Communicating Agent Intentions for Human-Agent Decision Making under Uncertainty

Paper ID: #101

ABSTRACT

Recent advances in visualisation technologies have opened up new possibilities for human-agent communication. For systems where agents use automated planning, visualisation of agent intentions, i.e., agent planned actions, can assist human understanding and decision making (e.g., deciding when human control is required or when it can be delegated to an agent). We are working in an application area, shipbuilding, where branched plans are often essential, due to the typical uncertainty experienced. Our focus is how best to communicate, using visualisation, the key information content of branched plans. It is important that such visualisations communicate the complexity and variety of the possible agent intentions i.e., executions, captured in a branched plan, whilst also connecting to the practitioner's understanding of the problem. Thus we utilise an approach to generate the complete branched plan, to be able to provide a full picture of its complexity, and a mechanism to select a subset of diverse traces that characterise the possible agent intentions. We have developed an interface which uses 3D visualisation to communicate details of these characterising execution traces. Using this interface, we conducted a study evaluating the impact of different modes of presentation on user understanding. Our results support our expectation that visualisation of branched plan characterising execution traces increases user understanding of agent intention and plan execution possibilities.

KEYWORDS

Virtual Agents, Human-Agent Decision Making, Explainable AI Planning

ACM Reference Format:

Paper ID: #101. 2023. Communicating Agent Intentions for Human-Agent Decision Making under Uncertainty. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), London, United Kingdom, May 29 – June 2, 2023*, IFAAMAS, 9 pages.

INTRODUCTION

In systems requiring joint human and AI agent decision making there is a need for human users to understand the intentions of agents, along with the agent rationale for different decisions. This requires the AI agent to be able to explain its reasoning to the human, something which remains a significant challenge [12, 25]. This is reflected in initiatives like DARPA's Explainable AI Program [13] and events such as [11, 31, 32].

For those application domains where AI agents (virtual or robot) use automated planning to control behaviour, the challenge is how to clearly communicate to the human the intentions of the agent which are encapsulated in its generated plans. It has been shown that 3D visualisation and simulation of agent plans can help human user understanding of agent intent [6, 33]. However, generating understandable visualisations is challenging because a plan sequence already implicitly encapsulates the balance made between dependency, constraint and choice, as well as the implied implementation of the plan steps themselves. This challenge is exacerbated when more complex plan structures are required, such as branched plans for partially observable domains.

This is the case for our application – shipbuilding – where branched plans are often essential, due to the typical uncertainty experienced. The advantage of branched plans is that they allow efficient action sequences to be captured for each of the possible worlds that might occur and thus capture a diverse space of alternative solutions. However, the size of the space of possible executions makes it challenging to communicate this to a practitioner, along with the intentions of the agent whose behaviour is underpinned by the plan (e.g., Figure 1, the branched plan used in our evaluation).

Thus, the problem we address in this work is how to communicate key information content of branched agent plans to human decision makers. This comprises a number of sub-problems: (i) how to communicate the complexity and variety of the possible executions captured within a branched plan; (ii) how to select subsets of execution traces that capture the scope of possibilities in branching plan structures to communicate to practitioners - something which is essential, as it is not desirable, or possible, to present all linearisations of a branched plan to a practitioner; and (iii) how to communicate the complexity of the branched plan and the selected execution traces to practitioners, in ways that connect with their understanding of the problem and increases their ability to understand the agents intention (their planned actions).

Our contribution addresses these sub-problems. We have: (i) developed a full branched plan generation mechanism, using [2], to branch on sensor action values and emphasize key action points; (ii) developed a mechanism to select subsets of execution traces capturing the scope of possibilities in branching plan structures; and (iii) demonstrated increased user understanding of agent intention and plan execution possibilities resulting from the way in which diverse trace information is communicated through visualisation.

BACKGROUND

A partially observable planning problem, e.g., [2], can be defined by a tuple, $P = \langle F, A, M, I, G \rangle$, with fluents F, actions A, sensor model M, the initial state clauses I, over F, and goal, G. The clauses of the initial state provide both the known positive and negative literals, as well as constraints over the currently unknown parts of the initial state. An action is defined by its preconditions and effects. An action is applicable if its preconditions are satisfied in the agent's partial state and the application of an action causes its effects to be applied to the agent's current state. Sensing actions are triggered whenever they become applicable and their observations

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom.

update the agent's state. A solution to the problem is a branched plan, π, which has both deterministic action nodes and branching nodes, such that every branch of the tree results in a goal state. The branching nodes are labelled with a proposition and have branches for each of the possible valuations. The application of a branched plan requires traversing the plan tree applying the deterministic actions and selecting the appropriate branches by detecting the value of the proposition in the environment.

Partially observable planning problems, P, with a certain subset of exclusive-or knowledge can be compiled into deterministic classical planning problem, $\mathbb{P}_{DET}$, following the approach in [2]. A key aspect of this encoding is that each sensing action is replaced by a pair of standard actions: one captures the effect of the sensor in the case that its proposition holds in the world and the other for the negative case. As a result the valuation of the sensors becomes a choice for the planner to make. A solution for $\mathbb{P}_{DET}$ is therefore an optimistic and partial solution for $\mathbb{P}$, which we denote, $\pi_{optimistic}$.

Case Study: Shipbuilding

This work is related to the Shipbuilding 4.0 initiative [30] which will see a move to "smart" shipbuilding, with deployment of semi-autonomous agents, such as construction robots, for use in limited access areas.

Automated Planning technologies are able to generate agent plans of actions that maximise the achievement of mission goals whilst at the same time adhere to strict operational safety constraints. Thus we adopt a planning approach and represent the shipbuilding planning problem as a partially observable planning problem, which captures various typical aspects of construction, including preparation, movement of robots and materials and the actual construction. Scenarios feature uncertainty in both the required preparation of the ground to permit construction and movement, and in the integrity of building materials. The model includes actions for movement, block-placing and sensing actions for identifying debris in the environment. Given the typical uncertainty which is a feature of domains such as this, we have found branched plans are often essential.

