Anticipating Responsibility for Plan Selection

Timothy Parker ;*

IRIT, CNRS, University of Toulouse, France

Abstract. In this paper we present two methods for plan comparison in a multi-objective planning setting with multiple agents. We use responsibility anticipation to define a form of regret-minimising plan selection. We also define a notion of complete responsibility to better handle responsibility for multiple outcomes. Finally, we perform a formal and experimental analysis of our comparison methods.

1 Introduction

Responsibility attribution [1, 2, 5, 10] is the process of determining which agent or set of agents can be held responsible for a particular outcome. This is a backward-looking process, meaning that while it is useful for allocating praise or blame, it cannot be used for plan selection, since responsibility for an outcome cam only be determined once the outcome has occurred. Responsibility anticipation [18] is the process of predicting which outcomes an agent *may* be responsible for if it performs a particular plan. This means it can be used in plan selection, such as by ensuring that an agent cannot be responsible for some negative outcome.

In this paper we present multiple methods for plan comparison in a multi-agent, multi-objective setting that make use of anticipated responsibility. We use responsibility to define a notion of regret, based on the goals or values that an agent is responsible for violating (and responsible for satisfying). This allows us to define a symbolic equivalent to regret-minimising plan comparison. We also introduce a notion of complete responsibility to better handle responsibility for multiple outcomes. We examine the axiomatic properties of our comparison rules as well as an experimental evaluation.

The rest of the paper is organised as follows, section 2 locates our paper in its field and compares our work with some related approaches from the literature. Section 3 introduces our planning model and section 4 defines our notions of responsibility and regret. Section 5 contains our plan comparison methods, section 6 does some analysis of these methods and section 7 describes our experiment and discusses the results. Finally, section 8 concludes the paper and outlines directions for future work.

2 Related Work

This paper builds on our previous work [18] in which we define attribution and anticipation for active, passive and contributive responsibility. This is built on work by Lorini et al. [15] and Braham and van Hees [4]. The current paper uses the notion of passive responsibility to create plan comparison methods in a multiagent planning setting with multiple objectives or values. Other approaches to formalising responsibility have been developed by Alechina et al. [1], Chockler and Halpern [5, 10] and Baier et al. [2]. However, to the best of our knowledge, ours is the only work to consider responsibility anticipation and its application to plan comparison.

We also show how passive responsibility can be used to define a notion of regret, for use in regret minimisation, a notion first introduced (independently) in decision theory by Savage [19] and Niehans [17]. It was later introduced to game theory by Linhart and Radner [14].

Our concept of regret is also similar to the notion of guilt introduced by Lorini and Mühlenbernd [16]. Their model uses separate numerical values to track both the individual utility of an agent for some particular history, as well as the degree of ideality of that history (such as the utility of the worst-off agent). The guilt of an agent in a history is the difference between the ideality achieved and the best possible ideality (fixing the actions of all other agents). This is similar to our notion of regret based on passive responsibility, but is purely numeric instead of symbolic.

3 Model

In this section we introduce the planning framework for our model.

3.1 Agents, Actions and Histories

Our model requires a finite set of agents Agt and a countable set of propositions $Prop = \{p, q, \ldots\}$. From $Prop$ we define a set of states $S = 2^{Prop}$, with elements $s, s', \ldots$ Let $Act = \{a, b, \ldots\}$ be a finite non-empty set of action names. To describe the actions taken by all agents at a single time we introduce the notion of a joint action, which is a function $J : Agt \longrightarrow Act$. The set of all joint actions is $JAct$.

To trace the actions of agents and changing states over time we define a k-history to be a pair $H = (H_{st}, H_{act})$ with $H_{st} : \{0, \ldots, k\} \longrightarrow S$ and $H_{act} : \{0, \ldots, k-1\} \longrightarrow JAct$. The set of k-histories is noted $Hist_k$. The set of all histories is $Hist = \bigcup_{k \in \mathbb{N}} Hist_k$.

3.2 Multi-Agent Action Theory

For simplicity, we define our action theory as a function:

$$\tau : S \times JAct \longrightarrow S$$

We say that history H is a τ-compatible history for action theory τ if each state respects the actions performed in the previous state. This means that for all $t \in \{0, \ldots, k-1\}$, $H_{st}(t+1) = \tau(H_{st}(t), H_{act}(t))$. The set of τ-compatible histories is noted $Hist(\tau)$.

