



# Experiential Explanations for Reinforcement Learning

Amal Alabdulkarim<sup>1</sup> · Madhuri Singh<sup>1</sup> · Gennie Mansi<sup>1</sup> · Kaely Hall<sup>1</sup> · Upol Ehsan<sup>1</sup> · Mark O. Riedl<sup>1</sup>

Received: 1 July 2024 / Accepted: 16 December 2024  
© The Author(s) 2025

## Abstract

Reinforcement learning (RL) systems can be complex and non-interpretable, making it challenging for non-AI experts to understand or intervene in their decisions. This is due in part to the sequential nature of RL in which actions are chosen because of their likelihood of obtaining future rewards. However, RL agents discard the qualitative features of their training, making it difficult to recover user-understandable information for “why” an action is chosen. We propose a technique *Experiential Explanations* to generate counterfactual explanations by training *influence predictors* along with the RL policy. Influence predictors are models that learn how different sources of reward affect the agent in different states, thus restoring information about how the policy reflects the environment. Two human evaluation studies revealed that participants presented with Experiential Explanations were better able to correctly guess what an agent would do than those presented with other standard types of explanation. Participants also found that Experiential Explanations are more understandable, satisfying, complete, useful, and accurate. Qualitative analysis provides information on the factors of Experiential Explanations that are most useful and the desired characteristics that participants seek from the explanations.

**Keywords** Explainable reinforcement learning · Explainable artificial intelligent · Human–artificial intelligence interaction · Explainable sequential decision-making

## 1 Introduction

Reinforcement learning (RL) techniques are becoming increasingly popular in critical domains such as robotics and autonomous vehicles. However, applications such as autonomous cars and socially assistive robots in health-care and home settings are expected to interact with non-AI experts. People often seek explanations for abnormal events or observations to improve their understanding of the agent

and better predict and control its actions [1]. This need for an explanation will likely arise when interacting with these systems, whether it is a car taking an unexpected turn or a robot performing a task unconventionally. Without an explanation, users can find it difficult to trust the agent’s ability to act safely and reasonably [2], especially if they have limited artificial intelligence expertise. Even in the case where the agent is operating optimally and without failure, the agent’s optimal behavior may not match the user’s expectations, resulting in confusion and lack of trust. Therefore, RL systems need to provide good explanations without compromising their performance. To address this need, *explainable RL* has become a rapidly growing area of research with many open challenges [3–9].

But what makes a good explanation for end-users who cannot change the underlying model? How do we generate explanations if we cannot modify the underlying RL model or degrade its performance? In RL, an agent learns dynamically to maximize its reward through a series of experiences interacting with the environment [10]. Through these interactions, the agent builds a utility model,  $V(s, a)$ , which estimates the future reward that can be expected to be achieved if a particular action  $a$  is performed from a particular state  $s$ .

✉ Amal Alabdulkarim  
amal@gatech.edu

Madhuri Singh  
msingh365@gatech.edu

Gennie Mansi  
gennie.mansi@gatech.edu

Kaely Hall  
khall33@gatech.edu

Upol Ehsan  
ehsanu@gatech.edu

Mark O. Riedl  
riedl@gatech.edu

<sup>1</sup> School of Interactive Computing, Georgia Institute of Technology, Atlanta, USA



**Fig. 1** In an interaction between the user and the agent, the user expected the agent to go up, but the agent went down instead. The user asks the system for an explanation. The figure shows two types of explanations the user can receive. **A** is an explanation that can be

While an RL agent learns from its experiences during training, those experiences are inaccessible after their training is over [4]. This is because  $V(s, a)$  summarizes future experiences as a single, real-valued number; this is all that is needed to execute a policy  $\pi = \text{argmax}_a V(s, a)$ . For example, consider a robot that receives a negative reward for being close to the stairs and thus learns that states along an alternative route have higher utility. When it comes time to figure out *why* the agent executed one action trajectory over another, the policy is devoid of any information from which to construct an explanation beyond the fact that some actions have lower expected utility than others.

Explanations need not only address agents' failures. The user's need for explanations can also arise when the agent makes an unexpected decision. Requests for explanations thus often manifest themselves as "why-not" questions referencing a counterfactual: "why did not the agent make a different decision?" Explanations in sequential decision-making environments can help users update their mental models of the agent or identify how to change the environment so that the agent performs as expected.

We propose a technique, **Experiential Explanations**, in which deep RL agents generate explanations in response to on-demand, local, counterfactuals proposed by users. These explanations qualitatively contrast the agent's intended action trajectory with a trajectory proposed by the user and specifically link agent behavior to environmental contexts. The agent maintains a set of models—called **influence predictors**—of how different sparse reward states influence the agent's state-action utilities and, thus, its policy. Influence predictors are trained alongside an RL policy with a specific focus on retaining details about the training experience. The influence models are "outside the black box," but provide information with which to interpret the agent's black-box policy in a human-understandable fashion.

derived directly from the agent's policy. **B** is an explanation that can be derived from the agent's policy with the assistance of additional models called negative and positive *influence predictors*

Consider the illustrative example in Fig. 1 where the user observes the agent go down around the wall. The user expects the agent to go up, but the agent goes down instead. The user asks why the agent did not choose up, which appears to lead to a shorter route to the goal. One possible and correct explanation would be that the agent's estimated expected utility for the down action is higher than that for the up action. However, this is not information that a user can easily act upon. The alternative, Experiential Explanation, states that up will pass through a region that is in proximity to dangerous locations. The user can update their understanding that the agent prefers to avoid stairs. The user can also understand how to take an explicit action to change the environment, such as blocking stairs—an action a non-expert can take to influence agent behavior when they do not have the ability to directly change the agent. Our technique focuses on post hoc, real-time explanations of agents operating in situ, instead of pre-explanation of agent plans [11], although the two are closely related.

We evaluate our Experiential Explanations technique with studies with human participants. We show that explanations generated by our technique allow users to better understand and predict agent actions and are found to be more useful. We additionally perform a qualitative analysis of participant responses to elicit *how* users use the information in the explanations in our technique and baseline alternatives to reason about the agent's behavior. This not only provides evidence for Experiential Explanations, but provides a blueprint to understand the human factors of other explanation techniques.

## 2 Background and related work

Reinforcement learning is an approach to learning in which an agent attempts to maximize its reward through feedback from a trial-and-error process. RL is suitable for Markov

decision processes, which can be represented as a tuple  $M = \langle S, A, T, R, \gamma \rangle$  where  $S$  is the set of possible world states,  $A$  is the set of possible actions,  $T$  is the state transition function, which determines the probability of the following state  $P(S)$  as a function of the current state and action.  $T : S \times A \rightarrow P(S)$ ,  $R$  is the reward function  $R : S \times A \rightarrow \mathbb{R}$ , and  $\gamma$  is a discount factor  $0 \leq \gamma \leq 1$ . RL first learns a policy  $\pi : S \rightarrow A$ , which defines which actions should be taken in each state. Deep RL uses deep neural networks to estimate the expected future utility such that  $\pi(s) = \text{argmax}_a V_\theta(s, a)$  where  $\theta$  are the model parameters.

As deep reinforcement learning is increasingly used in sensitive areas such as healthcare and robotics, Explainable RL (XRL) is becoming a crucial research field. However, the reinforcement learning paradigm differs substantially from other machine learning paradigms, such as supervised learning, in that they solve sequential decision-making problems where each decision an agent makes is not an end to itself but establishes conditions for future actions; understanding any one decision may be dependent on understanding prior decisions as well as future potential decisions. For an overview of XRL, please refer to Milani et al. [8]. We highlight the most relevant work on XRL to our approach.

Explanation via reward decomposition breaks down rewards into a sum of semantically meaningful reward categories, allowing actions to be evaluated based on the trade-offs between these categories [12]. Many XRL techniques generate explanations by decomposing utilities in the agent’s network and showing information about the components of the utility score for an agent’s state [13–15]. Our technique is related to decompositions in that we learn the influence of different sources of reward on utilities. However, we differ in that we approximate the decomposition using external models, avoiding the trade-off between using a semi-interpretable system or a better black-box system.

Explanations are inherently contrastive, as they are framed in the context of particular counterfactual comparisons. Rather than merely questioning why event P took place, individuals ask why event P occurred rather than event Q [1]. Contrastive approaches such as van der Waa et al. [16] leverage an interpretable model and a learned transition model. Madumal et al. [17] explain local actions with a causal structural model as an action influence graph to encode the causal relations between variables. Sreedharan et al. [18] generate contrastive explanations using a partial symbolic model approximation, including knowledge about concepts, preconditions, and action costs. These methods provided answers to user’s local why-not questions. While our approach still produces counterfactual explanations, we minimize the use of predefined symbolic or *a priori* environment knowledge. Other work looked at utilizing counterfactual explanations for other explanation goals. For example, Frost et al. [19] trained an exploration policy

to generate test time trajectories to explain how the agent behaves in unseen states after training. Explanation trajectories help users understand how the agent would perform under new conditions. While this method gives users a global understanding of the agent, our method focuses on providing local counterfactual explanations of an agent’s action. Olson et al. [20] explain the local decisions of an agent by demonstrating a counterfactual state in which the agent takes a different action to illustrate the minimal alteration necessary for a different result. Huber et al. [21] further this concept by creating counterfactual states with an adversarial learning model. Our technique, however, concentrates on the qualitative distinctions between trajectories and not on the states.

*Rationale generation*, introduced by Ehsan et al. [22], is an approach to explanation generation for sequential environments in which explanations approximate how humans would explain what they would do in the same situation as the agent. Rationale generation does not “open the black box” but looks at the agent’s state and action and applies a model trained from human explanations. Rationales do not guarantee faithfulness, though studies show they can still provide actionable insights. Like rationale generation, our models train alongside the agent. However, our technique also achieves a high degree of faithfulness. State2Explanation (S2E), introduced in Das et al. [23], demonstrates generating more faithful explanations by using them to inform the agent’s reward shaping and requires the derivation of the concepts from mathematical representations and expert domain knowledge. By aligning state-action pairs with high-level concepts, S2E uses natural language templates to provide actionable insights during training and deployment. In our approach, we avoid intervening with the learning of the agent while preserving a high level of faithfulness by confining the learning only from the agent’s transitions and trying to minimize the reliance on domain knowledge to only the names and effects of reward components.

### 3 Experiential Explanations

*Experiential Explanations* are contrastive explanations generated for local decisions with minimal reliance on structured inputs and without imposing limitations on the agent or the RL algorithm. Our explanation technique uses additional models called *influence predictors* that learn how the sources of the sparse rewards affect the agent’s utility predictions. The agent’s learned policy learns the utilities with respect to all rewards. Influence predictors can tell us how strongly or weakly any source of reward (positive or negative) is impacting states in the state space that the agent believes it will pass through. The influence predictors provide the additional context that explains the agent’s choices through finer-grained

utility detail. As in Fig. 1, we see the explanation referencing the agent's relationship to the environment in terms of negative and positive elements.

The main components of Experiential Explanations are the trained influence predictors and the two modules of explanation generation. The first module starts with the user proposing an alternative action and asking why the agent did something else, seeking a counterfactual explanation. This triggers simulation of the trajectories of both actions. For example, in Fig. 1 the agent's action is down, but the user asks about going up. The generated trajectories account for those actions and continue until the end of the respective episodes. In the second module, we compare the states along a specified horizon of the simulated trajectories using influence predictors.

In Sect. 3.1, we dive into the deeper details of the influence predictors. In Sect. 3.2, we demonstrate how they can be used in the explanation generation modules to generate explanations.

### 3.1 Influence predictors

*Influence predictors* reconstruct the effects of rewards received on the utility learned by the agent during training. These models are trained alongside the agent to predict the strength of the influence of different sources of reward on the agent by learning a value function that outputs the expected utility from each source of reward:  $U_c(s, a)$  where each  $c \in C$  is a distinct source (or class) of reward. For example, a terminal goal state might be a source of positive reward, stairs might be a source of negative reward, and other dangerous objects may be other sources of negative reward. In a more complex environment, multiple sources of positive and negative rewards can be helpful in describing the agent's plan. For instance, a robot attempting to make a cup of coffee would receive positive rewards for obtaining a clean mug, hot milk, and coffee, then combining the milk and coffee in the mug and stirring them. Negative influences could include getting a dirty cup or adding salt.

Each source of reward has its own influence predictor model. Typically, one class is associated with a goal state (or states required for task completion) and produces positive rewards. Sometimes an intermediate positive reward is given when it is known in advance that the task requires particular actions or states to be traversed, as in the case of reward shaping [24]. Other classes are associated with states that are *a priori* known to be detrimental to the task and produce negative rewards.

In principle, influence predictors can be used with any agent architecture and optimization algorithm as long as we can observe its transitions during training. Influence predictors are trained similarly to an RL agent (e.g., standard Deep  $Q$ -networks (DQN) [25], or A2C [26]) but with two dif-

ferences. First, influence predictors are *not* used to guide the agent's exploration during training or inference. They rely on the main policy's training loop to generate transitions. Second, each influence predictor model can receive only rewards from one class of rewards. Influence predictors are trained using loss based on a modified Bellman equation:

$$U_c(s_t, a_t) = |r_{t+1}^c| + \gamma \max_{a'} U_c(s_{t+1}, a') \quad (1)$$

where  $r_{t+1}^c$  is the reward associated with  $c$  delivered to the agent upon transition to state  $s_{t+1}$ . The absolute value reinterprets the reward as an influence, since the influence predictor is now predicting the utility of attaining the reward, even if it is negative. As a consequence of ever receiving only the absolute value of a class of rewards,  $U_c(\cdot)$  is an estimate of the strength of  $c$  in  $s$  and  $a$ .

The training process for influence predictors is shown in Algorithm 1. The agent, controlled by its base policy learning algorithm, interacts with the environment, executing actions, and receiving state observations. State transitions, actions, and rewards are recorded and used to populate separate buffers for each influence predictor. For model updates, we sample the transitions from each designated buffer separately and backpropagate loss in each influence predictor.

---

#### Algorithm 1 Influence predictor and agent training

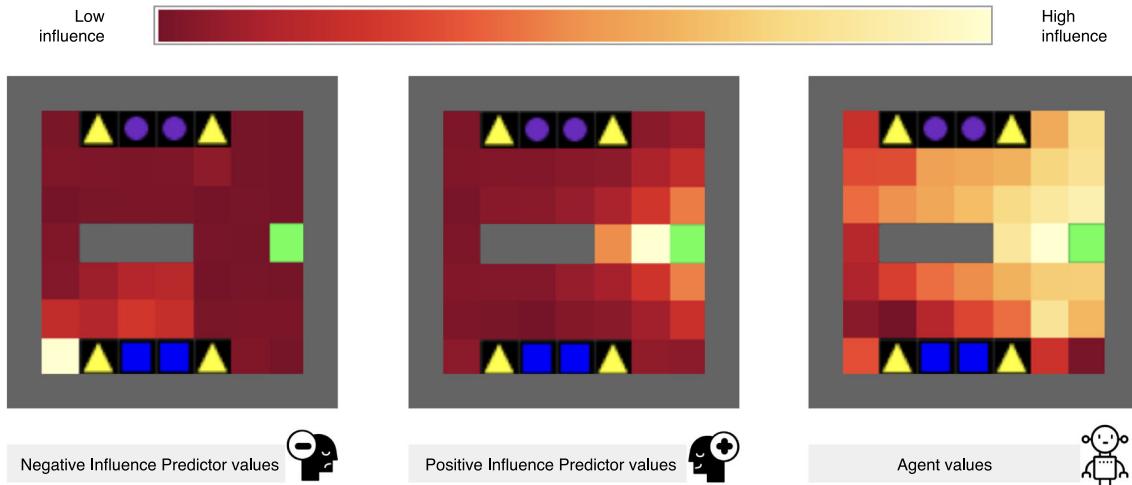
```

for each  $c \in C$  do ▷  $C$  is the reward sources
    Initialize buffer memory  $D_c$  to capacity  $N$ 
    Initialize action-value function  $U_c$  with random weights
end for
for episode = 1, M do
     $transitions_T \leftarrow agent.transitions(T)$  ▷ The agent's transition
    data from its buffer
    for t=0, T do
        for each  $c \in C$  do
             $transition_t[c] \leftarrow transitions_T[t]$ 
            if  $reward_t \in c$  then
                 $transition_t[c][reward] \leftarrow$  ←
                 $|transitions_T[t][reward]|$ 
            else
                 $transition_t[c][reward] \leftarrow 0$ 
            end if
             $(transition_t[c]) \rightarrow D_c$ 
        end for
    end for
    for each  $c \in C$  do
         $U_c \leftarrow update(D_c)$ 
    end for
     $agent \leftarrow update(transitions_T)$ 
end for

```

---

To illustrate, Fig. 2 shows an agent operating in an environment with two reward classes: a positive reward class for the terminal goal and a negative reward class for the blue squares—hence, two influence predictors. The influence map on the left was generated by querying the negative influ-



**Fig. 2** This figure shows an example of the agent’s learned maximum learned utility value at each state and the two influence predictors values visualized on Heatmaps. The blue squares are the source of negative rewards, whereas a green square is the source of positive rewards. The

purple and yellow shapes have no effect. Intuitively, one can see that the agent’s policy (right) is closely associated with the positive influence values and the *opposite* of the negative influence values

ence predictor for each state and maxing across all possible actions. The negative influence trends are stronger closer to the blue squares. The influence map in the middle is generated from the positive influence predictor. The map on the right is the agent’s main policy model, trained on all reward classes as normal. Even though the visualization aggregates action utilities, one can see that the agent will prefer a trajectory in the upper part of the map.