An important aspect of such domains is the need to keep the human(s) in-the-loop for joint decision making. Thus our focus is the use of visualisation to facilitate collaborative decision making between humans and agents (e.g., via manipulation and interrogation of objects and affordances in a virtual environment). In particular in this work our interest is how best to communicate the key information content of branching agent plans – the agents intention – to human decision makers.

BRANCHED PLAN GENERATION

We are interested in application domains that require branched plans, due to the uncertainty that is experienced. Branched plans allow efficient action sequences to be captured for each of the possible worlds that might be encountered during execution (with respect to the model). As a result branched plans can capture a diverse space of alternative solution sequences. Moreover sensor values and traces are not associated with likelihoods, leading to an interpretation that each of the executions is as likely as any other.

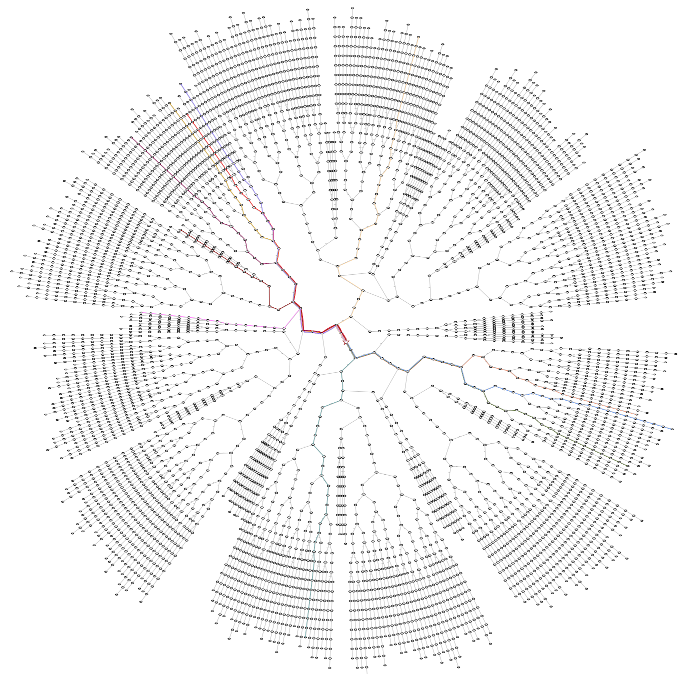

Figure 1: Visualisation of the branched plan for the problem used in the evaluation (see text for further detail).

Given our focus in this work, we require a complete branching tree structure to be generated, so that we can provide a practitioner with a full picture of the complexity of the branched plan. Our approach to this generation uses the K-Replanner [2]: an online approach to partially observable planning which supports efficient plan generation through a compilation to classical planning. The underlying classical planner is used to generate an optimistic plan, $\pi_{optimistic}$ and the K-Replanner approach follows this plan until an inconsistency is discovered, at which point it replans. However, as K-Replanner only explores individual real worlds, we extend it to generate the full plan [18]. At each sensing action encountered in the optimistic plan, the plan is branched for each of the possible values and each of these branches is explored iteratively. Although not explored in this work, we observe that if the entire tree is prohibitively large it would be possible to explore a partial plan, by bounding the number of branching points. The action sequences from this partial plan could be bounded and used as input for our visualisation tools (this is particularly suited to the Interleaved mode. Discussed further in: "Empirical Evaluation").

In addition, the K-Replanner's optimistic plan, $\pi_{optimistic}$, plays an important role in our approach to communicating, through visualisation, the branched plan possibilities to the user. As this optimistic plan is used to drive the planner's strategy during planning it means that our visualisation accurately represents the intention of the system. Whereas with more direct approaches to partially observable planning it might prove challenging to accurately reflect their strategies, K-Replanner's strategy is particularly amenable. For further detail on optimistic plan visualisation see: "User Interface: Presenting 3D Visualisations".

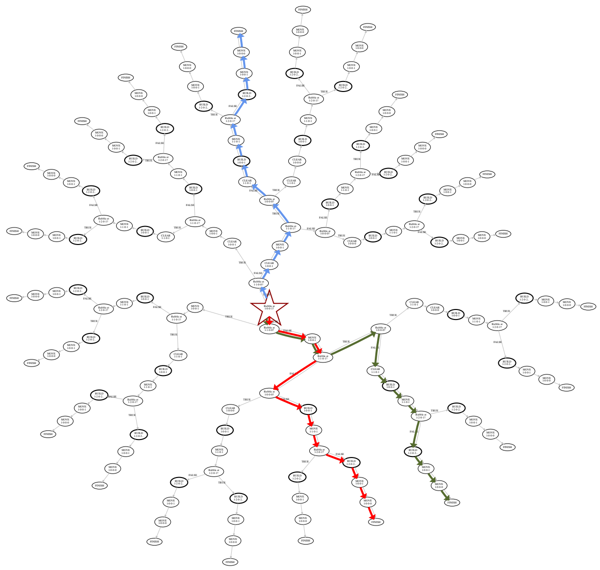

Figure 2: Radial visualisation of branched plan for sample problem: indicating actions and branching points; extended with selected execution traces (highlighted, view in colour).

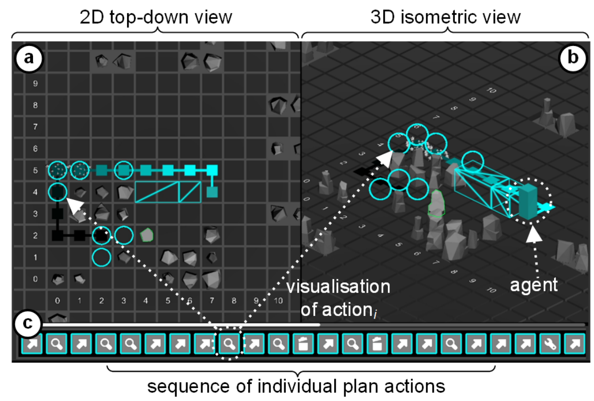

Figure 3: GUI: top-down (a) and isometric (b) simultaneous views of agent action 3D visualisation (simulation); sequence of agent actions (c) displayed as a timeline using icons.