* Corresponding Author. Email: timothy.parker@irit.fr

3.3 Planning Domains

Definition 1 (Multi-Agent Planning Domain)**.** *Multi-agent Planning Domain (MPD) is a tuple* $\nabla = (\tau, s_0)$ *where* τ *is an action theory and* $s_0 \in S$ *is an initial state.*

3.4 Action Sequences and Joint Plans

Now that we have defined a planning domain, we can define the notions of action sequence and plan. Given $k \in \mathbb{N}$, a k-action-sequence is a function

$$\pi : \{0, \dots, k-1\} \longrightarrow Act.$$

The set of k-action-sequences is noted Seq_k. The set of all action sequences is $Seq = \bigcup_{k \in \mathbb{N}} Seq_k$. For a (non-empty) coalition of agents $J \in 2^{Agt} \setminus \emptyset$ we can define a joint k-plan as a function $\Pi : J \longrightarrow Seq_k$ (if J is a singleton coalition we call Π an individual plan). The set of joint k-plans for a coalition J is written $Plan^J_k$. The set of all joint plans for J is $Plan^J = \bigcup_{k \in \mathbb{N}} Plan^J_k$.

Given a joint plan Π for coalition J and another coalition $J' \subseteq J$, we can write the sub-plan of Π corresponding to J' as $\Pi^{J'}$, we can also write $\Pi^{-J'}$ for sub-plan corresponding to $J \setminus J'$. Given two k-plans Π_1 and Π_2 for disjoint coalitions J_1, J_2, we write $\Pi_1 \cup \Pi_2$ for the joint plan for $J_1 \cup J_2$ such that $(\Pi_1 \cup \Pi_2)^{J_1} = \Pi_1$ and $(\Pi_1 \cup \Pi_2)^{J_2} = \Pi_2$. Finally, given two plans Π_1 and Π_2, if there exists some plan Π_3 such that $\Pi_2 = \Pi_1 \cup \Pi_3$ then we say that Π_1 *is compatible with* Π_2.

We can now define the notion of the history generated by a joint k-plan Π in the planning domain $\nabla = (\tau, s_0)$. It is the τ-compatible k-history along which the agents jointly execute the plan Π starting at state s_0. We write this as $H^{\Pi,\nabla}$.

3.5 Linear Temporal Logic

In our model, histories are temporal entities that are always finite in length, therefore the most natural choice to describe properties of histories is Linear Temporal Logic over Finite Traces [6, 7]. This allows us to describe temporal properties such as "φ never occurs" or "φ always occurs immediately after ψ". We write the language as $\mathcal{L}_{\mathsf{LTL}_f}$, defined by the following grammar:

$$\varphi \quad ::= \quad p \mid \neg\varphi \mid \varphi \wedge \varphi \mid \mathsf{X}\varphi \mid \varphi \,\mathsf{U}\, \varphi,$$

with p ranging over $Prop$. Atomic formulas in this language are those that consist of a single proposition p. X and U are the operators "next" and "until" of LTL_f. Operators "henceforth" (G) and "eventually" (F) are defined in the usual way: $\mathsf{G}\varphi \overset{\text{def}}{=} \neg(\top \,\mathsf{U}\, \varphi)$ and $\mathsf{F}\varphi \overset{\text{def}}{=} \neg\mathsf{G}\neg\varphi$. We define the semantics for X and U as follows, the rest is the same as $\mathcal{L}_{\mathsf{PL}+}$ (for $t \in \{0, \dots, k\}$).

$$\begin{aligned} H, t \models \mathsf{X}\varphi \quad &\Longleftrightarrow \quad t < k \text{ and } H, t+1 \models \varphi, \\ H, t \models \varphi_1 \,\mathsf{U}\, \varphi_2 \quad &\Longleftrightarrow \quad \exists t' \geq t : t' \leq k \text{ and } H, t' \models \varphi_2 \text{ and} \\ &\qquad \forall t'' \geq t : \text{ if } t'' < t' \text{ then } H, t'' \models \varphi_1. \end{aligned}$$

