

Interestingness elements for explainable reinforcement learning: Understanding agents' capabilities and limitations

Pedro Sequeira*, Melinda Gervasio

SRI International, 333 Ravenswood Avenue, Menlo Park, CA 94025, United States

ARTICLE INFO

Article history:

Received 12 November 2019

Received in revised form 28 July 2020

Accepted 6 August 2020

Available online 12 August 2020

Keywords:

Explainable AI

Reinforcement learning

Interestingness elements

Autonomy

Video highlights

Visual explanations

ABSTRACT

We propose an explainable reinforcement learning (XRL) framework that analyzes an agent's history of interaction with the environment to extract interestingness elements that help explain its behavior. The framework relies on data readily available from standard RL algorithms, augmented with data that can easily be collected by the agent while learning. We describe how to create visual summaries of an agent's behavior in the form of short video-clips highlighting key interaction moments, based on the proposed elements. We also report on a user study where we evaluated the ability of humans to correctly perceive the aptitude of agents with different characteristics, including their capabilities and limitations, given visual summaries automatically generated by our framework. The results show that the diversity of aspects captured by the different interestingness elements is crucial to help humans correctly understand an agent's strengths and limitations in performing a task, and determine when it might need adjustments to improve its performance.

© 2020 Elsevier B.V. All rights reserved.

1. Introduction

Reinforcement learning (RL) is a popular computational approach for autonomous agents facing a sequential decision problem in dynamic and often uncertain environments [1]. The goal of any RL algorithm is to learn a *policy*, i.e., a mapping from states to actions, given trial-and-error interactions between the agent and an environment. Typical approaches to RL focus on memoryless (reactive) agents that select their actions based solely on their current observation [2]. This means that by the end of learning, an RL agent can select the most appropriate action in each situation—the learned policy ensures that doing so will maximize the reward received by the agent during its lifespan, thereby performing according to the underlying task assigned by its designer.

A trained RL agent does not explicitly reason about its future to select actions. All it knows is that it *should* select actions as dictated by its learned policy, which makes it hard to explain its behavior. At most, agents know that choosing one action is preferable over others, or that some actions are associated with a higher value—but not why that is so or how it came to be. As a consequence, the *why* behind decision-making is lost during the learning process as the policy converges to an optimal action-selection mechanism. RL further complicates explainability by enabling an agent to learn from *delayed rewards* [1]—the reward received after executing some action is “propagated” back to the states and actions that led to that situation, meaning that important actions may not be associated with any (positive) reward.

* Corresponding author.

E-mail addresses: pedro.sequeira@sri.com (P. Sequeira), melinda.gervasio@sri.com (M. Gervasio).

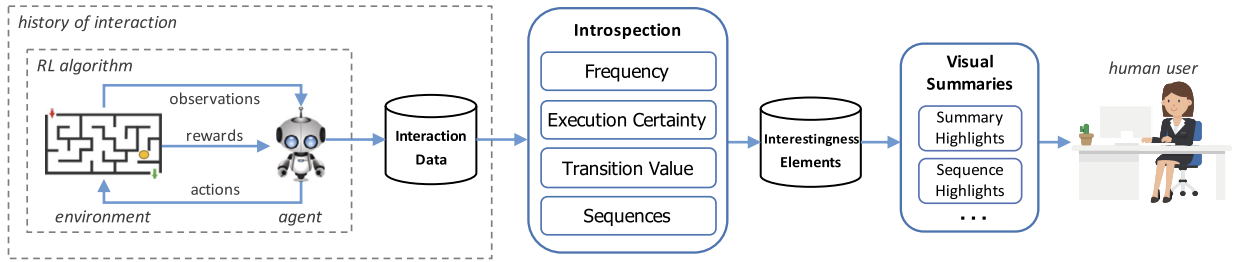


Fig. 1. Our framework for explainable autonomous RL agents. During its interaction with the environment, the agent collects relevant statistical information. Our introspection framework analyzes the collected data and identifies interestingness elements of the history of interaction. These elements are then used to summarize an agent's learned behavior and its capabilities and limitations in the task to a human user.

Ultimately, RL agents lack the ability to know why some actions are preferred over others, to identify the goals that they are currently pursuing, to recognize what elements are more desirable, to identify situations that are “hard to learn”, or to summarize the strategy learned to solve the task.

This lack of self-explainability can be detrimental to establishing trust with human collaborators who may need to delegate critical tasks to agents. An eXplainable RL (XRL) system that enables humans to correctly understand the agent's *aptitude* in a specific task, i.e., both its capabilities and limitations, whether innate or learned, will enable human collaborators to delegate tasks more appropriately as well as identify situations where the agent's perceptual, actuating, or control mechanisms may need to be adjusted prior to deployment.

1.1. Approach overview

In this paper, we present a framework towards making autonomous RL agents more explainable through introspective analysis (Fig. 1). The agent examines its history of interaction with the environment to extract *interestingness elements*¹ denoting meaningful situations, i.e., potentially “interesting” properties of the interaction. The goal of the elements is to allow humans to understand an agent's underlying characteristics and aptitude in a task, thus making an RL system more *explainable*.² As depicted in Fig. 1, four introspection dimensions are analyzed: frequency, execution certainty, transition-value and sequences.³ Each element is generated by applying statistical analysis and machine learning methods to data collected during the interaction. The introspection framework is the basis of a summarization system following the approach in [5], that extracts visual summaries highlighting relevant moments of the agent's interaction with an environment according to the different interestingness elements.

We conducted a user study to evaluate the usefulness of the proposed interestingness elements in helping characterize the behavior of different RL agents. The goal was to understand *why* and *for what* each element might be useful. We designed agents with diverse goals and perceptual characteristics in a video game task. The goal was to simulate players with different capabilities and limitations, and needing improvement in different aspect of the games. We trained each agent using standard RL, resulting in distinct behaviors and levels of game performance. We then conducted an online study where we asked subjects to observe several videos comprising visual summaries, each depicting the behavior of a particular agent based on specific interestingness elements.

Overall, our results show that different agent characteristics require different elements to be presented to a user to enable a correct understanding of proficiency in different parts of a task. Namely, presenting only “good” or only “bad” instances of agent behavior can lead users to form incorrect expectations about agent performance. The results also show that by capturing the diversity of an agent's interaction with the environment, frequently and infrequently visited states enable a correct understanding of the agent's general aptitude. Moreover, short summaries showing how an agent overcomes challenges and achieves its goals can provide a short but accurate depiction of its learned strategy.

The paper is organized as follows. Sec. 2 discusses related works and compares them with our approach. Sec. 3 introduces the necessary concepts behind the RL framework. Sec. 4 details the proposed interestingness elements while Sec. 5 shows how visual summaries can be derived from them. Sec. 6 presents the experimental study, including the design of the different RL agents and the survey. Sec. 7 details the results of the user study and provides an in-depth discussion of the main insights stemming from the results. Finally, Sec. 8 summarizes the main findings and describes current and future developments.

¹ We borrow the term from the association rule mining literature [see e.g., 3] to denote potentially relevant aspects of the agent's history of interaction with the environment.

² Although the elements and visual summaries generated by the framework *are not* explanations, they allow users to interpret an agent's behavior through the observation of relevant interactions.

³ In this paper we detail only the elements evaluated in our user study. We refer to Appendix A.A for additional elements of different types and to [4] for a preliminary version.

2. Related work

In this section we review recent works within XRL that have addressed some of the problems identified in this paper. We summarize the technical approaches, highlight the main contributions and compare them to our framework.

2.1. Deriving language templates

One approach is to use language templates to translate elements of the problem or the agent's policy into human-understandable explanations. For example, Khan et al. [6] used the discounted state-visitation frequency according to a policy to form contrastive explanations about the expected return in a given state by executing the optimal action versus alternative ones. In particular, their minimal sufficient explanations (MSE) populate templates and convey to the user the probability or frequency of reaching certain states with high rewards by using the optimal action but not other actions. However, explanations can be long and ill-suited to non-technical users, for whom the concept of utility may not be easy to understand.

Elizalde et al. [7] developed an intelligent assistant for operator training (IAOT) relying on factored state scenarios. The system is capable of providing explanations using templates tailored to novice, intermediate, and advanced users. Explanations are generated whenever a mistake is made and include the optimal action that should be performed, a textual justification according to the user level, and a (domain-dependent) visual diagram highlight the most relevant variable. The latter corresponds to the state feature providing the highest positive difference in the value function. A drawback of this method is that it requires creating a knowledge base for the justifications in the templates, e.g., obtained from written documentation or domain experts.

Wang et al. [8] generate explanations for partially-observable MDPs that expose users to the agent's policy, beliefs over states of the environment, transition probabilities, rewards, etc. in human-robot interaction (HRI) scenarios. The agent can refer to properties of the environment and its behavior but cannot identify the most pertinent situations or justify its actions.

Unlike these approaches which require a large amount of domain knowledge to create explanations, our framework relies on standard RL data and is domain-agnostic. In addition, techniques like ours that use visual summaries of behavior make understanding the agent's behavior more accessible to non-expert users.

2.2. Deep RL over image-based inputs

Other approaches aim to provide insights about policies learned using deep RL techniques on image-based input tasks. For example, Zahavy et al. [9] aggregate the high-dimensional state space of ATARI games into a hierarchical structure using t-Distributed Stochastic Neighbor Embedding (t-SNE) for dimensionality reduction. A graphical user interface then allows the manual selection and visualization of the aggregation of different states according to several types of features. However, the representation and provided information make it difficult for non-RL experts to extract meaningful information.

Greydanus et al. [10] use perturbation-based techniques to identify image regions that result in significant policy changes. The result is a video of agent performance where the relevant regions of the image are overlaid with different colors. This approach is computationally demanding as it requires sweeping whole input images at every step, and users of the system would be required to observe full-length videos of an agent's performance to understand its policy.

An advantage of these approaches is that they allow scaling to large input space domains. A drawback is that they cannot generalize to tasks with non-image-based inputs. Our approach is agnostic to the type of observations to derive the different elements. In addition, unlike these approaches that expose the entirety of an agent's behavior, we follow others who select only the moments of the agent's behavior that are relevant according to different criteria.

2.3. Abstracting explainable representations

Koul et al. [11] learn finite state representations of policies learned using deep recurrent neural networks (RNNs) that are agnostic to the input type. The idea is to abstract the learned policy into a simpler, more interpretable structure. First, quantized bottleneck networks (QBNs) are learned from a trained RNN policy that can reconstruct the observations and hidden states using a lower-dimensional representation space. From this, a finite state machine is extracted and discretized. Although the resulting structures provide insights on the nature of the learned policy, e.g., that it contains certain cycles, it is hard to interpret the semantics behind the discretized observations and hidden states.

