

Responsibility-aware Strategic Reasoning in Probabilistic Multi-Agent Systems

Chunyan Mu¹, Muhammad Najib², Nir Oren¹

¹Department of Computing Science, University of Aberdeen

²Department of Computer Science, Heriot-Watt University
Chunyan.Mu@abdn.ac.uk, M.Najib@hw.ac.uk, N.Oren@abdn.ac.uk

Abstract

Responsibility plays a key role in the development and deployment of trustworthy autonomous systems. In this paper, we focus on the problem of strategic reasoning in probabilistic multi-agent systems with responsibility-aware agents. We introduce PATL+R, a variant of Probabilistic Alternating-time Temporal Logic. PATL+R's novelty lies in its incorporation of modalities for causal responsibility, providing a framework for responsibility-aware multi-agent strategic reasoning. We present an approach to synthesise joint strategies that satisfy an outcome specified in PATL+R while optimising the share of expected causal responsibility and reward. This provides a notion of balanced distribution of responsibility and reward gain among agents. To this end, we utilise the Nash equilibrium as the solution concept for our strategic reasoning problem and demonstrate how to compute responsibility-aware Nash equilibrium strategies via a reduction to parametric model checking of concurrent stochastic multi-player games.

Introduction

Strategic decision-making and reasoning in multi-agent systems (MASs) operating in a dynamic and uncertain environment is a challenging research problem. Formalisms such as ATL (Alur, Henzinger, and Kupferman 2002) and PATL (Chen and Lu 2007) provide important frameworks for reasoning about strategic decision making in MASs. These formalisms have also found notable applications in the field of multi-agent planning (Van Der Hoek and Wooldridge 2002; Jamroga 2004; Schnoor 2010; Demri and Fervari 2023; Chrupa and Karpas 2024). Recently there has been a surge in research focused on the trustworthiness of autonomous systems, highlighting the importance of responsibility as a crucial element of AI systems (Kaur et al. 2022; Dignum 2020). Consequently, the notion of a *responsibility-aware agent* has emerged (Yazdanpanah et al. 2023; Kobayashi, Bondu, and Ishikawa 2023; Cosner et al. 2023). Such an agent must consider not only its own reward but also its responsibility with respect to the tasks assigned to it.

This paper explores the problem of strategic reasoning where agents are aware of *causal responsibility* (Chockler

and Halpern 2004; Parker, Grandi, and Lorini 2023), which captures the impact of an agent's or coalition's actions, or inaction, on outcomes. Analysing causal responsibility enables a balanced distribution of rewards or penalties among agents based on their contribution to outcomes¹. This provides a way for ensuring fairness among agents. For example, consider a scenario where two agents are involved in carrying out a *plan* (i.e., a joint strategy) assigned by a principal. If both agents share the same degree of responsibility for the execution of this plan, they should receive a fair share of reward (e.g., each receives the same amount of money). Moreover, agents may pursue different plans based on their preferences, meaning that individual agents may deviate from an overarching plan. Inspired by the concept of *equilibria for multi-agent planning* (Bowling, Jensen, and Veloso 2003), we utilise the Nash equilibrium (NE) as the solution concept to compute stable plans. If a plan is a NE, then no agent can benefit from unilaterally deviating from the plan.

Contributions. This paper focuses on the strategic reasoning problem with *step-bounded temporal properties*, where agents have access to *randomised memoryless* strategies and seek to optimise the balance between responsibility and reward (i.e., agents are responsibility-aware). We present a parametric model that captures both the uncertainty and strategic interactions in multi-agent systems of such a setting. We then introduce the PATL+R logic as a language for specifying properties that encompass responsibility and temporal objectives. By utilising the approach from the parametric model checking paradigm, we demonstrate how to model check PATL+R formulae against our proposed model. We show that such a model checking problem can be solved in PSPACE. We further demonstrate how we can utilise the parametric model to compute NE strategies/plans with respect to agents' *utility functions* that consider both reward and responsibility. These NE strategies/plans are in a sense optimal for each agent, considering trade-offs between responsibility and reward. We show that such a computation can also be done in PSPACE. Omitted proofs can be found in supplementary material.

¹In particular, when combined with *norms* (Henne et al. 2021), although we emphasise that the study of normative systems is beyond the scope of this paper.

Related work. Previous research has explored the incorporation of responsibility into automated planning. The work of Alechina et al. (Alechina, Halpern, and Logan 2017) formalises the attribution of responsibility (or blame) for joint plan failures. Our work diverges from this by focusing on how the degree of responsibility can actively influence plan generation rather than post-hoc analysis of a pre-determined plan. Parker et al.’s work (Parker, Grandi, and Lorini 2023) is more closely related to ours, as it considers causal responsibility for plan selection. Specifically, they consider both *Causal Active Responsibility* and *Causal Passive Responsibility*. Our concept of responsibility is based on theirs. The primary difference between their work and ours lies in the model used. They consider an imperfect information setting (we use a perfect information setting), and they only consider a deterministic setting, while ours is probabilistic. The work of Baier et al. (Baier, Funke, and Majumdar 2021) bears similarity to ours in the sense that they consider games and strategies, unlike the more traditional model in classical planning. However, unlike our work, none of these works consider game-theoretic (specifically, NE) solution concepts. In this regard, our work is more related to the paradigm of *rational verification and synthesis* (Abate et al. 2021; Fisman, Kupferman, and Lustig 2010), where a given (temporal logic) property is checked within the set of system equilibria, or equilibrium strategies (plans) are synthesised. There are works in this research strand in the domain of probabilistic systems, e.g., (Kwiatkowska et al. 2021; Gutierrez et al. 2021; Mittelman et al. 2023; Hyland et al. 2024). However, to the best of our knowledge, no work in this research line incorporates causal responsibility.

From the perspective of logical specification language, we mention the work that uses logical framework to reason about responsibility, e.g., (Yazdanpanah et al. 2019; Naumov and Tao 2021, 2019; Ciuni and Mastop 2009). All of these works differ from ours in the model (deterministic vs probabilistic, one-shot vs temporally extended), the concept of responsibility, or the research aims (c.f., attribution vs planning).

Parametric Model of Stochastic Game

In this section, we introduce the model utilised in this paper. We begin with the conventional model of a *concurrent stochastic multi-player game* (CSG) (Shapley 1953; Kwiatkowska et al. 2021). Then, we define the corresponding *parametric stochastic MAS*, which captures the dynamics of a given CSG.

Definition 1. A *concurrent stochastic multi-player game* (CSG) is a tuple $\mathcal{G} = (\text{Ag}, S, s^0, (\text{Act}_i)_{i \in \text{Ag}}, \delta, \text{Ap}, L)$ where:

- $\text{Ag} = \{1, \dots, n\}$ is a finite set of agents;
- S is a finite non-empty set of states;
- $s^0 \in S$ is the initial state;
- Act_i is a finite set of actions for i . With each agent i and state $s \in S$, we associate a non-empty set $\text{Act}_i(s)$ of available actions that i can perform in s . Write $\text{Act}^{\text{Ag}} = \text{Act}_1 \times \dots \times \text{Act}_n$.

- $\delta : S \times \text{Act}^{\text{Ag}} \rightarrow \text{Dist}(S)$ is a probabilistic transition function;
- Ap is a finite set of atomic propositions;
- $L : S \rightarrow 2^{\text{Ap}}$ is the state labelling function mapping each state to a set of atomic proposition drawn from Ap .

Following (Kwiatkowska et al. 2021), we augment CSGs with *reward structures* of the form $r = (r_s, r_a)$, where $r_s : S \rightarrow \mathbb{R}$ is the state reward and $r_a : \text{Act}^{\text{Ag}} \rightarrow \mathbb{R}$ is the action reward, and consider *cumulative rewards*, that is the sum of payoffs accumulated during the run until a specific point.

