP can be viewed as having special states called *terminal states*, which terminate an episode. Such states ease the mathematical notations as they could be viewed as a single state with single action that results in a reward of 0 and transition to itself. The return at a time instance t is defined as the discounted sum of rewards after time t :

$$G_t = R_{t+1} + \gamma R_{t+2} + \gamma^2 R_{t+3} + \dots$$

$$G_t = \sum_{i=1}^{\infty} \gamma^{i-1} R_{t+i}$$

where R_{t+1} denotes the reward received after taking an action A_t in state S_t .

Actions are taken at discrete time steps $t = 0, 1, 2, \dots$ according to a *policy* $\pi : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ which defines a selection probability for each action conditioned on the state. Each policy π has a corresponding state-value function $v_\pi(s)$, that maps each state $s \in \mathcal{S}$ to the expected return G_t from that state by following the policy π ,

$$v_\pi(s) = \mathbf{E}\{G_t | S_t = s, \pi\}$$

The state-value functions are significant when the given task requires prediction. On the contrary, if the given task requires control, then it is important to use the action-value functions $q_\pi(s, a)$ which gives the expected return G_t by taking an action a from state s and then following the policy π :

$$q_\pi(s, a) = \mathbf{E}\{G_t | S_t = s, A_t = a, \pi\}$$

4 Grip Selection Task

To evaluate the face valuing approach, we introduce a grip selection task that was inspired by a natural problem in a prosthetic arm domain where the agent needs to select an appropriate grip pattern for grasping a given object. The task consists of a set of n grips and m objects. Depending on the experiment, there could be many possible grips for a given object, with the correct grip being defined according to the user’s preference. The task could also consist of uncountable number of objects, making grip selection with pure trial-and-error a non-trivial problem.

This task was formulated as an undiscounted episodic MDP with a reward of 0 for every time step, and with 0 reward for completing an episode. At the beginning of each episode, a single object was randomly picked from a large set of objects and the agent was tasked with choosing a grip from a limited set of grips; once the agent selected a grip it needed to move a fixed number of steps towards the object to finish the grasping motion, thereby completing an episode. A human user provided reward to the agent by pressing a single button which delivered a reward of -5 for the corresponding time step. Moreover, pressing the button forced the agent to return back to its initial position regardless of its current position. The experimental setup is shown in Fig. 1.

As in real-world grasping tasks, a user could have personal preferences over which grip was suitable to grasp a given object. These preferences could change from episode to episode

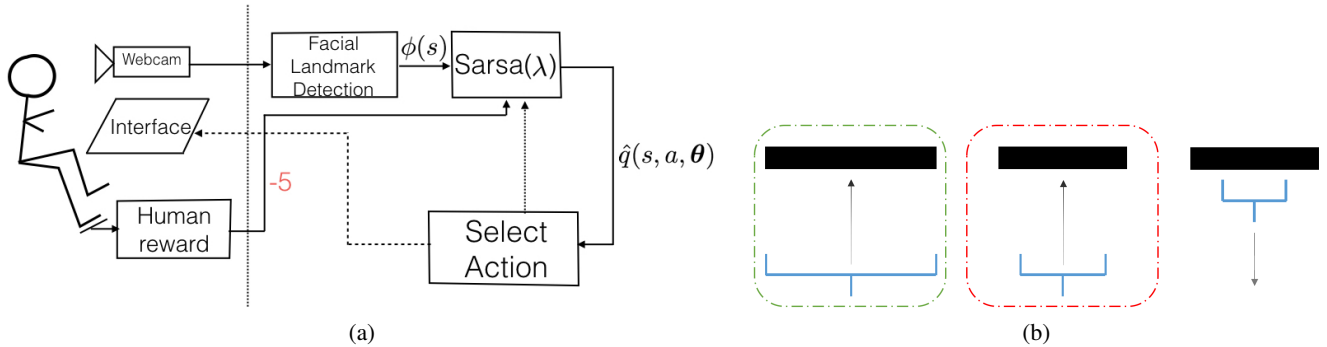


Figure 1: (a) Overview of the face valuing agent: the user observes the simulated task and can provide a negative reward to the SARSA agent during the task. The agent can observe user’s face via a webcam, and learns to behave based on the user’s preferences. (b) Grip selection task: Solid colored box denotes the object; the thin blue lined object denotes the grip selected by an agent. The initial position of the agent is called “grip-changing station”. The agent has to learn to select, based on the user’s preference, an appropriate grip — the grip’s width needs to be equal to the object’s width. The figure on the left denotes one correct combination while the figure on the center denotes an incorrect grip for the given object. The figure on the right denotes a scenario where the agent is forced to return to its initial position, as the user has pushed the reward button. To complete an episode, the agent has to select one of the correct grips; move forward and grasp the object.

