

Integrating Counterfactual Simulations with Language Models for Explaining Multi-Agent Behaviour

Bálint Gyevnár
University of Edinburgh
Edinburgh, United Kingdom
balint.gyevnar@ed.ac.uk

Stefano V. Albrecht
DeepFlow
London, United Kingdom
stefano.albrecht@deepflow.com

Christopher G. Lucas
University of Edinburgh
Edinburgh, United Kingdom
c.lucas@ed.ac.uk

Shay B. Cohen
University of Edinburgh
Edinburgh, United Kingdom
scohen@inf.ed.ac.uk

ABSTRACT

Explainability is vital for users' trust calibration in multi-agent systems (MAS), but explainable MAS face challenges due to complex environments, the human factor, and non-standardised evaluation. Leveraging the counterfactual effect size model and LLMs, we propose *Agentic eXplanations via Interrogative Simulation (AXIS)*. AXIS generates human-centred action explanations for multi-agent policies by having an LLM interrogate an environment simulator using prompts like *whatif* and *remove* to observe and synthesise counterfactual information over multiple rounds. We evaluate AXIS on autonomous driving across ten scenarios for five LLMs with a comprehensive methodology combining robustness, subjective preference, correctness, and goal/action prediction with an external LLM as evaluator. Compared to baselines, AXIS improves perceived explanation correctness by at least 7.7% across all models and goal prediction accuracy by 23% for four models, with comparable action prediction accuracy, achieving the highest scores overall. We release all code and an appendix at: <https://github.com/gyevnarb/axis>.

KEYWORDS

Multi-Agent Systems; Explainable AI; Causality; Large Language Models; Autonomous Driving

ACM Reference Format:

Bálint Gyevnár, Christopher G. Lucas, Stefano V. Albrecht, and Shay B. Cohen. 2026. Integrating Counterfactual Simulations with Language Models for Explaining Multi-Agent Behaviour. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 11 pages. <https://doi.org/10.65109/MCSD1905>

1 INTRODUCTION

Autonomous agents in multi-agent systems (MAS), e.g. self-driving cars and trading algorithms, are increasingly capable. They also contribute to the AI-fuelled erosion of human autonomy and oversight, as exemplified by fatal self-driving car incidents [44] and multiple stock market flash crashes [47]. Recently, concerns over emergent

risks from advanced artificial intelligence (AI) in MAS, e.g. collusion and goal misalignment, have also gained attention [9]. As such, the trustworthiness and reliability of these technologies are called into question [9, 38], fuelling urgent calls for AI governance [e.g. 20, 43].

A core principle of AI governance is transparency, partly facilitated by systems' ability to be explainable [34, 45, 51, 81]. Explanations are useful for showing regulatory compliance, as they can help with calibrating people's trust [27, 65, 85], debugging faulty systems [12, 42, 78, 83], and facilitating effective human-agent teaming [33, 46, 63]. For autonomous agents, explainable sequential decision-making (SDM) has been motivated by these benefits to develop algorithms for SDM and control tasks, with most methods focusing on post-hoc feature attribution [61, 80].

However, there are at least three roadblocks that hinder progress on explainable SDM. **(1) Complexity of environments** [61, 71, 80], which includes high-dimensional state/action spaces, partial observability, non-stationary dynamics, and (in MAS) coupled decision-making, complicating the disentanglement of causal relationships between inputs and actions. **(2) The human factor** [22, 28, 61, 62], which involves the varied cognitive capabilities, domain expertise, and explanation preferences of stakeholders. **(3) Lack of standardised evaluation** [21, 35, 41, 55, 64], which limits cross-study comparisons and creates methodological fragmentation where researchers use disparate definitions of what a good explanation is.

For MAS, our key insight in tackling these roadblocks is that the actions of agents can be explained in terms of behaviour considering the *causal relationships* among agents' actions, beliefs, desires, intentions and outcomes [16, 17, 56, 73, 84]. We refer to these explanations as *action explanations*, which are a complementary approach to the lower-level, mechanistic explanations of SDM policies, e.g. saliency maps. For creating action explanations, we make three observations. First, given a joint action policy in a MAS, we can simulate interventions on the actions of agents, leading to counterfactual worlds where we learn about the reactions of other agents. By summarising counterfactual information, useful explanations can be created [2, 7, 30, 31, 36]. Second, people prefer explanations that align with their cognitive processes [33, 46, 62, 63], and contemporary computational models of causal reasoning suggest that people simulate counterfactual worlds when explaining phenomena [26, 57, 72]. Third, LLMs excel at summarising information [11, 15, 82, 86] and can be useful as an evaluator [23, 32].



This work is licensed under a Creative Commons Attribution International 4.0 License.

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). <https://doi.org/10.65109/MCSD1905>

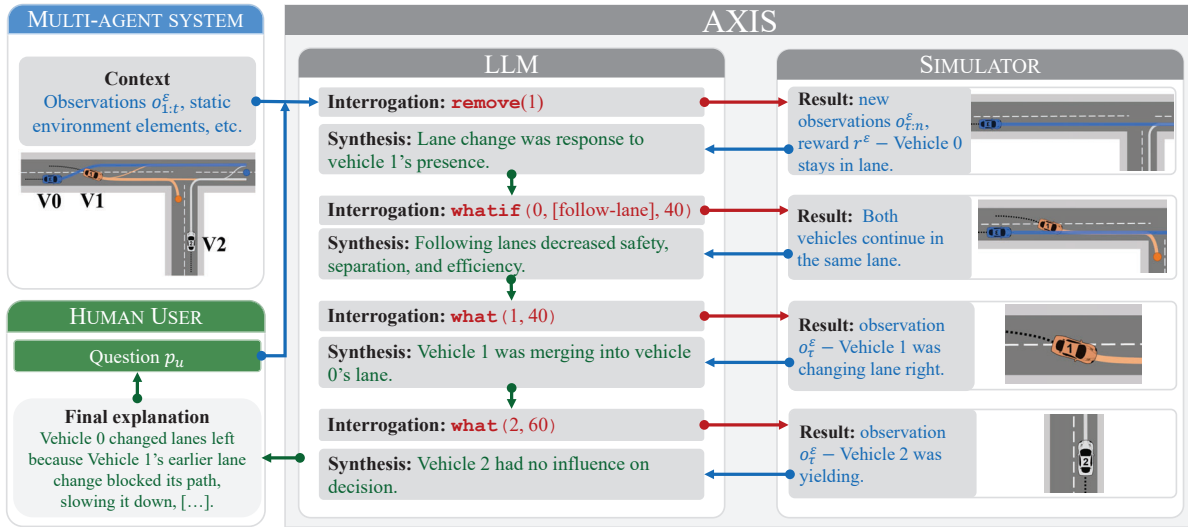


Figure 1: When a user asks a question about an agent’s actions, AXIS retrieves the current context, which the LLM uses to interrogate a simulator of the environment, eliciting counterfactual information about the actions of agents. The LLM synthesises these counterfactuals into explanations, repeating the process over multiple rounds. At the end, AXIS prompts the LLM to synthesise all counterfactuals into a final explanation, which is returned to the user as the answer to their question. Arrow legend: blue pass symbolic data with conversion to text, green pass explanations, red pass interrogation prompts.

We operationalise these observations in a new framework called *Agentic eXplanations via Interrogative Simulation (AXIS)*, which generates human-centred action explanations of a given multi-agent action policy. The overview and an example trace of AXIS execution is shown in Figure 1. We evaluate AXIS on the safety-critical task of motion planning for autonomous driving (AD) [75] across five LLMs using ten scenarios, including ones with irrational agent behaviour and partial observability, with up to three user prompts for each scenario. We propose a comprehensive evaluation methodology for explainable SDM, combining scenarios with qualitative and computational analyses, subjective preference and perceived correctness scores, and objective accuracy on the downstream tasks of goal and action prediction; using an external LLM as evaluator.

In our evaluation, AXIS improved the perceived correctness by at least 7.7% across all models and goal prediction accuracy by 23% for four models, correctly focusing on the most informative `remove` and `whatif` counterfactuals. AXIS also improved or gave comparable results on action prediction for four models and subjective preference for three models. In summary our contributions are:

- (1) We formalise the **problem of action explanation** for MAS described as partially observable stochastic games;
- (2) We propose **AXIS: Agentic Explanations via Interrogative Simulation** for generating action explanations for any MAS based on counterfactual simulations and LLMs;
- (3) We design an **evaluation methodology for explainable SDM** combining scenarios, qualitative correctness analysis, computational robustness to hyperparameters, and measures of subjective quality and objective actionability;
- (4) We conduct experiments in **simulated autonomous driving scenarios** showing higher perceived explanation correctness and prediction accuracy compared to baselines.

2 BACKGROUND

We focus on multi-agent sequential decision-making systems, given as partially observable stochastic games (POSG), using the definitions of Albrecht et al. [4] and illustrative examples from AD. We extend POSGs with the notion of goals following [3]. The notation $s_{a:b}$ is used for the ordered sequence (s_a, \dots, s_b) .