In this work, we assume that players have memoryless strategies. Informally, a memoryless strategy for player i prescribes, from each state s , the probability of each action $a \in \text{Act}_i(s)$ being chosen. The set of all memoryless strategies for player i from state s can be encoded by a set of variables $V_s^i = \{x_a : a \in \text{Act}_i(s)\}$. Intuitively, the value of $x_a \in V_s^i$ corresponds to the probability of action a being chosen by player i in state s . Let $V^i = \bigcup_{s \in S} V_s^i$. A memoryless (mixed) strategy for i in \mathcal{G} thus corresponds to an *evaluation* of such a set of variables represented by a function $C^i : V^i \rightarrow \mathbb{R}$.


Definition 2. Given a CSG $\mathcal{G} = (\text{Ag}, S, s^0, (\text{Act}_i)_{i \in \text{Ag}}, \delta, \text{Ap}, L)$, we construct the corresponding parametric stochastic multi-agent system (PSMAS) as a tuple $\mathcal{M} = (\text{Ag}, S, s^0, V, \Delta, \text{Ap}, L)$, where:

- $\text{Ag} = \{1, \dots, n\}$ is a finite set of agents;
- S is a finite non-empty set of states;
- $s^0 \in S$ is the initial state;
- $\text{Act} = \{a_1, a_2, \dots, a_m\}$ is a non-empty finite set of actions;
- $V^i = \{x_{i,1}, \dots, x_{i,m}\} \subseteq V$ is a finite set of variables (parameters) over \mathbb{R}^m for each agent i ;
- $\Delta : S \times \text{Act}^{\text{Ag}} \times S \rightarrow \mathcal{F}_v$ is the probabilistic transition function, where \mathcal{F}_v is the set of polynomials over V with rational coefficients which can be viewed as a parametric transition probability matrix that respects the distribution in δ ;
- Ap is a finite set of atomic propositions;
- $L : S \rightarrow 2^{\text{Ap}}$ is the state labelling function mapping each state to a set of atomic proposition drawn from Ap .

We introduce the notion of *admissible* evaluations in a given PSMAS. An evaluation C^i is admissible if

1. $\Delta_{C^i}(s, \alpha, s) \in [0, 1]$ for all $s, s' \in S$ and $\alpha \in \text{Act}^{\text{Ag}}$, i.e., each action profile has a probability between 0 and 1;
2. $C^i(x) \in [0, 1]$ for all $i \in \text{Ag}$ and $x \in V^i$, i.e., for each player, each action probability is between 0 and 1;
3. $\sum_{x \in V_s^i} C^i(x) = 1$ for all $i \in \text{Ag}$ and $s \in S$, i.e., For each player and each state, the total of action probability values should equal 1.

Henceforth, we assume that the model takes the form of a PSMAS. When evaluating solutions, we further assume that such solutions are admissible.



$$\begin{aligned}
\text{Ag} &= \{A_1, A_2\} \\
\text{Act} &= \{\text{Catch}, \text{Skip}\} \\
\text{Ap} &= \{\text{caught}_1, \text{caught}_2\} \\
S &= \{s_0, s_1, s_2, s_3\} \\
V^1 &= \{x_{1,1}, x_{1,2}\} \triangleq \{x_1, 1 - x_1\} \\
V^2 &= \{x_{2,1}, x_{2,2}\} \triangleq \{x_2, 1 - x_2\}
\end{aligned}$$

Figure 1: Example: catching balls with parametric probabilistic transitions

Example 1. Consider two agents (A_1 and A_2) catching a repeatedly thrown ball, as shown in Figure 1. Each agent can either ‘catch’ or ‘skip’ catching, and possible states are:

$$\begin{aligned}
s_0 &= \text{dropped} \triangleq \neg \text{caught}_1 \wedge \neg \text{caught}_2 \\
s_1 &= \text{collision} \triangleq \text{caught}_1 \wedge \text{caught}_2 \\
s_2 &= \text{score}_1 \triangleq \text{caught}_1 \wedge \neg \text{caught}_2 \\
s_3 &= \text{score}_2 \triangleq \neg \text{caught}_1 \wedge \text{caught}_2
\end{aligned}$$

Let the probability of A_1 (A_2) taking action *skip* is $x_{1,1}$ ($x_{2,1}$) and taking action *catch* is $x_{1,2} = 1 - x_{1,1}$ ($x_{2,2} = 1 - x_{2,1}$) respectively, transition labels and the relevant probabilities are: $p_1 = x_1 x_2$, $p_2 = (1 - x_1)(1 - x_2)$, $p_3 = x_1(1 - x_2)$, $p_4 = x_1(1 - x_2)$, $\alpha_1 = \text{skip}_1 \text{skip}_2$, $\alpha_2 = \text{catch}_1 \text{catch}_2$, $\alpha_3 = \text{catch}_1 \text{skip}_2$, $\alpha_4 = \text{catch}_1 \text{catch}_2$. Let the reward for A_1 (A_2) taking action *catch* be $r_a^1(\text{Catch}) = 2$ ($r_a^2(\text{Catch}) = 1$) and taking action *skip* be $r_a^1(\text{Skip}) = 1$ ($r_a^2(\text{Skip}) = 2$).

Definition 3. A history ρ is a non-empty finite sequence $s_0 \alpha_0 s_1 \alpha_1 \dots s_k$ of states and joint actions, where $\alpha_i \in \text{Act}^{\text{Ag}}$ is the i^{th} joint action, and $\Delta(s_i, \alpha_i, s_{i+1}) > 0$. $\rho_s(i)$ denotes the i^{th} state of ρ , and $\rho_\alpha(i)$ denotes the i^{th} joint action of ρ . In this case, we may write $\rho_s(i) \xrightarrow{\rho_\alpha(i)} \rho_s(i+1)$. Let $\text{Hist}_{\mathcal{M}}(s)$ denote the set of histories of \mathcal{M} starting from state s .

Example 2. An example history of Example 1 starting from state s_0 can be:

$$\rho = s_0 \xrightarrow{(1-x_1).\text{Catch}_1 \quad x_2.\text{Skip}_2} s_2 \xrightarrow{x_1.\text{Skip}_1 \quad (1-x_2).\text{Catch}_2} s_3$$

$x.a$ denotes an agent taking action a with probability x .

Definition 4. A memoryless (mixed)² strategy for i is a function from the set of states to a probability distribution over agent’s set of actions $\sigma_i : S \rightarrow \text{Dist}(\text{Act}_i)$. A strategy profile is a tuple of strategies for a set of agents, it is denoted by $\vec{\sigma} = (\sigma_1, \sigma_2, \dots, \sigma_n)$. For a set (or coalition) of agents $J \subseteq \text{Ag}$, write $\vec{\sigma}_J$ for $(\sigma_i)_{i \in J}$.

Definition 5. A state s and a strategy profile $\vec{\sigma}$ induce a set of histories denoted by $\text{Hist}_{\mathcal{M}}^{\vec{\sigma}}(s)$. For a history $\rho \in \text{Hist}_{\mathcal{M}}^{\vec{\sigma}}(s)$, the probability of $\rho = s_0 \xrightarrow{\alpha_0} s_1 \dots \xrightarrow{\alpha_{k-1}} s_k$, with $s_0 = s$, is given by:

$$\mathbb{P}(\rho) \triangleq \prod_{j=1}^{k-1} \left(\prod_{i=1}^n (\sigma_i(s_0 \xrightarrow{\alpha_0} s_1 \dots \xrightarrow{\alpha_{j-1}} s_j)(\alpha_j(i))) \right)$$

Definition 6. The payoff function defined as a map from a set of histories to a real value $\wp : \text{Hist}_{\mathcal{M}}(s) \rightarrow \mathbb{R}^{|\text{Ag}|}$. \wp^i denotes the payoff function of $i \in \text{Ag}$ and is defined as:

$$\wp^i(\rho) \triangleq \left(\sum_{j=0}^{t-1} (r_a^i(\alpha_j) + r_s^i(s_j)) \cdot \Delta(s_j, \alpha_j, s_{j+1}) \right)$$

where $\rho = s_0 \xrightarrow{\alpha_0} s_1 \xrightarrow{\alpha_1} \dots \xrightarrow{\alpha_{t-1}} s_t \in \text{Hist}_{\mathcal{M}}(s_0)$.