### 3.2 Explanation generation

At any time during the execution of the policy, the user can request a counterfactual explanation by proposing an action ( $a_{\text{user}}$ ) rather than the one the agent chose ( $a_{\text{agent}}$ ). Given the action chosen by the agent and the alternative action proposed by the user, the explanation generation process proceeds in three phases.

First, we generate and record the trajectory  $\tau_{\text{agent}}$  the agent will take from this point forward.<sup>1</sup> Second, we generate a counterfactual trajectory  $\tau_{\text{user}}$  as if the agent had executed the user’s proposed alternative action. To generate  $\tau_{\text{user}}$ , we force the agent to follow the user’s suggested alternative action (or sequence of actions in the case that the user wants to specify a partial or complete counterfactual trajectory) and complete the trajectory using the same planning approach.

Finally, we use influence predictors to compare the two trajectories. Recall, in Sect. 3.1, we have an influence predictor  $IP_c$  per reward source  $c \in C$ . We have access to all the influence predictor values in both trajectories at this stage,

which gives us control over the types of explanation we can generate depending on our use case.

We can produce local, detailed, global, or aggregated explanations in natural language. To do so, we require one modification to the typical Markov decision process formulation: we use an *enhanced* reward function,

$$R : S \times A \rightarrow \mathbb{R} \times \mathcal{L}^* \quad (2)$$

where  $\mathcal{L}$  is a natural language vocabulary and  $\mathcal{L}^*$  is the set of all possible natural language strings. The enhanced reward function produces both a real-valued number and a natural language string that describes the circumstances under which the reward is given. For example, for the scenario in Fig. 1, the enhanced reward function might produce the value  $-1$  and “fall down stairs” when the agent encounters the location in the top center of the environment.

These textual descriptions of reward circumstances can be embedded in the explanation templates to automatically add stance or other details about the reward influences to the explanations. In this work, we explore two variants of the explanations: an *aggregated explanation* and a detailed *local explanation*.

**Aggregated Explanation.** An aggregated explanation presents an overview of the effect of each reward class on the agent throughout a complete trajectory. To generate these explanations, we calculate the mean value of each influence predictor  $IP_c$  for all states  $s \in \tau_{\text{agent}}$  and compare this mean with the other trajectories.

$$\operatorname{argmax}_{a \in \{a_{\text{user}}, a_{\text{agent}}\}} \left( \frac{\sum_s \max(IP_c(\tau_a[s]))}{\text{length}(\tau_a)} \right), \forall c \in C. \quad (3)$$

<sup>1</sup> We assume access to a simulator or a world model (e.g., [27]) with sufficient fidelity to predict action outcomes, though a world model-based learner was not used to create a minimalistic experiment.

The resulting explanation assigns each influence predictor to the trajectory in which they were the most dominant. We then take these values and translate them into natural language (English) explanations. If the values generated by  $IP_c$  were similar for both trajectories, then they can be excluded from the explanation, as they do not provide valuable, contrastive information. For example, in the case of two influence predictors in a simple environment, if we assume that we only had a difference in the negative influence predictor, the explanation would be as follows: “*If I go up, I fear falling down the stairs; going down feels safer.*” This case is demonstrated in Fig. 1. Suppose that there is a difference in an additional  $IP_c$  mean. In that case, we add it to the explanation by appending its relevant sentence to explain all influence predictors if we observe a significant difference in their mean values.

**Local Explanation.** In more complex environments where there are many influences, the explanations can shed light on the agent’s near-horizon and concentrate more on exploring local values. These local explanations zoom in on the most influential factors related to actual and counterfactual actions over their corresponding *segments*, which are denoted by  $t_{\text{agent}} \subseteq \tau_{\text{agent}}$ . Each segment starts at the beginning of the designated action and includes a pre-determined number of additional steps (we used five steps) after the end of the action to capture some of the longer-term effects.

To generate a local explanation, we first identify the most significant factors for each local segment by calculating the maximum influences of each state  $s$  during a segment:

$$\{\underset{c \in \{C\}}{\text{argmax}}(\max(IP_c(s))) \mid \forall a \in \{a_{\text{user}}, a_{\text{agent}}\}, \forall s \in t_{\text{agent}}\} \quad (4)$$

From Eq. (4), we get two sets of influence predictors that had the maximum value at any state within the segments  $t_{\text{agent}}$  and  $t_{\text{user}}$ . If there is no difference between the segments of  $a_{\text{agent}}$  and  $a_{\text{user}}$  with respect to their sets of maximum influences  $IP_c$ , we resolve to compare the mean values of the influence predictors during the same period and obtain the three influences with the highest means in their segments. First, we get the mean of every influence predictor in each segment:

$$\text{mean}(IP_c, a) = \frac{\sum_s \max(IP_c(t_{\text{agent}}[s]))}{\|t_{\text{agent}}\|}, \quad \forall c \in C, \forall a \in \{a_{\text{user}}, a_{\text{agent}}\} \quad (5)$$

Then, we obtain the top three influences for each segment  $t_{\text{agent}}$  and  $t_{\text{user}}$ , ordered by  $\text{mean}(IP_c)$ .

We provide this explanation even if there is no difference in the top influences, which can happen if the alternative action considered is too similar to the actual action taken.

To generate English explanations, we use the list of the top influence predictors obtained through the maximum or highest mean method. And then embed them into the templates. An example of a generated explanation can be: “*If I got the nearby cup, I would still have to clean it before making my coffee. But if I get a cup from the cupboard, I will not need to clean it and I can immediately make coffee.*” where “cleaning” and “making coffee” are terms derived from the enhanced reward sources, as described above.

The explanation strategies mentioned above provide information about how external environmental rewards affect the choices made by the agent, even if they are not directly linked to the agent’s actions. These explanations are in line with the definition of first-order explanations [28], as they reveal the underlying inclinations of the agent toward the environment and other factors that influence its decisions.

## 4 Evaluation overview

Our evaluations focus on exploring how beneficial Experiential Explanations are to participants in understanding RL agents and how they use these explanations to reason about their own decisions through quantitative and qualitative analyzes. We conducted two studies to evaluate our explanations with human participants in *simple* and *complex* environments. Our simple environment is a grid world built on MiniGrid [29] environment. Our complex environment is the Crafter game [30], a 2D version of the popular commercial game, Minecraft.

Our evaluations address the following research questions in both settings.

- RQ 1.** Do explanations derived from influence predictors accurately represent the actual values of the agent?
- RQ 2.** How do Experiential Explanations improve users’ ability to predict the agent’s actions compared to baselines?
- RQ 3.** How does user satisfaction with Experiential Explanations compare to baseline explanations?
- RQ 4.** How do users who got Experiential Explanations compare in their reasoning with those who received baseline explanations?

In both simple and complex environments, we first address **RQ1**, which focuses on examining the faithfulness of the explanation, a key consideration of explanation generation, which measures the extent to which the explanations are faithful to the underlying system [31]. Influence predictors sit “outside of the black box,” whereas the main policy model that drives the agent is “inside.” However, influence predictors are trained alongside the main policy model, and thus, if they learned correctly, should be able to estimate the agent’s

actions with a high degree of accuracy. We don't expect them to replicate the agent's actions as they are learning influences instead of generating actions. To assess the faithfulness of influence predictors, we evaluated the accuracy of the combined decision of influence predictors in predicting the agents' actions and how closely the combined influence predictor's values estimate the agent's.

We then designed and conducted our human participant study following the framework proposed by Hoffman et al. [32], evaluating our explanations for actionability, satisfaction, and utilization. We examine the actionability of the given explanations through a prediction and ordering tasks of the agent's actions and sub-goals (**RQ2**). Then we examine the participants' perception of the explanations and their satisfaction ratings (**RQ3**). Lastly, both the prediction tasks and the satisfaction survey are followed by reasoning questions that we qualitatively analyze to explore the use of explanation elements in reasoning for the different groups (**RQ4**).

The first experiment focused on understanding the overall utility of the explanations in understanding and predicting the agent's actions. This experiment is set in a MiniGrid environment, where we examine the performance of the aggregated version of the explanations in a sparse reward setting with two reward classes.

In the second study, we evaluate the usefulness and actionability of explanations in a more complex environment. We set our experiment in the Crafter environment to explore the effectiveness of the local explanations generated by our Experiential Explanations method in a more complex environment with sequential multistep goals, with more granular reward classes. Moreover, we want to see if the local Experiential Explanations can convey these complex values to the participants coherently.

In the next sections, we provide a comprehensive overview of the two experiments we designed and conducted to evaluate our explanation technique and address our research questions.

## 5 Experiment 1: MiniGrid environment

In this study, we evaluate our explanations in an environment with two simple reward sources and simplified grid navigation settings, with only one possible counterfactual action at the decision point (see Fig. 3). To answer our research questions, we first automatically validated the faithfulness of the explanations (RQ1), and then, we evaluated our explanation method with quantitative and qualitative analyses through a human evaluation study (RQ2, RQ3 and RQ4). We built our experimental domain on MiniGrid [29], a popular minimalistic grid world environment that offers a variety of tasks to test RL algorithms.

### 5.1 Model setup and environment details

**Environment customization.** We require an environment with multiple reward classes—positive rewards as well as one or more sources of negative rewards—thus, we created a custom task where the agent is placed in a room with a goal (a green square) and other objects represented by colored shapes that we could assign to neutral or negative reward values when approached. The environment provides sparse rewards, a positive reward of  $(1 - \frac{\text{steps}}{\text{max allowed steps}})$  for reaching a green goal tile (the MiniGrid default) and  $-1$  for stepping into a dangerous object. Both the goal and stepping on the dangerous object terminate the episode. All other states and actions produce zero reward.

**Agent training.** Our base agent policy model is the advantage actor-critic (A2C) [33] optimized using the exponential policy optimization algorithm (PPO) [34]. A state observation is an image of the grid constrained by a limited field of view. We include more technical details about the agent's architecture and training in Appendix A.

**Influence predictors design and training.** We are using DQNs as our influence predictors (as described in [25]). Specifically, we have two predictors, one for negative rewards and one for positive/goal rewards. We provide more details on the training and architecture of the influence predictors in Appendix B.

**Explanation generation.** We use the aggregated explanations described in Sect. 3.2. These explanations contrast the aggregated difference in the means between the agent's paths, the one associated with the proposed counterfactual, and the one the agent did. Below are some examples of these explanations:

#### 1. Negative rewards heavily influence the counterfactual trajectory.

**Template:** “If the agent  $[a_{user}]$ , it will pass through regions influenced by the [negative influence]; going  $[a_{agent}]$  feels safer.”

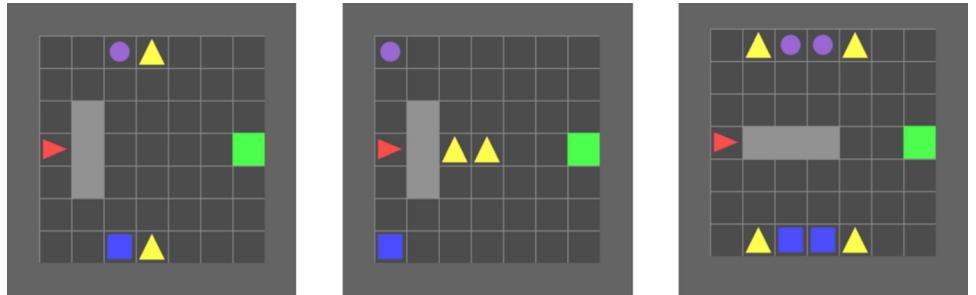
**Explanation Example:** “If the agent goes down, it will pass through regions influenced by the dangerous blue obstacles; going up feels safer.”

#### 2. Goal rewards have a lower influence on the counterfactual trajectory.

**Template:** “If the agent  $[a_{user}]$ , it will pass through regions less influenced by the [positive influence]; going  $[a_{agent}]$  feels better.”

**Explanation Example:** “If the agent goes right, it will pass through regions less influenced by the green goal; going left is better.”

**Fig. 3** Three episodes from the human study, in order presented to participants



## 5.2 Baselines

In this experiment, we focus on similar baseline explanations that the agent can generate. We compared our method with the following baselines:

- **Heatmap explanation:** Showing the agent's learned maximum values for each state. We generate the explanations for this baseline by getting the maximum agent value in each possible state on the map. This is similar to the agent values map in Fig. 2.
- **Q-Value explanation:** Presenting the average expected reward for both paths. These explanations utilize the simulator from our explanation generation pipeline, but not the influence predictors. This baseline is a comparison of the mean agent's values between the paths. An example of such an explanation is *The Q average if the agent goes up 9.84, while the Q average if the agent goes down is 9.66*.
- **No explanation:** The users only saw the correct answer without further explanation.

## 5.3 Faithfulness evaluation methodology

To evaluate the reliability of our explanation method, we resort to addressing the faithfulness (RQ1); we combine the predictions of the influence models using two methods. The first is to combine the values of the influence predictors and using their cumulative value to estimate the agent's action  $action = \text{argmax}_a (\sum_i P_i(s, a) - \sum_j N_j(s, a))$ , where  $P_i$  is a positive influence predictor and  $N_j$  is a negative influence predictor. The second method is to train a classifier on the predicted positive and negative influence values as features and the agent's actions as the ground truth. We used a support vector machine (SVM) classifier with a radial basis function (RBF) kernel. We also hypothesized that in the sparse reward setting, it is more meaningful to look at the influences around the emitting rewards rather than further away from them, as it would likely be more influential to the agent's decision in those areas rather than further away, and therefore, the values would likely be better at predicting the agent's actions when those influence signals are high enough to impact the agent.

To do this, we filter our dataset, by influence values at a set of different fixed intervals and examine the prediction accuracy at each threshold, we start at 0.100, and we keep increasing by +0.1 until it is no longer meaningful to do so. The faithfulness methodology does not make use of the baselines since it is the comparison of the aggregation of information from influence predictors to the main policy model.

