

Explainable Reinforcement Learning through a Causal Lens

Prashan Madumal, Tim Miller, Liz Sonenberg, Frank Vetere

Victoria, Australia

pmathugama@student.unimelb.edu.au, {tmiller, l.sonenberg, f.vetere}@unimelb.edu.au

Abstract

Prominent theories in cognitive science propose that humans understand and represent the knowledge of the world through causal relationships. In making sense of the world, we build *causal models* in our mind to encode cause-effect relations of events and use these to *explain* why new events happen by referring to counterfactuals — things that did not happen. In this paper, we use causal models to derive causal explanations of the behaviour of model-free reinforcement learning agents. We present an approach that learns a *structural causal model* during reinforcement learning and encodes causal relationships between variables of interest. This model is then used to generate explanations of behaviour based on counterfactual analysis of the causal model. We computationally evaluate the model in 6 domains and measure performance and task prediction accuracy. We report on a study with **120** participants who observe agents playing a real-time strategy game (Starcraft II) and then receive explanations of the agents' behaviour. We investigate: 1) participants' understanding gained by explanations through task prediction; 2) explanation satisfaction and 3) trust. Our results show that causal model explanations perform better on these measures compared to two other baseline explanation models.

Driven by lack of trust from users and proposed regulations, there are many calls for Artificial Intelligence (AI) systems to become more transparent, interpretable and explainable. This has renewed the interest in Explainable AI (XAI), which has been explored since the expert systems era (Chandrasekaran, Tanner, and Josephson 1989). A key pillar of XAI is *explanation*, a justification given for decisions and actions of the system.

However, much research and practice in XAI pays little attention to *people* as intended users of these systems (Miller 2018b). If we are to build systems that are capable of providing 'good' explanations, it is plausible that explanation models should mimic models of human explanation (De Graaf and Malle 2017). Thus, to build XAI models it is essential to begin with a strong understanding of how people define, generate, select and evaluate explanations.

There is a wealth of pertinent literature in cognitive psychology that explore the nature of explanations and how

people understand them. As humans, we view the world through a causal lens (Sloman 2005), building mental models with causal relationships to act in the world, to understand new events and also to *explain* events. Importantly, causal models give people the ability to consider *counterfactuals* — events that did not happen, but could have under different situations. Although this notion of causal explanation is also backed by literature in philosophy and social psychology (Hilton 2007), causality and counterfactuals are only just becoming more prevalent in XAI. Further, compared to the burst of XAI research in supervised learning, explainability in model-free reinforcement learning is hardly explored.

We introduce an *action influence* model for model-free reinforcement learning (RL) agents and provide a formalisation of the model using structural causal models (Halpern and Pearl 2005). Action influence models approximate the causal model of the environment relative to actions taken by an agent. Our approach differs from previous work in explainable RL in that we use causal models to generate *contrastive* explanations for *why* and *why not* questions, which previous models lack. Given assumptions about the direction of causal relationships between variables, during the policy learning process, we also learn the quantitative influences that actions have on variables. Which enable our model to reason approximately about counterfactual states and actions. We define how to generate explanations for 'why?' and 'why not?' questions from the action influence model. We define *minimally complete* explanations taking inspiration from social psychology literature (McClure and Hilton 1997).

We computationally evaluated our approach on 6 RL benchmarks domains using 6 different RL algorithms. Results indicate that these models are robust and accurate enough to perform task prediction (Hoffman et al. 2018, p.12) with a negligible performance impact. We conducted a human study using the implemented model for RL agents trained to play the real-time strategy game *Starcraft II*. Experiments were run for **120** participants, in which we evaluated the participants' performance in task prediction, explanation satisfaction, and trust. Results show that our model performs better than the tested baseline, but its impact on

trust is not statistically significant.

The main contribution of this paper is twofold: 1) We introduce and formalise the *action influence* model based on structural causal models and present definitions to generate explanations; 2) We conduct a between-subject human study to evaluate the proposed model with baselines.

Related Work

There exists a substantial body of literature that explores explaining the policies and actions of Markov Decision Processes (MDP), though most of them do not explicitly focus on reinforcement learning. Elizalde et al. (2009) generated explanations by selecting and using ‘relevant’ variables of states of factored MDPs, evaluated by domain experts. Taking the long term effect an action has, Khan, Poupart, and Black (2009) proposed generating sufficient and minimal explanations for MDPs using domain independent templates.

Policy explanations in human-agent interaction settings have been used to achieve transparency (Hayes and Shah 2017) and provide summaries of the policies (Amir and Amir 2018). Explanation in reinforcement learning has been explored, using interactive RL to generate explanations from instructions of a human (Fukuchi et al. 2017) and to provide contrastive explanations (van der Waa et al. 2018). Soft decision trees have been used to generate more interpretable policies (Coppens et al. 2019), and reward decomposition has been utilized to provide minimum sufficient explanations in RL (Juozapaitis et al. 2019). However, these explanations are not based on an underlying causal model.

