

Balancing Explicability and Explanations

Emergent Behaviors in Human-Aware Planning

Tathagata Chakraborti, Sarath Sreedharan, Subbarao Kambhampati*

School of Computing, Informatics, and Decision Systems Engineering

Arizona State University, Tempe, AZ 85281 USA

tchakra2,ssreedh3,rao@asu.edu

ABSTRACT

Human aware planning requires an agent to be aware of the intentions, capabilities and mental model of the human in the loop during its decision process. This can involve generating plans that are explicable to a human observer as well as the ability to provide explanations when such plans cannot be generated. In this paper, we bring these two concepts together and show how an agent can account for both these needs and achieve a trade-off during the plan generation process itself by means of a model-space search method MEGA. This in effect provides a comprehensive perspective of what it means for a decision-making agent to be “human-aware” by bringing together existing principles of planning under the umbrella of a single plan generation process. We situate our discussion in the context of recent work on explicable planning and explanation generation, and illustrate these concepts in modified versions of two well-known planning domains, as well as in a demonstration of a robot involved in a typical search and reconnaissance task with an external supervisor. Human factor studies in the latter highlight the usefulness of the proposed approaches.

CCS CONCEPTS

• **Computing methodologies** → **Artificial intelligence; Planning and scheduling**; Cognitive robotics; • **Human-centered computing**; • **Computer systems organization** → *Robotics*;

KEYWORDS

Human-Aware Planning, Explicable Planning, Plan Explanations, Explanation as Model Reconciliation, Minimal Explanations.

1 INTRODUCTION

It is often useful for a planning agent while *interacting* with a human in the loop to use, in the process of its deliberation, not only the model \mathcal{M}^R of the task it has on its own, but also the model \mathcal{M}_h^R that the human thinks it has (refer to Figure 1a). This mental model of the human [9] is in addition to the physical model of the human. This is, in essence, the fundamental thesis of the recent works on plan explanations [10] and explicable planning [35], summarized under the umbrella of *multi-model planning*, and is in addition to the originally studied *human-aware planning* (HAP) problems where actions of the human (and hence the actual *human model* and the robot’s belief of it) are also involved in the planning process. The need for explicable planning or plan explanations in fact occur when these two models – \mathcal{M}^R and \mathcal{M}_h^R – diverge. This means that the optimal plans in the respective models – $\pi_{\mathcal{M}^R}^*$ and $\pi_{\mathcal{M}_h^R}^*$ – may

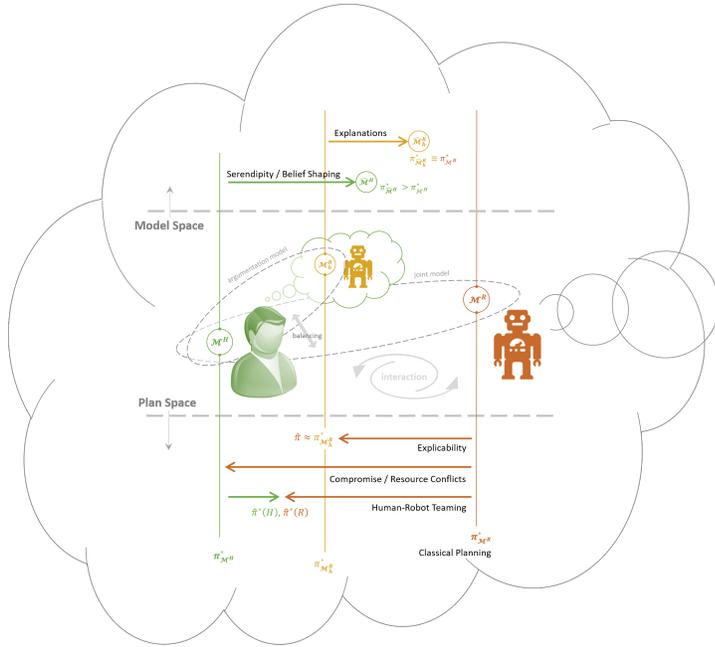
not be the same and hence optimal behavior of the robot in its own model is inexplicable to the human in the loop. In the explicable planning process, the robot produces a plan $\hat{\pi}$ that is closer to the human’s expected plan, i.e. $\hat{\pi} \approx \pi_{\mathcal{M}_h^R}^*$. In the explanation process, the robot instead attempts to update the human’s mental model to an intermediate model $\widehat{\mathcal{M}}_h^R$ in which the robot’s original plan is *equivalent* (with respect to a metric such as cost or similarity) to the optimal and hence explicable, i.e. $\pi_{\mathcal{M}^R}^* \equiv \pi_{\widehat{\mathcal{M}}_h^R}^*$.

