Human engagement providing evaluative and informative advice for interactive reinforcement learning

Adam Bignold¹ · Francisco Cruz^{2,3} · Richard Dazeley² · Peter Vamplew¹ · Cameron Foale¹

Received: 10 February 2021 / Accepted: 12 December 2021
© The Author(s) 2022

Abstract

Interactive reinforcement learning proposes the use of externally sourced information in order to speed up the learning process. When interacting with a learner agent, humans may provide either evaluative or informative advice. Prior research has focused on the effect of human-sourced advice by including real-time feedback on the interactive reinforcement learning process, specifically aiming to improve the learning speed of the agent, while minimising the time demands on the human. This work focuses on answering which of two approaches, evaluative or informative, is the preferred instructional approach for humans. Moreover, this work presents an experimental setup for a human trial designed to compare the methods people use to deliver advice in terms of human engagement. The results obtained show that users giving informative advice to the learner agents provide more accurate advice, are willing to assist the learner agent for a longer time, and provide more advice per episode. Additionally, self-evaluation from participants using the informative approach has indicated that the agent's ability to follow the advice is higher, and therefore, they feel their own advice to be of higher accuracy when compared to people providing evaluative advice.

Keywords Interactive reinforcement learning · Assisted reinforcement learning · Evaluative and informative advice · Reward shaping · Policy shaping · User study

1 Introduction

Reinforcement learning (RL) aims at the creation of agents and systems that are capable of functioning in real-world environments [1]. A common RL task involves decision-making and control, which given some information about

the current state of the environment must determine the best action to take in order to maximise long-term success. In this regard, RL allows improving the decision-making process while operating, to learn without supervision, and adapt to changing circumstances [2]. In classical, autonomous RL [1], the agent interacts with its environment learning by trial and error. The agent explores the environment and learns solely from the rewards it receives (see grey box within Fig. 1). RL has shown success in different domains such as management [3, 4], chemical processes [5, 6], robot scenarios [7–9], and game environments [10, 11], among others. However, RL has difficulties to learn in large state spaces. As environments become larger, the agent's training time increases and finding a solution can become impractical [12, 13].

Interactive reinforcement learning (IntRL) is an alternative to RL in which an advisor interacts with an RL agent in real-time [14]. The advisor can provide extra information to the agent regarding its behaviour or future actions it should perform. In this regard, the advice can be either evaluative or informative [15]. The former is an evaluation the advisor gives to the agent indicating how good or bad

Adam Bignold and Francisco Cruz have contributed equally to this work.

This work has been partially supported by the Australian Government Research Training Program (RTP) and the RTP Fee-Offset Scholarship through Federation University Australia.

✉ Francisco Cruz
francisco.cruz@deakin.edu.au

¹ School of Engineering, IT and Physical Sciences, Federation University, Ballarat, Australia

² School of Information Technology, Deakin University, Geelong, Australia

³ Escuela de Ingeniería, Universidad Central de Chile, Santiago, Chile

Fig. 1 Interactive reinforcement learning. The autonomous agent performs an action a_t from a state s_t , and the environment produces a response leading the agent to a new state s_{t+1} and receiving a reward

r_{t+1} . The interactive approach adds a human advisor for assistance. The advisor also observes the environment's response and can provide either evaluative or informative advice to the learner agent

was the last action performed. The latter is a suggestion given to the agent indicating what action to perform next from the current state. From the agent's perspective, these styles of interaction are commonly referred to as reward shaping and policy shaping, respectively. Human advisors are usually used in IntRL since they achieve good performance in areas such as problem-solving, forward planning, and teaching. Moreover, they have a large collection of knowledge and experiences to draw upon when encountering new environments and problems [16]. IntRL utilises these skills of humans to assist the agent with its own learning and decision-making. This approach has been shown to considerably improve the agent's learning speed and can allow RL to scale to larger or more complex problems [17]. Figure 1 shows the IntRL approach with a human advisor included providing either evaluative or informative advice to the learner agent.

There are two major barriers to humans providing information to RL agents. The first is the time required by the human. In this regard, it is important that the mechanisms used to provide advice to the agent serve to reduce the number of interactions required [18]. The second barrier is the skill needed by the human to provide the information. Humans usually need both programming skills and knowledge of the problem dynamics to encode information relevant to the agent's learning [19, 20]. A principle of IntRL is that the method to provide information to the agent should be understandable and usable by people without programming skills or deep problem domain expertise [14, 21]. Therefore, the time required by a human advisor should remain as low as possible to reduce the burden on the human and methods for providing information to an agent should be accessible to users without programming or machine learning expertise.

Current studies on human engagement when teaching artificial agents have been mainly focused on assessing human commitment independent of the type of advice. For instance, Thomaz and Breazel [22] have shown that human

tutors tend to have a positive bias when teaching agents, opting to reward rather than punish RL agents. It has been also shown that humans providing advice to a system over an extended period experienced frustration and boredom when receiving too many questions from the agent [23]. In this regard, studies suggest that participants do not like being prompted for input repeatedly, particularly when the input can be repetitive [24]. Current IntRL systems require the users to provide advice on a state-by-state basis [25], leaving current systems susceptible to the same issues of frustration and interruption as the active learning systems reported.

In this work, we aim to reduce the obligation of the human advisor while improving the learning speed of the agent. We address the question of which of the approaches, evaluative or informative, is the preferred instructional approach for humans. To this aim, we carry out an analysis of human engagement with twenty participants with no prior knowledge of machine learning techniques in order to avoid any previous bias. In our experiments, ten users give evaluative advice to the RL agent while ten users give informative advice in a simulated scenario. From the performed interactions, we analyse the advice accuracy and the advice availability of each assistive approach. We also present an analysis of how evaluative advice may be affected by reward bias when teaching the RL agent.

Therefore, the distinction between advice delivery styles, i.e. evaluative or informative (also known as reward-shaping and policy-shaping, respectively), and how humans engage and prefer to teach artificial agents, is studied in this work. While evaluative and informative approaches are about the method used to instruct the agent, reward-shaping and policy-shaping methods are about how the agent incorporates the provided advice, thus considering the agent's viewpoint. The main contributions of this paper can be summarised as follows:

- Introduction of the concepts of evaluative and informative advice as methods to instruct autonomous agents.
- Analysis of human engagement when providing advice to a learning agent.
- Analysis of how evaluative advice is affected by reward bias.
- Comparison of evaluative and informative methods as preferred by human advisors.
- Evidence that users giving informative advice provide more accurate advice is willing to assist the learner agent for a longer time and provides more advice per episode.

2 Reinforcement learning and interactive human-sourced advice

Learning from the ground up can be a challenging task. While humans and artificial agents using RL are both capable of learning new tasks, it is evident that any extra information regarding the task can significantly reduce the learning time [26–28]. For humans, we can get advice from peers, teachers, the Internet, books, or videos, among other sources. By incorporating advice, humans can learn what the correct behaviour looks like, build upon existing knowledge, evaluate current behaviour, and ultimately reduce the amount of time spent performing the wrong actions [29]. For artificial agents, the benefits of advice are the same. For instance, advice may be used to construct or supplement the reward function, resulting in an improved evaluation of the agent's actions or increased the utility of the reward function requiring fewer experiences to learn a behaviour [30, 31]. The advice can also be used to influence the agent's policy, either directly or through the action selection method, in order to reduce the search space [32, 33].