SELECTING CHARACTERISING PLAN TRACES

Our focus is communicating branched plans to human practitioners, however it is not typically possible, or desirable, to present all linearisations. Thus we propose using a small subset of the execution traces to provide examples from across the broad scope of the alternatives captured in the plan. The intention is to provide the practitioner with the intuition of what is captured within the plan without the burden of fully examining every trace. Thus, we aim to select a small number of execution traces that characterise the range of executions represented by the branched plan.

Our selection mechanism uses a dissimilarity measure to estimate the difference between two alternative action sequences. Using this measure on action sequences ignores other features in the trace, and means that the similarity of two traces is determined by the difference in the agent's actions and not differences in such things as sensor readings, which can be irrelevant. Using the dissimilarity measure allows the full set of linearisations to be clustered to identify its key groupings. Clustering provides flexibility, allowing a balance between the number of clusters and loss of detail.

Dissimilarity Measure: To estimate the dissimilarity between two execution sequences we use the Levenshtein distance [20]: the distance between two word sequences which provides a measure of the edit difference between the sequences, while also respecting ordering. In our case we use unique words for each ground action. This measure was used as it is directly applicable to partially observable planning domains and it has been demonstrated that the approach leads to the identification of diverse plans [9].

Clustering Execution Traces: Clustering identifies groups of elements, which are similar (or close) to the elements in their own group, while being dissimilar (or far away from) elements in other groups. We therefore aim to break the space of possibilities into clusters, each representing similar execution traces.

We used the Partitioning Around Medoids (PAM) implementation of the k-medoids method [17]. This approach partitions the data into k clusters, each associated with a representative data point (the *medoid*), considered the most central in the cluster. This approach operates from a dissimilarity matrix, which can be computed by comparing each pair of traces using the dissimilarity measure. The medoids are central members of their respective clusters, and so we use them as the representatives of their clusters. For an appropriate value of k, this set of medoids will identify diverse execution traces, characterising the execution traces in the plan.

Selecting the Number of Clusters: The appropriate choice of k is likely to depend on the application domain and the depth of understanding that is appropriate for the user, which might relate to the seriousness of inappropriate action, e.g., safety and security concerns. We therefore prefer the relatively lightweight k-medoids algorithm, which provides the necessary flexibility. One way to select a reasonable value for k is to calculate the average *silhouette* score for the clusters [27], which evaluates the clusters by averaging the similarity within clusters and dissimilarity between clusters, with respect to the distance measure. This can only be evaluated accurately for at least 2 clusters, so we first test to determine whether more than 1 cluster is appropriate [10]. The optimal average silhouette score indicates a good trade-off between the size of k and the amount of dissimilarity in each cluster. We therefore see this as a suitable default value.

COMMUNICATING BRANCHED PLANS

The size of branched plans makes it challenging to communicate the space of possible executions to practitioners and consequently the intentions of agents whose behaviours are underpinned by such a plan. Our approach to address this exploits two visualisation methods to: (i) communicate the space of alternative execution traces captured in the plan; and (ii) highlight the set of characterising execution traces (selected using the approach discussed earlier).

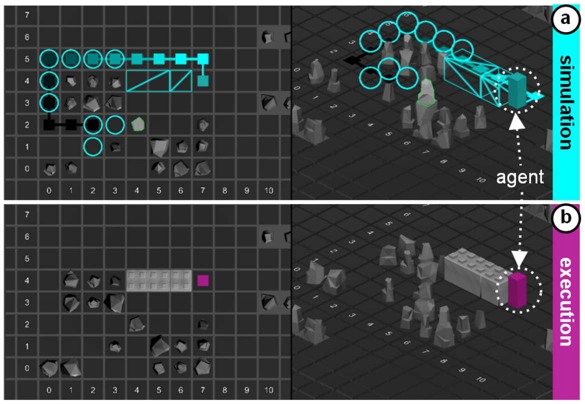

Figure 4: Visualisation example showing difference between: (a) *simulation*; (b) *execution* of selected plan traces.

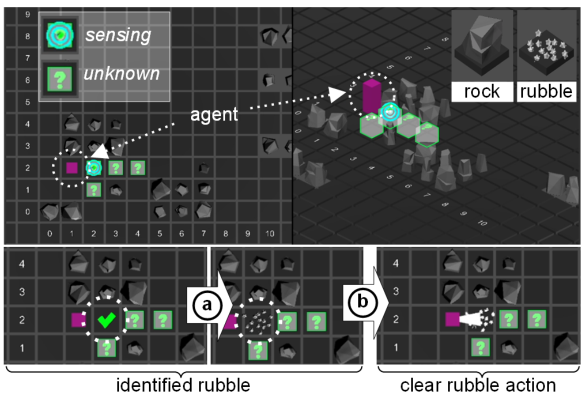

Figure 5: Example *sensing action* with agent in (1,2) sensing in (2,2). Here, the agent identifies rubble (a) and clears it (b).

(i) Communicating Space of Executions

We use a visualisation of the complete space of alternative executions to provide an indication of the number and complexity of the possible alternatives. Although there is no intention that this will lead to an understanding of the actual traces, we extend the visualisation to allow practitioners to explore the possible executions.

Branched plans allow efficient solutions to be captured for each possible concrete state. For our branched plan generator, the branching nodes are associated with sensing actions allowing an appropriate course of action to be selected for each sensor valuation. In scenarios with a small amount of uncertainty the plan might be captured by a concise tree. As the level of (relevant) uncertainty increases, the size of the tree grows and in some cases will be very large. We observe that there is a natural similarity between a branched tree and a classical planning state space. This allows us to exploit recent results in state space visualisation [22, 24] to effectively communicate the size/complexity of the branched plan.

We use a radial layout to visualise the tree. The root of the tree naturally sits in the centre of the visualisation and the branches of the tree expand from the root outwards. The graph has action and sensor action nodes and edges from sensor actions are labelled with the associated sensor value (i.e., *True* or *False*). An example visualisation of the branched plan for a small construction problem is presented in Figure 2 and the branched plan for the problem used in our evaluation is presented in Figure 1. Whereas in [22, 24] the search spaces branch on alternative choices, our plans branch on the valuation of sensor actions. It is therefore appropriate that the distance from the centre reflects the length of the execution, so that clearly longer executions can easily be identified.