Until now, these two processes of plan explanations and explicability have remained separate in so far as their role in an agent’s deliberative process is considered - i.e. a planner either generates an explicable plan to the best of its ability or it produces explanations of its plans where they required. However, there may be situations where a combination of both provide a much better course of action – if the expected human plan is too costly in the planner’s model (e.g. the human might not be aware of some safety constraints) or the cost of communication overhead for explanations is too high (e.g. limited communication bandwidth). Consider, for example, a human working with a robot that has just received a software update allowing it to perform new complex maneuvers. Instead of directly trying to conceive all sorts of new interactions right away that might end up spooking the user, the robot could instead reveal only certain parts of the new model while still using its older model (even though suboptimal) for the rest of the interactions so as to slowly reconcile the drifted model of the user. This is the focus of the current paper where we try to attain the sweet spot between plan explanations and explicability.

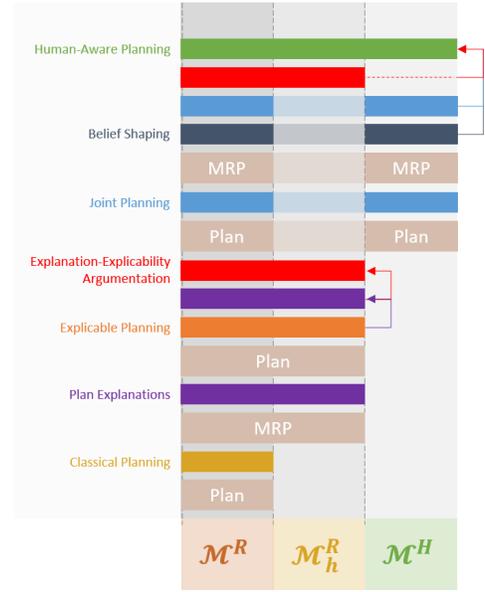
1.1 Related Work

As AI agents become pervasive in our daily lives, the need for such agents to be cognizant of the beliefs and expectations of the humans in their environment has been well documented [21]. From the perspective of task planning, depending on the extent of involvement of the human in the life cycle of a plan, work in this direction has ranged on a spectrum of “human-aware planning” [2, 3, 8, 11–13, 22, 33] where a robot passively tries to account for the plans of humans cohabiting its workspace, to “explicable planning” [14, 23, 34, 35] where a robot generates plans that are explicable or predictable to a human observer, to “plan explanations” [10, 17, 24] where the agent uses explanations to bring the human (who may have a different understanding of the agent’s abilities) on to the same page, to “human-in-the-loop planning” [1, 4, 16, 25, 30] in general where humans and planners are participating in the plan generation and/or execution process together.

*The first two authors contributed equally.



(a) The evolving scope of Human-Aware Planning (HAP)



(b) A Subsumption Architecture for HAP

Figure 1: The expanding scope of *human-aware planning* (HAP) acknowledging the need to account for the mental model of the human in the loop in the deliberative process of an autonomous agent. The planner can, for example, choose to bring the human’s model closer to the ground truth using explanations via a process called *model reconciliation* (MRP) so that an otherwise inexplicable plan makes sense in the human’s updated model or it can compute explicable plans which are closer to the human’s expectation. These capabilities can be stacked to realize more and more complex behavior – in this paper we will concentrate on the explicability versus explanation trade-off as a form of argumentation during human-aware planning.