Other work on causal explanation has focused on scientific explanations (Salmon 1984) and explanations using causal trees (Nielsen, Pellet, and Elisseeff 2012). Although some recent work has emphasized the importance of causal explanation for explainable AI systems (Miller 2018b; 2018a; Madumal et al. 2019; Madumal 2019), work on generating explanations from causal explanation models for MDPs and RL agents have been absent.

Causal Models for Explanations

In this section, we introduce the *action influence model*, which is based on *structural causal models* of Halpern and Pearl (2005). We first introduce the Starcraft II domain, a partially observable real-time strategy game environment with a large state and action space. For the purpose of implementing RL agents for explanation, we use a toned-down version of the full Starcraft II 1v1 match (an adversarial scenario) with 4 actions and 9 state variables for the agent’s model (see Figure 1). In the following sections we use this Starcraft II scenario accompanied by Figure 1 as our running example.

Preliminaries : Structural Causal Models

Structural causal models (SCMs) (Halpern and Pearl 2005) represent the world using random variables, divided into exogenous (external) and endogenous (internal), some of which might have causal relationships which each other.

These relationships can be described with a set of *structural equations*. Formally, a *signature* S is a tuple $(\mathcal{U}, \mathcal{V}, \mathcal{R})$, where \mathcal{U} is the set of exogenous variables, \mathcal{V} the set of endogenous variables, and \mathcal{R} is a function that denotes the range of values for every variable $\mathcal{Y} \in \mathcal{U} \cup \mathcal{V}$.

Definition 1. A *structural causal model* is a tuple $M = (\mathcal{S}, \mathcal{F})$, where \mathcal{F} denotes a set of structural equations, one for each $X \in \mathcal{V}$, such that $F_X : (\times_{U \in \mathcal{U}} \mathcal{R}(U)) \times (\times_{Y \in \mathcal{V} - \{X\}} \mathcal{R}(Y)) \rightarrow \mathcal{R}(X)$ give the value of X based on other variables in $\mathcal{U} \cup \mathcal{V}$. That is, the equation F_X defines the value of X based on some other variables in the model.

A *context* \vec{u} is a vector of unique values of each exogenous variable $u \in \mathcal{U}$. A *situation* is defined as a model/context pair (M, \vec{u}) . An *instantiation* is defined by assigning variables the values corresponding to those defined by their structural equations. An *actual cause* of an event φ is a vector of endogenous variables and their values such that there is some counterfactual context in which the variables in the cause are different and the event φ does not occur. An explanation is those causes that an explainee does not already know. For a more complete review of SCM’s we direct the reader to (Halpern and Pearl 2005).

Causal Models for Reinforcement Learning Agents

Our intent in this paper is not to provide explanations of *evidence* from the environment, but to provide explanations of the agent’s behaviour based on the knowledge of how actions influence the environment. As such, we extend the notion of SCMs to include actions as part of the causal relationships.

We incorporate *action influence* models for MDP-based RL agents, extending SCMs with the addition of actions. An MDP is a tuple $(\mathcal{S}, \mathcal{A}, \mathcal{T}, \mathcal{R}, \gamma)$, where \mathcal{S} and \mathcal{A} give state and action spaces respectively (here we assume the state and action space is finite and state features are described by a set of variables ϕ); $\mathcal{T} = \{P_{sa}\}$ a set of state transition functions (P_{sa} denotes state transition distribution of taking action a in state s); $\mathcal{R} : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is a reward function and $\gamma = [0, 1)$ is a discount factor. The objective of an RL agent is to find a policy π that maps states to actions maximizing the expected discounted sum of rewards. We define the action influence model for RL agents as follows.

Formally, a signature S_a for an action influence model is a tuple $(\mathcal{U}, \mathcal{V}, \mathcal{R}, \mathcal{A})$, in which \mathcal{U} , \mathcal{V} , and \mathcal{R} are as in SCMs, and \mathcal{A} is the set of actions.

Definition 2. An *action influence model* is a tuple (S_a, \mathcal{F}) , where S_a is as above, and \mathcal{F} is the set of structural equations, in which we have multiple for each $X \in \mathcal{V}$ — one for each *unique* action set that influences X . A function $F_{X,A}$, for $A \in \mathcal{A}$, defines the causal effect on X from applying action A . The set of *reward variables* $X_r \subseteq \mathcal{V}$ are defined by the set of nodes with an out-degree of 0; that is, the set of sink nodes.

We define the *actual instantiation* of a model M as $M_{\vec{v} \leftarrow \vec{g}}$, in which \vec{S} is the vector of state variable values from an MDP. In an actual instantiation, we set the values of all state variables in the model, effectively making the exogenous variables irrelevant.

Figure 1 shows the graphical representation of Definition 2 as an action influence graph of the Starcraft II agent described in the previous section, with exogenous variables hidden. These *action influence models* are SCMs except that each edge is associated with an action. In the action influence model, each state variable has a *set* of structural equations: one for each *unique* incoming action. As an example, from Figure 1, variable A_n is causally influenced by S and B only when action A_m is executed, thus the structural equation $F_{A_n.A_m}(S, B)$ captures that relationship.