There are many possible information sources for agents to use. For instance, external information can come from databases [34], labelled sets [35, 36], cases [37, 38], past experiences [39], other agents [40, 41], contextual perception [42], and from humans [43]. Human-supplied advice is contextually relevant information that comes from a human as a result of observation or awareness of the agent's current behaviour or goal. This information is commonly used to supplement, construct, or alter the RL process. Human-sourced advice can be more noisy, inaccurate, and inconsistent than other information sources. However, the critical benefit is that the advice is contextually relevant and can be applied to aid the agent in its current situation or goal.

IntRL may use human-sourced advice [44] or simulated users [45] to directly interact with the agent while it is learning/operating [14]. The focus for IntRL is limited to the use of advice during the learning process, not before or after. This limitation requires interactive techniques to be easy for an agent to get information from, and for humans to add information to so that the learning process is not slowed down. This limitation also means that the agent or policy should not be reset when new information is provided, as that is conceptually similar to creating a new agent rather than interacting with an existing one. When humans interact with the agent, they may either provide additional rewards in response to the agent's performance [46] or recommend actions to the agent to guide the exploration process [25].

2.1 Evaluative advice

Evaluative advice is information that critiques current or past behaviour of an agent [47, 48]. Advice that supplements, improves, or creates a reward function is considered to be evaluative as it is a reaction to an agent's behaviour rather than a direct influence on an agent's decision-making. The source of the advice is what separates evaluative advice from the reward function. A typical reward function is defined for an specific environment, whereas evaluative advice originates from an observer of the agent or other external sources [20, 49]. Figure 2 shows in green colour evaluative advisors supplementing the reward received from the environment.

Humans providing evaluative advice do not need to know the solution to a problem [50, 51]; it is enough for them to be able to assess the result of an action and then

Fig. 2 Interactive reinforcement learning feedback. While the informative advisor may suggest an action to be performed by the agent, the evaluative advisor may suggest a reward to supplement the reward obtained from the environment

decide whether it was the correct action to take. For instance, in the training an agent manually via evaluative reinforcement (TAMER) framework [52, 53], a human user continually critiques the RL agent's actions. The human observes the agent and in response to the agent's actions provides a simple yes/no evaluation of its choice of action. This Boolean evaluation acts as an additional reward signal, supplementing the reward function from the environment. This bare minimum of human influence is enough to significantly decrease the time required by the agent to learn the required task [52].

Another example of evaluative advice is the convergent actor-critic by humans (COACH) approach [54]. In this approach, a human trainer may give positive or negative feedback to a virtual dog learning to reach a goal position. The human feedback was divided into punishment and reward and labelled with different levels as 'mild electroshock,' 'bad dog,' 'good dog,' and 'treats.' Using COACH, the agent was able to learn the task facing multiple feedback strategies. Recently, this approach has been extended as Deep COACH [55] to represent the agent policy by deep neural networks.

2.2 Informative advice

Informative advice is information that aids an agent in its decision-making [56, 57]. Advice that recommends actions to take or avoid, suggests exploration strategies, provides information about the environment or proactively alters what action an agent may take is considered to be informative. Informative methods primarily focus on transferring information from the human and encoding it into the agent's policy, either directly, by altering the policy, or indirectly by influencing the agent's decision-making process [58]. Figure 2 shows in brown colour informative advisors suggesting an action to be taken.

Providing informative advice can be challenging for two reasons, the first of which is the human factor. Informative advice typically requires the human to know what the correct action is for a given state ahead of time. Not only does this require a greater understanding of the environment and the agent's position within it, but it also requires a more substantial commitment of time and effort to provide the advice. The time and effort required increase as the size of the environment, and the available actions increases [59]. The second reason utilising informative advice is challenging is that encoding information sourced from a human into a form an agent can understand can be a complicated process, as it is more informationally dense than evaluative advice [60].

For instance, an implementation of informative advice in IntRL is the ADVISE algorithm [61]. In ADVISE, a human observing an agent in operation can recommend

actions to take at any given step, which the agent may choose to follow. This methodology allows the human to guide the agent through parts of the environment which they are familiar with. This can result in a significant improvement over existing IntRL methods and a reduced need for exploration.

Another example of informative advice was presented by Cruz et al. [26] in which a robot learned a cleaning task using human-provided interactive feedback. In this domestic scenario, seven actions could be advised to the agent using multi-modal audiovisual feedback. The provided advice was integrated into the learning process with an affordance-driven [42] IntRL approach. After experiments, the robot collected more and faster reward, tested against different minimal confident level thresholds and different levels of affordance availability.

2.3 Evaluative versus informative

Evaluative advice has been more widely utilised in prior research as implementations are simpler to encode as the focus tends to be on the result of a decision rather than on what decision should be made [62]. This is due to it being easier to determine if an action was the correct or incorrect action to take once the result of the action is available. Most implementations of evaluative advice alter or supplement the reward function of the environment. Encoding information to alter the reward function is generally straightforward, as the primary focus is on whether to increase or decrease the reward given to the agent, as opposed to informative implementations that attempt to alter the decision-making policy [63]. Additionally, providing an evaluation requires less human effort than determining what information or action is relevant for a given state, as the information sought is typically a Boolean or a scalar measurement. Overall, evaluative advice is more direct to obtain, implement, and encode than the informative counterpart.

Informative advice tends to be more informationally dense than evaluative advice. While this does make sourcing and encoding the information difficult, it does provide more benefit to the agent [62]. Evaluative advice only reinforces behaviour after that behaviour has been exhibited, whereas informative advice can promote or discourage behaviour before it is presented. Advice that recommends taking or avoiding actions will reduce the search space for the agent, resulting in improved learning time. The downside of this is that if the agent never performs actions that are preemptively discouraged, and the advice is not optimal, then the optimal policy may not be found [64].

A direct comparison of the two styles is difficult as the implementations of human-sourced advice vary. Griffith

et al. [61] compared the effects of informative versus evaluative advice on artificial agents using their informative algorithm ADVISE, against the evaluative algorithm TAMER. Both algorithms utilise IntRL agents, and advice is given on a step-by-step basis. The ADVISE algorithm prompts the advisor for a recommended action which the agent can then follow, while TAMER prompts the advisor for a binary evaluation on the previously taken action. In the experiments, each agent is assisted by a simulated human, making the advice comparable.

The ADVISE algorithm allows the advisor to recommend an action, and therefore, the number of bits of information provided is equal to $\log_2(n_a)$ where n_a is the number of possible actions (e.g. if there are eight possible actions $n_a = 8$; then, each piece of informative advice provides three bits of information). In contrast, TAMER allows the human to provide a binary evaluation (i.e. correct/incorrect) which provides only a single bit of information. Therefore, the information gain from ADVISE is greater than TAMER and may bias the results. However, the experiments show that informative advice is more beneficial to the agent regardless of advice accuracy for the majority of cases. The use of a simulated human as an oracle in these experiments allowed for the provision of consistent advice that does not suffer from biases introduced by real humans. However, if the behaviour of actual human advice givers differs from that of the simulated human in terms of either accuracy and/or engagement, then the impact on agent behaviour may not reflect that observed in this study. Therefore, it is important to develop an understanding of the properties of actual human advice.