We have extended visualisation by emphasising the key action nodes, i.e. those that achieve subgoals (shown with emboldened border in Figures 1 and 2). In order to reduce the complexity of the visualisation, the nodes are annotated with simplified representations of the actions and sensor actions. Further information is provided through tooltips, which provide longer descriptions of the actions and decision points, as well as key state information.

(ii) Highlighting Characterising Execution Traces

We use these selected execution traces to enhance the branched plan radial visualisation and provide meaningful guidance to assist practitioners to navigate the tree and understand its alternatives. Figure 2 illustrates the visualisation of diverse alternative execution traces for our construction problem, with differently coloured lines added to the radial plan visualisation for each of the diverse traces.

Importantly, these characterising execution traces are the ones which are communicated to the practitioner, using 3D visualisation, as discussed in the next section.

COMMUNICATING EXECUTION TRACES

The 3D visualisation of a characterising execution trace provides an effective mechanism for clearly communicating the agent intention captured in that particular trace. We contrast this *3D visualisation* of an individual trace to the *radial visualisation* of a branched plan discussed earlier, as it refers to the use of 3D graphics to provide a visual run through of the sequence of actions, via animations, in a virtual environment. This 3D visualisation can be either: prior to execution, i.e. *simulation*; or the actual *execution* itself.

We have developed a GUI for presenting such 3D visualisations to practitioners and for uses in our evaluation. The interface is implemented using the Unity3D game engine. An example is shown in Figure 3. It provides side-by-side synchronous views of the current action being visualised: a top-down view of the agent acting in the world (left-hand side), and a 3D isometric view (right-hand side). In addition, the GUI has an icon-based representation of the sequence of actions in the execution trace (shown along the bottom).

3D Visualisation of Action Sequences

The generation of 3D visualisations of agent actions relies upon the ability of the interface to convey realistic and semantically meaningful animations within the virtual environment. To maximise practitioner understanding of the execution trace visualisations, we created contextually identifiable discrete sets of 3D animations constituting meaningful representations of agent actions.

An important requirement for the 3D visualisation is to provide visual representations that clearly differentiate between *simulation* and *execution*, and to ensure the animations are consistent and graphically similar. Thus, actions use the same graphical animations

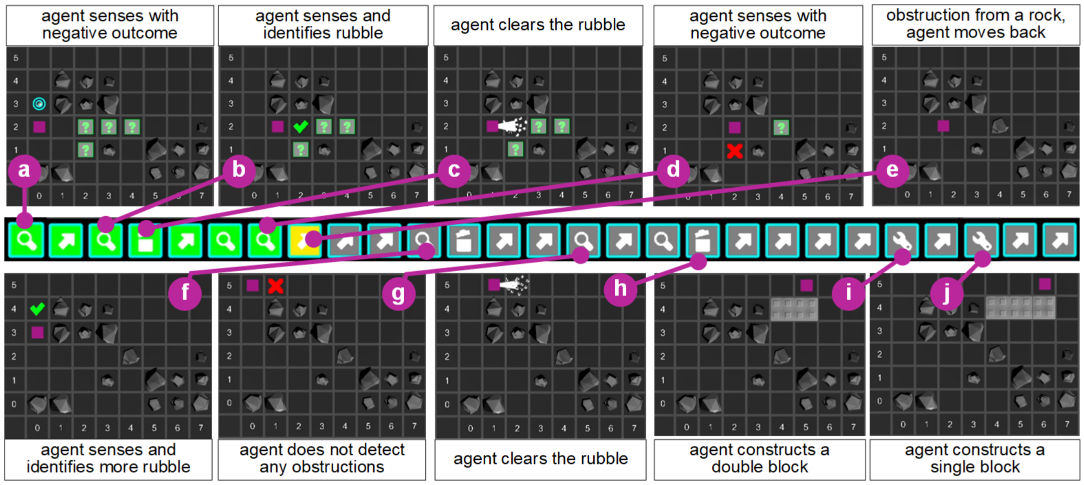

Figure 6: Example 3D visualisation of agent action sequences showing: a number of sensing actions (actions a, b, d, f and g in the figure); agent clearing of rubble (actions c and h); and agent construction of walls (actions i and j).

in simulation and execution but use a different rendering style: simulation actions are rendered in turquoise-coloured wireframe render; whilst execution actions are rendered in grey-coloured fully-shaded render (Figure 4).

A key concept in branched plan generation is the use of sensing actions (see section "Background"), which must be visualised so that practitioners can understand the implication of the valuation process. For instance, in Figure 5, as the agent has identified rubble, the agent's plan from that point is to clear the rubble before proceeding further.

As an illustration of 3D visualization of an agent sequence, as shown to a practitioner, Figure 6 shows a full trace of an *execution* of actions. The figure includes several sensing actions, actions clearing rubble and construction. The icons across the centre are coloured as follows: grey if the action is yet to be visualised; yellow if currently being visualised; or green if fully visualised.

GUI: PRESENTING 3D VISUALISATIONS

An important aspect of eXplainable AI Planning (XAIP) is helping users understand what decisions have been made in a plan, and why. As it is not always feasible to present all branched plan linearisations, our approach is to select a diverse set of execution traces, that characterise the possibilities captured within the full branched plan, and to present these to the practitioner via a series of 3D visualisations within the user interface. Here, we describe the ways in which presentations can be organised to exploit plan structure, as appropriate, to assist practitioners understanding by giving some transparency to the agents' planning strategy.

Communicating the Agent's Intent

Building on [2], our approach to branching plan generation is based on a specific strategy: construct the branched plan starting from a single optimistic plan, $\pi_{optimistic}$, and iteratively branch for alternative sensor values. In order to make this planning strategy transparent to the practitioner, and help them to understand the assumptions made generating the plan, we use this optimistic plan as the "backbone" for presentation of 3D visualisations.

Thus, within the user interface, both the agent's optimistic plan and the selected set of characterising plan traces can be presented to the practitioner – in order to establish common ground between the practitioner and the agent (the agent could follow one of the diverse plans; however, because it is not used to guide plan construction it may lack the rationality and focus of the optimistic plan).