and from experiment to experiment. Further, these preferences were hidden from the learning agent and the only way the agent could infer this is from the changing facial expressions of the human user. Therefore, in this work, we asked the human user to be as expressive as possible so as to provide clear cues for the learning agent to begin forming its behavioral choices.

5 Experimental Setup

For our experiments, two Sarsa(λ) (Rummery & Niranjan, 1994; Sutton & Barto, 1998) agents (one that uses face valuing and one without face valuing) are compared on the above described task. All the experimental results in this paper were performed by a well-trained user in a blind setting, i.e., the user did not know which of the two methods was currently being evaluated. The user provided the same form of rewards to both learning agents via their button pushes. The two main instructions we gave to the human user were 1) to express their pleasure or displeasure with the agent via any simple, repeatable, and minutely distinguishable facial expressions, and 2) to push the button whenever the learning agent was not behaving as per the user’s expectations.

5.1 State space

Agent	State space
w/o face valuing	current grip and object ids + bias term
w/ face valuing	23 feature points + bias term

Table 1: State space of the agents compared in the experiments are displayed.

The state spaces for both the agents are briefly described in Table 1. The agent without face valuing has the id of the current grip and id of the current object in its state space along with a bias term. The id of the current grip chosen by the

agent is one-hot encoded to form vector of length n . Similarly, the id of the object is also one-hot encoded to a vector of length m and concatenated with the vector corresponding to the current grip. So, the entire state space for this method is of length $m + n + 1$ where m is the total number of objects during the entire experiment and n is the total number of grips available to the agent during an experiment.

For the face valuing agent, 68 key points from a frame containing a human’s face are detected through a popular facial landmark detection algorithm (Kazemi et al., 2014). These key points are simple two dimensional coordinates that denote the position of certain special locations of a human’s face. These points from each frame were normalized and 23 points that correspond to the positions of eye brows and mouth of a human’s face were selected. Each of these 23 points, were tile-coded with 4 tilings and each tiling of size 10×10 resulting in a feature vector of size 9200. These key points seems to produce sufficient variations between different facial expressions.




Figure 2: Features extracted from face: facial landmarks corresponding to eye brows and mouth of a user’s face are the only features extracted from the face. These are marked as yellow circles.

