

Local and global explanations of agent behavior: integrating strategy summaries with saliency maps


Tobias Huber, Katharina Weitz, Elisabeth André, Ofra Amir

Angaben zur Veröffentlichung / Publication details:

Huber, Tobias, Katharina Weitz, Elisabeth André, and Ofra Amir. 2021. "Local and global explanations of agent behavior: integrating strategy summaries with saliency maps." *Artificial Intelligence* 301: 103571. <https://doi.org/10.1016/j.artint.2021.103571>.

Nutzungsbedingungen / Terms of use:

CC BY 4.0



Local and global explanations of agent behavior: Integrating strategy summaries with saliency maps [☆]

Tobias Huber ^{a,*}, Katharina Weitz ^a, Elisabeth André ^a, Ofra Amir ^b

^a Universität Augsburg, Universitätsstraße 6a, Augsburg, Germany

^b Technion - Israel Institute of Technology, Israel

ARTICLE INFO

Article history:

Received 20 April 2020

Received in revised form 14 June 2021

Accepted 26 July 2021

Available online 3 August 2021

Keywords:

Explainable AI

Strategy summarization

Saliency maps

Reinforcement learning

Deep learning

ABSTRACT

With advances in reinforcement learning (RL), agents are now being developed in high-stakes application domains such as healthcare and transportation. Explaining the behavior of these agents is challenging, as the environments in which they act have large state spaces, and their decision-making can be affected by delayed rewards, making it difficult to analyze their behavior. To address this problem, several approaches have been developed. Some approaches attempt to convey the *global* behavior of the agent, describing the actions it takes in different states. Other approaches devised *local* explanations which provide information regarding the agent's decision-making in a particular state. In this paper, we combine global and local explanation methods, and evaluate their joint and separate contributions, providing (to the best of our knowledge) the first user study of combined local and global explanations for RL agents. Specifically, we augment strategy summaries that extract important trajectories of states from simulations of the agent with saliency maps which show what information the agent attends to. Our results show that the choice of what states to include in the summary (global information) strongly affects people's understanding of agents: participants shown summaries that included important states significantly outperformed participants who were presented with agent behavior in a set of world-states that are likely to appear during gameplay. We find mixed results with respect to augmenting demonstrations with saliency maps (local information), as the addition of saliency maps, in the form of raw heat maps, did not significantly improve performance in most cases. However, we do find some evidence that saliency maps can help users better understand what information the agent relies on during its decision-making, suggesting avenues for future work that can further improve explanations of RL agents.

© 2021 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

The maturing of artificial intelligence (AI) methods has led to the introduction of intelligent systems in areas such as healthcare and transportation [67]. Since these systems are used by people in such high-stakes domains, it is crucial for users to be able to understand and anticipate their behavior. For instance, a driver of an autonomous vehicle will need to anticipate situations in which the car fails and hands over control to her, while a clinician will need to understand the treatment regime recommended by an agent to determine whether it aligns with the patient's preferences.

[☆] This paper is part of the Special Issue on Explainable AI.

* Corresponding author.

E-mail address: tobias.huber@uni-a.de (T. Huber).

The recognition of the importance of human understanding of agents' behavior, together with the complexity of current AI systems, has led to a growing interest in developing "explainable AI" methods [21,28,2]. The idea of making AI systems explainable is itself not new and was already discussed since the early days of expert systems [69,18]. However, state-of-the-art AI algorithms use more complex representations and algorithms (e.g., deep neural networks), making them harder to interpret. For example, in contrast to classical planning approaches such as the belief-desire-intention (BDI) framework [56] in which the goals of the agent are explicitly defined, current agents often use policies trained using complex reward functions and feature representations that are difficult for people to understand.

In this paper, we focus on the problem of describing and explaining the behavior of agents operating in sequential decision-making settings, which are trained in a deep reinforcement learning framework. In particular, we explore the usefulness of *global* and *local* post-hoc explanations [51] of agent behavior. Global explanations describe the overall policy of the agent, that is, the actions it takes in different regions of the state space. An example of such global explanations are strategy summaries [7], which show demonstrations of the agent's behavior in a carefully selected set of world states. Local explanations, in contrast, aim to explain specific decisions made by the agent. For instance, saliency maps are used to show users what information the agent is attending to [27].

We explore the combination of global and local information describing agent policies. The motivation for integrating the two approaches is their complementary nature: while local explanations can help users understand what information the agent attends to in specific situations, they do not provide any information about its behavior in different contexts. This is reinforced by a previous study conducted by Alqaraawi et al. [4] who evaluated local explanations and came to the conclusion that sole instance-level explanations are not sufficient and should be augmented with global information. Similarly, while demonstrating what actions the agent takes in a wide range of scenarios can provide users with a sense of the overall strategy of the agent, it does not provide any explanations as to what information the agent considered when choosing how to act in a certain situation.

To examine the benefits of these two complementary approaches and their relative usefulness, we integrate strategy summaries with saliency maps. Specifically, we adapt the HIGHLIGHTS-DIV algorithm for generating strategy summaries from our previous work [5] such that it can be applied to deep learning settings, and integrate it with saliency maps that are generated based on Layer-Wise Relevance Propagation (LRP) (using a method we previously published in [35]). We combine these two approaches by adding saliency maps, which show what information the agent attends to, to the summary generated by HIGHLIGHTS-DIV.

We evaluate this combination of global and local explanations in a user study in which we explore both the benefits of HIGHLIGHTS-DIV summaries and the benefits of adding saliency maps to strategy summaries. Specifically, we compare likelihood-based summaries (which select states for the summaries based on the likelihood of visiting them) and HIGHLIGHTS-DIV summaries, both with and without the addition of saliency maps. Study participants completed two types of tasks requiring the analysis of different agents trained to play the game of Pacman: an agent comparison task in which they compare the performance of two agents, and a retrospection task, in which they reflect on an agent's strategy. We chose those tasks to investigate whether the users trusted the right agent and to evaluate their mental models of the agents, respectively.

Our results show that participants who were shown HIGHLIGHTS-DIV summaries performed better on both tasks compared to participants who were shown likelihood-based summaries, and were also more satisfied with HIGHLIGHTS-DIV summaries. We find mixed results with respect to the benefits of adding saliency maps to summaries, which improved participants' ability to identify some aspects of agents' strategies, but in most cases did not lead to improved performance.

The paper makes the following contributions:

- It demonstrates that the HIGHLIGHTS-DIV algorithm, which was so far only used on classic reinforcement learning, can be applied to deep reinforcement learning agents with slight adjustments.
- It proposes a joint local and global explanation approach for RL agents by integrating LRP saliency maps and HIGHLIGHTS-DIV summaries.
- It evaluates the combination of global and local summaries in a user study, demonstrating the benefits of HIGHLIGHTS-DIV summaries and the potential benefits and limitations of local explanations based on saliency maps.

The remainder of this article is structured as follows: Section 2 reviews prior work on explainable intelligent agents, Sections 3 and 4 describe our previous works on local and global explanations, respectively. Section 5 details our combined implementation of those two methods, including the adaptation of HIGHLIGHTS-DIV to deep reinforcement learning. We describe the empirical evaluation we conducted in Section 6, and its results are summarized in Section 7. Finally, we discuss the results of the study and future directions in Section 8, and conclude in Section 9.

2. Related work

In this section, we review related works on explainable AI. We begin with a short review of global and local explanation methods of machine learning models, elaborating on the use of saliency maps, which we also make use of. We then discuss in more depth prior works on global and local explanations of policies of agents operating in sequential decision-making settings such as RL agents.

Global and local methods for interpretable machine learning Broadly, our work relates to the problem of interpretable machine learning, that is, explanations for the decisions of prediction models [21]. Few interpretable machine learning approaches provide global explanations, e.g., by showing examples of a set of instances and specifying how they were classified [57,39] or by generating prototypical images that maximize the activation of specific neurons [64].

The majority of methods focus on local explanations that explain single decisions of the model. To this end, various methods to measure the relevance of a part of the input for the model's decision have been proposed. For visual input, this information is often displayed as saliency maps that highlight how relevant each pixel is for a particular decision of the agent. Since the input for the Atari agents we use in this study is visual, we will use the word saliency map method even if the very same algorithm can be used on non-visual input data.

Gradient-based saliency map generation methods [64,66,68,62] utilize the derivative with respect to the input to estimate how much a small change in this input's value would change the prediction.

Occlusion-based methods ([79,57,60]), occlude areas inside the input and measure how much this changes the model's prediction. The idea behind this is to introduce uncertainty to the occluded area and to see how much the model is influenced by the loss of information in that area. Occlusion-based methods often come with the advantage of being independent of the model's structure but with the drawback of not being as precise as some model-specific methods.

In contrast to the aforementioned methods for generating saliency maps, Bach et al. [12] proposed Layer-wise Relevance Propagation (LRP), which uses the intermediate activations of the neurons during the forward pass to estimate the contribution of each input pixels to prediction.

Common to all interpretable ML approaches is that they focus on one-shot decisions. Thus, they do not fully address the problem of explaining behavior in sequential decision-making settings, where the agent takes actions, earns rewards, and affects the state of the world.

Local explanations of agent behavior Several approaches have been introduced for explaining specific decisions in the context of Markov Decision Processes (MDP). Some works attempt to provide justifications for a policy [38,37,20] by making statements about particular action choices (e.g., an action was chosen because it will lead to a state that has a higher value with higher probability). Others provide causal explanations by integrating a causal structure of the domain [61,72]. Krarup et al. [40] propose methods for generating contrastive explanations to explain action choices.

In this paper, we focus on the use of saliency maps for local explanations. Several works have implemented saliency maps in the context of Deep Reinforcement Learning (DRL). Because many DRL algorithms utilize CNNs it is possible to directly use the methods we covered in the previous paragraph on those algorithms. Zahavy et al. [78] and Wang et al. [74] for example used gradient-based saliency maps on traditional and Dueling Deep Q-Network (DQN) algorithms.

Greydanus et al. [27] and Iyer et al. [36] propose novel occlusion-based algorithms, where Greydanus et al. use Gaussian blur instead of complete occlusion and Iyer et al. utilize template matching to identify objects in the input. This allows them to train a new agent on this additional information and then selectively occlude those objects.

Lapuschkin et al. [43] used LRP to visualize the classical DQN architecture. In this paper, we use a more selective LRP variant which we tested on RL agents in our previous work [35] (see Section 3).

Global explanations of agent behavior Several global explanation methods describing what actions an agent takes in different states have been proposed. Hayes et al. [30] developed a system that allows users to “debug” an agent's strategy by querying its decisions in situations specified by the user. In contrast to this approach, strategy summarization methods select a set of important states to share with the user, such that the user does not need to query the agent with respect to specific states. We note that the two approaches are complementary. Booth et al. [14] compile logical formulas that specify when certain behaviors occur, e.g., by stating for which regime in the state space an agent will perform a particular action. However, this approach requires a state representation that is understandable to the user, which may not be the case in many complex domains, especially when DRL is used.

Our work takes the approach of summarizing agent policies (which we refer to as “strategy summaries”) by demonstrating the behavior of an agent in a subset of world states which are considered important by the agent [7,6]. Several methods have been proposed for selecting the subset of demonstrations to present in a summary. Some methods choose states that best enable the reconstruction of the original policy, using computational models such as inverse reinforcement learning (inferring the agent's reward function) or imitation learning (constructing a mapping from states to agents' actions) [34,42]. An alternative approach uses heuristics for identifying “interesting” situations. The HIGHLIGHTS-DIV algorithm we utilize falls into this category, as it selects states based on the distribution of Q-values of different actions. We chose to use this approach since it does not make any assumptions about people's reasoning, is simpler computationally, and was shown to improve users' understanding of agent behavior. Similar approaches have been developed in parallel [33,63], varying in the specific formulation of the interestingness criteria used to determine which states to include in the summary.

Another recent line of work explored the problem of generating plans that are more understandable to people [41,17,16]. The idea underlying this approach is that by having a model of human plans in a domain, it is possible to generate plans that achieve the desired goal while being as consistent as possible with people's mental models. However, in contrast to the strategy summarization approach, these approaches have only considered goal-based plans for short-term tasks. Furthermore, they require a model of how people plan in the domain, which might not always be feasible to obtain.

Evaluation of explanation methods for RL agents Some recent user studies examined the use of saliency maps and strategy summaries to explain the behavior of RL agents to people.

Alqaraawi et al. [4] and Selvaraju et al. [62] found that participants, who saw saliency maps, were able to predict the decision of an image classification model better than participants who did not see them. However, the participants were still only correct in about 60% of the cases and Alqaraawi et al. proposed to look beyond instance-level explanations in the future. For actual RL agents, Iyer et al. [36] and Anderson et al. [9] also used an action prediction task to evaluate saliency maps but found no clear advantage of saliency maps. In addition to the prediction task, Anderson et al. used a retrospection task to get an even better understanding of participants' mental models and, in addition to saliency maps, investigated reward decomposition [22] and a combination of both methods. Here, they found significant positive effects for reward decomposition and the combined approach and a marginally significant ($p = 0.086$) effect in favor of saliency maps.

Strategy summaries have been evaluated using several different tasks. Huang et al. [34] and Lage et al. [42] asked participants to predict what actions an agent would take based on summaries optimized for policy reconstructions. Their results show that summary methods that better match with people's computational models lead to improved action prediction, but that people may use different models in different contexts. Summaries generated by a variety of interestingness criteria were shown to improve people's ability to identify regions of the state space in which an agent spends more time and regions of the state space in which an agent requires additional training [63]. Importance-based summaries (e.g. HIGHLIGHTS-DIV) were shown to improve people's ability to identify the better performing agent in an agent comparison task [5] and their ability to decide whether to trust an agent in specific world states [33].

In sum, this work extends the existing state-of-the-art in explanations of RL agents, by proposing an integrated global and local explanation method, which enhances HIGHLIGHTS-DIV summaries (global) with LRP saliency maps (local) and conducting a user study to examine the joint and separate contributions of the local and global information to people's understanding of the behavior of RL agents.

3. Saliency maps

In this section, we describe the local explanation method which we use in our combined local and global explanation approach. While the development of the local explanation method is not the focus of this paper, we include the details of the approach for completeness. We revisit the foundations of Layer-wise Relevance Propagation (LRP) and show how to use it on the DQN architecture. Then we describe our previously published *argmax*-rule, an adjustment to this algorithm, which generates more selective saliency maps and which we use in this work. In addition to some previously published illustrations of the selectivity of the *argmax*-rule, we implemented new sanity checks for our saliency maps and report their results.

3.1. Foundations

LRP does not describe a specific algorithm but a concept that can be applied to any classifier f that fulfills the following two requirements. First, f has to be decomposable into several layers of computation where each layer can be modeled as a vector of real-valued functions. Secondly, the first layer has to be the input x of the classifier containing, for example, the input pixels of an image, and the last layer has to be the real-valued prediction of the classifier $f(x)$. Any DRL agent fulfills those requirements if we only consider the output value that corresponds to the action we want to analyze.

For a given input x , the goal of any method following the LRP concept is to assign relevance values R_j^l to each computational unit j of each layer of computation l , in such a way that R_j^l measures the local contribution of the unit j to the prediction $f(x)$. A method of calculating those relevance values R_j^l is said to follow the LRP concept if it sets the relevance value of the output unit to be the prediction $f(x)$ and calculates all other relevance values by defining

$$R_j^l := \sum_{k \in \{j \text{ is input for neuron } k\}} R_{j \leftarrow k}^{l+1}, \quad (1)$$

for **messages** $R_{j \leftarrow k}^{l+1}$, such that

$$R_k^{l+1} = \sum_{j \in \{j \text{ is input for neuron } k\}} R_{j \leftarrow k}^{l+1}. \quad (2)$$

In this way a LRP variant is determined by choosing messages $R_{j \leftarrow k}^{l+1}$. Through Definition (1), it is then possible to calculate all relevance values R_j^l in a backward pass, starting from the prediction $f(x)$ and going towards the input layer. Furthermore, Equation (2) gives rise to

$$\sum_k R_k^{l+1} = \sum_k \sum_{j \in \{j \text{ is input for neuron } k\}} R_{j \leftarrow k}^{l+1}$$

$$= \sum_j \sum_{k \in \{j \text{ is input for neuron } k\}} R_{j \leftarrow k}^{l,l+1} = \sum_j R_j^l.$$

This ensures that the relevance values of each layer l are a linear decomposition of the prediction

$$f(x) = \dots = \sum_{j=1}^{\dim(l)} R_j^l = \dots = \sum_{j=1}^{\dim(input)} R_j^{input}.$$

Such a linear decomposition is easier to interpret than the original classifier because we can think of positive values R_j^l to contribute evidence in favor of the decision of the classifier and of negative relevance values to contribute evidence against the decision.