3.6 Values and Goals

We assume that our agents will have multiple goals and/or values that they wish to satisfy, which may have different priority levels. We combine these into a single value base $\overline{\Omega}$ which is simply a set of LTL_f-formulas. We let $\overline{\Omega}^+$ be the set $\overline{\Omega} \cup \{\neg\omega : \omega \in \overline{\Omega}\}$. To represent the structure of $\overline{\Omega}$ we assume the existence of a relation $\preceq$ on subsets of $\overline{\Omega}^+$ that indicates whether the satisfaction or violation of a certain set of values is preferable to another. We also define $X \prec Y$ as $X \preceq Y$ and $Y \not\preceq X$. For example $\{\omega_1, \neg\omega_2\} \prec \{\neg\omega_1, \omega_2\}$ indicates that it is more important to satisfy ω_2 than ω_1. In order for our comparison methods to function correctly, we require that $\preceq$ is transitive, reflexive, and strongly connected.

We define the relation on $\overline{\Omega}^+$ instead of $\overline{\Omega}$ to allow us to distinguish cases where the value is violated and cases where the status of the value is irrelevant (such as if the satisfaction or violation of the value is already guaranteed).

Weighted Values A simple way of creating $\preceq$ for some value base $\overline{\Omega}$ is to assign each value $\omega \in \overline{\Omega}$ a numerical weighting $Val(\omega)$, [13]. The procedure to calculate $\preceq$ from this weighting is quite straightforward. Given a value set X we can calculate the utility of each set by $Val(X) = \Sigma_{\omega \in X} Val(\omega) - \Sigma_{\neg\omega \in X} Val(\omega)$. Then we say that $X \preceq Y$ if and only if $Val(X) \leq Val(Y)$.

Note that there are many alternative ways to rank sets of values. For some approaches from single-agent planning, see the work of Bienvenu et al. [3], Dennis et al. [8] and Grandi et al. [9].

4 Responsibility, Regret and Anticipation

In this paper we focus purely on the causal aspects of responsibility. In our previous work [18] we defined notions of active, passive and contributive causal responsibility. Active responsibility means to cause an outcome to occur by guaranteeing that it happens. Passive responsibility means to allow an outcome to occur while being able to prevent it (fixing the actions of all other agents). Contributive responsibility means to be part of a coalition of agents who guarantee the outcome. For simplicity, this paper only considers passive responsibility.

Definition 2 (Passive Responsibility)**.** *Let* $\nabla = (\tau, s_0)$ *be an MPD,* $i \in Agt$ *an agent, and* Π_1 *a joint plan. Let* $\omega \in \mathcal{L}_{\mathsf{LTL}_f}$*. Then, we say that* i *bears* Causal Passive Responsibility *(CPR) for* ω *in* (Π_1, ∇) *if* $H^{\Pi_1,\nabla} \models \omega$ *and there exists some* Π_2 *compatible with* $\Pi_1^{-\{i\}}$ *such that* $H^{\Pi_2,\nabla} \not\models \omega$*.*

An agent i is passively responsible for some outcome ω if by acting differently it could have prevented ω, keeping fixed the initial state and the actions of all other agents.

The idea of "regret" is meant to represent the gap between what "could have happened" and what actually happened from the perspective of some agent. A simple notion of regret is to take the set of all values (and violations of values) that the agent is passively responsible for. If the agent is passively responsible for a formula ω this means that, fixing the actions of all other agents, they could have ensured $\neg\omega$.

Definition 3 (Regret)**.** *Given a history* H*, a planning domain* ∇*, a value base* $(\overline{\Omega}, \preceq)$*, and an agent* i*, the regret set of* i *in* H *is the lowest-ranked (according to* $\preceq$*) maximal subset of* $\overline{\Omega}^+$ *that* i *is passively responsible for.*

An issue with the simple notion of regret is that if an agent is passively responsible for φ_1, φ_2 and φ_3, this does not mean that they could have brought about $\neg\varphi_1 \wedge \neg\varphi_2 \wedge \neg\varphi_3$, merely that they could have brought about each $\neg\varphi_i$ individually. The following definition of responsibility more closely matches the sense of "responsible for φ_1 *and* φ_2".

