Facial feedback for reinforcement learning: a case study and offline analysis using the TAMER framework

Guangliang Li¹ · Hamdi Dibeklioglu² · Shimon Whiteson³ · Hayley Hung⁴

Published online: 12 February 2020
© The Author(s) 2020

Abstract

Interactive reinforcement learning provides a way for agents to learn to solve tasks from evaluative feedback provided by a human user. Previous research showed that humans give copious feedback early in training but very sparsely thereafter. In this article, we investigate the potential of agent learning from trainers' facial expressions via interpreting them as evaluative feedback. To do so, we implemented TAMER which is a popular interactive reinforcement learning method in a reinforcement-learning benchmark problem—Infinite Mario, and conducted the first large-scale study of TAMER involving 561 participants. With designed CNN–RNN model, our analysis shows that telling trainers to use facial expressions and competition can improve the accuracies for estimating positive and negative feedback using facial expressions. In addition, our results with a simulation experiment show that learning solely from predicted feedback based on facial expressions is possible and using strong/effective prediction models or a regression method, facial responses would significantly improve the performance of agents. Furthermore, our experiment supports previous studies demonstrating the importance of bi-directional feedback and competitive elements in the training interface.

Keywords Reinforcement learning · Facial expressions · Human agent interaction · Interactive reinforcement learning

✉ Guangliang Li
guangliangli@ouc.edu.cn

Hamdi Dibeklioglu
dibeklioglu@cs.bilkent.edu.tr

Shimon Whiteson
shimon.whiteson@cs.ox.ac.uk

Hayley Hung
h.hung@tudelft.nl

¹ Ocean University of China, Qingdao, China

² Bilkent University, Ankara, Turkey

³ University of Oxford, Oxford, UK

⁴ Delft University of Technology, Delft, The Netherlands

1 Introduction

Socially intelligent autonomous agents have the potential to become our high-tech companions in the family of the future. The ability of these intelligent agents to efficiently learn from non-technical users to perform a task in a natural way will be key to their success. Therefore, it is critical to develop methods that facilitate the interaction between these non-technical users and agents, through which they can transfer task knowledge effectively to such agents.

Interactive reinforcement learning has proven to be a powerful technique for facilitating the teaching of artificial agents by their human users [17, 22, 49]. In interactive reinforcement learning, an agent learns from human reward, i.e., evaluations of the quality of the agent's behavior provided by a human user, in a reinforcement learning framework. Compared to learning from demonstration [1], interactive reinforcement learning does not require the human to be able to perform the task well herself; she needs only to be a good judge of performance. Nonetheless, agent learning from human reward is limited by the quality of the interaction between the human trainer and agent.

From the human user's point of view, humans may get tired of giving *explicit feedback* (e.g., button presses to indicate positive or negative reward) as training time progresses. In fact, several TAMER studies—a popular interactive reinforcement learning method for enabling autonomous agents to learn from human reward [22], have shown that humans give copious feedback early in training but very sparsely thereafter [21, 28]. Instead of button presses on which TAMER relies, facial expressions have been often used by humans to consciously or subconsciously encourage or discourage specific behaviors they want to teach, e.g., smiling to indicate good behavior and frowning to indicate bad behavior [50]. Therefore, in our study, we investigate the potential of using facial expressions as reward signals.

To examine this potential, we conducted the first large-scale study of TAMER by implementing it in the Infinite Mario domain. Our study, involving 561 participants, at the NEMO science museum in Amsterdam using museum visitors (aged 6–72). We recorded the facial expressions of all trainers during training and, in some conditions, told participants that their facial expressions would be used as encouraging explicit feedback, e.g., happy and sad expressions would map to positive and negative reward respectively, in addition to keypresses, to train the agent. However, due to the significant challenge of processing facial expressions sufficiently accurately online and in real time in a fairly unconstrained non-laboratory setting, only keypresses were actually used for agent learning.

In our experiment, we test two independent variables: 'facial expression'—whether trainers were told that their facial expressions would be used in addition to keypresses to train the agent, and 'competition'—whether the agent will inform competitive feedback to the trainer. The main idea of the facial expression condition is to examine the effect that the additional modality of facial expressions could have on the agent's learning and the relationship between the reward signal and the nature of facial expressions. In addition, factors like environmental stress from outside environment might affect the expressiveness of facial expressions, which might have a further impact on the prediction accuracy of the model trained based on these facial feedback. And our previous research [30, 31] showed that if an agent informs the trainer socio-competitive feedback, the trainer will provide more feedback and the agent will ultimately perform better. In this study, we want to see how an agent's competitive feedback will affect the trainer's facial expressiveness and the agent's learning, especially when it is coupled with facial expression condition.

We investigate how ‘facial expression’ and ‘competition’ affect the agent’s learning performance and trainer’s facial expressiveness in four experimental conditions in our study: the *control condition*—without ‘competition’ or ‘facial expression’, the *facial expression condition*—without ‘competition’ but with ‘facial expression’, the *competitive condition*—with ‘competition’ but without ‘facial expression’, and the *competitive facial expression condition*—with both. We hypothesize that ‘competition’ will result in better performing agents, and ‘facial expression’ will result in worse agent performance.

Our preliminary work on this topic was presented in [26, 27]. This article significantly extends upon our initial work by providing a more extensive analysis of participants’ facial feedback and testing the potential of agent learning from them. The experimental results in this article show that telling trainers to use facial expressions makes them inclined to exaggerate their expressions, resulting in higher accuracies for estimating their corresponding positive and negative feedback keypresses using facial expressions. Moreover, competition can also elevate facial expressiveness and further increase the predicted accuracy. Furthermore, with designed CNN–RNN model, our results in a simulation experiment show that it is possible for an agent to learn solely from predicted evaluative feedback based on facial expressions. Our results also indicate that using strong/effective prediction models or regression models, facial responses would significantly improve the performance of agents. To our knowledge, it is the first time facial expressions have been shown to work in TAMER, opening the door to a much greater potential for learning from human reward in more natural, personalized and possibly more long term learning scenarios.

The rest of this article starts with a review of the related work in Sect. 2. Section 3.1 presents an introduction on interactive reinforcement learning. In Sect. 3.2 we provide the background and details about TAMER framework and Sect. 3.3 describes the Infinite Mario domain we used in our user study and the implemented representation of it for TAMER agent learning. Section 4 describes the experimental setup and Sect. 5 describes the proposed experimental conditions. Section 6 reports and discusses the experimental results. Section 7 discusses the open questions for learning from facial expressions. Finally, Sect. 8 concludes.

2 Related work

Our work contributes to a growing literature on interactive reinforcement learning, which deals with how an agent should learn the behavior from reward provided by a live human trainer rather than from the usual pre-coded reward function in a reinforcement learning framework [17, 22, 40, 45, 48]. Reward provided by a live human trainer is termed “human reward” and reward from the usual pre-coded reward function is termed “environmental reward” in reinforcement learning. In interactive reinforcement learning, a human trainer evaluates the quality of an agent’s behavior and gives the agent feedback to improve its behavior. This kind of feedback can be restricted to express various intensities of approval and disapproval and mapped to numeric “reward” for the agent to revise its behavior. We will provide details about interactive reinforcement learning in Sect. 3. In this section, we will review literatures on reinforcement learning from human reward, and machine learning systems or agents learning from facial expressions.

2.1 Reinforcement learning from human reward

Clicker training [3] is a related concept that involves using only positive reward to train an agent. Isbell et al. [17] developed the first software agent called Cobot that learns from both reward and punishment by applying reinforcement learning in an online text-based virtual world where people interact. The agent learns to take proactive verbal actions (e.g. proposing a topic for conversation) from ‘reward and punish’ text-verbs invoked by multiple users. Later, Thomaz and Breazeal [49] implemented an interface with a tabular *Q-learning* [53] agent. In their interface, a separate interaction channel is provided to allow the human to give the agent feedback. The agent aims to maximize its total discounted sum of human reward and environmental reward. They treat the human’s feedback as additional reward that supplements the environmental reward. Their results show an improvement in agent’s learning speed with additional human reward. In addition, the work of Thomaz and Breazeal shows that by allowing the trainer to give action advice on top of human reward, the agent’s performance was further improved as a result. Suay and Chernova [45] extend their work to a real-world robotic system using only human reward. However, the work of Suay and Chernova treats human reward in the same way as the environmental reward in traditional RL and does not model the human reward as the TAMER framework that will be described below.