2.4 Human engagement

Studies on human engagement and teaching styles when engaging with interactive machine learning agents have previously been studied [21, 65]; however, they have been mainly focused on assessing human commitment independent of the type of advice. For instance, Amershi et al. [21] presented a comprehensive study looking at the engagement between humans and interactive machine learning. The study included some case studies demonstrating the use of humans as information sources in machine learning. This work highlighted the need for increased understanding of how humans engage with machine learning algorithms and what teaching styles the users preferred.

A study by Thomaz and Breazeal [22], later confirmed by Knox and Stone [66], found that human tutors tend to have a positive bias when teaching machines, opting to reward rather than punish RL agents. This bias leads to agents favouring the rewards provided by the human over the reward function of the environment. The positive bias

was observed in humans providing evaluative advice, as it tends to be provided as a reward [22]. Due to its characteristics, no such bias has been tested for or observed yet in informative-assisted agents. Knox and Stone [67] later mitigated the consequence of the positive bias in RL agents by developing an agent that valued human reward gained in the long term rather than the short term.

Another study performed by Cakmak and Thomaz [23] investigated the strategy of teachers when tutoring machine learning agents. The study found that humans providing advice to a system over an extended period experienced frustration and boredom when bombarded with questions from the agent. The stream of questions to the teachers caused some participants to “turn their brain off” or “lose track of what they were teaching” according to self-reports [68]. Similar results were obtained using a movie recommendation system developed for Netflix, where participants were repeatedly asked to state if the system was right or wrong [69, 70].

The previous studies suggest that participants do not like being prompted for input repeatedly, particularly when the input can be repetitive. Current IntRL systems do not prompt the user for information, instead, allowing the advisor to step in whenever they wish. Nevertheless, input into these systems is repetitive and requires the users to provide advice on a state-by-state basis [25], leaving current systems susceptible to the same issues of frustration and interruption as the active learning systems reported. Regardless, it is still not clear whether these issues will be translated into the IntRL scenarios. Therefore, the remainder of this paper reports details and results of an experiment carried out to establish the characteristics of advice provided by humans interacting with an IntRL agent and to assess whether these properties alter depending on whether evaluative or informative advice is provided.

3 Experimental methodology

In this section, we describe the IntRL methodology used during the experiments and frame the approach within an assisted RL framework. Moreover, we outline the method to collect human advice including participants’ characteristics, induction process, experiment details, and after-experience questionnaire.

3.1 Interactive reinforcement learning methodology

Assisted reinforcement learning (ARL) [20] is a general framework proposed to incorporate external information into traditional RL. The framework uses a conceptual taxonomy including processing components and

communications links to describe transmission, modification, and modality of the sourced information. The processing components comprise information source, advice interpretation, external model, and assisted agent, whereas the communications links are temporality, advice structure, and agent modification. ARL agents aim to gather as much information from an external source as possible, as this can lead to improved performance within the environment. A concrete example of an ARL agent is an IntRL agent. As previously mentioned, an IntRL agent can be advised with externally sourced information to support the learning process at any time of the training.

In this work, two different learner agents attempt to solve the Mountain Car problem [1] using IntRL (more details about the experimental problem are given in the next section): the first agent accepts evaluative advice and the other receives informative advice. Figure 3 shows the IntRL approach framed within the ARL framework [20] using both evaluative and informative advice. The figure shows the processing components using dotted red squares and the communication links using green parallelograms with underlined text. Using the ARL taxonomy, there are some common processing components and communication links that are adopted similarly by both approaches. The common elements are information source,

temporality, external model, and the assisted agent, which are adopted by the ARL framework as human-sourced advice, interactive assistance, an immediate model, and a Q-learning agent. All the other processing components and communication links differ to each other for evaluative and informative advice. For the evaluative approach, advice interpretation, advice structure, and agent modification are adopted by the ARL framework as binary advice to reward conversion, state–action pair value, and reward shaping, respectively. For the informative approach, they are adopted as advice to action selection conversion, state–action lookup, and policy shaping, respectively. As this approach relies on human trainers as an external information source, the higher the people engagement, the higher the opportunity to transfer knowledge to the agent. The accuracy of the advice and information gain as a result of the advice provided is also important, as they contribute to the policy being learned by the agent [64].

We aim to measure the human engagement, accuracy of advice, and the information gain for evaluative and informative advice for IntRL. To this aim, we perform experiments using two IntRL agents implemented with the temporal difference learning method Q-learning. The performance of the agent, or its ability to solve the problem, is not the main focus of this paper. A comparison of evaluative and informative advice, in terms of the performance of the agents, has been investigated in a prior study [61].

In the context of this work, human engagement is a measure of the number of interactions, the total time spent constructing interactions, and the distribution of interactions over the time the agent is operating. The observing human is given an opportunity to provide information once per step of the agent, and if the human does provide some advice during that step, then the interaction is recorded. However, a measure of the number of interactions is not sufficient, as the time and effort required to provide an interaction may differ between informative and evaluative advice methods. As a result, the interaction time is also recorded. Moreover, the accuracy of the information provided to the agent affects its performance within the environment [64]. In this regard, advice accuracy is a measure of how accurate the information provided by the human is, compared to the optimal action to take for each state the agent encounters. This can be calculated by comparing the advice provided by the human against the known optimal policy for this task.

Fig. 3 Interactive reinforcement learning method used to compare human engagement in evaluative and informative advice. The method is presented using the assisted reinforcement learning taxonomy [20], defining processing components (dotted red squares) and communication links (underlined green parallelograms) for each advice delivery style

3.2 Human-sourced advice

During the experiments, twenty people participated, ten for each advice delivery style¹. Each participant was able to communicate with an RL agent while observing its current state and performance. A participant interacting with the evaluative agent had the option of providing an agreement or disagreement (yes/no) to the agent's choice of action for the last time step. This binary evaluation was then used by the agent to supplement the reward it receives from the environment. A positive evaluation added +1 to the reward, while a negative evaluation subtracted -1 from the reward. Likewise, a participant interacting with the informative agent had the option of suggesting an action for the agent's next step, either left or right. If the agent was recommended an action, then that action was taken; otherwise, the agent operates as a usual RL agent. Each participant, regardless of teaching style, had three possible options each step. For the evaluative advice participants, the options were: agree, disagree, or do nothing, whereas, for the informative advice participants, the options were to recommend: left, right, or to do nothing.

The participants chosen for the experiment had not had significant exposure to machine learning and were not familiar with the Mountain Car environment. Before beginning the experiment, each participant was given a 5-min induction to the Mountain Car problem and then asked to complete a short questionnaire. The induction introduces the aim of the agent, the dynamics of the environment, the action space, and most significantly what the optimal solution to the problem is. The solution for the environment is described to the participant to give all participants an equal understanding and to reduce the time that they spend exploring the environment themselves so that they may focus on assisting the agent.