Let \mathcal{I} be the set of indexed agents in the system, e.g. vehicles on a road section. At timestep $t \in \mathbb{N}^+$, each agent $i \in \mathcal{I}$ is in local state $s_t^i \in \mathcal{S}^i$, e.g. position coordinates and velocity, such that the environment is in joint state $s_t \in \mathcal{S}$, where $\mathcal{S} = \times_{i \in \mathcal{I}} \mathcal{S}^i$. Agent i receives a local observation $o_t^i \in \mathcal{O}^i$, e.g. states and actions of other vehicles, which probabilistically depends on s_t through $p(o_t^i | s_t)$, with the joint observation written as $o_t \in \mathcal{O}$. In reaction to observations, agent i selects an action $a_t^i \in \mathcal{A}^i$, e.g. acceleration and steering, following its policy $\pi^i(a_t^i | o_{1:t}^i)$. Once all agents have acted, resulting in joint action $a_t \in \mathcal{A}$, the environment transitions to state $s_{t+1} \sim T(\cdot | s_t, a_t)$ and each agent receives an immediate reward $R^i(s_t, a_t, s_{t+1})$, which may be decomposed into a weighted reward component vector $\vec{R}^i(s_t, a_t, s_{t+1})$, e.g. acceleration, presence of collision. The system terminates once it reaches a terminal state in $\bar{\mathcal{S}} \subset \mathcal{S}$ or when a maximum number of steps is reached. Assuming that local observations and local actions are mutually independent, that is $p(o_t | s_t) = \prod_i p(o_t^i | s_t)$ and $\pi(a_t | o_{1:t}) = \prod_i \pi^i(a_t^i | o_{1:t}^i)$, then the probability of any joint state sequence $p(s_{1:t})$ is given by:

$$p(s_{1:t}) = \prod_{n=1}^{t-1} \int_{\mathcal{O}} \int_{\mathcal{A}} T(s_{n+1} | s_n, a_n) \pi(a_n | o_{1:n}) p(o_n | s_n) d o_n d a_n.$$

In the goal-extended version of a POSG, agent i is aiming to reach a goal state $g^i \in \mathcal{G}^i \subset \mathcal{S}^i$, defined as any partial local state description, e.g. destination coordinates. We say that the joint state

sequence $s_{1:t}$ reaches g^i and write $\text{Reach}(s_{1:t}, g^i)$ if $s_t^i \in g^i$ and $\forall n < t : s_n^i \notin g^i$, i.e. the goal is never achieved sooner than t . The goal g^i is assumed to be unobservable to other agents $j \neq i$. If agent i reaches g^i , we may calculate the agent’s cumulative reward as $R^i(s_{1:t}) = \sum_{n=1}^{t-1} R^i(s_n, a_n, s_{n+1})$.

2.1 A Definition of Action Explanation

We formally define here what it means to generate *action explanations* for a given POSG, taking a human-centred perspective [29, 62].

Assume that a user is observing the POSG and asks a *query* $q \in Q \subseteq \mathcal{A}^*$ about a sequence of actions of some agents, following a distribution $p(q|s_{1:t})$, where Q is a finite set of possible queries. For example, the query q may be a sequence of negative acceleration values of a vehicle that suddenly braked. Our goal is to generate an *explanation* $e \in \mathcal{E}$ for the query q , where \mathcal{E} is a finite set of possible explanations with some syntactic (but not semantic) constraints, e.g. natural language sentences or 64x64 px images. These explanations are generated by an *explanation function* $E : Q \times \mathcal{S}^* \mapsto \mathcal{E}$ that maps a query q and a state trajectory $s_{1:t}$ to an explanation $e = E(q, s_{1:t})$.

Our problem is to find an optimal explanation function \hat{E} which maximizes some measure(s) of the expected “quality” of its generated explanations, where “quality” is informally defined to include aspects such as fidelity, human preferences, or actionability, and is formally captured by an external reward model $R : Q \times \mathcal{E} \mapsto \mathbb{R}$ which maps a query q and explanation e to a scalar value $R(e || q)$.

Given the above definitions, the optimal explanation function is obtained by maximizing the expected external reward over trajectories of all lengths for all possible queries:

$$\hat{E} \in \arg \max_{E \in \Delta(E)} \sum_{t=1}^{\infty} \sum_{q \in Q} \int_{\Omega_t} p(q, s_{1:t}) R(E(q, s_{1:t}) || q) ds_{1:t}, \quad (1)$$

where $\Delta(E)$ is a search space over explanation functions, Ω_t is a set of trajectories we want to explain, and $p(q, s_{1:t}) = p(q|s_{1:t})p(s_{1:t})$ is the joint probability of a query and state trajectory.

The definition of a query q allows a user to inquire about the actions of only one agent (by replacing the actions of other agents with the empty set). In our experiments, we focus on such queries about a single agent $\varepsilon \in \mathcal{I}$, because single-agent queries are often how people would investigate complex systems [37, 59, 66]. Importantly, the generated explanations may still involve all agents. Finally, our formulation produces action explanations that are in terms of elements that can be derived from the state trajectory and queries, e.g. rewards or observations. Our formalisation is not immediately suited for giving a mechanistic interpretation of the functional parametrisation of the policies of agents (e.g. saliency maps). For that, other approaches may be more appropriate [61].

3 AGENTIC EXPLANATIONS VIA INTERROGATIVE SIMULATIONS

We present a novel method, *Agentic eXplanations via Interrogative Simulations* (AXIS), to approximate \hat{E} , giving an overview in Figure 1 and the full algorithm in Algorithm 1. Our appendix with full method details and the codebase are in the supplementary materials and will be open-sourced. AXIS is based on the empirically validated *counterfactual effect size model* (CESM) [72], which posits that a cause C is more likely to be selected for an explanation of an

event E if C co-occurs with E under different but *similar* background conditions (i.e. counterfactuals). For example, if a car crashes into a tree, then “the driver being drunk” might be more often selected as an explanation than “the presence of the tree”, because under similar background conditions (e.g. weather, road type), whenever the driver is drunk the car crashes, and conversely, whenever the driver is not drunk, the crash is absent. Meanwhile, if the tree is not present (but the driver is still drunk), a crash would occur.

AXIS converts the CESM into a computational method by using an LLM to propose interventions on the actions of agents and then forward-simulating the counterfactual system to observe whether the queried action still occurs. As such, the main assumption of AXIS is the existence of a simulator \hat{T} which approximates the state transition function T and the joint action policy π . When data is available, a parameterised model of future trajectories could be learned as \hat{T} . When data is not available, a range of modelling and planning algorithms can be deployed to obtain \hat{T} [6].

Options. A question with using LLMs to propose interventions is whether the representations of a POSG are understandable to the LLM, and vice versa; whether the instructions of an LLM mean anything to the simulator of a POSG. For example, actions in AD are sequences of real-valued vectors which are not easily parsed by an LLM for longer sequences [52]. In addition, LLM reasoning over low-level actions may not be effective for long-horizon explanations [52], unless $|\mathcal{A}|$ is finite and small, but this is not generally the case.

To bridge any representational gaps between LLM and simulator, we may use action abstractions based on the popular formalism of *options* [77]. Given a finite set of options \mathcal{M} , an option $m \in \mathcal{M}$ is a (named) triple (S_m, b_m, π_m^i) where $S_m \subseteq \mathcal{S}$ is a set of initial states, $b_m \subseteq \mathcal{S}$ is a set of termination states, and $\pi_m^i : \mathcal{S} \times \mathcal{A}^i \mapsto [0, 1]$ is the policy used by the option. An option groups action sequences that begin when the environment reaches an initial state in S_m , then the agent follows the (stochastic) policy π_m^i until a termination state in b_m is reached. In AXIS, the LLM proposes interventions in the form of options, which are then executed by the simulator. To wrap observed action sequences into options for the LLM, written as $\text{wrap}(a_{1:t}^i)$, we check the initial and termination states of all options against the sequence until a match is found.

Simulator & Interrogation Prompts. Building on options, we facilitate communication from the LLM to the simulator through a set of *interrogation prompts* shown in Table 1. These prompts allow the LLM to propose interventions to agents’ behaviour and the simulator to interpret those proposals. An interrogation prompt p , such as ‘What if vehicle 2 had changed lanes at time 40?’, consists of a keyword followed by arguments in parentheses with arguments referred to as $p.\text{arg}$, e.g. $\text{whatif}(2, [\text{change-lane}], 40)$.

The simulator takes as input the current sequence of observations $o_{1:t}^\varepsilon$ from the perspective of the queried agent ε . The LLM is then requested to propose an interrogation prompt p . In case $p.\text{time} \neq t$, we either roll back time by truncating $o_{1:t}^\varepsilon$ or simulate forward the environment until the starting simulator-time, written as τ , matches $q.\text{time}$. The environment is then forward simulated using \hat{T} , overriding agents’ policy with the specified intervention in prompt p as applicable, until the simulation terminates, returning new joint states $s_{\tau:n}$ and actions $a_{\tau:n}$, where n is a final timestep.