Knox and Stone [22] propose the *TAMER* framework that allows an agent to learn from only human reward signals instead of environmental rewards by directly modeling it. With TAMER as a tool, Knox et al. [21] study how humans teach agents by examining their responses to changes in their perception of the agent and changes in the agent’s behavior. They deliberately reduce the quality of the agent’s behavior whenever the rate of human feedback decreases, and found that the agent can elicit more feedback from the human trainer but with lower agent performance. In addition, Li et al. [28, 29] investigate how the trainer’s behavior in TAMER is affected when an agent gives the trainer feedback. For example, they allow the agent to display informative feedback about its past and present performance, and competitive feedback about the agent’s performance relative to other trainers. However, in this article, we investigate the effect of agent’s competitive feedback on the trainer’s training behavior in a different setting where a small group of closely related subjects train at the same time in the same room.

Similar to the TAMER framework, Pilarski et al. [40] proposed a continuous action actor-critic reinforcement learning algorithm [14] that learns an optimal control policy for a simulated upper-arm robotic prosthesis using only human-delivered reward signals. Their algorithm does not model the human reward signals and tries to learn a policy to receive the most discounted accumulated human reward.

Recently, MacGlashan et al. [34] propose an Actor-Critic algorithm to incorporate human-delivered reinforcement. Specifically, they assume that the human trainer employs a diminishing returns strategy, which means the initial human feedback for taking the optimal action a in state s will be positive, but goes to zero as the probability of selecting action a in state s goes to 1. Based on this assumption, they take the human reward as an Advantage Function (Temporal Difference in traditional Reinforcement Learning is an unbiased estimate of advantage function), which describes how much better or worse an action selection is compared to the current expected behavior. Then they use human reward to directly update the policy function.

While the work mentioned above interprets human feedback as a numeric reward, Lofstin et al. [32] interpret human feedback as categorical feedback strategies that depend both

on the behavior the trainer is trying to teach and the trainer's teaching strategy. Then they proposed an algorithm to infer knowledge about the desired behavior from cases where no feedback is provided. The experimental results of Loftin et al.'s work show that their algorithms could learn faster than algorithms that treat the feedback as a numeric reward. In addition, Griffith et al. [13] propose an approach called 'policy shaping' by formalizing the meaning of human feedback as a label on the optimality of actions and using it directly as policy advice, instead of converting feedback signals into evaluative rewards.

Therefore, there are several possibilities to take facial expressions as evaluative feedback for an autonomous agent to learn to perform a task, e.g., numeric reward as in the work of Thomaz and Breazeal [49] and TAMER [22], or action feedback as in SABL [32] and policy advice as in policy shaping [13]. In this article, we choose TAMER as the foundation and starting point to investigate the potential of agent learning from human trainer's facial expressions, via interpreting them as human reward, and do not claim that it is superior to other methods that learn from human evaluative feedback.

2.2 Learning from facial expressions

Emotions including expression, motivation, feelings etc., play an important role in information processing, behavior, decision-making and learning in social animals, especially humans [2, 10, 39, 43]. Much research has been done on the role of emotion in learning. Some classic works on affect, i.e., the direction of an emotional state, emphasize cognitive and information processing aspects in a way that can be encoded into machine-based rules [37, 39]. However, little work has been done to investigate the relation between emotion and learning with computational models especially with reinforcement learning as context, except [4, 11, 12, 25, 51].

Gadanho used an emotion system to calculate a well-being value that was used as reinforcement. The system was with capabilities analogous to those of natural emotions and used Q-learning to learn behavior selection, i.e., to decide when to switch and reinforce behavior [11].

Broekens [4] examines the relationship between emotion, adaptation and reinforcement learning by taking human's real emotional expressions as social reinforcement. Their results show that affective facial expressions facilitate robot learning significantly faster compared to a robot trained without social reinforcement. However, in their work, the social reinforcement is simply added to the environmental reward to form a composite reinforcement. Moreover, affective facial expressions are mapped to a predefined fixed numeric as social reinforcement. In addition, a mechanism with 9 stickers on the face were used to help recognize facial expressions. By contrast, our work tries to build a model with data collected from 498 people to predict the human trainer's feedback based on her facial expressions during the time of giving keypress feedback without any physical help to recognize them. Moreover, we test agent's learning from these predictive feedback without taking the environmental reward into account.

Recently, Veeriah et al. [51] propose a method—face valuing, with which an agent can learn how to perform a task according to a user's preference from facial expressions. Specifically, face valuing learns a value function that maps facial features extracted from a camera image to expected future reward. Their preliminary results with a single user suggest that an agent can quickly adapt to a user's changing preferences and reduce the amount of explicit feedback required to complete a grip selection task. The motivation of 'face valuing' to learn from facial expressions is similar to ours. However, in their experiments, the

user is well-trained and the user's expressed pleasure or displeasure of the agent's action are directly used to train the agent by mapping to predefined numeric values. In our paper, we collect the data of facial expressions and keypress feedback from 498 ordinary people and build a model to predict the human trainer's evaluative feedback based on her facial expressions during the time of giving keypress feedback. Moreover, while 'face valuing' interprets facial expressions as human reward and seeks the largest accumulated discounted human reward, in our work, we investigate the potential of agent's learning from facial expressions by interpreting them as immediate human reward.

In addition, Peeled et al. [38] propose a method for predicting people's strategic decisions based on their facial expressions. Their experiment is conducted in a controlled environment with 22 computer science students. They ask the participants to play several games and record videos of the whole process. At the same time, they log the participants' decisions throughout the games. The video snippet of the participants' faces prior to their decisions is represented and served as input to a classifier that is trained to predict the participant's decision. Their results show that their method outperforms standard SVM as well as humans in predicting subjects' strategic decisions.

Gordon et al. [12] develop an integrated system with a fully autonomous social robotic learning companion for affective child-robot tutoring. They measure children's valence and engagement via an automatic facial expression analysis system. The measured valence and engagement were combined into a reward signal and fed into the robot's affective reinforcement learning algorithm. They evaluate their system with 34 children in preschool classrooms for a duration of two months. Their results show the robot can personalize its motivational strategies to each student using verbal and non-verbal actions. However, in their work, the detected valence and engagement are weighted and summed with predefined weights as social reinforcement, while in our work we intend to directly predict the reward value from the detected facial expression.

To endow a chess companion robot for children with empathic capabilities, Leite et al. [25] use a multimodal framework to model the user's affective states and allow the robot to adapt its empathic responses to the particular preferences of the child who is interacting with it. They combine visual and task-related features to measure the user's valence of feeling. The change of valence before and after the robot taking the empathic strategy is calculated as rewards for a multi-armed bandit reinforcement learning algorithm. Their preliminary study with 40 children show that robot's empathic behavior has a positive effect on users.

3 Background

This section briefly introduces interactive reinforcement learning, technical details on the TAMER framework and the Infinite Mario testing domain used in our experiment.

3.1 Interactive reinforcement learning

In traditional reinforcement learning [19, 46], the agent learns from rewards provided by a predefined reward function not a human user. Different from traditional reinforcement learning, interactive reinforcement learning (Interactive RL) was developed to allow an ordinary human user to shape the agent learner by providing evaluative feedback [22, 32, 33, 48, 49]. The objectives of Interactive RL are to facilitate the agent to learn from a non-expert human user

in agent design and even programming, and use the human's knowledge to speed up the agent learning. In interactive reinforcement learning, based on the observed state in the environment, the agent will take an action. Then the human teacher who observed the agent's behavior will evaluate the quality of agent's action based on her knowledge, as shown in Fig. 1. The evaluation is used as feedback for the agent to update the learned behavior. Therefore, the agent's optimal behavior is decided by the evaluation provided by the human teacher.

As in traditional reinforcement learning (RL), an interactive RL agent learns to make sequential decisions in a task, represented by a policy deciding the action to be taken by the agent in an environmental state. A sequential decision task is modeled as a Markov decision process (MDP), denoted by a tuple $\{S, A, T, R, \gamma\}$. In MDP, time is divided into discrete time steps, and S is a set of states in the environment that can be encountered by the agent and A is a set of actions that the agent can perform. At each time step t , the agent observes the state of the environment, $s_t \in S$. Based on the observation, the agent will take an action $a_t \in A$. The experienced state-action pair will take the agent into a new state s_{t+1} in the environment, decided by a transition function $T: S \times A \times S$, which tells the probability of the agent transitioning to one state based on the action selection in a given state, $T(s_t, a_t, s_{t+1}) = Pr(s_{t+1}|s_t, a_t)$. The agent will receive an evaluative feedback r_{t+1} , provided by the human observer by evaluating the quality of the action selection based on her knowledge. That is to say, there is no predefined reward function in interactive RL— $R: S \times A \times S \rightarrow \mathfrak{R}$, which decides a numeric reward value at each time step based on the current state, action chosen and the resultant next state. Instead, the reward function is in the human teacher's mind.