When the induction was complete, the participant was asked to complete a questionnaire. The full questionnaire consists of seven questions, the first two of which aim to assess the level of general knowledge about machine learning techniques and understanding of the Mountain Car problem of the participants. After completing the first two questions, the participant is ready to begin the experiment. The remaining five questions were answered after the subject had completed their interaction with the agent.

The participant was given 500 ms to provide advice to the agent each step. To provide advice to the agent, the participant pressed one of two keys on the keyboard to indicate either approval/disapproval of the agent's last choice in action when using evaluative advice, or to

recommend the left/right action for the agent to take next when using informative advice. Therefore, the input mechanism was dependent on the advice delivery style being tested. If the human provided advice within the 500ms window, an interaction had taken place and the time taken to create that interaction was recorded. If the human did not provide advice within the time window provided, then no interaction was recorded, and the agent operated as usual. Additionally, the human could change the duration of the time window by 25% during the experiment by pressing the +/- keys. The experiments ran until the participant believed the agent had learned the correct behaviour, or until they tired of providing advice at which point the agent was terminated.

After the participant had chosen to stop providing advice, they were asked to complete the remainder of the questionnaire. The remaining five questions aimed to assess understanding of the Mountain Car problem now that the participants have experienced the environment. It also aimed to capture their perception about their level of engagement, the accuracy of their advice, and the agent's understanding of the advice supplied.

4 Interactive reinforcement learning scenario

In this section, we describe the key features of the experimental environment including the agent's representation, state and action representation, and reward function. Furthermore, we complement the human-agent interactive methodology described in the previous section by indicating the script given to the participants.

4.1 Features of the environment

The Mountain Car environment is a standard continuous-state testing domain for RL [1, 71]. In the environment, an underpowered car must drive from the bottom of a valley to the top of a steep hill. Since the gravity in the environment is stronger than the engine of the car, the car cannot drive straight up the side of the mountain. In order for the car to reach the top of the mountain, it must build up enough inertia and velocity. Figure 4 illustrates the mountain car environment and its features.

In our experiments, an RL agent controls the actions of the car. The car begins at a random position and with a low velocity somewhere within the starting position. In order to reach the goal position, the agent must build up enough momentum. To do so, the agent accelerates towards the goal until its velocity is reduced to zero by gravity. At this point, the agent turns and accelerates towards the other direction towards the highest possible position, again, until

¹ Ethics approval letter provided by the Human Research Ethics Committee of *Anonymous* University, research project number *Anonymous*.

Fig. 4 Mountain Car environment. The agent begins on the line at a random position within the yellow box and must travel to the green goal state. To do so, the agent accelerates towards the first (1) key position until its velocity is reduced to zero by gravity. At this point, the agent turns and accelerates towards the second (2) key position, again, until its velocity is reduced to zero. Finally, the agent accelerates to the right again, building up velocity to reach the goal state

its velocity is reduced to zero. Finally, the agent accelerates down the hill again, building up velocity to reach the goal state. Should the agent not reach high enough up the mountain to reach the goal position, it should repeat the actions of accelerating in the opposite direction until a zero velocity is reached and turning around.

The key to the agent solving the Mountain Car problem is to increase its own velocity (v). The agent's mass (m), the magnitude of acceleration (a), and the force of gravity (G) are constant. As the agent's acceleration is lower than the gravity acting upon it, pulling the agent to the lowest point of the environment, the agent must accelerate at the correct moments, and in the correct direction, to increase its velocity. The optimal solution to the Mountain Car problem is to accelerate in the current direction of travel and take a random action when velocity is zero. An example of a rule formulation denoting this behaviour is shown in Eq. (1). The policy states the agent's next action (A_t) should be to accelerate right if its velocity is greater than 0, i.e. keep right movement, to accelerate left if its velocity is less than 0, i.e. keep left movement, and to take a random action if velocity is 0.

$$A_t = \begin{cases} +1 & v > 0 \\ -1 & v < 0 \\ \in \{-1, 1\} & v = 0 \end{cases} \quad (1)$$

The agent controlling the car has three actions to choose from in any state: to accelerate left, to accelerate right, or not to accelerate at all. The graphical representation of these possible actions is shown in Fig. 5. At each step, the agent receives a reward of -1 , and no reward when

Fig. 5 A graphical representation of the Mountain Car agent. The entire rectangle (blue and red) represents the car. The blue box indicates which action the agent has chosen to perform, either to accelerate left, to accelerate right, or not to accelerate at all and continue moving in its current direction of travel

reaching the goal state. This reward encourages the agent to reach the goal in as few steps as possible to maximise the reward.

The agent's state consists of two state variables, position and velocity, which are represented as real numbers. The position variable p represents the agent's position within the environment, and ranges linearly from -1.2 to 0.6 , i.e. $p \in [-1.2, 0.6]$, with the lowest point of the environment residing at -0.53 . The velocity of the agent v has a range of -0.07 and 0.07 , i.e. $v \in [-0.07, 0.07]$. A velocity greater than zero indicates the agent is travelling to the right or increasing its position. If the agent collides with the edge of the environment on the left ($p = -1.2$), then the agent's velocity is set to 0.

In this work, the RL agent utilises discrete state variables. Therefore, twenty bins for each state variable have been used, creating a total of 400 (20×20) states. Of these 400 states, there are some that may never be visited by the RL agent, for example, it is impossible that the agent will be at top of the left mountain ($p = -1.2$) and have a high positive velocity ($v = 0.07$).

4.2 Interaction with experiment's participants

As indicated in the previous section, twenty persons participated as trainers. The participants were university students with no experience in machine learning. Although we were aware that the number of participants was rather small, we were still able to draw significant conclusions for future experiments. During the experiments, the agents were given a low learning rate, manually tuned to extend the time which the agent would take to learn a suitable behaviour on its own. This was chosen so that the focus would be on the human's input rather than on the agent's capabilities. Both the evaluative and informative agents were given a learning rate of $\alpha = 0.25$, a discounting factor of $\gamma = 0.9$, and used an ϵ -greedy action selection strategy with an epsilon of $\epsilon = 0.05$.

In order to avoid the participants get too familiar with the environment and biased the training, each person ran only one learning episode. The Mountain Car environment

has been chosen since an optimal policy for the problem is known. To have an optimal policy for the environment allows the accuracy of the human-sourced information to be measured. Additionally, the Mountain Car problem has a low state and action space, allowing for the humans to observe the impact of their interactions relatively quickly, as the agent is likely to encounter the same state–action pairs frequently.

At the beginning of the experiments, the script given to the participants for describing the optimal solution is outlined below:

The car [agent] begins at the bottom of the valley, between two mountains. The car aims to drive to the top of the mountain on the right side. However, the car does not have the power to drive directly up the mountainside; instead, it needs to build up momentum. Momentum is gained by driving as high as possible on one side of the mountain, then turning around and accelerating in the opposite direction. When the car reaches the highest point it can on the opposite side, the process is repeated. Eventually, the car will gain enough speed to reach the top of the mountain.