Table 1: Interrogation prompts used to facilitate communication between the LLM and the simulator.

| Prompt | Arguments and Effect | Example |
|--------|--|--|
| add | Args: start $\subset \mathcal{S}$, goal $\subset \mathcal{S}$ Effect: Add agent with initial state start and goal state goal. | add([2, 68], [-3, 14]) “Agent at [2,68] with goal [-3,14].” |
| remove | Args: agent $\in \mathcal{I}$ Effect: Remove specified agent from start of the simulation. | remove(1) “Remove the agent with ID 1.” |
| whatif | Args: agent $\in \mathcal{I}$, option $\in \mathcal{M}^*$, time $\in \mathbb{N}^+$ Effect: Simulates the agent executing option starting at time. | whatif(1 [turn], 40) “What if agent 1 turns at 40?” |
| what | Args: agent $\in \mathcal{I}$, time $\in \mathbb{N}^+$ Effect: Query agent’s state/action at observed/simulated time. | what(1, 40) “What was agent 1 doing at 40?” |

Verbalisation. We call *verbalisation*, and write $\text{verbalise}(\cdot)$, the process of converting into text the various representations of a POSG for the LLM. For AXIS, we may verbalise the observation sequences (V_o), the options (V_m), and the reward signal (V_r), as well as optionally the static environment elements and metadata (V_{env}). We refer to these objects as *context features* (or features), because the verbalised objects are used to provide context to the LLM.

To obtain V_o , we convert $o_{1:\tau}^\epsilon$ to the textual representations of the state and action vectors in the observations as provided by our programming language, noting that for complex state spaces or very long sequences, state filtering [5] or state abstraction [1, 36] may be necessary to produce useful observation verbalisations. In addition, to get V_m , we wrap the joint action sequence $a_{1:\tau}$ to options and use their associated names and start and end time steps. For verbalising rewards V_r , we calculate the cumulative reward $R^\epsilon(s_{1:n})$, and if available, decompose it to reward components $\tilde{R}^i(s_{1:n})$. An example for each verbalisation element is given in Appendix A.

Finally, we define five LLM prompt templates, `system`, `context`, `interrogation`, `explanation`, and `final`, to dynamically create prompts from verbalisation. Each template contains a number of placeholders, e.g. for the verbalised observations or the user prompt, which are replaced with the appropriate strings during runtime. The complete text of AXIS templates is given in Appendix B.

AXIS Algorithm. The pseudocode of AXIS is given in Algorithm 1. AXIS starts when a user prompts it for an action explanation about the actions of an agent ϵ . We write p_u for the user’s natural language question, with t denoting the time of prompting and ϵ denoting the queried agent. AXIS then resets its prompt history H and calls the `VERBALISECONTEXT` function to obtain the verbalised observations and options. A context template is filled using the verbalised information along with p_u , then added to H .

The core of Algorithm 1 is a multi-round interrogation and synthesis process that runs for up to N_{\max} rounds. During each round N , an interrogation template is filled based on p_u . The LLM is then prompted to generate an interrogation prompt p based on H . If p is `DONE`, the loop terminates, which allows the LLM to stop the

Algorithm 1 The explanation function of AXIS. Note, whenever the LLM is prompted, the prompt is appended to a prompt history.

Input: User question p_u , observations $o_{1:t}^\epsilon$, simulator \hat{T} , maximum rounds N_{\max} , optional verbalised static elements V_{env} .
Output: Natural language action explanation e .

- 1: $H \leftarrow []$ empty prompt history.
- 2: $(V_o, V_m, \emptyset) \leftarrow \text{VERBALISECONTEXT}(o_{1:t}^\epsilon, \emptyset)$
- 3: Fill context template from V_{env}, V_o, V_m, p_u and add to H .
- 4: **for** N **in** $[1, \dots, N_{\max}]$ **do**
- 5: Fill interrogation template with p_u .
- 6: $p \leftarrow$ Prompt LLM for interrogation prompt with H .
- 7: **if** $p == \text{DONE}$ **then break** ▶ Early termination
- 8: $s_{\tau:n}, a_{\tau:n}, r^\epsilon \leftarrow$ Run simulation with \hat{T} and p .
- 9: $(V_o, V_m, V_r) \leftarrow \text{VERBALISECONTEXT}([s_{\tau:n}, a_{\tau:n}], r^\epsilon)$
- 10: Fill explanation template with V_o, V_m, V_r, p_u , add to H .
- 11: $e \leftarrow$ Prompt LLM for explanation with H .
- 12: **end for**
- 13: $e \leftarrow$ Prompt LLM with final template and H .
- 14: **return** e
- 15: **function** `VERBALISECONTEXT`($o_{1:t}^\epsilon, r^\epsilon$)
- 16: $a_{1:t} \leftarrow$ Extract joint actions of visible vehicles from $o_{1:t}^\epsilon$.
- 17: $V_o \leftarrow \text{verbalise}(o_{1:t})$
- 18: $V_m \leftarrow \text{verbalise}(\{\text{wrap}(a_{1:t}^i) \mid i \in I\})$
- 19: $V_r \leftarrow \text{verbalise}(r^\epsilon)$ **if** $r^\epsilon \neq \emptyset$ **else** \emptyset
- 20: **return** (V_o, V_m, V_r)
- 21: **end function**

action explanation generation process early. Otherwise, a simulation is run with \hat{T} based on the interrogation query, producing new states $s_{\tau:n}$, actions $a_{\tau:n}$, and rewards $r^\epsilon = R^\epsilon(s_{\tau:n})$.

Following interrogation, the explanation synthesis phase begins. The `VERBALISECONTEXT` function is called again, this time including rewards r^ϵ . An explanation template is filled with the verbalised components and p_u , and added to H . The LLM then generates an explanation e based on the updated prompt history. After completing all rounds, the LLM is prompted with the final prompt template and the full prompt history for a final action explanation e .

4 EVALUATION METHODOLOGY

We evaluate AXIS on the task of motion planning for AD [75], a safety-critical decision-making problem, where explanations may be particularly useful for safety analysis and calibrating trust [48]. Our simulator \hat{T} is based on the method called IGP2[3], which operates in a 2D environment on a static road layout with vehicles following either A* search or Monte Carlo Tree Search algorithms.

The evaluation of explanations is notoriously difficult [61, 63, 64, 74]. The perceived utility of an explanation depends not only on its causal content but also on the context and background of people interpreting the explanation [28, 33, 62]. To mitigate these issues, we use scenario-based evaluation with a novel combination of subjective metrics of user preferences and objective metrics of actionability for goal and next action prediction.

Scenarios & User Prompts. We use ten scenarios with up to five agents (i.e. vehicles) based on previous work [3, 33, 39] where explanations were shown to be useful for users. Example scenarios are




Figure 2: Example scenarios with the queried vehicle shown in blue, other vehicles in orange/white, and occluded vehicles in red. Scenario #2 (left; rational): orange cuts in front of blue and rapidly slows to turn right, so blue changes left to avoid slowing down. Scenario #8 (right; occlusion): blue sees orange on a priority lane coming to a stop, inferring that red must be behind the building on the left, so blue stops for safety.

shown in Figure 2, with full descriptions in Appendix C. Scenarios capture tactical driving situations, e.g. lane changing, unprotected turn, lane merging, and roundabout entry, and involve either safety or efficiency-related reactions from the agents. Scenarios have three categories: five have normal, rational driving behaviour (rational scenarios #1-5), two contain irrational behaviour from a non-ego vehicle (irrational scenarios #6-7), and three contain occluded vehicles (occlusion scenarios #8-10). For each scenario, we define up to three user questions p_u , e.g. ‘Why did vehicle 0 change lanes left?’ and generate explanations for each question separately.

Baselines & Models. We compare against two baselines. **(1) Model only (ModelOnly)**, where we give the full initial context to the LLM and ask it to explain the scenario. **(2) No explanation (NoExp)**, only applicable for actionability metrics, where we do not provide any explanation to the evaluator, to establish a baseline capacity to predict the agent’s goal and next action. We are not aware of relevant previous work that is directly comparable to AXIS.

Five LLMs are used to test the robustness of AXIS to the type and size of LLMs: OpenAI GPT-4.1-2025-04-14 [69], o1-2024-12-17 [68], DeepSeek-V3-0324 [25], DeepSeek-R1 [24], and Meta Llama 3.3-70B-Instruct [54]. We also experimented with LLMs up to 7B parameters, but the results were inaccurate as the models only copied text from the context, without actually proposing meaningful interventions. For sampling from the LLMs, we use default hyperparameters across all models, given in Appendix D, and set a maximum output token limit of 512 for instruct and 16,384 for reasoning models. We use $N_{max} = 10$ maximum rounds in AXIS.