User Interface Modes

The User Interface (GUI) has the following presentation modes:

- *Simulation:* presentation of 3D visualisations of selected traces from the full branched plan prior to execution. In simulation mode, each trace implies a set of assumed sensor valuations for the sensing actions, which are used to generate the appropriate visualisations (differentiated using a different render mode as discussed earlier).
- *Execution:* presentation of 3D visualisation of the actual execution from a concrete starting state through to a *breakpoint*, or the final goal. A breakpoint is a state where the actual sensor valuations differ from the assumed values in the plan trace being executed.

The appropriate mode of presentation of information will often vary between applications. For example, in some scenarios, where the practitioner must have complete understanding of the agent plan before execution starts (e.g., high risk applications), *simulation* can be used to safely explore different possible plan alternatives prior to execution. Whereas in other situations it might be appropriate to interleave the presentation of information, via *simulation*, throughout the *execution* (e.g., slower execution applications).

Based on the different modes of presentation the following observations have guided the design of our empirical evaluation: (i) we observe that breakpoints, where the actual sensor valuations differ from the assumed sensor valuations during execution, provide a useful opportunity to supply further information to the practitioner; and (ii) as an agent executes its plan, the uncertainty in

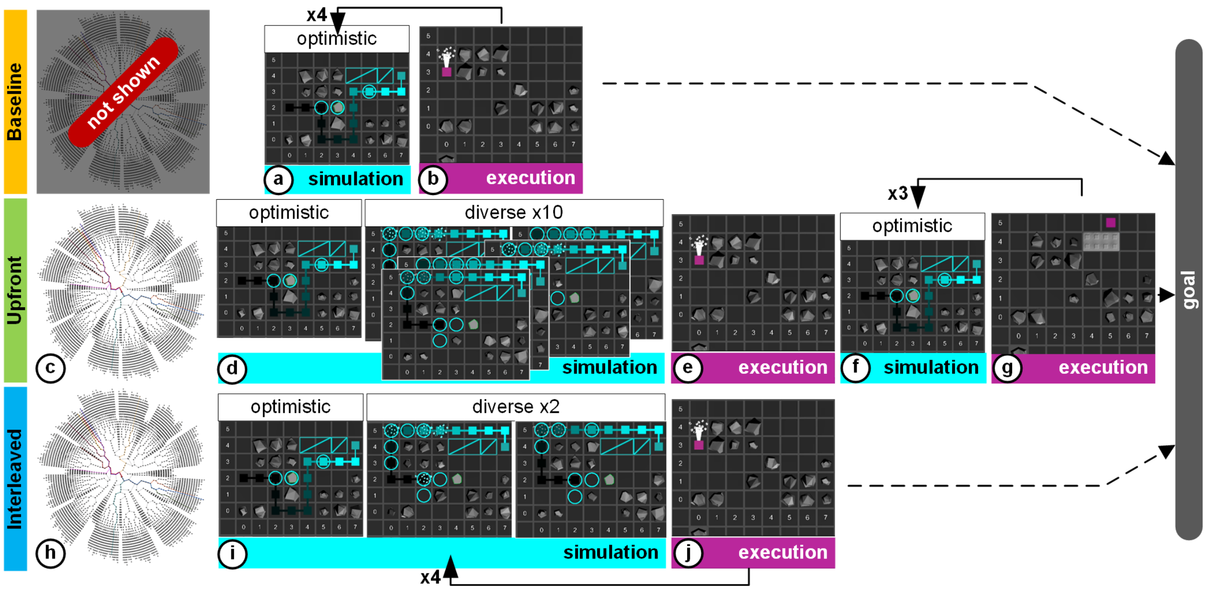

Figure 7: Overview of execution modes for the User Study with visualisation sequences depending on condition: Baseline, Upfront **or** Interleaved. **Full branched plans were shown to** Upfront **or** Interleaved **conditions (c or d respectively), but not** Baseline. **All conditions were shown a combination of** *simulation* **and** *execution* **visualisation sequences, with** k **diverse plans, depending on condition:** Baseline **a-b;** Upfront **with** k=10 **d-g; and** Interleaved **with** k=2 **i-j (see text for details).**

the concrete state decreases, thus isolating a smaller portion of the overall branched plan. Therefore, where appropriate, presenting information at stages during execution, in an Interleaved mode, can provide sets of execution continuations that are more focused towards the unfolding execution. Based on this observation, our working hypothesis is that Interleaved presentation will result in increases in participant understanding of agent intent.

EMPIRICAL EVALUATION

For evaluation we developed a prototype interface featuring: (i) Branched plan generation, based on the K-Replanner, extended to output the full contingency tree; (ii) automated selection of a set of execution traces characterising the scope of possibilities within generated branched plans; (iii) radial visualisation of branched plans, using the approach of [22, 24]; and (iv) a virtual environment for presentation to practitioners via 3D visualisation of selected traces. This was a virtual construction world which is representative of the class of problems we are interested in. It features an agent, performing construction tasks, with uncertainty in the required preparation of the ground to allow for construction and movement.

User Study

We set up a user study to investigate the impact of different modes of information presentation and exploration on practitioner understanding of agent plans, in terms of awareness of agent intended actions in the presence of uncertainty, the agents' overall goal and how confident and prepared they felt to answer the questions.