To use LRP on a DQN agent we first have to look at its network architecture. The DQN f , as introduced by Mnih et al. [50], consists of three convolutional layers $conv_1, \dots, conv_3$ followed by two fully connected layers fc_1 and fc_2 . For an input x we write $fc_i(x)$ and $conv_i(x)$ for the output of the layers fc_i and $conv_i$, respectively, during the forward pass that calculates $f(x)$. In this notation, the Q-Values (i.e., the output of the whole DQN) are $fc_2(x)$.

Following the LRP notation, we denote the relevance value of the j -th neuron in the layer l with R_j^l . As described above, we have to define messages $R_{j \leftarrow k}^{l,l+1}$ for any two consecutive Layers $l, l+1$ to determine a LRP variant. For now we assume that $l+1$ is one of the fully connected layers fc_i . The convolutional case works analogously and will be covered in more detail in the next section. $R_{j \leftarrow k}^{l,l+1}$ should measure the contribution of the j -th neuron of fc_{i-1} to the k -th neuron of fc_i , therefore we have to look at the calculation of $fc_i(x)_k$. The fully connected layer fc_i uses a weight matrix W_i , a bias vector b_i and an activation function σ_i as parameters for its output. Let W_i^k be the k -th row of W_i and b_i^k the k -th entry of b_i . Then the activation of the k -th neuron in $fc_i(x)$ is

$$\sigma_i(W_i^k \cdot fc_{i-1}(x) + b_i^k),$$

where \cdot denotes the dot product and fc_0 is the flattened output of $conv_3$.

Usually the ReLU function $\sigma(x) = \max(0, x)$ is used as activation function σ_i in the DQN architecture. Bach et al. [12] argue that any monotonous increasing function σ with $\sigma(0) = 0$, like the ReLU function, conserves the relevance of the dot product $W_i^k \cdot fc_{i-1}(x)$. Newer LRP variants, like the one used by Montavon et al. [52], also omit the bias when defining $R_{j \leftarrow k}^{l,l+1}$. With those two assumptions the relevance of each neuron of fc_{i-1} to $fc_i(x)_k$ is the same as their contribution to the dot product $W_i^k \cdot fc_{i-1}(x) = \sum_j w_{jk} fc_{i-1}(x)_j$. This is a linear decomposition, so we can use $w_{jk} fc_{i-1}(x)_j$ to measure the contribution of the j -th neuron of fc_{i-1} .

Since we want to find the parts of the input that contributed evidence in favor of the decision of the DQN agent, we restrict ourselves to the positive parts of that sum. That is, we set

$$z_{jk}^+ := \begin{cases} w_{jk} fc_{i-1}(x)_j & \text{if } w_{jk} fc_{i-1}(x)_j > 0 \\ 0 & \text{if } w_{jk} fc_{i-1}(x)_j \leq 0 \end{cases}.$$

With this, we define the messages as $R_{j \leftarrow k}^{l,l+1} := \frac{z_{jk}^+}{\sum_j z_{jk}^+} R_k^{l+1}$. This method is called the z^+ -rule (without bias) and satisfies the LRP Equation (2).

3.2. An argmax approach to LRP

In this subsection, we introduce our adjustment to the LRP variant called z^+ -rule which we revisited in the previous subsection. Recent work [36,23] indicates that DRL agents focus on certain objects within the visual input. With our approach, we aim to generate saliency maps that reflect this property by focusing on the most relevant parts of the input instead of giving too many details. For this purpose, we propose to use an *argmax* function to find the most contributing neurons in each convolutional layer.

This idea is inspired by Mopuri et al. [53], who generated visualizations for neural networks solely based on the positions of neurons that provide evidence in favor of the prediction. During this process, they follow only the most contributing neurons in each convolutional layer. Our method adds relevance values to the positions of those neurons and therefore expands the approach of Mopuri et al. by an additional dimension of information. Since those relevance values follow the LRP concept, they also possess the advantageous properties of the LRP concept like conservation of the prediction value.

As we have seen in the Foundations Section 3.1, an LRP method is defined by its messages $R_{j \leftarrow k}^{l,l+1}$ which propagate the relevance from a layer $l+1$ to the preceding layer l . If $l+1$ is a fully connected layer fc_i of the DQN (see Section 3.1 for our notation of the DQN architecture), we use the same messages that are used in the z^+ -rule. In the case that l and $l+1$ are convolutional layers $conv_{i-1}$ and $conv_i$, we propose new messages based on the *argmax* function. To define those messages we analyze how the activation of a neuron $conv_i(x)_k$ was calculated during the forward pass. Let W and A denote the




Fig. 1. A visualization of how our *argmax* approach differs from the z^+ rule.

weight kernel and part of $\text{conv}_{i-1}(x)$ respectively that were used to calculate $\text{conv}_i(x)_k$ during the forward pass. If we write W and A in the appropriate vector form, we get

$$\text{conv}_i(x)_k = \sigma\left(\sum_j w_j a_j + b\right),$$

where σ denotes the activation function of conv_i and b the bias corresponding to W . Analogously to the z^+ -rule we assume that the activation function and the bias can be neglected when determining the relevance values of the inputs a_i . We propose to use an *argmax* function to find the most relevant input neurons by defining the messages in the following way

$$R_{j \leftarrow k}^{l,l+1} := \begin{cases} R_k^{l+1} & \text{if } j = \text{argmax}\{w_j a_j\} \\ 0 & \text{if not.} \end{cases}$$

This definition satisfies the LRP condition given by Equation (2) because the only non-vanishing summand of the sum

$$\sum_{j \in \{j \text{ is input for neuron } k\}} R_{j \leftarrow k}^{l,l+1}$$

is R_k^{l+1} .

If we use the same *argmax* approach to propagate relevance values from conv_1 to the input conv_0 , then we get very sparse saliency maps where only a few neurons are highlighted. If we highlight the entire areas of the input conv_0 that were used to calculate relevant neurons of conv_1 , then we lose information about the relevance values inside those areas. Therefore, we draw inspiration from the guided Grad-CAM approach introduced in [62]. Guided Grad-CAM uses one thorough relevance analysis for the neurons of the last convolutional layer to get relevant areas for the specific prediction and another thorough relevance calculation for the input pixels to get fine granular relevance values inside those areas. We already did a thorough analysis of the neurons of the last convolutional layer by using the z^+ -rule on the fully connected layers. By following the most relevant neurons through the convolutional layers we keep track of the input areas that contributed the most to those values. Mimicking the second thorough analysis of the Guided Grad-CAM approach we propose to use the z^+ -rule to propagate relevance values from conv_1 to conv_0 . This generates fine granular relevance values inside the areas identified by following the most contributing neurons and ascertains that those relevance values follow the LRP concept.

Fig. 1 visualizes the differences between our *argmax* approach and the z^+ -rule. An implementation of our algorithm that builds upon the iNNvestigate framework [3] can be found here: https://github.com/HuTobias/LRP_argmax.

3.3. Illustration of the selectivity of the *argmax*-rule

In order to verify that our *argmax* approach, described in Section 3, creates more selective saliency maps than the z^+ -rule (see Section 3.1), we tested our approach on three different Atari 2600 games. For all games, we trained an agent using the DQN implementation of the OpenAI baselines framework [19]. The results of all experiments are shown in our previous work [35]. We review the Pacman results here since we use this game in the user study evaluating our combined explanation approach.




Fig. 2. The left image shows a screen of Pacman. The player (green circle) has to collect pellets (blue area) while avoiding ghosts (red circles). The saliency map created for this game-state by the z^+ -rule (middle) highlights a huge area as relevant while our *argmax* approach (right) focuses on the vicinity of the player. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)




Fig. 3. Example for how the LRP-argmax saliency maps change when the network's layers are randomized cascadingly, beginning with output layer fc_2 .

In the game Pacman, the player has to navigate through a maze and collect pellets while avoiding enemy ghosts. Because this game contains many important objects and gives the agent a huge variety of possible strategies, DQN agents struggle in this environment and perform worse than the average human player (see [50]). Explainable AI methods are especially desirable in environments like this, where the agent is struggling because they help us to understand where the agent had difficulties. The saliency maps created with the z^+ -rule (Fig. 2) reflect the complexity of Pacman by showing that the agent tries to look at nearly all of the objects in the game. This information might be helpful to optimize the DRL agent, but it also distracts from the areas which influenced the agents' decision the most. Fig. 2 shows that the saliency map created by the *argmax* approach is more focused on the vicinity of the agent and makes it clearer what the agent is focusing on the most. Fig. 2 also illustrates that a fine-granular saliency map in the vicinity of the agent is necessary to see that the agent will most likely decide on moving to the right as his next action.

3.4. Sanity checks

It is not yet possible to verify whether a saliency map algorithm perfectly reflects what a model learned. However, a basic prerequisite for this is that the saliency maps depend on the weights learned by the model. To verify this, Adebayo et al. [1] proposed sanity checks that cascadingly randomize each layer of the network, starting with the output layer. If the saliency maps depend on the learned weights, then this will lead to increasingly different visualizations. Sixt et al. [65] applied the sanity checks¹ for our *argmax*-rule and test it on the regular Pacman agents described in Section 6. An example of these tests for a single state is shown in Fig. 3.

To measure how similar two saliency maps are we use three different metrics proposed by Adebayo et al. [1]: Spearman rank correlation, structural similarity (ssim), and Pearson correlation of the histogram of oriented gradients. To account for a possible change of sign in the saliency maps, we adopt an approach by Sixt et al. [65] and use the maximum similarity of the original and the inverted saliency map. That means that for two saliency maps $S, S' \in \mathbb{R}^{m \times n \times c}$ and a similarity measurement $sim : \mathbb{R}^{m \times n \times c} \times \mathbb{R}^{m \times n \times c} \rightarrow \mathbb{R}$ we calculate the actual similarity with

$$\max(sim(S, S'), sim(\mathbf{1} - S, S')) \quad (3)$$

where $\mathbf{1} \in \mathbb{R}^{m \times n \times c}$ is filled with 1s. Fig. 4 shows the average similarities per randomized layer for a gameplay stream of 1000 states.

The relatively high values for the structural similarity (ssim) can be explained by the high amount of intersecting zeros in all saliency maps. Apart from that, we see the same trends already observed by Sixt et al. [65] and Adebayo et al.

¹ The code we used for the sanity checks can be found here: https://github.com/HuTobias/HIGHLIGHTS-LRP/tree/master/sanity_checks.




Fig. 4. The average similarities between saliency maps for the fully trained agent and agents where the layers have been randomized cascadingly, starting with the last layer fc_2 . The values are based on a stream of 1000 actions in the Atari 2600 Pacman game.

[1]: the saliency maps do analyze the learned weights but the fully connected layers are not sufficiently analyzed. As a consequence, the saliency maps are not class discriminatory. However, class discriminatory saliency maps often come with other drawbacks like being noisy [65] or not analyzing all layers [62].

4. Strategy summarization

This section describes the strategy summarization approach to global explanations, and the HIGHLIGHTS algorithm, and its extension HIGHLIGHTS-DIV, which we developed and evaluated in prior work [5].

Our formalization of the summarization problem assumes that the agent uses a Markov Decision Process (MDP), where A is the set of actions available to the agent, S is the set of states, $R: S \times A \rightarrow \mathbb{R}$ is the reward function, which maps each state and action to a reward, and Tr is the transition probability function, i.e., $Tr(s', a, s)$ defines the probability of reaching state s' when taking action a in state s . The agent has a policy π which specifies which action to take in each of the states.

We formalize the problem of summarizing an agent's behavior as follows: from execution traces of an agent, choose a set $T = \{t_1, \dots, t_k\}$ of trajectories to include in the summary, where each trajectory is composed of a sequence $((s_i, a_i), \dots, (s_{i+l-1}, a_{i+l-1}))$ of l consecutive states and the actions taken in those states. We consider trajectories rather than single states because seeing what action was taken by the agent in a specific state might not be meaningful without a broader context (e.g., watching a self-driving car for one second will not reveal much useful information). Because it is infeasible that people will be able to review the behavior of an agent in all possible states, we assume a limited budget k for the size of the summary, such that $|T| = k$. This budget limits the amount of time and cognitive effort that a person needs to invest in reviewing the agent's behavior.

There are several factors that could be considered when deciding which states to include in a summary, such as the effect of taking a different action in that state, the diversity of the states that are included in the summary, and the frequency at which states are likely to be encountered by the agent. The approach we describe here focuses on the first factor, which we refer to as the "importance" of a state. Intuitively, a good summary should provide a person reviewing the summary with a sense of the agent's behavior in states that the person considers important (e.g., when making a mistake would be very costly). The importance of states included in the summary could substantially affect the ability of a person to assess an agent's capabilities. For example, imagine a summary of a self-driving car that only shows the car driving on a highway with no interruptions. This summary would provide people with very little understanding of how the car might act in other, more important, scenarios (e.g., when another car drives into its lane when there is road construction). In contrast, a summary showing the self-driving car in a range of more interesting situations (e.g., overtaking another car, breaking when a person enters the road) would convey more useful information to people reviewing it.

4.1. The "Highlights" algorithm

The HIGHLIGHTS algorithm generates a summary of an agent's behavior from simulations of the agent in an online manner. It uses the notion of state *importance* [71] to decide which states to include in the summary. Intuitively, a state is considered important if taking a wrong action in that state can lead to a significant decrease in future rewards, as determined by the agent's Q-values. Formally, the importance of a state denoted $I(s)$, is defined as:

$$I(s) = \max_a Q_{(s,a)}^\pi - \min_a Q_{(s,a)}^\pi \quad (4)$$

This measure has been shown to be useful for choosing teaching opportunities in the context of student-teacher reinforcement learning [71,8].

Before providing a detailed pseudo-code of the algorithm, we describe its operation at a high-level. HIGHLIGHTS generates a summary that includes trajectories that capture the most important states that an agent encountered in a given number of simulations. To do so, at each step it evaluates the importance of the state and adds it to the summary if its importance value is greater than the minimum value currently represented in the summary (replacing the minimal importance

Algorithm 1: The HIGHLIGHTS algorithm.

Input: $\pi, k, l, \text{numSimulations}, \text{intervalSize}, \text{statesAfter}$
Output: T

```

1 runs = 0
2  $T \leftarrow \text{PriorityQueue}(k, \text{importanceComparator})$ 
3  $t \leftarrow$  empty list
4  $c = 0$ 
5 while (runs < numSimulations) do
6   sim = InitializeSimulation()
7   while (!sim.ended()) do
8      $(s, a) \leftarrow \text{sim.advanceState}(\pi)$ 
9     if ( $|t| == l$ ) then
10       $t.remove()$ 
11      $t.add((s, a))$ 
12     if ( $c > 0$ ) then
13       $c = c - 1$ 
14      $I_s \leftarrow \text{computeImportance}(\pi, s)$ 
15     if ( $\text{intervalSize} - c == \text{statesAfter}$ ) then
16      lastSummaryTrajectory.setTrajectory(t)
17     if ( $(|T| < k)$  or ( $I_s > \text{minImportance}(T)$ )) and ( $c == 0$ ) then
18       if  $|T| == k$  then
19          $T.pop()$ 
20        $st \leftarrow \text{new summaryTrajectory}(I_s)$ 
21        $T.add(st)$ 
22       lastSummaryTrajectory  $\leftarrow st$ 
23        $c = \text{intervalSize}$ 
24   runs = runs+1

```

Table 1

Parameters of the HIGHLIGHTS algorithm and the values assigned to them in the experiments reported in [5] (in parentheses).