Metrics. Unlike previous work in explainable SDM which usually centres on subjective preference and, infrequently, on downstream actionability [61], we use a combination of five different approaches. This evaluation setup allows us to qualitatively assess the explanations’ correctness and measure their robustness to variation, as well as understand users’ subjective preferences, perceived explanation correctness, and downstream performance (*require user study):

- (1) **Qualitative correctness** is assessed through case studies and by identifying which causes are selected by AXIS.
- (2) **Robustness to variations** is measured by varying the experimental conditions, e.g. hyperparameters and LLM models, to see whether correct explanations are generated even under different conditions.

- (3) **Subjective preference*** is measured along four 5-point Likert scales following Hoffman et al. [41]: perceived completeness, satisfaction, sufficient detail, and trustworthiness of the explanation. We calculate the geometric mean of the four scales and report a single score between 1 and 5.
- (4) **Perceived correctness*** is calculated by comparing our expert causal explanation of the scenario, given in Appendix D, to an explanation by AXIS, reported on a scale between 1, completely incorrect, to 5, completely correct.
- (5) **Goal and next action prediction*** measures the accuracy of people predicting the goal and immediate next action of agents given an explanation based on Gyevar and Towers [35], reported as accuracy in the interval $[0, 1]$.

When presenting aggregate results, we select best-round explanations which have the highest geometric mean of preference and perceived correctness scores across interrogation-synthesis rounds.

LLM-as-a-Judge. User studies are expensive and hard to control. In order to scale our evaluation setup, we use the popular approach of LLM-as-a-judge [32]. This is possible, because AXIS is targeted at end users in an accessible domain such as driving for which ample training data may be available. Previous work also showed that LLM-based substitutes of crowdsourced human data for explanation evaluation are feasible [23], although there may be limitations of this setup, which we will discuss in detail in Sections 5 and 6.

For our LLM-as-a-judge approach, we design prompt templates for each user-study-based metric (metrics 3 to 5 above) and use Anthropic Claude 3.5 [10], an LLM not used for explanation generation, as evaluator. Full prompt templates used for evaluation are in Appendix D. We manipulate three independent within-subjects variables in this experimental setup: **(1) Model:** the LLM model used in AXIS, **(2) Scenario:** the scenario explained by AXIS, and **(3) Query:** the natural language user question as listed in Appendix C.

Shapley Analysis of Context Features. We use Shapley value analysis [76] for context feature selection based on perceived correctness to optimise explanation correctness and to shorten prompt lengths. Due to the factorial number of evaluations required for a complete Shapley analysis, we evaluate on a subset of models and scenarios, while covering a broad range of experimental conditions. We select two LLMs, one smaller (Llama 3.3-70B) and one larger (GPT-4.1), to compare the differences in model size. For the two LLMs, we choose three scenarios (#3, #7, #8), each covering one of the three scenario types (rational, irrational, occlusion).

- **Layout (V_{env}):** The semantic road layout, driving rules, and metadata (e.g. units, speed limit) described in text.
- **States (V_o):** The states contained in the observations, i.e. position and velocity vectors, and scalar heading.
- **Actions (V_a):** The actions contained in the observations, i.e. acceleration and steering vectors.
- **Options (V_m):** The high-level options based on IGP2 [3], i.e. Continue, Change-Left, Right, Exit, and Stop
- **Linguistic Complexity:** Additional feature controlling linguistic complexity of the explanation, i.e. concise or complex.

We select features for our full experiments based on the Shapley analysis, with aggregate results over scenarios shown in Figure 3 and per-scenario results in Appendix E. Reward components were

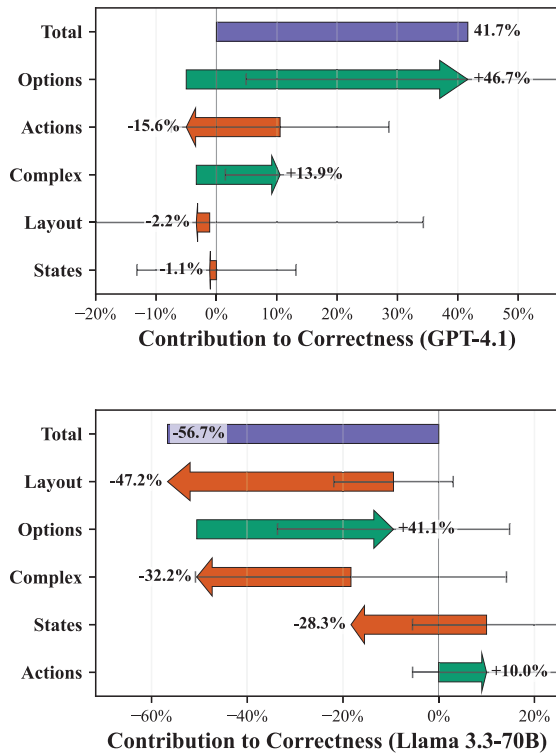


Figure 3: Shapley values calculated from perceived correctness for GPT-4.1 (top) and Llama 3.3-70B (bottom) models. The x-axis shows the Shapley value contribution to the final correctness score. The blue Total bar shows the cumulative Shapley value with all features. Results are aggregated across scenarios #3, #7, #8. Error bars show standard error of mean.

always included in the prompt as they are essential for explanations in terms of the desires and goals. We find that options are consistently helpful for both models, with the larger model using state and layout information more effectively, and the smaller model using action information more effectively. The larger model also benefits from linguistic complexity, and the smaller model from concise explanations. Therefore, we select options, complex linguistic content, and states for larger models, and options, actions, and concise linguistic complexity for the smaller model.

5 RESULTS

We ask the following two research questions to understand whether AXIS is effective at generating useful action explanations:

- (RQ1) Does the interrogation and synthesis process help with selecting the correct causes behind the actions of the queried vehicle in each scenario?
- (RQ2) Are action explanations of AXIS preferred and perceived correct by users, and do they help users predict the agent’s goal and next action?

Table 2: Verbatim excerpts from AXIS for scenario #2, given to the user prompt “Why did vehicle 0 change lanes left?”

GPT-4.1 (ModelOnly): Vehicle 0 changed left because vehicle 1, originally in the adjacent lane, completed a rightward lane change and moved ahead, creating space. With no immediate obstruction and an opportunity to overtake or maintain a higher speed, vehicle 0 initiated the lane change to improve its driving conditions.

GPT-4.1: Vehicle 0 changed lanes left to avoid a *potential conflict with vehicle 1*. [...] This maneuver enabled vehicle 0 to *maintain its speed and avoid being impeded* or forced to slow down. The *actions and timing of vehicle 1*, rather than vehicle 2, were the primary cause for vehicle 0’s lane change. [...]

DeepSeek-V3: Vehicle 0 changed lanes left to avoid a *potential conflict with Vehicle 1*, which was actively merging right into its path. The counterfactual tests showed: 1) When Vehicle 0 stayed in its lane, it had to *slow down significantly* due to Vehicle 1’s merge 2) When *Vehicle 1 didn’t merge*, *Vehicle 0 maintained smooth progress* without lane changes. [...]

Llama 70B: Vehicle 0 changed lanes left *in response to Vehicle 1’s lane change*, likely to *avoid a potential conflict* or make space, indicating a causal relationship between their actions.

o4-mini: Vehicle 0’s left-lane move wasn’t a random drift but a deliberate way to *maintain speed and clear the upcoming junction*. In the original run, the car beside it cut in and was going to slow or turn, so changing lanes let Vehicle 0 *avoid braking*.

We assess RQ1 through our qualitative and robustness analyses in and RQ2 through the subjective and actionability metrics.

5.1 Qualitative Assessment

Our results show that AXIS extracts the most important causes for rational scenarios. For example, looking at explanations of scenario #2 in Table 2, all models cite correct causes in terms of other agents’ actions and the queried agent’s rewards, including effects on speed, strategic maneuvering, and safety, while adding useful contextual information, e.g. highlighting irrelevant vehicles and road layout information. Explanations are often contrastive as well, which are argued to be most intelligible to people [62].

In comparison, the ModelOnly baseline is generic, not contrastive and fails to identify the most important causes, incorrectly attributing the lane change of the queried vehicle to vehicle 1 “completing a rightward lane change and moving ahead, creating space.” This highlights the benefit of AXIS’s interrogation and synthesis process in selecting the most relevant causes for the explanation. Similarly, intelligible and correct explanations are generated by AXIS for scenarios #1 to #5, as shown in Appendix F.

However, AXIS encounters the limitations of LLMs when having to reason about irrational behaviour or when occluded factors are present. In the case of scenarios with irrational agents, the LLMs mostly ignore the irrational behaviour, focusing entirely on the actions of the queried agent. The LLMs produce incomplete explanations because they seem to assume all vehicles to obey the driving rules and act rationally, so they never propose interventions




Figure 4: Proportion of queries across scenarios and prompts used by different models. Most models focus on **remove** and **whatif** queries which are best to extract counterfactuals.

where agents are forced to act irrationally. In the case of occlusion scenarios, the LLMs were explicitly told that occluded vehicles may be present, see Appendix B, yet they either failed to extract occlusion-related causes or assumed that all non-queried vehicles were occluded. These issues are a result of the limitations of LLMs, in spite of which AXIS still manages to improve overall scores, as we show in Section 5.2. It is also worth noting that reasoning about partial observability and irrationality is a very difficult challenge for any decision-making system in general [39, 53, 60].