Definition 4 (Complete Responsibility)**.** *Given a history H, a planning domain ∇, a value base $(\overline{\Omega}, \preceq)$, and a conjunction $\omega^+ = \omega_1 \wedge \ldots \wedge \omega_n$ of elements of $\overline{\Omega}^+$, we say that i is completely responsible for ω^+ if $H \models \omega^+$ and i is passively responsible for $\omega^- = \omega_1 \vee \ldots \vee \omega_n$ in H under ∇.*

This means that if an agent is completely responsible for $\varphi_1 \wedge \varphi_2 \wedge \varphi_3$ if $\varphi_1 \wedge \varphi_2 \wedge \varphi_3$ occurred and, fixing the actions of all other agents, the agent could have brought it about that $\neg\varphi_1 \wedge \neg\varphi_2 \wedge \neg\varphi_3$. We can then use this notion of responsibility to define a notion of regret.

Definition 5 (Consistent Regret)**.** *Given a history H, a planning domain ∇, a value base $(\overline{\Omega}, \preceq)$, and an agent i, the consistent regret set of i in H is the lowest-ranked (according to $\preceq$) maximal subset of $\overline{\Omega}^+$ that i is completely responsible for.*

All of the previous notions are defined retrospectively, meaning that they can only be evaluated after a plan has been executed and a history has been generated. This means that they are not useful for plan comparison, as we want to be able to compare plans before executing them. Therefore we introduce the notion of anticipation, which means to perform a retrospective evaluation on one or more possible outcomes. In other words, rather than executing a plan and then evaluating the result, we evaluate all possible results of executing the plan. We can define anticipation for both regret and passive responsibility.

Definition 6 (Regret Anticipation)**.** *Given an individual plan Π for an agent i, a planning domain ∇ and a value base $(\overline{\Omega}, \preceq)$, the anticipated (consistent) regret set of i in H is the lowest-ranked (according to $\preceq$) (consistent) regret set out of all histories H compatible with Π in ∇.*

Definition 7 (Passive Responsibility Anticipation)**.** *Given an individual plan Π for an agent i and a planning domain ∇, we say that i anticipates passive responsibility for some outcome φ if i is passively responsible for φ in some history H compatible with Π in ∇.*

5 Plan Comparison

We now present two notions of plan comparison based on notions of responsibility (as regret is based on responsibility). Anticipation-minimising comparison aims to minimise the set of values that the agent could be responsible for violating (and maximise the set that they could be responsible for satisfying) *in all possible outcomes* for a given plan. Regret-minimising comparison (and consistent-regret-minimising) comparison aims to minimise the set of values that the agent is (consistently) responsible for *in any single outcome*.

Definition 8 (Anticipation-minimising comparison)**.** *Given two plans Π_1 and Π_2 for agent i, a planning domain ∇ and a value base $(\overline{\Omega}, \preceq)$, we say that Π_2 is anticipation-minimising preferred to Π_1 if $X_1 \prec X_2$ where X_1 is the set of elements of $\overline{\Omega}^+$ that i anticipates responsibility for in Π_1 and* mutatis mutandis *for X_2 and Π_2.*

Definition 9 (Regret-minimising comparison)**.** *Given two plans Π_1 and Π_2 for agent i, a planning domain ∇ and a value base $(\overline{\Omega}, \preceq)$, we say that Π_2 is (consistent-)regret-minimising preferred to Π_1 if $X_1 \prec X_2$ where X_1 is the anticipated regret set of Π_1 and X_2 is the anticipated (consistent) regret set of Π_2.*

By minimising the worst-possible regret in the selected plan, the agent should select a plan either matches or is close to the best outcome for every possible joint plan of the other agents. This means that the plan selected is robustly good regardless of the actions of the other agents. Furthermore, while a regret-minimising agent may sometimes get bad outcomes, they should only get bad outcomes when the best possible outcome was also not very good.

For later comparison, we also present two very simple methods of plan comparison, called optimistic and pessimistic comparison. Optimistic comparison compares any two plans according to their best-case outcome, and pessimistic comparison compares any two plans according to their worst-case outcome.