Parameter	Description (value used in experiments)
k	Summary budget, i.e., number of trajectories (5)
l	Length of each trajectory (40)
numSimulations	The number of simulations run by HIGHLIGHTS (50)
intervalSize	Minimal number of states between two trajectories in the summary (50)
statesAfter	Number of states following s to include in the trajectory (10)

state). To provide more context to the user, for each such state HIGHLIGHTS also extracts a trajectory of states neighboring it and the actions taken in those states.

A pseudo-code of the HIGHLIGHTS algorithm is given in Algorithm 1. Table 1 summarizes the parameters of the algorithm. HIGHLIGHTS takes as input the policy of the agent π which is used to determine the agent's actions in the simulation and state importance values, the budget for the number of trajectories to include in the summary (k), and the length of each trajectory surrounding a state (l). Each such trajectory includes both states preceding the important state and states that were encountered immediately after it. The number of subsequent states to include is determined by the *statesAfter* parameter (the number of preceding states can be derived from this parameter and l). We also specify the number of simulations that can be run (*numSimulations*), and the minimal “break” interval between trajectories (*intervalSize*) which is used to prevent overlaps between trajectories. HIGHLIGHTS outputs a summary of the agent's behavior, which is a set of trajectories (T).

The algorithm maintains two data structures: T is a priority queue (line 2), which will eventually hold the trajectories chosen for the summary; t is a list of state-action pairs (line 3), which holds the current trajectory the agent encounters. The procedure runs simulations of the agent acting in the domain. At each step of the simulation, the agent takes an action based on its policy and advances to a new state (line 8). That state-action pair is added to the current trajectory (line 11). If the current trajectory reached its maximal length, the oldest state in the trajectory is removed (lines 9–10). HIGHLIGHTS computes the importance of s based on the Q-values of the agent itself, as defined in Equation (4) (line 14).

If a sufficient number of states were encountered since the last trajectory was added to the summary, state s will be considered for the summary (the $c == 0$ condition in line 17). State s will be added to the summary if one of two conditions holds: either the size of the current summary is smaller than the summary size budget, or the importance of s is greater than the minimal importance value of a state currently represented in the summary (line 17). If one of these conditions holds, a trajectory corresponding to s will be added to the summary. The representation of a trajectory in the summary (a *summaryTrajectory* object) consists of the set of state-action pairs in the trajectory (which will be presented in the summary), and the importance value I_s based on which the trajectory was added (such that it could be compared with the

importance of states encountered later). This object (st) is initialized with the importance value (line 20) and is added to the summary (line 21), replacing the trajectory with minimal importance if the summary reached the budget limit (lines 18–19). Because the trajectory will also include states that follow s , the final set of state-action pairs in the trajectory is updated later (lines 15–16). Last, we set the state counter c to the interval size, such that the immediate states following s will not be considered for the summary. At the end of each simulation, the number of runs is incremented (line 24). The algorithm terminates when it reaches the specified number of simulations.

Originally, HIGHLIGHTS was implemented in an online algorithm because it is less costly, both in terms of runtime and in terms of memory usage. In addition, such an algorithm can be incorporated into the agent's own learning process without additional cost. In this paper, we adapt the algorithm to work offline, as described in Section 5.

4.2. Considering state diversity: the HIGHLIGHTS-DIV algorithm

Because HIGHLIGHTS considers the importance of states in isolation when deciding whether to add them to the summary, the produced summary might include trajectories that are similar to each other. This could happen in domains in which the most important scenarios tend to be similar to each other. To mitigate this problem, we developed a simple extension to the HIGHLIGHTS algorithm, which we call HIGHLIGHTS-DIV. Similar to HIGHLIGHTS, this algorithm also determines which states to include in the summary based on their importance. However, it also attempts to avoid including a very similar set of states in the summary, thus potentially utilizing the summary budget more effectively.

HIGHLIGHTS-DIV takes into consideration the diversity of states in the following way: when evaluating a state s , it first identifies the state most similar to s that is currently included in the summary,² denoted s' . Then, instead of comparing the importance of a state to the minimal importance value that is currently included in the summary, HIGHLIGHTS-DIV compares I_s to $I_{s'}$. If I_s is greater than $I_{s'}$, the trajectory that includes s' in the summary will be replaced with the current trajectory (which includes s). This approach allows less important states to remain represented in the summary (because they will not be compared to some of the more important states that differ from them), potentially increasing the diversity of trajectories in the summary and thus conveying more information to users.

4.3. Empirical evaluation of HIGHLIGHTS and HIGHLIGHTS-DIV

We summarize the main results of the study conducted in our previous work, which demonstrated the usefulness of HIGHLIGHTS and HIGHLIGHTS-DIV summaries. For complete details of the study design and its results see Amir & Amir [5]. The performance of the basic HIGHLIGHTS algorithm was compared with that of two baselines: (1) likelihood-based summaries generated by sampling k trajectories uniformly from the agent's execution trace (such that states that are more likely to be visited are also more likely to appear in the summary), and (2) summaries generated from the first k trajectories the agent encounters. The task used in the study was identifying the agent that performs better in pairwise comparisons, based on the summaries. Three Ms. Pacman agents were trained to vary in their quality: a high-quality agent, medium-quality agent, and low-quality agent. This was achieved by varying the number of training episodes.

In the first experiment, 40 participants recruited from Amazon Mechanical Turk (23 female, mean age = 35.35, STD = 10.4), were asked to make the pairwise agent comparisons based on summaries generated by either the basic HIGHLIGHTS algorithm or one of the two baselines (Likelihood-based or First). The study used a within-subject design, such that each participant completed nine comparison tasks showing all combinations of pairs of agents and the summary method (e.g., comparing the high-quality agent to the low-quality agent based on the HIGHLIGHTS summary). In the second experiment 48 additional participants (25 female, mean age=36, STD=11.6), performed the same task, but this time summaries were generated either by HIGHLIGHTS-DIV, basic HIGHLIGHTS or the likelihood-based baseline (since the "first" baseline led to the worst performance in the first experiment). In both experiments, participants were incentivized to answer correctly as they received a bonus payment depending on their performance.

Results aggregated from both experiments are shown in Fig. 5. Both HIGHLIGHTS and HIGHLIGHTS-DIV summaries led to significantly improved performance of participants compared to the baselines. HIGHLIGHTS-DIV further led to improved performance compared to HIGHLIGHTS, especially when comparing the medium quality agent with the high-quality agent, which was the hardest comparison to make as their actual performance did not differ by much. Participants also expressed a subjective preference to HIGHLIGHTS summaries compared to baselines.

5. Integrating local and global information

In this section, we describe our integration of the local LRP-argmax saliency maps described in Section 3 into the global HIGHLIGHTS-DIV summaries described in Section 4.2. To this end, we describe how the agents were trained, what adjustments we made to HIGHLIGHTS-DIV for the deep reinforcement learning algorithm we used, how we generated saliency maps, and how the combined information is displayed.

² We assume that distance metric to compare states can be defined. This can be done in many domains, e.g., by computing Euclidean distance if states are represented by feature vectors.




Fig. 5. Correctness rates of participants (aggregated from both experiments) in choosing the better performing agent. The x-axis shows the three different pairwise agent comparison tasks (low quality agent vs. medium quality agent, etc.). In all cases, HIGHLIGHTS and HIGHLIGHTS-DIV outperformed the two baselines. HIGHLIGHTS-DIV led to significant improvement over HIGHLIGHTS only in the Low Vs. Medium agent comparison, which was the most difficult comparison as the two agents were most similar to each other in performance.

Agent training To evaluate our combined explanation approach, we trained several Pacman agents using the OpenAI baselines [19] implementation of the DQN-algorithm [50]. The network architecture used in this implementation is described in Section 3.1. The environment we use is the Atari 2600 game MsPacman included in the Arcade Learning Environment (ALE) [13], which we refer to in this work as Pacman for simplicity.

For each step in a game, the input state consists of the four last frames (the raw pixel values of a single screen of the game) f_1 to f_4 . Each frame f_i is converted to grayscale and scaled down to 84×84 pixels. The frames are then stacked to enable the agent to see temporal differences (i.e., movement). The agent chooses an action only every four frames and this action is repeated for the next four frames. In Pacman, the agent has nine different actions to choose from, which correspond to the meaningful actions that can be achieved with an Atari 2600 controller (do nothing, up, down, left, right, up-left, up-right, down-left, down-right).

The reward is based on the ALE [13] reward function that uses the increase of the in-game score between the beginning of the four frames of a state compared to the score after those frames. The final reward functions we used are detailed in Section 6, when we describe the agents used in the empirical evaluation.

Generating gameplay streams and saliency maps Since deep neural networks increase the time that the agent needs for each prediction and the LRP-analysis of this decision requires additional computation time, we recorded a stream of 10,000 steps for each agent and used them to create our summaries. These streams also increase the reproducibility of our experiments.³

We computed the average in-game score of each trained agent over the entire stream. This allows us to objectively say which agent achieved the most points during the simulations used for our summaries and therefore gives us a ground truth for the agent comparison task (see Section 6).

Since the Atari 2600 version of Pacman does not respond to input for the first 250 frames (empirically tested) after the game starts, we exclude those frames from the streams. Furthermore, we force the agent to repeat the ‘do nothing’ action for a random amount of steps between 0 and 30, until it is allowed to choose actions based on its policy. This method introduces randomness into the deterministic Pacman game and is also used during training by the DQN algorithm [50,19]. Saliency maps are created using the LRP-argmax algorithm described in Section 3.

Adjustments to HIGHLIGHTS-DIV For the summaries, we make several adjustments to the HIGHLIGHTS-DIV algorithm described in Section 4.2, to adapt it to the DQN settings. First, we change the way importance is calculated. Instead of using Equation (4), which calculates the importance by comparing the highest with the lowest Q-value, we use the difference between the highest and second-highest Q-values. Let second-highest be the operation that finds the second-highest value in a set, then this can be written as:

$$I(s) = \max_a Q_{(s,a)}^\pi - \text{second-highest}_a Q_{(s,a)}^\pi \quad (5)$$

While examining the gap between the best and worst actions worked well in a simpler Pacman environment in which there were only four possible actions, it did not generalize well to the Atari environment where there is a larger number of actions. One possible explanation for this is that some of the 9 actions of the Pacman environment overlap. For example, “left” and “top left” can be used interchangeably in many states. Therefore the agent might ignore some of the actions completely. To verify this, we examined the frequency of choosing each action and found that two of the three agents

³ Since the streams are fairly big we did not upload them. They are available upon request from the authors.

we trained were clearly biased against certain actions.⁴ Therefore, some Q-values are largely uninformed by exploration and might have arbitrarily low values, making the worst Q-value non-informative. For the diversity computation in HIGHLIGHTS-DIV, we use Euclidean distance over the raw $84 \times 84 \times 4$ input states.

Since we pre-generated a stream of 10,000 states, we implement an offline version of HIGHLIGHTS-DIV that selects the states for the summary retrospectively from the generated stream. The procedure begins by sorting the states based on their Q-values and adding them to the summary according to this ordering. To reduce the number of overlapping trajectories, we compare each new state with all states in the current summary and corresponding context states (this is equivalent to the HIGHLIGHTS-DIV variant). To find a suitable threshold that determines when a state is too similar to the states that were already selected for the summary, we randomly pick a subset of 1,000 random states from the recorded stream and calculate the similarity between each pair of states in this set. Then, we set the threshold to be a percentile of the distribution of those similarity values. We empirically found (by manually examining a sample of states) that using a threshold of 3% led to no obvious duplicate trajectories for any of the agents.

Video generation The videos we generate from the states chosen by the summary show 30 frames per second. To emphasize that demonstrations show different trajectories, they are separated by a black screen that appears for 1 second (inspired by the fade-out effect used in [63]). To prevent the users from using the in-game score to gauge how good an agent is, we mask the bottom half of the screen with black pixels. In pilot studies, participants complained that the videos were flickering too much. One of the reasons for this is that the Atari 2600 implementation of Pacman does not show every object in every frame to save computing power. Since we showed all frames after each other these objects appeared to be blinking and distracted the viewers. To combat this problem, we do not display the current frame f_i . Instead we display $\max(f_i, f_{i-1})$, the maximum of each pixel over the current frame f_i and the preceding frame f_{i-1} . While this introduces some artifacts (e.g. red pellets showing through blue ghosts) it considerably reduces the flickering.

Another measure we take against this flickering is to interpolate between the different saliency maps instead of showing a completely different saliency map for each frame. Let f_1 to f_4 be the four frames of an input state and let s_1 to s_4 be the saliency maps for each of these frames that analyze the agent's decision in this state. For $i < 4$ the action that Pacman will take after frame f_i is not related to the saliency map s_i since the agent only decides on a new action every four frames and is still repeating the action that he decided on based on the last state (composed of the 4 frames before f_1). Therefore we show the saliency map s_4 over the frame f_4 and for the other frames ($i < 4$) we interpolate between the last shown saliency map and s_4 .

Before this interpolation, we normalize the saliency maps to have a maximum of 1 and a minimum of 0. We do this across all 4 frames of the states s_1, \dots, s_4 to avoid losing information that might be transported in the magnitude of relevance values between the frames.

Finally, we add the interpolated saliency maps to the green channel of the original screen frame. Our complete implementation can be found here: <https://github.com/HuTobias/HIGHLIGHTS-LRP>.

6. Empirical evaluation

To evaluate our hypothesis that there is benefit to combining global and local explanations of RL agents, we conducted a user study. In this study, participants were asked to compare different agents and to reflect on the strategies of agents based on the information they were shown. We next describe in detail the study design, the specific hypotheses we tested, and the metrics we used to evaluate the results.

6.1. Study design

Empirical domain We used the Atari game Pacman for our experiments (see Section 5 for the specific implementation). Atari games are a common benchmark for state-of-the-art reinforcement learning algorithms [13,19,50,74] and to test explanation methods for those algorithms [5,27,35,43,75]. We chose Pacman since it is not as reaction-based as some other Atari games (e.g. Breakout or Enduro) and allows the RL agents to develop different strategies. Furthermore, no additional domain knowledge is necessary to understand Pacman and the rules are not too complicated. This enables us to conduct a study with a wide range of participants by simply explaining the rules at the beginning of the study.

In the game, Pacman obtains points by eating food pellets while navigating through a maze and escaping ghosts. There are two types of pellets: regular pills for which Pacman receives 10 points, and power pills that are worth 50 points and also turn the ghosts blue, which makes them edible by Pacman. Pacman receives 200, 400, 800, 1600 points for each ghost it eats successively. At random intervals, cherries spawn and move through the labyrinth. Eating a cherry gives 100 points.

To evaluate participants' ability to differentiate between alternative agents and analyze their strategies, we trained agents that behave qualitatively differently. To this end, we modified the reward function used for training (similar to the approach used by Sequeira et al. [63]), resulting in three types of agents. As mentioned in Section 5, we based all of those reward functions on the default ALE [13] reward function, which measures the increase in the in-game score (as described above) between the first and last frame of a state.

⁴ The results can be seen in https://github.com/HuTobias/HIGHLIGHTS-LRP/tree/master/action_checks.

Table 2
The four study conditions.

	Likelihood-based summaries	HIGHLIGHTS-DIV
No saliency maps	<i>L</i>	<i>H</i>
LRP saliency maps	<i>L+S</i>	<i>H+S</i>

- *Regular agent*: This agent was trained using the default reward function of the ALE.⁵
- *Power pill agent*: This agent was trained using a reward function that only assigned positive rewards to eating power pills.⁶
- *Fear-ghosts agent*: This agent used the default ALE reward function but was given an additional negative reward of -100 when being eaten by ghosts, causing it to more strongly fear ghosts (which is implicitly learned by the other agents due to the lack of future rewards caused by being eaten).

Each agent was trained for 5 Million steps with the algorithm described in Section 5. At the end of this training period, the best-performing policy is restored.