Nevertheless, Figure 4 shows that LLMs do focus on using the **whatif** and **remove** prompts across all scenarios. These prompts are the most useful for obtaining causal information, further highlighting the benefit of the interrogation phase in AXIS. For example, removing a vehicle can impact the behaviour of surrounding vehicles by allowing them to take simpler actions without the need to adjust to the removed vehicle. When a model failed to use these prompts, for example with DeepSeek-V3, it achieved the lowest perceived correctness (cf. Table 3), suggesting that it failed to understand the appropriate causal relationships.

5.2 Quantitative Analysis

Aggregate results for each model are shown in Table 3 with per-scenario results and best-round explanations shown in Appendix G. AXIS improves the perceived correctness of explanations across all models, by at least 7.67% and as much as 17.83%, with improved accuracy on goal prediction from 23.73% to 46% for four models, and improved or comparable results on action prediction. This means that AXIS-generated explanations better capture the causality in the scenario than baselines and are more actionable for downstream tasks. AXIS also achieves the highest scores for all metrics, although no one model is dominant over others.

In order to understand how evaluation scores depend on the interrogation-synthesis rounds N_{max} , we evaluate intermediate explanations, plotting aggregate results in Figure 5 and model-level results in Appendix H. For perceived correctness and goal/action prediction accuracy, even the intermediate AXIS explanations may improve on the baseline. In addition, except Llama 3.3-70B, models usually output DONE — that is, they deemed the available information exhaustive — before reaching N_{max} rounds, with average maximum $N = 5.03 \pm 0.27$. However, LLMs’ performance varies over N_{max} , exhibiting biases observed in some previous work, e.g. prematurely proposing solutions, relying too heavily on previous

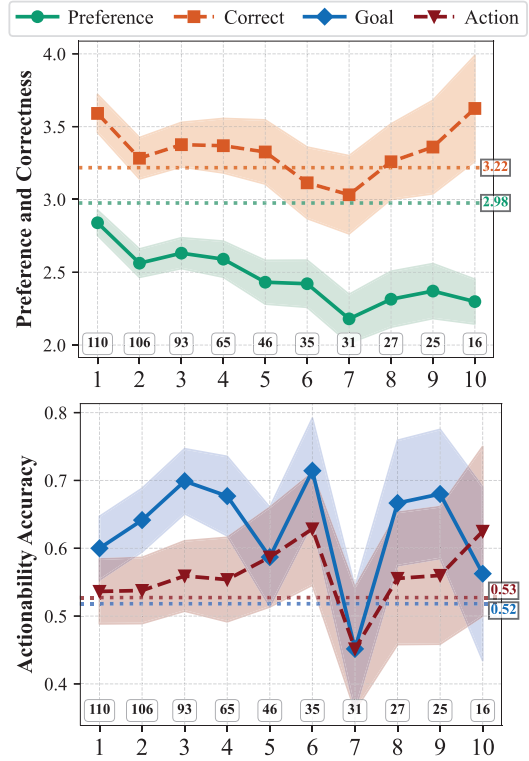


Figure 5: Evolution of preference/correctness (top), and goal/action prediction accuracy (bottom) of intermediate explanations over interrogation-synthesis rounds, aggregated across models, scenarios, and user prompts. The number of explanations for each round is shown above the x-axis in boxes. Horizontal dotted lines show the *ModelOnly* baseline.

explanation attempts, and a recency bias [18, 49]. An example of this bias is the drop in scores at $N = 7$. Here, Llama 3.3-70B interrogates counterfactuals completely unrelated to any previous lines of interrogation, e.g. the removal of a previously untested vehicle, and synthesises an explanation based solely on this latest information, ignoring everything that came before that. This results in a drop in performance, which is mostly recovered in the following round, when the model is re-prompted with the full context.

A similar example of LLM bias can be seen in the lower subjective preference scores for some models using AXIS. Our evaluation is based on previous work designed originally for human participants [41], for example, giving a score to the question “This explanation of why the self-driving car behaved as it did seems complete.” However, when we inspect the reasoning of the evaluator LLM, it often gives low scores due to a lack of reference to raw scenario data, for example exact timing or distance information, something that people may consider superfluous [62]. This also happens across all four axes of the preference metric, and is most pronounced for Llama 3.3-70B, which was instructed to generate concise explanations based on our Shapley analysis.

Finally, some stronger reasoning models (o4-mini, DeepSeek-R1) show significant improvements in perceived correctness, but

Table 3: Comparison of baselines and AXIS on metrics aggregated across scenarios and prompts, with AXIS using best-round explanations. Red/green values show the change of mean relative to the model-only baseline. AXIS improves the correctness of explanations for all models and actionability for most models. It also achieves the highest scores across all metrics.

| Metric | NoExp | GPT-4.1 | +AXIS | DSV3 | +AXIS | Llama 70B | +AXIS | o4-mini | +AXIS | DSR1 | +AXIS |
|-------------|-----------|-----------|--|-----------|--|-----------|---|-----------|--|-----------|---|
| Preference | N/A | 2.86±0.14 | 2.50±0.19 -0.36↓ | 2.69±0.14 | 3.49±0.17 +0.80↑ | 2.81±0.15 | 2.36±0.15 -0.45↓ | 3.19±0.20 | 3.15±0.19 -0.04↓ | 3.33±0.20 | 3.37±0.21 +0.04↑ |
| Correctness | N/A | 3.00±0.27 | 3.23±0.34 +0.23↑ | 2.64±0.22 | 3.18±0.32 +0.55↑ | 3.14±0.34 | 3.59±0.29 +0.45↑ | 3.59±0.33 | 4.23±0.25 +0.64↑ | 3.73±0.31 | 4.00±0.25 +0.27↑ |
| Goal Acc. | 0.56±0.05 | 0.59±0.11 | 0.73±0.10 +0.14↑ | 0.50±0.11 | 0.73±0.10 +0.23↑ | 0.45±0.11 | 0.64±0.10 +0.18↑ | 0.55±0.11 | 0.68±0.10 +0.14↑ | 0.50±0.11 | 0.50±0.11 |
| Action Acc. | 0.53±0.05 | 0.59±0.11 | 0.73±0.10 +0.14↑ | 0.59±0.11 | 0.45±0.11 -0.14↓ | 0.27±0.10 | 0.64±0.10 +0.36↑ | 0.59±0.11 | 0.55±0.11 -0.05↓ | 0.59±0.11 | 0.50±0.11 -0.09↓ |

they have lower goal and action prediction accuracy compared to other models. This is because reasoning models provide much more context and mention more causes in their explanations, which seem to “confuse” the evaluator LLM. Less advanced models were more succinct in their responses, so the evaluator model could extract the correct goal and next action more easily from the explanation. This is similar to how people might prefer explanations that are shorter because they might reduce cognitive overload.

6 RELATED WORK AND DISCUSSION

To answer our research questions, we found that the interrogation-synthesis process of AXIS greatly improves the relevance and correct extraction of causes in rational scenarios, but struggles with irrational behaviour and occluded factors due to the limitations of LLMs (RQ1). We also found that AXIS improves the perceived correctness of explanations for all models and actionability for most models, achieving the highest scores across all metrics, suggesting that AXIS is effective at generating action explanations (RQ2).

6.1 Related Work

Counterfactual explanation generation has received significant attention but almost exclusively in the single-agent setting, often focusing on mechanistic explanations of neural networks given only a single state/action [30]. For example, Olson et al. [67] developed a deep generative model to produce counterfactual state explanations for RL agents, illustrating minimal changes needed for alternative actions, but their method is limited to explaining a single state instead of full trajectories. In contrast, action explanation in SDM takes into account the entire trajectory of an agent’s actions and their outcomes. Madumal et al. [58] proposed an action influence model grounded in structural causal models to provide counterfactual explanations for model-free RL agents, but their method requires knowing *a priori* which features may have a causal effect. AXIS instead uses an LLM to automatically propose interventions which may have a causal effect. Admoni et al. [2] introduce SysLLM, which synthesises global textual summaries of single-agent policies by prompting LLMs to describe behavioural patterns based on a dataset of observed trajectories, showing alignment with expert insights, but their evaluation is limited to simple GridWorld environments. In contrast, we evaluated AXIS using the complex continuous action space of autonomous driving.

In the multi-agent setting, Boggess et al. [13] introduced policy summarisation and query-based language explanations for multi-agent RL (MARL) [4], extending this work to tackle temporal user queries [14]. Heuillet et al. [40] explored the application of Shapley

values [76] to cooperative MARL, although limited to explaining single actions. Gyevnar et al. [36] introduced a method which explains actions in MAS using a probabilistic forward model to identify salient causes, performing well in AD evaluations, although their requirement of a probabilistic model is not always easy to satisfy.