1.1.1 *The Evolving Scope of Human-Aware Planning.* The ongoing efforts to make planning more “human-aware” is illustrated in Figure 1a – initial work on this topic had largely focused on incorporating an agent’s understanding of the human model M^H into its decision making process. Since then the importance of considering the human’s understanding M_h^R of the agent’s actual model M^R in the planning process has also been acknowledged, sometimes implicitly [3] and later explicitly [10, 35]. These considerations engender interesting behaviors both in the space of plans and models. For example, in the model space, the modifications to the human mental model M_h^R is used for explanations in [10] while reasoning over the actual model M^H can reveal interesting behavior by affecting the belief state of the human, such as in planning for serendipity [8]. In the plan space, a human-aware agent can use M^H and M_h^R to compute joint plans for teamwork [32] or generate behavior that conforms to the human’s preferences [11, 22] and expectations [23, 34, 35]. From the point of view of the planner, this is, in a sense, an asymmetric epistemic setting with single-level nested beliefs over its models. Indeed, existing literature on epistemic reasoning [19, 27, 28] can also provide interesting insights in the planning process of an agent in these settings.

1.1.2 *A Subsumption Architecture for HAP.* These different forms of behavior can be *composed* to form more and more sophisticated

forms of human-aware behavior. This hierarchical composition of behaviors can be viewed in the form of a subsumption architecture for human-aware planning, similar in motivation to [7]. This is illustrated in Figure 1b. The basic reasoning engines are the Plan and MRP (Model Reconciliation) modules. The former accepts model(s) of planning problems and produces a plan, the latter accepts the same and produces a new model. The former operates in plan space and gives rise to classical, joint and explicable planning depending on the models it is operating on, while the latter operates in model space to produce explanations and belief shaping behavior. These are then composed to form argumentation modules for trading of explanations and explicability (which is the topic of the current paper) and human-aware planning in general.

1.1.3 *The Explicability-Explanation Trade-off.* From the perspective of design of autonomy, this trade-off has two important implications – (1) the agent can now not only explain but also *plan* in the multi-model setting with the trade-off between compromise on its optimality and possible explanations in mind; and (2) the argumentation process is known to be a crucial function of the reasoning capabilities of humans [26], and now by extension of autonomous agents as well, as a result of algorithms we develop here to incorporate the explanation generation process into the agent’s decision making process itself. General argumentation frameworks

for resolving disputes over plans have indeed been explored before [6, 15]. Our work can be seen as the specific case where the argumentation process is over a set of constraints that prove the correctness and quality of plans by considering the cost of the argument specifically as it relate to the trade-off in plan quality and the cost of explaining that plan. This is the first of its kind algorithm that can achieve this in the scope of plan explanations and explicable in presence of model differences with the human.

HUMAN-AWARE PLANNING REVISITED

The problem formulation closely follows that introduced in [10], reproduced here for clarity of methods built on the same definitions.

A *Classical Planning Problem*. is a tuple $\mathcal{M} = \langle \mathcal{D}, \mathcal{I}, \mathcal{G} \rangle^1$ with domain $\mathcal{D} = \langle F, A \rangle$ - where F is a set of fluents that define a state $s \subseteq F$, and A is a set of actions - and initial and goal states $\mathcal{I}, \mathcal{G} \subseteq F$. Action $a \in A$ is a tuple $\langle c_a, pre(a), eff^+(a) \rangle$ where c_a is the cost, and $pre(a), eff^+(a) \subseteq F$ are the preconditions and add/delete effects, i.e. $\delta_{\mathcal{M}}(s, a) \models \perp$ if $s \not\models pre(a)$; else $\delta_{\mathcal{M}}(s, a) \models s \cup eff^+(a) \setminus eff^-(a)$ where $\delta_{\mathcal{M}}(\cdot)$ is the transition function. The cumulative transition function is $\delta_{\mathcal{M}}(s, \langle a_1, a_2, \dots, a_n \rangle) = \delta_{\mathcal{M}}(\delta_{\mathcal{M}}(s, a_1), \langle a_2, \dots, a_n \rangle)$.

The solution to the planning problem is a sequence of actions or a (satisficing) *plan* $\pi = \langle a_1, a_2, \dots, a_n \rangle$ such that $\delta_{\mathcal{M}}(\mathcal{I}, \pi) \models \mathcal{G}$. The cost of a plan π is $C(\pi, \mathcal{M}) = \sum_{a \in \pi} c_a$ if $\delta_{\mathcal{M}}(\mathcal{I}, \pi) \models \mathcal{G}$; ∞ otherwise. The cheapest plan $\pi^* = \arg \min_{\pi} C(\pi, \mathcal{M})$ is the (cost) optimal plan. We refer to this cost as \mathcal{M} as $C_{\mathcal{M}}^*$.