Experimental conditions To evaluate the potential benefits of integrating global and local explanations, and their relative importance, we assigned participants to four different conditions (summarized in Table 2). The first two conditions included only global information, while the remaining two conditions integrated local explanations as well:

- **Likelihood-based Summaries (*L*)**: The summaries in this condition consisted of states that the agent was likely to encounter during gameplay. To generate these summaries, we randomly select state-action pairs from the streams of the Pacman agents playing the game. Since each state encountered in the game had the same probability of being chosen, states that are encountered more frequently will be more likely to be included. Because of the random component of these summaries, it is possible that a single summary is, by chance, particularly good or particularly bad. Therefore, we generated 10 different likelihood-based summaries and randomly assigned them to participants in this condition.
- **HIGHLIGHTS-DIV Summaries (*H*)**: In this condition, participants were shown summaries generated by the HIGHLIGHTS-DIV algorithm. The specific implementation of this algorithm and the parameters we used for diversity are described in Section 5.
- **Likelihood-based Summaries+Saliency (*L+S*)**: These summaries included the same states as those shown in the *L* summaries, but each image was overlayed with a saliency map generated by the LRP-argmax algorithm described in Section 3.
- **HIGHLIGHTS-DIV Summaries+Saliency (*H+S*)**: These summaries included the same states as those shown in the *H* summaries, where each image was overlayed with a saliency map generated by the LRP-argmax algorithm described in Section 3.

We used a budget of $k = 5$ for the summaries. That is, each summary included 5 base states chosen either based on likelihood or by HIGHLIGHTS-DIV, where for each state we included a surrounding context window of 10 states that occurred right before and after the chosen state. We used an interval size of 10 states to prevent directly successive states in the summary.

The video creation and saliency map overlay process are described in detail in Section 5. All video summaries used in the study are available online.⁷

We note that we did not include a condition that shows only local explanations, since by definition a local explanation is given for a specific state, forcing us to make some choice about which states to show (which means making a global decision). However, the *L+S* condition simulates a scenario where local explanations are shown for states selected based on the likelihood of the agent encountering them during gameplay.

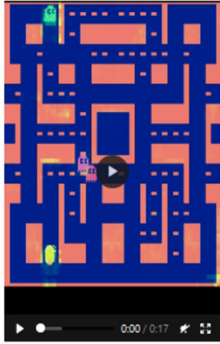
Participants We recruited participants through Amazon Mechanical Turk ($N = 133$ (*L*: 33, *H*: 33, *L+S*: 34, *H+S*: 33)), the majority of participants were between the ages of 25 and 44, 47 females). Participation was limited to people from the US, UK, or Canada (to ensure sufficient English level) with a task approval rate greater than 97%. Each participant was assigned to one of the four conditions (*L*: 33, *H*: 33, *L+S*: 34, *H+S*: 33). Since saliency maps are not designed for color blind people, the participants were also asked if they were color blind and stopped from participating if they are.

Procedure Participants were first asked to answer demographic questions (age and gender) and questions regarding their experience with Pacman and their views on AI. Then, they were shown a tutorial explaining the rules of the game Pacman

⁵ To remove unnecessary magnitude we divided the rewards by the factor 10, such that a regular pill gives a reward of 1.

⁶ We achieved this by only giving the agent a reward if the increase in score was between 50 and 99. The range is necessary since Pacman is forced to eat at least one regular pill directly before it eats a power pill.

⁷ https://github.com/HuTobias/HIGHLIGHTS-LRP/tree/master/Survey_videos.

Task Description


Please briefly describe the strategy of the AI agent shown in the video above:

Based on this video, select the objects that you think were most important for the strategy of this particular agent.

☐ Pacman
☐ Normal pills
☐ ...

How confident are you that you chose the better agent?

Not at all confident
☐
☐
☐
☐
☐
☐
☐
Very confident

Please briefly explain your selection:

Fig. 6. A sketch of the retrospection task: participants were asked to analyze the behavior of each agent by providing a textual description of its strategy and identifying the objects that are most important to its decision-making. The full task can be seen in Appendix D.

and were asked to play the game to familiarize themselves with it. To verify that participants understood the rules, they were asked to complete a quiz and were only allowed to proceed with the survey after answering all questions correctly. After completing the quiz, they were given information and another quiz regarding the Pacman agent video summaries. In conditions *L+S* and *H+S*, this also included an explanation and a quiz about saliency maps. Then, they proceeded to the main experimental tasks. See Appendix D for the complete questionnaire. Participants were compensated as follows: they received a \$4 base payment and an additional bonus of 10 cents for each correct answer. The study protocol was approved by the Institutional Review Board at the Technion.

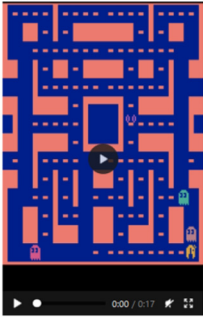
Main tasks We aimed to investigate three aspects related to the participants in the study: (1) the mental model of the participant about the agent, (2) participants' ability to assess agents' performance (appropriate trust), and (3) participants' satisfaction with respect to the explanations presented.

Task 1: Eliciting Mental Models through Retrospection. By mental model, we understand the cognitive representation that the participant has about a complex model [29,54], in our case, the agent. Humans automatically form mental models of agents based on their behavior [10]. These mental models help users understand and explain an agent's behavior. The examination of participants' mental models and their correctness helps to verify if explainable AI has been successfully applied [59,11]. To evaluate which mental models participants have formed about the agent's behavior, we designed a **retrospection task**. Here, we used a task reflection method inspired by prior studies [9,63], which is recommended by Hoffman et al. [31]. This task asked the participants to analyze the behavior of the three different AI agents: *Regular agent*, *Power pill agent* and *Fear-ghosts agent*. The ordering of the agents was randomized. Specifically, participants were shown the video summary (according to the condition they were assigned to) and were asked to briefly describe the strategy of the AI agent (textual), and to select up to 3 objects that they think were most important to the strategy of the agent (the possible objects were Pacman, power pills, normal pills, ghosts, blue ghosts, and cherries). They were also asked how confident they were in their responses, and to justify their reasoning. Fig. 6 shows a sketch of a retrospection task.

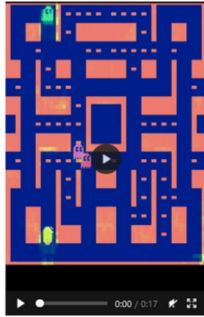
Task 2: Measuring Appropriate Trust through Agent Comparison. We use the term appropriate trust, based on the work of Lee and See [44] who present a conceptual 'trust in automation' framework. They define appropriate trust as a

Task Description

Agent A



Agent B



Which of the Pacman agents do you choose to play on your behalf?
☐ Agent A ☐ Agent B

How confident are you that you chose the better agent?
Not at all confident ☐ ☐ ☐ ☐ ☐ ☐ Very confident

Please briefly explain your selection:

Fig. 7. A sketch of the agent comparison task: participants were asked to choose which agent they would like to play on their behalf (i.e., identify the better performing agent) according to the two summary videos. The full task can be seen in Appendix D.

well-calibrated trust that matches the true capabilities of a technical system. We measure the appropriate trust using an **agent comparison task**. Here, the participants were shown summaries of two of the three agents at a time and were asked to indicate which agent performs better in the Pacman game (similar to tasks used in [5,62]). They thus made three comparisons (*Regular agent* Vs. *Power pill agent*, *Regular agent* Vs. *Fear-ghosts agent* and *Power pill agent* Vs. *Fear-ghosts agent*). We do not ask the participant directly about their trust in the two agents shown. Instead, the participants have to choose one of the two agents that they would like to play on their behalf (see Fig. 7). If they choose the correct agent then they receive a bonus payment as described in the previous *Procedure* paragraph. This implicit question reveals which agent participants consider more reliable and qualified for the task. As in the retrospection task, they were asked to indicate their level of confidence and to provide a textual justification for their decision. The ordering of the three agent comparisons was randomized.

Explanation satisfaction questions. Miller [48,49] argues that the end users' impressions about the agent should be queried and included in the evaluations of the explainable AI methods. This would ensure that the developed explanation methods are comprehensible not only to ML experts but also to end-users. We address this concern in our study by measuring participants' subjective satisfaction. To this end, we used **explanation satisfaction questions** adapted from the questionnaire proposed by Hoffman et al. [31]. We did this separately for the retrospection task (immediately after completing the three retrospection tasks) and for the agent comparison task (after completing the three comparisons), as we hypothesized there may be differences in the usefulness of the summaries for these two different types of tasks. Specifically, participants were asked the following questions using a 5-point Likert scale:

1. From watching the videos of the AI agents, I got an idea of the agents' strategies.
2. The videos showing the AI agents play contain sufficient detail about the agents' behavior.
3. The videos showing the AI agents play contain irrelevant details.
4. The videos showing the AI agents play were useful for *the task*. (only shown in groups *L* and *H*).
5. The gameplay scenarios shown in the videos were useful for *the task*. (only shown in groups *L+S* and *H+S*).
6. The green highlighting in the videos was useful for *the task*. (only shown in groups *L+S* and *H+S*).

We substituted *the task* with either *analyzing the agents' behavior* or *choosing the agent that performs better*, depending on the task they had just completed.

6.2. Hypotheses

Overall, we hypothesized that HIGHLIGHTS-DIV summaries will be more useful than likelihood-based summaries in both the retrospection and agent comparison tasks and that adding saliency maps will further improve participants' performance. More specifically, we state the following hypotheses:

- H1: For both tasks, participants shown summaries generated by HIGHLIGHTS-DIV will perform better than participants shown likelihood-based summaries. That is, performance in H will be better than performance in L and performance in $H+S$ will be better than performance in $L+S$. We expect HIGHLIGHTS-DIV summaries to be more useful as they demonstrate the agent's behavior in more meaningful states, which should help both in identifying which agent performs better (in line with prior findings [5,33]), as well as in determining whether an agent is capable of performing well in certain scenarios [33]. We expect similar effects in terms of participants' explanation satisfaction in each task.
- H2: For both tasks, adding saliency maps will improve participant's performance and satisfaction. That is, we expect the performance in $L+S$ will be better than in L and similarly, that performance in $H+S$ will be better than in H . Here, too, we expect similar effects in terms of participants' explanation satisfaction in each task. We expect this to be the case as the saliency maps allow people to see not only what actions the agent chooses, but also what information it attends to. Previous studies also found positive effects of saliency maps on participants' mental models [9,4] and on their ability to choose the better performing prediction model [62].
- H3: The effect of the summary generation method on satisfaction and performance will be greater than that of the inclusion of saliency maps in the agent comparison task. That is, we expect that global information will be more crucial for identifying the better performing agent, as it explicitly demonstrates how the agents act.
- H4: The effect of adding saliency maps on satisfaction and performance will be stronger than that of the summary generation method in the retrospection task. Since saliency maps explicitly show what information the agent attends to, we hypothesize it will contribute more to identifying the agent's strategy. However, this is complicated by the fact that likelihood-based summaries might not include interesting scenarios, making saliency maps less helpful in this case. Therefore, our more specific hypotheses are:
 - H4.1: Participants in the saliency conditions will be more likely to identify Pacman, the main source of information for our agents, as an important object.
 - H4.2: Participants in the HIGHLIGHTS conditions will be more likely to identify objects that relate to agent goals, such as power pills and blue ghosts. Therefore, they will also more accurately describe the agents' strategies.

6.3. Analysis

We analyze the main hypotheses using the non-parametric Mann-Whitney test [47], as our dependent variables are not normally distributed. We report effect sizes using rank biserial correlation [70]. Additionally, we report the mean values and the 95% confidence interval (CI) computed using the bootstrap method. In all plots, the error bars correspond to the 95% confidence intervals.

To make sure that the participants involved in our analysis did in fact watch the videos of the agents, we recorded whether they clicked play on each video in addition to how often each video was paused. We did not force them to watch the videos to filter out participants that would have just pressed play to avoid the forcing mechanism. Since we saw from the raw data that some participants only stopped watching videos after the retrospection task, we checked each task separately. As a heuristic to measure how attentively a user watched the videos of a task, we took the sum of pauses of the videos in this task, where watching a video until the end was recorded as a pause and not clicking play was counted as -1 pause. Based on this heuristic we removed all participants from the retrospection task who did not have at least three pauses (5 participants) and all participants from agent comparison task who did not have at least six pauses (11 participants). The number of necessary pauses in each task is equal to the number of videos in this task.

For evaluating the retrospection task we use a scoring system, where two of the authors involved in the training of the agents assigned a score to each item for each agent before the study started (see Appendix C for details). For example for the *Power pill agent*, which was only rewarded when it ate a Power pill, selecting the Power pill or Pacman increased the score by 1 point and including any other item reduced the score by 1 point. Furthermore, selecting more than three items resulted in a score of zero, since the participants were told to select a maximum of three items.

Inspired by Anderson et al. [9] we use summative content analysis [32] to evaluate participants' textual responses. An independent coder (not one of the authors) classified responses to the questions "Please briefly describe the strategy of the AI agent shown in the video above" in the retrospection task, and the question "Please briefly explain how you came to your selection" in both the retrospection task and the agent comparison task. Each question was asked three times (once for each agent description or agent comparison) resulting in 402 answers per question. For the first question, the coder identified 67 different concepts in the answers. For example, the answer "The strategy of this Pacman agents seems to be to mainly avoid the ghosts as it eats the normal pills on the screen. Although it can be seen eating a power pill, the clip still does not show Pacman seeking out and eating the ghosts" was coded to "prioritizing normal pills", "avoiding ghosts" and "do not care about blue ghosts". We aggregated those concepts to 16 groups by combining similar concepts like "eating normal pills" and "prioritizing normal pills".

Table 3

Summary of all significance tests (calculated with Mann-Whitney tests). The * denotes statistically significant differences and † denotes a p-value < 0.1.

Task	Variable	Effect of strategy summarization		Effect of saliency maps	
		$H > L$	$H+S > L+S$	$L+S > L$	$H+S > H$
retrospection task	score	0.008*	3.3e-05*	0.965	0.514
	satisfaction	0.021*	0.035*	0.677	0.710
	text score				0.088†
agent comparison task	score	0.014*	0.180	0.062†	0.307
	satisfaction	0.147	0.235	0.627	0.833




Fig. 8. Comparison of participants' average performance in each task, by condition. Participants in the HIGHLIGHTS conditions H and $H+S$ outperformed the likelihood-based conditions L and $L+S$. Saliency maps only had a slight positive effect when added to likelihood-based summaries in the agent comparison task.

To evaluate the correctness of participants' answers we implemented a simple scoring system. For each agent and for each answer group, we decided whether it is correct, irrelevant, or wrong, based on predefined 'ground-truth' answers that two of the authors, who were involved in the training of the agents, wrote for each agent before the study started. The exact groups and their assigned scores can be found in Appendix C and the open-sourced code.

The answers to the second question regarding participants' justifications of their responses were classified into six categories (the answer could be based on the game rules, the saliency maps, the gameplay, participants' interpretation, and two categories for unjustified or unrelated justifications which we grouped into one "unjustified" category) and an additional seventh category for the agent comparison task, that encoded that the user could not decide between the two agents and guessed.

We note that the classifications assigned by the coder are not mutually exclusive.

7. Results

In this section, we report the results of our study. We first describe the characteristics of the participant population with respect to their AI experience, attitude towards AI, and Pacman experience. Then we assess the main hypotheses (H1–H4) (results summarized in Table 3) and further provide a descriptive analysis of additional variables such as participants' confidence and analysis of mistakes.

AI and Pacman experience We verify that participants in different conditions did not differ much in their AI experience and views, and in their experience with the game Pacman. To this end, we asked them when they played Pacman for the last time. Across all four conditions, the majority of participants answered: 'I played Pacman more than 5 years ago'. After receiving a short description of what AI is (using a formulation based on Russel [58]), 104 participants stated that they had experience with AI. The exact kind of experience ranged from 'I know AI from the media' (78 participants) to 'I do research on AI related topics' (14 participants). On average the users had a positive attitude towards AI (mean of 3.95 on a 5-point Likert scale). There are no meaningful differences between the conditions (see Appendix A for more details).