5 Results and discussion

In this section, we analyse the main results obtained from the experimental scenario. First, we present user's self-evaluations in terms of the level of task understanding, engagement with the interactive agent, self-reported accuracy, and ability to follow advice. Thereafter, we discuss the characteristics of the given advice, such as frequency, accuracy, and availability.

5.1 User's self-evaluation

As previously mentioned, before each participant began interacting with the agent, they were asked to answer two questions from the questionnaire. The purpose of the questionnaire is to assess the participants understanding of the problem environment and their interactions with the agent. The first question asked was whether the participant had previously been involved in a machine learning study. None of the twenty participants reported being involved in a machine learning study previously.

Participant were then provided with a brief explanation of the dynamics of the environment and what would be the optimal behaviour. Subsequently, before starting the experiment, they were then asked to rate their level of understanding of the environment on a scale of 0–10. After interacting with the agent, the participants were asked the

Fig. 6 Participants' self-reported level of understanding of the solution and dynamics of the Mountain Car environment. The standard deviation is shown over the bars for each approach and group

same question again. Figure 6 shows the average self-reported level of understanding from the two groups of participants, i.e. evaluative and informative groups, and both before and after the experiments.

Interestingly, there is a small difference in the participant self-reported understanding of the environment before they begin interacting with the agent. The only difference in the explanation given to the two groups was the details on how they can interact with the agent. The participants giving evaluative advice were asked to rate the agent's choice of action as good or bad, while the participants giving informative advice were asked to recommend an action, either left or right. The difference in reported understanding before the experiment may indicate that evaluative advice delivery is easier to understand.

Additionally, a change in the level of participants self-reported understanding is observed after the experiment. Although the informative group shows a greater change of understanding than the evaluative group after the experiment, this is due to the initial self-reported understanding. After assisting the agent, the two groups reported a greater understanding of the environment showing no significant difference between both of them.

Moreover, after finishing the experiment, participants were also asked to report how they felt about their level of engagement with the agent. They were given three different options to answer.

- I could have spent more time interacting with the agent.
- I am happy with how much time I interacted with the agent.
- I spent too much time interacting with the agent.

Figure 7 shows the participants' reported level of engagement with the agent indicating no significant difference between the two groups. In both cases, the majority of

Fig. 7 Participants' self-reported level of engagement with the agent. Participants reported that they either *could have spent more time with the agent*, *were happy with how much time they provided*, or *spent too much time with the agent*

Fig. 8 Participants' self-reported level of advice accuracy. The informative group shows more confidence in the advice they give to the agent

participants were content with the level of engagement they had with the agent.

The participants were asked to report what they thought their level of accuracy was throughout the experiment. Participants were given six different options to answer, ranging from always incorrect to always correct. Figure 8 shows the self-reported results. The results obtained indicate that participants in the informative group were more confident in the advice they provided to the agent.

Finally, participants were asked to rate how well they thought the agent followed their advice. On a scale from 0 (never), to 10 (always), participants scored the agent's ability to follow the given advice. The obtained results, summarised in Fig. 9, show that participants using informative advice perceived the agent as better able to follow advice when compared to participants using evaluative advice.

We have computed the Student's t test to test the statistical difference between the self-reported results from

Fig. 9 Average of participants' self-reported feeling of how well the agent followed the advice provided. The participants score the agent's ability to follow advice using a scale from 0 (never), to 10 (always). The informative group perceives the agent to better follow the provided advice

the two groups. We have used the threshold of 0.1146 for the p values to consider the test statistically significant mainly due to the small number of participants. Table 1 shows the obtained t -scores along with the p values for the level of understanding of the environment, before and after the experiment, as well as the reported agent's ability to follow advice. While the differences in the self-reported understanding of the environment previous to the experiments and the perceived agent's ability to follow the provided advice are statistically significant, the difference between the two groups in the self-reported understanding of the environment after the experiments is not significant, confirming the results reported in Fig. 6.

5.2 Characteristics of the provided advice

From the assistance provided to the agent, we kept a record of the number of interactive steps and the percentage relative to the total amount of steps. Figure 10 displays the number of steps that each set of participants interacted with the agent to provide assistance. In the boxplot, the cross-marker represents the mean, dots are outliers, and the quartile calculation uses exclusive median. Overall, both groups provided advice in 9.15 steps on average; however, the data collected show a large variation in the engagement between the two types of advice. Participants providing informative advice assisted over twice as many steps than participants providing evaluative advice.

As demonstrated in previous work [61], agents assisted by informative advice learn quicker than agents assisted by evaluative advice. The increase in learning speed results in fewer steps per episode for environments with a termination condition. This decreases in steps per episode for informative-assisted agents gives fewer opportunities for the user to provide advice, as only one interaction may occur each step. As a result, the number of interactions per episode is not necessarily a suitable measure of engagement. Therefore, the number of steps in which interaction

Table 1 Student's t test for comparison of self-reported results for evaluative and informative advisors

Evaluation	t -score	p value
Understanding of the environment (before)	$t = 2.0608$	$p = 0.0541$
Understanding of the environment (after)	$t = 0.3943$	$p = 0.6980$
Agent's ability to follow advice	$t = 1.6584$	$p = 0.1146$

Fig. 10 Number of steps that participants provided advice to the learner agent on the Mountain Car environment. The amount of interactive steps is over two times for participants providing informative advice in comparison with evaluative advice **Fig. 11** Percentage of steps that participants provided advice to the learner agent on the Mountain Car problem. The percentage is computed as the ratio of interactions to interaction opportunities. The informative advice rate is almost twice as high in comparison to evaluative advice

occurred relative to the total amount of steps is used to measure engagement. Figure 11 shows the interaction rate as a percentage for the two sets of participants. As before, the boxplot uses cross-markers to represent the mean and exclusive median for quartile calculation. The interaction percentage is the ratio of interactions to interaction opportunities. Using this measurement, the length of the episode is disregarded. The results show that participants using an informative advice delivery method interact almost twice as often as their evaluative counterparts. Despite the higher rate of interaction shown by participants using informative advice, both groups self-reported they were happy with their level of engagement with the agent, as shown in Fig. 7.

Fig. 12 Percentage of interactions in which the advice provided was optimal for the state-action. Participants providing informative advice were around two times more accurate and presented less variability in comparison with participants using evaluative advice

While training the agent, the availability and accuracy of the provided assistance by the advisors were recorded. Figure 12 displays the accuracy percentage of the advice provided by each of the groups of participants. Cross-markers represent the mean and exclusive median is used for quartile calculation. An accurate interaction is one that provided the optimal advice for the agent in that state. Therefore, accuracy is a measurement of the number of correct interactions compared to the total interactions. Informative interactions are almost twice as accurate in comparison to evaluative interactions and also show much less variability. These results also reflect the self-reported level of advice accuracy shown in Fig. 8.

We have also computed the Student's t test to test the statistical difference between the obtained results in terms of the advice provided from the two groups. Table 2 shows the obtained t -scores along with the p values for the average interactive steps, the average interactive rate, and the accuracy of the advice provided. Although there exist

Table 2 Student's t -test for comparison of the provided advice from evaluative and informative advisors

Evaluation	t -score	p value
Average interactive steps	$t = 2.2530$	$p = 0.0370$
Average interactive rate	$t = 1.6828$	$p = 0.1097$
Accuracy of the advice provide	$t = 14.5772$	$p = 2.0778e-11$

statistical differences between the two groups for the average interactive steps and the average interactive rate, this is much clearer in the accuracy of the advice provided given the low p value.