The agent's learned behavior is described as a *policy*, $\pi: S \times A$, where $\pi(s, a) = Pr(a_t = a|s_t = s)$ is the probability of selecting a possible action $a \in A$ in a state s . The goal of the agent is to maximize the accumulated discounted reward the agent receives, denoted as $\sum_{k=0}^{\infty} \gamma^k r_{t+k+1}$ at time step t , where γ is the discount factor (usually $0 \leq \gamma < 1$). γ determines the present value of rewards received in the future: a reward received k time steps in the future is worth only γ^{k-1} times what it would be worth if it were received immediately. The return for a policy π is denoted as $\sum_{k=0}^{\infty} \gamma^k R(s_{t+k}, \pi(s_{t+k}), s_{t+k+1})$. There are usually two associated value functions for each learned policy π . One is the *state-value function*, referred to as the value of a state, $V^{\pi}(s)$, which is the expected return when an agent starts in a state s and follows a *policy* π thereafter, where

$$V^{\pi}(s) = E_{\pi} \left[\sum_{k=0}^{\infty} \gamma^k r_{t+k+1} | s_t = s \right]. \quad (1)$$

Fig. 1 Interaction in the interactive reinforcement learning framework

Similarly, another value function is the *action-value function*, referred to as the value of a state–action pair, $Q^\pi(s, a)$, which is the expected return after taking an action a in a state s , and thereafter following a *policy* π , where

$$Q^\pi(s, a) = E_\pi \left[\sum_{k=0}^{\infty} \gamma^k r_{t+k+1} | s_t = s, a_t = a \right]. \quad (2)$$

For each MDP, there exists a set of optimal policies π^* , which share the same optimal *state-value function*, V^* , defined as $V^*(s) = \max_\pi V^\pi(s)$, and *action-value function*, Q^* , defined as $Q^*(s, a) = \max_\pi Q^\pi(s, a)$. The goal of the agent is to learn an optimal policy π^* from its interaction with the human teacher.

3.2 TAMER framework

In this article, we use the TAMER framework [22] as the agent’s learning algorithm. The TAMER framework was built for a variant of the Markov decision process (MDP), a model of sequential decision-making addressed via dynamic programming [16] and reinforcement learning [46]. In the TAMER framework, there is no reward function encoded before learning. An agent implemented according to TAMER learns from real-time evaluations of its behavior, provided by a human teacher who observes the agent. These evaluations are taken as human reward signals. Therefore, TAMER is a typical interactive reinforcement learning method.

Knox and Stone [22] first proposed the original TAMER framework which learns the reward function and selects actions with it. They then proposed VI-TAMER which learns a value function from the learned human reward function via value iteration and selects actions with the value function [23]. In the original TAMER framework, the agent learns myopically from human reward, i.e., only taking the immediate reward into account by setting the discount factor γ to 0. In VI-TAMER, the discount factor γ is set close to 1, i.e., the agent seeks the largest accumulated discounted human reward, which is the same as in traditional reinforcement learning. Therefore, the original TAMER is equivalent to VI-TAMER when the discount factor γ is set to 0. In such case, the learned value function in VI-TAMER is equivalent to the learned human reward function.

In this article, we rephrase TAMER as a general model-based method for interactive reinforcement agent learning from human reward, as shown in Fig. 2. In this case, the TAMER agent learns a model of the human reward and then uses the learned human reward function to learn a value function model. The TAMER agent will select actions with the value function model to get the most accumulated discounted human reward. Figure 2 shows the diagram of agent learning in the TAMER framework.

Specifically, in TAMER, the human teacher observes the agent’s behavior and can give reward corresponding to its quality. There are four key modules for an agent learning with TAMER. The first one is to learn a predictive model of human reward. Specifically, the TAMER agent learns a function:

$$\hat{R}_H(s, a) = \mathbf{w}^T \Phi(s, a), \quad (3)$$

where $\mathbf{w} = (w_0, \dots, w_{m-1})^T$ is the parameter vector, and $\Phi(\mathbf{x}) = (\phi_0(\mathbf{x}), \dots, \phi_{m-1}(\mathbf{x}))^T$ are the basis functions, and m is the total number of parameters. $\hat{R}_H(s, a)$ is a function for approximating the expectation of human rewards received in the interaction experience, $R_H: S \times A \rightarrow \mathfrak{R}$.

Fig. 2 Agent learning in the TAMER framework (modified from [20])

Since it takes time for the human teacher to assess the agent's behavior and deliver her feedback, the agent is uncertain about which time steps the human reward is targeting at. The second module is the credit assigner to deal with the time delay of human reward caused by evaluation of agent's behavior and delivering it. TAMER uses a credit assignment technique to deal with the delay of human reward and multiple rewards for a single time step. Specifically, inspired by the research on the delay of human's response in visual searching tasks of different complexities [15], TAMER defines a probability density function to estimate the probability of the teacher's feedback delay. This probability density function provides the probability that the feedback occurs within any specific time interval and is used to calculate the probability (i.e. the credit) that a single reward signal is targeting a single time step. If a probability density function $f(t)$ is used to define the delay of the human reward, then at the current time step t , the credit for each previous time step $t-k$ is computed as:

$$c_{t-k} = \int_{t-k-1}^{t-k} f(x) dx. \quad (4)$$

If the human teacher gives multiple rewards, the label h for each previous time step (state-action pair) is the sum of all credits calculated with each human reward using Eq. 4. The TAMER agent uses the calculated label and state-action pair as a supervised learning sample to learn a human reward model— $\hat{R}_H(s, a)$ by updating its parameters, e.g., with incremental gradient descent. If at any time step t the human reward label h received by the agent is not 0, temporal difference error δ_t is calculated as

$$\begin{aligned} \delta_t &= h - \hat{R}_H(s, a) \\ &= h - \mathbf{w}^T \Phi(s_t, a_t). \end{aligned} \quad (5)$$

Based on the gradient of least square, the parameter of $\hat{R}_H(s, a)$ is updated with incremental gradient descent:

$$\begin{aligned}\mathbf{w}_{t+1} &= \mathbf{w}_t - \alpha \nabla_{\mathbf{w}} \frac{1}{2} \{h - \hat{R}_H(s_t, a_t)\}^2 \\ &= \mathbf{w}_t + \alpha \delta_t \Phi(s_t, a_t),\end{aligned}\quad (6)$$

where α is the learning rate.

The third one is the value function module. The TAMER agent learns a state value function model $V(s)$ or an action value function model $Q(s, a)$ from the learned human reward function $\hat{R}_H(s, a)$. At each time step, the agent will update the value function as:

$$Q(s, a) \leftarrow \hat{R}_H(s, a) + \gamma \sum_{s' \in S} T(s, a, s') \times \max_{a'} Q(s', a'), \quad (7)$$

or

$$V(s) \leftarrow \max_a [\hat{R}_H(s, a) + \gamma \sum_{s' \in S} T(s, a, s') V(s')]. \quad (8)$$

where $T(s, a, s')$ is the transition function, s and a are current state and action, s' and a' are the next state and action.

The fourth module is the action selector with the predictive value function model. As a traditional RL agent which seeks the largest discounted accumulated future rewards, the TAMER agent can also seek the largest accumulated discounted human reward by greedily selecting the action with the largest value, as:

$$a \leftarrow \arg \max_a Q(s, a), \quad (9)$$

or

$$a \leftarrow \arg \max_a \left[\hat{R}_H(s, a) + \sum_{s' \in S} T(s, a, s') V(s') \right]. \quad (10)$$

The TAMER agent learns by repeatedly taking an action, sensing reward, and updating the predictive model $\hat{R}_H$ and corresponding value function model. The trainer observes and evaluates the agent's behavior. In our experiment, she can give reward by pressing two buttons on the keyboard, which are assigned to the agent's most recent actions. Each press of the two buttons is mapped to a numeric reward of -1 or $+1$ respectively. In our experiment, we set the discount factor γ to 0. Then the learned value function in TAMER is equivalent to the learned human reward function $\hat{R}_H$. And the action selector chooses actions with $\hat{R}_H$. Note that unlike [20], when no feedback is received from the trainer, learning is suspended until the next feedback instance is received.

3.3 Infinite Mario domain

Super Mario Bros is an extremely popular video game, making it an excellent platform for investigating how humans interact with agents that are learning from them. To establish the generalizability of TAMER to more complex domains, and to make the experiment appealing to trainers of all ages, we implemented TAMER in the Infinite Mario domain from the Reinforcement Learning Competition [8, 55]. The Infinite Mario domain was adapted from the classic Super Mario Bros video game.