(H1) Participants shown HIGHLIGHT-DIV summaries performed better than participants shown likelihood-based summaries Participants' correctness rates for the agent comparison task are shown in Fig. 8(b). These results support H1, which states that

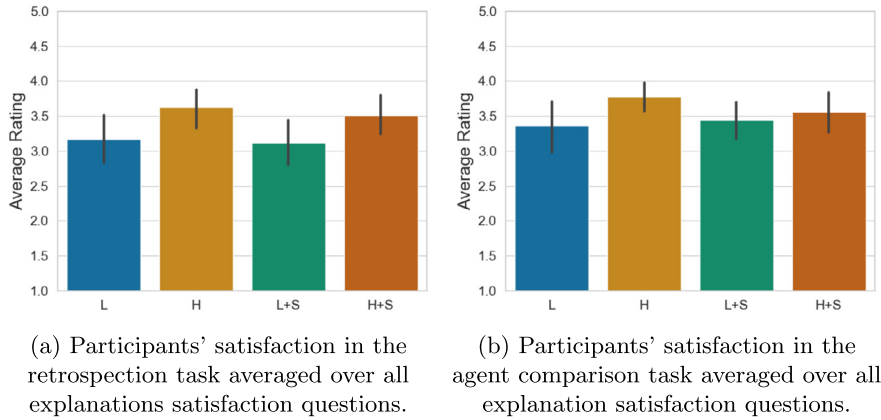


Fig. 9. Comparison of participants' average explanation satisfaction in each task, by condition. Each participant rated their agreement with several statements adapted from the explanation satisfaction questions proposed by Hoffman et al. [31] on a 5-point Likert scale (see Section 6.1). Participant's final rating was averaged over all those ratings, reversing the rating of the negative statements. Overall, participants in the HIGHLIGHTS conditions *H* and *H+S* rated the explanations highest.

HIGHLIGHTS-DIV summaries will lead to improved performance in both the agent comparison task and the retrospection task. The exact definition of performance per task is described in more detail in Section 6.3. Specifically, in the agent comparison task we find that participants in condition *H* significantly outperformed participants in condition *L* (*H*: mean=2.1, 95% CI=[1.83, 2.33], *L*: mean=1.63, 95% CI=[1.34, 1.91], Mann-Whitney test $U=334.5$, $p=0.014$, $r_{rb}=0.3$).⁸ While participants in the *H+S* condition achieved higher mean correctness rates than participants in the *L+S* condition, this difference is not statistically significant (*H+S*: mean=0.71, 95% CI=[0.6, 0.82], *L+S*: mean=0.65, 95% CI=[0.54, 0.75], Mann-Whitney test $U=391$, $p=0.180$, $r_{rb}=0.13$). Similarly, participants' average explanation satisfaction ratings, shown in Fig. 9(b), indicate that participants in condition *H* were more satisfied with the videos they received than the other participants. However, this difference is not significant (see Table 3).

With respect to participants' performance during the retrospection task, we find even stronger results (Fig. 8(a)) than in the agent comparison task, further supporting H1. Here too, participants in condition *H* obtained a higher score in the object selection sub-task than participants in condition *L* (*H*: mean=2.5, 95% CI=[1.89, 3.03], *L*: mean=1.5, 95% CI=[0.92, 2.06], Mann-Whitney test $U=346.5$, $p=0.008$, $r_{rb}=0.34$) and participants in the *H+S* condition received a higher score than participants in the *L+S* condition (*H+S*: mean=2.55, 95% CI=[2.02, 3.06], *L+S*: mean=0.73, 95% CI=[0.13, 1.31], Mann-Whitney test $U=206.5$, $p=0.00003$, $r_{rb}=0.58$). We found analogous significant differences in participants' explanation satisfaction during the retrospection task (Fig. 9(a)). Here, participants in condition *H* were more satisfied than participants in condition *L* (*H*: mean=3.63, 95% CI=[3.35, 3.88], *L*: mean=3.17, 95% CI=[2.82, 3.5], Mann-Whitney test $U=373.0$, $p=0.021$, $r_{rb}=0.29$) and participants in the *H+S* condition were more satisfied than participants in the *L+S* condition (*H+S*: mean=3.52, 95% CI=[3.25, 3.78], *L+S*: mean=3.12, 95% CI=[2.81, 3.43], Mann-Whitney test $U=364.5$, $p=0.035$, $r_{rb}=0.27$).

(H2) *Adding saliency maps improved performance in some areas depending on the task* There were no significant differences supporting our second hypothesis H2 which predicted that adding saliency maps will improve participants' performance in both tasks. Nevertheless, we report two positive effects of saliency maps that are only marginally⁹ significant and which might guide future research in this area. For the agent comparison task, we find that the saliency maps only improved performance when added to likelihood-based summaries (*L*: mean=0.54, 95% CI=[0.45, 0.64], *L+S*: mean=0.65, 95% CI=[0.54, 0.75], Mann-Whitney test $U=390.5$, $p=0.062$, $r_{rb}=0.21$). Fig. 8(a) shows that the saliency maps did not help participants identify the most important objects in the retrospection task. However, the summative content analysis of participants' textual descriptions of the agents' strategies, shown in Fig. 10, indicates that saliency maps helped participants to correctly describe how the agents use those objects. The descriptions of the agents' strategies written by participants in condition *H+S* received a higher score than the ones by participants in condition *H* (*H*: mean=1.50, 95% CI=[0.97, 2.0], *H+S*: mean=2.13, 95% CI=[1.55, 2.71], Mann-Whitney test $U=400$, $p=0.088$, $r_{rb}=0.195$).

(H3 + H4) *The effect of the summary generation method was greater than that of adding saliency maps* We hypothesized that the summary generation method will affect the performance of participants more than the addition of saliency maps in the agent comparison task (H3) and that the saliency maps will have a greater effect than the summary method in the retrospection task (H4). The study results support H3: we found that participants shown HIGHLIGHTS-DIV summaries sig-

⁸ Here 95% CI is the 95% confidence interval and r_{rb} is Rank biserial correlation.

⁹ In accordance with convention (Vogt et al. [73]), we use *marginally significant* to describe $0.05 \leq p < 0.1$.




Fig. 10. Participants' total score for their textual descriptions of the agents strategy during the agent comparison task (summed over all three agents). The scoring function is described in Section 6.3. The descriptions of participants in the HIGHLIGHTS-DIV conditions *H* and *H+S* received a higher score than those of participants in the likelihood-based conditions. The addition of saliency maps (*H+S*) slightly improved this effect further.

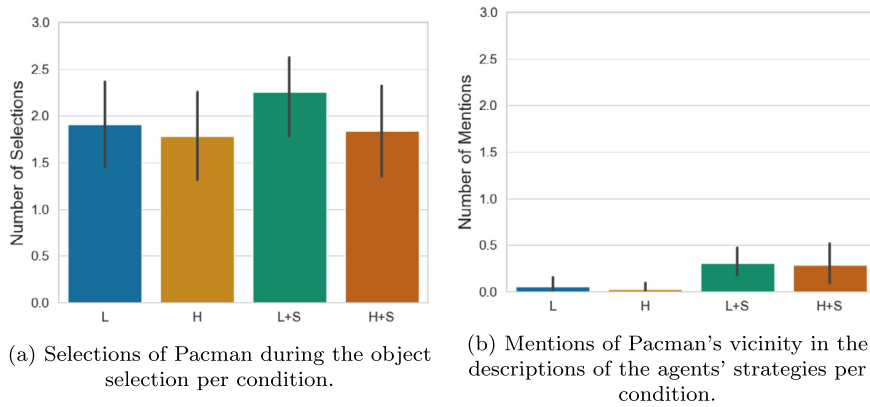


Fig. 11. The average number of times that participants correctly selected Pacman during the object selection (a), or referred to its vicinity in their textual descriptions (b) of the agents' strategies (sum over all three agents). The results indicate that saliency maps help the participants to identify what information the agents use.

nificantly outperformed participants shown likelihood-based summaries in the agent comparison task while adding saliency maps only improved performance for the likelihood-based summaries, and to a lesser extent.

For selecting the most important objects for the agent's strategy in the retrospection task, the addition of saliency maps did not improve performance, while HIGHLIGHTS-DIV summaries did improve performance compared to the likelihood-based summaries. Therefore we reject H4, even though the results shown in Fig. 10 indicate that saliency maps improved the textual descriptions of the agent's strategy written by participants in *H+S* compared to *H*.

In line with Hypothesis H4.1, Fig. 11 indicates that the improvement of the descriptions of the agents' strategies mainly stems from participants in the saliency groups *L+S* and *H+S* identifying that the agent mostly paid attention to the vicinity of Pacman. This effect was not as strong in the object selection question, since it did not capture the participants' reasoning.

Sub-Hypothesis H4.2 stated that strategy summarization would help participants identify the goals of the agents. The results shown in Fig. 12 support this Hypothesis, since participants in the HIGHLIGHTS-DIV conditions *H* and *H+S* identified the correct goals of the agent more often.

Participants with AI expertise benefited more from saliency maps To investigate the effect of AI expertise, we looked at participants with advanced AI experience. In total, 19 participants stated that they completed at least one course on AI or did research in AI (the latter being more common). These 19 participants were divided among conditions as follows: 5 per condition, except for condition *L+S*, which had only 4 (see Appendix A for more information). For the retrospection task (Fig. 13 a)), their results were in line with the results of the general population, albeit a little bit higher in condition *L*. For the agent comparison task (Fig. 13 b)), two participants did not watch the videos leaving 17 participants with advanced AI experience (4 per condition except for 5 in the *H* condition). Compared to the general population (Fig. 8 b)), the participants with AI expertise performed better in conditions *L+S* (mean = 2.25) and *H+S* (mean = 2.5) but worse in condition *H* (mean = 1.6). Moreover, the participants with AI expertise actually performed better in the saliency map conditions *L+S* and *H+S* than in the conditions without saliency maps (*L* and *H*). Since these observations were not part of our initial hypotheses and we do not have a sufficient number of participants with AI expertise, we did not investigate them further.




Fig. 12. The number of times that participants identified the agent's specific goal in the object selection (a) and strategy description (b) components of the retrospection task. The results are in line with Hypothesis H4.2 that strategy summarization helps to identify the agents' goals.

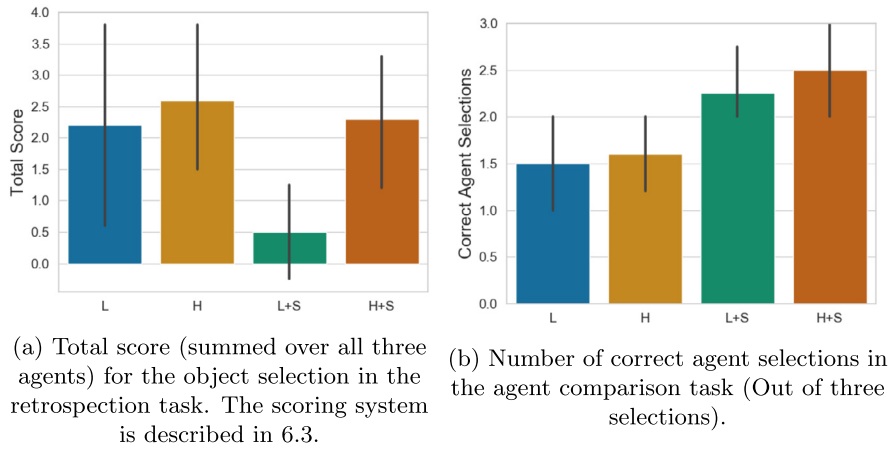


Fig. 13. The average per task performance of participants with AI expertise, by condition.

Participants' justifications Across all groups, most participants mainly based their justifications on the agents' gameplay (Fig. B.22). In the saliency conditions, most participants did not mention the saliency maps in their justifications. On average, less than one out of 3 justifications in *H+S* and in *L+S* referred to the green highlighting during the retrospection task. During the agent comparison task even fewer participants mentioned them (see Fig. B.23 for more details).

Another interesting point we found in participants' justifications during the retrospection task is that participants in condition *H* gave more unjustified explanations than any other condition (*H*: mean=0.66, compared to the second-highest condition *L+S*: mean=0.38). This is just an observation and did not repeat in the agent comparison task but it might be interesting to investigate further in the future. The values for all conditions can be seen in Fig. B.24.

Participants' confidence and viewing dynamics In addition to the main metrics used in our study, we further measured participants' confidence (in particular whether they were more confident when they answered correctly), and their viewing dynamics of the summaries (time and number of pauses). However, apart from a slight positive effect for the participants in condition *H*, there were no interesting differences in the three aforementioned variables (see Fig. B.19 to B.21 and Appendix B for additional details).

8. Discussion & future work

With the increasing use of RL agents in high-stakes domains, there is a growing need in developing and understanding methods for describing the behavior of these agents to their human users. In this paper, we explored the combination of global information describing agent behavior, in the form of strategy summaries, with local information in the form of saliency maps. To this end, we augmented HIGHLIGHTS-DIV [5] summaries, which select important and diverse states (adapted to DQN agents), with saliency maps generated using the LRP-argmax algorithm [35].

We implemented the combined approach in the Atari Pacman environment and evaluated the separate and joint benefits of showing users global and local information about the agent. We used two types of tasks: a retrospection task about the agent's strategy and a agent comparison task.

Strategy summarization The results of this study reinforce our prior findings [5] showing that summaries generated by HIGHLIGHTS-DIV lead to significantly improved performance of participants in the agent comparison task compared to likelihood-based summaries, and show that this result generalizes to RL agents based on neural networks. Furthermore, they show that HIGHLIGHTS-DIV summaries were more useful for analyzing agent strategies and were preferred by participants. Overall, in our study, the choice of states that are shown to participants was more important than the inclusion of local explanations in the form of saliency maps.

Limitations of saliency maps With respect to the addition of saliency maps, we found mixed results. In contrast to previous studies about saliency maps for image classification tasks, which found weak positive effects for saliency maps [4,62], there were no significant differences between the saliency and non-saliency conditions in our study. When examining participants' answer justifications, we observed that most participants did not mention utilizing the saliency maps, which may provide a partial explanation to their lack of contribution to participants' performance. Especially in the agent comparison task, participants seldom mentioned the saliency maps even though there was a marginally significant difference between the performance of participants in condition *L* and in condition *L+S*. Participants' comments also reflect their dissatisfaction with saliency maps, e.g., "I do not believe that the green highlighting was useful or relevant" and "The green highlights didn't seem to help much". This suggests that saliency maps in their current form may not be accessible enough to the average user.

Based on the comments from the participants and in-depth feedback we received in pilot studies, we note some possible accessibility barriers. First, when saliency maps are shown as part of a video, it may be difficult for users to keep track of the agent's attention, compared to displays of static saliency maps, as done in previous user studies [62,9,4]. For instance, one participant reported that "[i]t wasn't so easy to see the green area, it needed to be bigger or more prominent to be of more use." We tried to take measures against this by using a selective saliency map generation algorithm (LRP-argmax) and interpolating between selected saliency maps to reduce the amount of information, as well as allowing participants to pause the video at any time. However, this does not seem to be enough.

Second, participants were not accustomed to interpreting saliency maps, which can be non-intuitive to non-experts. One participant even commented that "[he/she] feel[s] as though this came with somewhat of a learning curve". In our pilot studies, we noticed that people who were familiar with reinforcement learning or deep learning could more easily interpret saliency maps than those who were not. For example, some participants said that they thought the agent was good when its attention was spread to different areas because they inferred it considered more information, while in fact, the agent was attending to different regions because it did not yet learn what the important information is. Similarly, one study participant commented: "...I don't know if I would prefer an AI that 'looked' around more at the board, or focused more in a small area to accomplish a task". Similar effects can be observed when examining the participants with AI expertise in our main study. Participants with AI expertise in the saliency conditions *L+S* and *H+S* received higher scores in the agent comparison task than the general population, while the participants with AI expertise in the *H* and *L* conditions received comparable or even lower scores than the general population. Even though this is only a small group, it further indicates that saliency maps, in the form of raw heat maps, might be better suited for debugging purposes than for actual end-users. It is possible that prior studies, which used raw saliency maps for interpreting image classification [4,62], did not encounter this problem due to the more intuitive nature of the task. Interpreting a visual highlighting for image classification only requires identifying objects that contributed to the classification, while in RL there is an added layer of complexity as interpretation also requires making inferences regarding how the highlighted regions affect the agent's long-term sequential decision-making policy.