## 5.4 User study methodology

We designed our human participant evaluation following the framework proposed by Hoffman et al. [32], evaluating our system for understandability, performance, and user satisfaction. We conducted our online evaluation between subjects using Prolific<sup>2</sup> with 81 participants aged 21–71 years old ( $M=37.41$ ,  $SD=11.03$ ); 39.5% of the participants identified as women. The average study completion time was approximately 10 min, and we compensated each participant with \$3.75. The study consists of two parts: answering a series of prediction tasks (RQ2, RQ4) and a satisfaction survey (RQ3, RQ4). The explanation group selection was fully randomized and balanced using Qualtrics.<sup>3</sup>

The study involved participants observing an agent (represented as a red triangle) navigate toward a goal (green square) across three episodes, with certain shapes yielding negative rewards. Initially, participants made uninformed predictions about the agent's movements, which were manipulated to ensure incorrectness. This established the conditions for which the participant must now learn to correct their understanding of the agent's interactions with the environment. This setup was used to teach the participants to adjust their understanding based on consistent negative reward assignments in subsequent episodes. The last two episodes were considered for the prediction task. The expectation was that the participants in the experimental condition would demonstrate improved prediction accuracy over those in the baseline conditions. In addition, a satisfaction survey assessed participants' perceptions of the explanations provided for the agent's behavior. Each of the questions in the prediction

<sup>2</sup> <http://prolific.co>.

<sup>3</sup> <https://www.qualtrics.com/>.

**Table 1** Percentages of correct answers for first prediction task and second prediction task and the combined percentages

	2nd Correct	3rd Correct	Both Correct
Experiential Exp.	<b>90.48%</b>	<b>95.24%</b>	<b>85.71%</b>
Heatmap Exp.	78.95%	78.95%	63.16%
Q-value Exp.	80.95%	80.95%	66.67%
No explanation	75.00%	70.00%	55.00%

Bold text highlights the highest percentage in each column

task we followed by a text entry question to capture the participants rationales for making these predictions. In addition, another text entry question was added to get an overall impression of the helpfulness of these explanations. These are the responses that we analyze in our qualitative section. We include screenshots and a detailed description of the study in Appendix D.

## 5.5 Quantitative results

**RQ1: Faithfulness.** For measuring faithfulness, we look at the accuracy of influence models in predicting agent actions; we found that the accuracy of the prediction of the direct aggregation is 72% and the accuracy of the classifier 76%. To examine the alignment of the influence predictors with the agent at states with higher influence, we gradually examine the alignment between the highest action as predicted by the influence predictors and the agent’s actions. At a threshold of 0.100, the aggregate of influence predictors agreed with the main policy 83%, covering at least 20% of the environment. At higher thresholds, agreement increases above 90%, indicating that when decisions are the most important, faithfulness is very high. Looking at the positive influence predictor alone, at a threshold of 0.200 and higher, it alone can predict the agent with a precision of 100%. This is to be expected; influence predictors cannot dictate what states and actions get explored—it is controlled by the main policy and the main training algorithm—so states that are less critical to positive reward attainment are under-sampled and prone to more error in both the main policy and the influence predictors. In contrast, more important decision points are sampled more often, so the influence predictors get more training on those states and actions.

**RQ2: Prediction Correctness.** We measure whether the participant can correctly predict the agent trajectory after interacting with the second and third episodes (Fig. 3).

Table 1 shows the percentage of correct answers for each type of explanation after the second and third episodes (the first was a control), as well as the combined average. Participants provided with Experiential Explanations had the highest success rate. Experiential Explanations were significantly higher than the no explanation condition, validated

with a logistic regression ( $p < 0.05$ ,  $df = 3$ ,  $\chi^2 = 4.404$ ). Experiential Explanations led to a 9.5+% improvement over other baselines, but were not statistically significant at the same threshold.

**RQ3: Satisfaction Survey.** For each type of explanation, we analyzed the ratings of the participants for the Likert scale questions on the satisfaction survey. We include a visualization of the results in Fig. 22 in Appendix. Experiential Explanations had the highest average scores on all dimensions, except trust. A one-way analysis of variance (ANOVA) followed by Tukey’s honestly significant difference test revealed that the Experiential Explanations ( $M = 3.62$ ,  $SD = 1.32$ ) were preferred over the Heatmaps ( $M = 2.42$ ,  $SD = 1.35$ ) in completeness ( $F(3, 81) = 3.032$ ,  $p < 0.05$ ). Experiential Explanations ( $M = 3.71$ ,  $SD = 1.0$ ) were preferred over the Heatmaps ( $M = 2.63$ ,  $SD = 1.38$ ) in satisfaction ( $F(3, 81) = 3.106$ ,  $p < 0.05$ ). Experiential Explanations ( $M = 4.38$ ,  $SD = 0.59$ ) were preferred over Q-Value Explanations ( $M = 3.28$ ,  $SD = 1.01$ ) and Heatmap Explanations ( $M = 3.0$ ,  $SD = 1.37$ ) in the dimension of understandability ( $F(3, 81) = 5.572$ ,  $p < 0.05$ ).

## 5.6 Qualitative analysis and findings

To address RQ4, we thematically analyzed open response questions for the tasks and satisfaction survey to better understand how participants used the explanations. Thematic analysis is one of the most common qualitative research methods and involves identifying, analyzing, and reporting patterned responses or meanings (i.e., *themes*) [35]. It is most appropriate for understanding a set of experiences, thoughts, or behaviors across data [36]. To conduct the thematic analysis for each of these codes, three authors performed the following established process, as outlined by [35]: (1) We familiarized ourselves with the data through repeated active readings. (2) We developed a coding framework consisting of a set of *codes* organizing the data at a granular level and the corresponding well-defined and distinct definitions. (3) Each of the authors independently coded the entirety of the data. If a code was unclear or a new code was needed, the authors consulted with each other to modify the coding framework. If a change was made, each author independently re-analyzed the entirety of the data. (4) We jointly examined the codes to derive themes and findings. We compared our application of the codes for each response to ensure consistency. Discrepancies were counted and used to calculate the inter-rater reliability. The process iterated until all discrepancies were resolved. We identified three themes: Usefulness, Format, and Utilization. The codes, definitions, and examples are summarized in Figs. 4 and 5.

**Explanation Usefulness.** The theme of Usefulness emerged in the open-ended comments about how participants per-

## Usefulness

Comments around whether participants found the explanation useful

Useful	The participant found the explanations useful	"The explanations <b>helped me</b> understand and achieve the task."
Not Useful	The participant found the explanations not useful/relied on other factors for their predictions	"The explanations <b>did not help</b> me achieve the task. I had to <b>rely on my self-interpretations</b> . (e.g., The agent's actual routes/consistency)"
Unclear	Usefulness couldn't be inferred from the answer	"I think the agent could have made it to the goal safely either way he went."

## Format

Comments around the explanation formats and the difficulties they encountered

Need more information	Mentions that the explanation is too brief, needed more details.	"The explanations <b>should be longer</b> and answer <b>only some of my questions</b> ."
Hard to understand	Mentions that the explanation was not straight forward/needed time to understand.	"The explanations were <b>hard to understand</b> ."
Suggestions	Suggestions on what they wanted to see in the explanation	"I wanted to <b>know more about why the agent avoided the purple circle</b> ."

**Fig. 4** Summary tables of the codes used to understand how users across the explanation groups perceived the explanations in terms of usefulness and format

## Utilization

Reasoning around the information used in reasoning about the agent

Object Related	Reference to the object shapes or colors.	"The agent hates the color <b>purple</b> . The agent wants to avoid <b>circles</b> ."
Reward Related	Reference to safer, better or higher rewards.	"This path has <b>higher reward</b> . Going up is <b>better</b> ."
Goal Related	Reference to distance from the goal or direction.	"This path is <b>closer to the goal</b> ." "This path leads to the <b>goal</b> ."
Random Choice	Arbitrary justification.	" <b>50/50 chance</b> . I think the agent will go up."
Prior Behavior	Reasoning relied solely on past directions	"The agent went <b>up previously</b> ."
Other	Other reasons (e.g., obstruction)	"The agent go around walls."

**Fig. 5** Summary table of the codes used to understand how users across the explanation groups utilized information from the explanations

ceived the explanations. Because these comments were made alongside users' explanations of their reasoning, this provides additional information on how useful users found the explanations. We identified three Usefulness codes: **Useful**, **Not Useful**, or **Unclear**. We coded a response **Useful** if a participant explicitly mentioned that the explanation was useful, for example, "I understood that purple was dangerous and that the agent should avoid getting near

it (EE8)." Responses were coded **Not Useful** if a participant explicitly stated that the explanation was not useful, for example, "I tried to read the chart but found it confusing and unclear. I didn't understand it very well (HE10)." Responses were coded **Unclear** if the user's perception was ambiguous, for instance, "I found the explanations a bit too brief. For instance, what makes the purple circles feel safer than

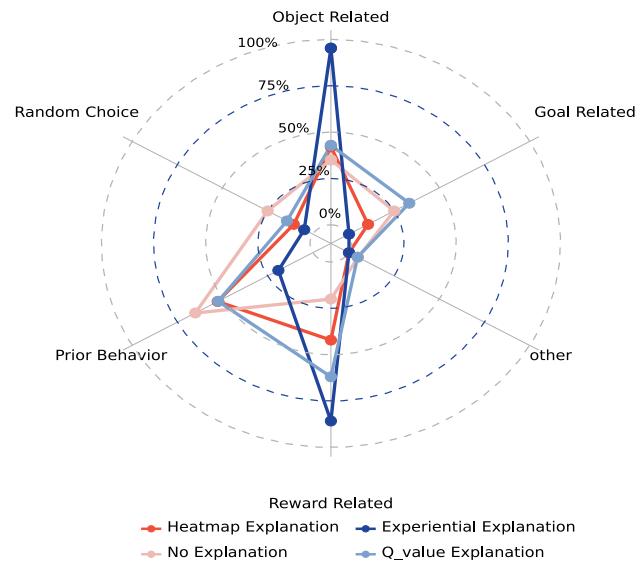
**Table 2** Participant's rates of usefulness for each explanation type

	Useful	Unclear	Not useful
Experiential Exp.	80.95%	19.05%	0.00%
Heatmap	31.58%	31.58%	36.84%
Q-value Exp.	47.62%	14.29%	38.10%
No explanation	N/A	N/A	N/A

the blue squares (EE7)?” The initial inter-rater reliability for Usefulness was 81.5%.

Table 2 shows participants' perceived Usefulness across explanation types according to usefulness codes. Experiential Explanations were considered **Useful** most often. Heatmap explanations were considered **Not Useful** most often, even more than when no explanation was provided. A logistic regression comparing the Usefulness of each explanation showed Experiential Explanations were 9 times more likely to be found **Useful** than Heatmap explanations and 4 times more likely than Q-value explanations ( $p < 0.05$ ,  $df = 3$ ,  $\chi^2 = 14.99$ ). We also note that 88.89% of the participants who thought the explanations were useful predicted correctly in both questions, indicating the correlation between correctness and usefulness ( $p < 0.05$ ,  $df = 2$ ,  $\chi^2 = 9.33$ ).

**Explanation Format.** Participant responses were also coded for Format. We derived three codes for Format: **Need More Information**, **Hard to Understand**, and **Suggestion**. Responses coded as **Need More Information** explicitly stated why the explanation was lacking and what data was desired. If the response said that the explanation was not clear, it was coded as **Hard to Understand**, and if the response offered another way to explain, it was coded as **Suggestion**. Only responses from individuals who received explanations and referenced the explanation's content could be coded for Format, 41.98% of our dataset; therefore, our reported observations are mainly comparisons of the percentages and could not be statistically tested due to the small data size. Participants often expressed a **Hard to Understand** 40% when shown Heatmap explanations, followed by 29.41% who got Q-value explanations while this was not a problem for the Experiential Explanation group. On the other hand, we observed **Need More Information** between the three explanation groups with higher concentrations of 35.29% being from the Q-value explanation group and 26.67% are from the Heatmap explanation groups and 21.05% are from the Experiential Explanation group. Participants who did not find the explanations useful had more Format comments, especially requesting more **Need More Information** 30.77% and **Hard to Understand** 69.23%. The initial inter-rater reliability was 76.47%.



**Fig. 6** Distribution of themes over the different explanation types. Participants who saw Experiential Explanations, which presented Object-Related and Reward-Related information, relied more heavily on the kinds of information it provided than other explanation types. Q-Value and Heatmap explanations both presented Reward-Related information. The breakdown of themes for the Heatmap and q-value explanations more closely resemble the no explanation case

**Explanation Utilization.** Utilization refers to whether the participant made use of the information available within the explanation to reason about what the agent will do. We extracted the following codes for the types of information participants referenced in their rationales: **Object-Related** (e.g., shapes or colors); **Reward-Related** (e.g., “safer,” “higher,” or “lower” rewards); **Goal-Related** (e.g., distance / direction to goal); **Prior Behavior** (e.g., references behavior from prior outcomes); **Random Choice** (e.g., no clear justification); and **Other** (e.g., references other factors such as obstructions). The inter-rater reliability of the initial coding was 75.6%.

We compared the frequency with which utilization codes appeared in open responses. If a participant used the same types of information as the explanation provided, we considered the participant as “Utilizing” the explanation. Experiential Explanations provide **Object- and Reward-Related** information. Heatmap and Q-Value Explanations provide **Reward-Related** Explanations. No explanation only provides for **Prior Behavior**. 71% of the participants who were shown Experiential Explanations utilized the exact kinds of information presented in the explanation, compared to 21% and 28% for the Heatmap and Q-Value explanations, respectively. Although the Q-Value and Heatmap explanations both presented **Reward-Related** information, participants still relied heavily on **Prior Behavior** to understand the agent.

Figure 6 shows the distribution of the reasons the participants referenced for each type of explanation. Participants who saw Experiential Explanations most used **Object-Related** and **Reward-Related** information. Participants who saw Q-Value explanations referenced **Goal-Related** reasons the most. Participants who saw no explanation most often utilized **Prior Behavior**, followed by **Object-Related** reasons. Notably, participants who saw Experiential Explanations were 4 times *less* likely to rely on **Prior Behavior**, indicating they did not rely as much on watching the agent to understand it. Those who saw Heatmap and Q-Value Explanations relied heavily on **Prior Behavior** and Utilization patterns mimicking those of the no explanation condition.

Experiential Explanations were 13 and 15 times more likely to have the information they provided be *Utilized* and *Useful*, respectively. *Utilized* refers to what they used in their reasoning, while *Useful* focuses on whether they thought the explanations were useful or not in the post-task satisfaction survey. Yes, it is the code from the previous section with respect to Heatmap and Q-Value explanations, respectively. We also found that among those who used components of explanations in their reasoning, 55.88% answered correctly on both questions ( $p < 0.05$ ,  $df = 1$ ,  $\chi^2 = 3.849$ ). But within the Experiential Explanations group, the percentage was 84.62% and that group was more statistically significant to get both questions correctly than other baseline explanation groups ( $p < 0.05$ ,  $df = 3$ ,  $\chi^2 = 12.22$ ).

## 6 Experiment 2: crafter environment

In this experiment, we address our research questions (RQ2, RQ3, and RQ4) with a focus on evaluating the actionability of the explanations in a more complex environment, as this experiment is designed to evaluate the influence predictors and Experiential Explanations in a visually richer environment, with more decision points and options, along with multiple reward classes. And, as before, we will validate the faithfulness through automated testing designed to address (RQ1).

To create such a setting, we used the Crafter environment developed by Hafner [30], a 2D version of the game Minecraft with the same rules for interacting with the environment and crafting new tools. Crafter can generate 2D open-world survival environments with randomized maps. These environments consist of varying terrains, a day–night cycle, as well as interactable materials, objects, and living entities. Crafter also supports an on-screen inventory, allowing an agent to collect items in order to construct a variety of items later. An agent playing Crafter needs to complete various tasks in sequential order while maintaining its health by obtaining resources and avoiding threats.