Definition 7. Given a non-empty coalition of agents $J \subseteq \text{Ag}$, a strategy profile $\vec{\sigma}_J$, and a state $s \in S$, a joint plan is a function $\pi : J \rightarrow \text{Hist}_{\mathcal{M}}^{\vec{\sigma}_J}(s)$. $\text{Hist}_{\pi}(s)$ is the history consistent with π from s . Let $\pi^{-i} = \pi^{\text{Ag} \setminus \{i\}} = (\pi^1, \dots, \pi^{i-1}, \pi^{i+1}, \dots, \pi^n)$ denote joint plan without i .

Definition 8. Two joint plans π_1 and π_2 are $\langle J \rangle$ -compatible if the actions taken by coalition $J \subseteq \text{Ag}$ along the histories consistent with π_1 and π_2 are equivalent, denoted as $\pi_1 \sim_{\langle J \rangle} \pi_2$. $\text{Plan}_{\pi}^{\langle J \rangle}(s)$ represents the set of plans that are $\langle J \rangle$ -compatible with π starting from state s .

Example 3. Consider two joint plans π_1, π_2 in Example 1 when the ball is thrown twice, and which result in the following joint traces.

$$\begin{aligned}
\pi_1 &= (\text{Catch}_1 \text{Skip}_2), (\text{Skip}_1 \text{Catch}_2) \\
\pi_2 &= (\text{Catch}_1 \text{Catch}_2), (\text{Skip}_1 \text{Skip}_2)
\end{aligned}$$

Clearly, we have that $\pi_1 \sim_{\{A_1\}} \pi_2$.

Definition 9 (Causal Active Responsibility (CAR) (Parker, Grandi, and Lorini 2023)). Given a PSMAS \mathcal{M} , we say i bears CAR for outcome³ φ in joint plan π at state s , if φ holds for all possible histories consistent with $\text{Plan}_{\pi}^{\{i\}}(s)$ while violated in some histories consistent with $\text{Plan}_{\pi}^{\{\text{Ag}\}}(s)$.

Intuitively, i takes CAR for the occurrence of the outcome φ in π starting at s if keeping i ’s actions fixed the other agents could not avoid the outcome by choosing different actions. Note that multiple agents can individually have CAR.

²This specific type of strategy is also called *behavioural* (Cristau, David, and Horn 2010).

³Outcome is specified in PATL+R formula, presented in the next section.

Definition 10 (Degree of CAR). *Given \mathcal{M} , the degree of the responsibility of $i \in \text{Ag}$ bearing CAR for outcome φ under joint plan π at state s is defined as the ratio of the probability of behaviours following $\text{Plan}_\pi^{\{i\}}(s)$ and leading to ω , relative to that of all possible behaviours leading to φ , if the occurrence of φ is avoidable (i.e., there exists some behaviours following π and not leading to φ); and 0 otherwise.*

Definition 11 (Causal Passive Responsibility (CPR) (Parker, Grandi, and Lorini 2023)). *Given \mathcal{M} , we say $i \in \text{Ag}$ bears CPR for outcome φ in joint plan π at state s , if φ holds for all possible behaviours consistent with π , while is violated in some behaviours consistent with $\text{Plan}_\pi^{\{\text{Ag} \setminus \{i\}\}}(s)$.*

Intuitively, agent i takes CPR for the occurrence of outcome φ in π starting at s if keeping all other agents' actions fixed i could avoid the outcome by choosing different actions.

Definition 12 (Degree of CPR). *Given \mathcal{M} , the degree of the responsibility of i bearing CPR for outcome φ under joint plan π at state s is defined as the ratio of the probability of the behaviours following $\text{Plan}_\pi^{\{\text{Ag} \setminus \{i\}\}}(s)$ and not leading to φ relative to that of all possible behaviours violating φ if the occurrence of φ following π is achievable; and 0 otherwise.*

The Logic PATL+R

We introduce PATL+R, a variant of PATL that incorporates quantified reward and responsibility formulae. In this paper, we specifically consider *bounded path semantics*.

PATL+R over Finite (Bounded) Paths

Definition 13. *Let $\mathcal{M} = (\text{Ag}, S, s^0, V, \Delta, \text{Ap}, L)$. The syntax of PATL+R is made up of state formulae and path formulae represented by ϕ and ψ , respectively.*

$$\begin{aligned} \phi &::= a \mid \neg\phi \mid \phi \wedge \phi \mid \langle A \rangle \mathbf{P}_{\bowtie p}[\psi] \mid \langle A \rangle \mathbf{R}_{\bowtie q}[\diamond_{\leq k}\phi] \\ &\quad \mid \langle A \rangle \mathbf{D}_{\bowtie d}[\text{CAR}_{i,\pi}(\psi)] \mid \langle A \rangle \mathbf{D}_{\bowtie d}[\text{CPR}_{i,\pi}(\psi)] \\ \psi &::= \bigcirc\phi \mid \phi \mathbf{U}_{\leq k}\phi \end{aligned}$$

Here $a \in \text{Ap}$ is an atomic proposition, $A \subseteq \text{Ag}$ is a set of agents, $\langle A \rangle$ is the strategy quantifier, $i \in A$ is an agent, $\bowtie \in \{\leq, <, \geq, >\}$, $p \in [0, 1]$, $q, d \in \mathbb{R}$ are probability, reward and responsibility degree bounds, respectively, and $k \in \mathbb{N}$ is a time bound.

Note that a PATL+R formula is defined relative to a state; path formula only appear within the probabilistic operator $\langle A \rangle \mathbf{P}_{\bowtie p}[\cdot]$, the reward operator $\langle A \rangle \mathbf{R}_{\bowtie q}[\diamond_{\leq k}\phi]$, and the responsibility degree operator $\langle A \rangle \mathbf{D}_{\bowtie d}[\gamma_i]$, where γ_i denotes the responsibility operator $\text{CAR}_{i,\pi}(\psi)$ or $\text{CPR}_{i,\pi}(\psi)$. The formula $\langle A \rangle \mathbf{P}_{\bowtie p}[\psi]$ expresses that the coalition A has a strategy such that the probability of satisfying path formula ψ is $\bowtie p$ when the strategy is followed. The formula $\langle A \rangle \mathbf{R}_{\bowtie q}[\diamond_{\leq k}\phi]$ expresses that the coalition A has a strategy such that the expected rewards of satisfying path formula $\diamond_{\leq k}\phi$ is $\bowtie q$ when the strategy is followed, where $\diamond_{\leq k}\phi = \text{true } \mathbf{U}_{\leq k}\phi$. The formula $\langle A \rangle \mathbf{D}_{\bowtie d}[\gamma_i]$ expresses that the degree of responsibility of i of satisfying path formula ψ under coalition A is $\bowtie d$.