One hypothesis for the large difference in accuracy is latency. In this context, latency is the time it takes for the human to decide on the advice to provide and then input it into the agent. It is possible that if the human is too late in providing advice, then the advice will inadvertently be provided to the state after the one intended. For the Mountain Car environment, a late interaction is more likely to remain accurate in the next state for informative advice than it is for evaluative advice. This is due to the layout of the state space and the nature of untrained agents. The optimal action for a state in the Mountain Car environment is likely to be the same as its neighbouring states. This is due to the optimal behaviour being to accelerate in a single direction until velocity reaches 0. Therefore, a recommended action that is received in the state after the one intended is likely to be the correct action, regardless of latency. This does not apply to evaluative advice. The participants assisting the evaluative agent do not provide a recommended action; instead, they critique the agent's last choice. An untrained agent has a largely random action selection policy and is therefore not likely to choose the same action twice in a row. As the agent's chosen action may have changed by the time it receives advice from the user, the accuracy is more affected.

This hypothesis is supported by the state breakdown of the advice accuracy. Figure 13 shows the accuracy of participants' advice for each state in the environment for

(a) informative and (b) evaluative advice, respectively. The darker the colour, the more accurate the advice supplied by the participants for that state. The comparison of the two heatmaps supports the earlier observations of the accuracy shown in Fig. 12; informative is much more accurate than evaluative advice. The informative advice method (Fig. 13a) shows that the states with the most inaccuracy are in the middle of the environment, where the optimal action changes. This inaccuracy is likely not due to poor participant knowledge, but rather providing delayed advice, after the agent has moved beyond the centre threshold.

The evaluative advice method (Fig. 13b) shows that accuracy differs wildly across the environment and does not have an obvious pattern. The poor result for accuracy of evaluative advice is likely due to the latency of advice delivery coupled with the lower probability that the advice will still be accurate to the following state compared to informative advice. Additionally, evaluative advice may have lower accuracy as it requires the human assessing each state–action pair. On the other hand, informative advice may require less time assessing each state, as the human may be following a set of rules for action recommendation and that the next state is easier to predict compared to the agent's next action choice.

Figure 14 shows the availability of participants' advice for each state in the environment for (a) informative and (b) evaluative advice, respectively. Availability in this context is a measure of how often the user provides advice in a state compared to the number of times the agent visited the state. The darker a state is on the heatmaps, the more

Fig. 13 State-based accuracy of informative and evaluative participants for the Mountain Car environment. Informative advice is in general more accurate than evaluative advice, except in states in the middle of the environment, where the optimal action changes

Fig. 14 State-based availability of informative and evaluative participants for the Mountain Car environment. Agents following participants using informative advice achieved higher velocities in the

often the user provides advice for that state. The agent that was assisted by informative advice (Fig. 14a) was able to achieve higher velocities in the environment and, as a result, visited more states in comparison with the evaluative advice method (Fig. 14b). One pattern that can be observed in the results is that the states on the edges show higher advice availability than those in the centre. These edge states are visited when the agent has learned a suitable behaviour, making the evaluation and recommendation of actions easier on the user, and increasing engagement. The edge states tend to be the last states the users provided advice, before voluntarily ending the experiment.

Finally, we tested the presence of the reward bias of the participants providing evaluative advice as it has been reported in existing literature [21]. In this regard, a deviation from fifty percent indicates reward bias, i.e. above 50%, means that the advisor provided more positive evaluation than negative evaluation. On average, participants provided 66.22% of positive advice, with a minimum rate of 57.14% and a maximum rate of 100.00% of positive evaluations. Figure 15 shows the reward bias of the participants providing evaluative advice. In general, participants provided much more positive advice, confirming prior findings that people are more likely to provide feedback on actions they view as correct than on incorrect actions.

environment, and as a consequence, more states were visited in comparison to the evaluative advice approach

Fig. 15 Reward bias of evaluative advice. Above 50% means that the advisor provided more positive evaluation than negative evaluation

6 Conclusions

The human trial performed in this work has investigated the engagement of human advice givers when assisting interactive reinforcement learning agents. The assessment was performed using two methods for providing assistance, namely evaluative and informative advice. Evaluative advice assesses the past performance of an agent, while informative advice supplements future decision-making. Previous work in the field has compared the performance of interactive reinforcement learning agents under the

influence of each assistance method, finding that informative-assisted agents learn faster. However, to the best of our knowledge, studies on human engagement when providing advice using each assistance method have not been performed.

The results obtained from the human trial show that advice givers providing informative advice outperformed those that used evaluative advice. Humans using an informative advice-giving method demonstrated more accurate advice, assisted the agent for longer, and provided advice more frequently. Additionally, informative advice givers rated the ability of the agent to follow advice more highly, perceived their own advice to be of higher accuracy, and were similarly content with their engagement with the agent as the evaluative advice-giving participants.

In this work, we have used a well-known environment and the advice was simply selecting the very next action. Therefore, this reduces considerably the subjectivity; however, there is a possibility of incorrect advice. The inconsistent advice problem has been previously studied [64, 72], and it has been shown that even a small amount of incorrect advice may be quite detrimental for the learning process. This is especially relevant when the problem is more complex or the trainer is not familiar with the problem domain. In such cases, human advice givers need special additional preparation in order to avoid subjectivity. In this paper, advice is applied only to the current state and time-step, and so incorrect advice will not have a long-term impact. However, the problem becomes a major issue when general advice is given. For instance, in the persistent rule-based IntRL approach [24] advice is retained by the agent and generalised over other similar states. To deal with such a problem, that approach incorporates probabilistic policy reuse [73] which uses a decay factor in order to forget the given advice and trust more in the learned agent's policy, unless new advice is given.

Future work will also consider the use of simulated users as a method of replicating the general behaviour of participants from this experiment. Including simulated users would allow for faster experiments, keeping experimental conditions under control, and repeat the process as many times as needed. The findings from this study can be used to create simulated users which more closely reflect the behaviour of actual human advisers.

Funding Open Access funding enabled and organized by CAUL and its Member Institutions.

Declarations

Conflict of interest The authors declare that they have no conflict of interest.