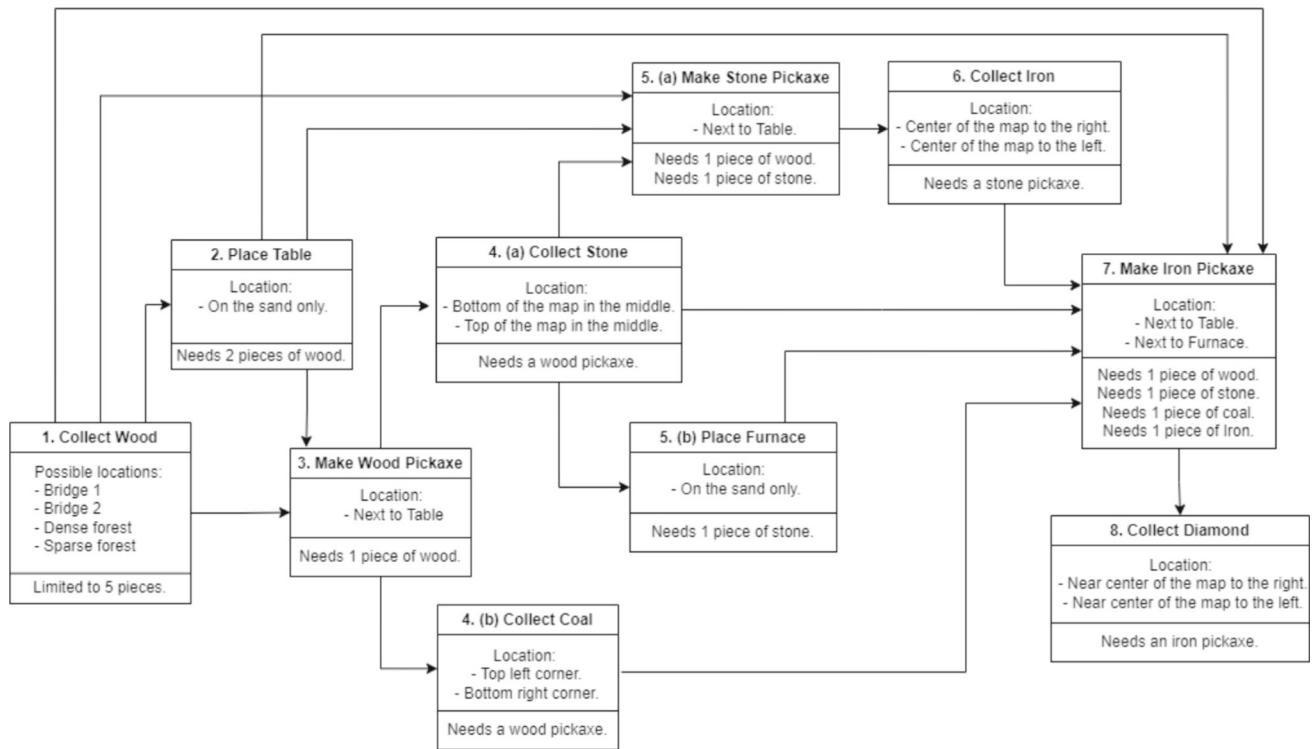
### 6.1 Model setup and environment details

**Environment Customization.** We needed a good distribution of positive and negative reward sources. Crafter provides multiple sources of positive rewards, but the only negative rewards come from health depletions. To counter this and to add a conceptual narrative, we created an agent persona, called the Digger, whose aim is to obtain a diamond. The Digger’s reward system provides positive rewards for achievements related to digging, such as acquiring resources to construct digging tools, creating these tools, and digging for minerals. We provide negative rewards if the agent strays from the Digger persona, for example, by constructing weapons, or killing an entity. Figure 7 depicts the sequence of achievements required to reach a diamond.

Additional customizations were added to optimize the training process. The world map was made static, so that the agent could find all required resources in a limited number of steps while making it easier for a participant to understand the world map. No rewards were given for performing actions to maintain health, such as drinking water, but improvement or depletion of health had corresponding positive and negative rewards. While all other achievement rewards are in the range of +1 to -1, diamond collection has a +50 reward to focus the agent’s training exploration toward getting a diamond. Finally, a small bonus is given for reaching the diamond in fewer timesteps, to improve agent performance.

**Agent training.** The agent needs to learn to traverse the map, completing the achievements required to reach the diamond in a small number of timesteps. At any given point, its observation consists of an image of its surroundings inside a limited field of view along with its current inventory. Given a static map, the agent is trained from various starting points to aid in the training exploration. We use on-policy learning algorithms for the agent, advantage actor-critic (A2C) [33] optimized using the exponential policy optimization algorithm (PPO) [34], implementation details for which are available in Appendix C.

**Influence Predictor design and training.** For the Crafter experiment, we attached influence predictors (IP) to achievement categories. Each predictor was responsible for tracking either favorable achievements leading to positive rewards or unfavorable achievements leading to negative rewards. The health-tracking IP was the only exception here, as it tracked both improvement and depletion of health, resulting in rewards that are sometimes positive and sometimes negative. Table 3 lists all influence predictors, along with the achievements they track. The health-tracking influence predictor is a special case of a reward signal that is both negative and positive. We opted not to divide the health rewards coming from the environment, because we didn’t want to change the health signal artificially; we would need to significantly



**Fig. 7** Depiction of the task dependencies of the Digger's quest, listing the achievements that the Digger needs to attain and the order in which that needs to be done

intervene to separate the signal into its positive and negative components. To handle this special case, we don't use the absolute function for the health-tracking influence predictor; we took both the sign and the magnitude of the values when it came to the health IP. This still works because we still standardized all the values before the explanation generation which corrects any inconsistency this inclusion of the sign might have caused.

The influence predictors were also trained using on-policy learning; the technical details are found in Appendix C. The predictors were given full access to the agent's rollout buffer allowing them to see the agent's observation space, action taken, and reward obtained, along with whether the episode just started and the probability distribution of agent's actions. IPs would only read the rewards pertaining to their own achievements and consider rewards from other sources as 0. Also, instead of taking the actual reward, IPs read the magnitude of the reward to track the intensity of influence of their achievements, irrespective of whether they are positive or negative. The health-tracking IP would also consider the sign of its reward, given that health can be both positive and negative. The predictor models calculated the value and subsequently the advantage and results using their own respective critics. They would then populate their own rollout buffers with these details and train on the same. Each IP

hence would learn its individual policy based on the agent's explorations.

We explored the IP values by standardizing them to compensate for reward variations and obtained graphs such as Fig. 8. We see that the influence of rewarded states associated with achievements keep rising until the achievement was obtained and drop immediately after, leading to creation of spikes as visual indicators of when an achievement was made. The IPs hence expose the user to knowledge about the agent's goal hierarchy because the positive rewards are attached to achievements. This is something that we don't see in the simpler MiniGrid and a useful emergent property of the technique.

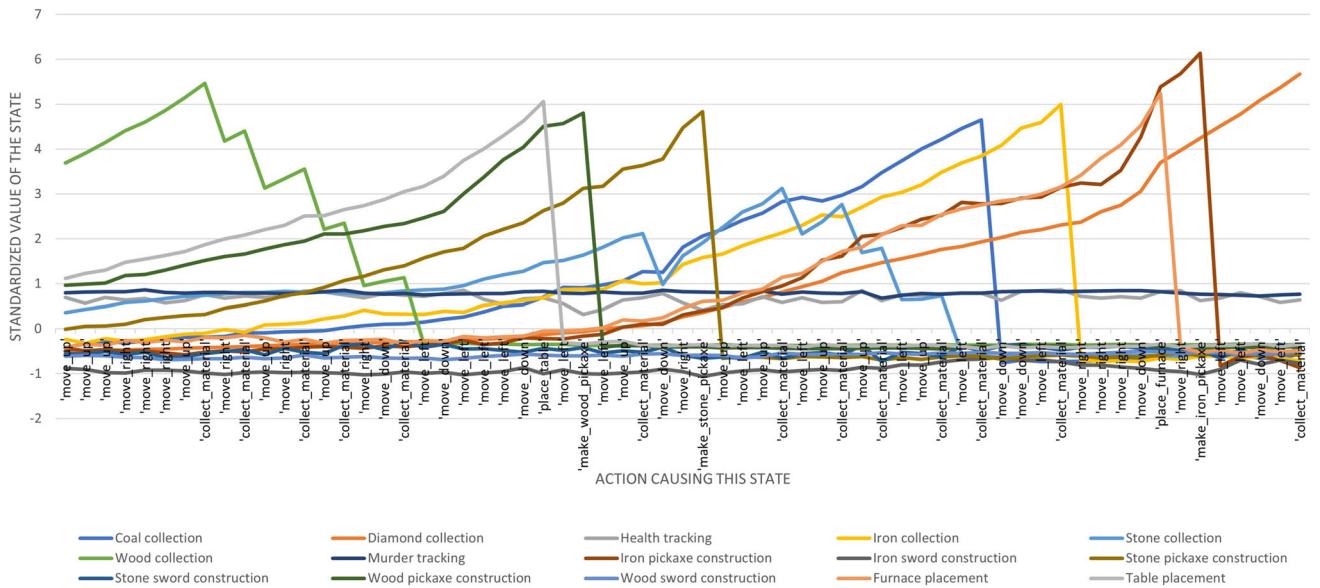
**Explanation generation.** To explain the actions of the agent, we provide local explanations that narrate what the agent plans to do next after the actual or counterfactual action, as described in Sect. 3.2 and using the following template.

**Template:** “[The agent] needed [highest influence predictor at decision point] and chose to [a<sub>agent</sub>]. This led to the prioritization of [list of top influences for t<sub>agent</sub>] next. However, if it had [a<sub>user</sub>], it would have to think about [list of top influences for t<sub>user</sub>] instead.”

**Example:** “Jade needed wood and chose to get it from the sparse forest. This led to the prioritization of building a table, building a stone pickaxe, and getting stones next. However,

**Table 3** Complete list of influence predictors with the achievements that they track and their corresponding reward types

Influence predictor	Reward type	Achievements
Coal collection	Positive	Collect-coal
Diamond collection	Positive	Collect-diamond
Iron collection	Positive	Collect-iron
Stone collection	Positive	Collect-stone
Wood collection	Positive	Collect-wood
Iron pickaxe construction	Positive	Make-iron-pickaxe
Stone pickaxe construction	Positive	Make-stone-pickaxe
Wood pickaxe construction	Positive	Make-wood-pickaxe
Furnace placement	Positive	Place-furnace
Table placement	Positive	Place-table
Murder tracking	Negative	Defeat-zombie
Iron sword construction	Negative	Make-iron-sword
Stone sword construction	Negative	Make-stone-sword
Wood sword construction	Negative	Make-wood-sword
Health tracking	Positive and negative	Life-maintenance, collect-drink, eat-cow



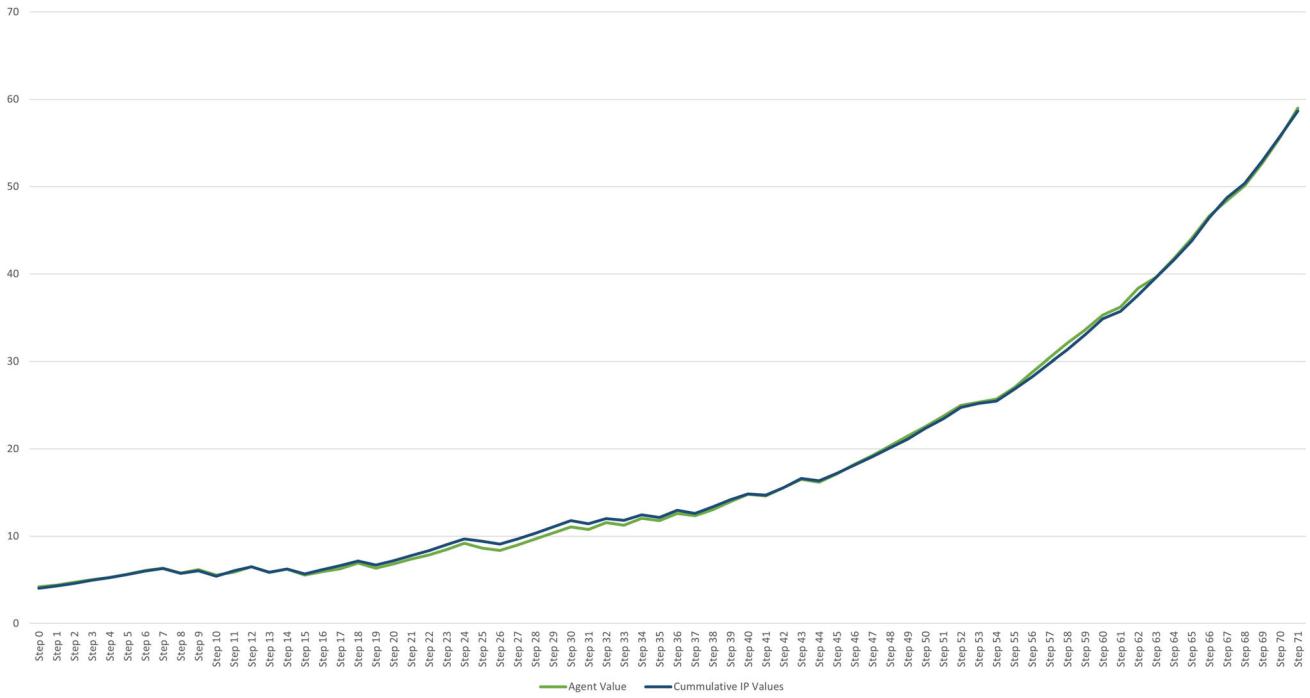
**Fig. 8** A plot of an agent's full trajectory over an episode and the standardized values of every state along the trajectory according to individual influence predictors. The x axis shows the temporally ordered actions preceding each state whose influence predictor values are represented on the y axis

*if it had gotten to the dense forest, it would have to think about owning a wooden pickaxe, fearing for its health and avoiding dangerous zombies instead.”*

This explanation conveys that the dense forest is a more dangerous option for the agent, because the agent knows, while the user might not, that there are zombies roaming around that forest and it would be wiser to get wood from the safer sparse forest. We anthropomorphize our agent; calling it “Jade” in all explanation groups and make minor edits for readability as needed (e.g., change the tense or add articles) to improve the survey experience of the participants.

## 6.2 Faithfulness evaluation methodology

We look at faithfulness to evaluate how good our explanation method is at approximating the agent. We calculate faithfulness (RQ1), which looks at how the influence predictors (*IP*) models can align with the agent's model. We used the two faithfulness evaluation methods from the first experiment as described in Sect. 5.3 to estimate an agent's actions based on influence predictors. The first method combines the values of all the positive and negative influence predictors to calculate a cumulative value, allowing



**Fig. 9** Plot depicts the values of different states as traversed by the agent during a successful test run. The green line represents the values according to the agent’s learned policy. The deep blue line depicts the cumulative values according to the policies learned by the influence

predictors. Both plots match almost exactly, meaning that the policies learned by the influence predictors align closely with the policy learned by the agent

us to determine the agent’s action as follows:  $[action = \text{argmax}_a \left( \sum_i P_i(s, a) - \sum_j N_j(s, a) \right)]$  where  $(P_i)$  represents a positive influence predictor and  $(N_j)$  represents a negative influence predictor.

The second method involves training a classifier using the standardized positive and negative influence values as features, with the agent’s actions as the ground truth. We used the same setup we had in the first experiment, an SVM classifier with an RBF kernel.