Definition 14. *Given a model \mathcal{M} , semantics for PATL+R are interpreted as follows. For a state $s \in S$ of \mathcal{M} , the satisfaction relation $s \models_{\mathcal{M}} \phi$ for state formula denotes “ s satisfies ϕ ”:*

- $s \models_{\mathcal{M}} a$ iff $a \in L(s)$.
- $s \models_{\mathcal{M}} \neg\phi$ iff $s \not\models_{\mathcal{M}} \phi$.
- $s \models_{\mathcal{M}} \phi \wedge \phi'$ iff $s \models_{\mathcal{M}} \phi$ and $s \models_{\mathcal{M}} \phi'$.
- $s \models_{\mathcal{M}} \langle A \rangle \mathbf{P}_{\bowtie p}[\psi]$ iff there exists a strategy profile of coalition A such that the probability of paths from state s which are consistent with the strategy and satisfy ψ is $\bowtie p$, i.e.,

$$\exists \vec{\sigma}_A. \mathbb{P}(\{\rho \in \text{Hist}_{\mathcal{M}}^{\vec{\sigma}_A}(s) \mid \rho \models_{\mathcal{M}} \psi\}) \bowtie p.$$

- $s \models_{\mathcal{M}} \langle A \rangle \mathbf{R}_{\bowtie q}[\diamond_{\leq k}\phi]$ iff there exists a strategy profile of coalition A such that the expected accumulated reward of paths from state s which are consistent with the strategy and satisfy ψ is $\bowtie q$, i.e.,

$$\exists \vec{\sigma}_A. \mathbb{E}_{\mathcal{M}}^{\vec{\sigma}_A}(s, \text{rew}(r, \diamond_{\leq k}\phi)) \bowtie q.$$

where:

$$\text{rew}(r, \diamond_{\leq k}\phi)(\rho) =$$

$$\begin{cases} \infty & \text{if } \forall j \leq k. \rho_s(j) \not\models_{\mathcal{M}} \phi \\ \sum_{i=0}^{k_\phi} (r_a(\rho_a(i)) + r_s(\rho_s(i))) & \text{otherwise} \end{cases}$$

where $k_\phi = \min\{j - 1 \mid \rho_s(j) \models_{\mathcal{M}} \phi\}$, and k denotes the time bound of $\diamond_{\leq k}\phi = \text{true } \mathbf{U}_{\leq k}\phi$.

- $s \models_{\mathcal{M}} \langle A \rangle \mathbf{D}_{\bowtie d}[\text{CAR}_{i,\pi}(\psi)]$ iff $\mathbb{D}^A(s, \text{CAR}_{i,\pi}(\psi)) \bowtie d$ where:

$$\mathbb{D}^A(s, \text{CAR}_{i,\pi}(\psi)) =$$

$$\begin{cases} 0 & \text{if } \nexists \pi'' \in \text{Plan}_\pi^{\{\text{Ag}\}}(s). (\forall \vec{\sigma}_A'' \in \pi''(A). \\ & \quad (\forall \rho \in \text{Hist}_{\mathcal{M}}^{\vec{\sigma}_A''}(s). \rho \not\models_{\mathcal{M}} \psi)) \\ \frac{\mathbb{P}(\{\rho \in \text{Hist}_{\mathcal{M}}^{\vec{\sigma}_A}(s) \mid \rho \models_{\mathcal{M}} \psi, \vec{\sigma}_A'' \in \pi''(A), \pi'' \in \text{Plan}_\pi^{\{\text{Ag}\}}(s)\})}{\mathbb{P}(\{\rho \in \text{Hist}_{\mathcal{M}}(s) \mid \rho \models_{\mathcal{M}} \psi\})} & \text{otherwise} \end{cases}$$

- $s \models_{\mathcal{M}} \langle A \rangle \mathbf{D}_{\bowtie d}[\text{CPR}_{i,\pi}(\psi)]$ iff $\mathbb{D}^A(s, \text{CPR}_{i,\pi}(\psi)) \bowtie d$ where:

$$\mathbb{D}^A(s, \text{CPR}_{i,\pi}(\psi)) =$$

$$\begin{cases} 0 & \text{if } \nexists \pi' \in \text{Plan}_{\mathcal{M}}^{\{\text{Ag}\}}(s). (\forall \vec{\sigma}_A' \in \pi'(A). \\ & \quad (\forall \rho \in \text{Hist}_{\mathcal{M}}^{\vec{\sigma}_A'}(s). \rho \models_{\mathcal{M}} \psi)) \\ \frac{\mathbb{P}(\{\rho \in \text{Hist}_{\mathcal{M}}^{\vec{\sigma}_A'}(s) \mid \rho \not\models_{\mathcal{M}} \psi, \vec{\sigma}_A'' \in \pi''(A), \pi'' \in \text{Plan}_\pi^{\{\text{Ag} \setminus \{i\}\}}(s)\})}{\mathbb{P}(\{\rho \in \text{Hist}_{\mathcal{M}}(s) \mid \rho \not\models_{\mathcal{M}} \psi\})} & \text{otherwise} \end{cases}$$

For a history $\rho \in \text{Hist}_{\mathcal{M}}(s_0)$, the satisfaction relation $\rho \models_{\mathcal{M}} \psi$ for a path formula ψ denotes that “ ρ satisfies ψ ”:

- $\rho \models_{\mathcal{M}} \bigcirc\phi$ iff $\rho_s(1) \models_{\mathcal{M}} \phi$.
- $\rho \models_{\mathcal{M}} \phi \mathbf{U}_{\leq k}\phi'$ iff there exists $i \leq k$ such that: $\rho_s(i) \models_{\mathcal{M}} \phi'$, and $\rho_s(j) \models_{\mathcal{M}} \phi$ for all $j < i$.

Example 4. Consider the outcome expressed by

$$\varphi = \langle A_1, A_2 \rangle \bigcirc (\text{dropped } \vee \text{score}_2)$$

in Example 1, assume agent A_1 (A_2) executes action skip with probability x_1 (x_2), and catch with probability $1 - x_1$ ($1 - x_2$), and consider joint plan $\pi = (\text{skip}_1 \text{skip}_2)$ and agent A_1 . Note that $\forall \pi' \in \text{Plan}_\pi^{\{A_1\}}(s_0)$, $\text{Hist}_{\pi'}(s_0) \models_{\mathcal{M}} \varphi$, (i.e., for $\pi' = (\text{skip}_1 \text{catch}_2)$ or $\pi' = (\text{skip}_1 \text{skip}_2)$). Therefore, while keeping the initial state and actions of A_1 fixed, the other agents (i.e., A_2 in this example) could not have

acted differently to avoid the occurrence of φ . Now consider $\pi'' = (\text{Catch}_1 \text{Skip}_2)$, and note that $\text{Hist}_{\pi''}(s) \not\models_{\mathcal{M}} \varphi$, i.e., there exists a joint plan $\pi'' \in \text{Plan}_{\mathcal{M}}^{\langle\{A_1, A_2\}\rangle}(s_0)$ which avoids the occurrence of φ . Thus, $s_0 \models_{\mathcal{M}} \text{CAR}_{\mathcal{M}}(A_1, \pi, \varphi)$ and A_1 bears CAR for φ in π at s_0 , and clearly:

$$\begin{aligned} \mathbb{P}(\varphi) &= x_1 x_2 + x_1(1 - x_2) = x_1 \\ \mathbb{P}(\neg\varphi) &= (1 - x_1)(1 - x_2) + x_2(1 - x_1) = 1 - x_1 \\ \mathbb{D}^{\langle A_1, A_2 \rangle}[\text{CAR}_{A_1, \pi}(\varphi)] &= \frac{x_1 x_2 + x_1(1 - x_2)}{x_1 x_2 + x_1(1 - x_2)} \cdot 1 = 1 \end{aligned}$$

Example 5. Consider the outcome expressed by

$$\varphi = \langle A_1, A_2 \rangle \circ \text{collision}$$

in Example 1, assume each agent executes *Skip* and *Catch* with probability x and $1 - x$ respectively. given a joint plan $\pi = (\text{Catch}_1 \text{Catch}_2)$ and agent A_1 . Here, $\text{Hist}_{\pi}(s_0) \models_{\mathcal{M}} \varphi$. Consider $\pi' = (\text{Skip}_1 \text{Catch}_2)$, and note that $\pi' \in \text{Plan}_{\pi}^{\langle\{A_2\}\rangle}(s_0)$ and $\text{Hist}_{\pi'}(s_0) \not\models_{\mathcal{M}} \varphi$, i.e., by keeping the initial state and the actions of all other agents fixed, A_1 can instead execute action *Catch*₁ to avoid the realisation of φ . Thus, A_1 bears CPR for φ in π at s_0 . Clearly,

$$\begin{aligned} \mathbb{P}(\varphi) &= (1 - x)^2 \\ \mathbb{D}^{\langle A_1, A_2 \rangle}[\text{CPR}_{A_1, \pi}(\varphi)] &= \frac{x(1-x)}{x^2 + x(1-x) + (1-x)^2} = \frac{x(1-x)}{1-x+x^2} \end{aligned}$$