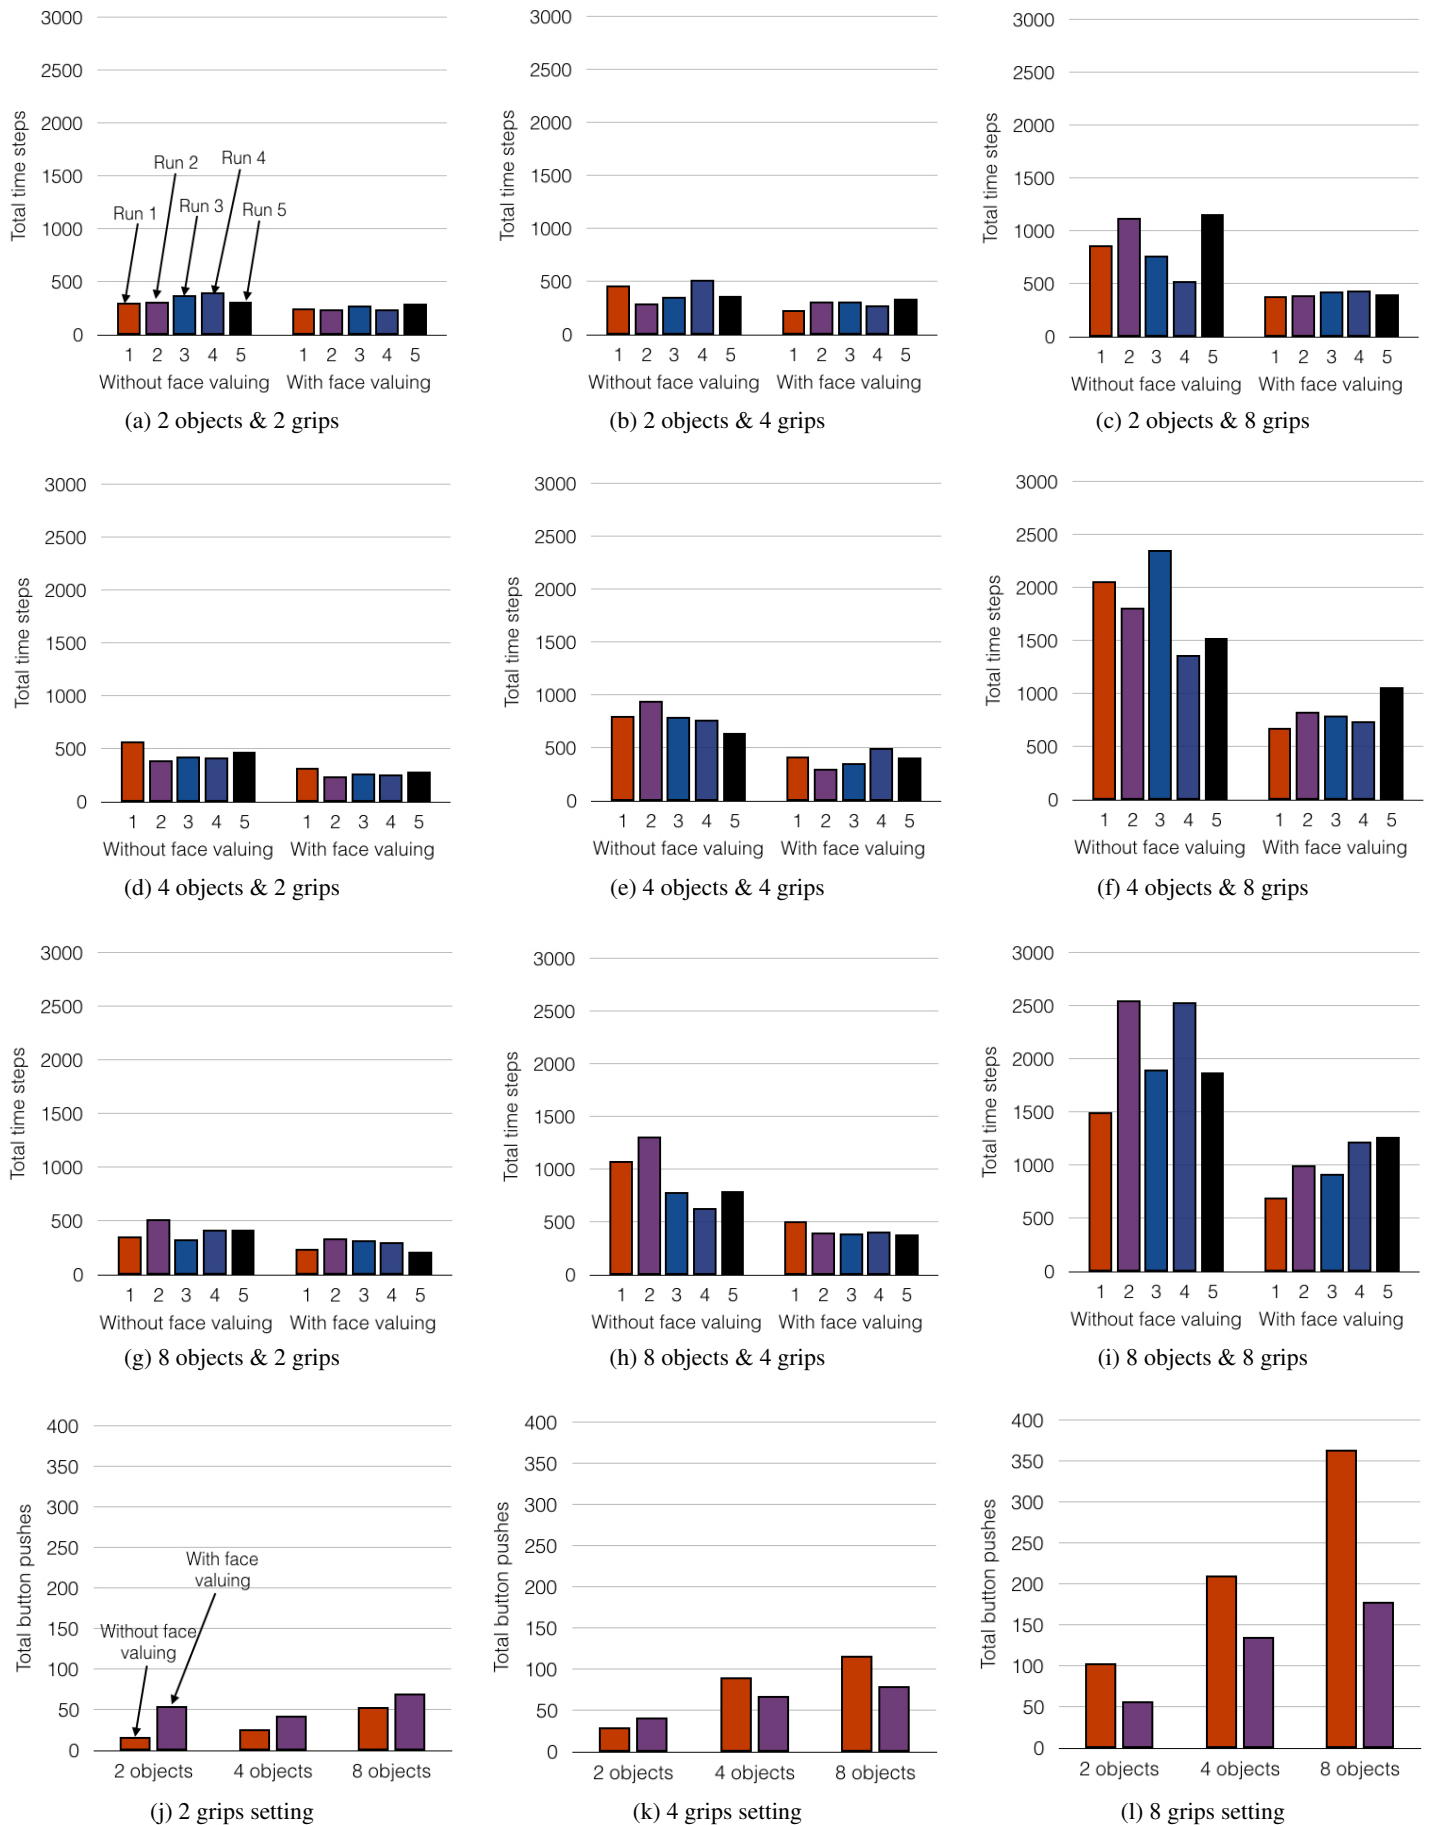


Figure 3: Total time steps taken for different grip and object settings: the face-valuing agent learned to adapt faster compared to a conventional agent in most settings. Moreover, in difficult settings where there are more number of appropriate object-grip combinations, the face-valuing agent required less human generated rewards for achieving this level of performance. Each plot was generated by averaging data obtained from 5 independent runs and each run consisted of 15 episodes.

5.2 Action space

The complete action space for both the agents consisted of $\{grip_1, grip_2, \dots, grip_n, \uparrow, \downarrow\}$ actions where the first n actions implied in choosing that particular grip. The remaining two actions were the move one step forward towards the object and move one step back towards the grip-changing station. The actions available to the agent depended on its position relative to the object and the grip-changing station. When the agent was in the grip-changing station, the available actions were $\{grip_1, grip_2, \dots, grip_n, \uparrow\}$ whereas when the agent had left the grip-changing station, only the $\{\uparrow, \downarrow\}$ actions were available. When the user pushed the reward button, the agent lost all its actions except $\{\downarrow\}$ until it reached the grip-changing station.

The agent observed the state space once every *one-tenth of a second* and had to take an action on every time step. The agent, however, had the freedom of choosing the same action for many consecutive time steps which allowed the human user to expressively respond to the learning agent.

6 Experiments

6.1 Experiment 1: Different object-grip settings

The first experiment compared the two agents with multiple grip & object settings. The plots of total time steps and total human generated rewards accumulated by the agents are shown in Fig. 3. The plots (Fig. 3 (a), (b), \dots , (i)) represent the total time taken by a learning agent to complete a successful grasp across episodes. The plots (Fig. 3 (j), (k), (l)) display the total number of human generated rewards given to an agent to successfully adapt to user’s preferences. These graphs (Fig. 3) were generated from the same user experiments conducted in a blind manner. A perfect agent would have no human generated reward in all these settings and would take only 11 time steps to complete an episode.