To learn more about how our influences perform in different parts of the trajectory, we take a slightly different approach than the one we did in the first study. In the first study we did so by focusing on how the influences values impact the agent at a close proximity. In this experiment, because our setting here is more complex and combining the influence predictors individual action predictions can be noisy, and it is not feasible to easily identify the strengths of their signals at different locations, as not all the 15 IPs are equally influential at all decision points, we resort to looking at the sum of their values at each state as a way to show their alignment with the agent’s. Then we compare this sum with the agent’s values for each state. The goal is to see how closely the cumulative IP values can match those of the agent. Our process begins by obtaining the cumulative IP values of

each state using the formula:

$$\text{Cumulative IP value for state } s = \sum_{IP_c \in C^+} IP_c(s) - \sum_{IP_c \in C^-} IP_c(s) \quad (6)$$

where  $C^+$  is the set of positive IPs including health-tracking IP and  $C^-$  is the set of negative IPs.

We then calculate the Root Mean Squared Percentage Error (RMSPE) [37] between the agent values and the cumulative IP values for a set of 9 test runs. These runs used the same trained agent on the same map, but varied the starting location. RMSPE is calculated as:

$$RMSPE = \sqrt{\frac{1}{N} \sum_{n=1}^N \left( \frac{\hat{y}_n - y_n}{y_n} \right)^2} \quad (7)$$

As before, the faithfulness evaluation does not use the baselines as the purpose is to evaluate the influence predictor models’ ability to approximate the policy model.

### 6.3 Baselines

The setup of this experiment allows the generation of localized explanations, so we used the following baselines to compare them with:

1. **Local agent value explanation:** For both options provided to the user, this explanation presents the agent value of the state reached right as the action sequence ends. This action sequence will end in some reward-producing state, but not necessarily in the goal state. Similar to the previous experiment, these values are generated using the simulator. We similarly used templates to generate the explanations encapsulating the values.

**Template:** Jade's value for  $[a_{\text{agent}}]$  is [agent value after  $a_{\text{agent}}$ ]. While the Jade's value for  $[a_{\text{user}}]$  is [agent value after  $a_{\text{user}}$ ]

**Example:** Jade's value for getting wood from the sparse forest is 6.99. While Jade's value for getting wood from dense forest is 5.59

2. **No explanation:** As in the first experiment, users only saw the correct answer, without further explanation.

The agent value explanation is similar to the Q-value explanation we had in the previous study, except now it is calculated on a local point rather than an aggregate mean over the full trajectory. We did not include the Heatmap baseline because the agent may need to traverse the same positions on the map multiple times to complete its task, which made Heatmaps inapplicable.

### 6.4 User study methodology

We follow a study design similar to our first experiment; we have similar main study components of prediction tasks (RQ1, RQ3) and a satisfaction survey (RQ2, RQ3) as in our previous study. We adjusted the prediction task to be more suitable for the new environment and local explanations. We also introduce a new ordering task to further explore the comprehension of the explanations that address RQ3 and RQ4.

Our study is an online evaluation between subjects using Prolific with 96 participants aged 18-82 years old ( $M=38.90$ ,  $SD=12.91$ ); 32.65% of the participants identified as women. The median study completion time is approximately 12 min, and each participant was compensated with \$2.50 with a bonus for thoughtful responses.

Initially, to introduce the participants to the map and receive the initial explanation, we show them a scenario and then ask them what the agent would choose between two options. After making a selection, they will see an explanation generated by our technique or one of the baselines. The explanation group selection was fully randomized and balanced using Qualtrics.

Then they will see another scenario, a **prediction task**, where they are asked to predict the agent's follow-up action; they are presented with one of two prediction tasks. The first examines the case of multiple positive influences, where they are asked which action will the agent take next. The other examines the negative influences, where they are asked which option the agent will avoid. In both versions, the correct options relate to the influence with the highest value. After making their predictions, we ask them to explain their reasoning.

The next task is an **ordering task**, where we show them a list of the factors affecting the agent for each of the action choices of the initial scenario. And ask them to order them according to their relative importance to the agent. We also ask them what factors they considered with the ordering. Then we ask if the explanations were helpful in this task and how.

Finally, they conduct a **satisfaction survey**, identical to experiment 1, where they rate the explanation on understandability, satisfaction, amount of details, completeness, usefulness, precision, and trust. In addition, a text entry question about the usefulness of the explanations to their decision-making. We provide a detailed description of the survey and a screenshot in Appendix G.

**Disguising of assets** The functionality of Crafter's materials, objects, and living entities is quite intuitive, and many people are familiar to different degrees with Minecraft. We needed to ensure that the explanations were the only source of a participant's understanding of concepts of this world, as opposed to prior knowledge about Minecraft or commonsense. We re-skinned the art assets to random alternatives which carry no relation to the thing they represented. Figure 10 displays the main map with a subset of original and re-skinned assets. The full list of re-skinned assets can be found in technical appendix.

### 6.5 Quantitative results

**RQ1: Faithfulness** For the first test, the accuracy of the influence predictors in predicting the agent's actions along the trajectory, we found that the accuracy of the prediction of the direct aggregation is 73.35%.

Using the classifier from experiment 1 with only movement actions (4 classes) resulted in 69.61% and with all possible actions (11 classes) is 60.71%. The classifier from experiment 1 only uses influence predictor values as features, not the entire observation; this was sufficient for experiment 1. However, the Crafter environment is much more complex in terms of action space, state space, and temporality—states are seen repeatedly but under different conditions such as different tasks. Furthermore, the size and complexity of the state-action space is such that certain trajectories through

**Fig. 10** Comparison of a 13x13 tile Crafter map with original art assets on the left and re-skinned art assets on the right



state space are vastly over-sampled by the policy learner while others are under-sampled. This makes it hard to learn a classifier without artificially forcing a more balanced sampling of states and actions. Other classifier architectures may perform better. Regardless, it shows that the influence predictors provide information that can be used to reconstruct the policy many times above random chance (9%) even using naive reconstruction means such as this classifier.

The cumulative values of the influence predictors, calculated with Eq. (6), were almost identical to those of the agent. We plot the results of an example test run of average length that shows how the values are aligned almost perfectly, as demonstrated in Fig. 9. This is verified by the RMSPE values that averaged 3.13% ( $SD = 0.52\%$ ,  $min = 2.45\%$ ,  $max = 3.82\%$ ), indicating a high similarity between the agent policy and the cumulative IP values.

## RQ2: Prediction Correctness and Order Correctness

We measure whether participants can correctly predict the agent's action at the prediction task decision point after interacting with the explanation of the initial decision point. The correctness rose up for the local Experiential Explanation group from  $0.4 \rightarrow 0.71$ , while there was no difference in the other groups as shown in table 4. The group that received local Experiential Explanations demonstrated a significant improvement in the accuracy of the prediction in all groups after exposure to an explanation. The results of the logistic regression model were statistically significant ( $X^2(2, N = 96) = 13.327, p < 0.05$ ). There was a positive correlation between the local Experiential Explanation group and the accurate prediction of influences ( $OR : 2.55, 95\%CI[1.785, 3.326]$ ).

For the ordering task, we measure the ability of the participants to order the factors that impact the agent the most for each action proposed at the initial decision point. We used Sum of Squared Error to compare the correctness of the sequences as shown in Fig. 11. The group provided with

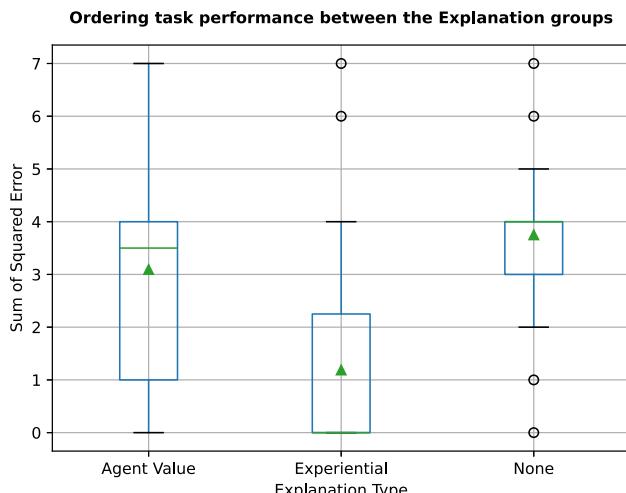
**Table 4** A comparison of the explanation groups prediction correctness for their initial predictions; before seeing any explanations; and after seeing the relevant explanation

	Initial predictions	Prediction task
Experiential Explanation	40.63%	71.88%*
Agent value explanation	28.13%	28.13%
No explanation	37.50%	34.38%

\*Statistical significance difference compared to other explanation groups ( $\chi^2(2, N = 96) = 13.327, p < 0.05$ , logistic regression)

local Experiential Explanations showed significantly reduced squared errors ( $M=1.19$ ,  $SD=2.27$ ) compared to other baseline groups which is found by conducting a one-way analysis of variance (ANOVA) followed by Tukey's honestly significant difference test ( $F(2, 96) = 14.1315$ ,  $p < 0.05$ ). This indicates that the local Experiential Explanation group was more effective in achieving accurate ordering compared to the agent value explanation group ( $M = 3.09$ ,  $SD = 2.89$ ) and the group without explanations ( $M = 3.63$ ,  $SD = 2.91$ ).

**RQ3: Satisfaction Survey** As with the first experiment, we analyzed the explanation goodness ratings the participants provided for each explanation group for the Likert scale questions. We include a visualization of the results in Fig. 28 in Appendix. The local Experiential Explanation group exhibits the most favorable mean values in all measured dimensions, with the exception of Trust. A one-way analysis of variance (ANOVA) followed by a Tukey’s honestly significant difference test revealed that these differences are not statistically significant, with the sole exception of the dimension of sufficient detail ( $F(2, 96) = 4.716, p < 0.05$ ), where the local Experiential Explanation ( $M = 3.33, SD = 1.45$ ) was perceived as significantly more detailed compared to the agent value explanation ( $M = 4.06, SD = 1.16$ ).



**Fig. 11** A comparison showing the distribution of the sum of squared errors for the ordering task between the different explanation groups

## 6.6 Qualitative analysis and findings

To understand how users in different explanation groups varied in the way they used the information provided in explanations, we thematically analyzed the open response questions asking participants why they selected their particular answers for the reasoning tasks. We followed the same procedure we describe in Sect. 5.6, familiarizing ourselves with the data, developing a coding framework, independently coding the data, and jointly resolving differences [35]. Discrepancies were counted and used to calculate the inter-rater reliability. The process was repeated until all discrepancies were resolved. The inter-rater reliability for all three questions we analyzed is 71%, 75.1%, and 80.1%.

Using this procedure, we identified the following codes that are broadly separated into: (1) *elements*, participants understanding the parts of the environment the agent would interact with; and (2) *logic*, participants reasoning around the circumstances leading to the agent's actions. The codes, definitions, and examples are summarized in Fig. 12. Significantly, the codes for Experiment 2 reflect and extend the codes identified for Utilization in Experiment 1, but they were developed independently, grounded in each set of data, demonstrating an additional layer of validity and consistency between the two studies' findings.

We also used the same procedure to conduct a thematic analysis of the open response question from the satisfaction survey asking participants about the usefulness of the explanations. We qualitatively coded their responses to identify the features that users perceived as present and missing from the explanations. We broadly separated our codes into features related to the broader *Game* (e.g., game rules) and features related to the *Agent* (e.g., agent strategies). The

codes, definitions, and examples are summarized in Fig. 13. The inter-rater reliability for this question is 83.5%.

We focus our discussion on the most salient themes we identified from our analysis. (1) Participants in the different explanation groups vary in how they utilize the information provided in the explanation. (2) Participants identified the features they perceived as missing from the explanations. (3) In the absence of a complete picture, how do participants reconcile their mental models of the agent?

### Comparing Elements & Logic Utilization Across Explanation Groups.

The group provided with local Experiential Explanations was the only group to demonstrate direct reliance, relying only on the Explanation Logic and Explanation Element with nothing else added. Of the 96 questions answered by the local explanation group, 41 (42.7%) directly relied on the explanation provided, compared to none in either of the other two baselines. Those who did not directly rely on the explanation most often relied on both the *Explanation Logic* and *Explanation Element*, but added other explanations or elements. The group provided with agent value explanations most heavily relied on *Real Life - Object Logic*, *Other Logic*, and *Info Elements*. The group provided with no explanations most heavily relied on *Real Life - Object Logic*, *Image Event Logic*, and *Info Elements*. This indicates that participants in the local Experiential Explanation group more heavily relied on the information provided in the explanation to predict the agent's actions, rather than drawing on other knowledge. It also confirms the experimental validity, as not all participants in the local Experiential Explanations solely depended on the explanations, indicating answers were not directly given to them in the explanation.

Overall, using influence predictors provided a significant improvement over other explanation methods in helping users predict the agent's downstream actions. Participants who received the local Experiential Explanations were significantly more likely ( $p < 0.05$ ) to correctly choose the influences on the model compared to the other explanation groups. Further, those who received local Experiential Explanations more frequently relied on the information provided in the explanation, rather than what was provided in the images or on connections to experiences outside of the experiment. 42.7% of the participants in the local Experiential Explanations group directly relied on only the explanation elements and explanation logic provided to make their decision.

The improvement in performance aligns with the provision of additional information through our method. Influence predictors are models that learn how sources of reward affect the agent in different stages, restoring information about how the policy reflects the environment. This information allows users to predict the downstream actions of the agent

## Elements

Reasoning around the parts of the environment with which the agent interacts

Explanation Element	References an object or item provided in the explanation	"It was mentioned [in the explanation] previously that the <b>yellow ducks</b> are dangerous, so I assume she should avoid them"
Info Element	References an object or item shown in the image or listed as an answer for a question	"I ranked the <b>flashlight</b> as a factor of highest importance since it seems to have been used for getting a coaster. It's gone in the 'Finish' picture"
Other Element	References anything other than an Explanation or Info Element	"I think Jade could have considered the coaster as a weapon similar to the hammer where she could have smashed any <b>obstacles</b> that get in her way"

## Logic

Reasoning around why the agent acted as they did

Explanation Logic	Reasoning based on the logic provided in the explanation	"It was mentioned [in the explanation] previously that the <b>yellow ducks</b> are dangerous, so I assume she should avoid them"
Real Life - Object Logic	Reasoning based on what they know about objects in real life	"I went from <b>usable tools</b> to <b>least usable</b> or beneficial"
Real Life - People Logic	Reasoning based on what they know about how people think or act in real life	"When Jade listed these options they did so in this order, when people recall objects they tend to recall the more fond objects first. So I rearranged the items in [this] order"
Real Life - Other Games Logic	Reasoning based on around what they know from playing other video games	"I figure in video games having tools is usually very important..."
Image Event Logic	Reasoning based on an event in the game depicted in the images or in the non-explanation portion of descriptions	"I ranked the <b>flashlight</b> as a factor of highest importance, since it seems to have been used for getting a coaster. It [the flashlight] is gone in the 'Finish' picture"
Other Logic	Any other reasoning used, including when no specific reason mentioned or if they say they chose at random	"It felt right to me, logical but with little actual facts behind it"

**Fig. 12** Summary table of the codes used to understand how users across the explanation groups utilized information from the explanations. Element codes refer to parts of the environment that participants

thought the agent would interact with. Logic codes refer to participants' beliefs around the circumstances leading to the agent's actions

as opposed to the other experimental groups that could not provide this information. Hence, those presented with local Experiential Explanations had significantly higher correct predictions and reliance on the elements and logic of the explanation given, compared to baselines. The agent value explanations baseline did not have information that could be translated to make downstream predictions, and so those participants relied more heavily on **Real Life Object** and **Other Logic** to fill in the gaps around what the agent might do. Similarly, participants who were not provided an explanation at all were left in the dark entirely, relying mostly on **Image Event Logic**, **Real Life Object**, **People**, and **Other Games Logic**. Further, no participants in either the agent value explanations or the no explanation groups were able to correctly predict the agent's actions and order the agent's priorities as compared to 13 of

the 26 local Experiential Explanation participants who were able to correctly do both.