Other works focus on abstracting state representations of the task and creating graph structures denoting the agent's behavior. For example, Hayes and Shah [12] and van der Waa et al. [13] frame explanations as summaries assembled from outputs of binary classifiers that characterize the agent's trajectories and decisions. The system in [12] was designed for HRI scenarios where users pose questions that are identified from a set of question templates, resolved to a relevant set of world states, and then summarized and composed into a natural language form to respond to the inquiring user.

The approach by van der Waa et al. [13] allows users to ask contrastive questions (why x instead of y ?) about the agent's policy. Similar to our element of sequences, a path is constructed from the most likely state-action pairs. An explanation is then constructed from a language template. A similar approach was proposed by Madumal et al. [14], where action influence

graphs for RL are extracted via structural causal models. This allows generating explanations for *why* and *why not* questions regarding the actions are selected over others. A user study was performed by applying the method to an agent policy learned on a simplified StarCraft II combat game, and results show that the method allowed for a better understanding of the agent's strategy.

In contrast to these approaches, our framework identifies relevant aspects of the agent's behavior according to different criteria and provides the semantics for the visual summaries.

2.4. Extracting behavior examples

Some approaches generate examples of agent behavior in the form of sequences of state-action pairs that maximize the understanding of the agent's policy by a human observer. Such approaches typically rely on inverse reinforcement learning (IRL) to generate optimal examples, i.e., trajectories that, when observed by humans, allow them to recover the reward function used by the agent, which subsequently allows them to reason about the agent's planning process.

For example, Huang et al. [15] model how humans make inferences about an agent's (robot) reward function from examples of its behavior. To select maximally informative examples, authors rely on algorithmic teaching techniques and different distance functions to select trajectories that best describe the reward function when compared to other reward functions, i.e., that maximize the probability of humans inferring the correct agent goals.

Lage et al. [16] use a similar machine teaching approach by producing summaries of policies by extracting trajectories that maximize the quality of their reconstruction according to the principle of maximum entropy. The approach also modeled the human reconstruction of the agent's policy using imitation learning instead of IRL. Human-subject studies on different RL domains showed the importance of selecting appropriate user models for summarization to domain context.

This approach is related to how we select examples of behavior by using the interestingness elements, but it's targeted at having people understand agent goals—or goal-directed behavior—rather than their competency or aptitude in a task. Therefore, the selection criteria of trajectories are based on how informative they are of the agent's typical behavior. In contrast, within our framework, behavior summaries are only *a means* of explaining the agent's behavior; the focus is on identifying situations that are relevant to how the agent performs the task.

2.5. Identifying key interaction moments

Approaches more similar to ours try to identify key moments of the agent's interaction. As proposed by Amir et al. [17], the idea is that behavior summaries or examples should cover states that are likely to be of interest to users. In that paper, the authors provide a conceptual model of the summarization process of agents and propose several strategies for selecting important states, many of which are realized by our proposed interestingness elements.

Huang et al. [18] present a system capturing *critical states*, in which the value of one action is comparatively much higher than others. They performed a user study where subjects were exposed to critical states extracted for two different policies in the game of Pong, where one policy had a much higher training time. Subjects were then asked whether they trusted each policy and whether they would take control over them in a series of query states. Results show that exposing users to critical states enables an appropriate level of trust with regard to the agents' performance.

Our approach for visual summaries is based on the highlights system proposed by Amir and Amir [5]. Their system captured video clips highlighting an agent's behavior in the game of Pacman using the concept of importance, dictated by the largest difference in the Q -values of a given state. This allows the identification of trajectories that are representative of the agent's best performance. Subjects in a user study were shown highlights of three agents with increasing amounts of training (and hence, proficiency) and the results show that the proposed importance metric enabled correct assessment of the agents' relative performance compared to other baseline summarization methods.

A first observation is that we go beyond the concept of importance for explaining the agent's behavior proposed in these papers as dictated solely by the highest value-difference states. In particular, we propose mathematical interpretations within RL for concepts such as frequency, uncertainty, predictability and contradiction. Our framework also includes an element that allows visualizing sequences of relevant states rather than single moments. Further, the prior user studies focused mainly on the perception of agent proficiency and whether users should trust them based on their performance. In this paper we are primarily concerned with establishing connections between the different elements and the correct understanding of both the agent's capabilities and limitations. Importantly, our results provided insights on which elements are better at indicating when and how agents might need intervention to improve their performance.

3. Reinforcement learning

We now introduce the necessary notation and terminology used throughout the paper related to the computational framework of RL. RL agents can be modeled using the *Markov decision process* (MDP) framework [19]. Formally, we denote an MDP as a tuple $\mathcal{M} = (\mathcal{S}, \mathcal{A}, P, R, \gamma)$, where: \mathcal{S} is the set of all possible environment states; \mathcal{A} is the action repertoire of the agent; $P(s' | s, a)$ is the probability of the agent visiting state s' after having executed action a in state s ; $R(s, a)$ represents the reward that the agent receives for performing action a in state s ; $\gamma \in [0, 1]$ is some discount factor denoting the importance of future rewards.

Table 1

Overview of the dimensions of analysis evaluated in the user study and the generated interestingness elements.

Dimension	Generated interestingness elements
Frequency	(in)frequent situations: <i>situations the agent finds very (un)common</i>
Execution certainty	(un)certain executions: <i>hard/easy to predict situations with regards to action execution</i>
Transition-value	minima and maxima: <i>favorable and adverse situations</i>
Sequence	most likely sequences to maxima: <i>capture the agent's learned strategy</i>

An MDP evolves as follows. At each time step $t = 0, 1, 2, \dots$, the environment is in some state $S_t = s$. The agent selects some action $A_t = a \in \mathcal{A}$ and the environment transitions to state $S_{t+1} = s'$ with probability $P(s' | s, a)$. The agent receives a reward $R(s, a) = r \in \mathbb{R}$, and the process repeats.

The goal of the agent can be formalized as that of gathering as much reward as possible throughout its lifespan discounted by γ . This corresponds to maximizing the value $v = \mathbb{E} [\sum_t \gamma^t r]$, where the reward r dictates the immediate utility of taking action a in state s at time-step t , in light of the underlying task that the agent must solve. In order to maximize v , the agent must learn a policy, denoted by $\pi : \mathcal{S} \rightarrow \mathcal{A}$, that maps each state $s \in \mathcal{S}$ directly to an action $\pi(s) \in \mathcal{A}$. In the case of MDPs, this corresponds to learning a policy $\pi^* : \mathcal{S} \rightarrow \mathcal{A}$ referred to as the *optimal policy* maximizing the value v .

Typical approaches to RL learn a function $Q^* : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ associated with π^* that verifies the recursive relation:

$$Q^*(s, a) = r + \gamma \sum_{s' \in \mathcal{S}} P(s' | s, a) \max_{b \in \mathcal{A}} Q^*(s', b). \quad (1)$$

$Q^*(s, a)$ represents the *value* of executing action a in state s and henceforth following the optimal policy. RL algorithms like Q -learning [20] assume that the agent has no knowledge of either P or R . Hence, they typically start by *exploring* the environment—selecting actions in some exploratory manner—collecting samples in the form (s, a, r, s') which are then used to successively approximate Q^* using Eq. (1). After exploring its environment, the agent can *exploit* its knowledge and select the actions that maximize (its estimate of) Q^* .

4. Introspection framework

In this section we describe our introspection framework, including the generated interestingness elements.

4.1. Interaction data

As shown in Fig. 1, the introspection framework relies on the following data collected by the agent during its interaction with the environment:

- $n(s)$: the number of times the agent visited s ; $n(s, a)$: the number of times it executed action a after observing s ; and $n(s, a, s')$: the number of times it visited s' after executing action a when observing s ;
- $\hat{P}(s' | s, a)$: the *estimated probability* of observing s' when executing action a after observing s . This can be modeled from the interaction according to $\hat{P}(s' | s, a) = n(s, a, s')/n(s, a)$;
- $Q(s, a)$: the agent's *estimate of the Q function*, corresponding to the expected value of executing a having visited s and henceforth following the current policy. This can be estimated using any value-based RL algorithm;
- $V(s)$: the agent's *estimate of the V function* that indicates the value of observing s and henceforth following the current policy. This corresponds to $V(s) = \max_{a \in \mathcal{A}} Q(s, a)$.

Some of this data is already collected by standard RL methods, namely the Q and V functions, and, in the case of model-based approaches, the \hat{P} and \hat{R} functions. The remainder can easily be collected by the agent during its interaction with the environment by updating counters and running averages.

As shown in Fig. 1, this data is used in various introspective analyses designed to extract statistical information from the interaction that might help explain the agent's behavior. We note that it is up to the end-user of the system to determine what constitutes the *history of interaction* during which the data is collected. Namely, if the interaction data is collected *during training*, it can retain information about the interaction beyond the learned policy, thereby potentially capturing interesting challenges encountered by the agent while learning. In contrast, capturing data only *after training* might miss some subtleties of the training phase but will be more representative of the agent's learned policy. Another possibility is to capture data during both training and testing to capture a set of interesting situations without losing aspects of the agent's optimal behavior.

Table 1 lists the analyses considered in this paper and a short description of the elements that each generates. Next, we present the purpose behind each dimension of analysis and discuss how each can be used to help explain the agent's interaction with an environment, providing a mathematical interpretation for the extraction of the interestingness elements given available interaction data.

4.2. Frequency analysis

This dimension of analysis can be used to expose the agent's (in)experience with its environment, denoting both commonly and rarely encountered situations. It extracts two interestingness elements: **frequent situations** and **infrequent situations**.

Infrequent elements may indicate states not sufficiently explored by the agent. E.g., imagine an agent learning to run through a maze. If during learning the agent finds locations that are hard to reach but are not relevant to solving the maze, after training these would likely be identified as infrequent. Another example is a situation that had such a negative impact on the agent's performance, e.g., a death situation in a game, that its action-selection mechanisms made sure they were rarely visited in subsequent visits to nearby states. This element can also denote novel or rare situations, which could indicate a highly dynamic environment. In an interactive setting, these elements provide opportunities for an agent to ask a user for guidance, e.g., in choosing the appropriate actions in infrequently encountered situations.

In contrast, frequent elements can be used to provide examples of the agent's typical behavior, which can help to understand its learned strategy, or how the agent faces situations that are recurrent, as occurs in many video-game scenarios where the player has available multiple "lives".

Elements' extraction. Frequent and infrequent situations correspond to states that were visited less or more frequently than others during interaction, respectively. We can extract these elements according to whether the frequency in the tables $n(s)$, $n(s, a)$ and $n(s, a, s')$ is below/above a given threshold, or by selecting the top/bottom k states.

4.3. Execution certainty analysis

This dimension can be used to reveal the agent's confidence in its decisions. Two complementary elements are extracted: **certain executions** and **uncertain executions**.