In Infinite Mario, the Mario avatar must move towards the right of the screen as fast as possible and at the same time collect as many points as possible. To facilitate comparisons of TAMER with other learning methods that have been applied to this domain, we used the standard scoring mechanism that was established for the Reinforcement Learning Competition. The standard scoring mechanism gives positive reward for killing a monster (+ 1), grabbing a coin (+ 1), and finishing the level (+ 100). It gives negative points for dying (− 10) and for each time step that passes (− 0.01). The actions available for Mario are (left, right, no direction), (not jumping, jumping) and (not sprinting, sprinting), resulting in 12 different combined actions for the agent at every time step. The state space is quite complex, as Mario observes a 16×21 matrix of tiles, each of which has 14 possible values.

To reduce the state space, in our TAMER implementation we take each visible enemy (i.e. monster) and each tile within a 8×8 region around Mario as one state feature. The most salient features of the observations will be extracted as state representation. For each state feature, a number of properties are defined, including whether it is a:

- pit,
- enemy,
- mushroom,
- flower,
- coin,
- smashable block,
- question block,

and the distance (x —horizontal direction, y —vertical direction and Euclidean distances) from Mario. We filter and select the top two state features by ranking all state features based on a priority of whether it is a pit, an entity (a monster, mushroom, flower or fire-ball), a block and the distance. The state representation includes the properties of the selected two state features and the properties of Mario. The properties of Mario include whether it is at the right of a wall and the speed of it (x -speed and y -speed). Thus, the feature vector Θ for the state representation is

$$\Theta = [\phi_1, \phi_2, \phi_M], \quad (11)$$

where ϕ_1 and ϕ_2 are two vectors for the two selected state features with each consisting of the above 10 properties, and ϕ_M is a vector consisting of the properties of Mario.

In this article, the TAMER agent learns a *model tree* [52] that constructs a tree-based piecewise-linear model to estimate $\hat{R}_H(s, a)$. The inputs to $\hat{R}_H$ are the above state representation and the combined action. The TAMER agent takes each observed reward signal as part of a label for the previous state–action pair (s, a) and then uses it as a supervised learning sample to update the model tree by the divide-and-conquer method. The model tree can have multivariate linear models at each node with only features tested in the subtree of the current node instead of using all features, analogous to piecewise linear functions. We use model tree because features for state representation are mostly binary and not all features are always relevant. Model tree can select the relevant subset of the features to predict the human reward, thus resulting in more accurate prediction.

In our study, we have 3 levels (0, 1 and 2) in Infinite Mario domain. Level 0 is from the Reinforcement Learning Competition generated with seed 121 difficulty 0. Note that the seed is a random integer value that was used by the level generator to generate levels by probabilistically choosing a series of idiomatic pieces of levels and fitting them together

[44]. We designed level 1 and 2 based on level 0 with increased difficulties, e.g., increasing the number of monsters, changing the type of monsters, adding one pit, changing the height of walls and length of flat stretches, etc. As in Super Mario Bros, Mario can enter the next level automatically if he finishes one level. The game goes back to level 0 if Mario dies or finishes level 2. A given game ends when Mario dies.

To see whether a TAMER agent can successfully learn to play the game and compare the learning performance with other methods, the first author trained the agent on level 0 in the Infinite Mario Domain for 10 trials with TAMER. An episode ends when Mario dies or finishes the level. The policy was frozen and recorded at the end of each episode of training. Then, each recorded policy was tested for 20 games offline. The performance for each episode was averaged over the 20 games and then over the 10 trials. The result shows that our TAMER agent can achieve 120 points in 12 episodes, while it takes about 500 episodes for a SARSA agent to achieve a similar performance level [47] and a hierarchical SARSA agent implemented with object-oriented representation about 50 episodes to reach a similar level and 300 episodes to achieve 149.48 points [36], which is almost optimal. Although the TAMER agent does not learn an optimal policy, it can successfully learn a good policy substantially faster than these other methods, making this set-up very suitable for our experiments with members of the public where each training session can only last up to 15 min.

4 Experimental setup

Our user study was conducted in conjunction with the research program of the NEMO science museum in Amsterdam. This program enables museum visitors to participate as experimental subjects in studies conducted by scientists from nearby universities. In our experiment, a group of up to four trainers, typically a family or group of friends, trained their own TAMER agents at the same time. In each group, each participant sat at her own table, facing away from the other members and their screens. There was also a camera on the screen in front of the trainer's face. Each participant signed a consent form (with parental consent for children under 18 years old) permitting recording the data and using it for research. Then participants in the group were asked to train their own agents for 15 min in the same room at the same time.

Each participant could quit at any time she wanted before the 15 min elapsed. Finally, we debriefed the participants and asked for feedback and comments. The experiment was carried out in the local language with English translations available for foreign visitors. We recorded the training data including state observation, action, human reward, the time of human reward being given, score, the ending time for each time step, and video data of facial expressions and keypresses on the keyboard for each trainer during training. Note that one time step corresponds to the execution of an action by the agent. Trainers were not given time to practice before the experiment because we were concerned that they might get tired of expressing facial emotions after the practice.

Our experiment is a between-subjects study with 561 participants from more than 27 countries participated and randomly distributed into our four experimental conditions (described below). Of them, 221 were female and 340 were male respectively, aged from 6 to 72. Figure 3 shows the distribution of participants across age ranges and genders. Data from 63 participants were disregarded: five participants lacked parental consent; three had not played Super Mario Bros before and were unable to judge Mario's behavior; one had an

Fig. 3 Demographic information of participants across age ranges and genders in the four conditions described in Section 5 of our study. Note that: **a** control condition, **b** facial expression condition, **c** competitive condition, **d** competitive facial expression condition, *F* female, *M* male

internet connection problem, one quit after only 5 min training; and the rest did not fully understand the instructions, got stuck and gave feedback randomly by alternating positive and negative feedback in quick succession, or interrupted their family members. After pruning the data, 498 participants remained: 109 participants in the control condition; 100 in the facial expression condition; 135 in the competitive condition; and 154 in the competitive facial expression condition.

5 Experimental conditions

In this section, we describe the four conditions we proposed and tested in our experiment. We investigate whether telling trainers to use facial expression as an additional channel to train agents will affect the trainer's training, agent's learning and the trainer's facial expressiveness. In addition, factors like environmental stress from outside environment might affect the expressiveness of facial expressions. This would have a further impact on the prediction accuracy of the model trained based on these facial feedback. Our previous research [30, 31] showed that an agent's socio-competitive feedback can motivate the trainer to provide more feedback and the agent will ultimately perform

better. In this study, we want to investigate how an agent's competitive feedback will affect the trainer's facial expressiveness and the agent's learning, especially when it is coupled with facial expression condition. We also want to know whether an agent can learn from human trainer's facial expressions via interpreting as evaluative feedback.

Note that as in the original TAMER, in all conditions, participants could give positive and negative feedback by pressing buttons on the keyboard to train the agent. Only the keypress signal was used for agent learning and videos of training by participants in all conditions were recorded.

5.1 Control condition

The interface for the control condition is the performance-informative interface replicated from [30] and implemented in the Infinite Mario domain, as shown in Fig. 4.

In the interface, each bar in the performance window above the game board indicates the agent's performance in one game chronologically from left to right. The agent's performance is the score achieved by the Mario agent in the game task. During training, the pink bar represents the score received so far for the current game, while the dark blue bars represent the performance of past games. When a game ends (i.e., Mario dies), the corresponding bar becomes dark blue and a new score received in the new game is visualized by a pink bar to its right. When the performance window is full, the window is cleared and new bars appear from the left. Trainers in this condition were told to use keypresses to train the agent.

Fig. 4 The training interface used in the control condition and facial expression condition. Note that a trainer can switch between the 'training' and 'not training' modes by pressing the space button, which will be shown in the interface. In the 'training' mode, the trainer can give rewards to train the agent. In the 'not training' mode, the agent's learned policy is frozen and used to perform the task. The trainer cannot give feedback to further train the agent any more

5.2 Facial expression condition

The interface used in this condition is the same as in the control condition except that trainers were told to use their facial expressions as encouraging explicit feedback, e.g., happy and sad expressions would map to positive and negative reward respectively, in addition to keypresses, to train the agent. That is to say, at the same time of giving keypress feedback, they were instructed to give positive or negative facial expressions corresponding to the keypress feedback. We told them this because we want to investigate whether telling trainers to use facial expressions as a separate channel to train the agent would affect the trainer's training and agent's learning, compared to trainers in the control condition. We hypothesize that telling trainers to use facial expressions to train the agent will result in worse performing agents than agents trained by those being told to use only keypress feedback to train agents. This is because telling participants to use facial expressions as separate reward signal could induce distraction from giving high quality keypress feedback. In addition, we hypothesize that because of more posed facial behaviors by trainers in this condition, the expressiveness of trainers' facial expression will be higher than those in the control condition.