The qualitative analysis of participants' responses to how well they felt they understood the agent revealed implications for the design of explanations in complex environments. Many people in the local Experiential Explanation group directly relied on the explanation and felt they understood the agent. For example, one participant clearly stated what several others echoed, "They made it clear what her thought[t] process was, including her priorities" (L31).<sup>4</sup> However, others expressed that they wanted more from the explanations. As one participant said, "The descriptions explain *what* Jade [the agent] is doing, but *not necessarily why*" (L28, emphasis added).

<sup>4</sup> The first letter in the participant code indicate their explanation group. (L) Local Experiential Explanations, (A) agent value explanation and (N) no explanation.

## Present & Missing Features

Reasoning around what features were present or missing from explanations

Game Goals	References what the objective of the game is	"While the descriptions give you a breakdown on her priorities, you have no idea her motives or the <b>motives of the game</b> . Not even sure what the goal is."
Game Rules	References the set of principles governing the way the game works or how values or states are assigned (cannot mention agent)	"None of it had the context of <b>how the game functions</b> ."
Game World	General reference to the environment in which the agent is acting	" <b>Everything just seemed so disconnected</b> and there wasn't much explanation at all about <b>what was going on</b> ..."
Agent Actions	References what the agent does in the game	"The descriptions explain <b>what Jade is doing</b> , but not necessarily why."
Agent Strategies	References why the agent is acting as she is and/or how the agent coordinates their actions to achieve a goal	"They made it clear <b>what her thought process was</b> , including her <b>priorities</b> ."
Agent Goals	References what the agent is trying to accomplish overall, the agent's motives, and/or a desired outcome	"While the descriptions give you a breakdown on her priorities, you have no idea her motives or the motives of the game. Not even sure what the goal is."

**Fig. 13** Summary table of the codes used to understand what features users perceived as present and missing from the explanations. Features are separated by if they tie to the game or the agent

**The Multi-Dimensional Nature of Users' Needs in Complex Environments** The nature of the task is extremely disguised to ensure that participants would not be able to rely on prior knowledge of the Crafter game or commonsense. This is important for understanding how much information can be derived from the explanations and put to use in making predictions about the agent and the environment. The extreme disguise of the task was also useful in that it elicited feedback about what users *wanted* regardless of whether they were successful or unsuccessful at the task. These are mainly reflected in the *Game (Goals, Rules, World)* and the *Agent (Actions, Strategies, and Goals)* codes. These desires were not only expressed in the local Experiential Explanation group, but were also expressed across the agent and no explanation groups. For example, a participant in the no explanation group who did not get any correct answers said "There is no context or logic to the decision-making process from the information given" (N16), which almost exactly mirrored the statement of a local Experiential Explanation participant who correctly responded to all the questions, "I don't know what she [the agent] chose or what her goal is, so her motivation doesn't make sense to me" (L19). Despite the ability to perform well on the task, the participant still desired additional information beyond what influence predictors could provide, revealing the multi-dimensional nature of the needed explanations. An explanation of the rewards influences on the agent through influence predictors helps clarify some of the causal dimensions, which were enough to predict the agent's functions, yet they fall short of providing a comprehensive understanding of

the agent's underlying motivations or broader environmental context that remains outside the scope of the reward function.

Participants' desire for additional information isn't a shortcoming of our method as participants in the Local Experiential Explanation group did well on the predictive task; it shows areas for future work addressing the multi-dimensional nature of users' needs in complex environments. Since influence predictors tracked the sources of reward for the agent, local Experiential Explanations extracted and provided participants with all the information possible from the model. The information we provided about the circumstances related to the agent's actions still significantly helped participants predict the circumstances leading to the agent's actions. Consequently, the request of some participants can help identify where new explanation methods can be developed to start meeting users' other needs. Looking to the codes around present and missing features helps highlight some of these directions, including users' desires around the *Game* (i.e., *Goals, Rules, World*) and their desires around the *Agent* (i.e., *Actions, Strategies, and Goals*), which were echoed across all experimental groups.

**Examining How Participants Attempt to Fill the Gaps to Complete their Understanding of the Agent.** Because of the extreme disguise of the tasks, participants expressed an unanticipated degree of creativity and curiosity to explain the actions of the agent. The environment in which we placed the participants was intentionally opaque to ensure that they relied on the information provided in the explanations, rather than relying on their knowledge from other games or experiences. Still, participants used their imaginations to fill in the

gaps around what was not provided in the explanations and reconcile the agent’s model with their knowledge. For example, one participant hypothesized that the gray ducks were actually “evil robot ducks” (L12). Another participant created a full narrative around the explanation about how Jade is trying to escape a cave (N16). There was also a plethora of hypotheses around healing apples (L25), or Jade’s desire for a new pet lizard or crab (N25, L26, A27, L31). We set up Jade as an agent in an environment that would not make intuitive sense, giving participants the job of predicting her task with an explanation. The participants filled out her story around the explanations they were given, adding motive and context where we did not.

In some contexts, creativity is not a desired feature. For example, it is (probably) not good for doctors (or others) to creatively invent reasons around AI recommendations for medical environments. Similarly, users should not be creatively attributing reasons their bank loan was rejected. However, in storytelling and games, creativity and curiosity are critical features of a positive and engaging experience. This opens up the direction of future work for these environments to understand how explanations can be a feature that encourages creativity and curiosity while maintaining high task performance and understanding of the agent and promoting an overall positive experience. We encourage future work to consider exploring explanation methods as places to which people can tether their imaginations, building on explanations to construct fun, interesting storytelling experiences.

## 7 Limitations and future work

Our experiments focus on the content of the explanation rather than optimizing the language delivery or comparing different styles of language explanations. Templates were sufficient to show a significant effect on human participants regardless of the acknowledged limitations of templates such as rigidity and formulaicity of text. As such, we evaluated against baselines with other presentational modalities, such as visual or numeric modalities, instead of language variation baselines. As more verbal and fluent explanation methods are developed, more research should consider how different presentational aspects of language generation can impact users’ performance and preference for explanation. Preliminary investigation by Tambwekar and Gombolay [38] suggests that surface-form language variation may not have a large impact.

Our approach highlights broader trends in the influences presenting an overview of how they collectively impact the agent’s actions over time. While this provides a helpful summary, there is potential to derive richer insights by examining these interactions in greater depth. In certain contexts, the

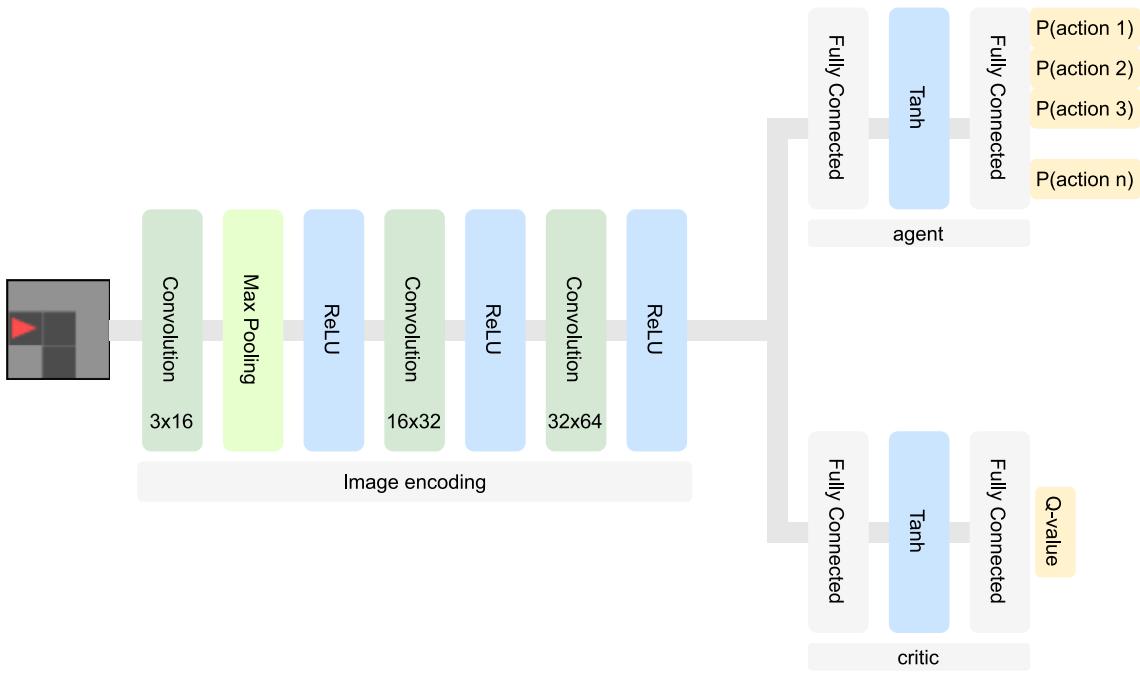
interactions between these influences may provide valuable insights for system users. Nonetheless, it remains unclear whether such detailed explanations are feasible or necessary.

We do not consider the question of *when* users should ask for explanations; we make the simplifying assumption that an explanation generation system should be able to generate an explanation at any point in response to any “why-not” question from an operator for any action performed by the agent. However, further research is required to understand the cognitive load involved in explanations that are presented in real-time in an interactive context. Determining the optimal timing and method for presenting these real-time insights will be crucial to ensure they enhance user understanding without causing undue strain. For example, explanations given in real time in an environment like Crafter might be disruptive, whereas an *after-action review* [39] may allow an operator to revisit a previous point in time after the fact to seek explanations.

Future research in this domain should explore the potential of the multi-dimensional nature of explanations to more effectively reconcile users’ mental models of agents by addressing the disconnects in knowledge that might arise in understanding any part of the system. This exploration can include the integration of various sources of explanation, different modalities, and innovative interaction techniques. For example, exploring how increased interactivity and allowing users to ask varying questions at different times would have impacts on the users’ performance and preference of explanations [40]. Explainable RL will be a crucial component as reinforcement learning algorithms begin to drive agents and rise of human-robots collaboration settings. Future research should investigate how this approach could be extended to continuous environments, robotics, or multi-agent systems to enhance its scope and applicability.

## 8 Conclusions

Explaining reinforcement learning to non-AI expert audiences is challenging. Actionable local explanations should reference future anticipated interactions with respect to the environment. Our Experiential Explanations technique attempts to capture the “experiences” of the agent in terms of reward influences that can be relayed to the user in the form of explanations that help them construct an appropriate mental model of how the agent interacts with the environment. Experiential Explanations provide context to the RL agent’s behaviors without changing the agent architecture while still providing a high degree of faithfulness. This is significant because any core RL algorithm can be used to train the influence predictors, as we demonstrated through our studies using both Deep *Q*-Networks and PPO. Quantitative and qualitative studies with non-AI expert study participants



**Fig. 14** Architecture of the agent A2C model

ponents show that Experiential Explanations help users better understand the agent and predict its behaviors in partially-observable long-horizon discrete planning problems with sparse rewards. They also show that Experiential Explanations are preferred and found more useful than the compared agent explanation baselines. The experimental results suggest that Experiential Explanations may provide for more actionable ways for non-AI expert users to understand the behavior of reinforcement learning agents.

## A Implementation details of the MiniGrid agent

Figure 14 shows the architecture of our agent. The state observation is encoded using a convolutional neural network (CNN) and then passed to the actor-critic network heads composed of two shallow MLPs with one hidden layer each. The MiniGrid agent was implemented using the torch-ac implementation presented in the official MiniGrid documentation.<sup>5</sup>

We train our agents with the following hyperparameters:

processes = 20  
frames = 1221120  
batch size = 256  
discount factor = 0.99  
learning rate =  $3.5 \times 10^{-4}$

<sup>5</sup> <https://github.com/lcs/willem/rl-starter-files/>

## PPO hyperparameters:

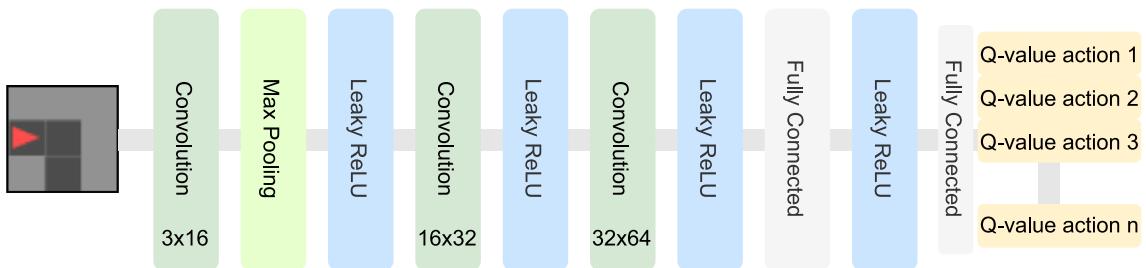
epochs = 10  
lambda coefficient in GAE = 0.95  
entropy term coefficient = 0.01  
value loss term coefficient = 0.5  
maximum norm of the gradient = 0.5  
optimizer = *Adam*  
optimizer epsilon =  $1 \times 10^{-8}$   
clipping epsilon for PPO = 0.2

## B Implementation details of the MiniGrid influence predictors

Our DQN influence predictors are implemented in Pytorch<sup>6</sup> and all the code is available in our git repository. Figure 15 details the architecture of the influence predictor's model. We train the influence predictor with the following hyperparameters:

replay buffer capacity = 50000  
gamma = 0.5  
target update = 5

<sup>6</sup> [https://pytorch.org/tutorials/intermediate/reinforcement\\_q\\_learning.html](https://pytorch.org/tutorials/intermediate/reinforcement_q_learning.html)



**Fig. 15** Architecture of the influence predictor network

bootstrap = 100000  
 batch size = 64  
 learning rate =  $1 \times 10^{-5}$   
 optimizer = RMSprop

## C Implementations details of Crafter's agent and influence predictors

As mentioned earlier, both the agent and the influence predictors are implemented using Stable-Baselines3's [41] PPO [34] implementation. Given the observations are in the form of images, the policy used here ActorCriticCnnPolicy (also from the Stable-Baselines3 [41]).

They use the same values as hyperparameters. The training was done on a single node of a single a40 processor. The training time can be reduced by parallelizing over multiple nodes/processors. The hyperparameters were obtained from [42].

Hyperparameters for the implementations:

nodes = 1  
 timesteps = 8994816

### PPO hyperparameters:

learning rate =  $3 \times 10^{-4}$

Number of steps per env per update = 4096

gamma = 0.95

Number of epochs = 4

batch size = 128

lambda coefficient in GAE = 0.65

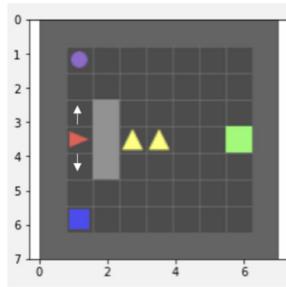
entropy term coefficient = 0.0

value loss term coefficient = 0.5

maximum norm of the gradient = 0.5

optimizer = Adam

clipping parameter = 0.2



Where do you think the agent [red triangle] will go?

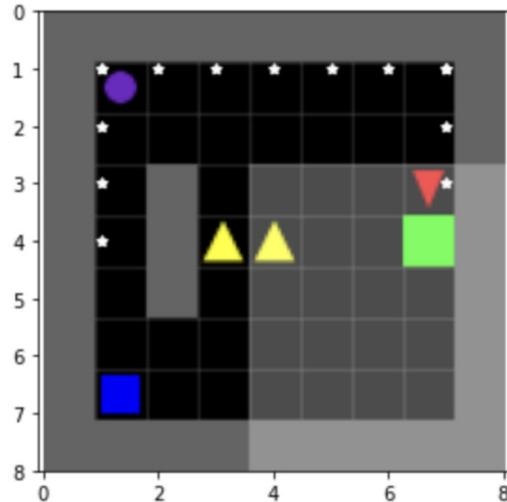
The agent will go up.  
 The agent will go down.