6.2 Limitations and Future Work

Certain limitations must be considered when interpreting our results. **(1) LLM as external reward model:** While previously De Bona et al. [23] have shown that LLM-based substitutes of crowd-sourced human data for explanation evaluation are feasible, their results are preliminary, and the assessment of LLMs, especially on subjective preference, is biased, as our results and previous work have shown [8, 18]. Future work should focus on running a human subjects evaluation for comparison against LLM-based scores. **(2) LLM brittleness in multi-turn conversation:** AXIS can be thought of as a multi-turn conversation between the LLM and the simulator. As a result, AXIS is subject to limitations of the underlying LLM, e.g. verbosity, premature ending of conversation, incorrect assumptions, and recency bias [50]. We also did not tailor individual prompts for specific LLMs besides our Shapley analysis, and we may see higher scores with better-tuned prompts. **(3) Baseline availability and experimental setup:** As with almost all work in explainable SDM, our results are limited by a lack of comparable baselines from the literature, which we tried to solve by comparing against LLM-only baselines. Subjective judgements are highly dependent on fluency [79], which further complicates cross-method comparison. In addition, explanations do not exist without context. In our case, this meant using the domain of autonomous driving in experiments, and future work should test other domains. **(4) Scaling with number of agents:** We tested AXIS with up to five agents in coupled scenarios with complex continuous action spaces. However, it is unclear how the method would scale to discrete domains with simpler action spaces where many more agents may be present, e.g. level-based foraging [19] or warehouse routing [70].

ACKNOWLEDGMENTS

This work was supported in part by the Edinburgh International Data Facility (EIDF) and the Data-Driven Innovation Programme at the University of Edinburgh, the UKRI Centre for Doctoral Training in Natural Language Processing, funded by the UKRI (grant EP/S022481/1) and the University of Edinburgh, School of Informatics and School of Philosophy, Psychology & Language Sciences.

REFERENCES

- [1] David Abel, Dilip Arumugam, Lucas Lehnert, and Michael Littman. 2018. State Abstractions for Lifelong Reinforcement Learning. In *Proceedings of the 35th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 80)*, Jennifer Dy and Andreas Krause (Eds.). PMLR, 10–19.
- [2] Sahar Admoni, Omer Ben-Porat, and Ofra Amir. 2025. SySLLM: Generating Synthesized Policy Summaries for Reinforcement Learning Agents Using Large Language Models. <https://doi.org/10.48550/arXiv.2503.10509> arXiv:2503.10509 [cs]
- [3] Stefano V. Albrecht, Cillian Brewitt, John Wilhelm, Balint Gyevnar, Francisco Eiras, Mihai Dobre, and Subramanian Ramamoorthy. 2021. Interpretable Goal-based Prediction and Planning for Autonomous Driving. In *IEEE International Conference on Robotics and Automation (ICRA)*, 1043–1049. <https://doi.org/10.1109/ICRA48506.2021.9560849>
- [4] Stefano V. Albrecht, Filippos Christianos, and Lukas Schäfer. 2024. *Multi-Agent Reinforcement Learning: Foundations and Modern Approaches*. MIT Press. <https://www.marl-book.com>
- [5] Stefano V Albrecht and Subramanian Ramamoorthy. 2016. Exploiting causality for selective belief filtering in dynamic Bayesian networks. *Journal of Artificial Intelligence Research* 55 (2016), 1135–1178.
- [6] Stefano V. Albrecht and Peter Stone. 2018. Autonomous agents modelling other agents: A comprehensive survey and open problems. *Artificial Intelligence* 258 (2018), 66–95. <https://doi.org/10.1016/j.artint.2018.01.002>
- [7] Ziyang An, Hendrik Baier, Abhishek Dubey, Ayan Mukhopadhyay, and Meiyi Ma. 2024. Enabling MCTS Explainability for Sequential Planning Through Computation Tree Logic. In *Proceedings of the 27th European Conference on Artificial Intelligence (ECAI)*.
- [8] Jacy Reese Anthis, Ryan Liu, Sean M. Richardson, Austin C. Kozlowski, Bernard Koch, James Evans, Erik Brynjolfsson, and Michael Bernstein. 2025. LLM Social Simulations Are a Promising Research Method. arXiv 2504.02234.
- [9] Anthropic. 2025. Agentic Misalignment. <https://www.anthropic.com/research/agentic-misalignment>
- [10] Anthropic. 2025. Claude 3.5 Haiku. <https://www.anthropic.com/claude> Large language model.
- [11] Ahsaas Bajaj, Pavitra Dangati, Kalpesh Krishna, Pradhiksha Ashok Kumar, Rhea Uppaal, Bradford Windsor, Eliot Brenner, Dominic Dotterrer, Rajarshi Das, and Andrew McCallum. 2021. Long Document Summarization in a Low Resource Setting Using Pretrained Language Models. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing: Student Research Workshop*, Jad Kabbara, Haitao Lin, Amandalynne Paullada, and Jananis Vamvas (Eds.). Association for Computational Linguistics, Online, 71–80. <https://doi.org/10.18653/v1/2021.acl-srw.7>
- [12] Tom Bewley and Jonathan Lawry. 2021. TripleTree: A Versatile Interpretable Representation of Black Box Agents and Their Environments. *Proceedings of the AAAI Conference on Artificial Intelligence* 35, 13 (May 2021), 11415–11422. <https://doi.org/10.1609/aaai.v35i13.17360>
- [13] Kayla Boggess, Sarit Kraus, and Lu Feng. 2022. Toward Policy Explanations for Multi-Agent Reinforcement Learning. In *Thirty-First International Joint Conference on Artificial Intelligence*, Vol. 1. 109–115. <https://doi.org/10.24963/ijcai.2022/16>
- [14] Kayla Boggess, Sarit Kraus, and Lu Feng. 2023. Explainable Multi-Agent Reinforcement Learning for Temporal Queries. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence (IJCAI '23)*, Macao, P.R.China, 55–63. <https://doi.org/10.24963/ijcai.2023/7>
- [15] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language Models Are Few-Shot Learners. In *Advances in Neural Information Processing Systems*, Vol. 33. Curran Associates, Inc., 1877–1901.
- [16] Ruth M.J. Byrne. 2023. Good explanations in explainable artificial intelligence (XAI): evidence from human explanatory reasoning. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence (IJCAI '23)*, International Joint Conferences on Artificial Intelligence (IJCAI), Macao, P. R. China, Article 733, 9 pages. <https://doi.org/10.24963/ijcai.2023/733>