5.3 Competitive condition

Factors like environmental stress from outside environment might affect the expressiveness of facial expressions. This would have a further impact on the prediction accuracy of the model trained based on these facial feedback. Our previous research [31] showed that putting people in a socio-competitive situation could further motivate them to give more feedback and improve the agent's performance. In this study, we want to investigate how an agent's competitive feedback will affect the trainer's facial expressiveness and the agent's learning, especially when it is coupled with facial expression condition. Therefore, in the competitive condition, we allow the agent to indicate the rank and score of the other members of the group, who are all training their own agents simultaneously in the same room, as described in the experimental setup in the previous section. The groups typically consist of family members or close friends, e.g., children and (grand) parents, brothers and sisters.

To implement this condition, we added a leaderboard to the right of the interface used in the control and facial expression conditions, as shown in Fig. 5. In the leaderboard, the first names, scores and ranks of all the group members currently training the agent are shown. When the trainer starts training for the first time, her agent's performance is initialized to 0 and ranked in the leaderboard. Whenever the trainer finishes a game (i.e. Mario dies), the new game score and rank are updated in the leaderboard. To create more movement up and down in the leaderboard, the last game score is always used. The trainer can directly check her score and rank in the leaderboard. Therefore, the trainer can keep track of both the agent's learning progress and the agent's performance relative to that of other members of her group.

5.4 Competitive facial expression condition

The final condition is a combination of the facial expression and competitive conditions. Specifically, the interface is the same as in the competitive condition but, as in the facial

Fig. 5 The training interface for the competitive condition and competitive facial expression condition

expression condition, trainers were told to use both keypresses and facial expressions to train agents. As in other conditions, only keypresses were actually used for agent learning. We hypothesize that the expressiveness of trainers' facial expressions in this condition will be higher than those in the competitive condition, since trainers in this condition were told to use facial expressions as additional channel to train the agent.

6 Experimental results

In this section, we present and analyze our experimental results. We consider one additional individual variable: mode, in addition to 'competition' and 'facial expression' the two tested independent variables. 'Mode' means the most frequent highest level in the game reached by the agent tested offline: level 0, level 1 and level 2 in the game design, as explained in Sect. 6.4. However, some other factors such as the distribution of the trainer's skill levels across conditions, experience in gaming especially in Super Mario, environmental stress, the domain stochasticity, participant's cultural difference etc., may still affect the results in our study. Nonetheless, we believe that the large number of participants can compensate for these variabilities encountered while running studies in the non-laboratory setting.

6.1 Feedback given

Figure 6 shows how feedback was distributed per 200 time steps over the learning process for the four conditions. From Fig. 6 we can see that, the number of time steps with feedback received by agents in the four conditions increased at the early training stage and decreased dramatically afterwards, which supports previous studies [21, 28] and our motivation for investigating methods of enabling agents to learn from the trainer's facial expressions. In addition, trainers in the competitive and competitive facial

Fig. 6 Mean number of time steps with feedback per 200 time steps for all four conditions during the training process. Note that black bars represent standard deviations

expression condition seem to have a trend to give more feedback than those in the control and facial expression conditions before 1000 time steps, though not significant (t test). Moreover, subjects in the facial expression condition tend to give a similar number of keypress feedback compared to those in the control condition and subjects in the competitive facial expression condition tend to give a similar number of keypress feedback compared to those in the competitive condition, even though they were told to give feedback via both keypress and facial expression.

6.2 Analysis of facial feedback

It is well known that facial expressions reflect inner feelings and emotions. Factors from outside environment might affect the expressiveness of facial expressions, which might have a further impact on the prediction accuracy of the model trained based on these facial feedback. Therefore, before testing whether an agent can learn from facial feedback, we first assess effects of ‘facial expression’ and ‘competition’ on the informativeness of facial expressions as a feedback signal. To this end, 3-D locations of 512 densely defined facial landmarks (see Fig. 7) are automatically detected and tracked using the state-of-the-art method proposed by Jeni et al. [18]. Videos with a down-sampled frame rate of 20 fps are used in tracking. Data from 31 participants (5.5% of the data analyzed) are discarded due to methodological problems such as face occlusion, talking, and chewing gum during the experiment. In total 9,356,103 frames are tracked. To eliminate rigid head movements, the tracked faces are shape-normalized by removing translation, rotation and scale. Since the normalized faces are approximately frontal with respect to the camera, we ignore the depth (z) coordinates of the normalized landmarks. Consequently, 1024 location parameters are obtained per frame, i.e., x and y coordinates of the tracked 512 facial landmarks.

Fig. 7 512 tracked facial landmarks

6.2.1 Facial expressiveness

To analyze facial activity levels of different conditions, we compute the standard deviation of the landmark positions within a 2 s interval around each keypress feedback (1 s before to 1 s after). The 2 s interval is chosen because the great majority of facial expressions of felt emotions last between 0.5 and 4 s [9] and the differences of standard deviations between conditions are largest in this case. Standard deviations are averaged for each subject in each condition. Then, we analyze the significant differences in vertical (y) and horizontal (x) facial movements, between different conditions using a one tailed t test. The computed p values for x and y coordinates are combined as $p_c = \sqrt{p_x^2 + p_y^2}$ to represent the significance level for each landmark position as one parameter.

Figure 8a–f visualize the significance level of differences (p_c) in expressiveness between different conditions, namely *control versus facial expression*, *competitive versus competitive facial expression*, *control versus competitive facial expression*, *control versus competitive*, *facial expression versus competitive facial expression*, and *facial expression versus competitive*. The visualized p_c values for transition locations between detected landmarks are computed by linear interpolation. It is important to note that the second condition in each comparison (except the *competitive condition* in *facial expression versus competitive*) displays a higher activeness for each landmark. This may be due to high correlation between the displacements of dense landmarks.

When we analyze the significant differences ($p_c < 0.05$) in facial activeness between different conditions, it is seen that the mouth region is more dynamic in the *facial expression condition* compared to the *control condition* (Fig. 8a). Similarly, the deviation of movements in the mouth, upper cheek, and forehead regions for the *competitive facial expression condition* are higher than those of the *competitive condition* (Fig. 8b). These results can be explained by the fact that subjects exaggerate their expressions in facial expression conditions.

In *control versus competitive*, *facial expression versus competitive facial expression*, and *control versus competitive facial expression*, competitive conditions display higher activity levels almost for the whole surface of the face in comparison to the control conditions as shown in Fig. 8c–e. These findings suggest that competitive conditions can elevate facial expressiveness. Furthermore, expressiveness in the *competitive condition* seems to

Fig. 8 Significance level (p_c) of differences in expressiveness between different conditions: **a** control versus facial expression, **b** competitive versus competitive facial expression, **c** control versus competitive facial expression, **d** control versus competitive, **e** facial expression versus competitive facial expression, and **f** facial expression versus competitive. Dark gray or black colored ($p_c < 0.05$) regions display significantly different activeness between the corresponding conditions. Notice that the second condition in each comparison (except in **f** facial expression versus competitive) displays a higher activeness for each landmark. In **f**, five of 512 landmarks (shown as “+”) have a lower activeness for the second condition (competitive)

be higher than that of the *facial expression* condition (Fig. 8f). However, significant differences are observed only around/on the eyebrows region (mostly on the right eyebrow). Activeness on the remaining facial surface does not differ significantly.

In summary, consistent with our hypotheses, telling trainers to use facial expressions as additional channel to train agents could increase the trainer’s facial expressiveness. Moreover, competition can also elevate the trainer’s facial expressiveness. As shown in Fig. 8f, higher facial expressiveness is observed in the competitive condition compared to the facial expression condition. In addition, large differences were observed between the control and competitive conditions (Fig. 8d, e). These findings indicate that the effect of the competitive situation on expressiveness could be larger than that effect of the facial expression condition (although their difference may be limited).

6.2.2 Classification of positive and negative feedback with facial expressions

Next, we investigate the discriminative power of facial responses for classifying positive and negative feedback. To this end, we implement a binary classification method using a convolutional neural network—recurrent neural network architecture (CNN–RNN) that can be trained jointly in an end-to-end manner. We opt for employing a deep architecture for this task since deep learning methods can effectively model complex data with high accuracy.

The CNN module in our architecture is composed of four convolutional layers followed by two fully connected layers. A set of 5×5 -pixel filters are used in all convolutional layers. Rectified linear unit (ReLU) is applied to the output of each convolutional layers. Max-pooling with a 2×2 windows is applied after each convolution except for the first convolutional layer. After the final max-pooling, two fully connected layers follow. The output of the last fully-connected layer can be thought as the spatial representation of face, and fed as input to RNN module (recurrent modules at the corresponding frame).