Finally, while the sanity checks reported in Section 3.4 showed that our saliency maps do analyze what the network learned, they were also found to be indifferent to specific actions. Since prior studies have shown that users find class discriminatory explanations more useful for understanding agents' decisions [26,45,15], the lack of discrimination between certain actions can be detrimental to the usefulness of saliency maps.

Potential of saliency maps Regarding the potential of saliency maps, we made encouraging observations. Even though saliency maps did not significantly increase participants' scores in the simple object selection part of the retrospection task, they did result in improved scores in the textual strategy description. The difference between our HIGHLIGHTS-DIV conditions *H+S* and *H* is similar to the one observed by Anderson et al. [9] ($p=0.086$ compared to our $p=0.088$), who also evaluated participants' mental models for RL agents utilizing a strategy description task. The poor result of condition *L+S* can be explained by the fact that Anderson et al. manually chose meaningful states, which we only did with our global explanation method in the HIGHLIGHTS-DIV conditions.

A possible reason for the difference between the object selection and the strategy description sub-tasks is the higher complexity of strategy description. It requires participants to not only identify the correct objects but also to describe how they are used. Under this assumption, the increased performance of participants in condition *H+S* suggests that saliency maps were useful for putting the objects in the correct context. For example, participants' textual descriptions showed that, while the non-saliency groups know that Pacman is important (most likely based on the fact that it is important for them as players), they did not identify it as a central source of information for the agent.

Second, we observed in the agent comparison task that saliency maps alone improved participants' ability to place appropriate trust into different agents when comparing conditions *L* and *L+S*. The, performance in condition *L+S* was comparable to the performance of participants in the HIGHLIGHTS-DIV conditions, *H* and *H+S*. This indicates that there is valuable in-

formation for this kind of task within saliency maps. The lacking improvement of condition $H+S$ compared to H might be explained by the accessibility issues of saliency maps mentioned earlier. When presented with strategy summaries, participants may have had less reason to rely on the non-intuitive saliency maps.

Combination of local and global explanations It is important to note that the positive effects of saliency maps in the retrospection task were only visible in the HIGHLIGHTS-DIV conditions H and $H+S$, reinforcing our claim that the choice of states is crucial for explaining RL agents. Therefore, even if the limitations of saliency maps mentioned above are addressed, the potential benefits might only be visible and likely reinforced by a combination with strategy summarization techniques. We note that studies that evaluate local explanations typically implicitly make a global decision about which states to present local explanations for [9,46]. Our results suggest that this implicit choice may have a substantial impact on participants' understanding of agent behavior.

In the retrospection task, we observed that local explanations in the form of saliency maps were useful for identifying what objects the agent attends to (see Fig. 11), while strategy summaries were more useful for identifying the agent's goals (see Fig. 12). This was reflected by participants' utterances such as: "The agent seemed to be paying attention to the area directly in front of it and partly to the areas directly to each side." and "Pacman wanted those ghosts! His goal was to move as fast as he could towards them." and suggests that the two approaches are indeed complementary. The local saliency maps contribute to users' understanding of the agents' *attention*, as they reflect the information the agent attends to, while strategy summaries contribute to users' understanding of the agent's *intentions*, as they reflect how the agent acts.

Taken together, our results suggest that there is potential for a combined explanation framework in the future if the accessibility issues of saliency maps are addressed.

Study limitations Our study has several limitations. First, we used a single domain in our user study. However, other recent work has used strategy summaries similar in spirit to HIGHLIGHTS-DIV in another domain [63] and several works have used saliency maps in other domains (e.g., several Atari games including Pong and Space invaders were used by Greydanus et al. [27]).

Second, while our combined explanation approach is easily adaptable to other global explanation methods which choose an informative subset of states, and local methods that highlight relevant information in those states, our study only explored one combination of a particular global explanation method and a particular local explanation method. We chose the HIGHLIGHTS-DIV summary method since strategy summary approaches that are based on policy reconstruction require making various assumptions about people's computational models, and that these models differ depending on context [42]. We chose saliency maps as a local method both because it is visual and thus can be integrated with a visual summary, and also because other methods typically require additional models or assumptions (e.g., causal explanations [46] require a causal graph of the domain). The specific choice of the LRP-argmax algorithm was motivated by its selectivity, which reduces the amount of information that participants have to process. The accessibility problems of saliency maps we identified were mainly related to the presentation of the information. This indicates that simply highlighting how relevant parts of the input are for the prediction of an agent is insufficient even when based on other saliency map algorithms.

Finally, it is not yet possible to obtain actual ground truth about an agent's reasoning for the retrospection task. To approximate this, we trained three agents with different optimization goals (i.e., reward functions) and checked whether the participants were able to differentiate between them, similar to [63]. We note that it is possible that the agents do not pay attention to the objects, which correspond to their goals (e.g., the power pills), within the visual input in a way in which a human would. In fact, our results indicate that the agents are not paying direct attention to their goals, since the saliency maps did not help users identify the differences in the goals. Creating agent variations in ways that do not depend on different optimization goals might be better suited to show the potential benefits of saliency maps.

Future work One of the main contributions of our results about the complementary of global and local XAI methods is that it can inform the future design of XAI systems that combine such information. For instance, we plan to use our results as a first step towards an interactive explanation system for RL agents. Based on the generally good results of strategy summaries, we will start by showing those summaries to the users. This should already give the users a relatively good understanding with respect to the agent's *intentions*. However, in order not to overwhelm the users with information, we propose an interactive design where the saliency maps are not always present. Instead, the user can ask for additional information about what the agent is paying *attention* to. While this reduces users' cognitive load, it does not yet address the users' problems with interpreting the saliency maps. To this end, we plan to augment the visual information with textual explanations that help users interpret the information correctly, similar to how Rabold et al. [55] did with LIME explanations. Furthermore, we plan to build on our previous work [77] and explore the presentation of those textual explanations through virtual agents.

Finally, to verify that our results generalize beyond simulated environments, we would like to conduct user studies in real-world domains such as healthcare. Explainability is crucial in AI systems deployed in the medical field (e.g., pain classification [76]) since possible errors could lead to dire consequences. RL methods face additional challenges and requirements in the healthcare domain where random exploration of the state space is not possible and evaluation is challenging [24,25], making explanation methods even more important. In recent work, we have begun exploring the use of strategy summaries in healthcare using an HIV simulator [42], and intend to further explore this direction.

9. Conclusion

This work is a first step toward the development of combined explanation methods for reinforcement learning (RL) agents that provide users with both global information regarding the agent's strategy, as well as local information regarding its decision-making in specific world-states. To this end, we present a joint global and local explanation method, building on our prior work on strategy summaries (HIGHLIGHTS-DIV) and on generating saliency maps for deep RL agents (LRP-argmax). This method is easily adaptable to other global and local algorithms.

To evaluate this combined global and local explanation method, as well as the contribution of each explanation type, we conducted a user study. Hereby, we examined participants' mental models through a retrospection task and used an agent comparison task to investigate whether their trust was appropriate given agents' capabilities.

Regarding the usefulness of *global strategy summaries*, our results show that HIGHLIGHTS-DIV summaries (1) help to establish appropriate trust in agents based on neural networks (extending prior results about classic RL agents [5]) and (2) improve participants' mental models of those agents.

The evaluation of *local explanations* in the form of LRP saliency maps reveals strengths as well as weaknesses. On the one hand, our analysis shows that reinforcement learning comes with additional usability challenges not present in previously evaluated image classification tasks. First, presenting saliency maps on videos instead of static images [9,4] overwhelms users with a lot of information in a short amount of time and increases the risk of overlooking crucial information. Second, compared to more intuitive image classification tasks [4,62], the average user lacks the experience to correctly infer how the highlighted regions affect the agent's long-term sequential decision-making.

On the other hand, the results indicate that saliency maps have the potential to (1) extend users' mental models beyond strategy summaries by providing insight into what information the agent used and (2) improve users' ability to choose the better agent even with likelihood-based summaries.

Taken together, the results support a combination of local and global explanations, since participants in the combined explanation condition received the highest scores during our survey. However, our evaluation suggests that simply highlighting pixels that are relevant for the agent's decision is insufficient for RL agents and that more work is needed to increase the accessibility of saliency maps.

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.

Acknowledgements

This work was partially funded by research grant #2185/20 from the Israeli Science Foundation and the Deutsche Forschungsgemeinschaft (DFG) under project DEEP (Grant Number 392401413).

We thank Otto Grothe for his help with analyzing the participants' textual responses.

Appendix A. Participants demographics

In this section, we provide more details regarding participants' demographics. As Fig. A.14 shows, most participants were between 18 and 34 years old. There were no major differences in gender distribution between the four conditions (Fig. A.15).




Fig. A.14. The number of participants in each age group per condition. The bars show from left to right: "18-24", "25-34", "35-44", "45-54", "55-64" and "65 or older". The categories "17 or younger" and "do not want to specify" were never selected.




Fig. A.15. Number of female participants per condition.




Fig. A.16. The average attitude towards AI, rated on a 5 point Likert scale.




Fig. A.17. The Pacman experience across all conditions where the bars depict when the participants played Pacman the last time. From left to right the bars represent: “never”, “more than 5 years ago”, “less than 5 years ago” and “less than 1 year ago”.

We verified that participants in different conditions did not differ much in their AI experience and views and in their Pacman experience. To this end, we asked them when they played Pacman for the last time (1 = “never”, 2 = “more than 5 years ago”, 3 = “less than 5 years ago”, 4 = “less than 1 year ago”). Across all four conditions, the median group was 2: “I played Pacman more than 5 years ago”. A comparison is shown in Fig. A.17.

For the AI experience we adapted a description of AI from Zhang et al. [80] and Russel [58] to “The following questions ask about Artificial Intelligence (AI). Colloquially, the term ‘artificial intelligence’ is often used to describe machines (or

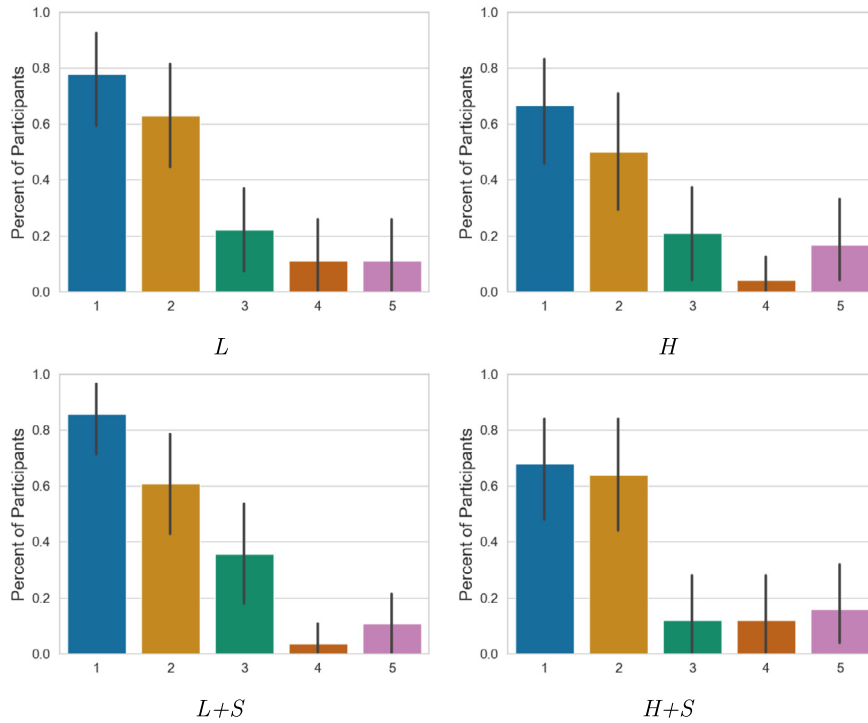


Fig. A.18. Distribution of the chosen AI experience items for each condition. The x-axis depicts the items described above.

computers) that mimic ‘cognitive’ functions that humans associate with the human mind, such as ‘learning’ and ‘problem solving’. AI agents are already able to perform some complex tasks better than the median human (today). Examples for such intelligent agents are search engines, chatbots, chess bots, and voice assistants.”

After that, every participant who stated to have AI experience (104 across all conditions) had to select one or more of the following items:

- 1: I know AI from the media.
- 2: I use AI technology in my private life.
- 3: I use AI technology in my work.
- 4: I took at least one AI related course.
- 5: I do research on AI related topics.
- Other:

The last free form option was used exactly once and read “work on MTurk”. The distribution of the other items for each condition is shown in Fig. A.18.

To measure the participants’ attitude towards AI we adapted a question from Zhang et al. [80] and asked them to rate their answer to the question “Suppose that AI agents would achieve high-level performance in more areas one day. How positive or negative do you expect the overall impact of such AI agents to be on humanity in the long run?” on a 5-point Likert scale from “Extremely negative” to “Extremely positive”. The results are shown in Fig. A.16.

Appendix B. Supplementary results

In this section, we present additional information about the results of the study that goes beyond the main hypotheses we explored and described in the paper.

Confidence, time and pauses To investigate whether participants were confident in their decisions, they had to rate the confidence in each of their selections (item selection in the retrospection task and agent selection in the agent comparison task) on a 7 point Likert scale. The results across each task are shown in Fig. B.19.

To evaluate whether participants were especially diligent or effective during the tasks, we measured the time that each participant stayed on each page of the survey and calculated the average time per task (each task consists of three pages). Furthermore, we kept track of each time a video was paused, as described in Section 6.3. The average completion times of




Fig. B.19. The average confidence that participants in each condition had in their answers during each task.

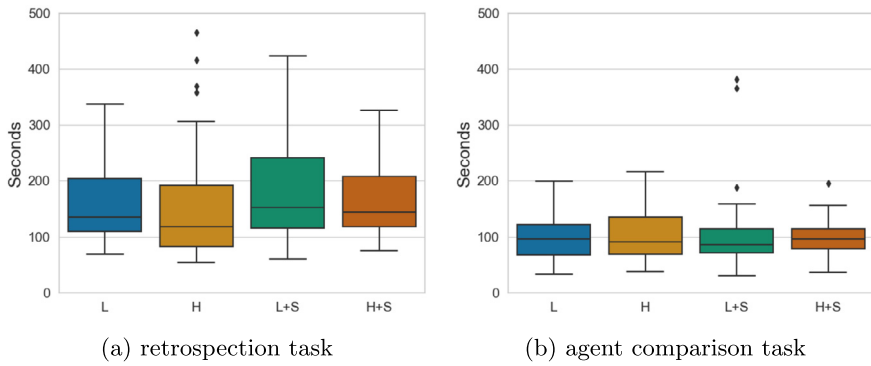


Fig. B.20. The average time taken by participants in each condition per agent analysis (a) and comparison of agent pairs (b).

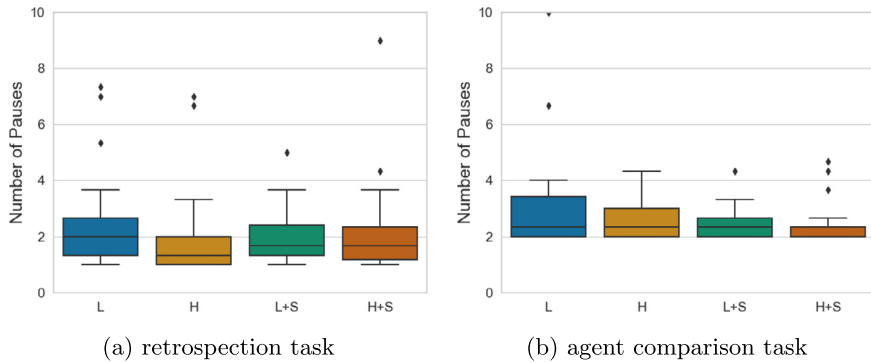


Fig. B.21. The average number of times that participants in each condition paused the videos during each agent analysis (a) and comparison of agent pairs (b).

participants and the average number of pauses are shown in Fig. B.20 and B.21, respectively (shown in boxplots due to the presence of several outliers that strongly affect the mean values).