A *Human-Aware Planning (HAP) Problem*. is given by the tuple $\Psi = \langle \mathcal{M}^R, \mathcal{M}^H, \mathcal{M}_h^R \rangle$ where $\mathcal{M}^R = \langle D^R, \mathcal{I}^R, \mathcal{G}^R \rangle$ is the planner’s model of a planning problem, while $\mathcal{M}^H = \langle D^H, \mathcal{I}^H, \mathcal{G}^H \rangle$ and $\mathcal{M}_h^R = \langle D_h^R, \mathcal{I}_h^R, \mathcal{G}_h^R \rangle$ are respectively the planner’s estimate of the human’s model and the human’s understanding of its own model.

The solution to the human-aware planning problem is a joint plan [8] $\pi = \langle a_1, a_2, \dots, a_n \rangle$; $a_i \in \{D^R \cup D_h^R\}$ such that $\delta_{\Psi}(\mathcal{I}^R \cup \mathcal{I}_h^R, \pi) \models \mathcal{G}^R \cup \mathcal{G}_h^R$. The robot’s component in the plan is referred to as $\pi(R) = \langle a_i \mid a_i \in \pi \wedge D^R \rangle$. For the purposes of this paper, we ignore the robot’s belief of the human model, i.e. $\mathcal{M}_h^R = \langle \{\}, \{\}, \{\} \rangle$ - in effect, making the human an observer only or a passive consumer of the plan - and focus instead on the challenges involves in planning with the human’s model of the planner. Planning with the human model has indeed been studied extensively in the literature, as noted above, and this assumption does not change in any way the relevance of the work here. Specifically, the following concepts are built on top of the joint planning problem - e.g. an explicable plan in this paper would, in the general sense, correspond to the robot’s component in the joint plan being explicable to the human in the loop. Thus, for the purposes of this paper, we have $\pi(R) \equiv \pi$; without loss of generality, we focus on the simplified setting with only the model of the planner and the human’s approximation of it.

Explicable Planning

In “explicable planning” a solution to the human-aware planning problem is a plan π such that (1) it is executable (but may no longer

be optimal) in the robot’s model but is (2) “closer” to the expected plan in the human’s model, *given a particular planning problem* -

- (1) $\delta_{\mathcal{M}^R}(\mathcal{I}^R, \pi) \models \mathcal{G}^R$; and
- (2) $C(\pi, \mathcal{M}_h^R) \approx C_{\mathcal{M}_h^R}^*$.

“Closeness” or distance to the expected plan is modeled here in terms of cost optimality, but in general this can be any preference metric like plan similarity. In existing literature [23, 34, 35] this has been usually achieved by modifying the search process so that the heuristic that guides the search is driven by the robot’s knowledge of the human’s mental model. Such a heuristic can be either derived directly [23] from the human’s model (if it is known) or learned [35] through interactions in the form of affinity functions between plans and their purported goals.

Plan Explanations

The other approach would be to (1) compute optimal plans in the planner’s model as usual, but also provide an explanation (2) in the form of a model update to the human so that (3) the same plan is now also optimal in the human’s updated model of the problem. Thus, a solution involves a plan π and an explanation \mathcal{E} such that -

- (1) $C(\pi, \mathcal{M}^R) = C_{\mathcal{M}^R}^*$;
- (2) $\widehat{\mathcal{M}}_h^R \leftarrow \mathcal{M}_h^R + \mathcal{E}$; and
- (3) $C(\pi, \widehat{\mathcal{M}}_h^R) = C_{\widehat{\mathcal{M}}_h^R}^*$.

Note that here a model update, as indicated by the + operator may include a correction to the belief (goals or state information) as well as information pertaining to the action model itself. In [10] the authors explored various ways of generating such solutions - including methods to minimize the lengths of the explanations given as a result. However, this was done in an after-the-fact fashion, i.e. the optimal plan was already generated and it was just a matter of finding the best explanation for it. This not only ignores the possibility of finding better plans (that are equally optimal) with smaller explanations, but also avenues of compromise in a manner we discussed previously whereby the planner sacrifices its optimality to further reduce overhead in the explanation process.

