

# EXPLORATORY NOT EXPLANATORY: COUNTERFACTUAL ANALYSIS OF SALIENCY MAPS FOR DEEP RL

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

Saliency maps are often used to suggest explanations of the behavior of deep reinforcement learning (RL) agents. However, the explanations derived from saliency maps are often unfalsifiable and can be highly subjective. We introduce an empirical approach grounded in counterfactual reasoning to test the hypotheses generated from saliency maps and show that explanations suggested by saliency maps are often not supported by experiments. Our experiments suggest that saliency maps are best viewed as an exploratory tool rather than an explanatory tool.

## 1 INTRODUCTION

Saliency map methods are a popular visualization technique that produce heatmap-like output highlighting the importance of different regions of some visual input. They are frequently used to explain how deep networks classify images in computer vision applications (Simonyan et al., 2014; Springenberg et al., 2014; Shrikumar et al., 2017; Smilkov et al., 2017; Selvaraju et al., 2017; Zhang et al., 2018; Zeiler & Fergus, 2014; Ribeiro et al., 2016; Dabkowski & Gal, 2017; Fong & Vedaldi, 2017) and to explain how agents choose actions in reinforcement learning (RL) applications (Bogdanovic et al., 2015; Wang et al., 2015; Zahavy et al., 2016; Greydanus et al., 2017; Iyer et al., 2018; Sundar, 2018; Yang et al., 2018; Annasamy & Sycara, 2019).

Saliency methods in computer vision and reinforcement learning use similar procedures to generate these maps. However, the temporal and interactive nature of RL systems presents a unique set of opportunities and challenges. Deep models in reinforcement learning select sequential actions whose effects can interact over long time periods. This contrasts strongly with visual classification tasks, in which deep models merely map from images to labels. For RL systems, saliency maps are often used to assess an agent’s *internal representations* and *behavior* over multiple frames in the environment, rather than to assess the importance of specific pixels in classifying images.

Despite their common use to explain agent behavior, it is unclear whether saliency maps provide useful explanations of the behavior of deep RL agents. Some prior work has evaluated the applicability of saliency maps for explaining the behavior of image classifiers (Adebayo et al., 2018; Kindermans et al., 2019; Samek et al., 2016), but there is not a corresponding literature evaluating the applicability of saliency maps for explaining RL agent behavior.

In this work, we develop a methodology grounded in counterfactual reasoning to empirically evaluate the explanations generated using saliency maps in deep RL. Specifically, we:

- C1** Survey the ways in which saliency maps have been used as evidence in explanations of deep RL agents.
- C2** Describe a new interventional method to evaluate the inferences made from saliency maps.
- C3** Experimentally evaluate how well the pixel-level inferences of saliency maps correspond to the semantic-level inferences of humans.

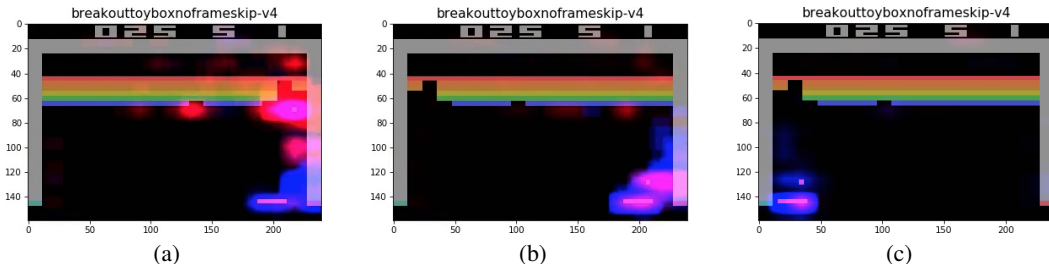


Figure 1: (a) A perturbation saliency map from a frame in Breakout, (b) a saliency map from the same model and frame with the brick pattern reflected across the vertical axis, and (c) a saliency map from the same model and frame with the ball, paddle and brick pattern reflected across the vertical axis. The blue and red regions represent their importance in action selection and reward estimation from the current state, respectively. The pattern and intensity of saliency around the channel is not symmetric in either reflection intervention.

## 2 INTERPRETING SALIENCY MAPS IN DEEP RL

Consider the saliency maps generated from a deep RL agent trained to play the Atari game Breakout. The goal of Breakout is to use the paddle to keep the ball in play so it hits bricks, eliminating them from the screen. Figure 1a shows a sample frame with its corresponding saliency. The red regions represent the importance of these pixels for action selection. Note the high saliency on the missing section of bricks (“tunnel”) in Figure 1a.

Creating a tunnel to target bricks at the top layers is one of the most high-profile examples of agent behavior being explained according to *semantic*, human-understandable concepts (Mnih et al., 2015). Given the intensity of saliency on the tunnel in 1a, it may seem reasonable to infer that this saliency map provides evidence that the agent has learned to aim at tunnels. If this is the case, moving the horizontal position of the tunnel should lead to similar saliency patterns on the new tunnel. However, Figures 1b and 1c show that the saliency pattern is not preserved. Neither the presence of the tunnel, nor the relative positioning of the ball, paddle, and tunnel, are responsible for the intensity of the saliency observed in Figure 1a.

### 2.1 SALIENCY MAPS AS INTERVENTIONS

Examining how some of the technical details of reinforcement learning interact with saliency maps can help understand both the potential utility and the potential pitfalls of interpreting saliency maps. RL methods enable agents to learn how to act effectively within an environment by repeated interaction with that environment. Certain states in the environment give the agent positive or negative reward. The agent learns a *policy*, a mapping between states and actions according to these reward signals. The goal is to learn a policy that maximizes the discounted sum of rewards received while acting in the environment (Sutton et al., 1998). Deep reinforcement learning uses deep neural networks to represent policies. These models enable interaction with environments requiring high-dimensional state inputs (e.g., Atari games).

Consider the graphical model in Figure 2a representing the deep RL system for a vision-based game environment. The environment maintains some (usually latent) game state. Some function  $F$  produces a high-dimensional pixel representation of game state (“Pixels”). The learned network takes this pixel image and produces logits used to select an action. Temporally extended sequences of this action selection procedure result in observed agent behavior.

Saliency maps are produced by performing some kind of intervention  $M$  on this system and calculating the difference in logits produced by the original and modified images. The interventions used to calculate saliency for deep RL are performed at the pixel level (red node and arrow in Figure 2a). These interventions change the conditional probability distribution of “Pixels” by giving it another parent (Pearl, 2000).

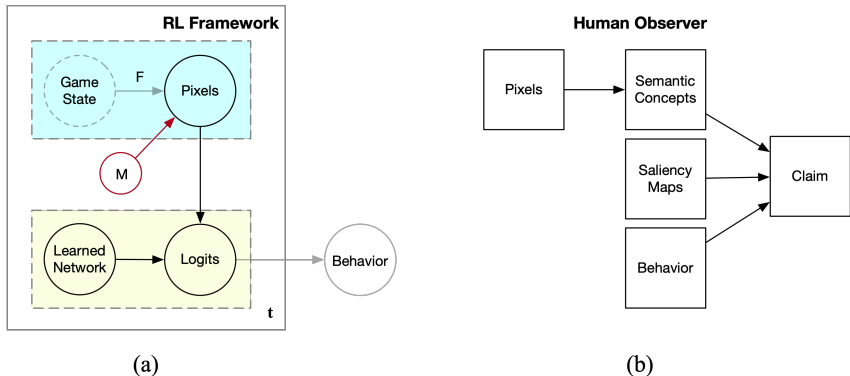


Figure 2: (a) Causal graphical model of the relationships between the agent (yellow plate) and an image-based environment (blue plate), and (b) diagram of how a human observer infers explanations.