The study was delivered via an online questionnaire with all participants viewing a system briefing, covering the functioning of the prototype using text, images and videos e.g. visual difference between *simulation* and *execution*, as in Figure 4. A single planning instance was used for the study with 4 breakpoints (points where the actual sensor values differ from those assumed in the optimistic plan). Depending on experimental condition, participants were shown and asked questions relating to radial visualisation of the full branched plan and 3D visualisations (videos) of selected traces in *simulation* or *execution*. We recruited 24 native english speakers who were randomly assigned to one of three conditions which differ with respect to the 3D visualisations shown and their organisation presentation mode, as follows, and shown in Figure 7:

- Baseline: participants were shown *simulation* of the optimistic plan a, followed by optimistic plan *execution*, to the next breakpoint, b, or the goal. After each breakpoint this continued, with *simulation* of the optimistic plan from the current state, followed by *execution* of the optimistic plan looping through to the goal.
- Upfront: participants were shown the radial visualisation of the full branched graph, c. Then *simulation* of each of k diverse traces (rationale for k below) and the optimistic plan, d, followed by *execution* of the optimistic plan through to the first breakpoint, e. Continuation from the first breakpoint, repeatedly loops, showing *simulation* of the optimistic plan from the current state, f, followed by *execution* of the optimistic plan to the next breakpoint or goal, g, repeating through to the goal.
- Interleaved: participants were shown the radial visualisation h. Then *simulation* of each of k diverse traces (rationale for k below) and the optimistic plan, i, followed by *execution* of the optimistic plan, through to the next breakpoint or the goal, j. Continuation from each breakpoint repeatedly loops starting from the optimistic plan from the new current state.

BREAKPOINT QUESTIONS:

Q1: Partipant Informedness:
"... sufficient information about the possible alternative executions for you to anticipate the execution steps ..?"

Q2: Participant Confidence:
"How confident are you about your answer?"

Q3: Awareness of Agent Intention:
"... what do you think the agent will do next?"

POST-EXECUTION QUESTIONS:

Q4: Participant Informedness:
"... How well prepared were you ..?"

Q5: Participant Confidence:
" ... do you know what the agent's goal was? ... How confident are you?

Q6: Awareness of Agent Intention (overall Goal):
"... do you know what the agent's goal was? ... State what the agent's goal was"

Figure 8: User-Study Questions: Q1-Q3 were asked after each breakpoint; Q4-Q6 were asked after the execution completed.

BREAKPOINT QUESTIONS									
	Q1: Participant **Informedness**			Q2: Participant **Confidence**			Q3: % Awareness of **Agent Intention**		
P	U	I	B	U	I	B	U	I	B
1	4.1	2.5	1.6	3.4	2.8	2.1	66.7	62.5	50.0
2	4.0	3.1	1.7	3.9	2.8	1.8	77.8	62.5	60.0
3	3.8	3.1	2.6	3.9	2.3	2.5	33.3	87.5	60.0
4	3.9	2.5	3.0	3.9	2.1	2.8	66.7	87.5	80.0
μ	**3.9** ✓	2.8	2.2	**3.8** ✓	2.6	2.3	61.1	**75.0** ✓	62.5

Figure 9: Responses for Breakpoint questions P1-4 and μ (average), for Upfront (U), Interleaved (I), Baseline (B). Q1 (6-point Likert scale: 0=Not at all, 5=Yes Fully), results confirm expectation that (U) receives more information earlier and thus yields higher informedness rankings (overall avg. 3.9 (U), 2.8 (I), and 2.2 (B)). Q2 (6-point Likert scale: 0=No idea, 5=High): as expected, results show high confidence for (U) with decreasing confidence for (I) and (B) respectively. Q3: results confirm expectation of increased participant awareness of Agent Intention for (I) than (U) and (B). Details in text.

Rationale for values of k: We wanted to ensure a similar number traces shown to Interleaved and Upfront participants, so didn't use the silhouette scores directly as this would mean different numbers at each breakpoint and would introduce too much variation. Instead, the following values of k were used: for Interleaved k=2, the mode of the silhouette scores across the breakpoints; and for Upfront, $k = 10$, to ensure that these participants saw a similar number to the overall number for Interleaved.

User-Study Questions: Participant questions were either *breakpoint* or *post-execution*. These questions are summarised in Figure 8. Breakpoint questions are asked during execution, at each breakpoint (points where the actual sensor values differ from the assumed values in the optimistic plan being executed), and relate to participants feelings of informedness, confidence and their actual awareness of agent intent (the agents next planned action or overall goal). For participant levels of informedness and confidence, ratings were given on a 6-point likert scale with higher values denoting higher levels, and for participant awareness of the agents next action (breakpoint) and overall goal (post-execution) participants were asked to select from multiple choice options.

Working Hypotheses: Regarding Informedness and Confidence, our expectation was that the Upfront participants would report high levels of both, with a strong sense of how the initial stages of *execution* would progress, but would perhaps lose confidence towards the end, if the *execution* deviated from the plans they observed at the start. With regard to participant awareness of agent intention, we expected that the Interleaved participants would most accurately select the agents next action or overall goal, but possibly feel less informed/confident, especially in early stages, or where *execution* deviated from the small collection of plans they observed.

RESULTS

Breakpoint Questions. During presentation of the 3D visualisation of plan execution, participants were asked a series of questions at each breakpoint: the points at which the actual sensor values differ from the assumed values in the optimistic plan. At each breakpoint, participants were asked to rate their informedness and confidence to answer questions and identify the agents intent (select the agents next action). Results are reported in Figure 9.

- *(Q1) Participant Informedness:* As expected the more informed conditions, Interleaved and Upfront, gave higher rankings than Baseline, with respect to receiving sufficient information to anticipate execution steps they observed. Overall ratings are highest for Upfront with an average rating of 3.9 (in comparison to 2.8 and 2.2 for Interleaved and Baseline respectively). Of interest is the increased rankings over breakpoints 1-4 for Baseline. We suspect this results from clarification through exposure to execution visualisations in the absence of diverse trace simulations.
- *(Q2) Participant Confidence*: Ratings, are high for Upfront, avg. 3.8, with decreasing levels of confidence for Interleaved and Baseline, 2.6 and 2.3 respectively. For Upfront, this is to be expected, given the upfront simulation of diverse traces. Interestingly, despite their confidence Upfront participants weren't necessarily correct in their answers to Q3, awareness of agent intention, as discussed below.
- *(Q3) Participant Awareness of Agent Intention (Next Action)* Results show that participant understanding of agent intentions (Next Action), is highest for those in Interleaved, with an overall avg. of 75%. Overall performance on this question is poorer for both Upfront (61.1%) and Baseline (62.5%). For Upfront, analysis of responses over the 4 breakpoints is interesting, with correctness of answers decreasing as the execution progresses. This supports our earlier conjecture that participants might lose confidence towards later stage of execution if this deviated from the plans presented at the start. From this interpretation, the results are consistent with our expectation: that Interleaved will most accurately select agent intention as the simulations shown fit more closely with their current understanding of the problem.