**Fig. 16** Example of the prediction task. The participants saw the map and predicted which direction they thought the agent preferred

## D Experiment 1: evaluation survey

The human evaluation study in MiniGrid is made up of three parts. The first part is a prediction task, as shown in the example in Fig. 16. We presented three episodes to each participant, as shown in Fig. 3. During each episode, the agent—the red triangle—must navigate to the green square. Some colored shapes are associated with negative rewards, but the participant is not told which ones. After each episode, the participant is asked to predict whether the agent will go up or down. The participant has no information at the beginning of the first episode to make an informed choice, and the participant choice is random—we do not evaluate on this choice. Further, upon the participant's choice, we assign which object is dangerous such that the participant is always wrong. This established the conditions for which the participant must now learn to correct their understanding of the agent's interactions with the environment. We use that assignment of dangerous object consistently throughout the second and third episodes. The participant tries to predict the

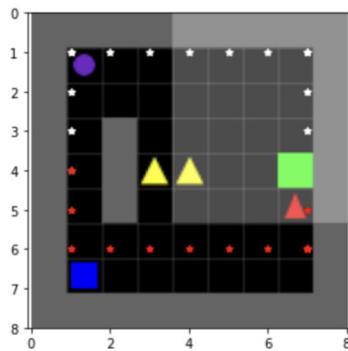
The correct answer is: The agent chose to go up.



**Fig. 17** After predicting the agent's following action, all participants will see the correct action. But for participants in the no explanation group, this is the only information they are given

#### Explanation:

If the agent goes down (the path shown in red stars), it will pass through regions influenced by dangerous blue obstacles; going up feels safer.



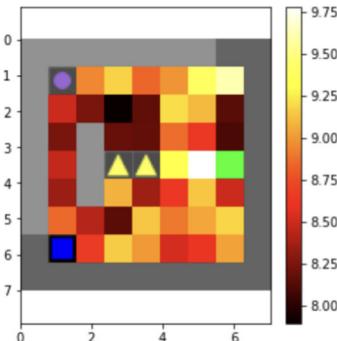
**Fig. 18** Experiential Explanation highlights the dangerous objects and compares the agent's actual and suggested path

agent's choice after the second and third episodes (without the object assignment switching); we expect to see the participants in the experimental condition to have higher accuracy rates than in the baseline conditions.

After each episode, participants fill out a brief text entry question explaining the rationale behind their choice. Then

#### Explanation:

This is the agent's maximum anticipated future rewards (Q values) distribution on the map. (The higher value is, the more favorable a state is).

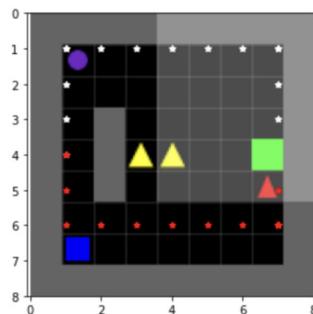


**Fig. 19** Heatmap explanation shows the maximum anticipated future reward for each state in the environment

#### Explanation:

The following is a comparison of the agent's anticipated future rewards (Q values) assigned to the path it took and the alternative path. (The higher this value is, the more favorable a path is).

- Q average if the agent goes up (shown in white): 9.85
- Q average if the agent goes down (shown in red): 9.69



**Fig. 20** Q-value explanation compares the actual and suggested paths regarding their average anticipated future reward

after the participants submit their predictions and reasoning, they see the correct answer as shown in Fig. 17. Afterward, they see one of the three types of explanations, Experiential Explanation (Fig. 18), Heatmap Explanation (Fig. 19), Q-value explanation (Fig. 20), or they see no explanation at all. Then they repeat this process two times.

After the participants finish the prediction task, they proceed to the second part of the study, a satisfaction survey,

Based on your experience with the agents explanations please evaluate them on the following points

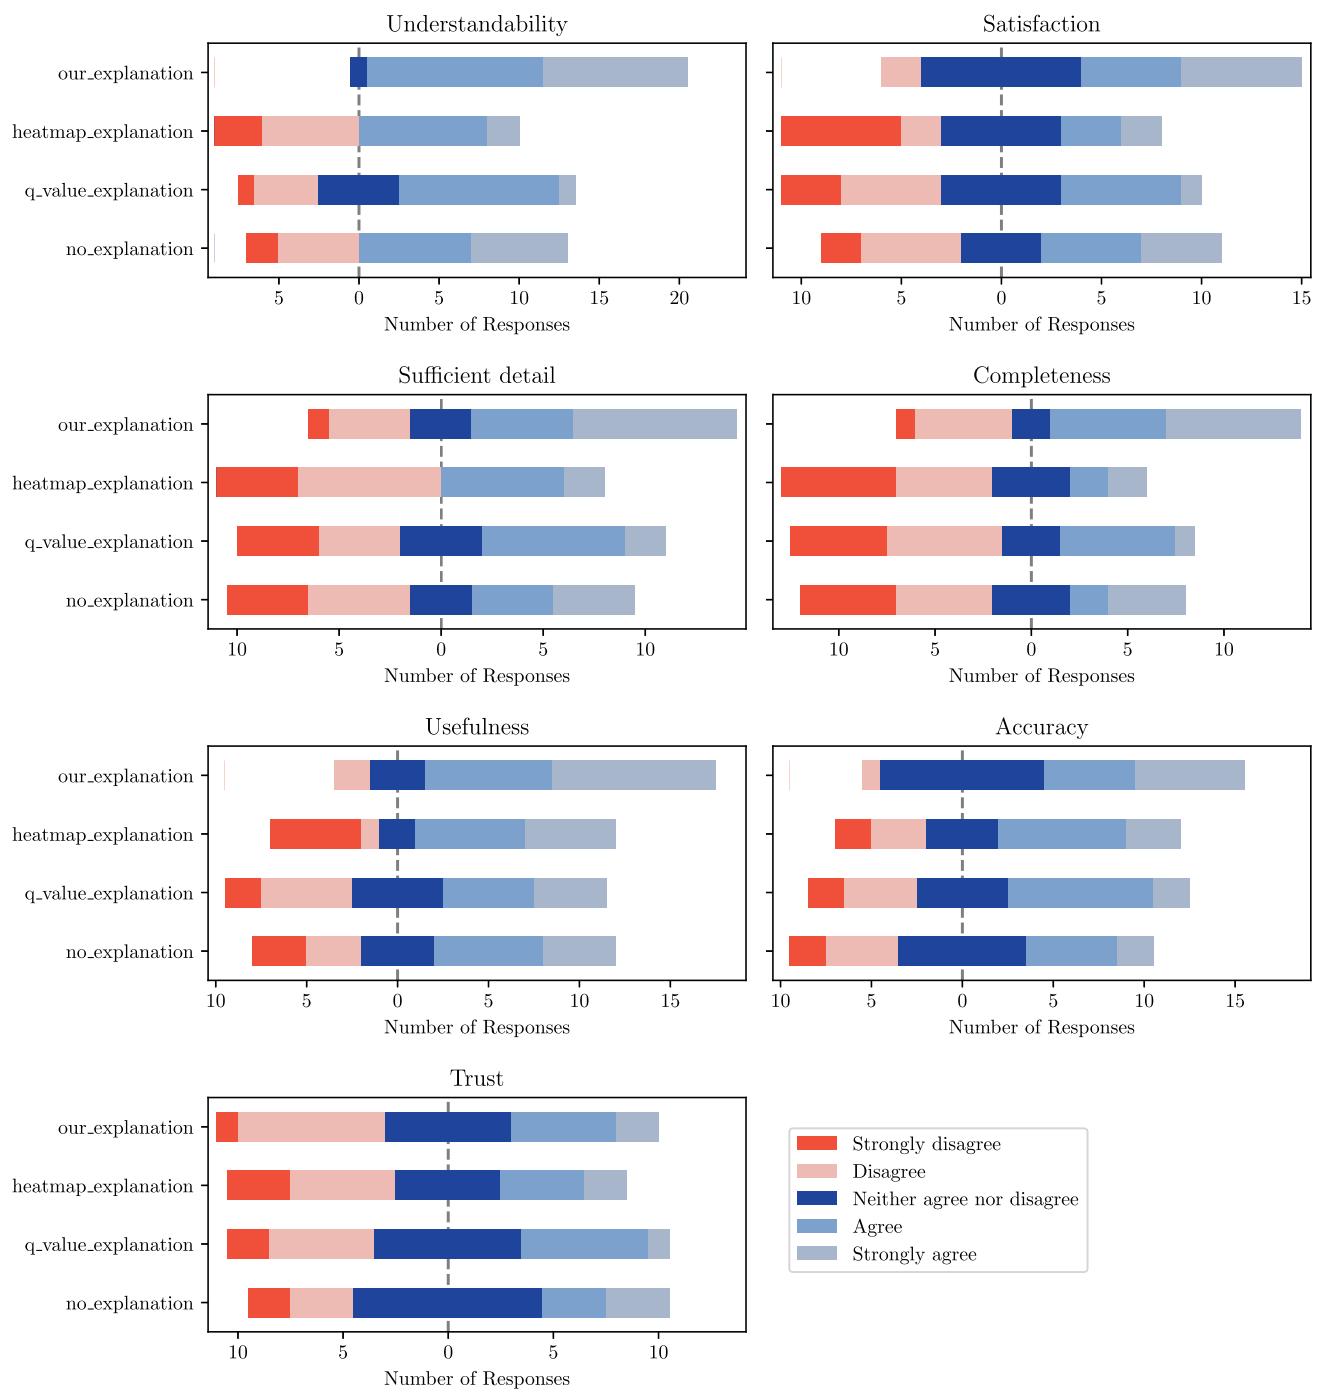
	Strongly agree	Somewhat agree	Neither agree nor disagree	Somewhat disagree	Strongly disagree
From the explanations I <b>understand</b> how the agent navigates the environment.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
These explanations of how the agent works are <b>satisfying</b> .	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
These explanations of how the agent works have <b>sufficient detail</b> .	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
These explanations of how the agent works seem <b>complete</b> .	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
These explanations of how the agent works are <b>useful to my goals</b> .	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
These explanations of how the agent works show me how <b>accurate</b> the agent is.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
These explanations lets me judge when I should <b>trust and not trust</b> the agent.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

**Fig. 21** After the experiments, the participants completed the shown satisfaction survey evaluating the explanation type they interacted with on a seven-point scale, measuring helpfulness, satisfaction, understandability, level of detail, completeness, accuracy, and trust

consisting of five-point Likert scale questions asking the user to rate understandability, satisfaction, amount of details, completeness, usefulness, precision, and trust for the explanations they saw. It also had an open-ended response question asking how the explanation helped or did not help them understand the agent. This survey is shown in Fig. 21.

## E Experiment 1: satisfaction survey analysis

Figure 22 shows the results of the satisfaction survey questions shown in Fig. 21. The discussion and analysis of Fig. 22 results are in the main paper.



**Fig. 22** Experiment 1 satisfaction survey results for each explanation quality dimensions grouped by explanation type

Q561



Timestep = 57

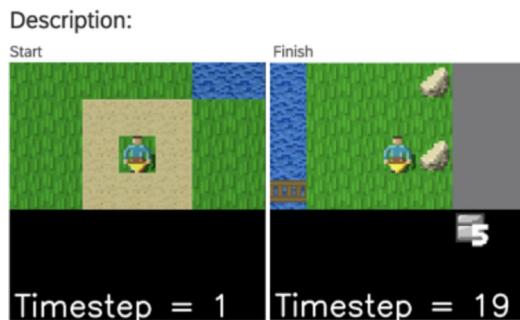
This picture displays Jade 🧑 standing in the middle, with her inventory shown in a black rectangle below. Each action may require one or more time steps to finish. She is currently at step 57.

Q562

Jade needs to decide what to do next, the options are to get a coaster or get a fan, which one would she do?

Get a coaster  
 Get a fan

**Fig. 23** First decision point the participants encounter during experiment 2



Jade went to the sparse rocks. Jade assigned a value to each state based on expected future rewards, reflecting her proximity to her goal. Below are the values of each state after each action. The higher the value, the better the expected outcome.

Jade's value for getting rock slabs from the sparse rocks is 6.99. While Jade's value for getting rock slabs from dense rocks is 5.59

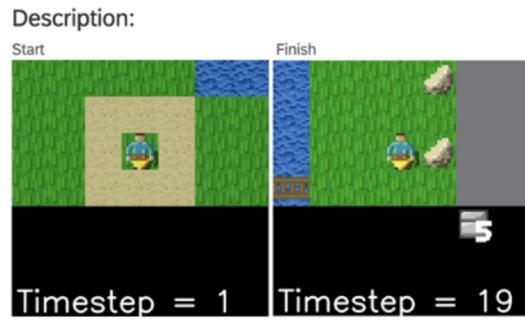
**Fig. 24** Example of the agent value explanation baseline

## F Crafter re-skin mappings

Table 5 lists a complete mapping of assets re-skinned assets

## G Experiment 2: evaluation survey

The human evaluation study in Crafter has three tasks, a prediction task and an ordering task, and a satisfaction survey. We picked four random decision points from the agent's trajectory, each paired with a counterfactual trajectory that the



Jade needed rock slabs and chose to get them from the sparse rocks cluster. This prompted her to prioritize adopting a parrot, owning a plier and getting scissors next.

However, if she had gotten to the dense rocks, she would have to think about owning a cutter, fearing for her health and avoiding dangerous yellow ducks instead.

**Fig. 25** Example of the local experiential explanation

agent can take by performing a counterfactual action at this decision point. We presented each participant with one decision point and asked them to predict the agent's action as shown in Fig. 23. After this first interaction, they are shown their assigned explanation. The explanations are one of the three baselines specified in Sect. 6.3. An agent value explanation (Fig. 24), the local experiential explanation (Fig. 25), or they do not see an explanation.

Then after seeing the explanation, we asked them to predict and reason about the agent's action at a future decision point displayed in Fig. 26. Lastly, the participants engage

Q567

Jade is now standing in the middle of three colored ducks: a brown one, a gray one, and a yellow one. Which direction should she avoid?