Functionally, this can be accomplished through a variety of means, including changing the color of the pixel (Simonyan et al., 2014), adding a gray mask (Zeiler & Fergus, 2014), blurring a small region (Greydanus et al., 2017), or masking objects with the background color (Iyer et al., 2018). The interventions  $M$  are used to simulate the effect of the absence of the pixel(s) on the network’s output. Note however that these interventions change the image in a way that is inconsistent with the generative process  $F$ . They are not “naturalistic” modifications. This type of intervention produces images for which the learned network function may not be well-defined.

## 2.2 EXPLANATIONS FROM SALIENCY MAPS

To form explanations of agent behavior, human observers combine information from saliency maps, agent behavior, and semantic concepts. Figure 2b shows a system diagram of how these components interact. Hypotheses about the semantic features identified by the learned policy are proposed by reasoning backwards about what representation might jointly produce the observed saliency pattern and agent behavior.

Counterfactual reasoning has been identified as a particularly effective way to present explanations of the decision boundaries of deep models (Mittelstadt et al., 2019). Humans use counterfactuals to reason about the enabling conditions of particular outcomes, as well as to identify situations where the outcome would have occurred even in the absence of some action or condition (Byrne, 2019). Saliency maps provide a kind of pixel-level counterfactual, but if the goal is to explain agent behavior according to semantic concepts, interventions at the pixel level may not be sufficient.

Since many semantic concepts may map to the same set of pixels, it may be difficult to identify the functional relationship between changes in pixels and changes in network output according to semantic concepts or game state (Chalupka et al., 2015). Researchers may be interpreting differences in network outputs as evidence of differences in semantic concepts. However, changes in pixels do not guarantee changes in semantic concepts or game state.

In terms of changes to pixels, semantic concepts, and game state, we distinguish among three classes of interventions: *distortion*, *semantics-preserving*, and *fat-hand* (see Table 1). Semantics-preserving and fat-hand interventions are defined with respect to a specific set of semantic concepts. Fat-hand interventions change game state in such a way that the semantic concepts of interest are also altered.

The pixel-level manipulations used to produce saliency maps primarily result in distortion interventions, though some saliency methods (e.g., object-based) may conceivably produce semantics-preserving or fat-hand interventions as well. As pixel-level interventions are not guaranteed to produce changes in semantic concepts, counterfactual evaluations applying semantics-preserving interventions may be a more appropriate approach for precisely testing hypotheses of behavior.

	Pixels	Game state	Semantic concepts	Examples
Distortion	✓	✗	✗	Adversarial ML (Szegedy et al., 2014)
Semantics-preserving	✓	✓	✗	Reflection across a line of symmetry
Fat-hand	✓	✓	✓	Teleporting the agent to a new position

Table 1: Categories of interventions on images. **Distortion interventions** change pixels without changing game state or semantic concepts. The pixel perturbations in adversarial ML, where adversarial noise is added to images to change the output of the network without making human-perceptible changes in the image. **Semantics-preserving interventions** are manipulations of game state that result in an image preserving some semantic concept of interest. Reflections across lines of symmetry typically alter aspects of game state, but do not meaningfully change any semantic information about the scene. **“Fat-hand” interventions** are manipulations intended to measure the effect of some specific treatment, but which unintentionally alter other relevant aspects of the system. The term is drawn from the literature on causal modeling (Scheines, 2005).

### 3 SURVEY OF USAGE OF SALIENCY MAPS IN DEEP RL LITERATURE

To assess how saliency maps are used to make inferences regarding agent behavior in practice, we surveyed recent conference papers in deep RL. We focused our pool of papers on those that use saliency maps to generate explanations or make claims regarding agent behavior. Our search criteria consisted of examining papers that cited work that first described any of the following four types of saliency maps:

**Jacobian Saliency.** Wang et al. (2015) extended gradient-based saliency maps to deep RL by computing the Jacobian of the output logits with respect to a stack of input images.

**Perturbation Saliency.** Greydanus et al. (2017) generate saliency maps by perturbing the original input image using a Gaussian blur of the image and measure changes in policy from removing information from a region.

**Object Saliency.** Iyer et al. (2018) use template matching, a common computer vision technique (Brunelli, 2009), to detect (template) objects within an input image and measure salience through changes in Q-values for masked and unmasked objects.

**Attention Saliency.** Most recently, attention-based saliency mapping methods have been proposed to generate interpretable saliency maps (Mott et al., 2019; Nikulin et al., 2019).

From a set of 90 papers, we found 46 claims drawn from 11 papers that cited and used saliency maps as evidence in their explanations of agent behavior. The full set of claims are given in Appendix A.

#### 3.1 SURVEY RESULTS

We found three categories of saliency map usage, summarized in Table 2. First, all claims interpret salient areas as a proxy for agent focus. For example, a claim about a Breakout agent notes that the network is focusing on the paddle and little else (Greydanus et al., 2017).

Second, 87% of the claims in our survey propose hypotheses about the features of the learned policy by reasoning backwards about what representation might jointly produce the observed saliency pattern and agent behavior. These types of claims either develop an *a priori* explanation of behavior and verify it using saliency, or they propose an *ad hoc* explanation after observing saliency to reason about how the agent is using salient areas. One *a priori* claim notes that

	Discuss Focus	Generate Explanation	Evaluate Explanation
Jacobian	21	19	0
Perturbation	11	9	1
Object	5	4	2
Attention	9	8	0
Total	46	40	3

Table 2: Summary of the survey on usage of saliency maps. The columns contains different types of saliency methods and the rows contain different traits regarding the claims.

the displayed score is the only differing factor between two states and verify from saliency that the agent is focusing on these pixels (Zahavy et al., 2016). An *ad hoc* claim about a racing game notes that the agent is recognizing a time-of-day cue from the background color and acting to prepare for a new race (Yang et al., 2018).

Finally, only 7% (3 out of 46) of the claims attempt to empirically verify the explanations generated from saliency maps. One of these attempts to corroborate the interpreted saliency behavior by obtaining additional saliency samples from multiple runs of the game. The other two attempt to manipulate semantics in the pixel input to assess the agent’s response by, for example, adding an additional object to verify a hypothesis about memorization (Annasamy & Sycara, 2019).

### 3.2 COMMON PITFALLS IN CURRENT USAGE

In the course of the survey, we also observed several more qualitative characteristics of how saliency maps are routinely used.

**Subjectivity.** Recent critiques of machine learning have already noted a worrying tendency to conflate speculation and explanation (Lipton & Steinhardt, 2018). Saliency methods are not designed for formalizing an abstract human-understandable concept such as “aiming” in Breakout, and they do not provide a means to quantitatively compare semantically meaningful consequences of agent behavior. This leads to subjectivity in the conclusions drawn from saliency maps.