Situations where the agent is uncertain of what to do indicate opportunities to ask a user for help, or indicate parts of the state space in which the agent requires more training. For example, in a combat scenario where the agent controls a group of units, the distribution over actions might be more even if the agent encounters a larger opposing force than encountered during training, meaning that the agent is not sure of what action to select. Uncertain situations are also significant because people tend to require explanations mostly for unusual behavior [21].

In contrast, certain situations denote parts of the task that the agent has learned well. These can also be used to reveal to a user the agent's behavior in situations where it is more predictable.

Elements' extraction. This dimension estimates the (un)certainly of each state with respect to action execution. Observations where many different actions are executed often (high dispersion) are considered uncertain, while those in which only a few actions are selected are considered certain. Given a state s , the execution certainty associated with s is measured by the concentration of the executions of the actions $a \in \mathcal{A}$. We use the *evenness* of the distribution over next states as given by its normalized true diversity (information entropy) [22]. Formally, let $p(X)$ be a probability distribution over $x_i \in X, i = 1, \dots, N$ of set X . The evenness of p over X is provided by:

$$\xi(X) = - \sum_{x_i \in X^+} p(x_i) \ln p(x_i) / \ln N, \quad (2)$$

where $X^+ \doteq \forall x_i \in X : p(x_i) > 0$. We then use this evenness measure to calculate the dispersion of the distribution over actions according to $\xi_s = \xi(\pi(s))$, where π is any policy of interest.

In our study we approximate the agent's interaction policy using $\hat{\pi}(s) = n(s, \mathcal{A})/n(s)$. We argue that this formulation retains information about the agent's history of interaction beyond the learned "optimal" policy by, e.g., capturing situations that were harder to learn. However, one could use the policy directly derived by the algorithm, in the case of policy-based methods. Similarly, we could compute the evenness based on the learned (normalized) Q -values instead.

4.4. Transition-value analysis

This dimension aims at enabling users to understand the desirability attributed by the agent to a given situation. The goal is to analyze how the value attributed to some state changes with respect to possible states visited at the next time-step. The elements extracted by this analysis are **minima** and **maxima** situations.

Maxima may denote acquired preferences or *favorable* situations for the agent. An example of such a situation can occur when the agent "wins" a game or moves to a higher level, in which case it has achieved the maximal value situation. Another situation can occur in more dynamic settings. E.g., in a game like Asteroids, an agent destroying an asteroid might represent a (local) maximum but, inevitably, the agent will encounter new asteroids to be destroyed, and will no longer be in a maximum state. Still, in such situations it achieves a locally-maximal value and its policy will likely lead the agent to visit similar states in the future. Therefore, presenting maxima elements might allow a user to observe the agent performing

well in the task according to what it has learned.⁴ This lets users verify whether performance is in line with expectations, or that adjustment of the agent's training, reward function, perceptual characteristics, etc. is needed.

In contrast, minima denote highly *adverse* situations that the agent wants to avoid and where taking any action leading to a different state is preferable. They can also denote undesirable situations visited by the agent not because of its behavior but due to the stochastic nature of the environment. By observing the agent's behavior in such situations, a user can understand how well the agent handles difficult situations, and whether further adjustments are required.

Elements' extraction. This analysis combines information from the agent's estimated $V(s)$ function and the transition function $\hat{P}(s' | s, a)$. These refer to states whose values are lower/greater than or equal to the values of all possible next states, as informed by the agent's learned predictive model. Formally, let $\mathcal{T}_s \doteq \{s' \in \mathcal{S} : \exists a \in \mathcal{A} \hat{P}(s' | s, a) > 0\}$ be the set of observed transitions starting from state s . The *local minima* are defined by $\mathcal{S}_{min} \doteq \forall s \in \mathcal{S} : \forall s' \in \mathcal{T}_s V(s) \leq V(s')$. Similarly, the *local maxima* are defined by $\mathcal{S}_{max} \doteq \forall s \in \mathcal{S} : \forall s' \in \mathcal{T}_s V(s) \geq V(s')$.

4.5. Sequence analysis

The goal of this analysis is to calculate interesting sequences of actions learned by the agent. In particular, these sequences of agent behavior start and end with important states as identified by the other analyses. Each sequence is associated with a certain likelihood of occurring. In the experiments reported in this paper, we specifically selected sequences starting from a (local) minimum, and then performing actions until a (local) maximum is observed, as they provide good indicators of an agent's aptitude in the task. We thus refer to this element as **most likely sequences to maxima**. Alternatively, the initial and final state sets can include the (in)frequent, (un)certain, or other states of interest discovered by the different analyses. Each sequence would then have a different semantic meaning attached to it, revealing the agent's behavior when transitioning to distinct situations.

This element can be used to reveal more purposeful agent behavior to a human user, e.g., performing actions from a low-valued situation to a learned maximum. This avoids requiring the user to view complete episodes while providing potentially more information for understanding the agent's behavior compared to viewing performance in the context of a single element. The extracted sequences can also be used by the user to query the agent about its future goals and behavior in *any* possible situation. Namely, given the flexibility of the sequence-finding procedure, the initial and final sets can be user-defined. This would allow providing the user with contrastive explanations about why the alternatives to some actions—the foils—are not as desirable as those chosen by the agent [21]. The starting points may denote the *origins* of behavior while the sequence likelihood and the target states denote the agent's *beliefs* and *reasons*, respectively—two crucial elements people use to explain intentional events [23].

Elements' extraction. We first create a state-transition graph where nodes are states $s \in \mathcal{S}$ and edges are actions $a \in \mathcal{A}$ denoting the observed transitions, weighted according to the probability $\hat{P}(s' | s, a)$. The resulting graph is directed with nonnegative weights, making it amenable to best-first search algorithms to find the most likely sequence of state-action pairs between two given states.

We then use a variant of [Dijkstra's algorithm](#) [24], whose input is an initial state $s_i \in \mathcal{S}_i$, where \mathcal{S}_i is a state set of interest, and a set of possible final states, \mathcal{S}_f . First, we determine the most likely paths between s_i and each target state $s_f \in \mathcal{S}_f$. Let $P_{if} = [s_0 = s_i, a_1, s_1, \dots, a_n, s_n = s_f]$ denote a path between s_i and s_f . The probability of the agent visiting s_f after s_i and following path P_{if} is thus given by $p(s_i, s_f) = \sum_{t=1}^{|P_{if}|} \hat{\pi}(s_{t-1}, a_t) \hat{P}(s_t | s_{t-1}, a_t)$, where the probability of selecting actions, given by $\hat{\pi}$, can be estimated as described in Sec. 4.3 for the extraction of execution certainty. After calculating the paths' probabilities, we then choose the *most likely path weighted by value*, connecting the initial state and an optimal final state, according to: $s_i^* = \operatorname{argmax}_{s_f \in \mathcal{S}_f} p(s_i, s_f) V(s_f)$.

5. Visual summaries

As depicted in Fig. 1, provided that the behavior of the agent can be readily visualized, the interestingness elements extracted by our analyses can be used to generate *visual summaries*. In particular, we follow an approach similar to [5], where short video-clips highlighting the behavior of the agent during key moments of interaction are recorded. This type of behavior summarization is therefore suitable for agents relying on a visual input such as robots or game-playing agents. Nonetheless, as noted in [5] the extraction of summaries will depend on the characteristics of the specific task domain.

The rationale behind this approach is that simply watching the behavior of the agent performing the task, e.g., during whole episodes, might be burdensome for a user and mostly uninformative. Therefore, the system is given a *budget* for a given summary. The budget will typically be application-dependent and can be defined in terms of, e.g., the maximum length of the video clip or the maximum number of time-steps or highlights per summary. We generate two types of video highlights in our framework: summaries and sequences.

⁴ Although our experimental study informed us that maxima and minima often occur when the agent performs the task well/poorly, additional experiments will be required to further characterize *all* situations in which they might occur.

5.1. Summary highlights

These highlights combine different examples of the agent's behavior. Like the approach in [5], given some history of interaction with an environment, e.g., a sequence of simulated episodes, we summarize agent behavior by selecting k highlights (agent trajectories) of length l time-steps, where $(l - 1)/2$ time-steps occur before and after the situation of interest. The highlights are then combined into a single video clip for visualization.

Departing from [5], we introduce novel ways in which the key moments are selected and presented to a user. First, as mentioned earlier, each interestingness element is defined according to some analysis dimension. For example, frequency elements can be sorted from the most to the least frequent states. During highlight selection, we try to sample situations closer to the extrema—e.g., when highlighting frequent states, a trajectory passing through the most frequent state is the most preferred; in the case of the sequences dimension, we select the largest sequence. The idea is to select the most representative and informative examples regarding each source of interestingness.

For summaries where $k > 1$, we use a secondary selection criterion of *diversity*, which is similar to the HIGHLIGHTS-DIV approach in [5], where here we use a different example selection procedure. Namely, whenever the maximum number of k highlights for a type of element has been reached, the framework decides whether to replace any of the existing highlights with the new state. Formally, let $d(s_1, s_2) \in [0, 1]$ be an state diversity/distance metric relevant to the task, e.g., how visually distinct the states are, or how distant in time the states were visited. Then, given the $k + 1$ highlights, we keep the set of k highlights maximizing:

$$\max_{s_i, s_j} d(s_i, s_j) \times \min_{s_i, s_j} d(s_i, s_j), \quad i \neq j, \quad i, j = 0, \dots, k - 1 \quad (3)$$

The idea is to maximize the total diversity of the final set of highlights such that the user can observe, e.g., different nuances of the agent's learned behavior, or contextually distinct situations. Therefore, Eq. (3) aims at including maximally-separated examples by maximizing both the maximal and the minimal distances between pairs of examples to be included in a summary.

Finally, we introduce the step of deciding how to visualize and compile the highlights. When experimenting with visualization, we noticed that, depending on the task, the user sometimes fails to understand which part of the agent's behavior is being highlighted. This can be critical in highly dynamic environments, where the behavior of the agent can vary dramatically at each time-step. To mitigate this problem, we introduce fade-in/fade-out effects around each highlight so that the user can focus on the important moment being highlighted while also being presented with the context in which it occurred.

5.2. Sequence highlights

Our framework provides an additional type of highlighting capability based on the element of likely sequences to favorable situations, which enables highlighting relevant sequences of behavior rather than single moments. In particular, in our user study we focus on sequences starting in a *minimum* and ending in a *maximum*, as informed by the transition-value analysis. This reveals examples of the agent's learned strategy, i.e., a trajectory showcasing its capabilities in overcoming perceived difficulties and in reaching its learned preferences.

6. Experimental study

To evaluate the usefulness of our XRL framework in helping humans correctly understand the aptitude of different RL agents, we conducted a user study using a simple video-game scenario based on Frogger.⁵ Fig. 2a shows a screenshot of the game. The player is responsible for controlling frogs, one at a time, to go from the grassy strip at the bottom to the lilypads at the top, with only one frog at a time allowed on a lilypad. The player has to first cross the road without getting hit by cars and then cross the river by jumping on logs. The control set is $\mathcal{A} = \{N, S, E, W\}$, where each action deterministically moves the frog by 40 pixels in the corresponding direction.