Definition 10 (Optimistic Comparison)**.** *Given two plans Π_1 and Π_2 for agent i, a planning domain ∇ and a value base $(\overline{\Omega}, \preceq)$, we say that Π_2 is optimistic-preferred to Π_1 if $X_1 \prec X_2$ where X_1 is a highest-ranked (according to $\preceq$) set of values satisfied by some history H compatible with Π_1 in ∇ and mutatis mutandis for X_2.*

Definition 11 (Pessimistic Comparison)**.** *Given two plans Π_1 and Π_2 for agent i, a planning domain ∇ and a value base $(\overline{\Omega}, \preceq)$, we say that Π_2 is pessimistic-preferred to Π_1 if $X_1 \prec X_2$ where X_1 is a lowest-ranked (according to $\preceq$) set of values satisfied by some history H compatible with Π_1 in ∇ and mutatis mutandis for X_2.*

6 Analysis

In this section we will analyse and compare the properties of some of our comparison methods. A minimal requirement for any comparison method is that it respects strong dominance, as outlined by Horty [12]. This means that for individual plans Π_1 Π_2 for agent i, if for every possible joint plan Π_X for $Agt \setminus \{i\}$, the outcome for $\Pi_1 \cup \Pi_X$ is better (according to $\preceq$) than the outcome of $\Pi_2 \cup \Pi_X$, then Π_1 should be preferred to Π_2. In other words, if the outcome for Π_1 is always better than Π_2 regardless of the actions of other agents, then Π_1 should be preferred.

We will now show that anticipation-minimising comparison does not respect strong dominance between plans. Consider the following example where we compare plans Π_1 and Π_2 for agent i where $Agt = \{i, j\}$. We suppose that there are exactly three plans (Π'_1, Π'_2 and Π'_3) available to j and that $\overline{\Omega} = \{\omega_1, \omega_2, \omega_3\}$ where (in terms of importance) $\omega_1 > \omega_2 > \omega_3$.

	Π_1	Π_2
Π'_1	$\neg\omega_1$	$\neg\omega_1 \wedge \neg\omega_2$
Π'_2	$\neg\omega_2$	$\neg\omega_1 \wedge \neg\omega_2$
Π'_3	$\neg\omega_3$	$\neg\omega_1 \wedge \neg\omega_2$

In this example Π_2 is strongly dominated by Π_1, since $\neg\omega_1 \wedge \neg\omega_2$ is worse than simply $\neg\omega_1$ or $\neg\omega_2$ and is worse than $\neg\omega_3$ as ω_3 is less important than ω_1 and ω_2. However, Π_2 is anticipation-minimising preferred to Π_1 because Π_2 only anticipates responsibility for $\neg\omega_1$ and $\neg\omega_2$ whereas Π_2 anticipates responsibility for $\neg\omega_1$, $\neg\omega_2$ and $\neg\omega_3$. This shows that the failure of anticipation-minimising comparison is that it fails to consider the values that may be violated or that the agent may be responsible for violating in specific individual outcomes, rather than overall.

We can also show that regret-minimising comparison and consistent-regret-minimising comparison also do not guarantee to respect strong dominance. Consider another two-agent example with plans Π_1, Π_2 for i and Π'_1 for j. We let $\overline{\Omega} = \{\omega_1, \omega_2, \omega_3\}$ and suppose that $\{\neg\omega_2, \omega_3\} \prec \{\omega_2, \neg\omega_3\}$ but $\{\omega_1, \omega_2, \neg\omega_3\} \prec \{\omega_1, \neg\omega_2, \omega_3\}$.

	Π_1	Π_2
Π'_1	$\omega_1 \wedge \omega_2$	$\omega_1 \wedge \omega_3$

In this case Π_2 dominates Π_3, since both forms of regret-minimising comparison do not take into account that the relative ranking of ω_2 and ω_3 is changed by the presence of ω_1. In scenarios where this behaviour does not occur both forms of regret-minimising comparison do respect strong dominance, though more work is needed to determine exactly when this happens.

Using a similar example, we will now show that regret-minimising comparison is different to consistent-regret-minimising comparison.

	Π_1	Π_2	Π_3	Π_4
Π'_1	$\omega_1 \wedge \omega_2$ $\wedge\omega_3 \wedge \omega_4$	$\omega_1 \wedge \omega_2$ $\wedge\omega_3$	$\omega_1 \wedge \omega_2$ $\wedge\omega_3$	$\omega_1 \wedge \omega_4$
Π'_2	ω_1	ω_2	ω_3	$\omega_1 \wedge \omega_4$

In this example we use numerical utilities for values where $\omega_1 = \omega_2 = \omega_3 = 10$ and $\omega_4 = 6$. Now we show regret:

	Π_1	Π_2	Π_3	Π_4
Π'_1	26	14	14	-14
Π'_2	-16	-16	-16	-4

Now we show consistent regret:

	Π_1	Π_2	Π_3	Π_4
Π'_1	0	-6	-6	-20
Π'_2	-6	-6	-6	0

In this example consistent-regret-minimising comparison would pick either Π_1, Π_2 or Π_3 whereas regret-minimising comparison would pick Π_4.