**Unfalsifiability.** One hallmark of a scientific hypothesis or claim is that it is falsifiable (Popper, 1959). If a claim is false, its falsehood should be identifiable from some conceivable experiment or observation. One of the most disconcerting practices identified in the survey is the presentation of unfalsifiable interpretations of saliency map patterns. An example: “A diver is noticed in the saliency map but misunderstood as an enemy and being shot at” (Zahavy et al., 2016). It is unclear how we might falsify an abstract concept such as “misunderstanding”.

**Assessment of Learned Representations.** Evidence indicates that humans rely on causal relationships between objects and their categorization into abstract classes to learn complex processes, such as video games (Tenenbaum & Niyogi, 2003; Dubey et al., 2018). Our survey suggests *researchers infer* that: (1) salient regions map to learned representations of semantic concepts (e.g., ball, paddle), and (2) the relationships between the salient regions map to high-level behaviors (e.g., channel-building, aiming). The researcher’s expectations impose a strong bias on both the existence and nature of these mappings.

## 4 METHODOLOGY

Our survey indicates that many researchers use saliency maps as an explanatory tool to infer the representations and processes behind an agent’s behavior. However, the extent to which such inferences are valid has not been empirically evaluated under controlled conditions.

In this section, we show how to generate falsifiable hypotheses from saliency maps and propose an intervention-based approach to verify the hypotheses generated from saliency maps. We intervene on game state to produce counterfactual semantic conditions. This provides a medium through which we might assess the relationship between saliency and learned semantic representations.

**Building Falsifiable Hypotheses from Saliency Maps.** Though saliency maps may not relate directly to semantic concepts, they may still be an effective tool for exploring hypotheses about agent behavior. Claims or explanations informed by saliency maps have three components: semantic concepts, learned representations, and behavior. Let  $X$  be a subset of the semantic concepts that can be inferred from the input image. Semantic concepts are often identified visually from the pixel output as the *game state* is typically latent. Let  $B$  represent behavior, or aggregate actions, over temporally extended sequences of frames, and let  $R$  be a representation that is a function of some pixels that the agent learns during training. As noted in Section 3, researchers often attempt to infer representations from saliency patterns and agent behavior.

To create scientific claims from saliency maps, we recommend using a relatively standard pattern which facilitates objectivity and falsifiability:

Concept set  $X$  is salient  $\implies$  Agent has learned representation  $R$  resulting in behavior  $B$ .

Consider the Breakout brick reflection example presented in Section 2. The hypothesis introduced (“the agent has learned to aim at tunnels”) can be reconstructed as: *bricks* are salient  $\implies$  agent has learned to *identify a partially complete tunnel* resulting in *maneuvering the paddle to hit the ball toward that region*. Forming hypotheses in this format constructs falsifiable claims amenable to empirical analysis.

**Counterfactual Evaluation of Claims.** As previously established in Figure 2, the learned representation and pixel input share a causal relationship with saliency maps generated over a sequence of frames. Given that the representation learned is static, the relationship between the learned representation and saliency should be invariant under different manipulations of pixel input. We use this property to assess saliency under counterfactual conditions.

We generate counterfactual conditions by intervening on the RL environment. Prior work has focused on manipulating the pixel input. However, this does not modify the underlying latent game state. Instead, we intervene directly on *game state*. In the do-calculus formalism (Pearl, 2000), this shifts the intervention node in Figure 2a to game state, which leaves the generative process of the pixel image  $F$  intact.

We employ TOYBOX, a set of fully parameterized implementation of Atari games (Foley et al., 2018), to generate interventional data under counterfactual conditions. The interventions are dependent on the mapping between semantic concepts and learned representations in the hypotheses. Given a mapping between concept set  $X$  and a learned representation  $R$ , any intervention would require meaningfully manipulating the state in which  $X$  resides to assess the saliency on  $X$  under the semantic treatment applied.

Since the learned policies should be semantically invariant under manipulations of the RL environment, by intervening on state, we can verify whether the counterfactual states produce expected patterns of saliency on the associated concept set  $X$ . If the counterfactual saliency maps reflect similar saliency patterns, this provides stronger evidence that the observed saliency indicates the agent has learned representation  $R$  corresponding to semantic concept set  $X$ .

## 5 EVALUATION OF HYPOTHESES ON AGENT BEHAVIOR

We conduct three case studies to evaluate hypotheses about the relationship between semantic concepts and semantic processes formed from saliency maps. Each case study uses observed saliency maps to identify hypotheses in the format described in Section 4. The hypotheses were generated by watching multiple episodes and noting atypical, interesting or popular behaviors from saliency maps. Using TOYBOX allows us to produce counterfactual states and to generate saliency maps in these altered states. The case studies are conducted on two Atari games, Breakout and Amidar. The deterministic nature of both games allows some stability in the way we interpret the network’s action selection. Each map is produced from an agent trained with A2C (Mnih et al., 2016) using the OpenAI Baselines implementation (Dhariwal et al., 2017).

**Case Study 1: Breakout Brick Translation.** Here we evaluate the behavior from Section 2:

**Hypothesis 1:** bricks are salient  $\implies$  agent has learned to {identify a partially complete tunnel} resulting in {maneuvering the paddle to hit the ball toward that region}

We intervene on the state by translating the brick configurations horizontally. We expect salience will be nearly invariant to the horizontal translation of the brick configuration. Figure 3a depicts saliency after intervention. Salience on the tunnel is less pronounced under left translation, and more pronounced under right translation. Since the paddle appears on the right, we additionally move the ball and paddle to the far left (Figure 3b).

*Conclusion.* Temporal association (e.g. formation of a tunnel followed by higher saliency) does not generally imply causal dependence. In this case at least, tunnel formation and salience appear to be confounded by location or, at least, the dependence of these phenomena are highly dependent on location.

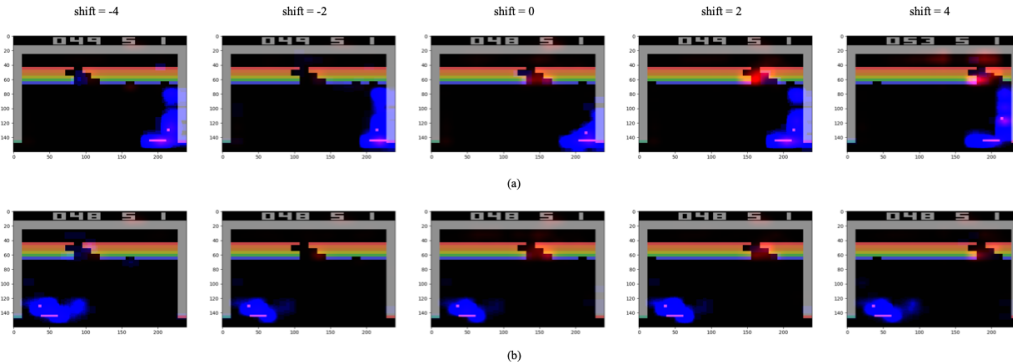


Figure 3: (a) saliency after shifting the brick positions where shift=0 represents the original frame; (b) saliency after shifting the brick positions along with shifting ball and paddle to the left.