MEGA

We bring the notions of explicability and explanations together in a novel planning technique MEGA (Multi-model Explanation Generation Algorithm) that trades off the relative cost of explicability to providing explanations during the plan generation process itself². The output of MEGA is a plan π and an explanation \mathcal{E} such that (1) π is executable in the robot’s model, and with the explanation (2) in the form of model updates it is (3) optimal in the updated human model while (4) the cost (length) of the explanations, and the cost of deviation from optimality in its own model to be explicable to the human, is traded off according to a constant α -

- (1) $\delta_{\mathcal{M}^R}(\mathcal{I}^R, \pi) \models \mathcal{G}^R$;
- (2) $\widehat{\mathcal{M}}_h^R \leftarrow \mathcal{M}_h^R + \mathcal{E}$;
- (3) $C(\pi, \widehat{\mathcal{M}}_h^R) = C_{\widehat{\mathcal{M}}_h^R}^*$; and

¹Note that the “model of a planning problem” includes the action model *as well as the initial and goal states of an agent*.

²As in [10] we assume that the human mental model is known and has the same computation power ([10] also suggests possible ways to address these issues, the same discussions apply here as well). Also refer to the discussion on model learning later.

$$(4) \pi = \arg \min_{\pi} \{ |\mathcal{E}| + \alpha \times |C(\pi, \mathcal{M}^R) - C_{\mathcal{M}^R}^*| \}.$$

Clearly, with higher values of α the planner will produce plans that require more explanation, with lower α it will generate more explicable plans. Thus, with the help of this hyperparameter α , an autonomous agent can deliberate over the trade-off in the costs it incurs in being explicable to the human (second minimizing term in (4)) versus explaining its decisions (first minimizing term in (4)). Note that this trade-off is irrespective of the cognitive burden of those decisions on the human in the loop. For example, for a robot in a collapsed building during a search and rescue task, or the rover on Mars, may have limited bandwidth for communication and hence prefer to be explicable instead.

We employ a *model space* A^* search (Algorithm 1) to compute the expected plan and explanations for a given value of α . Similar to [10] we define a state representation over planning problems with a mapping function $\Gamma : {}^a\mathcal{M} \mapsto \mathcal{F}$ which represents a planning problem by transforming every condition in it into a predicate. The set Λ of actions contains unit model change actions $\lambda : \mathcal{F} \rightarrow \mathcal{F}$ which make a single change to a domain at a time.

We start by initializing the min node tuple (N) with the human mental model and an empty explanation. For each new possible model we come across during our model space search, we test if the objective value of the new node is smaller than the current min node. We stop the search once we identify a model that is capable of producing a plan that is also optimal in the robot's own model. This is different from the stopping condition used by the original MCE-search³ in [10], where the authors are just trying to identify the first node where the given plan is optimal.

Property 1. MEGA yields the smallest possible explanation for a given human-aware planning problem.

This means that with a high enough α (see below) the algorithm is guaranteed to compute the best possible plan for the planner as well as the smallest explanation associated with it. This is by construction of the search process itself, i.e. the search only terminates after the all the nodes that allow $C(\pi, \widehat{\mathcal{M}}_h^R) = C_{\widehat{\mathcal{M}}_h^R}^*$ have been exhausted.

This is beyond what is offered by the model reconciliation search in [10], which only computes the smallest explanation *given* as a plan that is optimal in the planner's model.

Property 2. $\alpha = |\mathcal{M}^R \Delta \mathcal{M}_h^R|$ yields the most optimal plan in the planner's model along with the minimal explanation possible given a human-aware planning problem.

This is easy to see, since with $\forall \mathcal{E}, |\mathcal{E}| \leq |\mathcal{M}^R \Delta \mathcal{M}_h^R|$, the latter being the total model difference, the penalty for departure from explicable plans is high enough that the planner must choose from possible explanations only (note that the explicability penalty is always positive until the search hits the nodes with $C(\pi, \widehat{\mathcal{M}}_h^R) = C_{\widehat{\mathcal{M}}_h^R}^*$, at which point onwards the penalty is exactly zero). In general this works for any $\alpha \geq |MCE|$ but since an MCE will only be known retrospectively after the search is complete, the above condition suffices since the entire model difference is known up front and is the largest possible explanation in the worst case.