Model Checking PATL+R

As PATL+R extends PATL, which operates within a branching-time logic framework, the fundamental model checking algorithm shares a basic structure similar to those used in CTL (see (Baier and Katoen 2008) for details). This algorithm operates through the recursive computation of the set $\text{Sat}(\phi)$, which represents states satisfying the formula ϕ within the model. Given a state formula ϕ , the algorithm recursively evaluates the truth values of its subformulae ϕ' across all states, starting from the propositional formulae of ϕ and following the recursive definitions of each modality. The problem of model checking a PSMAS w.r.t. arithmetic term comparison (excluding responsibility operators introduced in this paper) is similar to PATL and rPATL (see (Chen and Lu 2007; Kwiatkowska et al. 2021) for details). Therefore, we focus on evaluating responsibility degree formulae $\mathcal{E}_{s, \bar{\sigma}_A}(\mathbf{D}[\gamma_i])$.

CAR formulae. If $\gamma_i = \text{CAR}_{i, \pi}(\psi)$, we compute the probability of behaviours leading to the path formula ψ when following the strategies devised by i within the joint plan π , and the probability of all possible behaviours satisfying ψ . The degree of the CAR formula is computed as the ratio of these probabilities. Algorithm 1 outlines the procedure.

CPR formulae. If $\gamma_i = \text{CPR}_{i, \pi}(\psi)$, we compute the probability of behaviours violating the path formula ψ where agent i varies its actions while other agents maintain compatible actions in π . We also compute the probability of all possible behaviours violating ψ . The degree of the CPR formula is computed as the ratio of these probabilities. Algorithm 2 outlines the procedure.

Algorithm 1: Calculate $\mathcal{E}_{s, \bar{\sigma}_A}(\mathbf{D}[\text{CAR}_{i, \pi}(\psi)])$

```

1: Input:  $\mathcal{M}, s, i, \pi, \psi, A$ 
2: Output:  $\mathcal{E}_{s, \bar{\sigma}_A}(\mathbf{D}[\text{CAR}_{i, \pi}(\psi)])$  with parameters of  $\mathcal{M}$ 
3:  $\kappa \leftarrow 0$ 
4:  $\mathcal{L}_{\psi}, \mathcal{L}^+, \mathcal{L}^- \leftarrow \{\}, \{\}, \{\}$ 
5: for each history  $\rho \in \text{Hist}_{\mathcal{M}}(s)$  do
6:   if  $\rho \models_{\mathcal{M}} \psi$  then
7:      $\mathcal{L}_{\psi} \leftarrow \mathcal{L}_{\psi} \cup \{\rho\}$ 
8:   for each  $\pi'$  in compatible plans  $\text{Plan}_{\pi}^{\langle\{i\}\rangle}(s)$  do
9:     for each consistent history  $\rho \in \text{Hist}_{\pi'}(s)$  do
10:      if  $\rho \models_{\mathcal{M}} \psi$  then
11:         $\mathcal{L}^+ \leftarrow \mathcal{L}^+ \cup \{\rho\}$ 
12:      for each  $\pi''$  in all possible plans  $\text{Plan}_{\mathcal{M}}^{\langle A \rangle}(s)$  do
13:        for each consistent history  $\rho \in \text{Hist}_{\pi''}(s)$  do
14:          if  $\rho \not\models_{\mathcal{M}} \psi$  then
15:             $\mathcal{L}^- \leftarrow \mathcal{L}^- \cup \{\rho\}$ 
16: if  $|\mathcal{L}^-| > 0$  then
17:    $\kappa \leftarrow 1$ 
18: Return  $\mathbb{P}(\mathcal{L}^+) / \mathbb{P}(\mathcal{L}_{\psi}) * \kappa$ 

```

Algorithm 2: Calculate $\mathcal{E}_{s, \bar{\sigma}_A}(\mathbf{D}[\text{CPR}_{i, \pi}(\varphi)])$

```

1: Input:  $\mathcal{M}, s, i, \pi, \varphi, A$ 
2: Output:  $\mathcal{E}_{s, \bar{\sigma}_A}(\mathbf{D}[\text{CPR}_{i, \pi}(\varphi)])$  with parameters in  $\mathcal{M}$ 
3:  $\kappa \leftarrow 0$ 
4:  $\mathcal{L}_{\neg\varphi}, \mathcal{L}^+, \mathcal{L}^- \leftarrow \{\}, \{\}, \{\}$ 
5: for each history  $\rho \in \text{Hist}_{\mathcal{M}}(s)$  do
6:   if  $\rho \not\models_{\mathcal{M}} \varphi$  then
7:      $\mathcal{L}_{\neg\varphi} \leftarrow \mathcal{L}_{\neg\varphi} \cup \{\rho\}$ 
8:   for each  $\pi'$  in compatible plans  $\text{Plan}_{\pi}^{\langle\{A\}\rangle}(s)$  do
9:     for each consistent history  $\rho \in \text{Hist}_{\pi'}(s)$  do
10:      if  $\rho \models_{\mathcal{M}} \varphi$  then
11:         $\mathcal{L}^+ \leftarrow \mathcal{L}^+ \cup \{\rho\}$ 
12:      for each  $\pi''$  in all possible plans  $\text{Plan}_{\pi}^{\langle A \setminus \{i\} \rangle}(s)$  do
13:        for each consistent history  $\rho \in \text{Hist}_{\pi''}(s)$  do
14:          if  $\rho \not\models_{\mathcal{M}} \varphi$  then
15:             $\mathcal{L}^- \leftarrow \mathcal{L}^- \cup \{\rho\}$ 
16: if  $|\mathcal{L}^-| > 0$  then
17:    $\kappa \leftarrow 1$ 
18: Return  $\mathbb{P}(\mathcal{L}^-) / \mathbb{P}(\mathcal{L}_{\neg\varphi}) * \kappa$ 

```

Algorithms 1 and 2 both run in polynomial space, thus overall, we obtain the following theorem.

Theorem 1. *Model checking PATL+R formula is in PSPACE.*

Proof. The proof will be in two parts. (i) First, we show that model checking PATL+R formula without **D** operators is in PSPACE. This follows from the fact that the fragment of PATL+R without **D** operators corresponds to the logic rPATL proposed in (Kwiatkowska et al. 2019), whose model checking over CSGs is shown to be in PSPACE. (ii) We show that model checking PATL+R formulae of the form $\langle A \rangle \mathbf{D}_{\triangleright \text{ad}}[\text{CAR}_{i, \pi}(\psi)]$ or $\langle A \rangle \mathbf{D}_{\triangleright \text{ad}}[\text{CPR}_{i, \pi}(\psi)]$ can be done in PSPACE under the memoryless strategy assumption. No-

tice that in Algorithms 1 and 2, the plans π', π'' and history ρ have size of at most polynomial in \mathcal{M} . Furthermore, the checks in lines 6, 12, and 19 can be done in polynomial time. Thus, the algorithms run in PSPACE. As such, we obtain PSPACE model checking complexity. \square

Finding Stable Joint Plans

In this section, we discuss how to compute stable joint plans. To this end, we first introduce the notion of a *utility function*, which intuitively is a function that considers both reward/payoff and a responsibility degree.