**Case Study 2: Amidar Score.** Amidar is a Pac-Man-like game in which an agent attempts to completely traverse a series of passages while avoiding enemies. The yellow sprite that indicates the location of the agent is almost always salient in Amidar. Surprisingly, the displayed score is salient as often as the yellow sprite throughout the episode with varying levels of intensity (see Figure 4a representing an object saliency map). This can lead to multiple hypotheses about the agent’s learned representation: (1) the agent has learned to associate increasing score with higher reward; (2) due to the deterministic nature of Amidar, the agent has created a lookup table that associates its score and its actions. We can summarize these hypotheses as follows:

**Hypothesis 2:** score is salient  $\implies$  agent has learned to {use score as a guide to traverse the board} resulting in {successfully following similar paths in games}.

To evaluate hypothesis 2, we designed four interventions on score:

- *intermittent\_reset*: modify the score to 0 every  $x \in [5, 20]$  timesteps.
- *random\_varying*: modify the score to a random number between [1,200] every  $x \in [5, 20]$  timesteps.
- *fixed*: select a score from [0,200] and fix it for the whole game.
- *decremented*: modify score to be 3000 initially and decrement score by  $d \in [1, 20]$  at every timestep.

Figures 4b and 4c show the result of intervening on displayed score on reward and saliency intensity, respectively, for the first 1000 timesteps of an episode. The mean is calculated over 50 episodes. If an agent died before 1000 timesteps, the last reward was extended for the remainder of the timesteps and saliency was set to zero.

Using reward as a summary of agent behavior, different interventions on score produce different agent behavior. Rewards differ over time for all interventions, typically due to early agent death. However, saliency intensity patterns of all interventions follow the original trajectory very closely. Different interventions on displayed score cause differing degrees of degraded performance (Figure 4b) despite producing similar saliency maps (Figure 4c), indicating that agent behavior is underdetermined by saliency. Specifically, the saliency intensity patterns are similar for the *control*, *fixed*, and *decremented* scores, while the non-ordered score interventions result in degraded performance.

Similar trends are noted for Jacobian and perturbation saliency methods in Appendix B.1.

*Conclusion.* The existence of a high correlation between two processes (e.g., incrementing score and continuance of saliency) does not imply causation. Interventions can be useful in identifying the common cause leading to the high correlation.

**Case Study 3: Amidar Enemy Distance.** Enemies are salient in Amidar at varying times. From visual inspection, we observe that enemies close to the player tend to have higher saliency. Accordingly, we generate the following hypothesis:

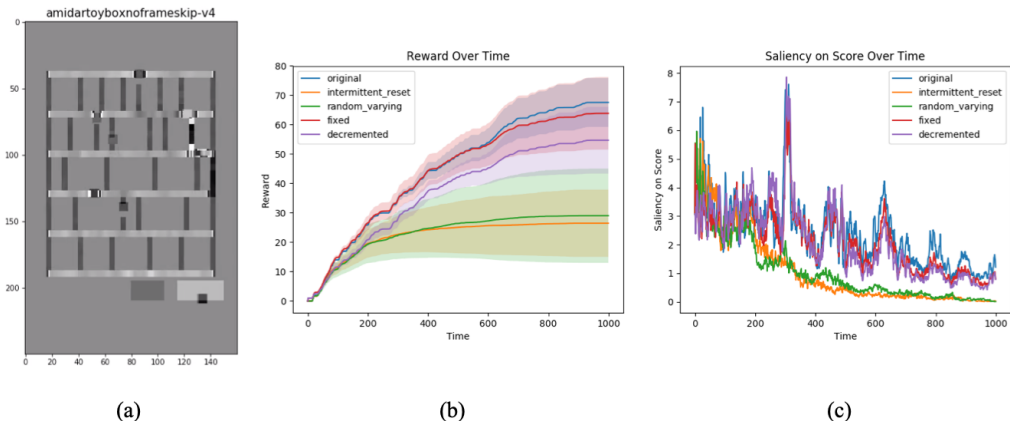


Figure 4: Interventions on displayed score in Amidar. (a) an example object saliency map for an episode of Amidar; (b) reward over time for different interventions on score; (c) saliency on displayed score over time.

**Hypothesis 3:** enemy is salient  $\implies$  agent has learned to {look for enemies close to it} resulting in {successful avoidance of enemy collision}.

Without directly intervening on the game state, we can first identify whether the player-enemy distance and enemy saliency is correlated using observational data. We collect 1000 frames of an episode of Amidar and record the Manhattan distance to the player and object saliency of each enemy. Figure 5a shows the distance of each enemy to the player over time with saliency intensity represented by the shaded region. Figure 5b shows the correlation between the distance to each enemy and the corresponding saliency. It is clear that there is no correlation between saliency and distance of each enemy to the player.

Given that statistical dependence is almost always a necessary pre-condition for causation, we expect that there will not be any causal dependence. To further examine this, we intervene on enemy positions of salient enemies at each timestep by moving the enemy closer and farther away from the player. Figure 5c contains these results. Given Hypothesis 3, we would expect see an increasing trend in saliency for enemies closer to the player. However, this is not the case.

In addition, we find no correlation in the enemy distance experiments for the Jacobian or perturbation saliency methods (included in Appendix B.2). Jacobian saliency performed the worst for the intervention-based experiment, suggesting that there is no impact of player-enemy distance on saliency.

*Conclusion.* Spurious correlations, or misinterpretations of existing correlation, can occur between two processes (e.g. correlation between player-enemy distance and saliency), and human observers are susceptible to identifying spurious correlations (Simon, 1954). Spurious correlations can sometimes be identified from observational analysis without requiring interventional analysis.

## 6 DISCUSSION AND RELATED WORK

Thinking counterfactually about the explanations generated from saliency maps facilitates empirical evaluation of those explanations. The experiments above show some of the difficulties in drawing conclusions from saliency maps. These include the tendency of human observers to incorrectly infer association between observed processes, the potential for experimental evidence to contradict seemingly obvious observational conclusions, and the challenges of potential confounding in temporal processes.

One of the main conclusions from this evaluation is that *saliency maps are an exploratory tool rather than an explanatory tool*. Saliency maps alone cannot be reliably used to infer explanations and instead require other supporting tools. This can include combining evidence from saliency






Figure 5: (a) shows the distance to player of each enemy in Amidar over time where saliency intensity is represented by the shade behind each enemy.; (b) shows the correlation between the distance to player and corresponding saliency for each enemy; (c) shows saliency upon intervening on enemy position.

maps with other explanation methods or employing a more experimental approach to evaluation of saliency maps such as the approach demonstrated in the case studies above.

The framework for generating falsifiable hypotheses suggested in Section 4 can assist with designing more specific and falsifiable explanations. The distinction between the components of an explanation, particularly the semantic concept set  $X$ , learned representation  $R$  and observed behavior  $B$ , can further assist in experimental evaluation.

**Generalization of Proposed Methodology.** We propose intervention-based experimentation as a primary tool to evaluate the hypotheses generated from saliency maps. Yet, alternative methods can identify a false hypothesis even earlier. For instance, evaluating statistical dependence alone can help identify some situations in which causation is absent (e.g., Case Study 3).

We employ TOYBOX, a high-performance suite of intervenable RL environments, designed for behavioral tests Tosch et al. (2019). However, limited forms of evaluation may be possible in non-intervenable environments, though they may be more tedious to implement. For instance, each of the interventions conducted in Case Study 1 can be produced in an observation-only environment by manipulating the pixel input (Brunelli, 2009; Chalupka et al., 2015). Developing more experimental systems for evaluating explanations is an open area of research.