When a frog reaches a lilypad, a new frog appears at the bottom of the screen. When two frogs reach the lilypads, the agent completes the current level and starts the next level. This clears the lilypads and makes cars and logs go faster. A player has a limit of 100 moves to pass a level and 3 lives in a game. The player loses a life when the maximum number of moves is reached, a frog is hit by a car, a frog jumps into the river, or a frog is on a log that goes off-screen. At the beginning of each game, the time interval at which logs are introduced in the environment is randomized.

6.1. Agent experiments

To make the Frogger game an interesting task from the perspective of behavior explanation, we designed agents with *partial observability*. Specifically, each state $s \in S$ provides a local view of the environment according to the frog's location.

⁵ We used the implementation in <https://github.com/pedrodb/frogger>.

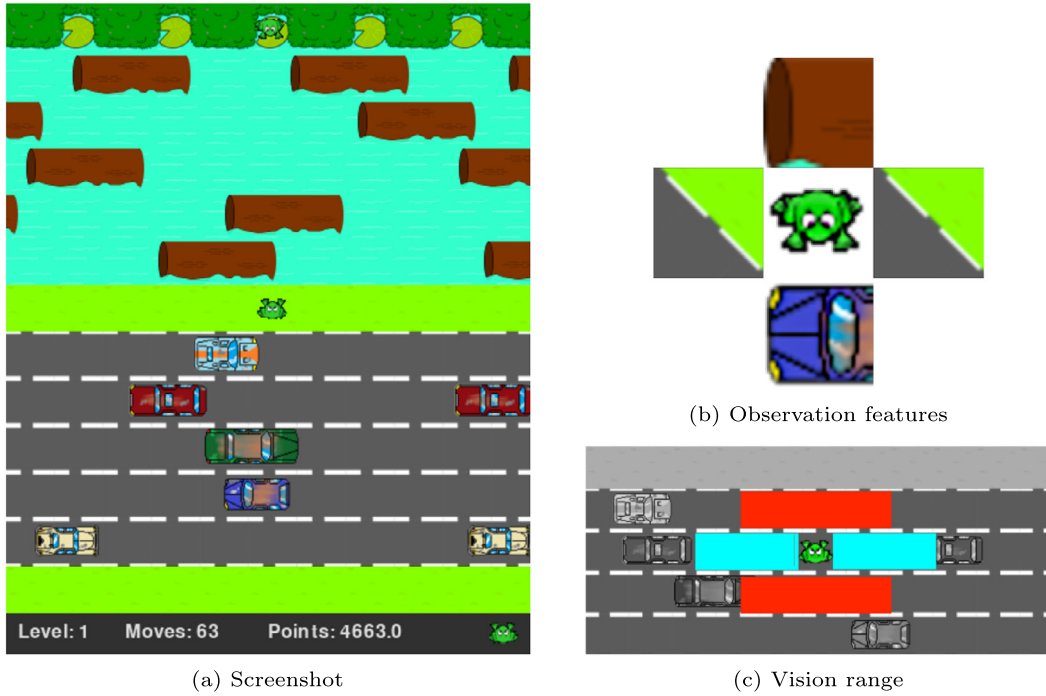


Fig. 2. The Frogger game task (left) and agent design aspects (right). (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

States are represented by four discrete features, denoted by $s = [\phi_N, \phi_S, \phi_E, \phi_W]$, each $\phi_i \in \Phi$ indicating the element that is visible in the corresponding direction, where $\Phi = \{\text{empty}, \text{water}, \text{car}, \text{log}, \text{lilypad}, \text{bounds}\}$ is the feature set, *bounds* means that the agent is near the environment's borders, and *empty* that no element is visible in that direction.⁶

This partial observability introduces uncertainty in the environment's dynamics *from the agent's perspective*. Fig. 2b shows what an agent observes given the true game state in Fig. 2a. As we can see, $\phi_S = \text{car}$ although the car is not directly beneath the frog. This has to do with the way we designed the perception of cars, where two parameters control for an agent's vision range, as illustrated in Fig. 2c. The blue squares indicate the agent's *horizontal vision range* vis_h , defined as the horizontal distance (in pixels) within which an agent can see cars in its *E* and *W* directions, i.e., a car in the same lane. Red squares indicate the *vertical vision range* vis_v , defined by the vertical distance (in pixels) within which an agent can see cars to its *N* and *S*, i.e., in the lanes *directly* above and below the agent, respectively. As for the logs, we designed the *log* feature such that an agent anticipates the movement of adjacent logs.⁷

The reward function R was defined as follows. The agent receives a reward $r(s, a) = -200$ in situations where executing a when observing s results in one of the aforementioned death conditions. A punishment of -300 occurs when the agent depletes its available lives and a reward of 5,000 is received whenever a frog arrives on a lilypad. A reward of -1 is also received at each time-step.

6.1.1. Agent types

To study the impact of the different interestingness elements in helping understand the aptitude of RL agents, we designed three different agents:

Optimized: this agent can observe cars at a moderate distance, i.e., when they are neither too close nor too far away. Its parameters were empirically fine-tuned to achieve a very high performance in the task. The agent uses the default reward function.

High-vision: this agent has very high vision capability, i.e., it can anticipate the presence of cars from afar. It was designed to simulate an agent with inadequate perceptual capabilities, i.e., whose sensing mechanisms were poorly calibrated to the task. It also uses the default rewards.

⁶ The agent does not distinguish between the road and grass for feature *empty*, nor the type of car for the *car* feature.

⁷ The agent's observations are extracted without access to the game engine's internal dynamics, including the logs' velocities, hence they are not perfect. This means there is still a chance ($p > 0$) that performing action *N* when $\phi_N = \text{log}$ will result in a death in the river.

Table 2

The parameterization of each agent in our experiments. Legend: r_{river} : reward received after a death in the river; q_{init} : Q-values initialization. See text for more details.

Agent	vis_h	vis_v	r_{river}	q_{init}
Optimized	60	40	−200	5,000
High-vision	140	120	−200	5,000
Fear-water	60	40	−10,000	0

Table 3

The performance of each agent in 2,000 test episodes. Legend: River: agent jumped into river; Car: agent was hit by a car; Time: maximum moves to pass a level (100) reached.

Agent	Mean level	Number of deaths				Mean steps
		River	Car	Time	Total	
Optimized	4.6 ± 0.7	2,395	1,051	14	3,460	288 ± 26
High-vision	2.0 ± 0.4	411	888	3,027	4,326	291 ± 22
Fear-water	1.0 ± 0.0	0	3,604	2,396	6,000	222 ± 46

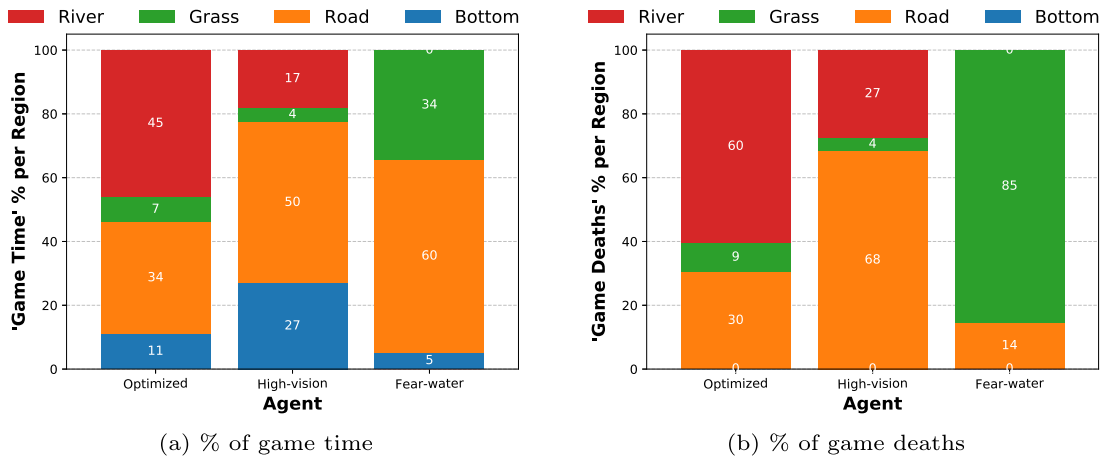


Fig. 3. The percentage of time spent and deaths occurring in each environment region over 2,000 test episodes (Grass: middle grassy strip; Bottom: grassy strip on the bottom).

Fear-water: this agent has the same visual capabilities as the optimized agent. It also uses the default reward function with one exception: when it dies on the river, it receives a punishment of $r = -10,000$. This simulates a motivational impairment—the rewards idealized by its designer are not appropriate for the agent to learn the intended task.

A detailed parameterization of each agent is listed in Table 2.

6.1.2. Agent RL training

Each agent was first trained using standard Q-learning [20] for 2,000 episodes,⁸ with each episode corresponding to one game, i.e., ending when the agent lost all of its 3 lives or when a limit of 300 time-steps was reached. We used a Softmax action-selection mechanism with exponentially-decaying temperature β , i.e., given training episode e , the temperature is given by $\beta = \beta_{\min} + \beta_{\max} \cdot 0.995^e$. In our experiments, we set $\beta_{\min} = 0.05$ and $\beta_{\max} = 20$. To promote exploration, we used optimistic initialization by setting $Q(s, a) = q_{init}$ for all states $s \in \mathcal{S}$ and actions $a \in \mathcal{A}$, where q_{init} controls the agent's optimism. As listed in Table 2, unlike other agents, the Fear-water agent was initialized with $q_{init} = 0$ to further prevent it from reaching lilypads as a result of pure exploration. After training, the learned policy was tested for another 2,000 episodes by setting $\beta = \beta_{\min}$, resulting in greedy action-selection.⁹

6.1.3. Agent performance results

Summary statistics of the agents' performance during the 2,000 test episodes after learning are listed in Table 3. Fig. 3 depicts the agents' performance relative to each region of the environment. Our goal was to design the agents such that, af-

⁸ Source code of the interestingness elements framework, the agents and algorithms used in the experimental study can be found at <https://github.com/SRI-AIC/InterestingnessXRL>.

⁹ The performance of each agent during learning is detailed and discussed in Appendix A.B.

Table 4

Composition and budget of the summarization techniques used in the user study.

Name	Summary composition	k
Max	all maxima	4
Min	all minima	4
Cert	all certain execution states	4
Uncert	all uncertain execution states	4
Freq	all frequent states	4
Infreq	all infrequent states	4
Max-Min	half minima, half maxima	4
Cert-Uncert	half certain-execution, half uncertain	4
Freq-Infreq	half frequent, half infrequent	4
All	one of each type	6
Seq	sequence from minimum to maximum	1

ter training, their underlying characteristics and limitations would lead to contrasting behaviors. As we can see, the different characteristics indeed resulted in the agents attaining different performances in the task.