Definition 15. Given $\mathcal{M} = (\text{Ag}, S, s^0, V, \Delta, \text{Ap}, L)$, a joint plan π , we define the payoff valuation function of $i \in \text{Ag}$ as the expected payoff of $\text{Hist}_\pi(s^0)$:

$$\mathcal{V}_\pi^i(s^0) \triangleq \sum_{\rho \in \text{Hist}_\pi(s^0)} \wp^i(\rho) \quad (1)$$

Example 6. The expected payoff of π_1 of A_1 given in Example 3 can be computed as:

$$\mathcal{V}_{\pi_1}^{A_1}(s^0) = 2 \cdot (1 - x_1)x_2 + 1 \cdot x_1(1 - x_2)$$

Definition 16. Given $\mathcal{M} = (\text{Ag}, S, s^0, V, \Delta, \text{Ap}, L)$, an outcome φ , and a joint plan π , we define the responsibility valuation function of $i \in \text{Ag}$ as the responsibility degree of $\text{Hist}_\pi(s^0)$:

$$\mathcal{R}_\pi^i(s_0) \triangleq \mathbb{D}^{\text{Ag}}[\text{CAR}_{i,\pi}(\varphi)] + \theta \cdot \mathbb{D}^{\text{Ag}}[\text{CPR}_{i,\pi}(\varphi)] \quad (2)$$

where θ is a Lagrange coefficient to adjust the weight of CAR degree and CPR degree.

The equation above defines the responsibility valuation function $\mathcal{R}_\pi^i(s^0)$ for agent i . By combining CAR and CPR with appropriate weights, the responsibility valuation function provides a measure of agent i 's overall responsibility in achieving the outcome φ within the specified joint plan π and a starting state s^0 .

Definition 17. The utility function is defined in terms of the payoff valuation and responsibility valuation as:

$$u_\pi^i \triangleq \lambda(\mathcal{V}_\pi^i(s^0), \mathcal{R}_\pi^i(s^0)) \quad (3)$$

Here, λ is a polynomial function allowing us to weigh the importance of responsibility and gaining a reward.

A simple instantiation may be a weighted linear function of the form $\lambda = \lambda_1 \mathcal{V}_\pi^i(s^0) - \lambda_2 \mathcal{R}_\pi^i(s^0)$. Different functions for λ can be viewed as different norms imposed on the agents in the system (Ciabatonni et al. 2023), in which case our approach can be viewed as ascribing a sanction or reward based on responsibility for achieving or violating the state of affairs associated with the norm.

Given a PSMAS \mathcal{M} , a PATL+R path formula φ , and utility function u_π^i for each i , our task is to compute an NE joint plan π that satisfies φ . The high-level procedure is as follows.

1. First we compute the set of strategy profiles satisfying the formula φ . We apply parametric model checking techniques on \mathcal{M} to compute rational valuation functions over parameters V^i for each agent $i \in \text{Ag}$, representing strategic transition probabilities in mixed strategies over each action $a \in \text{Act}$ in each state s .

2. Then, from this set of strategies, we restrict it to a subset of strategies that are NE. We do this by formulating a set of equations corresponding to NE, and solve them with respect to the set of strategies from step 1.

Nash Equilibrium for Responsibility and Utility

Definition 18. Given $\mathcal{M} = (\text{Ag}, S, s^0, V, \Delta, \text{Ap}, L)$, for each agent i , and $\pi^{\text{Ag} \setminus i}$, a plan π^i is a best response w.r.t. utility functions if it is the set

$$u\text{BR}_i(\pi^{\text{Ag} \setminus i}) \triangleq \{\pi^i \mid \max_{\pi^i} (\lambda(\mathcal{V}_{\pi^i, \pi^{\text{Ag} \setminus i}}^i(s^0), \mathcal{R}_{\pi^i, \pi^{\text{Ag} \setminus i}}^i(s^0)))\}$$

A joint plan π is considered a mixed Nash equilibrium (NE) if it belongs to the best response sets $\pi \in u\text{BR}_i(\pi^{\text{Ag} \setminus i})$ for all $i \in \text{Ag}$ in \mathcal{M} .

Computing expected payoff for reachability \mathcal{V} . The computation of the payoff valuation function, as given in (1), involves calculating the reachability rewards using the formula: $\mathcal{E}_{s, \vec{\sigma}_A}^i((A)\mathbf{R}[\diamond_{\leq k} \phi]) = \mathbb{E}(\wp^i(\{\rho \in \text{Hist}_{\mathcal{M}}(s, \vec{\sigma}_A) \mid \rho \models_{\mathcal{M}} \diamond_{\leq k} \phi\}))$, where $\wp^i(\rho) = \sum_{j=0}^{k_\phi} (r_a^i(\rho_a(j)) + r_s^i(\rho_s(j)))$, $k_\phi = \min\{k, k'\}$ and $k' \leq k$ s.t. $\rho_s(k') \models_{\mathcal{M}} \phi$.

Computing expected responsibility \mathcal{R} . The computation of the responsibility valuation function, as given in (2), involves calculating both the CAR and CPR degrees using the formula:

$$\begin{aligned} & \mathcal{E}_{s, \vec{\sigma}_A}(\mathbf{D}[\text{CAR}_{i,\pi}(\psi)]) \\ &= \begin{cases} 0 & \text{if } \exists \pi'' \in \text{Plan}_{\pi}^{\langle A \rangle}(s). (\forall \rho \in \text{Hist}_{\pi''}(s). \rho \not\models_{\mathcal{M}} \psi) \\ \frac{\mathbb{P}(\{\rho \in \text{Hist}_{\pi'}(s) \mid \rho \models_{\mathcal{M}} \psi \text{ for } \pi' \in \text{Plan}_{\pi}^{\langle i \rangle}(s)\})}{\mathbb{P}(\{\rho \mid \rho \models_{\mathcal{M}} \psi\})} & \text{otherwise} \end{cases} \\ & \mathcal{E}_{s, \vec{\sigma}_A}(\mathbf{D}[\text{CPR}_{i,\pi}(\psi)]) \\ &= \begin{cases} 0 & \text{if } \exists \pi' \in \text{Plan}_{\pi}^{\langle A \rangle}(s). (\forall \rho \in \text{Hist}_{\pi'}(s). \rho \models_{\mathcal{M}} \psi) \\ \frac{\mathbb{P}(\{\rho \in \text{Hist}_{\pi''}(s) \mid \rho \not\models_{\mathcal{M}} \psi \text{ for } \pi'' \in \text{Plan}_{\pi}^{\langle A \setminus i \rangle}(s)\})}{\mathbb{P}(\{\rho \mid \rho \not\models_{\mathcal{M}} \psi\})} & \text{otherwise} \end{cases} \end{aligned}$$

which can be obtained by Algorithm 1 and 2.

Finally, computing the best response set $u\text{BR}_i(\pi^{\text{Ag} \setminus i})$ corresponds to finding the set of plans π^i that maximise the utility.

Parametric best response expressions. When performing parametric model checking against PATL+R formulas to evaluate the best responses formulated above, we obtain parametric expressions representing the best response strategies. These parametric expressions contain parameters that represent the probabilities associated with different actions or strategies of the agents involved. The values of these parameters determine the optimal strategies for each agent within the joint plan. For example, in the case of utility-based best response, the parametric expression representing the probabilities of different actions that maximise the expected cumulative rewards for each agent. Similarly, for responsibility-based best response, the parameters may represent the probabilities associated with actions that minimise the expected responsibility degree for each agent. By varying the values of the parameters, we can explore different scenarios and analyse how changes in the probabilities of actions impact the overall effectiveness of the joint plan.

Example 7. Consider the scenario from Example 1, $\varphi = \langle A_1, A_2 \rangle \diamond_{\leq 2} (\text{collision} \vee \text{dropped})$. Assume the probability of A_1 and A_2 taking action skip is x_1 and x_2 respectively, and the reward value of A_1 taking skip and catch is 1 and 2 respectively, while the reward value of A_2 taking skip and catch is 2 and 1 respectively. We compute the expected payoff of the agents w.r.t φ as follows:

$$\begin{aligned} \mathbb{E}(\varphi^{A_1}(\llbracket \varphi \rrbracket)) &= 2(1-x_1)(1-x_2) + 2(1-x_1)x_2^2x_1 \\ &\quad + x_1^2(1-x_2)x_2 + 2(1-x_1)x_2(1-x_2) \\ &\quad + 2x_1(1-x_2)^2(1-x_1) + x_1x_2 \end{aligned}$$

And a similar calculation for $\mathbb{E}(\varphi^{A_2}(\llbracket \varphi \rrbracket))$.