Explanation Generation

In this section, we present definitions that generate explanations from an action influence model. The process of explanation generation has 3 phases: 1) defining the qualitative causal relationships of variables as an action influence model; 2) learning the structural equations during RL; and 3) generating *explanans* from SCMs using the definitions given below.

We define an *explanation* as a pair that consist of: 1) an *explanandum*, the event to be explained; and 2) an *explanan*, the subset of causes given as the explanation (Miller 2018b). Consider the example ‘Why did you do P ?’ and the explanation ‘Because of Q ’. Here, the *explanandum* is P and *explanan* is Q . Identifying the *explanandum* from a question is not a trivial task. In this paper, we define explanations for questions of the form ‘Why A ?’ or ‘Why not A ?’ , where A is an action. In the context of a RL agent we define a *complete explanan* below.

Definition 3. A *complete explanan* for an action a under the actual instantiation $M_{\vec{v} \leftarrow \vec{s}}$ is a tuple $(\vec{X}_r = \vec{x}_r, \vec{X}_h = \vec{x}_h, \vec{X}_i = \vec{x}_i)$, in which \vec{X}_r is the vector of reward variables reached by following the causal chain of the graph to sink nodes; \vec{X}_h the vector of variables of the head node of action a , \vec{X}_i the vector of intermediate nodes between head and reward nodes, and $\vec{x}_r, \vec{x}_h, \vec{x}_i$ gives the values of these variables under $M_{\vec{v} \leftarrow \vec{s}}$.

Informally, this defines a complete explanan for action a as the complete causal chain from action a to any future reward that it can receive. From Figure 1, the causal chain for action A_s is depicted in bold edges, and the extracted explanan tuple $([S = s], [A_n = a_n], [D_u = d_u, D_b = d_b])$ is shown as darkened nodes. We use depth-first search to traverse the graph until all the sink nodes are reached from the head node of the action edge.

‘Why?’ Questions

Lim, Dey, and Avrahami (2009) found that the most demanded explanatory questions are *Why* and *Why not* questions. To this end, we focus on explanation generation for *why* and *why not* questions in this paper.

Minimally Complete Explanations Striking a balance between *complete* and *minimal* explanations depend on the epistemic state of the explainee (Miller 2018b). In this paper, we assume that we know nothing about the epistemic state of the explainee.

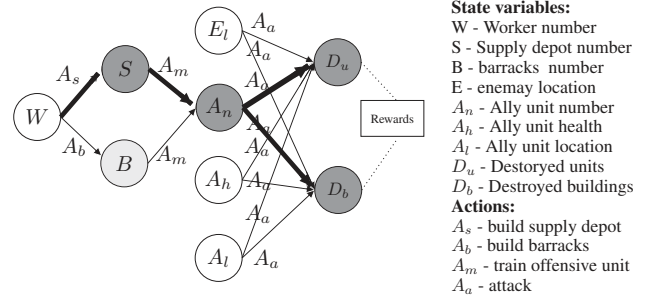


Figure 1: Action influence graph of a Starcraft II agent

Recall from the definition of *explanans* (Definition 3), a ‘complete’ explanation would include *explanans* of all the intermediate nodes between the head and reward node of the causal chain. Clearly, for a large graph, this risks overwhelming the explainee. For this reason, we define *minimally complete* explanations.

McClure and Hilton (1997) show that referring to the goal as being the most important for explaining actions. In our causal models, the rewards are the ‘goals’, but these alone do not form meaningful explanations because they are merely numbers. We define the human interpretable ‘goal’ using the variables in the predecessor nodes of the rewards. These define the immediate causes of the reward, and therefore which states will result in rewards. However, this alone is only a longer-term motivation for taking an action. As such, we also include the head node of the action edge as the immediate reason for doing the action. We use this model to define our *minimally complete* explanations.

Definition 4. A *minimally complete* explanation is a tuple $(\vec{X}_r = \vec{x}_r, \vec{X}_h = \vec{x}_h, \vec{X}_p = \vec{x}_p)$, in which $\vec{X}_r = \vec{x}_r$ and $\vec{X}_h = \vec{x}_h$ do not change from Definition 3, and $\vec{X}_p = \vec{x}_p$ is the vector of variables that are immediate predecessors of any variable in X_r within the causal chain, with \vec{x}_p the values in the actual instantiation.

Informally, for a complete causal chain, we take the first and last arcs of the causal chain, with their source and destination nodes, omitting intermediate nodes, as the minimal explanation. From Figure 1, for the action A_s , the minimally complete explanation is just the complete explanation, as there are no intermediate nodes.

Clearly, one could define other heuristics to decide which intermediate nodes to use as explanations, such as the knowledge of the explainee. However, for the purposes of this paper, we use this simple definition.

‘Why not?’ Questions

Why not questions let the explainee ask why an event has not occurred, thus allowing *counterfactuals* to be explained; something that is known to be a powerful explanation mechanism (Miller 2018b; Byrne 2019). Our model generates counterfactual explanations by comparing causal chains of the actual event occurred and the *explanandum* (counterfactual action). First, we define a *counterfactual instantiation*

that specifies the optimal state variable values under which the counterfactual action B would be chosen.