³An MCE or a *minimally complete explanation* is the shortest model update so that a given plan optimal in the robot model is also optimal in the updated human model.

Algorithm 1 MEGA

```

1: procedure MEGA-SEARCH
2: Input: HAP  $\Psi = \langle \mathcal{M}^R, \mathcal{M}_h^R \rangle, \alpha$ 
3: Output: Plan  $\pi$  and Explanation  $\mathcal{E}$ 
4: Procedure:
5:   fringe  $\leftarrow$  Priority_Queue()
6:   c_list  $\leftarrow$  {} ▷ Closed list
7:    $N_{min} \leftarrow \langle \mathcal{M}_h^R, \{\} \rangle$  ▷ Node with minimum objective value
8:    $\pi_h^* \leftarrow \pi^*$  ▷ Optimal plan being explained
9:    $\pi_h^R \leftarrow \pi$  s.t.  $C(\pi, \mathcal{M}_h^R) = C_{\mathcal{M}_h^R}^*$  ▷ Plan expected by human
10:  fringe.push( $\langle \mathcal{M}_h^R, \{\} \rangle$ , priority = 0)
11:  while True do
12:     $\langle \widehat{\mathcal{M}}, \mathcal{E} \rangle, c \leftarrow$  fringe.pop( $\widehat{\mathcal{M}}$ )
13:    if OBJ_VAL( $\langle \widehat{\mathcal{M}}, \mathcal{E} \rangle$ )  $\leq$  OBJ_VAL( $N_{min}$ ) then
14:       $N_{min} \leftarrow \langle \widehat{\mathcal{M}}, \mathcal{E} \rangle$  ▷ Update min node
15:    if  $C(\pi_h^*, \mathcal{M}^R) = C_{\mathcal{M}^R}^*$  then
16:       $\langle \mathcal{M}_{min}, \mathcal{E}_{min} \rangle \leftarrow N_{min}$ 
17:      return  $\langle \pi_{\mathcal{M}_{min}}, \mathcal{E}_{min} \rangle$  ▷ If  $\pi_h^*$  is optimal in  $\mathcal{M}^R$ 
18:    else
19:      c_list  $\leftarrow$  c_list  $\cup$   $\widehat{\mathcal{M}}$ 
20:      for  $f \in \Gamma(\widehat{\mathcal{M}}) \setminus \Gamma(\mathcal{M}^R)$  do ▷ Models that satisfy condition 1
21:         $\lambda \leftarrow \langle 1, \{\widehat{\mathcal{M}}\}, \{f\} \rangle$  ▷ Removes f from  $\widehat{\mathcal{M}}$ 
22:        if  $\delta_{\mathcal{M}_h^R, \mathcal{M}^R}(\Gamma(\widehat{\mathcal{M}}), \lambda) \notin$  c_list then
23:          fringe.push( $\langle \delta_{\mathcal{M}_h^R, \mathcal{M}^R}(\Gamma(\widehat{\mathcal{M}}), \lambda), \mathcal{E} \cup \lambda \rangle$ , c + 1)
24:        for  $f \in \Gamma(\mathcal{M}^R) \setminus \Gamma(\widehat{\mathcal{M}})$  do ▷ Models that satisfy condition 2
25:           $\lambda \leftarrow \langle 1, \{\widehat{\mathcal{M}}\}, \{f\}, \{\} \rangle$  ▷ Adds f to  $\widehat{\mathcal{M}}$ 
26:          if  $\delta_{\mathcal{M}_h^R, \mathcal{M}^R}(\Gamma(\widehat{\mathcal{M}}), \lambda) \notin$  c_list then
27:            fringe.push( $\langle \delta_{\mathcal{M}_h^R, \mathcal{M}^R}(\Gamma(\widehat{\mathcal{M}}), \lambda), \mathcal{E} \cup \lambda \rangle$ , c + 1)
28:  procedure OBJ_VAL( $\langle \widehat{\mathcal{M}}, \mathcal{E} \rangle$ )
29:  return  $|\mathcal{E}| + \alpha \times |C(\pi_h^*, \mathcal{M}^R) - C_{\mathcal{M}^R}^*|$ 

```

Property 3. $\alpha = 0$ yields the most explicable plan.