**Evaluation and Critiques of Saliency Maps.** Prior work in the deep network literature has evaluated and critiqued saliency maps. Kindermans et al. (2019) and Adebayo et al. (2018) demonstrate the utility of saliency maps by adding random variance in input. Seo et al. (2018) provide a theoretical justification of saliency and hypothesize that there exists a correlation between gradients-based saliency methods and model interpretation. Samek et al. (2017) and Hooker et al. (2018) present evaluation of existing saliency methods for image classification. Others have critiqued network attention as a means of explanation (Jain & Wallace, 2019; Brunner et al., 2019).

## 7 CONCLUSIONS

We conduct a survey of uses of saliency maps, propose a methodology to evaluate saliency maps, and examine the extent to which the agent’s learned representations can be inferred from saliency maps. We investigate how well the pixel-level inferences of saliency maps correspond to the semantic concept-level inferences of human-level interventions. Our results show saliency maps cannot be trusted to reflect causal relationships between semantic concepts and agent behavior. We recommend saliency maps to be used as an exploratory tool, not explanatory tool.

## REFERENCES

Julius Adebayo, Justin Gilmer, Michael Muelly, Ian Goodfellow, Moritz Hardt, and Been Kim. Sanity checks for saliency maps. In *Advances in Neural Information Processing Systems*, pp. 9505–9515, 2018.

- Raghuram Mandyam Annasamy and Katia Sycara. Towards better interpretability in deep q-networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 4561–4569, 2019.
- Miroslav Bogdanovic, Dejan Markovikj, Misha Denil, and Nando De Freitas. Deep apprenticeship learning for playing video games. 2015.
- Roberto Brunelli. *Template Matching Techniques in Computer Vision: Theory and Practice*. Wiley Publishing, 2009. ISBN 0470517069, 9780470517062.
- Gino Brunner, Yang Liu, Damin Pascual, Oliver Richter, and Roger Wattenhofer. On the validity of self-attention as explanation in transformer models, 2019.
- Ruth 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*, pp. 6276–6282. International Joint Conferences on Artificial Intelligence Organization, 7 2019. doi: 10.24963/ijcai.2019/876. URL <https://doi.org/10.24963/ijcai.2019/876>.
- Krzysztof Chalupka, Pietro Perona, and Frederick Eberhardt. Visual causal feature learning. In *Proceedings of the Thirty-First Conference on Uncertainty in Artificial Intelligence, UAI 15*, pp. 181–190, Arlington, Virginia, United States, 2015. AUAI Press. ISBN 978-0-9966431-0-8. URL <http://dl.acm.org/citation.cfm?id=3020847.3020867>.
- Piotr Dabkowski and Yarín Gal. Real time image saliency for black box classifiers. In *Advances in Neural Information Processing Systems*, pp. 6967–6976, 2017.
- Prafulla Dhariwal, Christopher Hesse, Oleg Klimov, Alex Nichol, Matthias Plappert, Alec Radford, John Schulman, Szymon Sidor, and Yuhuai Wu. OpenAI Baselines. <https://github.com/openai/baselines>, 2017.
- Rachit Dubey, Pulkit Agrawal, Deepak Pathak, Thomas L Griffiths, and Alexei A Efros. Investigating human priors for playing video games. *arXiv preprint arXiv:1802.10217*, 2018.
- John Foley, Emma Tosch, Kaleigh Clary, and David Jensen. Toybox: Better Atari Environments for Testing Reinforcement Learning Agents. In *NeurIPS 2018 Workshop on Systems for ML*, 2018.
- Ruth C Fong and Andrea Vedaldi. Interpretable explanations of black boxes by meaningful perturbation. *arXiv preprint arXiv:1704.03296*, 2017.
- Vikash Goel, Jameson Weng, and Pascal Poupart. Unsupervised video object segmentation for deep reinforcement learning. In *Advances in Neural Information Processing Systems*, pp. 5683–5694, 2018.
- Sam Greydanus, Anurag Koul, Jonathan Dodge, and Alan Fern. Visualizing and understanding atari agents. *arXiv preprint arXiv:1711.00138*, 2017.
- Sara Hooker, Dumitru Erhan, Pieter-Jan Kindermans, and Been Kim. Evaluating feature importance estimates. *arXiv preprint arXiv:1806.10758*, 2018.
- Rahul Iyer, Yuezhong Li, Huao Li, Michael Lewis, Ramitha Sundar, and Katia Sycara. Transparency and explanation in deep reinforcement learning neural networks. 2018.
- Sarthak Jain and Byron C. Wallace. Attention is not Explanation. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 3543–3556, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics.
- Pieter-Jan Kindermans, Sara Hooker, Julius Adebayo, Maximilian Alber, Kristof T Schütt, Sven Dähne, Dumitru Erhan, and Been Kim. The (un) reliability of saliency methods. In *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning*, pp. 267–280. Springer, 2019.
- Zachary C Lipton and Jacob Steinhardt. Troubling trends in machine learning scholarship. *arXiv preprint arXiv:1807.03341*, 2018.

- Brent Mittelstadt, Chris Russell, and Sandra Wachter. Explaining explanations in ai. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, FAT\* '19, pp. 279–288, 2019.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529, 2015.
- Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*, pp. 1928–1937, 2016.
- Alex Mott, Daniel Zoran, Mike Chrzanowski, Daan Wierstra, and Danilo J Rezende. Towards interpretable reinforcement learning using attention augmented agents. *arXiv preprint arXiv:1906.02500*, 2019.
- Dmitry Nikulin, Anastasia Ianina, Vladimir Aliev, and Sergey Nikolenko. Free-lunch saliency via attention in atari agents. *arXiv preprint arXiv:1908.02511*, 2019.
- Judea Pearl. *Causality: models, reasoning and inference*, volume 29. Springer, 2000.
- Karl Popper. *The logic of scientific discovery*. Routledge, 1959.
- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. Why should i trust you?: Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pp. 1135–1144. ACM, 2016.
- Christian Rupprecht, Cyril Ibrahim, and Chris Pal. Visualizing and discovering behavioural weaknesses in deep reinforcement learning. 2018.
- Wojciech Samek, Alexander Binder, Grégoire Montavon, Sebastian Lapuschkin, and Klaus-Robert Müller. Evaluating the visualization of what a deep neural network has learned. *IEEE transactions on neural networks and learning systems*, 28(11):2660–2673, 2016.
- Wojciech Samek, Alexander Binder, Grégoire Montavon, Sebastian Lapuschkin, and Klaus-Robert Müller. Evaluating the visualization of what a deep neural network has learned. *IEEE transactions on neural networks and learning systems*, 28(11):2660–2673, 2017.
- Richard Scheines. The similarity of causal inference in experimental and nonexperimental studies. *Philosophy of Science*, 72(5):927–940, 2005.
- Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, Dhruv Batra, et al. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *ICCV*, pp. 618–626, 2017.
- Junghoon Seo, Jeongyeol Choe, Jamyounng Koo, Seunghyeon Jeon, Beomsu Kim, and Taegyun Jeon. Noise-adding methods of saliency map as series of higher order partial derivative. *arXiv preprint arXiv:1806.03000*, 2018.
- Avanti Shrikumar, Peyton Greenside, and Anshul Kundaje. Learning important features through propagating activation differences. *arXiv preprint arXiv:1704.02685*, 2017.
- Herbert A. Simon. Spurious correlation: A causal interpretation. *Journal of the American Statistical Association*, 49(267):467–479, 1954. ISSN 01621459. URL <http://www.jstor.org/stable/2281124>.
- Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. *ICLR*, 2014.
- Daniel Smilkov, Nikhil Thorat, Been Kim, Fernanda Viégas, and Martin Wattenberg. Smoothgrad: removing noise by adding noise. *arXiv preprint arXiv:1706.03825*, 2017.
- Jost Tobias Springenberg, Alexey Dosovitskiy, Thomas Brox, and Martin Riedmiller. Striving for simplicity: The all convolutional net. *arXiv preprint arXiv:1412.6806*, 2014.