Since recent studies [5–7] suggest that temporal patterns of expressions provide discriminative information, we model the spatio-temporal dynamics of the expressions for distinguishing between negative and positive feedback classes. To this end, we employ a two-layer RNN (with 256 units at each layer) followed by 2 output neurons with ReLU activation function. We train the proposed CNN–RNN model minimizing the mean squared (classification) error. Details of our CNN–RNN architecture is shown in Table 1.

To obtain training and test samples (video segments) for our CNN–RNN model, we extract 2 s intervals from facial videos, around each positive and negative feedback key-press (1 s before to 1 s after). To deal with pose, scale, and translation variations as well as obtaining pixel-to-pixel comparable face images regardless of expression or identity variations, each face image (in videos) is warped onto a frontal average face shape using a piecewise linear warping. Note that, in this way, the facial landmark points aligned to the same location for each of the warped/normalized faces. Images are then converted to gray scale and the resulting video segments are fed to the CNN–RNN model.

In our experiments, we employ a fivefold cross-validation testing scheme with the same data split. Each time one fold is used as test set and the remaining fourfolds are used to train and validate the model. There is no subject overlap between training and test folds. Thus, our results are based on subject-independent training. 107,395 positive and 99,702 negative feedback instances were used in the experiment.

To show the effectiveness and accuracy of our proposed method, results of a random baseline are also reported for comparison. Class labels for random baseline are assigned by drawing a random class label according to the ratio of positive and negative class labels

Table 1 Configuration of the proposed CNN–RNN architecture

Layer	Kernel size	Stride	Output size
conv1-1	5×5	1	$92 \times 92 \times 16$
conv1-2	5×5	1	$88 \times 88 \times 32$
pool1	2×2	2	$44 \times 44 \times 32$
conv2	5×5	1	$40 \times 40 \times 64$
pool2	2×2	2	$20 \times 20 \times 64$
conv3	5×5	1	$16 \times 16 \times 128$
pool3	2×2	2	$8 \times 8 \times 128$
fc4	–	–	1×2048
fc5	–	–	1×2048
rn6	–	–	1×256
rn7	–	–	1×256
fc8	–	–	1×2

“conv”, “fc”, and “rn” denote convolution, full connection, and recurrent layers in the network, respectively. Numbers next to the name of each layer indicate the order of the corresponding layer. For brevity, ReLU layers are discarded in the table

Table 2 Accuracy of classifying positive and negative feedback using facial responses

Condition	Positive	Negative	Total
Proposed method			
Control	0.62	0.69	0.66
Facial expression	0.65	0.73	0.70
Competitive	0.75	0.70	0.73
Competitive facial expression	0.79	0.75	0.78
Random baseline			
Control	0.50	0.50	0.50
Facial expression	0.42	0.58	0.51
Competitive	0.52	0.48	0.50
Competitive facial expression	0.58	0.42	0.51

from the training set. As shown in Table 2, the use of facial expressions significantly (t test with $p < 0.001$) outperforms the random baseline in each condition. The highest accuracy is achieved for the *competitive facial expression* condition, followed by the *competitive* condition. This can be explained by the increased facial expressivity due to the competitive setting and posed facial expressions. As expected, the proposed method provides higher accuracies for facial expression conditions.

6.3 Learning from facial feedback

Obviously, a critical next step is to examine whether an agent can learn from the human trainer's facial expressions with the trained CNN–RNN model in Sect. 6.2.2 to predict human reward based on human trainer's facial expressions. Ideally, we would run a new experiment where the trainer gives new facial expressions while watching the agent learn and use the predictor to get reward that we trained with. But that's prohibitively expensive since there are still many problems to be solved before testing whether agents or social robots can successfully learn feedback signals from facial expressions in fairly unconstrained settings. For example, someone might be smiling for any number of reasons that have nothing to do with the agent. Therefore, to evaluate whether an agent can learn from predicted facial feedback with our trained CNN–RNN model in Sect. 6.2.2 from online interaction with human users is not ready yet. In the paper, we take a first step and do an evaluation with the data we collected. The closest approximation is to get predicted human reward based on facial expressions at the time when keypress feedback was given for the complete training trajectory of each trainer. Thus, we use the predicted feedback instead of the keypress feedback to train the agent for the complete trajectory. Please notice that we trained our CNN–RNN model in a subject-independent manner. In other words, our prediction model does not require or employ previously recorded data of the test subjects.

We compare the average learning performance of the four conditions in terms of learning from keypress feedback, learning from binary keypress feedback, learning from random feedback and learning from predicted binary feedback. The differences between these four kinds of feedback are illustrated in Fig. 9. In the case of learning from keypress feedback, the agent learns from weighted aggregate human reward, which was calculated with a credit assign technique on the original keypress feedback in the recorded training data, as shown in Sect. 3.2. In the case of learning from binary keypress feedback, we remove the weight for human reward in the recorded training data by

Fig. 9 Illustration of the difference between agent learning from keypress feedback, binary keypress feedback, predicted binary feedback and random feedback

formatting it into positive or negative (binary keypress) feedback. Then we trained the agent with these modified binary keypress feedback. In this case, learning from these binary keypress feedback is equivalent to a perfect binary prediction with a trained model on facial expression (100% accurate prediction). And learning from random feedback represent the worst case scenario for binary prediction (with probability of 0.5 for positive and negative feedback each, 50% prediction accuracy). Our model with 62–79% prediction accuracy for these four conditions are in the middle of learning from binary keypress and random feedback.

Agent learning from these four kinds of feedback were all trained with collected data. So the only difference between agent learning from keypress feedback, binary keypress feedback, predicted binary feedback and random feedback is the feedback for each time step in the training trajectory. Specifically, the agent continually learns from each kind of feedback, and we record the learning policy per 200 time steps and test the policy for 20 games. The offline performance is the average of the performance in the 20 games. Then the agent continues learning from each kind of feedback for another 200 time steps with policy recorded and tested for 20 games each. This process repeats up to 2800 time steps.

We compare the agent's learning performance from predicted binary feedback to random feedback, actual keypress feedback, and binary keypress feedback in all four conditions, as shown in Fig. 10. Our experimental results in Fig. 10 show that, when the prediction accuracy is low in the first three conditions (control, facial expression and competitive condition), agent learning from predicted binary feedback with our model is only a little better than learning from random feedback. However, when the prediction accuracy increased to 79% in the competitive facial expression condition, agent learning from predicted binary feedback with our model can reach to around 10 which is close to the performance of learning from binary keypress feedback (around 20). In this case, learning from binary keypress feedback is equivalent to learning from a trained model with 100 percent prediction accuracy. Both learning from predicted binary feedback with our model and learning from binary keypress feedback are in the same scenario. The only difference is the prediction accuracy. Therefore, this suggest that learning solely from predicted feedback based on facial expressions is possible and there is still much room for improvement in agent's performance using improved models with higher prediction accuracy.

Fig. 10 Offline performance of agent learning from predicted binary feedback with facial expressions, compared to learning from keypress feedback, binary keypress feedback and random feedback for all four conditions

In addition, from Fig. 10 we can see that for agent's learning from all four types of feedback, the agent's learning performance quickly goes up at the beginning of the training process and goes down afterwards. From our observation in the experiment, trainers started to train complex behaviors for Mario after training for a while, which caused the decrease of performance in the middle of training, even with keypress feedback. At the beginning, it is easy to train the agent to run to the right to complete the level, which will get the most score (+ 100). After training for some time, trainers started to train complex behaviors, e.g., picking up mushroom, which gets no point at all. In this case, they need to train the agent to unlearn the right running behavior, which cause Mario to stop and go left. In addition, the Mario agent will get -0.01 score for one time step longer in the game.

Using keypress feedback (weighted aggregate human reward), the agent can relearn the right running behavior and recover from the decrease of performance. However, with binary keypress feedback, it would be much difficult since the weight for human reward is removed, as shown in Fig. 10 for all four conditions. In this case, learning from binary keypress feedback is equivalent to learning from a trained model with 100% prediction accuracy. In our experiment, the highest prediction accuracy is 79%. So it would be much more difficult for agents learning from our predicted binary feedback to recover from the decrease of the performance. Moreover, the simulation responses (i.e., the resultant states

and actions of the agent) are from the recorded agent's trajectory that learns from actual keypress feedback and they are not based on the behavior that learns from predicted binary feedback using facial expressions. In other words, not only the accuracy of feedback prediction but also the temporal positioning of correctly predicted responses affects the simulation results. However, our results show that learning from binary keypress feedback cannot recover from the decrease of the performance, while learning from original keypress feedback (adding weight) can recover from it. This suggests that adding the weight of human reward could help the agent to recover from the decrease of performance. Therefore, a regression prediction model might further help agent learning from predicted feedback based on facial expressions, though further investigation into the prediction of evaluative feedback based on facial responses with regression method is needed.