Definition 5. A *counterfactual instantiation* for a counterfactual action B is a model $M_{\vec{Z} \leftarrow \vec{S}_Z}$, where \vec{Z} gives the instantiation of all predecessor variables of action B with current state values *and* the instantiation of all successor nodes (of B) of the causal chain by forward simulating, using the structural equations.

Informally, this gives the ‘optimal’ conditions (according to the action influence model) under which we would select counterfactual action B , simulated through structural equations. We unravel this further in the Example 1 discussion using the Starcraft II scenario.

In the following definition, we use $\vec{X} = \vec{x}$ to represent the tuple $(\vec{X}_p = \vec{x}_p, \vec{X}_h = \vec{x}_h, \vec{X}_r = \vec{x}_r)$, and similar for $\vec{Y} = \vec{y}$ for readability.

Definition 6. Given a minimally complete explanation $\vec{X} = \vec{x}$ for action A under the actual instantiation, and a minimally complete explanation $\vec{Y} = \vec{y}$ for action B under the counterfactual instantiation $M_{\vec{Z} \leftarrow \vec{S}_Z}$ (from Definition 5), we define a *minimally complete contrastive explanation* as the tuple $(\vec{X}' = \vec{x}', \vec{Y}' = \vec{y}', \vec{X}_r = \vec{x}_r)$ such that \vec{X}' is the maximal set of variables in \vec{X} in which $(\vec{X}' = \vec{x}') \cap (\vec{Y}' = \vec{y}') \neq \emptyset$, where \vec{x}' is then contrasted with \vec{y}' . That is, we only explain things that are different between the actual and counterfactual. This corresponds to the *difference condition* (Miller 2018a). And \vec{X}_r gives the reward nodes of action A .

Intuitively, a contrastive explanation extracts the actual causal chain for the taken action A , and the counterfactual causal chain for the B , and finds the differences.

Example 1. Consider the question, asking why a Starcraft II agent built supply depots, rather than choosing to build barracks:

Question Why not *build_barracks* (A_b)?

Explanation Because it is more desirable to do action *build_supply_depot* (A_s) to have more Supply Depots (S) as the goal is to have more Destroyed Units (D_u) and Destroyed buildings (D_b).

First we get the *actual instantiation* $m = [W = 12, S = 1, B = 2, A_n = 22, D_u = 10, D_b = 7]$ (instantiation should include all variables in the current state, only the required ones are shown for readability). The causal chain for the *actual* action ‘why A_s ?’ would be as in Figure 1, and for the *counterfactual* action ‘why not A_b ?’ , the causal chain nodes would be $B \rightarrow A_n \rightarrow [D_u, D_b]$. We then get the *counterfactual instantiation* $m' = [W = 12, S = 3, B = 2, A_n = 22, D_u = 10, D_b = 7]$ using Definition 5. Applying the difference condition here, we obtain the minimally complete contrastive explanation (from Definition 6) as the tuple $([S = 1], [S = 3], [D_u = 10, D_b = 7])$ and contrast $[S = 1]$ with $[S = 3]$ to obtain the explanation of Example 1 (generated using a simple NLP template).

Learning Structural Causal Equations

Our approach so far relies on knowing the structural model, in particular, to determine the effects of counterfactual actions. *Why not* questions are inherently counterfactual (Balke and Pearl 1995), and having just the policy of an agent is not enough to generate explanations as counterfactuals refers to *possible* worlds that did not happen. Consider the Example 1, to generate this explanation, the optimal/maximum value of the state variable S is needed in the given time instance.

However, in model-free reinforcement learning, such environment dynamics are not known. And learning a model of the environment is a difficult problem. Though, when given a graph of causal relations between variables, learning a set of structural equations that are approximate yet ‘good enough’ to give counterfactual explanations maybe feasible.

To this end, we assume that a DAG specifying causal direction between variables is given, and learn the structural equations as multivariate regression models during the training phase of the RL agent. We perform experience replay (Mnih et al. 2015) by saving $e_t = (s_t, a_t, r_t, s_{t+1})$ at each time step t in a data set $D_t = \{e_1, \dots, e_t\}$. Then we update the sub-set of structural equations $F_{X.A}$ using a regression learner $\hat{\mathbb{L}}_{(s,a,r,s') \sim U(D)}$, in that we *only* update structural equations associated with the specific action in the experience frame, drawn uniformly as mini-batches from D . For example, from Figure 1, for any experience frame with the action A_s , only the equation $F_{S.A_s}(W)$ will be updated. Any regression learner can be used as the learning model $\hat{\mathbb{L}}$, such as multi-layer perceptron regressors.

While this approach may seem similar to learning environment dynamics of model-based RL methods, we only learn the structural equations, and we are only after an approximation that is good enough for explaining instances. Thus they can be approximate but still useful for explanation. Further, specifying the assumptions about the causal direction between variables is a much easier problem to encode by hand, and can be tested with the data.