The optimized agent achieves the best performance among the three agents, reaching significantly higher levels. Compared to the high-vision agent, we see that this is due to the differences in the vision range parameters. The high-vision agent can see cars at a higher distance, resulting in more cautious behavior while crossing the road. While this leads the agent to spend more time on the road, it gets hit by cars significantly less than other agents.

For the fear-water agent, receiving a highly negative reward from dying in the river (two times the magnitude of the lilypad reward) prevents it from reaching any lilypad. The agent was also not optimistically initialized, making it less motivated to explore and inadvertently discover the lilypads. We see in Fig. 3 that the agent usually dies by getting hit by a car from the middle grass row. This is because the value of jumping into the river (action N) is significantly lower than that of any other action, i.e., the agent “fears” falling into the water more than anything else.¹⁰

6.2. Introspection and visual summaries

Referring back to Fig. 1, an agent’s history of interaction in our experiments comprises a total of 4,000 episodes (2,000 training and 2,000 testing). During that time, all the interaction data described in Sec. 4.1 was collected. As mentioned earlier, the goal was to capture both the challenges encountered by the agents during learning and aspects of their learned policy. We then applied our introspection framework to extract the interestingness elements in Table 1 for each agent.¹¹

After introspective analysis, visual summaries were produced by generating video clips highlighting the agents’ performance according to the procedures described in Sec. 5. The summaries were captured during the 2,000 test episodes.¹² For the diversity metric $d(s_1, s_2)$ over highlights, we used the absolute difference of game score between the moments in which s_1 and s_2 were observed, the goal being to capture the agents’ behavior at various stages of the game.

6.3. User study

To assess whether the visual examples derived from the different interestingness elements would help human users correctly understand the different agents’ underlying capabilities and limitations, we conducted a user study.¹³ We recruited participants using the psiTurk tool [25] for Amazon Mechanical Turk (MTurk). A total of 82 participants (40% female, age group mode [25 – 34] years) participated in the study, each receiving \$7.50 for their completion of the Human Intelligence Task (HIT).

6.3.1. Summarization techniques

Our objective was to understand the differential contribution of each interestingness element to understanding an agent’s behavior. We wanted to know whether single elements are sufficient for correctly perceiving an agent’s aptitude or if a combination works better to convey certain behavioral characteristics.

Table 4 lists the *summarization techniques* used in our study. The first six techniques generated summaries composed of highlights for a single type of interestingness element. The goal was to assess how each type exposed different characteristics of the agents’ behavior. The second group of techniques generated videos composed of trajectories highlighting the extrema of each dimension of analysis, enabling us to evaluate the explanatory power of each dimension. The All technique,

¹⁰ We note that this agent is still able to cross the road and avoid cars but could not jump to logs when reaching the river. This agent serves as a baseline for assessing the usefulness of the interestingness elements in identifying behavioral impairments.

¹¹ Appendix A.B provides statistics on the interestingness elements captured for each agent.

¹² While introspection was performed over all episodes, summaries were extracted only from the test episodes. This was both to avoid capturing mistakes due to the training process itself (exploration) and not because of the agents’ inherent limitations, and to show how the agent learned to deal with potentially rare or uncertain situations that occurred during training.

¹³ The study was approved by SRI International’s Institutional Review Board (IRB).

given a higher budget, generated highlights including 6 examples with the goal of assessing the effect of including one highlight per element. Finally, the Seq technique produces a single sequence highlight to determine whether a single trajectory is sufficient to reveal subtleties underlying the agents' behavior.

We used a trajectory length of $l = 21$ time-steps (≈ 5 secs.) for each highlight in a summary and a maximum length of 80 time-steps for sequence highlights. To avoid biases during the observation of highlights by human users, we removed all game information at the bottom of the screen from the videos, i.e., the level, remaining moves, points and available lives.¹⁴

6.3.2. Conditions

Each condition of our study corresponds to a *scenario* where we assess the subjects' perception of all three agents' characteristics given one summarization technique. Since we have 11 techniques and three agents, we have two independent variables: *agent*, with 3 levels, and *scenario*, with 11. Informed by a pilot study, we determined that requiring subjects to watch videos for all players and all summarization techniques (a full within-subjects design), would be too onerous. On the other hand, getting accustomed to the task requires time and effort, which made exposing each subject to only one condition (between-subjects design) equally undesirable. Therefore, we opted for a partially-balanced, incomplete block experimental design [26], where each subject was exposed to 6 randomly-selected scenarios, and in each scenario subjects were exposed to the behavior of all three agents.

6.3.3. Experimental procedure

After accepting the task, subjects started by agreeing to a consent form explaining the goals and protocol of the study. They were then redirected to a page describing the dynamics and rules of Frogger, after which they were presented with a 5-question quiz about the game.¹⁵ Failed questions were highlighted in red to reinforce the main concepts, but all participants were allowed to continue. The primary goal was to ensure that subjects understood the dynamics underlying the task (e.g., that reaching a new level speeds up cars and logs) rather than to prevent them from participating in the study.

An instruction page was then presented that explained what the participants had to do in each of the 6 scenarios. As seen in Fig. 4, each scenario showed videos of the three agents in parallel. The relative location of the video for each agent in the page (left, middle, right) was randomly assigned. As seen on the top of Fig. 4, participants were told that they were watching videos of the gameplay of 3 different *players*. Further, the instructions stated that the players were different for each scenario and that the order of appearance was not relevant. The reason for this is two-fold: first, we wanted subjects to believe they were watching the behavior of human players, and second, we wanted to avoid learning effects between scenarios.

Subjects could start each video when they pleased and after playing all videos (enforced programmatically), a questionnaire with the survey questions appeared below the videos, as depicted in Fig. 4. We designed a set of questions for each scenario to assess subjects' perception of the agents' underlying characteristics and also properties of their learned behavior, given the scenario's summarization technique.

We knew from the previous performance results (see Fig. 3) that the agents' capabilities in each region are quite different. Each region in Frogger provides different challenges and requires distinct learned skills to successfully overcome them. As such, we created two *region-related* questions:

Time: "Select the regions where you think each player normally spends more time playing the game (select at least one)".

Practice: "Select the regions where you think each player needs more practice (select none if you think the player doesn't need more practice)".

These questions assess the subjects' ability to perceive the difficulty agents had in different regions of the game. *Time* assesses relative presence while *practice* evaluates difficulties experienced in each region. For each question subjects could select from {river, middle grass row, road, bottom grass row}. The outcome of each question is thus a Boolean variable for each region.

We also designed two questions about the overall *aptitude* of the agents:

Level: "The player is capable of reaching advanced levels".

Help: "The player needs help to be better at the game".

Subjects were asked to rate their agreement to each statement on a 5-point Likert scale (1 strongly disagree – 5 strongly agree). The goal was to evaluate the subjects' perception of the overall capabilities of each agent. In particular, *level* estimates the perception of how well an agent plays the game, while *help* evaluates whether subjects could correctly determine if an agent needed intervention to improve its performance. Subjects were instructed to answer questions with respect to what they inferred about the players given what they had seen in the video clips. The goal was for subjects to evaluate the players' aptitude beyond their performance in the videos.¹⁶ A final question assessed the subjects' *confidence* in their responses on

¹⁴ The videos used in the study are at <https://github.com/pedrodbs/frogger-study>.

¹⁵ Screenshots of the several pages and transcriptions of all the instructions, questions and options are included in Appendix A.C.

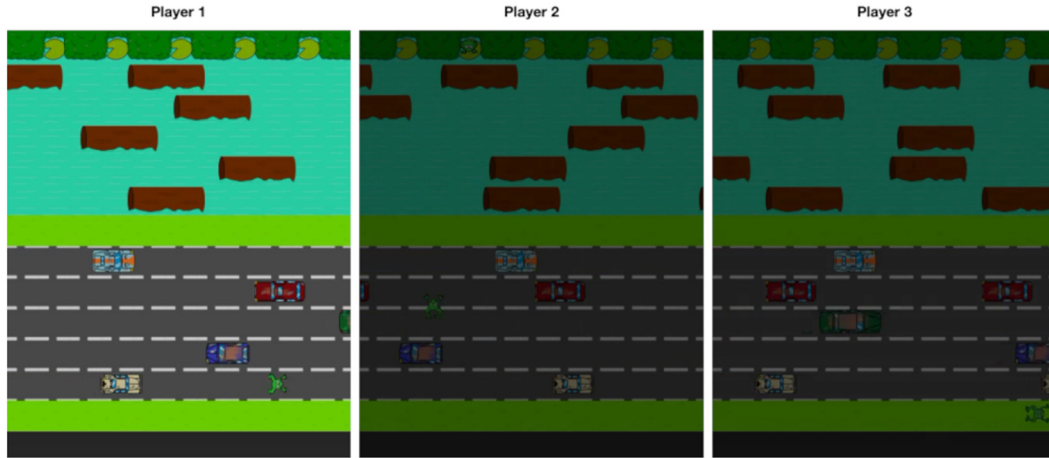
¹⁶ Subjects were allowed to play, stop and rewind videos as many times as they wanted.

Part 1 out of 6.

[Show Instructions Reminder](#)

Below are video clips showing game-play highlights of 3 different players.

Click the play button to (re)start each video. Once you have watched all three videos, some questions will appear below.



Select the regions where you think each player normally spends more time playing the game (select at least one):

	Player 1	Player 2	Player 3
On the river (jumping on logs)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
On the grass (between the road and the river)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
On the road (avoiding cars)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
On the bottom grass row	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Select the regions where you think each player needs more practice (select none if you think the player doesn't need more practice):

	Player 1	Player 2	Player 3
On the river (jumping on logs)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Fig. 4. Questionnaire used in the user study.

a 5-point Likert scale (1 not confident – 5 very confident). The outcomes of these questions are thus ordinal variables with values in [1, 5].

After answering the questions for all 6 scenarios, a final page collected subjects' demographics and general opinions (open question) about the study. Subjects then returned to MTurk to collect their compensation.

7. Analysis and results

We are mainly interested in answering these two research questions:

RQ1: Does the information conveyed by the various summarization techniques induce a different perception of each agent's aptitude?

RQ2: What summarization techniques enable a correct perception of the agents' aptitude in the task?

To begin, we removed the responses of subjects who did not seem to take the task seriously. On average, participants correctly answered 4.2 ± 1.0 out of 5 quiz questions and took an average of 19.5 ± 11.9 min. to complete the survey. We eliminated 5 subjects who had ≤ 2 correct answers or spent less than 6 min. on the survey. The data for the remaining 77 subjects resulted in each scenario being sampled 42 ± 4 times (min: 36, max: 50).