Fig. B.19 (a) shows that participants in condition *H* were slightly more confident on average in their analysis of the agents. This is also reflected by the lesser amount of time per analysis (Fig. B.20 (a)) and pauses (Fig. B.21 (a)). Apart from this, there are no obvious differences between the average confidence, time, and pause values for each task (Fig. B.19 to B.21).

Participants' justifications As described in Section 6.3, an independent coder identified different concepts inside the participants' justifications. Fig. B.23 shows the average number of mentions of *gameplay* and of *saliency maps* in the different tasks, across the different conditions. As discussed in Section 7, most participants mainly based their justifications on the agents'




Fig. B.22. Comparison of how often the participants referenced the agents' **gameplay** in their justifications for their answers.




Fig. B.23. Comparison of how often the participants referenced the green highlighting of the LRP-argmax **saliency maps** in their justifications for their answers.




Fig. B.24. Comparison of how often the participants justifications contained **unjustified** arguments.

gameplay (Fig. B.22). In the saliency conditions, participants seldom mention the saliency maps in their justifications (see Fig. B.23). Finally, Fig. B.24 shows that participants in condition *H* gave more unjustified explanations in the retrospection task. However, this observation did not repeat in the agent comparison task.

Appendix C. Evaluation of the retrospection task

As described in Section 6.3, we evaluated participants' scores in the object selection part of the retrospection task with a simple scoring function based on predefined answers by two of the authors involved in the training of the agents. Hereby, we assign a score of 1 to each object that is connected to the agents' specific goal and their source of information (Pacman's position for all agents), -1 for each object that was not related to the agents' reward function and -0.5 to objects that were related to the reward but on which the agent did not focus. The specific scores are shown in Fig. C.25.

selected object	<i>Power pill agent</i>	<i>Regular agent</i>	<i>Fear-ghosts agent</i>
“Pacman”	1	1	1
“normal pill”	−1	−0.5	−0.5
“power pill”	1	−0.5	−0.5
“normal ghost”	−1	−0.5	1
“blue ghost”	−1	1	1
“cherry”	−1	−0.5	−0.5

Fig. C.25. The specific scores that participants received for selecting each object during the retrospection task.

For the free form answers to the question “Please describe the strategy of the AI agent” an independent coder identified various not mutually exclusive concepts contained in the participants’ answers. We aggregated these concepts into the following 16 groups, where the coder used ‘G’ for ghosts, ‘PP’ for power pills and ‘NP’ for normal pills:

1. *eating power pills*: “eating PP”, “eating as many PP as possible”, “eat PP when ghosts are near”, “eat PP when ghosts are near”, “prioritizing PP”, “prioritizing PP to eat ghosts”, “prioritizing PP, but not eat ghosts”, “eat PP to get points”
2. *ignore power pills*: “do not care about PP”
3. *eat normal pills*: “eat NP to get points”, “eating NP”, “eating as many NP as possible”, “prioritizing NP”, “clearing the stage”
4. *ignore normal pills*: “do not care about NP”, “focus on areas without [sic] NP”
5. *avoid ghosts*: “avoiding G”, “avoiding G strongly”, “wait for G to go away”, “outmaneuvering G”, “hiding from G”, “mislead ghosts”, “avoids being eaten / caught”, “avoiding to lose / staying alive”, “stays away from danger”
6. *move towards ghosts*: “being close to G”, “trying to eat G NON blue”, “(easily) caught by G”, “easily caught by G”
7. *ignore ghosts*: “do not care about G”
8. *making ghosts blue*: “making G blue”
9. *eat blue ghosts*: “being close to blue G”, “eating as many G as possible”, “eat blue G to get points”, “chasing/going for G”, “eating the blue G”, “eating to jail many G” (jailing since the ghosts move back to jail after being eaten), “prioritizing PP to eat ghosts”
10. *avoid blue ghosts*: “avoiding blue G”
11. *ignore blue ghosts*: “do not care about blue G”, “prioritizing PP, but not eat ghosts”
12. *eat cherry*: “prioritizing cherry”, “eat cherry to get points”, “going for cherry”, “eating cherry”
13. *ignore cherry*: “do not care about cherry”
14. *random movement*: “moving randomly”, “move all over map”, “switching directions /back&forth”, “not moving / being stuck”, “sticking to walls / outside”, “confused”, “without strategy /random”, “not planning ahead”, “switching directions”
15. *focus on Pacman*: “focus on PM”, “focus on what’s in front of/around PM”, “stuck to itself”
16. *staying in corners*: “staying in corners”

These groups are used to define a simple scoring function. Depending on the agent, each group could either be positive, neutral, or negative. Positive groups contain concepts that are in line with the predefined descriptions of the agents’ strategies by two of the authors involved in the training. Neutral groups consist of correct observations, which are byproducts of the agent’s strategy, and negative concepts go against the agent’s strategy. Each positive group contained in an answer increased the participant’s score by 1 and each negative group decreased the score by −1. Here, we define a group to be “contained in an answer” if at least one concept of this group was included in the answer. Neutral groups did not affect the score.

Power pill agent:

- *positive*: “eat power pill”, “ignore normal pill”, “ignore ghosts”, “ignore blue ghost”, “ignore cherry”, “focus on Pacman”, “staying in corners”
- *neutral*: “eat normal pill”, “making ghosts blue”

Regular agent:

- *positive*: “ignore cherry”, “focus on Pacman”, “making ghosts blue”, “eat blue ghost”
- *neutral*: “eat normal pill”, “eat power pill”, “ignore ghosts”

Fear-ghosts agent:

- *positive*: “avoid ghost”, “focus on Pacman”, “making ghosts blue”, “eat blue ghost”, “ignore cherry”
- *neutral*: “eat normal pill”, “eat power pill”

Appendix D. Questionnaire

In this section, we provide the complete questionnaire used in the study.
On the first page the participants were asked to provide personal information:

Personal information

What is your age?

Choose one of the following answers

☐ 17 or younger
☐ 18 - 24
☐ 25 - 34
☐ 35 - 44
☐ 45 - 54
☐ 55 - 64
☐ 65 or older
☐ I prefer not to specify

To which gender identity do you most identify?

Choose one of the following answers

☐ Male
☐ Female
☐ I prefer not to answer.
☐ Other:

Do you have a colour vision impairment?

☒ Yes
 ☐ No

Did you play Pacman before?

Choose one of the following answers

☐ No
☐ The last time I played Pacman was less than 1 year ago.
☐ The last time I played Pacman was less than 5 years ago.
☐ The last time I played Pacman was more than 5 years ago.

The following questions ask about Artificial Intelligence (AI). Colloquially, the term "artificial intelligence" is often used to describe machines (or computers) that mimic "cognitive" functions that humans associate with the human mind, such as "learning" and "problem solving". AI agents are already able to perform some complex tasks better than the median human (today). Examples for such intelligent agents are search engines, chatbots, chessbots and voice assistants.

Do you have experience with AI (Artificial Intelligence)?

☒ Yes
 ☐ No

Do you have experience with AI (Artificial Intelligence)?

☐ Yes ☐ No

What kind of AI experience do you have?

Check all that apply

☐ I know AI from the media.

☐ I use AI technology in my private life.

☐ I use AI technology in my work.

☐ I took at least one AI related course.

☐ I do research on AI related topics.

☐ Other:

Suppose that AI agents would achieve high-level performance in more areas one day.

	1: Extremely negative	2	3	4	5: Extremely positive	I don't know
How positive or negative do you expect the overall impact of such AI agents to be on humanity in the long run?	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

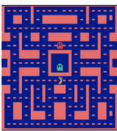
Information about Pacman:

Instructions

In this survey, you will be shown summaries of Pacman games. These videos will show you parts of games played by various AI agents trained to play Pacman. Later you will be asked to complete 6 tasks based on those videos. You will earn a 10 cent bonus for each task you solved correctly.

Please carefully read the following description of Pacman. There will be a short quiz to make sure you understood the information correctly.

A Pacman game takes place in a Labyrinth:



Pacman's goal is to eat as many "pills" (pink rectangles) as possible to earn points, while escaping the ghosts.

Pacman receives 10 points for each pill, and dies when it's eaten by a ghost.

In addition to the regular pellets, there are also four special "power pills" (large pink rectangles). Pacman receives 50 points for eating a power pill.

Eating a power pill also gives Pacman limited time during which the ghosts become blue and Pacman can eat the ghosts. Pacman receives 200, 400, 800, 1600 points for each ghost it eats successively.

After a ghost is eaten, its eyes move back to the ghost box in the middle of the screen and the ghost restarts from there in his normal form.

At random intervals cherries spawn and move through the labyrinth.

If PacMan eats the cherry he gets 100 points.

Pacman starts the game with three lives and loses one life each time it gets eaten by a ghost.

Please use this link to play the game and get familiar with it (While the objects look different in that version the rules are the same as described above.):

[Pacman](#)

This quiz tests whether the participants understood the information about Pacman. Participants were sent back to the previous page if they got an answer wrong.

Quiz Pacman

Please fill out this quiz to show that you understood the information.

*Eating  gives Pacman:


Choose one of the following answers

- ☐ No points
- ☐ 1 point
- ☐ 10 points

*What is shown in this image: 

Choose one of the following answers

- ☐ A ghost.
- ☐ A blue ghost.
- ☐ Pacman.

*What is shown in this picture? 

Choose one of the following answers

- ☐ a piece of wall.
- ☐ a power pill.
- ☐ a regular pill.

*When do ghosts become blue?

Choose one of the following answers

- ☐ Every 2 minutes
- ☐ When Pacman eats a power pill
- ☐ When Pacman gets close to a ghost

*What happens when ghosts are blue?

Choose one of the following answers

- ☐ Pacman can eat them
- ☐ Ghosts move faster
- ☐ Ghosts move slower


Additional information about the provided explainable AI methods. The information about saliency maps was only displayed if the participant was in one of the saliency conditions.

Additional Information

In the following questions you will be shown videos that **summarize the behavior** of an AI agent that was trained to play Pacman. These videos contain **several scenes** from a longer session of gameplay and aim to show you, in a limited amount of time (about 20 seconds), **how the agent plays the game**.

To aid you in your tasks, the screens will be augmented with green heatmaps that show **how relevant a pixel was for the decision of the AI agent** that controls Pacman ("what the agent is looking at"). On some screens the green color may appear yellowish.

The **brighter the green** highlighting of a pixel is, the **more relevance** it had for the decision of the AI agent.



In this example, Pacman is paying a lot of attention on the area in front of itself.



This quiz tests whether the participants understood the information about the provided explainable AI methods. Participants were sent back to the previous page if they got an answer wrong.

Second Quiz

Please fill out this quiz to show that you understood the information.

★ A green area is:

① Choose one of the following answers

- ☐ relevant for the Pacman agent.
- ☐ relevant for your decision.
- ☐ relevant for the ghosts.

★ If the green highlighting of an area is very bright then:

① Choose one of the following answers

- ☐ the area is worth more points.
- ☐ the area is more relevant.
- ☐ the area is less relevant.

★ The videos in the following questions show:

① Choose one of the following answers


- ☐ scenes where the Pacman agent performed very good.
- ☐ a summary of how the Pacman agent behaves.
- ☐ scenes where the Pacman agent performed very bad.

This is the retrospection task that was repeated for each of the three agents in a randomized order:

Analyze the agent

In these three tasks you will be shown **summaries of different AI agents** trained to play Pacman. For each AI agent, you will be asked to **analyze the strategy of that particular AI agent**. To this end, you will be asked to describe the AI agents' strategy and to select the objects that were the most important for the strategy of that particular AI agent. You will receive a bonus of 10 cents for each agent that you analyze correctly.

The video shows a summary of typical behavior of a Pacman agent. Please watch the whole video.



★ Please briefly describe the strategy of the AI agent shown in the video above:

★ Based on this video, select the objects that you think were most important for the strategy of this particular AI agent.

Select a **maximum of 3** (it does not necessarily have to be 3) Objects.

🟢 Check all that apply

- ☐ Pacman 
- ☐ Normal pills 
- ☐ Power pills 
- ☐ Ghosts 
- ☐ Blue Ghosts 
- ☐ Cherrys 

★ How confident are you in your selection?

	Not at all confident						very confident
How confident are you that you chose the right objects?	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

★ Please briefly explain how you came to your selection:

After all three agents, the participants were asked for their satisfaction:

Explanation Satisfaction

	1 : I disagree strongly.	2	3	4	5 : I agree strongly.
From watching the videos of the AI agents, I got an idea of the agents' strategies.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The videos showing the AI agents play contain sufficient detail about the agents' behavior.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The videos showing the AI agents play contain irrelevant details.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The gameplay scenarios shown in the videos were useful for analyzing the agents' behavior	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The green highlighting in the videos was useful for analyzing the agents' behavior	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Do you have additional comments?

This is the agent comparison task that was repeated for each combination of the three agents in a randomized order:


Compare the agents

Below are videos showing parts of games played by **two different Pacman agents**: Agent A (left) and Agent B (right).


Based on these videos, choose which Pacman agent you'd like to play on your behalf. You will receive a bonus of 10 cents for each time you select the Pacman agent that achieves higher points on average.

Note: Press play to start a video. You can pause and rewatch the videos by pressing the button again. Please watch the whole videos.

Agent A



Agent B



Which of the Pacman agents do you choose to play on your behalf?

Choose one of the following answers

☐ Player A

☐ Player B

How confident are you that you chose the better agent?

	Not at all confident						very confident
How confident are you that you chose the better player?	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Please briefly explain your selection:

After all three comparisons, the participants were asked for their satisfaction again:

Explanation Satisfaction

*

	1 : I disagree strongly.	2	3	4	5 : I agree strongly.
From watching the videos of the AI agents, I got an idea of the agents' strategies.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The videos showing the AI agents play contain sufficient detail about the agents' behavior.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The videos showing the AI agents play contain irrelevant details.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The gameplay scenarios shown in the videos were useful for choosing the agent that performs better.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The green highlighting in the videos was useful for choosing the agent that performs better.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Do you have additional comments?