From the plots (Fig. 3 (a), (b), \dots , (i)), it can be observed that the agent with face valuing quickly adapted with the user’s preferences in all the different object and grip settings. Also, from the plots (Fig. 3 (j), (k), (l)), the total number of human generated rewards for the face valuing agent was comparatively lower than the agent without face valuing in the difficult settings of this experiment.

During the initial phase of the experiment, the face valuing agent utilized the human generated reward to learn to perceive a value of its actions from the human user’s face. This learned value was leveraged to adapt the agent’s actions in later phases of the experiment. In simple experiment settings where there are fewer number of object-grip combinations, like the 2 grips experiment setup, the agent without face valuing could quickly learn the appropriate behavior from button pushes provided by the user and this resulted in a better performance compared to an agent with face valuing. However, in setups with large number of possible combinations of grips and objects, the agent without face valuing lost this advantage and failed to scale up. The face-valuing agent performed comparatively better in these scenarios as it learned to perceive a “goodness” of its actions which guided the agent in choosing the correct action at a given time instance.

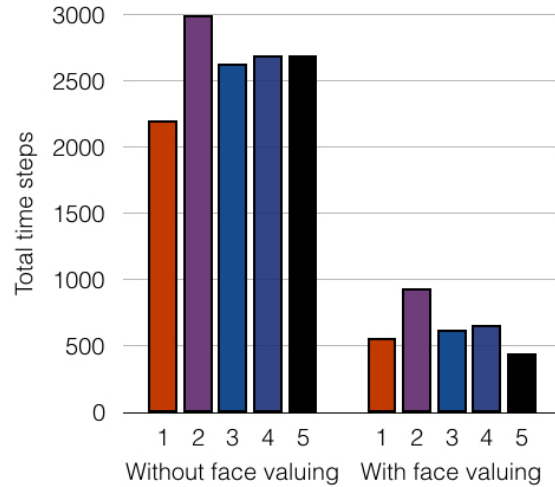


Figure 4: Experiment with Infinite objects & 5 grips: Plot of total time steps for limited grips and infinite objects setting. An experiment setting where a new object is introduced every episode and the agent needs to grasp this object from one of its grip. Each bar represents the total time steps taken by the agent to complete 15 episodes of this task. The plot was generated from data obtained from 5 independent runs. This is a key result of our approach as it clearly shows the performance improvement obtained through face valuing.

For the agent without face valuing, the user’s preferences could be communicated only through the reward channel. Naturally, this approach was more expensive in terms of the number of manual rewards compared to the face-valuing approach. On the other hand, the face-valuing approach utilized the user’s facial features, which conveyed the preferences over the grips—a simple approach observed in the user was to for them to have a neutral or a sad expression when the agent was not selecting the correct grip and to display a positive expression when the agent selected the correct grip. Interestingly, by learning a value function over the face, this agent learned to wait for an affirmative expression from the human user before moving forward to grasp the object. When there were no such expressions from the user, the agent switched from one grip to another until the user gave a “go ahead” expression.

6.2 Experiment 2: Infinite objects and finite grips

A second experiment (Fig. 4) showed the performance improvement obtained through our approach in a different and more difficult setting: one where a new object was generated for each episode and no object was ever seen more than once by the agent throughout the experiment. This experiment therefore explored the ability of face valuing to address new or changed tasks, and highlighted the importance of adapting quickly to a user’s preferences.

The Fig. 4 denotes the total time steps taken by the agents to complete this experiment whereas the Table 2 shows the total number of instances of human explicit feedback required by the agent to successfully complete this task. Data was generated from experiments with a single user.

From the plot in Fig. 4, it can be observed that the face-

valuing agent was much quicker in adapting to the user’s preferences. It learned to complete the task quicker than an agent without face valuing. From the Table 2, it is clear that that total number of instances of human generated feedback to complete this task was less for the face-valuing agent.

Since a new object was introduced in every episode, the agent without face valuing could not learn the possible grip/object combinations only from human-generated rewards. This was the cause for it requiring more human feedback in completing this task. In the face-valuing approach, as the learning agent relied on values related to facial features, it could adapt easily in these situations. Effectively, the agent with face valuing learned to keep switching the grips periodically until the user gave a “go ahead” expression. Unfortunately, the agent without face valuing did not have this advantage and could not perform effectively in this setting.