- Ramitha Sundar. *Transparency in Deep Reinforcement Learning Networks*. PhD thesis, Carnegie Mellon University, 2018.
- Richard S Sutton, Andrew G Barto, et al. *Reinforcement learning: An introduction*. MIT press, 1998.
- Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. In *ICLR 2014: Proceedings of the Third International Conference on Learning Representations*, 2014.
- Joshua B Tenenbaum and Sourabh Niyogi. Learning causal laws. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume 25, 2003.
- Emma Tosch, Kaleigh Clary, John Foley, and David Jensen. Toybox: A suite of environments for experimental evaluation of deep reinforcement learning, 2019.
- Douwe van der Wal, Bachelor Opleiding Kunstmatige Intelligentie, and Wenling Shang. Advantage actor-critic methods for carracing. 2018.
- Ziyu Wang, Tom Schaul, Matteo Hessel, Hado Van Hasselt, Marc Lanctot, and Nando De Freitas. Dueling network architectures for deep reinforcement learning. *arXiv preprint arXiv:1511.06581*, 2015.
- Zhao Yang, Song Bai, Li Zhang, and Philip HS Torr. Learn to interpret atari agents. *arXiv preprint arXiv:1812.11276*, 2018.
- Tom Zahavy, Nir Ben-Zrihem, and Shie Mannor. Graying the black box: Understanding dqns. In *International Conference on Machine Learning*, pp. 1899–1908, 2016.
- Matthew D Zeiler and Rob Fergus. Visualizing and understanding convolutional networks. In *European conference on computer vision*, pp. 818–833. Springer, 2014.
- Jianming Zhang, Sarah Adel Bargal, Zhe Lin, Jonathan Brandt, Xiaohui Shen, and Stan Sclaroff. Top-down neural attention by excitation backprop. *International Journal of Computer Vision*, 126(10):1084–1102, 2018.

## A SURVEY OF USAGE OF SALIENCY MAPS IN DEEP RL LITERATURE

We conducted a survey of recent literature to assess how saliency maps are used to interpret agent behavior in deep RL. We began our search by focusing on work citing the following four types of saliency maps: Jacobian (Wang et al., 2015), perturbation (Greydanus et al., 2017), object (Iyer et al., 2018) and attention (Mott et al., 2019). Papers were selected if they employed saliency maps to create explanations regarding agent behavior. This resulted in selecting 46 claims from 11 papers. These 11 papers have appeared at ICML (3), NeurIPS (1), AAAI (2), ArXiv (3), OpenReview (1) and a thesis (1). There are several model-specific saliency mapping methods that we excluded from our survey.

Following is the full set of claims. All claims are for Atari games. The *Reason* column represents whether an explanation for agent behavior was provided (Y/N) and the *Exp* column represents whether an experiment was conducted to evaluate the explanation (Y/N).

Claim	Game	Saliency Type	Reason	Exp
Greydanus et al. (2017)				
“The agent is positioning its own paddle, which allows it to return the ball at a specific angle.”	Pong	Perturbation	Y	N
“Interestingly, from the saliency we see that the agent attends to very little besides its own paddle: not even the ball.”	Pong	Perturbation	N	N
“After the agent has executed the kill shot, we see that saliency centers entirely around the ball. This makes sense since at this point neither paddle can alter the outcome and their positions are irrelevant.”	Pong	Perturbation	Y	N
“It appears that the deep RL agent is exploiting the deterministic nature of the Pong environment. It has learned that it can obtain a reward with high certainty upon executing a precise series of actions. This insight...gives evidence that the agent is not robust and has overfit to the particular opponent.”	Pong	Perturbation	Y	N
“[The agent] had learned a sophisticated aiming strategy during which first the actor and then the critic would ‘track’ a target. Aiming begins when the actor highlights a particular alien in blue...Aiming ends with the agent shooting at the new target.”	Space-Invaders	Perturbation	Y	N
“The critic highlights the target in anticipation of an upcoming reward.”	Space-Invaders	Perturbation	Y	N
“Notice that both actor and critic tend to monitor the area above the ship. This may be useful for determining whether the ship is protected from enemy fire or has a clear shot at enemies.”	Space-Invaders	Perturbation	Y	N
“We found that the agent enters and exits a ‘tunneling mode’ over the course of a single frame. Once the tunneling location becomes salient, it remains so until the tunnel is finished.”	Breakout	Perturbation	N	N
Zahavy et al. (2016)				
“A diver is noticed in the saliency map but misunderstood as an enemy and being shot at.”	Seaquest	Jacobian	Y	N

<p>“Once the agent finished the first screen it is presented with another one, distinguished only by the score that was accumulated in the first screen. Therefore, an agent might encounter problems with generalizing to the new screen if it over-fits the score pixels. Figure 15 shows the saliency maps of different games supporting our claim that DQN is basing its estimates using these pixels. We suggest to further investigate this, for example, we suggest to train an agent that does not receive those pixels as input.”</p>	Breakout	Jacobian	N	N
Bogdanovic et al. (2015)				
<p>“The line of cars in the upper right are far away and the agent correctly ignores them in favour of focusing on the much more dangerous cars in the lower left.”</p>	Freeway	Jacobian	Y	N
Yang et al. (2018)				
<p>“Firstly, the agent ignores irrelevant features from mountains, sky, the mileage board and empty grounds, and relies on information from the race track to make decisions. Specifically, the agent keeps separate two categories of objects on the race track, i.e., cars and the player.”</p>	Enduro	Binary Jacobian	Y	N
<p>“On the one hand, the agent locates the player and a local area around it for avoiding immediate collisions with cars. On the other hand, the agent locates the next potential collision targets at different locations, particularly the remote ones.”</p>	Enduro	Binary Jacobian	Y	N
<p>“Near the completion of the current goal, the agent celebrates in advance. As shown from Fig. 4(d) to Fig. 4(f), the left gaze loses its focus on cars and diverts to the mileage board starting when only 13 cars remain before completion. The previous car tracker now picks up on the important information that it is close to victory, and fixates on the countdown.”</p>	Enduro	Binary Jacobian	Y	N
<p>“Upon reaching the target, the agent does not receive reward signals until the next day starts. During this period the agent learns to output no-op actions, corresponding to not playing the game.”</p>	Enduro	Binary Jacobian	Y	N
<p>“When slacking happens, the agent considers the flag signs as important and the road not. The complete reverse in focus as compared to the normal case explains this shift in policy. The flags outweigh the road in importance, since they are signs of absolute zero return.”</p>	Enduro	Binary Jacobian	Y	N
<p>“(Prepping) As it turns out, the agent recognizes that the time is dawn (right before morning when race starts) from the unique colours of the light gray sky and orange mountains, therefore the agent gets ready early for a head start in the new race.”</p>	Enduro	Binary Jacobian	Y	N
<p>“When smog partially blocks the forward view, the left gaze loses its focus on cars. It strays off the road into some empty area.”</p>	Enduro	Binary Jacobian	Y	N