Post-Execution Questions. Following completion of execution all participants were asked to rate their informedness, confidence and

POST-EXECUTION QUESTIONS								
Q4: Participant **Informedness**			Q5: Participant **Confidence**			Q6: % Awareness of **Agent Intention**		
U	I	B	U	I	B	U	I	B
3.89 ✓	3.25	3.1	**4** ✓	2.75	3.1	66.7	**75.0** ✓	60.0

Figure 10: Post-Execution Responses for Q4 (6 point Likert scale: 0=Not Prepared, 5=Well Prepared), and Q5 (6 point Likert scale: 0=Not Confident, 5=Very Confident)) and Q6. Results consistent with breakpoint questions, supporting expectation that: (i) ratings of informedness and confidence would be higher for UPFRONT**; and (ii)** INTERLEAVED **would show increased awareness of agents overall goal intention.**

awareness of the agents overall goal (Q4-Q6 shown in Figure 8). Responses to these questions are shown in Figure 10.

- *(Q4-Q5) Participant Informedness and Confidence:* For UPFRONT this shows consistently high rankings of confidence and informedness, which is consistent with their responses to the breakpoint questions. Whilst the rankings for Q4-Q5 for INTERLEAVED and BASELINE are lower than for UPFRONT, reflecting perhaps that they are given less information initially (INTERLEAVED) or overall (BASELINE). Although this is increased for the post-execution reporting, reflecting the increase in information over the execution.
- *(Q6) Participant Awareness of Agent Intent (overall Goal)* In contrast to the breakpoint questions, users in this condition exhibited similar accuracy, with respect to the agent goal, to INTERLEAVED. This improvement suggests increase in level of informedness from exposure to the execution visualisation itself.
  In contrast to the execution questions, users in this condition exhibited similar accuracy, with respect to the agent goal, to INTERLEAVED. This improvement suggests increase in level of informedness from exposure to the execution visualisation itself.

Overall, results of the study are encouraging and support our expectations. For Informedness and Confidence, the UPFRONT participants reported higher levels of both, with a strong sense of how the initial stages of *execution* would progress, but slight decreases towards the end, as *execution* deviated from the plans they observed at the start. With regard to participant awareness of agent intention, the results supported our expectation that the INTERLEAVED participants would most accurately select the agents next action or overall goal, but possibly feel less informed/confident, especially in early stages, or where *execution* deviated from the small number of plans they observed.

RELATED WORK

Explainable AI Planning (XAIP) [12], is an area of growing importance and focus in planning, which is motivated by the need for trust, interaction and transparency between users and AI controlled agents [14, 15, 21, 28]. Communicating the intentions of the agent plays an important role in XAIP. The form of visualisation for communicating intention vary from annotations indicating objects that are involved in the agent's plan [7], the indication of intended movements of the robot [4] and a visualisation of the agent's internal decision making [5]. A common approach is to use specialised visualisations to present the intentions of an agent to a system user, e.g., through projection [4, 19] or augmented reality [7, 33]. In [4] it is demonstrated that projecting the robots intentions improves the user rating of the robot. In [7] they build a domain specific visualisation, using augmented reality to project a robot's intentions, which also allows manipulation.

In [3], glyphs are organised to represent actions on a timeline and uses domain dependent information to create action abstractions. Our approach is compatible with plan abstraction visualisation techniques (within sequences of non-sensing actions). The approach in [1] defines a measure of importance, which is appropriate for MDPs. Importance is based on comparisons of the expected rewards for each action, which there is not an equivalent generated during planning for branching plans (except solvability, which was not an interesting feature of the problems that we considered here). In [8], it's assumed that the user model is available so specific situations where the user's understanding of the current sequence might fail can be isolated. Thus, they consider explanations as model corrections. We do not assume a user model, rather, that a common ground can be established by exploiting sequential plan visualisations.

The approach for visualising the complete branched plan that we use here, is related to other domain independent visualisations [22–24]. In [23] they present a plan visualisation, which exploits a visual metaphor in order to communicate abstract planning concepts, such as action preconditions. Our branched plan visualisation was inspired by the state space visualisation of [22].

Selecting/generating sets of diverse plans has been investigated [16], and various applications have been identified, including risk assessment [29] and user preferences [26]. [9] adopt two diversity measures, including the measure we use, and demonstrate that both approaches lead to identifying diverse plans. They also demonstrate that a domain specific diversity measure was particularly effective. Our approach is compatible with any approach that can return a set of diverse plans (especially those able to vary set size).

DISCUSSION AND CONCLUSION

We have considered the problem of visualising a contingent plan and providing the user with visualisation of the intended plan and key information about the contingency tree. The aim is to provide access to the potential alternatives captured in the contingency tree, so users can better understand and assess time required and risk implied by the plan. The results of a user study assessing our approach are promising. They indicate that users can gain an awareness of agent intentions and the scope of alternative possibilities through exposure to selected traces which characterise the branched plan space.

In future work, the different modes of presentation and use of the silhouette score will be further explored, along with user preferences for different modes of presentation and the interface itself, especially in the context of in-situ visualisation of actions through mixed-reality, which will further support interactive exploration of the agents' plan traces.

REFERENCES

[1] Dan Amir and Ofra Amir. 2018. HIGHLIGHTS: Summarizing Agent Behaviors to People. In *17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.

[2] Blai Bonet and Hector Geffner. 2011. Planning under partial observability by classical replanning: Theory and experiments. In *International Joint Conference on Artificial Intelligence*.

[3] Gabriel Dias Cantareira, Gerard Canal, and Rita Borgo. 2022. Actor-Focused Interactive Visualization for AI Planning. In *Proceedings of the International Conference on Automated Planning and Scheduling*.

[4] Ravi Teja Chadalavada, Henrik Andreasson, Robert Krug, and Achim J Lilienthal. 2015. That's on my mind! robot to human intention communication through on-board projection on shared floor space. In *2015 European Conference on Mobile Robots (ECMR)*. IEEE, 1–6.