Computational Evaluation

We evaluate *action influence models* in 5 OpenAI RL benchmark domains (Brockman et al. 2016) and in the Starcraft II domain. The goal of this evaluation is to determine if learning action influence models leads to models that are faithful to the problem. Task prediction accuracy and training time for the structural causal equations are measured. The purpose of task prediction is to evaluate if the model is accurate enough to predict what an agent will do next, under the assumption that if it is not, then the model will not be of use to a human explainee.

We computationally simulate task prediction using Algorithm 1. Here we instantiate all the equations (which are the set of regression models \mathcal{L}) with the values of the current state S of the agent. We identify the equation that has maximum difference with the predicted state variable value and the actual, then get the action associated with it. This is informed by the reasoning that the agent will try to follow the

Algorithm 1 Task Prediction:Action Influence Model

Input: trained regression models \mathcal{L} , current state S_t **Output:** predicted action a

```
1:  $\vec{F}_p \leftarrow []$ ; vector of predicted difference
2: for every  $\hat{L} \in \mathcal{L}$  do
3:    $P_y \leftarrow \hat{L} \cdot \text{predict}(S_{x,t})$ ; predict variable  $S_y$  at  $S_{t+1}$ 
4:    $\vec{F}_p \leftarrow |S_y - P_y|$ ; difference with actual  $S_y$  value
5: end for
6: return  $\max(\vec{F}_p) \cdot \text{getAction}()$ 
```

| Env - RL | Size | Accuracy (%) | | | Performance (s) | | |
|-------------------|------|--------------|------|------|-----------------|-------|------|
| | | LR | DT | MLP | LR | DT | MLP |
| Cartpole-PG | 4/2 | 83.8 | 81.6 | 86.0 | 0.007 | 0.018 | 0.03 |
| MountainCar-DQN | 3/3 | 69.7 | 57.8 | 69.6 | 0.020 | 0.037 | 0.32 |
| Taxi-SARSA | 4/6 | 68.2 | 74.2 | 67.9 | 0.001 | 0.001 | 0.49 |
| LunarLander-DDQN | 8/4 | 68.4 | 63.7 | 72.1 | 0.002 | 0.002 | 0.33 |
| BipedalWalker-PPO | 14/4 | 56.9 | 56.4 | 56.7 | 0.010 | 0.015 | 0.41 |
| Starcraft-A2C | 9/4 | 94.7 | 91.8 | 91.4 | 0.144 | 0.025 | 3.33 |

Table 1: Action influence model evaluation in 6 benchmark reinforcement learning domains (using different RL algorithms, PG, DQN etc.), measuring mean task prediction accuracy and training time of the structural causal equations in 100 episodes after training.

optimal policy, and the action with the biggest impact to correct the policy will be executed. The impact is measured by the above mentioned difference. This is itself an approximation, but is a useful guide for task prediction.

We use linear SGD regression (LR), decision tree regression (DT) and multilayer perceptron regression (MLP) as the learners that approximate the structural equations. We choose benchmark domains based on varying levels of complexity, size (state features/number of actions) and train them using various RL algorithms to demonstrate the robustness of the model. Table 1 summarises the results of task prediction and time taken to train the structural equations given the replay data.

Overall, the results show the model did a reasonable job of task prediction, providing evidence that this could be useful for explanations. Domains that have a clear causal structure (e.g Starcraft) performs best in task prediction. Considering the performance cost it incurs, there was little gained by using MLP to approximate the equations, where in most cases linear regression is adequate. Apart from the BipedalWalker domain, our model performs well in task prediction with a negligible performance hit. The bipedalWalker domain has continuous actions, which our current model cannot handle accurately. We plan to extend our model to continuous actions in future work.

Empirical Evaluation: Human Study

A human-grounded evaluation is essential to evaluate the explainability of a system, thus we carry out human-subject experiments involving explaining RL agents. We present two main hypotheses for the empirical evaluation; **H1**) Causal-

model-based explanations build better mental models of the agent leading to a better **understanding** of its strategies (We make the assumption here that there is no intermediate effect on the mental model from other sources); and **H2**) Better understanding of an agent’s strategies promotes **trust** in the agent.

Methodology: We use StarCraft II, a real-time strategy game and a popular RL environment (Vinyals et al. 2017) as the domain. We implemented a RL agent for our experiment that competes in the default map.

To evaluate hypothesis (H1), we use the method of *task prediction* (Hoffman et al. 2018). Task prediction can provide a quick view of the explainee’s mental model formed through explanations, where the task is for the participant to predict ‘What will the agent do next?’. We use the 5-point Likert *Explanation Satisfaction Scale* developed by Hoffman et al. (2018, p.39) to measure the subjective quality of explanations. To evaluate hypothesis (H2), we use the 5-point Likert *Trust Scale* of Hoffman et al. (2018, p.49). We obtained ethics approval from The University of Melbourne human research ethics committee (ID-1953619.1).

Experiment Design: We use a recording of a full gameplay video (22 min) with the RL agents playing against in-game bot AI. The experiment has 4 phases.

Phase 1 involves collecting demographic information and training the participants. Using five gameplay video clips, the participant is trained to understand and differentiate the actions of the agent.