“In Fig. 5(a), the left gaze detects the two ghosts on the upper-right corner of the map. Therefore, ms pacman, as located by the right gaze, stays in the mid-left section to safely collect dense rewards.”	Pacman	Binary Jacobian	Y	N
“In Fig. 5(b), the left gaze locks in on all three vulnerable ghosts in the mid-right section, as ms pacman chases after them.”	Pacman	Binary Jacobian	Y	N
“In Fig. 5(c), the left gaze detects a newly-appeared cherry at the lower-left warp tunnel entrance. Ms pacman immediately enters the closest opposite tunnel entrance in the shortest path to the cherry.”	Pacman	Binary Jacobian	Y	N
“In Fig. 5(d), the right gaze locates ms pacman entering the upper-right tunnel. In this case, the left gaze no longer detects a moving object, but predicts the upper-left tunnel as the exiting point.”	Pacman	Binary Jacobian	Y	N
“As shown in Fig. 5(g), the left gaze locates the last pellet when ms pacman is in the mid-section of the maze. Therefore ms pacman moves towards the pellet.”	Pacman	Binary Jacobian	Y	N
“In Fig. 5(h), a red ghost appears in the left gaze close to the pellet, causing ms pacman to deviate to the right.”	Pacman	Binary Jacobian	Y	N
“After changing course, the ghosts approach ms pacman from all directions as shown in Fig. 5(i). Even though the agent detects all the ghosts (they appear in the gazes), ms pacman has no route to escape.”	Pacman	Binary Jacobian	Y	N
“The left gaze often focuses on white ice blocks that are the destinations of jumping.”	Frostbite	Binary Jacobian	N	N
“Fig. 6(b) shows the sub-task of the player entering the igloo, after jumping over white ice blocks for building it. The player must avoid the bear when running for the igloo.”	Frostbite	Binary Jacobian	Y	N
“As it shows, the igloo is two jumps away from completion, and the left gaze focuses on the igloo in advance for preparing to enter.”	Frostbite	Binary Jacobian	Y	N
Annasamy & Sycara (2019)				
“For example, in MsPacman, since visualizations suggest that the agent may be memorizing pacmans positions (also maybe ghosts and other objects), we simply add an extra pellet adjacent to a trajectory seen during training (Figure 7a). The agent does not clear the additional pellet and simply continues to execute actions performed during training (Figure 7b).”	Pacman	Object	Y	Y
“Similarly, in case of SpaceInvaders, the agent has a strong bias towards shooting from the leftmost-end (seen in Figure 4). This helps in clearing the triangle like shape and moving to the next level (Figure 7d). However, when triangular positions of spaceships are inverted, the agent repeats the same strategy of trying to shoot from left and fails to clear ships (Figure 7c).”	Space-Invaders	Object	Y	Y

Goel et al. (2018)				
“In our Breakout results, the network learns to split the paddle into a left and right side, and does not move the middle portion.”	Breakout	Object	Y	N
“On Beam Rider, our network achieves low loss after learning to segment the games light beams, however these beams are purely visual effects and are unimportant for action selection.”	Beam Rider	Object	Y	N
“Consequently, the game enemies, which are much smaller in size, are ignored by the network, and the resulting learned representation is not that useful for a reinforcement learning agent.”	Beam Rider	Object	N	N
van der Wal et al. (2018)				
“There is a very notable difference between the policy saliency between the two models, where the former one only pays limited attention to the road and almost no attention to the engine indicator, the opposite from fA3C-LSTM. Explicitly, it means masking any regions from the input does not cause much perturbation to the policy when trained with continuous space as targets, likely because the real consequence from a small change in action, e.g. no braking ( $a_3=0$ ) versus braking ( $a_3=0.3$ ), can be very substantial but numerically too subtle for the network to capture during optimization on the continuous spectrum.”	CarRacing	Perturbation	Y	N
Rupprecht et al. (2018)				
“Analyzing the visualizations on Seaquest, we make an interesting observation. When maximizing the Q-value for the actions, in many samples we see a low or very low oxygen meter. In these cases the submarine would need to ascend to the surface to avoid suffocation. Although the up action is the only sensible choice in this case, we also obtain visualized low oxygen states for all other actions. This implies that the agent has not understood the importance of resurfacing when the oxygen is low. We then run several roll outs of the agent and see that the major cause of death is indeed suffocation and not collision with enemies.”	SeaQuest	Perturbation	Y	Y
Mott et al. (2019)				
“The most dominant pattern we observe is that the model learns to attend to task-relevant things in the scene. In most ATARI games that usually means that the player is one of the foci of attention, as well as enemies, power-ups and the score itself (which is an important factor in the calculating the value function).”	SeaQuest	Attention	N	N



<p>“Figure 4 shows a examples of this in Ms Pacman and Alien in the both games the model scans through possible paths, making sure there are no enemies or ghosts ahead. We observe that when it does see a ghost, another path is produced or executed in order to avoid it.”</p>	Pacman	Attention	Y	N
<p>“In many games we observe that the agent learns to place trip-wires at strategic points in space such that if a game object crosses them a specific action is taken. For example, in Space Invaders two such trip wires are following the player ship on both sides such that if a bullet crosses one of them the agent immediately evades them by moving towards the opposite direction.”</p>	Space Invaders	Attention	Y	N
<p>“Another example is Breakout where we can see it working in two stages. First the attention is spread out around the general area of the ball, then focuses into a localized line. Once the ball crosses that line the agent moves towards the ball.”</p>	Breakout	Attention	Y	N
<p>“As can be seen, the system uses the two modes to make its decisions, some of the heads are content specific looking for opponent cars. Some are mixed, scanning the horizon for incoming cars and when found, tracking them, and some are location based queries, scanning the area right in front of the player for anything the crosses its path (a trip-wire which moves with the player).”</p>	Enduro	Attention	Y	N
<p>“Comparing the attention agent to the baseline agent, we see that the attention agent is sensitive to more focused areas along the possible future trajectory. The baseline agent is more focused on the area immediately in front of the player (for the policy saliency) and on the score, while the attention agent focuses more specifically on the path the agent will follow (for the policy) and on possible future longer term paths (for the value).”</p>	Pacman	Perturbation	Y	N

Nikulin et al. (2019)

<p>“Figs. 4(a)-(b) show Dense FLS digging a tunnel through blocks in Breakout. The model focuses its attention on the end of the tunnel as soon as it is complete, suggesting that it sees shooting the ball through the tunnel as a good strategy.”</p>	Breakout	Attention	Y	N
<p>“Figs. 4(c)-(d) depict the same concept of tunneling performed by the Sparse FLS model. Note how it focuses attention on the upper part of the screen after destroying multiple bricks from the top. This attention does not go away after the ball moves elsewhere (not shown in the images). We speculate that this is how the agent models tunneling: rather than having a high-level concept of digging a tunnel, it simply strikes wherever it has managed to strike already.”</p>	Breakout	Attention	Y	N