Example 8. Continuing Example 7, $\varphi = \langle A_1, A_2 \rangle \diamond_{\leq 2} (\text{collision} \vee \text{dropped})$, and consider a joint plan $\pi = \{(\text{catch}_1 \text{skip}_2)(\text{skip}_1 \text{skip}_2)\}$, We have:

$$\mathbb{P}(\mathcal{L}_\varphi) = 1 + 4x_1x_2^2 - 4x_1^2x_2^2 - x_1^2 - x_2^2 - 2x_1x_2 + 4x_1^2x_2$$

$$\mathbb{D}_{\mathcal{M}}^A(s_0, [\text{CAR}_{A_1, \pi}(\varphi)]) = \frac{(1-x_1)(x_2^2x_1) + (1-x_1)(1-x_2)}{\mathbb{P}(\mathcal{L}_\varphi)}$$

$$\mathbb{D}_{\mathcal{M}}^A(s_0, [\text{CAR}_{A_2, \pi}(\varphi)]) = \frac{x_1x_2(x_2 - x_1x_2 + 1)}{\mathbb{P}(\mathcal{L}_\varphi)}$$

Finding optimised plans by solving equilibria equations.

Theorem 2. Given $\mathcal{M} = (\text{Ag}, S, s^0, V, \Delta, \text{Ap}, L)$, if the utility function u^i is monotonic on i 's mixed strategies, then the (mixed) joint plan $\pi = (\pi_1, \dots, \pi_n)$ is a mixed NE of \mathcal{M} iff for each $i \in \text{Ag}$, every pure strategy (with probability 1) of π^i is a best response to $\pi^{\text{Ag} \setminus i}$.

The theorem establishes the equivalence between being a mixed Nash Equilibrium and every pure strategy being a best response to the complementary mixed strategies. The proof can be obtained by contradiction, details can be found in (Mu, Najib, and Oren 2024).

Theorem 2 has implications for computing NE. If a mixed plan π is a best response, then each pure strategy within the mix must also be a best response, implying that all pure strategies in the mix yield the same expected utility. In other words, every choice in the support of any agent's equilibrium mixed plan must result in the same valuation: $u^i(\pi^{-i}, \pi^i) = u^i(\pi^{-i}, \pi^{i'})$ for any two actions $a, a' \in \text{Act}$ corresponding to $\pi^i, \pi^{i'}$ with positive probabilities. This fact allows us to express the NE condition as follows:

$$\begin{cases} u^i(\pi^{-i}, \pi^i) = u^i(\pi^{-i}, \pi^{i'}) & \forall \pi^i, \pi^{i'} \\ \sum_{j=1}^m \pi_{ij} = 1 & \forall i \\ 0 \leq \pi_{ij} \leq 1 & \forall i, j. \end{cases} \quad (4)$$

Solving these equations identifies the NE and finds the optimal joint plan for given objectives. Thus, we obtain the following theorem. The proof is obtained from a reduction to the existential theory of the real numbers (Canny 1988).

Theorem 3. Computing NE joint plan satisfying PATL+R formula φ can be done in PSPACE.

Example 9. Continuing Example 7 (where we use a linear weighted function such that $\lambda_1 = 1, \lambda_2 = 0$) for best response calculation, by Theorem 2, we have:

$$\begin{cases} 2(1-x_2) + 4x_2(1-x_2) = x_2 + 2(1-x_2)x_2 \\ 2(1-x_1) = x_1 + 8(1-x_1)x_2 \\ 0 \leq x_1 \leq 1 \\ 0 \leq x_2 \leq 2 \end{cases}$$

by solving the above NE conditions, we get:

$$x_1 = \sqrt{17}/4 - 1/4 \approx 0.78, \quad x_2 = 2/3,$$

i.e., when A_1 takes action of skip with probability of 0.78 and A_2 takes action of skip with probability of 2/3, an NE equilibrium is reached.

Example 10. Continuing Example 8 for best response calculation, by Theorem 2, we have:

$$\begin{cases} \frac{(1-0)(0 \cdot x_2^2 + 1 - x_2)}{1 + 4 \cdot 0 \cdot x_2^2 - 4 \cdot 0^2 \cdot x_2^2 - x_1^2 - x_2^2 - 2 \cdot 0 \cdot x_2 + 4 \cdot 0^2 \cdot x_2} \\ = \frac{(1-1)(1 \cdot x_2^2 + 1 - x_2)}{1 + 4 \cdot 1 \cdot x_2^2 - 4 \cdot 1^2 \cdot x_2^2 - x_1^2 - x_2^2 - 2 \cdot 1 \cdot x_2 + 4 \cdot 1^2 \cdot x_2} \\ = \frac{x_1 \cdot 0(0 - x_1 \cdot 0 + 1)}{1 + 4x_1 \cdot 0^2 - 4x_1^2 \cdot 0^2 - x_1^2 - 0^2 - 2x_1 \cdot 0 + 4x_1^2 \cdot 0} \\ = \frac{x_1 \cdot 1(1 - x_1 \cdot 1 + 1)}{1 + 4x_1 \cdot 1^2 - 4x_1^2 \cdot 1^2 - x_1^2 - 1^2 - 2x_1 \cdot 1 + 4x_1^2 \cdot 1} \\ 0 \leq x_1 \leq 1 \\ 0 \leq x_2 \leq 1 \end{cases}$$

by solving the above NE conditions, we get: $x_1 = 0$ and $x_2 = 1$, i.e., when A_1 takes action of catch while A_2 takes action of skip under joint plan π , an NE equilibrium is reached. This meets our intuition, as the outcome considered in the example for responsibility analysis is $\varphi = \langle A_1, A_2 \rangle \diamond_{\leq 2} (\text{collision} \vee \text{dropped})$, to minimise their responsibility, the best option would be for one of the agents to catch the ball. Additionally, the joint plan considered in the example is $\pi = \{(\text{catch}_1 \text{skip}_2)(\text{skip}_1 \text{skip}_2)\}$, therefore, the optimised plan regarding minimised responsibility would have A_1 catch the ball while A_2 skips catching.

Conclusions and Future Work

In this paper we developed an approach to multi-agent strategic reasoning with responsibility-aware agents. We introduced PATL+R, a logic that can be used to reason about causal responsibility and temporal properties. We provide an approach to model-check PATL+R formulae against a parametric model of probabilistic MAS and show that it can be done in PSPACE and is thus no more difficult than the model checking problem of the extended rPATL logic as presented in (Kwiatkowska et al. 2021). We also show how to synthesise NE joint strategies/plans that satisfy some PATL+R property, which can also be done in PSPACE.

Building on this paper, several promising avenues for future work are immediately apparent. Perhaps the most obvious is to consider finite-memory strategy settings. As shown in (Kwiatkowska et al. 2018; Svorenová and Kwiatkowska 2016), bounded temporal properties often require finite-memory strategies. It would also be interesting to explore how to extend our approach to a more expressive logic, such as Probabilistic Strategic Logic (PSL) (Aminof et al. 2019). While this would enable us to model complex scenarios and reason about agent behavior more easily, we must also consider the associated increase in complexity. Another avenue for exploration is the following question: what can be done if no NE joint plan satisfies the required property? One approach to addressing this is to ‘‘repair’’ the system, for example by introducing explicit norms (Bulling and Dastani 2016) or modifying the reward structure using reward machines (Dann et al. 2022).