In phase 2, a clip of the gameplay video (15 sec) is played in a web-based UI, with a textual description of the scene. The participant can select the question type (why/why not) and the action, which together forms a question ‘Why/Why not *action A*?’ . Then, the textual explanation for the question with a figure of the relevant sub-graph of the agent’s action influence graph is displayed. Explanations are pre-generated from our implemented algorithm. The participant can ask **multiple** questions in a single gameplay video. After every gameplay video, the participant completes the *Explanation Satisfaction Scale*. This process is repeated so we have data for each participant from five videos.

The third phase measures the *understanding* the explainee has after seeing the gameplay and the explanations. We measure understanding using the task prediction method as follows: the participant is presented with another gameplay video (10 sec), and presented with three selections of textual descriptions of what *action* the agent will do in *next* step; the participant selects an option, which includes ‘I don’t know’. We expect the participant is projecting forward the *local strategy* of the agent using their mental model. This mental model is formed through (or helped by) explanations seen in phase 1. This process is repeated for 8 tasks. In 4 of the task predictions, the behaviour is explainable using a causal chain previously seen in the training, but with different variable values. In the other 4 tasks, the behaviour is novel, but can be inferred by combining causal chains from different training tasks. In the fourth phase, the participant completes a 5-point *Trust Scale*.

We conducted the experiments on *Amazon MTurk*,




Figure 2: Box plot of task prediction scores of explanation models, T=total score, F=familiar score, N=novel score (higher is better, means represented as bold dots).

a crowd-sourcing platform popular for obtaining data for human-subject experiments (Buhrmester, Kwang, and Gosling 2011). The experiment was fully implemented in an interactive web-based environment. We excluded noisy data of users in 3 ways. First, we tested participants to ensure they had learnt about the agent’s actions by prompting them to identify them. If the participant failed this, the experiment did not proceed. Then, for participants who completed, we omitted their data from analysis based on two criteria: 1) if the threshold of the time the participant spent on viewing explanations and answering tasks is below a few seconds, which was deemed too short to learn anything useful; and 2) if the participant’s textual responses to explain their task prediction choice were gibberish text or a 1-2 word response, as this indicated lack of engagement and care in the task. We controlled for language by only recruiting participants from the US.

Experiment Parameters: The experiment was run with 4 independent variables. We tested abstract (C) and detailed (D) versions of our action influence models and 2 baseline models described below: 1) Gameplay video without any explanations (N); 2) Relevant variable explanations (R). These explanations are generated using state relevant variables using *template 1* of Khan, Poupart, and Black (2009, p.3) and visualized through a state-action graph, e.g ‘Action A is likely to increase *relevant variable P*’; 3) Detailed action influence model explanations, where the causal graph is augmented to include atomic actions.

We ran experiments for **120** participants, allocated evenly to the independent variables. Each experiment ran for approximately 40 minutes. We scored each participant on task prediction, 2 points for a correct prediction; 1 for responding ‘I don’t know’ and 0 for an incorrect prediction for a total of 16 points. Scores were tallied. We compensated each participant with 8.5USD. Of the 120 participants, 36 were female, 82 male and 2 were not given. Participants were aged between 19 to 59 ($\mu = 34.2$) and had an average self-rated gaming experience and Starcraft II experience of 3.38 and 2.02 (5-point Likert) respectively.

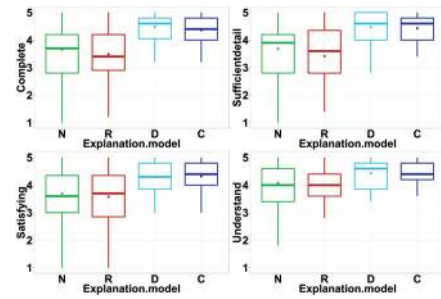


Figure 3: Box plot of explanation quality (likert scale 1-5, higher is better, means represented as dots).




Figure 4: Box plot of trust (likert scale 1-5, higher is better, means represented as dots).

Results

Task Prediction: For the first hypothesis, the corresponding null and alternative hypotheses are: 1) $H_0 : \mu_C = \mu_R = \mu_D = \mu_N$; 2) $H_1 : \mu_C \geq \mu_R$; 3) $H_2 : \mu_C \geq \mu_D$; 4) $H_3 : \mu_C \geq \mu_N$, where abstract causal explanations (our model), detailed causal explanations, relevant variable explanations, and no explanations are given by C, D, R, and N respectively.

We conduct one-way ANOVA (Figure 2 illustrates the task score variance with explanation models). We obtained a p-value of **0.003** ($\mu_C = 10.90$, $\mu_D = 10.20$, $\mu_R = 8.97$, $\mu_N = 8.53$), thus we conclude there are significant differences between models on task prediction scores. We performed Tukey multiple pairwise-comparisons to obtain the significance between groups. From Table 2, the differences between the causal explanation model paired with other explanation models are significant for C-R and C-N pairs with p-values of **0.006** and **0.034**. Additionally, we calculate the effect of the number of questions on the score, and obtain no statistical correlation using a correlation test (number of questions vs score, $p = 0.33$, model C) among same models. Because participants could select ‘I don’t know’ and receive 8 out of 16, we also further analyse scores based on 2 = ‘correct’, 0 = ‘incorrect or ‘I don’t know’, and obtain results that are still significant ($p=0.004$), means (C=10.90, D=10.10, R=8.93, N=8.47), for model pairs (C-N $p=0.005$, C-R $p=0.035$). We conducted Pearson’s Chi-Square as a non-parametric test on the task prediction scores, which showed significant results ($p\text{-value} = 0.008$, X-squared = 17.281).