6.4 Effect of competitive feedback on agent's performance

To examine the effect of agent's competitive feedback on its learning performance, we analyze the distribution of final offline performance for each condition with the learned final policy tested offline for 20 games and averaged over the 20 games for each subject. Figure 11a, which contains histograms of the final offline performance for the four conditions, shows that the distribution in the four conditions are all characterized by three modes. The gap between modes is caused by the score mechanism which gives credit + 100 for finishing one level and much less otherwise. Therefore, the three modes in Fig. 11a from left to right correspond to level 0, 1 and 2 in the game respectively. Then we compare the agent's performance trained by subjects affected by agent's competitive feedback to those unaffected by it across modes and within each mode separately. Figure 11b highlights the importance of 'competition'. From Fig. 11b we can see that, 'competition' can positively affect agent learning, especially those in mode 3 where the agent performs best. This support prior results [29, 31] demonstrating the importance of bi-directional interaction and competitive elements in the training interface, and show that 'competition' can significantly improve agent learning and help the best trainers the most.

7 Discussion and open questions

Our work contributes to the design of human-agent systems that facilitate the agent to learn more efficiently and be easier to teach. To our knowledge, we are the first to highlight clearly the relationship between the reward signal and the nature of facial expressions. We demonstrate that an understanding of how to design the interaction between the agent and the trainer allows for the design of the algorithms that support how people can teach effectively and be actively engaged in the training process at the same time. This is useful for personalizing interaction with a socially assistive robotics, e.g., an educational assistive robot tutor [12] trying to teach children a second language. Facial expression interaction can also be used for children with autism to improve these children's social interaction abilities [41]. In such cases, facial expression can be extracted as evaluative feedback for personalizing the interaction process for users with different abilities. Human rewards given without the intention to teach or otherwise affect behavior—possibly derived from smiles, attention, tone of voice, or other social cues are more abundantly broadcast and can be observed without adding any cognitive load to the human [20].

Fig. 11 **a** Distribution of final offline performance across the four conditions, **b** effect of competition on agent's learning performance across and within each mode. Note that *FE* facial expression

More recently, research has also focused on developing robots that can detect common human communication cues for more natural interactions. Social HRI is a subset of HRI that encompasses robots which interact using natural human communication modalities, including speech, body language and facial expressions. This allows humans to interact with robots without any extensive prior training, permitting desired tasks to be completed more quickly and requiring less work to be performed by the user [35]. The systems and techniques discussed above focus on the recognition of one single input mode in order to determine human affect. The use of multimodal inputs over a single input provides two main advantages: when one modality is not available due to disturbances such as occlusion or noise, a multimodal recognition system can estimate affective state using the remaining modalities, and when multiple modalities are available, the complementarity and diversity of information can provide increased robustness and performance. For example, similar to our work, Ritschel and Addre [42] and Weber et al. [54] used the audience's vocal laughs and visual smiles to calculate the reward via a predefined reward function to shape the humor of a robot.

In this article, we focus on one modality—human's facial expressions, and investigate the potential of extracting evaluative feedback from facial expressions to train an agent to perform a task. Our results show that an agent is able to learn given only facial expressions, though the learning performance is modest. There is still much room for improvement in agent's performance using improved models with higher prediction accuracy. Moreover, in our experiment, facial expressions were recorded when the trainer was focusing on training the agent and in some conditions they were asked to give facial expressions intentionally. While in the real world, the environment is complex and such rewards extracted from facial expressions might be untargeted, e.g., someone might be smiling for any number of reasons that have nothing to do with the agent. Consequently, interpretation and attribution of these social cues will be especially challenging. This is evidenced by the results in Table 2, where the reward signal of the facial expressions in the control condition were the most difficult to predict. Therefore, there is still much work to be done before agents or social robots begin to learn feedback signals from facial expressions in fairly unconstrained settings. For example, in our work, the prediction based on facial expressions was done at the time of giving keypress feedback, we also need to understand whether facial expression feedback might also be useful for agent learning even when no keypress feedback is given or when both are used for agent learning.

8 Conclusion

This article investigated the potential for agents to learn from human trainers' facial expressions. To this end, we conducted the first large-scale study with usable data from 498 participants (children and adults) by implementing TAMER in the Infinite Mario domain. With designed CNN–RNN model, our analysis shows that telling trainers to use facial expressions makes them inclined to exaggerate their expressions, resulting in higher accuracy for predicting positive and negative feedback using facial expressions. Competitive conditions also elevated facial expressiveness and further increased predicted accuracy. This has significant consequences for the design of agent learning systems that wish to take into account a trainer's spontaneous facial expressions as a reward signal. Moreover, our results with a simulation experiment show that learning solely from predicted feedback based on facial expressions is possible and there is still much room for improvement in

agent's performance using improved models with higher prediction accuracy or a regression method. In addition, our experiment supports previous studies demonstrating the importance of bi-directional feedback and competitive elements in the training interface. Finally, we believe that our approach could transfer to other domains and apply to other interactive learning algorithms, since TAMER succeeds in many domains including Tetris, Mountain Car, Cart Pole, Keepaway Soccer, Interactive Robot Navigation etc. [20, 24].

Acknowledgements This work was in part supported by Natural Science Foundation of China (Grant No. 51809246), Natural Science Foundation of Shandong Province (Grant No. ZR2018QF003). This research was part of Science Live, the innovative research program of Science Center NEMO in Amsterdam that enables scientists to carry out real, publishable, peer-reviewed research using NEMO visitors as volunteers. Science Live is partially funded by KNAW and NWO.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Argall, B. D., Chernova, S., Veloso, M., & Browning, B. (2009). A survey of robot learning from demonstration. *Robotics and Autonomous Systems*, 57(5), 469–483.
2. Berridge, K. C. (2003). Pleasures of the brain. *Brain and Cognition*, 52(1), 106–128.
3. Blumberg, B., Downie, M., Ivanov, Y., Berlin, M., Johnson, M. P., & Tomlinson, B. (2002). Integrated learning for interactive synthetic characters. *ACM Transactions on Graphics (TOG)*, 21(3), 417–426.
4. Broekens, J. (2007). Emotion and reinforcement: Affective facial expressions facilitate robot learning. In T. S. Huang, A. Nijholt, M. Pantic, & A. Pentland (Eds.), *Artificial intelligence for human computing*, pp. 113–132. Berlin: Springer.
5. Cohn, J. F., Kruez, T. S., Matthews, I., Yang, Y., Nguyen, M. H., Padilla, M. T., et al. (2009). Detecting depression from facial actions and vocal prosody. In *Affective computing and intelligent interaction and workshops, 2009. ACII 2009. 3rd international conference on* (pp. 1–7). IEEE.
6. Dibeklioglu, H., Hammal, Z., Yang, Y., & Cohn, J. F. (2015). Multimodal detection of depression in clinical interviews. In *Proceedings of the 2015 ACM on international conference on multimodal interaction* (pp. 307–310). New York: ACM.
7. Dibeklioglu, H., Salah, A. A., & Gevers, T. (2015). Recognition of genuine smiles. *IEEE Transactions on Multimedia*, 17(3), 279–294.
8. Dimitrakakis, C., Li, G., & Tziortziotis, N. (2014). The reinforcement learning competition 2014. *AI Magazine*, 35(3), 61–65.
9. Ekman, P., & Scherer, K. (1984). Expression and the nature of emotion. *Approaches to Emotion*, 3, 19–344.
10. Frijda, N. H., & Mesquita, B. (2000). Beliefs through emotions1. *Emotions and Beliefs: How Feelings Influence Thoughts*, 3, 45.
11. Gadanho, S. C. (2003). Learning behavior-selection by emotions and cognition in a multi-goal robot task. *The Journal of Machine Learning Research*, 4, 385–412.
12. Gordon, G., Spaulding, S., Westlund, J. K., Lee, J. J., Plummer, L., Martinez, M., et al. (2016). Affective personalization of a social robot tutor for children's second language skills. In *AAAI* (pp. 3951–3957).
13. Griffith, S., Subramanian, K., Scholz, J., Isbell, C., & Thomaz, A. L. (2013). Policy shaping: Integrating human feedback with reinforcement learning. In *Advances in neural information processing systems* (pp. 2625–2633).