References

- [1] J. Adebayo, J. Gilmer, M. Muelly, I. Goodfellow, M. Hardt, B. Kim, Sanity checks for saliency maps, in: S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, R. Garnett (Eds.), *Advances in Neural Information Processing Systems*, vol. 31, Curran Associates, Inc., 2018, pp. 9505–9515, <http://papers.nips.cc/paper/8160-sanity-checks-for-saliency-maps.pdf>.
- [2] D. Aha, T. Darrell, M. Pazzani, D. Reid, C. Sammut, P. Stone, IJCAI-17 workshop on explainable AI (XAI), in: *IJCAI-17 Workshop on Explainable AI, XAI*, 2017.
- [3] M. Alber, S. Lapuschkin, P. Seegerer, M. Hägele, K.T. Schütt, G. Montavon, W. Samek, K.-R. Müller, S. Dähne, P.-J. Kindermans, iNNvestigate neural networks!, arXiv preprint, arXiv:1808.04260, 2018.
- [4] A. Alqaraawi, M. Schuessler, P. Weiß, E. Costanza, N. Berthouze, Evaluating saliency map explanations for convolutional neural networks: a user study, arXiv preprint, arXiv:2002.00772, 2020.
- [5] D. Amir, O. Amir, Highlights: summarizing agent behavior to people, in: *Proc. of the 17th International conference on Autonomous Agents and Multi-Agent Systems, AAMAS*, 2018.
- [6] O. Amir, F. Doshi-Velez, D. Sarne, Agent strategy summarization, in: *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems, International Foundation for Autonomous Agents and Multiagent Systems*, 2018, pp. 1203–1207.
- [7] O. Amir, F. Doshi-Velez, D. Sarne, Summarizing agent strategies, *Auton. Agents Multi-Agent Syst.* 33 (5) (2019) 628–644.
- [8] O. Amir, E. Kamar, A. Kolobov, B.J. Grosz, Interactive teaching strategies for agent training, in: *International Joint Conferences on Artificial Intelligence*, 2016.
- [9] A. Anderson, J. Dodge, A. Sadarangani, Z. Juozapaitis, E. Newman, J. Irvine, S. Chattopadhyay, A. Fern, M. Burnett, Explaining reinforcement learning to mere mortals: an empirical study, in: *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, International Joint Conferences on Artificial Intelligence Organization, July 2019, pp. 1328–1334.
- [10] S. Anjomshoa, A. Najjar, D. Calvaresi, K. Främling, Explainable agents and robots: results from a systematic literature review, in: *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS '19*, International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 2019, pp. 1078–1088.
- [11] A.B. Arrieta, N. Díaz-Rodríguez, J. Del Ser, A. Bannetot, S. Tabik, A. Barbado, S. García, S. Gil-López, D. Molina, R. Benjamins, R. Chatila, F. Herrera, Explainable artificial intelligence (XAI): concepts, taxonomies, opportunities and challenges toward responsible AI, *Inf. Fusion* 58 (2020) 82–115, <http://www.sciencedirect.com/science/article/pii/S1566253519308103>.
- [12] S. Bach, A. Binder, G. Montavon, F. Klauschen, K.-R. Müller, W. Samek, On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation, *PLoS ONE* 10 (7) (07 2015) 1–46, <https://doi.org/10.1371/journal.pone.0130140>.
- [13] M.G. Bellemare, Y. Naddaf, J. Veness, M. Bowling, The arcade learning environment: an evaluation platform for general agents, *J. Artif. Intell. Res.* 47 (Jun 2013) 253–279.
- [14] S. Booth, C. Muise, J. Shah, Evaluating the interpretability of the knowledge compilation map: communicating logical statements effectively, in: *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, International Joint Conferences on Artificial Intelligence Organization, 7 2019, pp. 5801–5807.
- [15] R.M.J. Byrne, Counterfactuals in explainable artificial intelligence (XAI): evidence from human reasoning, in: *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, International Joint Conferences on Artificial Intelligence Organization, July 2019, pp. 6276–6282.
- [16] M. Cashmore, A. Collins, B. Krarup, S. Krivic, D. Magazzeni, D. Smith, Towards explainable AI planning as a service, arXiv preprint, arXiv:1908.05059, 2019.
- [17] T. Chakraborti, S. Sreedharan, S. Grover, S. Kambhampati, Plan explanations as model reconciliation, in: *2019 14th ACM/IEEE International Conference on Human-Robot Interaction, HRI*, IEEE, 2019, pp. 258–266.
- [18] B. Chandrasekaran, M.C. Tanner, J.R. Josephson, Explaining control strategies in problem solving, *IEEE Intell. Syst.* 1 (1989) 9–15.
- [19] P. Dhariwal, C. Hesse, O. Klimov, A. Nichol, M. Plappert, A. Radford, J. Schulman, S. Sidor, Y. Wu, P. Zhokhov, OpenAI baselines, <https://github.com/openai/baselines>, 2017.
- [20] T. Dodson, N. Mattei, J. Goldsmith, A natural language argumentation interface for explanation generation in Markov decision processes, in: *Algorithmic Decision Theory*, 2011, pp. 42–55.

- [21] F. Doshi-Velez, B. Kim, A roadmap for a rigorous science of interpretability, arXiv preprint, arXiv:1702.08608, 2017.
- [22] M. Erwig, A. Fern, M. Murali, A. Koul, Explaining deep adaptive programs via reward decomposition, in: *IJCAI/ECAP Workshop on Explainable Artificial Intelligence*, 2018.
- [23] V. Goel, J. Weng, P. Poupart, Unsupervised video object segmentation for deep reinforcement learning, in: *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, Montréal, Canada, 2018*, pp. 5688–5699, <http://papers.nips.cc/paper/7811-unsupervised-video-object-segmentation-for-deep-reinforcement-learning>.
- [24] O. Gottesman, F. Johansson, M. Komorowski, A. Faisal, D. Sontag, F. Doshi-Velez, L.A. Celi, Guidelines for reinforcement learning in healthcare, *Nat. Med.* 25 (1) (2019) 16–18.
- [25] O. Gottesman, F. Johansson, J. Meier, J. Dent, D. Lee, S. Srinivasan, L. Zhang, Y. Ding, D. Wihl, X. Peng, et al., Evaluating reinforcement learning algorithms in observational health settings, arXiv preprint, arXiv:1805.12298, 2018.
- [26] O. Goudet, D. Kalainathan, P. Caillou, I. Guyon, D. Lopez-Paz, M. Sebag, *Learning Functional Causal Models with Generative Neural Networks*, Springer International Publishing, Cham, 2018, pp. 39–80.
- [27] S. Greydanus, A. Koul, J. Dodge, A. Fern, Visualizing and understanding Atari agents, in: *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, 2018*, pp. 1787–1796.
- [28] D. Gunning, Explainable artificial intelligence (XAI), Defense Advanced Research Projects Agency (DARPA), 2017, Web.
- [29] F.G. Halasz, T.P. Moran, Mental models and problem solving in using a calculator, in: *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*, 1983, pp. 212–216.
- [30] 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, ACM, 2017*, pp. 303–312.
- [31] R.R. Hoffman, S.T. Mueller, G. Klein, J. Litman, Metrics for explainable AI: challenges and prospects, arXiv preprint, arXiv:1812.04608, 2018.
- [32] H.-F. Hsieh, S.E. Shannon, Three approaches to qualitative content analysis, *Qual. Health Res.* 15 (9) (2005) 1277–1288.
- [33] S.H. Huang, K. Bhatia, P. Abbeel, A.D. Dragan, Establishing appropriate trust via critical states, in: *2018 IEEE/RISJ International Conference on Intelligent Robots and Systems, IROS, IEEE, 2018*, pp. 3929–3936.
- [34] S.H. Huang, D. Held, P. Abbeel, A.D. Dragan, Enabling robots to communicate their objectives, in: *Robotics: Science and Systems XIII, Massachusetts Institute of Technology, Cambridge, Massachusetts, USA, July 12–16, 2017*, pp. 1–18, <http://www.roboticsproceedings.org/rss13/p59.html>.
- [35] T. Huber, D. Schiller, E. André, Enhancing explainability of deep reinforcement learning through selective layer-wise relevance propagation, in: *Joint German/Austrian Conference on Artificial Intelligence, Künstliche Intelligenz, Springer, 2019*, pp. 188–202.
- [36] R. Iyer, Y. Li, H. Li, M. Lewis, R. Sundar, K.P. Sycara, Transparency and explanation in deep reinforcement learning neural networks, in: *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society, AIES 2018, New Orleans, LA, USA, 2018*, pp. 144–150.
- [37] O. Khan, P. Poupart, J. Black, L. Sucar, E. Morales, J. Hoey, Automatically generated explanations for Markov decision processes, in: *Decision Theory Models for Applications in AI: Concepts and Solutions*, 2011, pp. 144–163.
- [38] O.Z. Khan, P. Poupart, J.P. Black, Minimal sufficient explanations for factored Markov decision processes, in: *ICAPS*, 2009.
- [39] B. Kim, R. Khanna, O.O. Koyejo, Examples are not enough, learn to criticize! criticism for interpretability, in: *Advances in Neural Information Processing Systems*, 2016, pp. 2280–2288.
- [40] B. Krarup, M. Cashmore, D. Magazzeni, T. Miller, Model-based contrastive explanations for explainable planning, in: *Proceedings of the ICAPS 2019 Workshop on Explainable Planning, XAIP, 2019*.
- [41] A. Kulkarni, Y. Zha, T. Chakraborti, S.G. Vadlamudi, Y. Zhang, S. Kambhampati, Explicable planning as minimizing distance from expected behavior, in: *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems, International Foundation for Autonomous Agents and Multiagent Systems*, 2019, pp. 2075–2077.
- [42] 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, IJCAI-19, International Joint Conferences on Artificial Intelligence Organization, July 2019*, pp. 1401–1407.
- [43] S. Lapuschkin, S. Wäldchen, A. Binder, G. Montavon, W. Samek, K.-R. Müller, Unmasking Clever Hans predictors and assessing what machines really learn, *Nat. Commun.* 10 (1) (2019) 1096.
- [44] J.D. Lee, K.A. See, Trust in automation: designing for appropriate reliance, *Hum. Factors* 46 (1) (2004) 50–80.
- [45] D. Lopez-Paz, R. Nishihara, S. Chintala, B. Scholkopf, L. Bottou, Discovering causal signals in images, in: *The IEEE Conference on Computer Vision and Pattern Recognition, CVPR, July 2017*.
- [46] P. Madumal, T. Miller, L. Sonenberg, F. Vetere, Explainable reinforcement learning through a causal lens, arXiv preprint, arXiv:1905.10958, 2019.
- [47] P.E. McKnight, J. Najab, Mann-Whitney U test, in: *The Corsini Encyclopedia of Psychology*, 2010, p. 1.
- [48] T. Miller, Explanation in artificial intelligence: insights from the social sciences, *Artif. Intell.* 267 (2018) 1–38.
- [49] T. Miller, P. Howe, L. Sonenberg, Explainable AI: beware of inmates running the asylum, in: *IJCAI-17 Workshop on Explainable AI (XAI)*, vol. 36, 2017.
- [50] V. Mnih, K. Kavukcuoglu, D. Silver, A.A. Rusu, J. Veness, M.G. Bellemare, A. Graves, M. Riedmiller, A.K. Fidjeland, G. Ostrovski, et al., Human-level control through deep reinforcement learning, *Nature* 518 (7540) (2015) 529–533.
- [51] C. Molnar, Interpretable machine learning: a guide for making black box models explainable, <https://christophm.github.io/interpretable-ml-book>, 2019.
- [52] G. Montavon, W. Samek, K. Müller, Methods for interpreting and understanding deep neural networks, *Digit. Signal Process.* 73 (2018) 1–15, <https://doi.org/10.1016/j.dsp.2017.10.011>.
- [53] K.R. Mopuri, U. Garg, R.V. Babu, CNN fixations: an unraveling approach to visualize the discriminative image regions, *IEEE Trans. Image Process.* 28 (5) (2019) 2116–2125.
- [54] D.A. Norman, Some observations on mental models, in: *Mental Models*, Psychology Press, 2014, pp. 15–22.
- [55] J. Rabold, H. Deininger, M. Siebers, U. Schmid, Enriching visual with verbal explanations for relational concepts - combining LIME with Aleph, in: *PKDD/ECML Workshops*, 2019.
- [56] A.S. Rao, M.P. Georgeff, et al., BDI agents: from theory to practice, in: *ICMAS*, vol. 95, 1995, pp. 312–319.
- [57] M.T. Ribeiro, S. Singh, C. Guestrin, Why should I trust you? Explaining the predictions of any classifier, in: *Proc. of ACM International Conference on Knowledge Discovery and Data Mining, ACM, 2016*, pp. 1135–1144.
- [58] S. Russell, P. Norvig, *Artificial Intelligence: A Modern Approach Global Edition*, Pearson, 2016.
- [59] H. Rutjes, M. Willemsen, W. Ijsselstein, Considerations on explainable AI and users' mental models, in: *Where Is the Human? Bridging the Gap Between AI and HCI*, vol. 5, Association for Computing Machinery, Inc, United States, 2019.
- [60] K. Schulz, L. Sixt, F. Tombari, T. Landgraf, Restricting the flow: information bottlenecks for attribution, *CoRR*, arXiv:2001.00396, <http://arxiv.org/abs/2001.00396>, 2020.
- [61] B. Seegerbarth, F. Müller, B. Schattenberg, S. Biundo, Making hybrid plans more clear to human users - a formal approach for generating sound explanations, in: *Twenty-Second International Conference on Automated Planning and Scheduling*, 2012.
- [62] R.R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, D. Batra, Grad-CAM: visual explanations from deep networks via gradient-based localization, *Int. J. Comput. Vis.* 128 (2) (2020) 336–359, <https://doi.org/10.1007/s11263-019-01228-7>.
- [63] P. Sequeira, E. Yeh, M.T. Gervasio, Interestingness elements for explainable reinforcement learning through introspection, in: *IUI Workshops*, 2019, p. 7.

- [64] K. Simonyan, A. Vedaldi, A. Zisserman, Deep inside convolutional networks: visualising image classification models and saliency maps, CoRR, arXiv:1312.6034.
- [65] L. Sixt, M. Granz, T. Landgraf, When explanations lie: why modified BP attribution fails, CoRR, arXiv:1912.09818, 2019.
- [66] J.T. Springenberg, A. Dosovitskiy, T. Brox, M.A. Riedmiller, Striving for simplicity: the all convolutional net, CoRR, arXiv:1412.6806, 2014.
- [67] P. Stone, R. Brooks, E. Brynjolfsson, R. Calo, O. Etzioni, G. Hager, J. Hirschberg, S. Kalyanakrishnan, E. Kamar, S. Kraus, K. Leyton-Brown, D. Parkes, P. William, S. AnnaLee, S. Julie, T. Milind, T. Astro, Artificial Intelligence and Life in 2030. One Hundred Year Study on Artificial Intelligence: Report of the 2015-2016 Study Panel, 2016.
- [68] M. Sundararajan, A. Taly, Q. Yan, Axiomatic attribution for deep networks, in: Proceedings of the 34th International Conference on Machine Learning, vol. 70, JMLR.org, 2017, pp. 3319–3328.
- [69] W.R. Swartout, Xplain: a system for creating and explaining expert consulting programs, Artif. Intell. 21 (3) (1983) 285–325.
- [70] 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 (1) (2014).
- [71] L. Torrey, M. Taylor, Teaching on a budget: agents advising agents in reinforcement learning, in: Proceedings of the 2013 International Conference on Autonomous Agents and Multi-Agent Systems, International Foundation for Autonomous Agents and Multiagent Systems, 2013, pp. 1053–1060.
- [72] J. van der Waa, J. van Diggelen, K. van den Bosch, M.A. Neerincx, Contrastive explanations for reinforcement learning in terms of expected consequences, CoRR, arXiv:1807.08706.
- [73] W. Vogt, Dictionary of Statistics & Methodology: A Nontechnical Guide for the Social Sciences, Sage Publications, 2005.
- [74] Z. Wang, T. Schaul, M. Hessel, H. van Hasselt, M. Lanctot, N. de Freitas, Dueling network architectures for deep reinforcement learning, in: Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, 2016, pp. 1995–2003.
- [75] L. Weitkamp, E. van der Pol, Z. Akata, Visual rationalizations in deep reinforcement learning for Atari games, CoRR, arXiv:1902.00566, Feb. 2019.
- [76] K. Weitz, T. Hassan, U. Schmid, J.-U. Garbas, Deep-learned faces of pain and emotions: elucidating the differences of facial expressions with the help of explainable AI methods, Tech. Mess. 86 (7–8) (2019) 404–412.
- [77] K. Weitz, D. Schiller, R. Schlagowski, T. Huber, E. André, “Do you trust me?": Increasing user-trust by integrating virtual agents in explainable AI interaction design, in: Proceedings of the 19th ACM International Conference on Intelligent Virtual Agents, IVA '19, ACM, New York, NY, USA, 2019, pp. 7–9.
- [78] T. Zahavy, N. Ben-Zrihem, S. Mannor, Graying the black box: understanding DQNs, in: Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, 2016, pp. 1899–1908.
- [79] M.D. Zeiler, R. Fergus, Visualizing and understanding convolutional networks, in: Computer Vision - ECCV 2014 - 13th European Conference, Proceedings, Part I, Zurich, Switzerland, 2014, pp. 818–833.
- [80] B. Zhang, A. Dafoe, Artificial intelligence: American attitudes and trends, Available at SSRN 3312874, 2019.