The yellow ducks  
 The gray ducks  
 The brown ducks

**Fig. 26** Second decision point the participants encounter during experiment 2. Which is used for the prediction task and followed by the reasoning question

**Table 5** Original crafter assets vs their rock-paper-scissors (RPS) asset collection counterparts

Actual asset	Re-skinned asset
Materials and objects	
Player (Jade)	<i>No change</i>
Tree	Rock
Grass	<i>No change</i>
Sand	<i>No change</i>
Water	<i>No change</i>
Stone	Scissors
Iron	Paper
Coal	Lizard
Diamond	Electric torch
Zombie	Duck
Table	Parrot
Furnace	Crab
Bridge	<i>No change</i>
Cow	Apple
Inventory Items	
Wood	Rock slabs
Wooden pickaxe	Cutter
Wooden sword	Throwing knives
Iron pickaxe	Coaster
Iron sword	Fan
stone pickaxe	Plier
stone sword	Hammer
Parameters	
Health	<i>No change</i>
Food	<i>No change</i>
Drink	<i>No change</i>
Energy	<i>No change</i>

### Getting Rock slabs from the sparse cluster.

parrot

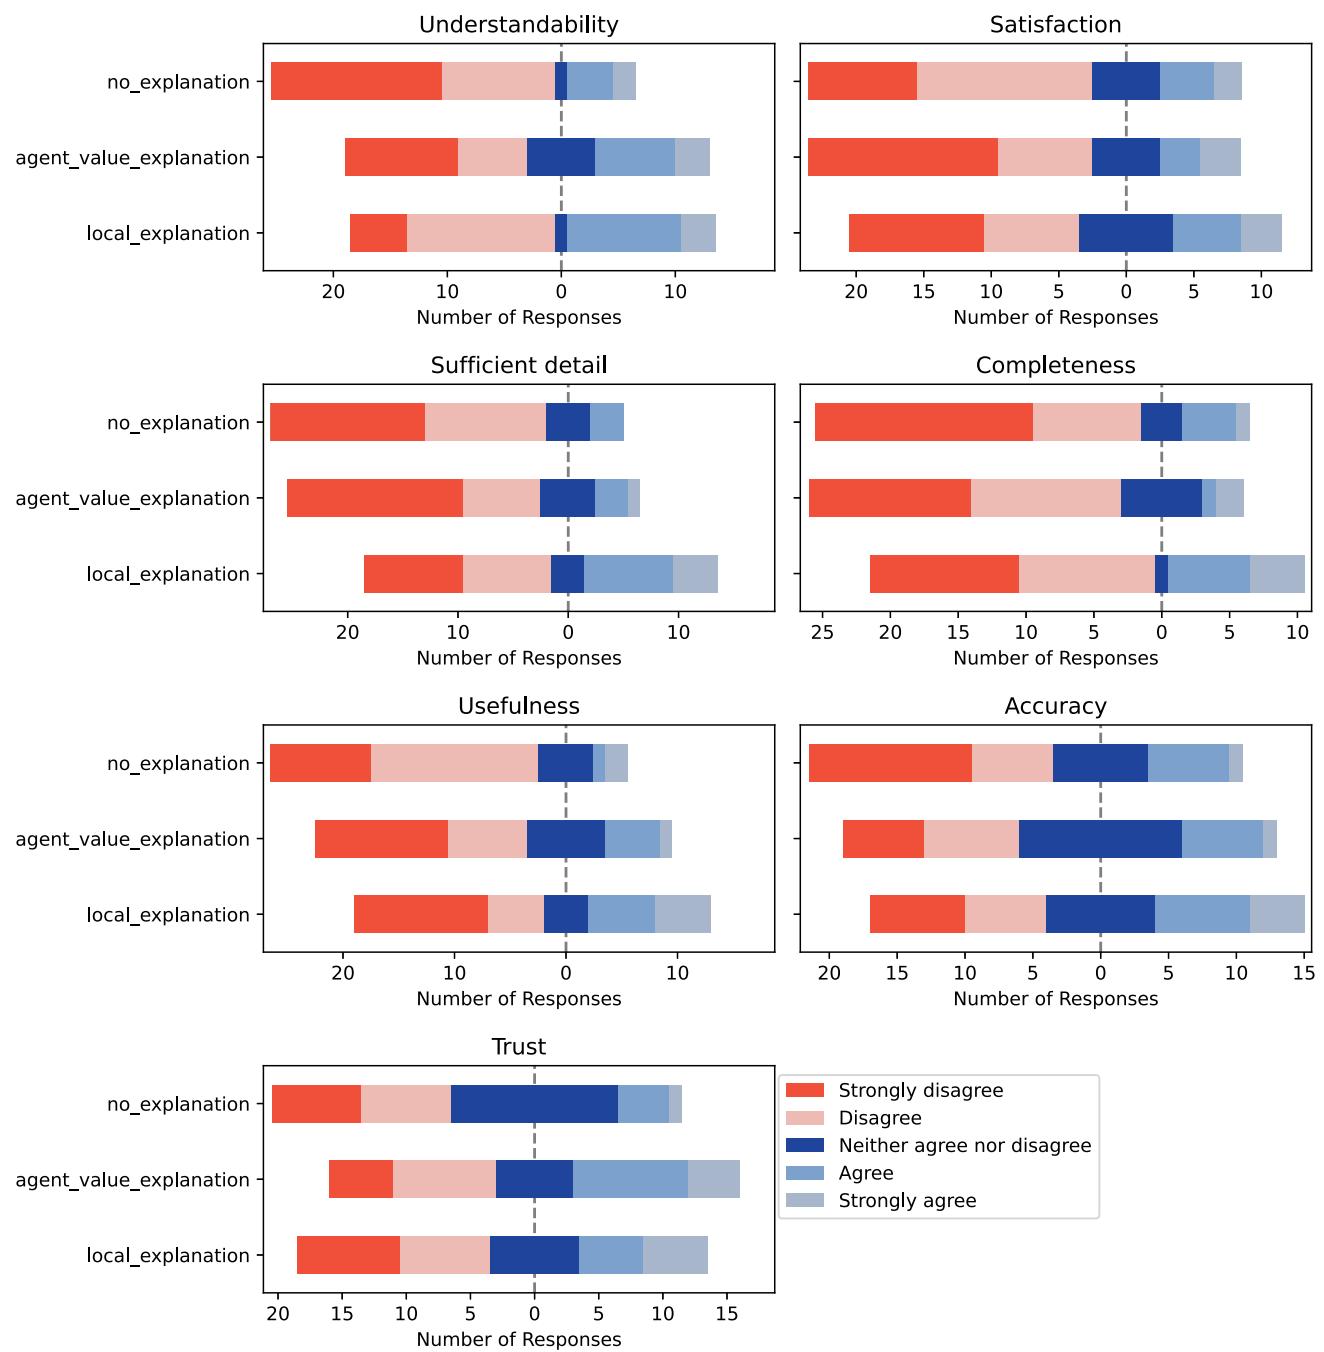
pliers

scissors

**Fig. 27** Ordering task which asks the participants to order the influences by their importance to the agent at the decision point

in an additional task in order of the influences by their importance to the agent when it considers performing the counterfactual action or the actual action as shown in Fig. 27. Both tasks are followed by a text entry question asking them to reason about their predictions.

After participants complete both tasks, they proceed to the last part of the study, a satisfaction survey identical to the one we used in experiment 1 as shown in Fig. 21 consisting of five-point Likert scale questions asking the user to rate understandability, satisfaction, amount of details, completeness, usefulness, precision, and trust for the explanations they saw. It also had an open-ended response question asking how the explanation helped or did not help them understand the agent. The results of the survey for Experiment 2 are shown in Fig. 22.



**Fig. 28** Experiment 2 satisfaction survey results for each explanation quality dimensions grouped by explanation type

**Acknowledgements** The authors thank the Entertainment Intelligence and Human-Centered AI Labs members at the Georgia Institute of Technology for their constructive suggestions and insightful discussions. Their contributions have been invaluable to the progress and completion of this research.

**Author Contributions** Conceptualization was performed by Amal Alabdulkarim, Mark O. Riedl; methodology by Amal Alabdulkarim, Madhuri Singh, Mark O. Riedl; formal analysis and investigation by Amal Alabdulkarim, Madhuri Singh, Gennie Mansi, Kaely Hall; writing—original draft preparation—by Amal Alabdulkarim, Madhuri Singh, Gennie Mansi; writing—review and editing—by Amal Alabdulkarim, Mark O. Riedl, Madhuri Singh, Gennie Mansi, Kaely Hall, Upol Ehsan; supervision by Mark O. Riedl

**Funding** The authors did not receive support from any organization for the submitted work.

**Data Availability** The authors declare that the code and data supporting the findings of this study are available in the manuscript and its supplementary information files.

## Declarations

**Conflict of interest** The authors have no conflict of interest to declare that are relevant to the content of this article.

**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

- Miller T (2019) Explanation in artificial intelligence: insights from the social sciences. *Artif Intell* 267:1–38
- Zelvelder AE, Westberg M, Främling K (2021) Assessing explainability in reinforcement learning. In: International workshop on explainable, transparent autonomous agents and multi-agent systems, pp 223–240. Springer
- Anjomshoae S, Najjar A, Calvaresi D, Främling K (2019) Explainable agents and robots: results from a systematic literature review. In: Proceedings of the 18th international conference on autonomous agents and MultiAgent systems. AAMAS '19, pp. 1078–1088. International foundation for autonomous agents and MultiAgent systems, Richland, SC
- Alharin A, Doan T-N, Sartipi M (2020) Reinforcement learning interpretation methods: a survey. *IEEE Access* 8:171058–171077
- Wells L, Bednarz T (2021) Explainable AI and reinforcement learning—a systematic review of current approaches and trends. *Front Artif Intell* 4:550030
- Heuillet A, Couthouis F, Díaz-Rodríguez N (2021) Explainability in deep reinforcement learning. *Knowl-Based Syst* 214:106685. <https://doi.org/10.1016/j.knosys.2020.106685>. (Accessed 2022-09-30)
- Puiutta E, Veith EMSP (2020) Explainable reinforcement learning: a survey. In: Holzinger A, Kieseberg P, Tjoa AM, Weippl E (eds) Machine learning and knowledge extraction. Lecture notes in computer science. Springer, Cham, pp 77–95. [https://doi.org/10.1007/978-3-030-57321-8\\_5](https://doi.org/10.1007/978-3-030-57321-8_5)
- Milani S, Topin N, Veloso M, Fang F (2023) Explainable reinforcement learning: a survey and comparative review. *ACM Comput Surv*. <https://doi.org/10.1145/3616864>
- Chakraborti T, Sreedharan S, Kambhampati S (2020) The emerging landscape of explainable automated planning & decision making. In: Proceedings of the twenty-ninth international joint conference on artificial intelligence, pp. 4803–4811. International joint conferences on artificial intelligence organization, Yokohama, Japan. <https://doi.org/10.24963/ijcai.2020/669>. <https://www.ijcai.org/proceedings/2020/669>. Accessed 11 Mar 2022
- Sutton RS, Barto AG (2018) Reinforcement learning: an introduction. A Bradford Book, Cambridge
- Chakraborti T, Sreedharan S, Kambhampati S (2021) The emerging landscape of explainable automated planning & decision making. In: Proceedings of the twenty-ninth international conference on international joint conferences on artificial intelligence, pp 4803–4811
- Juozapaitis Z, Koul A, Fern A, Erwig M, Doshi-Velez F (2019) Explainable reinforcement learning via reward decomposition. In: IJCAI/ECAI workshop on explainable artificial intelligence
- Das D, Chernova S (2020) Leveraging rationales to improve human task performance. In: Proceedings of the 25th international conference on intelligent user interfaces, pp 510–518
- Anderson A, Dodge J, Sadarangani A, Juozapaitis Z, Newman E, Irvine J, Chattopadhyay S, Fern A, Burnett M (2019) Explaining reinforcement learning to mere mortals: an empirical study. CoRR [arXiv:1903.09708](https://arxiv.org/abs/1903.09708)
- Septon Y, Huber T, André E, Amir O (2023) Integrating policy summaries with reward decomposition for explaining reinforcement learning agents. In: International conference on practical applications of agents and multi-agent systems, pp 320–332. Springer
- Waa J, Diggelen J, Bosch Kvd, Neerincx M (2018) Contrastive explanations for reinforcement learning in terms of expected consequences. [arXiv:1807.08706](https://arxiv.org/abs/1807.08706)
- Madumal P, Miller T, Sonenberg L, Vetere F (2020) Explainable reinforcement learning through a causal lens. In: Proceedings of the AAAI conference on artificial intelligence, vol. 34, pp 2493–2500
- Sreedharan S, Soni U, Verma M, Srivastava S, Kambhampati S (2022) Bridging the gap: providing post-hoc symbolic explanations for sequential decision-making problems with inscrutable representations. In: International conference on learning representations. <https://openreview.net/forum?id=o-1v9hdSult>
- Frost J, Watkins O, Weiner E, Abbeel P, Darrell T, Plummer B, Saenko K (2022) Explaining reinforcement learning policies through counterfactual trajectories. In: ICML workshop on human in the loop learning (HILL)
- Olson ML, Khanna R, Neal L, Li F, Wong W-K (2021) Counterfactual state explanations for reinforcement learning agents via generative deep learning. *Artif Intell* 295:103455
- Huber T, Demmler M, Mertes S, Olson ML, André E (2023) Counterfactual-RL: understanding reinforcement learning agents' strategies through visual counterfactual explanations. In: Proceedings of the 2023 international conference on autonomous agents and multiagent systems, pp 1097–1106
- Ehsan U, Tambwekar P, Chan L, Harrison B, Riedl MO (2019) Automated rationale generation: a technique for explainable AI and its effects on human perceptions. In: Proceedings of the 24th international conference on intelligent user interfaces, pp 263–274

23. Das D, Chernova S, Kim B (2024) State2explanation: concept-based explanations to benefit agent learning and user understanding. In: Proceedings of the 37th international conference on neural information processing systems. NIPS '23. Curran Associates Inc., Red Hook
24. Ng AY, Harada D, Russell S (1999) Policy invariance under reward transformations: theory and application to reward shaping. In: Icml, vol. 99, pp 278–287. Citeseer
25. Mnih V, Kavukcuoglu K, Silver D, Rusu AA, Veness J, Bellemare MG, Graves A, Riedmiller M, Fidjeland AK, Ostrovski G et al (2015) Human-level control through deep reinforcement learning. *Nature* 518(7540):529–533
26. Mnih V, Badia AP, Mirza M, Graves A, Lillicrap T, Harley T, Silver D, Kavukcuoglu K (2016) Asynchronous methods for deep reinforcement learning. In: International conference on machine learning, pp 1928–1937. PMLR
27. Hafner D, Lillicrap T, Norouzi M, Ba J (2020) Mastering Atari with discrete world models. [arXiv:2010.02193](https://arxiv.org/abs/2010.02193)
28. Dazeley R, Vamplew P, Foale C, Young C, Aryal S, Cruz F (2021) Levels of explainable artificial intelligence for human-aligned conversational explanations. *Artif Intell* 299:103525
29. Chevalier-Boisvert M, Willems L, Pal S (2018) Minimalistic Grid-world environment for OpenAI Gym. GitHub. <https://github.com/maximecb/gym-minigrid>
30. Hafner D (2021) Benchmarking the spectrum of agent capabilities. [arXiv:2109.06780](https://arxiv.org/abs/2109.06780)
31. Zhou J, Gandomi AH, Chen F, Holzinger A (2021) Evaluating the quality of machine learning explanations: a survey on methods and metrics. *Electronics* 10(5):593
32. Hoffman RR, Mueller ST, Klein G, Litman J (2023) Measures for explainable AI: explanation goodness, user satisfaction, mental models, curiosity, trust, and human-AI performance. *Front Comput Sci* 5:1096257
33. Willems L (2018) RL Starter Files for Mini-Grid. <https://github.com/lcswillems/rl-starter-files/tree/4205e05b7905fec16519bc0802596673d86af018>. Accessed 25 Sept 2022
34. Schulman J, Wolski F, Dhariwal P, Radford A, Klimov O (2017) Proximal policy optimization algorithms. [arXiv:1707.06347](https://arxiv.org/abs/1707.06347)
35. Braun V, Clarke V (2006) Using thematic analysis in psychology. *Qual Res Psychol* 3(2):77–101
36. Miles MB, Huberman AM, Saldaña J (2020) Qualitative data analysis: a methods sourcebook, 4th edn. SAGE, Los Angeles
37. Chan WT, Chow YK, Liu LF (1995) Neural network: an alternative to pile driving formulas. *Comput Geotech* 17(2):135–156. [https://doi.org/10.1016/0266-352X\(95\)93866-H](https://doi.org/10.1016/0266-352X(95)93866-H)
38. Tambwekar P, Gombolay M (2024) Towards reconciling usability and usefulness of policy explanations for sequential decision-making systems. *Front Robot AI* 11:1375490
39. Dodge J, Khanna R, Irvine J, Lam K-H, Mai T, Lin Z, Kiddle N, Newman E, Anderson A, Raja S et al (2021) After-action review for AI (AAR/AI). *ACM Trans Interact Intell Syst (TiiS)* 11(3–4):1–35
40. Qian P, Unhelkar V (2022) Evaluating the role of interactivity on improving transparency in autonomous agents. In: Proceedings of the 21st international conference on autonomous agents and multiagent systems, pp 1083–1091
41. Raffin A, Hill A, Gleave A, Kanervisto A, Ernestus M, Dormann N (2021) Stable-baselines3: reliable reinforcement learning implementations. *J Mach Learn Res* 22(268):1–8
42. Stanić A, Tang Y, Ha D, Schmidhuber J (2022) Learning to generalize with object-centric agents in the open world survival game crafter

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