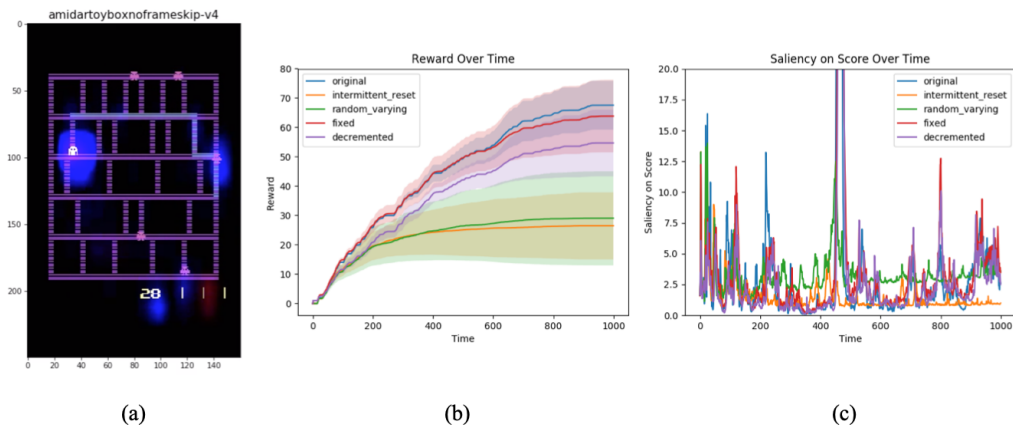


Figure 6: Interventions on displayed score in Amidar. (a) an example perturbation saliency map for an episode of Amidar; (b) reward over time for different interventions on score; (c) saliency on displayed score over time.

“Figs. 4(e)-(f) illustrate how the Dense FLS model playing Seaquest has learned to attend to in-game objects and, importantly, the oxygen bar at the bottom of the screen. As the oxygen bar is nearing depletion, attention focuses around it, and the submarine reacts by rising to refill its air supply.”

Seaquest      Attention      Y      N

“Figs. 4(g)-(h) are two consecutive frames where an agent detects a target appearing from the left side of the screen. The bottom part of the screenshots shows how attention in the bottom left corner lights up as soon as a tiny part of the target, only a few pixels wide, appears from the left edge of the screen. In the next frame, the agent will turn left and shoot the target (not shown here). However, the agent completely ignores targets in the top part of the screen, and its attention does not move as they move (also not shown).”

Breakout      Attention      Y      N

---

Wang et al. (2015)

“The value stream learns to pay attention to the road. The advantage stream learns to pay attention only when there are cars immediately in front, so as to avoid collisions.”

Enduro      Jacobian      Y      N

## B EVALUATION OF HYPOTHESES ON AGENT BEHAVIOR

### B.1 CASE STUDY 2: AMIDAR SCORE

We further evaluated the effects of interventions on displayed score (Section 5) in Amidar on perturbation and Jacobian saliency maps. These results are presented in Figures 6 and 7, respectively.

### B.2 CASE STUDY 3: AMIDAR ENEMY DISTANCE

We further evaluated the relationship between player-enemy distance and saliency (Section 5) in Amidar on perturbation and Jacobian saliency maps. These results are presented in Figures 8 and 9, respectively.

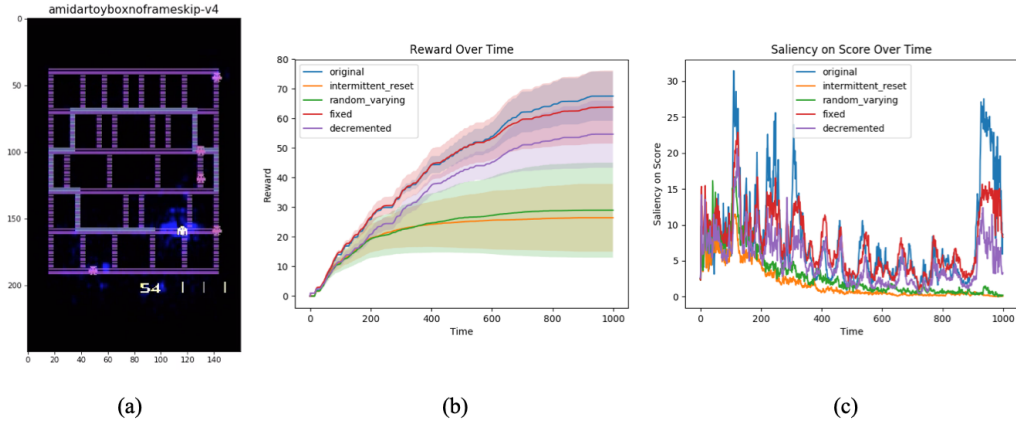


Figure 7: Interventions on displayed score in Amidar. (a) an example Jacobian saliency map for an episode of Amidar; (b) reward over time for different interventions on score; (c) saliency on displayed score over time.

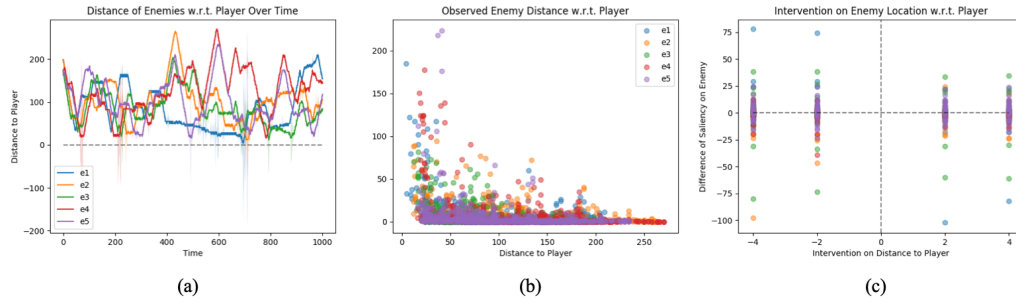


Figure 8: (a) shows the distance to player of each enemy in Amidar over time where saliency intensity is represented by the shade behind each enemy.; (b) shows the correlation between the distance to player and corresponding saliency for each enemy; (c) shows saliency upon intervening on enemy position.

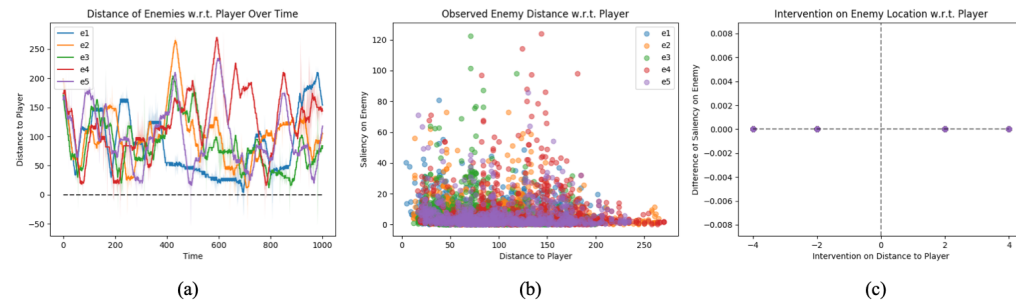


Figure 9: (a) shows the distance to player of each enemy in Amidar over time where saliency intensity is represented by the shade behind each enemy.; (b) shows the correlation between the distance to player and corresponding saliency for each enemy; (c) shows saliency upon intervening on enemy position.