| Model pair | mean-diff | lwr | upr | p-value |
|------------|--------------|--------|-------|--------------|
| C - N | 2.400 | 0.534 | 4.265 | 0.006 |
| C - R | 1.966 | 0.101 | 3.832 | 0.034 |
| D - N | 1.666 | -0.198 | 3.532 | 0.097 |
| D - R | 1.233 | -0.632 | 3.098 | 0.316 |
| C - D | 0.733 | -1.132 | 2.598 | 0.735 |
| R - N | 0.433 | -1.432 | 2.298 | 0.930 |

Table 2: Pairwise-comparisons of explanation models of task prediction scores (higher positive diff is better)

| Metric | Mdl-pair | Mean-dif | Median-dif | p-val |
|------------|----------|----------|------------|--------------|
| Complete | C-N | 0.707 | 0.700 | 0.061 |
| | C-R | 0.873 | 1.000 | 0.012 |
| Sufficient | C-N | 0.746 | 0.700 | 0.039 |
| | C-R | 1.013 | 1.000 | 0.002 |
| Satisfying | C-N | 0.633 | 0.800 | 0.082 |
| | C-R | 0.740 | 0.700 | 0.029 |
| Understand | C-N | 0.326 | 0.400 | 0.497 |
| | C-R | 0.400 | 0.400 | 0.316 |

Table 3: Explanation quality (likert scale data 1-5)

Therefore we reject H_0 and H_2 and accept all other alternative hypotheses. Our results show that causal model explanations lead to a significantly better **understanding** of agent’s strategies than the 2 baselines we evaluated, especially against previous models of relevant explanations. Participants did slightly worse on tasks with novel behaviour.

Explanation Quality: Figure 3 depicts the likert scale data of explanation metrics (understand, satisfying, sufficient detail and complete) for aggregated video explanations of explanation models. As before we performed a pair-wise ANOVA test, results are summarised in Table 3. Our model obtained statistically significant results and outperformed the benchmark ‘relevant explanation’ (R) for all metrics except ‘Understand’.

Trust: For the second main hypothesis (H2) that investigate whether explanation models promote trust, the obtained p-values for trust metrics *confident*, *predictable*, *reliable* and *safe* were not statistically significant (using pair-wise ANOVA). Although the difference is not significant we can see causal models have high means and medians (see Figure 4). We conclude that while the explanation quality and scores are significantly better for our model, to promote trust further interaction is necessary; or perhaps our RL agent is simply not a trustworthy Starcraft II player.

We further analysed self-reported demographic data to see if there is a correlation between task prediction scores and self-reported Starcraft II experience level (5-point Likert). Pearson’s correlation test was not significant ($p=0.45$) thus we conclude there is no correlation between scores and experience level. This can possibly be attributed to our Starcraft II scenario differing from the standard game.

A limitation of our experiment is that we made a strong linearity assumption for Starcraft II, which enabled linear

regression to learn SCMs for a relatively small number (9) of state variables.

Conclusion

In this paper, we introduced *action influence* models for model-free reinforcement learning agents. Our approach learns a structural causal model (SCM) during reinforcement learning and has the ability to generate explanations for *why* and *why not* questions by counterfactual analysis of the learned SCM. We computationally evaluated our model in 6 benchmark RL domains on *task prediction*. We then conducted a human study ($n=120$) to evaluate our model on 1) task prediction, 2) explanation ‘goodness’ and 3) trust. Results show that our model performs significantly better in the first 2 evaluation criteria. One weakness of our approach is that the causal model must be given beforehand. Future work includes using epistemic knowledge of the explainee to provide explanations that are more targeted, and extending the model to continuous domains.

Acknowledgment

This research was supported by the University of Melbourne research scholarship (MRS) and by Australian Research Council Discovery Grant DP190103414: *Explanation in Artificial Intelligence: A Human-Centred Approach*. The authors also thank Fatma Faruq for the valuable feedback on early drafts of this paper.