- [17] Ruth M. J. Byrne. 2005. *The rational imagination: how people create alternatives to reality*. The MIT Press, Cambridge, MA. <https://doi.org/10.7551/mitpress/5756.001.0001>
- [18] Yupeng Chang, Xu Wang, Jindong Wang, Yuan Wu, Linyi Yang, Kaijie Zhu, Hao Chen, Xiaoyuan Yi, Cunxian Wang, Yidong Wang, et al. 2024. A survey on evaluation of large language models. *ACM transactions on intelligent systems and technology* 15, 3 (2024), 1–45.
- [19] Filippos Christianos, Lukas Schäfer, and Stefano V Albrecht. 2020. Shared Experience Actor-Critic for Multi-Agent Reinforcement Learning. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- [20] Council of Europe. 2024. The Framework Convention on Artificial Intelligence. <https://www.coe.int/en/web/artificial-intelligence/the-framework-convention-on-artificial-intelligence> Accessed: 2025-04-20.
- [21] Richard Dazeley, Peter Vamplew, and Francisco Cruz. 2023. Explainable Reinforcement Learning for Broad-XAI: A Conceptual Framework and Survey. *Neural Computing and Applications* 35, 23 (Aug. 2023), 16893–16916. <https://doi.org/10.1007/s00521-023-08423-1>
- [22] Richard Dazeley, Peter Vamplew, Cameron Foale, Charlotte Young, Sunil Aryal, and Francisco Cruz. 2021. Levels of Explainable Artificial Intelligence for Human-Aligned Conversational Explanations. *Artificial Intelligence* 299 (2021), 103525. <https://doi.org/10.1016/j.artint.2021.103525>
- [23] Francesco Bombassei De Bona, Gabriele Dominic, Tim Miller, Marc Langheinrich, and Martin Gjoreski. 2024. Evaluating Explanations Through LLMs: Beyond Traditional User Studies. arXiv:2410.17781 [cs.AI] <https://arxiv.org/abs/2410.17781>
- [24] DeepSeek. 2025. DeepSeek-R1. <https://deepseekv3.org/deepseek-r1>
- [25] DeepSeek. 2025. DeepSeek-V3 Technical Report. <https://deepseekv3.org/paper>
- [26] Stephanie Droop and Neil R Bramley. 2023. Extending counterfactual reasoning models to capture unconstrained social explanations.
- [27] Mary T. Dzindolet, Scott A. Peterson, Regina A. Pomranky, Linda G. Pierce, and Hall P. Beck. 2003. The Role of Trust in Automation Reliance. *International Journal of Human-Computer Studies* 58, 6 (June 2003), 697–718. [https://doi.org/10.1016/S1071-5819\(03\)00038-7](https://doi.org/10.1016/S1071-5819(03)00038-7)
- [28] Upol Ehsan, Samir Passi, Q. Vera Liao, Larry Chan, I-Hsiang Lee, Michael Muller, and Mark O Riedl. 2024. The Who in XAI: How AI Background Shapes Perceptions of AI Explanations. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24)*, Association for Computing Machinery, New York, NY, USA, 1–32. <https://doi.org/10.1145/3613904.3642474>
- [29] Upol Ehsan and Mark O. Riedl. 2020. Human-Centered Explainable AI: Towards a Reflective Sociotechnical Approach. In *HCI International 2020 - Late Breaking Papers: Multimodality and Intelligence (Lecture Notes in Computer Science)*, Constantine Stephanidis, Masaaki Kurosu, Helmut Degen, and Lauren Reinerman-Jones (Eds.). Springer International Publishing, Cham, 449–466. https://doi.org/10.1007/978-3-030-60117-1_33
- [30] Jasmina Gajcin and Ivana Dusparic. 2024. Redefining Counterfactual Explanations for Reinforcement Learning: Overview, Challenges and Opportunities. *ACM Comput. Surv.* 56, 9 (April 2024), 219:1–219:33. <https://doi.org/10.1145/3648472>
- [31] Jasmina Gajcin, Jovan Jeromela, and Ivana Dusparic. 2024. Semifactual Explanations for Reinforcement Learning. In *Proceedings of the 12th International Conference on Human-Agent Interaction (HAI '24)*, Association for Computing Machinery, New York, NY, USA, 167–175. <https://doi.org/10.1145/3687272.3688324>
- [32] Jiawei Gu, Xuhui Jiang, Zhichao Shi, Hexiang Tan, Xuehao Zhai, Chengjin Xu, Wei Li, Yinghan Shen, Shengjie Ma, Honghao Liu, Saizhuo Wang, Kun Zhang, Yuanzhuo Wang, Wen Gao, Lionel Ni, and Jian Guo. 2025. A Survey on LLM-as-a-Judge. <https://doi.org/10.48550/arXiv.2411.15594>
- [33] Balint Gyevnar, Stephanie Droop, Tadeq Quillien, Shay B. Cohen, Neil R. Bramley, Christopher G. Lucas, and Stefano V. Albrecht. 2025. People Attribute Purpose to Autonomous Vehicles When Explaining Their Behavior: Insights from Cognitive Science for Explainable AI. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (Yokohama, Japan)*, Association for Computing Machinery, New York, NY, United States. <https://doi.org/10.1145/3706598.3713509>
- [34] Balint Gyevnar, Nick Ferguson, and Burkhard Schafer. 2023. Bridging the Transparency Gap: What Can Explainable AI Learn From the AI Act?. In *26th European Conference on Artificial Intelligence*. IOS Press, 964–971. <https://doi.org/10.3233/FAIA230367>
- [35] Balint Gyevnar and Mark Towers. 2025. Objective Metrics for Human-Subjects Evaluation in Explainable Reinforcement Learning. In *2025 Multi-Disciplinary Conference on Reinforcement Learning and Decision Making (Dublin, Ireland)*.
- [36] Balint Gyevnar, Cheng Wang, Christopher G. Lucas, Shay B. Cohen, and Stefano V. Albrecht. 2024. Causal Explanations for Sequential Decision-Making in Multi-Agent Systems. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems (Auckland, New Zealand) (AAMAS '24)*, International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 771–779.
- [37] Joseph Y Halpern. 2016. *Actual Causality*. MIT Press, Cambridge, MA.
- [38] Lewis Hammond, Alan Chan, Jesse Clifton, Jason Hoelscher-Obermaier, Akbir Khan, Euan McLean, Chandler Smith, Wolfram Barfuss, Jakob Foerster, Tomáš Gavenciak, The Anh Han, Edward Hughes, Vojtěch Kovařík, Jan Kulveit, Joel Z. Leibo, Caspar Oesterheld, Christian Schroeder de Witt, Nisarg Shah, Michael Wellman, Paolo Bova, Theodor Cimpanu, Carson Ezell, Quentin Feuillade-Montixi, Matija Franklin, Esben Kran, Igor Krawczuk, Max Lamparth, Niklas Lauffer, Alexander Meinke, Sumeet Motwani, Anka Reuel, Vincent Conitzer, Michael Dennis, Iason Gabriel, Adam Gleave, Gillian Hadfield, Nika Haghtalab, Atoosa Kasirzadeh, Sébastien Krier, Kate Larson, Joel Lehman, David C. Parkes, Georgios Piliouras, and Iyad Rahwan. 2025. *Multi-Agent Risks from Advanced AI*. Technical Report 1. Cooperative AI Foundation. <https://doi.org/10.48550/ARXIV.2502.14143> arXiv:2502.14143