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. Sutton RS, Barto AG (2018) Reinforcement learning: an introduction. MIT Press, Cambridge
2. Kaelbling LP, Littman ML, Moore AW (1996) Reinforcement learning: a survey. *J Artif Intell Res* 4:237–285
3. Giannoccaro I, Pontrandolfo P (2002) Inventory management in supply chains: a reinforcement learning approach. *Int J Prod Econ* 78(2):153–161
4. Lepenioti K, Bousdekis A, Apostolou D, Mentzas G (2021) Human-augmented prescriptive analytics with interactive multi-objective reinforcement learning. *IEEE Access* 9:100677–100693
5. Machalek D, Quah T, Powell KM (2021) A novel implicit hybrid machine learning model and its application for reinforcement learning. *Comput Chem Eng* 155:107496
6. Cruz F, Acuña G, Cubillos F, Moreno V, Bassi D (2007) Indirect training of grey-box models: application to a bioprocess. In: *International symposium on neural networks*. Springer, pp 391–397
7. Kitano H, Asada M, Kuniyoshi Y, Noda I, Osawa E, Matsuura H (1997) RoboCup: a challenge problem for AI. *AI Mag* 18(1):73
8. Churamani N, Cruz F, Griffiths S, Barros P (2016) iCub: learning emotion expressions using human reward. In: *Proceedings of the workshop on bio-inspired social robot learning in home scenarios*. IEEE/RSJ IROS, p 2
9. Lee K, Smith LM, Abbeel P (2021) PEBBLE: feedback-efficient interactive reinforcement learning via relabeling experience and unsupervised pre-training. In: *Proceedings of the 38th international conference on machine learning*. PMLR, pp 6152–6163
10. Tesauro G (1994) TD-Gammon, a self-teaching backgammon program, achieves master-level play. *Neural Comput* 6(2):215–219
11. Barros P, Tanevska A, Cruz F, Sciutti A (2020) Moody learners-explaining competitive behaviour of reinforcement learning agents. In: *2020 joint IEEE 10th international conference on development and learning and epigenetic robotics (ICDL-EpiRob)*. IEEE, pp 1–8
12. Mankowitz DJ, Dulac-Arnold G, Hester T (2019) Challenges of real-world reinforcement learning. In: *ICML workshop on real-life reinforcement learning*, p 14
13. Cruz F, Wüppen P, Fazrie A, Weber C, Wermter S (2018) Action selection methods in a robotic reinforcement learning scenario. In: *2018 IEEE Latin American conference on computational intelligence (LA-CCI)*. IEEE, pp 13–18
14. Thomaz AL, Hoffman G, Breazeal C (2005) Real-time interactive reinforcement learning for robots. In: *Proceedings of association for the advancement of artificial intelligence conference AAAI, workshop on human comprehensible machine learning*, pp 9–13

15. Li G, Gomez R, Nakamura K, He B (2019) Human-centered reinforcement learning: a survey. *IEEE Trans Hum Mach Syst* 49(4):337–349
16. Brod G, Werkle-Bergner M, Shing YL (2013) The influence of prior knowledge on memory: a developmental cognitive neuroscience perspective. *Front Behav Neurosci* 7:139
17. Subramanian K, Isbell Jr CL, Thomaz AL (2016) Exploration from demonstration for interactive reinforcement learning. In: *Proceedings of the 2016 international conference on autonomous agents & multiagent systems*. International Foundation for Autonomous Agents and Multiagent Systems, pp 447–456
18. Bignold A, Cruz F, Dazeley R, Vamplew P, Foale C (2021) An evaluation methodology for interactive reinforcement learning with simulated users. *Biomimetics* 6(1):13
19. Arzate Cruz C, Igarashi T (2020) A survey on interactive reinforcement learning: design principles and open challenges. In: *Proceedings of the 2020 ACM designing interactive systems conference*, pp 1195–1209
20. Bignold A, Cruz F, Taylor ME, Brys T, Dazeley R, Vamplew P, Foale C (2021) A conceptual framework for externally-influenced agents: an assisted reinforcement learning review. *J Amb Intell Hum Comput* 56:1–24
21. Amershi S, Cakmak M, Knox WB, Kulesza T (2014) Power to the people: the role of humans in interactive machine learning. *AI Mag* 35(4):105–120
22. Thomaz AL, Breazeal C (2008) Teachable robots: understanding human teaching behavior to build more effective robot learners. *Artif Intell* 172(6–7):716–737
23. Cakmak M, Thomaz AL (2010) Optimality of human teachers for robot learners. In: *2010 IEEE 9th international conference on development and learning (ICDL)*. IEEE, pp 64–69
24. Bignold A, Cruz F, Dazeley R, Vamplew P, Foale C (2021) Persistent rule-based interactive reinforcement learning. *Neural Comput Appl* 89:1–18
25. Moreira I, Rivas J, Cruz F, Dazeley R, Ayala A, Fernandes B (2020) Deep reinforcement learning with interactive feedback in a human-robot environment. *Appl Sci* 10(16):5574
26. Cruz F, Parisi GI, Wermter S (2018) Multi-modal feedback for affordance-driven interactive reinforcement learning. In: *2018 International joint conference on neural networks (IJCNN)*. IEEE, pp 1–8
27. Sharma M, Holmes MP, Santamaría JC, Irani A, Isbell CL Jr, Ram A (2007) Transfer learning in real-time strategy games using hybrid CBR/RL. *IJCAI* 7:1041–1046
28. Taylor ME, Stone P, Liu Y (2007) Transfer learning via inter-task mappings for temporal difference learning. *J Mach Learn Res* 8(Sep):2125–2167
29. Shin YS, Niv Y (2020) Biased evaluations emerge from inferring hidden causes. *PsyArXiv preprint psyarxiv:10.31234*
30. Grzes M (2017) Reward shaping in episodic reinforcement learning. In: *Proceedings of the sixteenth international conference on autonomous agents and multiagent systems (AAMAS 2017)*. ACM, pp 565–573
31. Marom O, Rosman BS (2018) Belief reward shaping in reinforcement learning. In: *Proceedings of the thirty-second AAAI conference on artificial intelligence*. AAAI, pp 3762–3769
32. Millán-Arias C, Fernandes B, Cruz F, Dazeley R, Fernandes S (2020) A robust approach for continuous interactive reinforcement learning. In: *Proceedings of the 8th international conference on human-agent interaction*, pp 278–280
33. Millán-Arias C, Fernandes B, Cruz F, Dazeley R, Fernandes S (2021) A robust approach for continuous interactive actor-critic algorithms. *IEEE Access* 9:104242–104260
34. Shah P, Hakkani-Tur D, Heck L (2016) Interactive reinforcement learning for task-oriented dialogue management. *Workshop on deep learning for action and interaction*. *Adv Neural Inf Process Syst* 2016:1–11
35. Goodfellow I, Bengio Y, Courville A, Bengio Y (2016) *Deep learning*, vol 1. MIT Press, Cambridge
36. LeCun Y, Bengio Y, Hinton G (2015) *Deep learning*. *Nature* 521(7553):436
37. Kang B, Compton P, Preston P (1995) Multiple classification ripple down rules: evaluation and possibilities. In: *Proceedings 9th Banff knowledge acquisition for knowledge based systems workshop*, vol 1, pp 17–21
38. Compton P, Edwards G, Kang B, Lazarus L, Malor R, Menzies T, Preston P, Srinivasan A, Sammut C (1991) Ripple down rules: possibilities and limitations. In: *Proceedings of the sixth AAAI knowledge acquisition for knowledge-based systems workshop*, Calgary, Canada, University of Calgary, pp 6–1
39. Taylor ME, Stone P (2009) Transfer learning for reinforcement learning domains: a survey. *J Mach Learn Res* 10(Jul):1633–1685
40. Littman ML (1994) Markov games as a framework for multi-agent reinforcement learning. In: *Machine learning proceedings 1994*. Elsevier, pp 157–163
41. Tan M (1993) Multi-agent reinforcement learning: independent versus cooperative agents. In: *Proceedings of the tenth international conference on machine learning*, pp 330–337
42. Cruz F, Parisi GI, Wermter S (2016) Learning contextual affordances with an associative neural architecture. In: *Proceedings of the European symposium on artificial neural network. Computational intelligence and machine learning ESANN, UCLouvain*, pp 665–670
43. Argall B, Browning B, Veloso M (2007) Learning by demonstration with critique from a human teacher. In: *Proceedings of the ACM/IEEE international conference on human-robot interaction*. ACM, pp 57–64
44. Millán C, Fernandes B, Cruz F (2019) Human feedback in continuous actor-critic reinforcement learning. In: *Proceedings of the European symposium on artificial neural networks, computational intelligence and machine learning ESANN. ESANN*, pp 661–666
45. Ayala A, Henríquez C, Cruz F (2019) Reinforcement learning using continuous states and interactive feedback. In: *Proceedings of the international conference on applications of intelligent systems*, pp 1–5
46. Thomaz AL, Breazeal C (2007) Asymmetric interpretations of positive and negative human feedback for a social learning agent. In: *The 16th IEEE international symposium on robot and human interactive communication*, 2007. *RO-MAN 2007*. IEEE, pp 720–725
47. Ng AY, Harada D, Russell S (1999) Policy invariance under reward transformations: theory and application to reward shaping. In: *Proceedings of the international conference on machine learning ICML*, vol 99, pp 278–287
48. Brys T, Nowé A, Kudenko D, Taylor ME (2014) Combining multiple correlated reward and shaping signals by measuring confidence. In: *Proceedings of the association for the advancement of artificial intelligence conference. AAAI*, pp 1687–1693
49. Marthi B (2007) Automatic shaping and decomposition of reward functions. In: *Proceedings of the international conference on machine learning ICML*. ACM, pp 601–608
50. Rosman B, Ramamoorthy S (2014) Giving advice to agents with hidden goals. In: *2014 IEEE international conference on robotics and automation (ICRA)*. IEEE, pp 1959–1964
51. Huang J, Juan R, Gomez R, Nakamura K, Sha Q, He B, Li G (2021) Gan-based interactive reinforcement learning from demonstration and human evaluative feedback. *arXiv preprint arXiv:2104.06600*
52. Knox WB, Stone P (2010) Combining manual feedback with subsequent MDP reward signals for reinforcement learning. In: *Proceedings of the 9th international conference on autonomous*