Table 5

Significant differences in the responses of the region-related questions (V : Cramér's V statistic; Scenarios: scenarios with the most significant pairwise differences. Empty cells denote non-significant differences where $p \geq 0.01$).

Region	Agent	Time		Practice	
		V	Scenarios	V	Scenarios
River	Optimized	0.52	Min	0.30	Infreq
	High-vision	0.65	Min, Uncert, Freq		
	Fear-water				
Grass	Optimized			0.26	
	High-vision			0.28	Infreq
	Fear-water	0.41	Min, Uncert		
Road	Optimized	0.50	Min	0.46	Seq, Cert
	High-vision	0.42	Uncert	0.49	Infreq
	Fear-water	0.36	Max	0.35	Min
Bottom	Optimized	0.19			
	High-vision	0.69	Min	0.27	Infreq
	Fear-water	0.42	Uncert		

Table 6

Pairing of response variables and agent performance data (ground-truth).

Response variable	Agent performance variable
Time	Mean percentage of game time spent in the region.
Practice	Mean number of deaths occurred in the region (excludes lives lost due to maximum number of moves achieved).

7.1. Analysis of perceived agent performance in different regions

7.1.1. Research question 1

To assess whether the different summarization techniques had an effect on subjects' perception of agent behavior in the different regions, we performed a χ^2 (chi-squared) test over the sum of positive responses and calculated the effect size using Cramér's V with the bias correction in [27]. A Bonferroni correction was then applied to test for significant pairwise differences.

Table 5 shows the significant effects ($p < 0.01$) of the summarization technique on the subjects' responses found for each agent.¹⁷ We also show which techniques had the most pairwise significant differences, as informed by the Bonferroni comparison—this indicates what interestingness elements and introspection dimensions exposed the most distinct characteristics of the agents. For the *time* response, significant effects were found for all agents in the road and bottom grass regions, whereas for the *practice* variable significant effects were found for all agents only in the road region.

7.1.2. Research question 2

To identify which summarization techniques allowed for a correct assessment of the agents' aptitude in each region, we first paired each response variable with a game performance variable, as listed in Table 6. The performance data serves as ground truth for the agents' behavior and characteristics against which we compare the subjects' interpretations. Then, for each response variable, we transformed the vector containing the percentage of subjects' responses for each agent and summarization technique across all regions into a Boltzmann distribution.¹⁸ Similarly, we derived a Boltzmann distribution from the mean values of the corresponding performance variable, resulting in the expected relative performance of a given agent in each region.¹⁹ To determine whether the subjects' responses diverged from the agents' actual performance (ground-truth) we used the Jensen-Shannon divergence (JSD) [28], which is essentially a distance metric with scores ranging from 0 (identical distributions) to 1 (maximally different).

Table 7 presents the results of this analysis. We highlight the summarization techniques that resulted in either a very correct (low JSD) or very incorrect (high JSD) interpretation of the agents' performance relative to each region.²⁰

¹⁷ Appendix A.D contains plots depicting subjects' responses for the region-related questions.

¹⁸ E.g., if all subjects selected all regions in the *practice* question for an agent in a given scenario, this means they believed that the agent needed to practice equally in each region, resulting in a probability of 1/4 for each element of the response distribution.

¹⁹ E.g., for the *practice* response variable, this corresponds to the probability of an agent losing a life in each region.

²⁰ We note that the JSD thresholds used in our analysis, marked with the red and blue colors, are used only for the purpose of isolating the relevant results of our study.

Table 7

JSD between response and performance variables for the region-related questions, for each scenario and agent (blue: low divergence, red: high divergence).

Scenario	Time			Practice		
	Opt.	High	Fear	Opt.	High	Fear
Max	0.09	0.70	0.60	0.01	0.47	0.54
Min	0.45	0.60	0.03	0.35	0.08	0.80
Cert	0.18	0.70	0.15	0.02	0.63	0.51
Uncert	0.37	0.04	0.18	0.29	0.18	0.61
Freq	0.03	0.55	0.07	0.02	0.16	0.60
Infreq	0.05	0.77	0.57	0.00	0.75	0.73
Max-Min	0.11	0.33	0.35	0.47	0.03	0.66
Cert-Uncert	0.18	0.14	0.13	0.35	0.23	0.74
Freq-Infreq	0.03	0.30	0.50	0.01	0.04	0.44
All	0.22	0.11	0.18	0.36	0.15	0.61
Seq	0.01	0.37	0.27	0.01	0.32	0.62

Table 8

Significant differences ($p < 0.01$) in the responses of the aptitude-related questions (ϵ^2 : epsilon-squared effect; Scenarios: scenarios with most significant pairwise differences).

Response var.	Agent	ϵ^2	Scenarios
Level	Optimized	0.32	Min
	High-vision	0.39	Min, Uncert
	Fear-water	0.07	Uncert
Help	Optimized	0.33	Min, Uncert
	High-vision	0.40	Min, Freq
	Fear-water	0.07	Uncert

Table 9

Pairing of response variables and agent performance data (ground-truth).

Response variable	Agent performance variable
Level	Mean game level achieved.
Help	Mean number of lives lost per game.

7.2. Analysis of perceived overall agent aptitude

7.2.1. Research question 1

To assess whether the summarization techniques resulted in different perceptions of the agents' aptitudes, we modeled the response for each agent as an ordinary least squares (OLS) regression model and performed a Kruskal-Wallis H-test. Effect sizes were calculated using the ϵ^2 (epsilon-squared) statistic [29], followed by a Bonferroni correction post-hoc pairwise comparison.

Table 8 shows a significant effect of the summarization technique (scenario) on the *level* and *help* responses for all agent types ($p < 0.01$). The results for effect size also show that the videos depicting the behavior of optimized and high-vision agents had a greater impact on the differences induced by the summarization techniques. The Bonferroni comparison shows that the Min and Uncert elements had the greatest impact in inducing different responses from the subjects.

7.2.2. Research question 2

To investigate the correctness of the subjects' assessment of the agents' capabilities and limitations we again paired each response variable with a corresponding agent performance variable, as listed in Table 9. We linearly rescaled (min-max scale) the response variables for all scenarios and the corresponding performance variables so that they fell in the [0, 1] interval, and then modeled them as Gaussian distributions.²¹

Then, for each agent, we calculated the Hellinger distance (H) [30] between the distributions of the response and performance variables for each scenario to obtain a divergence measure with scores ranging from 0 (matching distributions), to 1 (diverging distributions). Table 10 highlights the largest and smallest Hellinger distances between the response and game performance variables for each agent and summarization technique.²²

²¹ For example, for the *level* performance variable, this models the probability of an agent reaching high levels.

²² The thresholds for the H distance were chosen empirically for purposes of analysis, i.e., to highlight some results and compare the influence of the different summarization techniques.

Table 10

H distance between response and performance variables for the aptitude-related questions, for each scenario and agent (blue: low divergence, red: high divergence).

Scenario	Level			Help		
	Opt.	High	Fear	Opt.	High	Fear
Max	0.27	0.95	0.99	0.31	0.57	0.99
Min	0.56	0.48	0.99	0.40	0.30	0.99
Cert	0.45	0.91	0.99	0.41	0.50	0.99
Uncert	0.47	0.52	0.99	0.37	0.20	0.98
Freq	0.32	0.53	0.99	0.26	0.28	0.99
Infreq	0.39	0.96	0.99	0.11	0.63	0.99
Max-Min	0.31	0.73	0.99	0.14	0.27	0.99
Cert-Uncert	0.35	0.71	0.99	0.04	0.20	0.99
Freq-Infreq	0.31	0.71	0.99	0.04	0.26	0.99
All	0.38	0.75	0.99	0.09	0.32	0.99
Seq	0.40	0.91	0.99	0.36	0.55	0.99

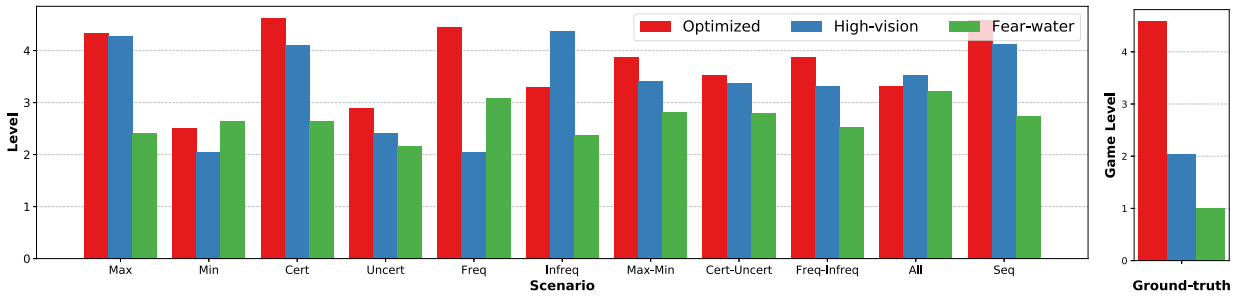
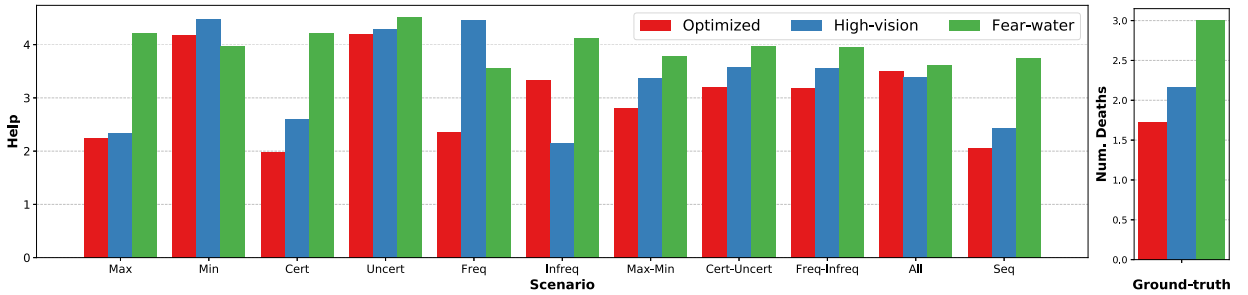
(a) *level* response vs. game level(b) *help* response vs. number of game deaths

Fig. 5. Comparison of subjects' mean responses for each agent and summarization technique (Scenario, left) and the agents' game performance across the 2,000 test episodes (Ground-truth, right). We note the different scales between responses and ground-truth data.