References

- Amir, D., and Amir, O. 2018. Highlights: Summarizing agent behavior to people. In *Proc. of the 17th International conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*.
- Balke, A., and Pearl, J. 1995. Counterfactuals and policy analysis in structural models. In *Proceedings of the Eleventh conference on Uncertainty in artificial intelligence*, 11–18. Morgan Kaufmann Publishers Inc.
- Brockman, G.; Cheung, V.; Pettersson, L.; Schneider, J.; Schulman, J.; Tang, J.; and Zaremba, W. 2016. Openai gym.
- Buhrmester, M.; Kwang, T.; and Gosling, S. D. 2011. Amazon’s mechanical turk: A new source of inexpensive, yet high-quality, data? *Perspectives on psychological science* 6(1):3–5.
- Byrne, R. M. J. 2019. Counterfactuals in explainable artificial intelligence (xai): Evidence from human reasoning. In *Proceedings of the 28th International Joint Conferences on Artificial Intelligence*, 6276–6282.
- Chandrasekaran, B.; Tanner, M. C.; and Josephson, J. R. 1989. Explaining control strategies in problem solving. *IEEE Intelligent Systems* (1):9–15.
- Coppens, Y.; Efthymiadis, K.; Lenaerts, T.; and Nowe, A. 2019. Distilling deep reinforcement learning policies in soft decision trees. In Miller, T.; Weber, R.; and Magazzeni, D., eds., *Proceedings of the IJCAI 2019 Workshop on Explainable Artificial Intelligence*, 1–6.

- De Graaf, M. M. A., and Malle, B. F. 2017. How People Explain Action (and Autonomous Intelligent Systems Should Too). *AAAI 2017 Fall Symposium on "AI-HRI"* 19–26.
- Elizalde, F., and Sucar, L. E. 2009. Expert evaluation of probabilistic explanations. In *ExaCt*, 1–12.
- Elizalde, F.; Sucar, L. E.; Luque, M.; Diez, J.; and Reyes, A. 2008. Policy explanation in factored markov decision processes. In *Proceedings of the 4th European Workshop on Probabilistic Graphical Models (PGM 2008)*, 97–104.
- Fukuchi, Y.; Osawa, M.; Yamakawa, H.; and Imai, M. 2017. Autonomous self-explanation of behavior for interactive reinforcement learning agents. In *Proceedings of the 5th International Conference on Human Agent Interaction*, 97–101. ACM.
- Halpern, J. Y., and Pearl, J. 2005. Causes and explanations: A structural-model approach. part II: Explanations. *The British journal for the philosophy of science* 56(4):889–911.
- Hayes, B., and Shah, J. A. 2017. Improving robot controller transparency through autonomous policy explanation. In *Proceedings of the 2017 ACM/IEEE international conference on human-robot interaction*, 303–312. ACM.
- Hilton, D. 2007. Causal explanation. *Social psychology: Handbook of basic principles* 232–253.
- Hoffman, R. R.; Mueller, S. T.; Klein, G.; and Litman, J. 2018. Metrics for explainable AI: Challenges and prospects. *arXiv preprint arXiv:1812.04608*.
- Juozapaitis, Z.; Koul, A.; Ferm, A.; Erwig, M.; and Doshi-Velez, F. 2019. Explainable reinforcement learning via reward decomposition. In Miller, T.; Weber, R.; and Magazzini, D., eds., *Proceedings of the IJCAI 2019 Workshop on Explainable Artificial Intelligence*, 47–53.
- Khan, O. Z.; Poupart, P.; and Black, J. P. 2009. Minimal sufficient explanations for factored markov decision processes. In *ICAPS*.
- Lim, B. Y.; Dey, A. K.; and Avrahami, D. 2009. Why and why not explanations improve the intelligibility of context-aware intelligent systems. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2119–2128. ACM.
- Madumal, P.; Miller, T.; Sonenberg, L.; and Vetere, F. 2019. A grounded interaction protocol for explainable artificial intelligence. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, 1033–1041.
- Madumal, P. 2019. Explainable agency in intelligent agents: Doctoral consortium. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, 2432–2434.
- McClure, J., and Hilton, D. 1997. For you can't always get what you want: When preconditions are better explanations than goals. *British Journal of Social Psychology* 36(2):223–240.
- Miller, T. 2018a. Contrastive explanation: A structural-model approach. *arXiv preprint arXiv:1811.03163*.
- Miller, T. 2018b. Explanation in artificial intelligence: Insights from the social sciences. *Artificial Intelligence*.
- Mnih, V.; Kavukcuoglu, K.; Silver, D.; Rusu, A. A.; Veness, J.; Bellemare, M. G.; Graves, A.; Riedmiller, M.; Fidjeland, A. K.; Ostrovski, G.; et al. 2015. Human-level control through deep reinforcement learning. *Nature* 518(7540):529.
- Nielsen, U.; Pellet, J.-P.; and Elisseff, A. 2012. Explanation trees for causal bayesian networks. *arXiv preprint arXiv:1206.3276*.
- Salmon, W. C. 1984. *Scientific explanation and the causal structure of the world*. Princeton University Press.
- Slooman, S. 2005. *Causal models: How people think about the world and its alternatives*. Oxford University Press.
- van der Waa, J.; van Diggelen, J.; Bosch, K. v. d.; and Neerinx, M. 2018. Contrastive explanations for reinforcement learning in terms of expected consequences. *arXiv preprint arXiv:1807.08706*.
- Vinyals, O.; Ewalds, T.; Bartunov, S.; Georgiev, P.; Vezhnevets, A. S.; Yeo, M.; Makhzani, A.; Küttler, H.; Agapiou, J.; Schrittwieser, J.; et al. 2017. Starcraft ii: A new challenge for reinforcement learning. *arXiv preprint arXiv:1708.04782*.