Agent	Total no. of button pushes
w/o face valuing	623.6
w/ face valuing	137.4

Table 2: Total number of button pushes provided for adapting an agent to the user’s preference in experiment 2. The values represent the total number of button pushes given by the human user to an agent for the complete experiment setting that lasted 15 episodes. Average obtained from 5 independent runs.

7 Discussion

In our experiments, the learning agent with face valuing had the ability to perceive a human user’s face and, eventually, learned to perceive a value of its behavior from its user’s facial expressions. Our preliminary results therefore suggest that, by learning to value a human user’s facial expressions, the agent could adapt quickly to its user’s preferences with minimal explicit corrective feedback. This learning occurred as follows: during the initial phase of the experiment, the agent used the explicit corrective feedback to learn a value function from the user’s facial gestures; these gestures served as useful clues about future rewards based on the agent’s current behavior, and guided its behavior.

Several studies have shown that users, to a certain extent, are willing to teach machines to perform tasks automatically. For example, in medical domains, it is already common for people with amputations to extend their capabilities or limitations through a partnership with machines (Williams, T. W., 2011). However, currently available technology does not identify and adapt quickly to the different preferences of their users; this is a serious bottleneck to intelligence or physical amplification in human-machine partnerships. Our work helps begin to address these limitations.

For evaluating our approach, we introduced a grip selection task wherein the learning agent had to figure out the goal through its interaction with the user; this agent can be readily termed a *goal-seeking agent* (Pilarski et al., 2015). To demonstrate the significance of our approach, we performed two sets of experiments: the first one involved multiple object-grip settings on the grip selection task; we termed the second

experiment as the *infinite objects setting*, because one new object was generated for every episode and the agent needed to grasp this object by selecting one grip from its limited set of grips. This infinite objects settings is pertinent to real-world scenarios, where there are uncountable number of objects which can be grasped from a limited set of grip patterns.

The results from the first user experiment (Fig. 3 (a), (b), ..., (i)) suggest that the face-valuing agent learns to adapt quicker to its user’s preferences in this task. From the plots in Fig. 3 (j), (k), (l), it can be observed that the face-valuing agent learns to adapt to its user’s preferences with significantly lesser number of explicit human generated feedback signals, particularly in difficult experiment settings. From the second user experiment (Fig. 4), we empirically show a scenario where conventional methods can fail. From both the experiments, it can be observed that the face-valuing agent successfully learns to adapt and completes one episode after another by relying only on facial expressions, specifically the value learned from facial expressions. On the other hand, the agent without face valuing could rely only on human generated feedback for identifying the correct grip for a given object, which resulted in a greater number of button pushes being given by the user. Moreover, we observed that the face valuing agent learned to wait for an affirmative facial expression before moving towards the object. Otherwise, the agent would switch from one grip to another at the grip-changing station until the user provided an affirmative expression.

Though our experiments were simulated, we believe that our approach can be much more valuable in a realistic robot setting—we expect a robot’s behavior would elicit more expressive facial feedback from the user than our simple simulated domain, and thus more powerful features for a face-valuing agent. In a robot setting, the user can observe the robot’s actions and their consequences in a real-world environment, where it is natural for the user to implicitly emote through various facial cues. Robotic experiments are needed to help quantify the advantage of a face-valuing approach to human-machine interaction.

8 Conclusions

We introduced a new and a promising approach called *face valuing* for adapting an agent to a user’s preferences, and showed that it can produce large performance improvements over a conventional agent that learns only from human-generated rewards. By allowing the agent to perceive a value from a user’s facial expressions, the total number of expensive human generated-rewards delivered during a task was substantially reduced and the agent was quickly able to adapt to its user’s preferences. Face valuing learns a mapping from facial expression to user satisfaction; it formalizes satisfaction as a value function and learns this value function through temporal-difference methods. Most work on the use of facial features in human-machine interaction uses facial features as control signals for an agent; surprisingly, our work seems to be the first to use facial expressions to instead train a learning system. Face valuing is general and largely task agnostic, and we believe it will therefore extend well to other settings and other forms of human body language.