7.2.3. Assessing relative performance

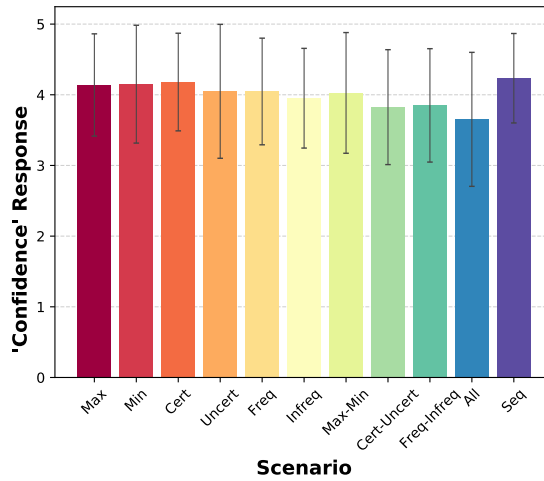
The results in the previous section indicate, for some scenarios, a lack of agreement between subjects' responses and the agents' game performance. However, this might be due to our experimental design, where each scenario presented highlights for the three agents simultaneously. Thus, we also analyzed which summarization techniques resulted in a correct perception of the *relative performance* between the agents.

Fig. 5 illustrates the idea behind our approach. By comparing the shape of the individual scenario plots to the ground truth plot, we see that some techniques appear to support a correct perception of the agents' relative abilities, while others do not. To quantify this difference, we compared the distribution of subjects' responses across the three agents with the distribution of their actual performance. Specifically, for each scenario, we derived Boltzmann probability distributions for the agents from the mean values of the responses and the corresponding performance variables. We then computed the JSD between each response distribution and the ground-truth performance distribution. The results are listed in Table 11, where the *diverging agents* column shows the agents responsible for the distributions being highly divergent.

Table 11

JSD between the relative response and performance for the aptitude-related questions for each scenario (blue: low divergence, red: high divergence). For high JSD we indicate the agents mostly responsible for the subjects' incorrect perceptions.

Scenario	Level		Help	
	JSD	Diverging agents	JSD	Diverging agents
Max	0.28		0.03	
Min	0.51	Fear	0.53	High, Fear
Cert	0.09		0.02	
Uncert	0.04		0.21	
Freq	0.00		0.73	High, Fear
Infreq	0.92	Opt., High	0.03	
Max-Min	0.10		0.04	
Cert-Uncert	0.23		0.06	
Freq-Infreq	0.06		0.06	
All	0.52	High	0.24	
Seq	0.12		0.02	

**Fig. 6.** The subjects' mean confidence in their responses in each scenario.

7.3. Confidence analysis

Finally, we wanted to determine whether some summarization techniques influenced the subjects' confidence response. A Kruskal-Wallis analysis of variance did not reveal a significant effect of the summarization technique (scenario) on the subjects' confidence response ($p = 0.047$). Fig. 6 shows the mean responses per scenario. We see that, in general, the subjects' confidence in their answers was not greatly affected by exposure to particular scenarios. A Bonferroni correction post-hoc comparison revealed that the largest difference was found between Seq (most confident) and All (least confident) ($p = 7.5 \times 10^{-4}$, $\epsilon^2 = 0.34$).

7.4. Discussion

In this section we discuss in depth the results of our experimental study and provide the main insights stemming from using each summarization technique.

7.4.1. Importance of "negative" moments

As detailed in Section 2, most previous works within XRL identify key moments of interaction resulting in examples of behavior where the agent excels in the task—this has been shown to be useful e.g., to assess relative performance when comparing between different agents. In contrast, our dimensions of analyses result in the presentation of situations where the agent is also less confident, thereby providing opportunities for intervention, e.g., for the agent designer to improve the agent's perceptions, improve learning, etc.

The results of our user study show the importance of showing agent behavior in such situations. Namely, Tables 5 and 8 show that the most significant pairwise differences were produced in the Min and Uncert scenarios, denoting the importance of exposing "hard" situations for assessing the agents' relative presence in different regions and their overall aptitude in the task. Further, Table 10 shows that for the high-vision agent, who could succeed in the task despite perceptual limitations,

exposing subjects to these situations allowed for a more accurate assessment of the agent's capabilities in overcoming such flaws.

7.4.2. Importance of diversity in behavior summaries

Despite the importance of showing situations where the agents are less confident, alone they do not allow for a correct understanding of the agents' aptitude in the task. In fact, a main conclusion of the study is that no single technique can induce a correct understanding of the underlying capabilities and limitations of *all* agents in *all* situations. In particular, exposing only "good" or "bad" moments can lead users to an incorrect perception of aptitude.

For example, Tables 10 and 11 show that providing highlights of only favorable situations (Max and Cert) does not lead to a completely correct perception of the agents' capabilities. Meanwhile, showing only minima and infrequent moments can lead users to incorrectly perceive relative agent limitations and to believe that all agents fail equally. Two cases worth noting here are the optimized agent, whose performance appears to have been underestimated by subjects in the Min and Uncert scenarios, and the fear-water agent, whose limitations were perceived to be very similar to that of the other two agents, which we know to be incorrect given these agents' true game performance.

Primarily, the problem is that by providing examples only of where an agent excels or where it fails, users may not be exposed to the parts of the task that required more of the agent's attention (time spent), nor where and how the agent needs to be improved. In contrast, our results show that techniques providing examples captured at the two extrema of the dimensions of analysis can provide a more accurate view of the agents' competence in the task. For example, Table 7 shows that the aptitude of the high-vision agent was better depicted by the moments captured in the Max-Min and Freq-Infreq scenarios while Table 10 showed similar results for the optimized agent in the Freq-Infreq and Cert-Uncert scenarios. Overall, these techniques provide a balance of "good" and "bad" moments of performance, which can allow for a better understanding of agents' capabilities and limitations in the task.

7.4.3. Importance of the frequency dimension

One dimension of analysis providing good overall results was frequency. In particular, the results of the region-related responses in Table 5 show that the frequency dimension, by providing a balance between common and rare situations, provides a good indication of the improvements needed in each subtask. Further, Table 7 shows a low divergence when comparing subjects' responses with the agents' true performance in the Freq-Infreq scenario, thus emphasizing the importance of the frequency dimension for correctly assessing where an agent spends most of its time and where it needs more (or less) practice.

On one hand, the results show the importance of frequent elements for correctly determining where an agent spends most of its time. On the other, infrequent situations are also relevant in that they can help identify when the agent fails. Moreover, frequency seems to be useful for showing an agent's overall behavior characteristics, while the transition-value dimension and other more complex elements relying on the learned value function help reveal its difficulties and capabilities in hard, challenging situations.

7.4.4. Importance of the sequence dimension

Another technique worth mentioning is that of most likely sequences to maxima—i.e., capturing agent trajectories that go from a difficult situation (local minimum) to a favorable situation (local maximum). An advantage of this technique is that an observer gets to see the "normal" behavior of the agent in the task, which is related to its learned strategy. The results in Table 11 show that even with a single trajectory, subjects could correctly perceive the unique characteristics of the agents' behavior and infer their underlying relative aptitude. Although the Seq scenario did not get the lowest divergence scores overall, Fig. 6 shows that subjects were more confident in their responses. Hence, sequences may enable a reasonably accurate assessment of how the agent performs in the task, particularly with a limited explanation budget (time).

7.4.5. Richness of an agent's behavior

A related insight is that the summarization techniques, and the associated interestingness elements, are only as good at capturing different characteristics of RL agents as the richness of the agents' underlying capabilities. For example, from Tables 7 and 10 we see that the behavior exhibited by the fear-water agent due to its "motivational impairments" is so limited that no single interestingness element captured qualitatively distinct behaviors. Notably, the summaries gave subjects the impression that the agent needs practice in all the regions equally when in fact it has the most difficulty in jumping onto logs when reaching the river. In contrast, the agents exhibiting more diverse behavior enabled subjects to differentiate the agents' capabilities under different conditions, as evidenced by their responses (see e.g., Table 8, where a significant effect of the summarization technique was found).

In general, agents that have predictable, monotonous behavior may not generate the necessary numerical nuances targeted by our introspection framework. For example, if the values associated with the states are all high, one cannot distinguish between minima and maxima. Likewise, if the observed behavior is highly deterministic, then uncertain elements cannot be determined.

7.4.6. Problems with combining all elements

Our results also captured an interesting result regarding the All scenario, where summaries contain one highlight for each interestingness element. One might think this technique would be optimal, given that it highlights all the different aspects

of the agent's interaction. However, in our study it led to the lowest confidence in responses (Fig. 6) and poor perceptions overall, e.g., Tables 7, 10 and 11 show that this technique resulted in overall high divergence between subjects' responses and true agent performance. We believe that the diversity of moments highlighted by each element may have confused the subjects, making it difficult for them to determine an agent's overall aptitude. Also, half the highlights show infrequent situations or where agents are uncertain of what to do, leading subjects to underestimate their performance.

8. Future work and conclusions

8.1. Future work

A problem with visual summaries based on highlights is that, by definition, they do not capture all aspects of an agent's interaction with the environment. Because of that, users may extrapolate the behavior of an agent beyond what they see, potentially leading them to incorrectly perceive its overall performance. Further, this can be exacerbated when user expectations are not met. For example, in our study subjects may have formed beliefs about how "normal" players behave in the Frogger game prior to seeing the videos, which may have led to incorrect perceptions of aptitude in some scenarios. Also, by design, our experiment involved visualizing agents simultaneously, which could have led to the creation of false expectations based on the best and worst performance observed among all agents.

One possible solution would be to augment highlights with more quantitative measures about the agents' overall performance in the task. This may help users form a better assessment of the agents' capabilities and limitations, and of when and how they should intervene. For example, observing a high-performing agent failing to perform a particular subtask can direct a user to correct specific aspects of the agent's design. Such information could also help users understand that not all agent limitations are the same and that different agents may require distinct adjustments.

Another solution to mitigate problems with unmet expectations would be to provide users with short descriptions about what the highlights are conveying about an agent's behavior. This would be particularly useful in our framework given the variety of interestingness elements and the qualitatively distinct aspects of an agent's experience that each captures. This way, users could better understand agents' limitations and set more realistic expectations about their behavior.

Our introspection framework was designed to be used throughout the different stages of an agent's lifecycle and for different explanation modes. For example, it can be used *during learning*, to track the agent's learning progress and acquired preferences or *after learning*, to summarize the most relevant aspects of the interaction. It can be used *passively*, to let a user ask the agent about its current goals or to justify its behavior in a given situation or *proactively*, to let the agent consult the user in situations where its decision-making is more uncertain or unpredictable. We are currently developing a platform where the competency of RL agents, including competency in learning, can be assessed for different tasks and in novel situations.

We are also currently extending the introspection framework to support analyzing high-dimensional domains that require the use of deep-RL techniques by defining interpretations of the concepts behind the different interestingness elements that can leverage deep network representations. Some elements allow for a direct translation to deep-RL architectures, e.g., estimating execution certainty from the output distribution of a policy network or transition value from situations where executing different actions leads to strictly increasing or decreasing value. Others may require more specialized structures, e.g., the transition certainty of a situation can be estimated through generative models [e.g., 31,32] that predict the likelihood of next states, while a model for characterizing the familiarity of inputs [e.g., 33] could be used to determine the frequency of observations.