[5] Tathagata Chakraborti, Kshitij P Fadnis, Kartik Talamadupula, Mishal Dholakia, Biplav Srivastava, Jeffrey O Kephart, and Rachel KE Bellamy. 2017. Visualizations for an explainable planning agent. *arXiv preprint arXiv:1709.04517* (2017).

[6] Tathagata Chakraborti, Sarath Sreedharan, Anagha Kulkarni, and Subbarao Kambhampati. 2018. Projection-Aware Task Planning and Execution for Human-in-the-Loop Operation of Robots in a Mixed-Reality Workspace. In *IEEE International Conference on Intelligent Robots and SystemsIROS*.

[7] Tathagata Chakraborti, Sarath Sreedharan, Anagha Kulkarni, and Subbarao Kambhampati. 2018. Projection-aware task planning and execution for human-in-the-loop operation of robots in a mixed-reality workspace. In *2018 IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*.

[8] Tathagata Chakraborti, Sarath Sreedharan, Yu Zhang, and Subbarao Kambhampati. 2017. Plan explanations as model reconciliation: moving beyond explanation as soliloquy. In *26th International Joint Conference on Artificial Intelligence*. 156–163.

[9] Alexandra Coman and Hector Munoz-Avila. 2011. Generating Diverse Plans Using Quantitative and Qualitative Plan Distance Metrics.. In *AAAI*. Citeseer, 946–951.

[10] Richard O Duda, Peter E Hart, et al. 1973. *Pattern classification and scene analysis*. Vol. 3. Wiley New York.

[11] EXTRAAMAS. 2020. 2nd Int. Workshop on EXplainable, TRansparent Autonomous Agents and Multi-Agent Systems. https://extraamas.ehealth.hevs.ch/topics.html.

[12] Maria Fox, Derek Long, and Daniele Magazzeni. 2017. Explainable planning. *arXiv preprint arXiv:1709.10256* (2017).

[13] David Gunning and David Aha. 2019. DARPA's Explainable Artificial Intelligence (XAI) Program. *AI Magazine* 40, 2 (2019).

[14] Jörg Hoffmann and Daniele Magazzeni. 2019. Explainable AI Planning (XAIP): Overview and the Case of Contrastive Explanation. In *Reasoning Web. Explainable Artificial Intelligence*. Springer, 277–282.

[15] Subbarao Kambhampati. 2019. Synthesizing explainable behavior for human-AI collaboration. In *Proc. of the 18th Int. Conference on Autonomous Agents and MultiAgent Systems*.

[16] Michael Katz and Shirin Sohrabi. 2020. Reshaping Diverse Planning.. In *AAAI*. 9892–9899.

[17] Leonard Kaufmann and Peter Rousseeuw. 1987. Clustering by Means of Medoids. *Journal of ML Research* (1987).

[18] Radimir Komarnitsky and Guy Shani. 2016. Computing Contingent Plans Using Online Replanning. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence* (Phoenix, Arizona) *(AAAI'16)*. AAAI Press, 3159–3165.

[19] Florian Leutert, Christian Herrmann, and Klaus Schilling. 2013. A spatial augmented reality system for intuitive display of robotic data. In *2013 8th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 179–180.

[20] V I Levenshtein. 1966. Binary codes capable of correcting deletions, insertions and reversals. *Cybernetics and Control Theory* 10 (1966), 707–710.

[21] Alan Lindsay. 2019. Towards Exploiting Generic Problem Structures in Explanations for Automated Planning. In *10th International Conference on Knowledge Capture*. 235–238.

[22] Maurício Cecílio Magnaguagno, Ramon Frage Pereira, Martin Duarte Móre, and Felipe Rech Meneguzzi. 2017. Web planner: A tool to develop classical planning domains and visualize heuristic state-space search. In *2017 Workshop on User Interfaces and Scheduling and Planning (UISP@ ICAPS), 2017, Estados Unidos*.

[23] Mauricio Cecilio Magnaguagno, Ramon Fraga Pereira, and Felipe Meneguzzi. 2016. DOVETAIL–An Abstraction for Classical Planning Using a Visual Metaphor. In *The Twenty-Ninth International Flairs Conference*.

[24] Maurício C Magnaguagno, Ramon Fraga Pereira, Martin D Móre, and Felipe Meneguzzi. 2020. Web Planner: A Tool to Develop, Visualize, and Test Classical Planning Domains. In *Knowledge Engineering Tools and Techniques for AI Planning*. Springer, 209–227.

[25] Tim Miller. 2019. Explanation in artificial intelligence: Insights from the social sciences. *Artificial Intelligence* 267 (2019), 1–38.

[26] Tuan Anh Nguyen, Minh Do, Alfonso Emilio Gerevini, Ivan Serina, Biplav Srivastava, and Subbarao Kambhampati. 2012. Generating diverse plans to handle unknown and partially known user preferences. *Artificial Intelligence* 190 (2012), 1–31.

[27] Peter J Rousseeuw. 1987. Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *Journal of computational and applied mathematics* 20 (1987), 53–65.

[28] Shirin Sohrabi, Jorge A Baier, and Sheila A McIlraith. 2011. Preferred explanations: Theory and generation via planning. In *25th AAAI Conference on Artificial Intelligence*.

[29] Shirin Sohrabi, Anton V Riabov, Michael Katz, and Octavian Udrea. 2018. An AI Planning Solution to Scenario Generation for Enterprise Risk Management.. In *AAAI*. 160–167.

[30] Venesa Stanić, M. Hadjina, N. Fafandjel, and T. Matulja. 2018. Toward shipbuilding 4.0 - an industry 4.0 changing the face of the shipbuilding industry. *Brodogradnja* 69 (2018), 111–128.

[31] XAI. 2020. IJCAI-PRICAI Workshop on Explainable Artificial Intelligence (XAI). https://sites.google.com/view/xai2020/home.

[32] XAIP. 2020. 3rd International Workshop on Explainable AI Planning (XAIP). http://xaip.mybluemix.net/#.

[33] M Zolotas and Y Demiris. 2019. Towards explainable shared control using augmented reality. In *IEEE International Conference on Intelligent Robots and Systems, IROS*.