- [39] Josiah P. Hanna, Arrasy Rahman, Elliot Fosong, Francisco Eiras, Mihai Dobre, John Redford, Subramanian Ramamoorthy, and Stefano V. Albrecht. 2021. Interpretable Goal Recognition in the Presence of Occluded Factors for Autonomous Vehicles. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. <https://doi.org/10.1109/IROS51168.2021.9635903>
- [40] Alexandre Heuillet, Fabien Couthouis, and Natalia Díaz-Rodríguez. 2022. Collective eXplainable AI: Explaining Cooperative Strategies and Agent Contribution in Multiagent Reinforcement Learning With Shapley Values. *Comp. Intell. Mag.* 17, 1 (Feb. 2022), 59–71. <https://doi.org/10.1109/MCI.2021.3129959>
- [41] Robert R. Hoffman, Shane T. Mueller, Gary Klein, and Jordan Litman. 2023. Measures for Explainable AI: Explanation Goodness, User Satisfaction, Mental Models, Curiosity, Trust, and Human-AI Performance. *Frontiers in Computer Science* 5 (2023).
- [42] Tobias Huber, Katharina Weitz, Elisabeth André, and Ofra Amir. 2021. Local and Global Explanations of Agent Behavior: Integrating Strategy Summaries with Saliency Maps. *Artificial Intelligence* 301 (Dec. 2021), 103571. <https://doi.org/10.1016/j.artint.2021.103571>
- [43] IASEAI. 2025. IASEAI Call to Action for Lawmakers, Academics, and the Public Ahead of AI Summit in Paris. <https://www.iaseai.org/conference/statement> Accessed: 2025-04-20.
- [44] Nidhi Kalra and Susan M. Paddock. 2016. Driving to Safety: How Many Miles of Driving Would It Take to Demonstrate Autonomous Vehicle Reliability? *Transportation Research Part A: Policy and Practice* 94 (2016), 182–193. <https://doi.org/10.1016/j.tra.2016.09.010>
- [45] Margot E Kaminski. 2021. The right to explanation, explained. *Berkeley Technology Law Journal* 34 (2021), 189–218. <https://doi.org/10.15779/Z38TD9N83F>
- [46] Sunnie S. Y. Kim, Elizabeth Anne Watkins, Olga Russakovsky, Ruth Fong, and Andrés Monroy-Hernández. 2023. "Help Me Help the AI": Understanding How Explainability Can Support Human-AI Interaction. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. Association for Computing Machinery, New York, NY, USA, 1–17. <https://doi.org/10.1145/3544548.3581001>
- [47] Andrei Kirilenko, Albert S. Kyle, Mehrdad Samadi, and Tugkan Tuzun. 2017. The Flash Crash: High-Frequency Trading in an Electronic Market. *Journal of Finance* 72, 3 (2017), 967–998. <https://doi.org/10.1111/jofi.12498>
- [48] Anton Kuznietsov, Balint Gyevar, Cheng Wang, Steven Peters, and Stefano V. Albrecht. 2024. Explainable AI for Safe and Trustworthy Autonomous Driving: A Systematic Review. *IEEE Transactions on Intelligent Transportation Systems* 25, 12 (Dec. 2024), 19342–19364. <https://doi.org/10.1109/ITITS.2024.3474469>
- [49] Philippe Laban, Hiroaki Hayashi, Yingbo Zhou, and Jennifer Neville. 2025. LLMs Get Lost In Multi-Turn Conversation. arXiv:arXiv 2505.06120 <https://arxiv.org/abs/2505.06120>
- [50] Philippe Laban, Hiroaki Hayashi, Yingbo Zhou, and Jennifer Neville. 2025. LLMs Get Lost In Multi-Turn Conversation. <https://doi.org/10.48550/arXiv.2505.06120> arXiv:2505.06120 [cs].
- [51] Bruno Lepri, Nuria Oliver, Emmanuel Letouze, Alex Pentland, and Patrick Vinck. 2018. Fair, transparent, and accountable algorithmic decision-making processes. *Philosophy & Technology* 31, 4 (2018), 611–627. <https://doi.org/10.1007/s13347-017-0279-x>
- [52] Nelson F. Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2024. Lost in the Middle: How Language Models Use Long Contexts. *Transactions of the Association for Computational Linguistics* 12 (2024), 157–173. https://doi.org/10.1162/tacl_a_00638
- [53] Qinghua Liu, Alan Chung, Csaba Szepesvári, and Chi Jin. 2022. When is partially observable reinforcement learning not scary?. In *Conference on Learning Theory*. PMLR, 5175–5220.
- [54] Llama Team, AI at Meta. 2024. The Llama 3 Herd of Models. arXiv 2407.21783 [cs].
- [55] Helena Löfström, Karl Hammar, and Ulf Johansson. 2022. A Meta Survey of Quality Evaluation Criteria in Explanation Methods. In *Intelligent Information Systems*, Jochen De Weerd and Artem Polyvyanyy (Eds.). Springer International Publishing, 55–63.
- [56] Tania Lombrozo. 2006. The structure and function of explanations. *Trends in cognitive sciences* 10, 10 (2006), 464–470. <https://doi.org/10.1016/j.tics.2006.08.004>
- [57] Christopher G Lucas and Charles Kemp. 2015. An improved probabilistic account of counterfactual reasoning. *Psychological review* 122, 4 (2015), 700.
- [58] Prashan Madumal, Tim Miller, Liz Sonenberg, and Frank Vetere. 2020. Explainable Reinforcement Learning through a Causal Lens. *Proceedings of the AAAI Conference on Artificial Intelligence* 34, 03 (April 2020), 2493–2500. <https://doi.org/10.1609/aaai.v34i03.5631>
- [59] Bertram F. Malle and Joshua Knobe. 1997. Which Behaviors Do People Explain? A Basic Actor–Observer Asymmetry. *Journal of Personality and Social Psychology* 72, 2 (1997), 288–304. <https://doi.org/10.1037/0022-3514.72.2.288>
- [60] Peta Masters and Sebastian Sardina. 2019. Goal recognition for rational and irrational agents. In *Proceedings of the 18th international conference on autonomous agents and multiagent systems*. 440–448.
- [61] Stephanie Milani, Nicholay Topin, Manuela Veloso, and Fei Fang. 2024. Explainable Reinforcement Learning: A Survey and Comparative Review. *Comput. Surveys* 56, 7 (April 2024), 168:1–168:36. <https://doi.org/10.1145/3616864>
- [62] Tim Miller. 2019. Explanation in Artificial Intelligence: Insights from the Social Sciences. *Artificial Intelligence* 267 (Feb. 2019), 1–38. <https://doi.org/10.1016/j.artint.2018.07.007>
- [63] Tim Miller. 2023. Explainable AI Is Dead, Long Live Explainable AI! Hypothesis-driven Decision Support Using Evaluative AI. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency (FAccT '23)*. Association for Computing Machinery, New York, NY, USA, 333–342. <https://doi.org/10.1145/3593013.3594001>
- [64] Sina Mohseni, Niloofar Zarei, and Eric D. Ragan. 2021. A Multidisciplinary Survey and Framework for Design and Evaluation of Explainable AI Systems. *ACM Transactions on Interactive Intelligent Systems* 11, 3-4 (Aug. 2021), 24:1–24:45. <https://doi.org/10.1145/3387166>
- [65] Mahsan Nourani, Samia Kabir, Sina Mohseni, and Eric D. Ragan. 2019. The Effects of Meaningful and Meaningless Explanations on Trust and Perceived System Accuracy in Intelligent Systems. *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing* 7 (Oct. 2019), 97–105.
- [66] Matthew J. O’Laughlin and Bertram F. Malle. 2002. How People Explain Actions Performed by Groups and Individuals. *Journal of Personality and Social Psychology* 82, 1 (2002), 33–48. <https://doi.org/10.1037/0022-3514.82.1.33>
- [67] Matthew L. Olson, Roli Khanna, Lawrence Neal, Fuxin Li, and Weng-Keen Wong. 2021. Counterfactual State Explanations for Reinforcement Learning Agents via Generative Deep Learning. *Artificial Intelligence* 295 (June 2021), 103455. <https://doi.org/10.1016/j.artint.2021.103455>
- [68] OpenAI. 2024. OpenAI o1 System Card. <https://openai.com/index/openai-o1-system-card/>
- [69] OpenAI. 2025. Introducing GPT-4.1 in the API. <https://openai.com/index/gpt-4-1/>
- [70] Georgios Papoudakis, Filippos Christianos, Lukas Schäfer, and Stefano V. Albrecht. 2021. Benchmarking Multi-Agent Deep Reinforcement Learning Algorithms in Cooperative Tasks. In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks (NeurIPS)*. <http://arxiv.org/abs/2006.07869>
- [71] Erika Puiutta and Eric MSP Veith. 2020. Explainable reinforcement learning: A survey. In *International cross-domain conference for machine learning and knowledge extraction*. Springer, 77–95.
- [72] Tadeq Quillien and Christopher G. Lucas. 2023. Counterfactuals and the Logic of Causal Selection. *Psychological Review* 131, 5 (2023), 1208–1234. <https://doi.org/10.31234/osf.io/ts76y>
- [73] Anand S Rao, Michael P Georgeff, et al. 1995. BDI agents: From theory to practice.. In *JCMAS*, Vol. 95, 312–319.
- [74] Cynthia Rudin. 2019. Stop Explaining Black Box Machine Learning Models for High Stakes Decisions and Use Interpretable Models Instead. *Nature Machine Intelligence* 1, 5 (May 2019), 206–215. <https://doi.org/10.1038/s42256-019-0048-x>
- [75] Wilko Schwarting, Javier Alonso-Mora, and Daniela Rus. 2018. Planning and Decision-Making for Autonomous Vehicles. *Annual Review of Control, Robotics, and Autonomous Systems* 1, 1 (2018), 187–210. <https://doi.org/10.1146/annurev-control-060117-105157>
- [76] L. Shapley. 1953. A Value for N-Person Games. In *Classics in Game Theory*. Princeton University Press, Chapter Classics in Game Theory, 69–79.
- [77] Richard S. Sutton, Doina Precup, and Satinder Singh. 1999. Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning. *Artificial Intelligence* 112, 1 (1999), 181–211. [https://doi.org/10.1016/S0004-3702\(99\)00052-1](https://doi.org/10.1016/S0004-3702(99)00052-1)
- [78] Mark Towers, Yali Du, Christopher Freeman, and Timothy J. Norman. 2024. Explaining an Agent’s Future Beliefs Through Temporally Decomposing Future Reward Estimators. In *ECAI 2024*. IOS Press, 2790–2797. <https://doi.org/10.3233/FAIA240814>
- [79] Christian Unkelbach and Rainer Greifeneder. 2013. A general model of fluency effects in judgment and decision making. In *The experience of thinking*. Psychology Press, 11–32.
- [80] George A. Vouros. 2022. Explainable Deep Reinforcement Learning: State of the Art and Challenges. *ACM Comput. Surv.* 55, 5, Article 92 (2022), 39 pages. <https://doi.org/10.1145/3527448>
- [81] Sandra Wachter, Brent Mittelstadt, and Chris Russell. 2017. Counterfactual explanations without opening the black box: Automated decisions and the GDPR. *Harvard Journal of Law & Technology* 31, 2 (2017), 841–887. <https://doi.org/10.2139/ssrn.3063289>
- [82] Qingyue Wang, Yanhe Fu, Yanan Cao, Shuai Wang, Zhiliang Tian, and Liang Ding. 2025. Recursively Summarizing Enables Long-Term Dialogue Memory in Large Language Models. *Neurocomputing* 639 (July 2025), 130193. <https://doi.org/10.1016/j.neucom.2025.130193>
- [83] Michael Winikoff. 2017. Debugging Agent Programs with Why? Questions. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems (AAMAS '17)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 251–259.
- [84] Michael Winikoff, Galina Sidorenko, Virginia Dignum, and Frank Dignum. 2021. Why bad coffee? Explaining BDI agent behaviour with valuations. *Artificial Intelligence* 300 (2021), 103554.

- [85] Ming Yin, Jennifer Wortman Vaughan, and Hanna Wallach. 2019. Understanding the Effect of Accuracy on Trust in Machine Learning Models. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/3290605.3300509>
- [86] Tianyi Zhang, Faisal Ladhak, Esin Durmus, Percy Liang, Kathleen McKeown, and Tatsunori B. Hashimoto. 2024. Benchmarking Large Language Models for News Summarization. *Transactions of the Association for Computational Linguistics* 12 (2024), 39–57. https://doi.org/10.1162/tacl_a_00632