7 Experiment

7.1 Setup and Results

To test these plan comparison methods against each other we set up a simple experiment. The experiment involved two agents with four actions and four shared values. For simplicity, we consider only plans of length 1. To simulate the action theory and initial state we randomly determine which values are satisfied for each possible pair of actions (this is re-randomised for every iteration).

Experiment	Agent	A	B	C	D
1	1	40	10	0	-50
	2	40	10	0	-50
2	1	40	10	0	-50
	2	40	0	10	-50
3	1	10	40	0	-50
	2	10	0	40	-50

Table 1. Value weightings for each agent in each experiment.

We run three different experiments with three different value weightings. The first is identical for both agents, designed to promote maximum cooperation. The second weighting has some disagreement between values while the third has significant disagreement.

The tables below are generated by running 100,000 iterations with each possible pairing of agent types (optimistic, pessimistic, regret-minimising, consistent-regret-minimising and random). The random agent is included both to see how well each type of agent deals with unpredictable or irrational agents, as well as to provide a baseline result. The code for this experiment can be found at: https://pastebin.com/KUbQTwLa.

Agent Type	Opt	Pess	CRM	RM	Rand
Opt	33.2	21.4	29.3	24.3	9
Pess	21.4	14.6	17.8	17.3	12.3
CRM	29.4	18	21.9	20.4	13.7
RM	24.2	17.3	20.2	19.5	13.1
Rand	9.2	12.4	13.6	13.2	-0.1

Table 2. Results of Experiment 1 (agent 1 on left, agent 2 on top)

Agent Type	Opt	Pess	CRM	RM	Rand
Opt	24	21.1	26.9	23.7	9.1
Pess	20.5	15.5	18	17.7	12.4
CRM	26.4	18.1	21.5	20.4	13.6
RM	23.2	17.8	20.2	19.6	13.3
Rand	8.3	11.9	13.3	12.7	0

Table 3. Results of Experiment 2

Agent Type	Opt	Pess	CRM	RM	Rand
Opt	19	18	21	19.3	8.9
Pess	16.4	15.9	17.3	16.7	12.3
CRM	20.1	18.7	20.3	19.3	13.7
RM	18.2	18.3	19.6	19	13.3
Rand	5.5	8.1	8.9	8.3	0

Table 4. Results of Experiment 3

We also ran a brief experiment (table 5) with a joint planning setup that was able to choose a single joint action. This planner would seek to maximise the utility of the worst-off agent, breaking ties randomly. This is provided as an approximate upper bound for single-agent planning.

7.2 Discussion

For experiments 1 and 2, the optimistic agent was the best performing against every non-random agent, suggesting that optimistic agents perform effectively when their values align exactly or closely with those of other agents and the other agents behave at least somewhat rationally. However, consistent-regret-minimising agents performed the best in experiment 3 (except that Opt-CRM performed better than CRM-CRM, which deserves investigation), and consistently performed the best when paired with random agents, suggesting that consistent-regret-minimising agents perform effectively when their values differ significantly from other agents or when other agents act irrationally.

Finally, the joint planning experiment in table 5 shows that even in the best case for individual planning (Opt-Opt in table 2), joint planning is significantly better.

8 Conclusions and Future Work

In this paper we have introduced two methods for plan comparison in a multiagent multi-objective setting using anticipated passive responsibility. This approach uses regret minimisation without requiring numerical utilities to be assigned to values. We have also introduced the notion of complete responsibility to better handle responsibility for multiple values. We have formally analysed our methods and performed an experimental evaluation relative to other options.

There are several directions open for future research. In particular, we can further develop our methods for plan comparison using responsibility, or consider alternatives. For example, instead performing regret minimisation once we could adapt our methods to use iterated regret minimisation as introduced by Halpern and Pass [11].