References

- Alizadeh, T., Calinon, S. and Caldwell, D.G., 2014, May. Learning from demonstrations with partially observable task parameters. In *Robotics and Automation (ICRA), 2014 IEEE International Conference on* (pp. 3309-3314). IEEE.
- Breazeal, C., 1998, May. Regulating human-robot interaction using emotions, drives, and facial expressions. In *Proceedings of Autonomous Agents* (Vol. 98, pp. 14-21).
- Breazeal, C., Love, B.C. and Mooney, R.J., 2012. Learning from human-generated reward.
- Edwards, A.L., Dawson, M.R., Hebert, J.S., Sherstan, C., Sutton, R.S., Chan, K.M. and Pilarski, P.M., 2015. Application of real-time machine learning to myoelectric prosthesis control: A case series in adaptive switching. *Prosthetics and orthotics international*, p.0309364615605373.
- Kazemi, V. and Sullivan, J., 2014. One millisecond face alignment with an ensemble of regression trees. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 1867-1874).
- Knox, W.B., Fasel, I. and Stone, P., 2009, March. Design Principles for Creating Human-Shapable Agents. In *AAAI Spring Symposium: Agents that Learn from Human Teachers* (pp. 79-86).
- Knox, W.B., Stone, P. and Breazeal, C., 2013, March. Teaching agents with human feedback: a demonstration of the tamer framework. In *Proceedings of the companion publication of the 2013 international conference on Intelligent user interfaces companion* (pp. 65-66). ACM.
- Knox, W.B. and Stone, P., 2015. Framing reinforcement learning from human reward: Reward positivity, temporal discounting, episodicity, and performance. *Artificial Intelligence*, 225, pp.24-50.
- Koenig, N.P. and Mataric, M.J., 2012, October. Training Wheels for the Robot: Learning from Demonstration Using Simulation. In *AAAI Fall Symposium: Robots Learning Interactively from Human Teachers*.
- Kuhlmann, G., Stone, P., Mooney, R. and Shavlik, J., 2004, July. Guiding a reinforcement learner with natural language advice: Initial results in RoboCup soccer. In *The AAAI-2004 workshop on supervisory control of learning and adaptive systems*.
- Liu, K. and Picard, R.W., 2003, April. Subtle expressivity in a robotic computer. In *CHI 2003 Workshop on Subtle Expressiveness in Characters and Robots*.
- Loftin, R., Peng, B., MacGlashan, J., Littman, M.L., Taylor, M.E., Huang, J. and Roberts, D.L., 2016. Learning behaviors via human-delivered discrete feedback: modeling implicit feedback strategies to speed up learning. *Autonomous Agents and Multi-Agent Systems*, 30(1), pp.30-59.
- Mead, R., Atrash, A. and Mataric, M.J., 2013. Automated proxemic feature extraction and behavior recognition: Applications in human-robot interaction. *International Journal of Social Robotics*, 5(3), pp.367-378.
- Peng, B., MacGlashan, J., Loftin, R., Littman, M.L., Roberts, D.L. and Taylor, M.E., 2016, May. A Need for Speed: Adapting Agent Action Speed to Improve Task Learning from Non-Expert Humans. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems* (pp. 957-965). International Foundation for Autonomous Agents and Multiagent Systems.
- Pilarski, P.M. and Sutton, R.S., 2012, October. Between Instruction and Reward: Human-Prompted Switching. In *AAAI Fall Symposium: Robots Learning Interactively from Human Teachers*.
- Pilarski, P.M., Sutton, R.S. and Mathewson, K.W., 2015. Prosthetic Devices as Goal-Seeking Agents.
- Rummery, G.A. and Niranjan, M., 1994. On-line Q-learning using connectionist systems.
- Schulman, J., Ho, J., Lee, C. and Abbeel, P., 2013. Learning from demonstrations through the use of non-rigid registration. In *Proceedings of the 16th international symposium on robotics research (ISRR)*.
- Sutton, R.S. and Barto, A.G., 1998. Introduction to reinforcement learning (Vol. 135). Cambridge: MIT Press.
- Thomaz, A.L. and Breazeal, C., 2008. Teachable robots: Understanding human teaching behavior to build more effective robot learners. *Artificial Intelligence*, 172(6), pp.716-737.
- Williams, T. W. (2011). Guest editorial: progress on stabilizing and controlling powered upper-limb prostheses. *Journal of Rehabilitation Research and Development* 48(6): ix-xix.