14. Grondman, I., Busoniu, L., Lopes, G. A., & Babuska, R. (2012). A survey of actor-critic reinforcement learning: Standard and natural policy gradients. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 42(6), 1291–1307.
15. Hockley, W. E. (1984). Analysis of response time distributions in the study of cognitive processes. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 10(4), 598.
16. Howard, R. A. (1960). *Dynamic programming and Markov processes*. Cambridge: The MIT Press.
17. Isbell, C., Shelton, C. R., Kearns, M., Singh, S., & Stone, P. (2001). A social reinforcement learning agent. In *Proceedings of the fifth international conference on autonomous agents* (pp. 377–384). New York: ACM.
18. Jeni, L. A., Cohn, J. F., & Kanade, T. (2015). Dense 3D face alignment from 2D videos in real-time. In *Automatic face and gesture recognition (FG), 2015 11th IEEE international conference and workshops on* (Vol. 1, pp. 1–8). IEEE.
19. Kaelbling, L. P., Littman, M. L., & Moore, A. W. (1996). Reinforcement learning: A survey. *Journal of Artificial Intelligence Research*, 4, 237–285.
20. Knox, W. (2012). *Learning from human-generated reward*. Ph.D thesis, University of Texas at Austin.
21. Knox, W. B., Glass, B. D., Love, B. C., Maddox, W. T., & Stone, P. (2012). How humans teach agents. *International Journal of Social Robotics*, 4(4), 409–421.
22. Knox, W. B., & Stone, P. (2009). Interactively shaping agents via human reinforcement: The tamer framework. In *Proceedings of the fifth international conference on Knowledge capture* (pp. 9–16). New York: ACM.
23. Knox, W. B., & Stone, P. (2015). Framing reinforcement learning from human reward: Reward positivity, temporal discounting, episodicity, and performance. *Artificial Intelligence*, 225, 24–50.
24. Knox, W. B., Stone, P., & Breazeal, C. (2013). Training a robot via human feedback: A case study. In G. Herrmann, M. J. Pearson, A. Lenz, P. Bremner, A. Spiers, U. Leonards (Eds.), *Proceedings of the international conference on social robotics*, pp. 460–470. Berlin: Springer.
25. Leite, I., Pereira, A., Castellano, G., Mascarenhas, S., Martinho, C., & Paiva, A. (2011). Modelling empathy in social robotic companions. In *International conference on user modeling, adaptation, and personalization* (pp. 135–147). Berlin: Springer.
26. Li, G., Dibeklioglu, H., Whiteson, S., & Hung, H. (2016). Towards learning from implicit human reward. In *Proceedings of the 2016 international conference on autonomous agents & multiagent systems* (pp. 1353–1354). International Foundation for Autonomous Agents and Multiagent Systems.
27. Li, G., Hung, H., & Whiteson, S. (2015). A large-scale study of agents learning from human reward. In *Proceedings of the 2015 international conference on autonomous agents and multiagent systems* (pp. 1771–1772). International Foundation for Autonomous Agents and Multiagent Systems.
28. Li, G., Hung, H., Whiteson, S., & Knox, W. B. (2013). Using informative behavior to increase engagement in the tamer framework. In *Proceedings of the 2013 international conference on autonomous agents and multi-agent systems* (pp. 909–916).
29. Li, G., Hung, H., Whiteson, S., & Knox, W. B. (2014). Learning from human reward benefits from socio-competitive feedback. In *4th Joint IEEE international conferences on development and learning and epigenetic robotics* (pp. 93–100). IEEE.
30. Li, G., Whiteson, S., Knox, W. B., & Hung, H. (2016). Using informative behavior to increase engagement while learning from human reward. *Autonomous Agents and Multi-Agent Systems*, 30(5), 826–848.
31. Li, G., Whiteson, S., Knox, W. B., & Hung, H. (2018). Social interaction for efficient agent learning from human reward. *Autonomous Agents and Multi-Agent Systems*, 32(1), 1–25.
32. Loftin, R., Peng, B., MacGlashan, J., Littman, M. L., Taylor, M. E., Huang, J., et al. (2015). Learning behaviors via human-delivered discrete feedback: Modeling implicit feedback strategies to speed up learning. In *Autonomous agents and multi-agent systems* (pp. 1–30).
33. MacGlashan, J., Ho, M. K., Loftin, R., Peng, B., Roberts, D., Taylor, M. E., & Littman, M. L. (2017). Interactive learning from policy-dependent human feedback. arXiv preprint [arXiv:1701.06049](https://arxiv.org/abs/1701.06049).
34. MacGlashan, J., Littman, M. L., Roberts, D. L., Loftin, R., Peng, B., & Taylor, M. E. (2016). 2016. Convergent actor critic by humans. In *Workshop on human-robot collaboration: Towards co-adaptive learning through semi-autonomy and shared control at IROS*.
35. McColl, D., Hong, A., Hatakeyama, N., Nejat, G., & Benhabib, B. (2016). A survey of autonomous human affect detection methods for social robots engaged in natural HRI. *Journal of Intelligent & Robotic Systems*, 82(1), 101–133.
36. Mohan, S., & Laird, J. E. (2011). An object-oriented approach to reinforcement learning in an action game. In *Proceedings of the seventh AAAI conference on artificial intelligence and interactive digital entertainment* (pp. 164–169).

37. Ortony, A. (1990). *The cognitive structure of emotions*. Cambridge: Cambridge University Press.
38. Peled, N., Bitan, M., Keshet, J., & Kraus, S. (2013). Predicting human strategic decisions using facial expressions. In *IJCAI*.
39. Picard, R. W., Papert, S., Bender, W., Blumberg, B., Breazeal, C., Cavallo, D., et al. (2004). Affective learning—A manifesto. *BT Technology Journal*, 22(4), 253–269.
40. Pilarski, P. M., Dawson, M. R., Degris, T., Fahimi, F., Carey, J. P., Sutton, R. S. (2011). Online human training of a myoelectric prosthesis controller via actor-critic reinforcement learning. In *Rehabilitation robotics (ICORR), 2011 IEEE international conference on* (pp. 1–7). IEEE.
41. Pour, A. G., Taheri, A., Alemi, M., & Meghdari, A. (2018). Human–robot facial expression reciprocal interaction platform: Case studies on children with autism. *International Journal of Social Robotics*, 10(2), 179–198.
42. Ritschel, H., & André, E. (2018). Shaping a social robot's humor with natural language generation and socially-aware reinforcement learning. In *Proceedings of the workshop on NLG for human–robot interaction* (pp. 12–16).
43. Scherer, K. R. (2001). Appraisal considered as a process of multilevel sequential checking. *Appraisal Processes in Emotion: Theory, Methods, Research*, 92, 120.
44. Smith, G., Treanor, M., Whitehead, J., Mateas, M. (2009). Rhythm-based level generation for 2D platformers. In *Proceedings of the 4th international conference on foundations of digital games* (pp. 175–182). ACM.
45. Suay, H. B., & Chernova, S. (2011). Effect of human guidance and state space size on interactive reinforcement learning. In *20th IEEE international symposium on robot and human interactive communication* (pp. 1–6). IEEE.
46. Sutton, R., & Barto, A. (1998). *Reinforcement learning: An introduction*. Cambridge: Cambridge University Press.
47. Taylor, M. E. (2011). Teaching reinforcement learning with Mario: An argument and case study. In *Proceedings of the second symposium on educational advances in artificial intelligence* (pp. 1737–1742).
48. Tenorio-Gonzalez, A. C., Morales, E. F., & Villaseñor-Pineda, L. (2010). Dynamic reward shaping: Training a robot by voice. In *Advances in artificial intelligence—IBERAMIA 2010* (pp. 483–492). Berlin: Springer.
49. Thomaz, A. L., & Breazeal, C. (2008). Teachable robots: Understanding human teaching behavior to build more effective robot learners. *Artificial Intelligence*, 172(6), 716–737.
50. Vail, P. L. (1994). *Emotion: The on/off switch for learning*. Rosemont: Modern Learning Press.
51. Veeriah, V., Pilarski, P. M., Sutton, R. S. (2016). Face valuing: Training user interfaces with facial expressions and reinforcement learning. arXiv preprint [arXiv:1606.02807](https://arxiv.org/abs/1606.02807).
52. Wang, Y., & Witten, I. H. (1996). *Induction of model trees for predicting continuous classes*. Hamilton, New Zealand: Department of Computer Science, University of Waikato.
53. Watkins, C. J., & Dayan, P. (1992). Q-learning. *Machine Learning*, 8(3–4), 279–292.
54. Weber, K., Ritschel, H., Aslan, I., Lingenfelder, F., & André, E. (2018). How to shape the humor of a robot-social behavior adaptation based on reinforcement learning. In *Proceedings of the 2018 on international conference on multimodal interaction* (pp. 154–162). ACM.
55. Whiteson, S., Tanner, B., & White, A. (2010). The reinforcement learning competitions. *AI Magazine*, 31(2), 81–94.