- agents and multiagent systems, vol 1. International Foundation for Autonomous Agents and Multiagent Systems, pp 5–12
53. Knox WB, Stone P (2009) Interactively shaping agents via human reinforcement: the TAMER framework. In: Proceedings of the fifth international conference on knowledge capture. ACM, pp 9–16
 54. MacGlashan J, Ho MK, Loftin R, Peng B, Wang G, Roberts DL, Taylor ME, Littman ML (2017) Interactive learning from policy-dependent human feedback. In: Proceedings of the 34th international conference on machine learning, vol 70, pp 2285–2294
 55. Arumugam D, Lee JK, Saskin S, Littman ML (2019) Deep reinforcement learning from policy-dependent human feedback. arXiv preprint [arXiv:1902.04257](https://arxiv.org/abs/1902.04257)
 56. Kessler Faulkner T, Gutierrez RA, Short ES, Hoffman G, Thomaz AL (2019) Active attention-modified policy shaping: socially interactive agents track. In: Proceedings of the international conference on autonomous agents and multiagent systems AAMAS. International Foundation for Autonomous Agents and Multiagent Systems, pp 728–736
 57. Behboudian P, Satsangi Y, Taylor ME, Harutyunyan A, Bowling M (2020) Useful policy invariant shaping from arbitrary advice. In: AAMAS adaptive and learning agents workshop ALA 2020, p 9
 58. Lin J, Ma Z, Gomez R, Nakamura K, He B, Li G (2020) A review on interactive reinforcement learning from human social feedback. *IEEE Access* 8:120757–120765
 59. Cruz F, Wüppen P, Magg S, Fazrie A, Wermter S (2017) Agent-advising approaches in an interactive reinforcement learning scenario. In: 2017 joint IEEE international conference on development and learning and epigenetic robotics (ICDL-EpiRob). IEEE, pp 209–214
 60. Grizou J, Lopes M, Oudeyer P-Y (2013) Robot learning simultaneously a task and how to interpret human instructions. In: Proceedings of the joint IEEE international conference on development and learning and epigenetic robotics ICDL-EpiRob. IEEE, pp 1–8
 61. Griffith S, Subramanian K, Scholz J, Isbell C, Thomaz AL (2013) Policy shaping: integrating human feedback with reinforcement learning. In: Advances in neural information processing systems, pp 2625–2633
 62. Pilarski PM, Sutton RS (2012) Between instruction and reward: human-prompted switching. In: AAAI fall symposium series: robots learning interactively from human teachers, pp 45–52
 63. Amir O, Kamar E, Kolobov A, Grosz B (2016) Interactive teaching strategies for agent training. In: Proceedings of the international joint conference on artificial intelligence IJCAI, pp 804–811
 64. Cruz F, Magg S, Nagai Y, Wermter S (2018) Improving interactive reinforcement learning: what makes a good teacher? *Connect Sci* 30(3):306–325
 65. Kamar E, Hacker S, Horvitz E (2012) Combining human and machine intelligence in large-scale crowdsourcing. In: Proceedings of the 11th international conference on autonomous agents and multiagent systems. International Foundation for Autonomous Agents and Multiagent Systems, vol 1, pp 467–474
 66. Knox WB, Stone P (2012) Reinforcement learning from human reward: discounting in episodic tasks. In: 2012 IEEE RO-MAN: the 21st IEEE international symposium on robot and human interactive communication. IEEE, pp 878–885
 67. Knox WB, Stone P (2013) Learning non-myopically from human-generated reward. In: Proceedings of the 2013 international conference on intelligent user interfaces. ACM, pp 191–202
 68. Cakmak M, Chao C, Thomaz AL (2010) Designing interactions for robot active learners. *IEEE Trans Auton Ment Dev* 2(2):108–118
 69. Guillory A, Bilmes JA (2011) Simultaneous learning and covering with adversarial noise. *ICML* 11:369–376
 70. Guillory A, Bilmes JA (2011) Online submodular set cover, ranking, and repeated active learning. In: Advances in neural information processing systems, pp 1107–1115
 71. Moore AW, Birnbaum L, Collins G (1991) Variable resolution dynamic programming: efficiently learning action maps in multivariate real-valued state-spaces. In: Proceedings of the eighth international conference on machine learning, pp 333–337
 72. Kessler Faulkner TA, Thomaz A (2021) Interactive reinforcement learning from imperfect teachers. In: Companion of the 2021 ACM/IEEE international conference on human-robot interaction, pp 577–579
 73. Fernández F, Veloso M (2006) Probabilistic policy reuse in a reinforcement learning agent. In: Proceedings of the fifth international joint conference on autonomous agents and multi-agent systems. ACM, pp 720–727

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.