8.2. Summary and main conclusions

In this paper we introduced a novel framework for introspective XRL that extracts different *interestingness elements*, each capturing a particular aspect of the agent's experience with an environment. We showed how to generate visual summaries from these elements in the form of video-clips highlighting diverse moments of the agent's experience while using its learned policy. To assess the usefulness of the elements in helping humans understand RL agents' competency, we performed a user study using the game of Frogger. First, we created agents with different perceptual capabilities and motivations in the task to simulate potential problems in the agents' design. We trained the agents using RL, resulting in different game behaviors and levels of performance. We then performed an online survey where we presented subjects with videos of the agents' behavior generated using different summarization techniques, each using a distinct combination of the interestingness elements.

Our results show that no single summarization technique—and hence no single interestingness element—provides a complete understanding of every agent in all possible situations of a task. Ultimately, a combination of elements enables the most accurate understanding of an agent's aptitude in a task. Different combinations may be required for agents with distinct capabilities and levels of performance. The results highlighted in particular the element of frequency, which provides a good understanding of an agent's overall behavior characteristics. As for the transition-value dimension, it can denote difficulties and capabilities of agents in hard, challenging situations. In addition, the element producing the most likely sequences to desirable situations can provide a short but accurate understanding of an agent's learned strategy. Another

insight stemming from the results is that combining many types of elements in a single video summary can confound users.

A few general notes about the proposed framework are worth mentioning. Our framework for explainable reinforcement learning extracts a diverse set of interestingness elements capturing significant moments of agent behavior. The introspection framework, comprising the dimensions of analysis that extract the interestingness elements, is *domain-independent*: the data that is collected and analyzed is agnostic to the specific learning scenario. In addition, it is *algorithm-independent* in that it can be used in conjunction with standard RL methods without having to modify the learning mechanism itself.

We note also that the dimensions of analysis proposed here provide mathematical interpretations of different *concepts* for analyzing an RL agent's interaction with the environment. Namely, each concept provides different insights over the agent's behavior that are useful for judging its performance and to understand how it can be improved. We do not claim that the list presented in this paper is complete or definitive—other aspects of an RL agent's interaction can be defined, and different interpretations, possibly applicable to different RL algorithms, can be proposed for the same concepts. Similarly, if one RL algorithm does not learn the interaction data required by some element, then only a subset of the interestingness dimensions will be available for analysis.

Moreover, the analyses in our framework do not make direct assertions about potential problems with the agent's training or with the resulting policy. Neither do they prescribe any specific intervention by a human user to improve the performance of the agent in the task. Rather, they help indicate important aspects of the agent's behavior. The purpose of the user study was precisely to help make the connection between the different elements and the perception of the agent's characteristics by human observers.

Regarding state-space dimensionality, we chose to develop our initial framework for simpler tasks because, given the inherent complexity of sequential decision-making, it is important to first determine how to generate the critical *elements* for explainable RL. Thus, the concepts behind our elements—and their effects in enabling users to understand the aptitude of agents—are more important than the particular evaluation domain or the RL architecture used to train the agents. Furthermore, we note that the subjects of our study were unaware that the players were in fact autonomous agents, let alone that RL was used to derive their behavior. As such, our evaluation of the usefulness of the explanation concepts is agnostic to the underlying representations and RL algorithm used.

The results of this paper show that the *diversity* of aspects captured by the proposed interestingness elements can help humans correctly identify an RL agent's aptitude in a task. Further, not only do they help reveal the underlying characteristics of agents, but as a consequence, we believe they can be good indicators of which adjustments might be required to improve their performance.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Supplementary material

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.artint.2020.103367>.

References

- [1] R.S. Sutton, A.G. Barto, Reinforcement Learning: An Introduction, MIT Press, 1998.
- [2] M.L. Littman, Memoryless policies: theoretical limitations and practical results, in: Proceedings of the Third International Conference on Simulation of Adaptive Behavior: From Animals to Animats 3, SAB94, MIT Press, Cambridge, MA, USA, 1994, pp. 238–245.
- [3] R.J. Bayardo Jr., R. Agrawal, Mining the most interesting rules, in: Proceedings of the Fifth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '99, ACM, New York, NY, USA, 1999, pp. 145–154.
- [4] P. Sequeira, E. Yeh, M. Gervasio, Interestingness elements for explainable reinforcement learning through introspection, in: Joint Proceedings of the ACM IUI 2019 Workshops, ACM, 2019, p. 7.
- [5] D. Amir, O. Amir, Highlighting agent behavior to people, in: Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS '18, IFAAMAS, Richland, SC, 2018, pp. 1168–1176.
- [6] O.Z. Khan, P. Poupart, J.P. Black, Minimal sufficient explanations for factored Markov decision processes, in: Proceedings of the 19th International Conference on International Conference on Automated Planning and Scheduling, ICAPS'09, AAAI Press, 2009, pp. 194–200.
- [7] F. Elizalde, E. Sucar, J. Noguez, A. Reyes, Generating explanations based on Markov decision processes, in: MICAI 2009: Advances in Artificial Intelligence, Springer, Berlin, Heidelberg, 2009, pp. 51–62.
- [8] N. Wang, D.V. Pynadath, S.G. Hill, The impact of pomdp-generated explanations on trust and performance in human-robot teams, in: Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems, AAMAS '16, IFAAMAS, Richland, SC, 2016, pp. 997–1005.
- [9] T. Zahavy, N. Ben-Zrihem, S. Mannor, Graying the black box: understanding dqns, in: Proceedings of the 33rd International Conference on Machine Learning, in: Proceedings of Machine Learning Research, vol. 48, PMLR, New York, New York, USA, 2016, pp. 1899–1908.
- [10] S. Greydanus, A. Koul, J. Dodge, A. Fern, Visualizing and understanding Atari agents, in: Proceedings of the 35th International Conference on Machine Learning, in: Proceedings of Machine Learning Research, vol. 80, PMLR, Stockholmsmässan, Stockholm, Sweden, 2018, pp. 1792–1801.
- [11] A. Koul, A. Fern, S. Greydanus, Learning finite state representations of recurrent policy networks, in: International Conference on Learning Representations, ICLR 2019, 2019, <https://openreview.net/forum?id=S1gOpsCctm>.
- [12] B. Hayes, J.A. Shah, Improving robot controller transparency through autonomous policy explanation, in: Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction, HRI '17, ACM, New York, NY, USA, 2017, pp. 303–312.

- [13] J. van der Waa, J. van Diggelen, K. van den Bosch, M. Neerincx, Contrastive explanations for reinforcement learning in terms of expected consequences, in: *Proceedings of the 2nd Workshop on Explainable Artificial Intelligence, XAI 2018*, 2018, pp. 165–170.
- [14] P. Madumal, T. Miller, L. Sonenberg, F. Vetere, Explainable reinforcement learning through a causal lens, *arXiv:1905.10958*, 2019.
- [15] S.H. Huang, D. Held, P. Abbeel, A.D. Dragan, Enabling robots to communicate their objectives, *Auton. Robots* 43 (2019) 309–326.
- [16] I. Lage, D. Lifschitz, F. Doshi-Velez, O. Amir, Exploring computational user models for agent policy summarization, in: *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, International Joint Conferences on Artificial Intelligence Organization, California*, 2019, pp. 1401–1407.
- [17] O. Amir, F. Doshi-Velez, D. Sarne, Summarizing agent strategies, *Auton. Agents Multi-Agent Syst.* 33 (2019) 628–644.
- [18] S.H. Huang, K. Bhatia, P. Abbeel, A.D. Dragan, Establishing appropriate trust via critical states, in: *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018, pp. 3929–3936.
- [19] M.L. Puterman, *Markov Decision Processes: Discrete Stochastic Dynamic Programming*, 1st ed., John Wiley & Sons, Inc., New York, NY, USA, 1994.
- [20] C.J.C.H. Watkins, P. Dayan, Q-learning, *Mach. Learn.* 8 (1992) 279–292.
- [21] T. Miller, Explanation in artificial intelligence: insights from the social sciences, *Artif. Intell.* 267 (2019) 1–38.
- [22] C.P.H. Mulder, E. Bazeley-White, P.G. Dimitrakopoulos, A. Hector, M. Scherer-Lorenzen, B. Schmid, Species evenness and productivity in experimental plant communities, *Oikos* 107 (2004) 50–63.
- [23] M.M.A. De Graaf, B.F. Malle, How people explain action (and autonomous intelligent systems should too), Technical Report FS-17-01, in: *AAAI 2017 Fall Symposium on Artificial Intelligence for Human-Robot Interaction*, 2017.
- [24] E.W. Dijkstra, A note on two problems in connexion with graphs, *Numer. Math.* 1 (1959) 269–271.
- [25] J. McDonnell, J. Martin, D. Markant, A. Coenen, A. Rich, T. Gureckis, *psiTurk (Version 1.02) [Software]*, <https://github.com/NYUCCL/psiTurk>, 2012.
- [26] R.C. Bose, W.H. Clatworthy, S.S. Shrikhande, *Tables of Partially Balanced Designs with Two Associate Classes*, Technical Report, North Carolina State University. Dept. of Statistics, 1953.
- [27] W. Bergsma, A bias-correction for Cramér's V and Tschuprow's T, *J. Korean Stat. Soc.* 42 (2013) 323–328.
- [28] B. Fuglede, F. Topsøe, Jensen-Shannon divergence and Hilbert space embedding, in: *Proceedings of the International Symposium on Information Theory, 2004, SIT 2004*, 2004, pp. 31–37.
- [29] M. Tomczak, E. Tomczak, The need to report effect size estimates revisited. An overview of some recommended measures of effect size, *Trends Sport Sci.* 21 (2014) 19–25.
- [30] E. Hellinger, Neue Begründung der Theorie quadratischer Formen von unendlichvielen Veränderlichen, *J. Reine Angew. Math.* 136 (1909) 210–271.
- [31] D. Ha, J. Schmidhuber, World models, *arXiv preprint, arXiv:1803.10122*, 2018.
- [32] D. Hafner, T. Lillicrap, I. Fischer, R. Villegas, D. Ha, H. Lee, J. Davidson, Learning latent dynamics for planning from pixels, in: K. Chaudhuri, R. Salakhutdinov (Eds.), *Proceedings of the 36th International Conference on Machine Learning, PMLR, Long Beach, California, USA*, in: *Proceedings of Machine Learning Research*, vol. 97, 2019, pp. 2555–2565, <http://proceedings.mlr.press/v97/hafner19a.html>.
- [33] S. Akcay, A. Atapour-Abarghouei, T.P. Breckon, Ganomaly: semi-supervised anomaly detection via adversarial training, in: C.V. Jawahar, H. Li, G. Mori, K. Schindler (Eds.), *Computer Vision – ACCV 2018, Springer International Publishing, Cham*, 2019, pp. 622–637.