References

- Abate, A.; Gutierrez, J.; Hammond, L.; Harrenstein, P.; Kwiatkowska, M.; Najib, M.; Perelli, G.; Steeples, T.; and Wooldridge, M. J. 2021. Rational verification: game-theoretic verification of multi-agent systems. *Appl. Intell.*, 51(9): 6569–6584.
- Alechina, N.; Halpern, J. Y.; and Logan, B. 2017. Causality, Responsibility and Blame in Team Plans. In *AAMAS*, 1091–1099. ACM.
- Alur, R.; Henzinger, T. A.; and Kupferman, O. 2002. Alternating-time temporal logic. *J. ACM*, 49(5): 672–713.
- Aminof, B.; Kwiatkowska, M.; Maubert, B.; Murano, A.; and Rubin, S. 2019. Probabilistic Strategy Logic. In *IJCAI*, 32–38.
- Baier, C.; Funke, F.; and Majumdar, R. 2021. A Game-Theoretic Account of Responsibility Allocation. In Zhou, Z., ed., *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI 2021, Virtual Event / Montreal, Canada, 19-27 August 2021*, 1773–1779. ijcai.org.
- Baier, C.; and Katoen, J. 2008. *Principles of model checking*. MIT Press. ISBN 978-0-262-02649-9.
- Bowling, M.; Jensen, R.; and Veloso, M. 2003. A formalization of equilibria for multiagent planning. In *IJCAI*, 1460–1462.
- Bulling, N.; and Dastani, M. 2016. Norm-based mechanism design. *Artificial Intelligence*, 239: 97–142.
- Canny, J. 1988. Some algebraic and geometric computations in PSPACE. In *Proceedings of the twentieth annual ACM symposium on Theory of computing*, 460–467.
- Chen, T.; and Lu, J. 2007. Probabilistic Alternating-time Temporal Logic and Model Checking Algorithm. In *FSKD*, 35–39. IEEE Computer Society.
- Chockler, H.; and Halpern, J. Y. 2004. Responsibility and Blame: A Structural-Model Approach. *J. Artif. Intell. Res.*, 22: 93–115.
- Chrapa, L.; and Karpas, E. 2024. On Verifying Linear Execution Strategies in Planning Against Nature. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 34, 86–94.
- Ciabattoni, A.; Horty, J. F.; Slavkovik, M.; van der Torre, L.; and Knoks, A. 2023. Normative Reasoning for AI (Dagstuhl Seminar 23151). *Dagstuhl Reports*, 13(4): 1–23.
- Ciuni, R.; and Mastop, R. 2009. Attributing distributed responsibility in stit logic. In *Logic, Rationality, and Interaction: Second International Workshop, LORI 2009, Chongqing, China, October 8-11, 2009. Proceedings 2*, 66–75. Springer.
- Cosner, R. K.; Chen, Y.; Leung, K.; and Pavone, M. 2023. Learning responsibility allocations for safe human-robot interaction with applications to autonomous driving. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 9757–9763. IEEE.
- Cristau, J.; David, C.; and Horn, F. 2010. How do we remember the past in randomised strategies? In Montanari, A.; Napoli, M.; and Parente, M., eds., *GANDALF*, volume 25 of *EPTCS*, 30–39.
- Dann, M.; Yao, Y.; Alechina, N.; Logan, B.; and Thangarajah, J. 2022. Multi-Agent Intention Progression with Reward Machines. In Raedt, L. D., ed., *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*, 215–222. International Joint Conferences on Artificial Intelligence Organization. Main Track.
- Demri, S.; and Fervari, R. 2023. Model-checking for ability-based logics with constrained plans. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 6305–6312.
- Dignum, V. 2020. Responsibility and artificial intelligence. *The oxford handbook of ethics of AI*, 4698: 215.
- Fisman, D.; Kupferman, O.; and Lustig, Y. 2010. Rational synthesis. In *Tools and Algorithms for the Construction and Analysis of Systems: 16th International Conference, TACAS 2010, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2010, Paphos, Cyprus, March 20-28, 2010. Proceedings 16*, 190–204. Springer.
- Gutierrez, J.; Hammond, L.; Lin, A. W.; Najib, M.; and Wooldridge, M. 2021. Rational Verification for Probabilistic Systems. In *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning*, volume 18, 312–322.
- Henne, P.; O’Neill, K.; Bello, P.; Khemlani, S.; and Brigard, F. D. 2021. Norms Affect Prospective Causal Judgments. *Cogn. Sci.*, 45(1).
- Hyland, D.; Gutierrez, J.; Shankaranarayanan, K.; and Wooldridge, M. 2024. Rational Verification with Quantitative Probabilistic Goals. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems*, 871–879.
- Jamroga, W. 2004. Strategic planning through model checking of ATL formulae. In *Artificial Intelligence and Soft Computing-ICAISC 2004: 7th International Conference, Zakopane, Poland, June 7-11, 2004. Proceedings 7*, 879–884. Springer.
- Kaur, D.; Uslu, S.; Rittichier, K. J.; and Durrezi, A. 2022. Trustworthy artificial intelligence: a review. *ACM computing surveys (CSUR)*, 55(2): 1–38.
- Kobayashi, T.; Bondu, M.; and Ishikawa, F. 2023. Formal modelling of safety architecture for responsibility-aware autonomous vehicle via event-b refinement. In *International Symposium on Formal Methods*, 533–549. Springer.
- Kwiatkowska, M.; Norman, G.; Parker, D.; and Santos, G. 2018. Automated verification of concurrent stochastic games. In *Quantitative Evaluation of Systems: 15th International Conference, QEST 2018, Beijing, China, September 4-7, 2018. Proceedings 15*, 223–239. Springer.
- Kwiatkowska, M.; Norman, G.; Parker, D.; and Santos, G. 2019. Equilibria-based probabilistic model checking for concurrent stochastic games. In *International Symposium on Formal Methods*, 298–315. Springer.

Kwiatkowska, M.; Norman, G.; Parker, D.; and Santos, G. 2021. Automatic verification of concurrent stochastic systems. *Formal Methods in System Design*, 58(1): 188–250.

Mittelmann, M.; Maubert, B.; Murano, A.; and Perrussel, L. 2023. Formal verification of bayesian mechanisms. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 11621–11629.

Mu, C.; Najib, M.; and Oren, N. 2024. Responsibility-aware Strategic Reasoning in Probabilistic Multi-Agent Systems. *CoRR*, abs/2411.00146.

Naumov, P.; and Tao, J. 2019. Blameworthiness in strategic games. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 3011–3018.

Naumov, P.; and Tao, J. 2021. Two Forms of Responsibility in Strategic Games. In *IJCAI*, 1989–1995. ijcai.org.

Parker, T.; Grandi, U.; and Lorini, E. 2023. Anticipating Responsibility in Multiagent Planning. In *ECAI*, volume 372, 1859–1866.

Schnoor, H. 2010. Strategic planning for probabilistic games with incomplete information. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: volume 1-Volume 1*, 1057–1064.

Shapley, L. S. 1953. Stochastic games. *Proceedings of the national academy of sciences*, 39(10): 1095–1100.

Svorenřová, M.; and Kwiatkowska, M. 2016. Quantitative verification and strategy synthesis for stochastic games. *European Journal of Control*, 30: 15–30.

Van Der Hoek, W.; and Wooldridge, M. 2002. Tractable multiagent planning for epistemic goals. In *Proceedings of the first international joint conference on Autonomous agents and multiagent systems: part 3*, 1167–1174.

Yazdanpanah, V.; Dastani, M.; Alechina, N.; Logan, B.; and Jamroga, W. 2019. Strategic responsibility under imperfect information. In *18th International Conference on Autonomous Agents and Multiagent Systems AAMAS 2019*. IFAAMAS.

Yazdanpanah, V.; Gerding, E. H.; Stein, S.; Dastani, M.; Jonker, C. M.; Norman, T. J.; and Ramchurn, S. D. 2023. Reasoning about responsibility in autonomous systems: challenges and opportunities. *AI & SOCIETY*, 38(4): 1453–1464.