Experiment	Utility
1	48.5
2	47.3
3	44.4

Table 5. Average utilities achieved with joint planning.

We would also like to perform a more thorough axiomatic analysis of our comparison methods, as well as a more thorough experimental analysis. As the current paper is a work-in-progress the experiment is quite preliminary, as it tests only one-shot plans in a fully randomised domain. It would be interesting to run the experiment in a multi-step planning with more realistic domains and more than two agents. There are also other approaches in the literature that we would like to evaluate against.

Finally, while the model is currently in a purely theoretical state, it would be good to investigate how it could be applied in real-world planning scenarios. This will require both a complexity analysis of the framework (likely to be at least NP-hard, in line with classical planning [20]) and the creation of efficient methods for computing regret-minimising plans.

References

[1] N. Alechina, J. Halpern, and B. Logan, 'Causality, responsibility and blame in team plans', in *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems, (AAMAS)*, (2017).

[2] C. Baier, F. Funke, and R. Majumdar, 'A game-theoretic account of responsibility allocation', in *Proceedings of the 30th International Joint Conference on Artificial Intelligence, (IJCAI)*, (2021).

[3] M. Bienvenu, C. Fritz, and S. McIlraith, 'Planning with qualitative temporal preferences.', in *Proceedings of the 10th Conference on Principles of Knowledge Representation and Reasoning, (KR)*, (2006).

[4] M. Braham and M. van Hees, 'An anatomy of moral responsibility', *Mind*, **121**(483), 601–634, (2012).

[5] H. Chockler and J. Halpern, 'Responsibility and blame: A structural-model approach', *Journal of Artificial Intelligence Research*, **22**, (2004).

[6] G. De Giacomo and M. Vardi, 'Linear temporal logic and linear dynamic logic on finite traces', in *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI)*, (2013).

[7] G. De Giacomo and M. Vardi, 'Synthesis for LTL and LDL on finite traces', in *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*, (2015).

[8] L. Dennis, M. Fisher, M. Slavkovik, and M. Webster, 'Formal verification of ethical choices in autonomous systems', *Robotics and Autonomous Systems*, **77**, 1–14, (2016).

[9] U. Grandi, E. Lorini, T. Parker, and R. Alami, 'Logic-based ethical planning', in *Proceedings of the 21st International Conference of the Italian Association for Artificial Intelligence (AIxIA)*, (2022).

[10] J. Halpern, 'A modification of the halpern-pearl definition of causality', in *Proceedings of the 24th International Joint Conference on Artificial Intelligence, (IJCAI)*, (2015).

[11] J. Halpern and R. Pass, 'Iterated regret minimization: A new solution concept', *Games and Economic Behavior*, **74**(1), 184–207, (2012).

[12] J. Horty, *Agency and deontic logic*, Oxford University Press, 2009.

[13] F. Juma, E. Hsu, and S. McIlraith, 'Preference-based planning via maxsat', in *Proceedings of the 25th Canadian Conference on Artificial Intelligence*, (2012).

[14] P.B. Linhart and R. Radner, 'Minimax-regret strategies for bargaining over several variables', *Journal of Economic Theory*, **48**(1), 152–178, (1989).

[15] E. Lorini, D. Longin, and E. Mayor, 'A logical analysis of responsibility attribution: emotions, individuals and collectives', *Journal of Logic and Computation*, **24**(6), (2014).

[16] E. Lorini and R. Mühlenbernd, 'The long-term benefits of following fairness norms under dynamics of learning and evolution', *Fundamenta Informaticae*, **158**(1-3), 121–148, (2018).

[17] J. Niehans et al., 'Zur preisbildungen bei ungewissen erwartungen', *Swiss Journal of Economics and Statistics (SJES)*, **84**(V), 433–456, (1948).

[18] T. Parker, U. Grandi, and E. Lorini, 'Anticipating responsibility in multiagent planning', in *Proceedings of the 26th European Conference on Artificial Intelligence (ECAI) (2023)*, (accepted for publication).

[19] L. J. Savage, 'The theory of statistical decision', *Journal of the American Statistical Association*, **46**(253), 55–67, (1951).

[20] H. Turner, 'Polynomial-length planning spans the polynomial hierarchy', in *Proceedings of Logics in Artificial Intelligence, European Conference, (JELIA)*, (2002).