Under this condition, the planner has to minimize the cost of explanations only. Of course, at this point it will produce the plan that requires the shortest explanation, and hence the most explicable plan. Note that this is distinct from just computing the optimal plan in the human's model, since such a plan may not be executable in the planner's model so that some explanations are required even in the worst case. This is also a welcome additions to the explicability only view of plan generation introduced in [23, 34, 35], where the human model only also guides the plan generation process instead of doing so directly, though none of these works provided any insight into how to make the remainder of the model reconciliation possible in such cases, as done here with the explanations associated with the generated plans.

Property 4. MEGA-search is required only once per problem, and is independent of α .

Algorithm 1 terminates only after all the nodes containing a minimally complete explanation³ have been explored. This means that for different values of α , the agent only needs to post-process the nodes with the new objective function in mind. Thus, a large part of the reasoning process for a particular problem can be pre-computed.

2 EVALUATIONS

We will now provide internal evaluations of MEGA in modified versions of two well-known IPC domains Rover and Barman [20]

Domain Name	Problem	\mathcal{E}	$\Delta = 2$		$\Delta = 7$		$\Delta = 10$	
			Time (secs)	\mathcal{E}	Time (secs)	\mathcal{E}	Time (secs)	\mathcal{E}
Rover	p1	0	1.22	1	5.83	3	143.84	
	p2	1	1.79	5	125.64	6	1061.82	
	p3	0	8.35	2	10.46	3	53.22	
Barman	p1	2	18.70	6	163.94	6	5576.06	
	p2	2	2.43	4	57.83	6	953.47	
	p3	2	45.32	5	4183.55	6	5061.50	

Table 1: Computation time for human-aware plans in Rover and Barman domains along with the length of explanations.

demonstrating the trade-off in the cost and computation time of plans with respect to varying size of the model difference and the hyper-parameter α and follow it up with a demonstration of MEGA in action on a robot in a search and reconnaissance domain. Finally, we will report on human factor studies on how this trade-off is received by users. The code and the domain models will be available after the double-blind review process is over.

2.1 Empirical Results: Cost Trade-off

The value of α determines how much an agent is willing to sacrifice its own optimality versus the cost of explaining a (perceived) suboptimal plan to the human. In the following, we illustrate this trade-off on modified versions of two well-known IPC domains.

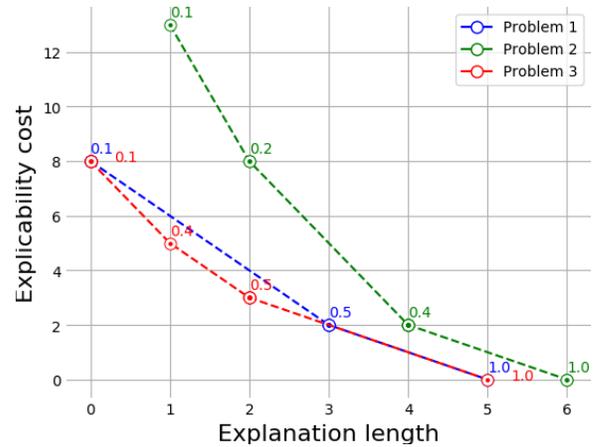
2.1.1 The Rover (Meets a Martian) Domain. Here the Mars Rover has a model as described in the IPC domain, but has gone an update whereby it can carry all the rock and soil samples needed for a mission at the same time. This means that it does not need to empty the store before collecting new rock and soil samples anymore so that the new action definitions for `sample_soil` and `sample_rock` no longer contain the precondition (`empty ?s`).

During its mission it runs across a Martian who is unaware of the robot’s expanded storage capacity, and has an older, extremely cautious, model of the rover it has learned while spying on it from its cave. It believes that any time we collect a rock sample, we also need to collect a soil sample and need to communicate this information to the lander. It also believes that before the rover can perform `take_image` action, it needs to send the soil data and rock data of the waypoint from where it is taking the image. Clearly, if the rover was to follow this model in order not to spook the Martians, it will end up spending a lot of time performing unnecessary actions (like dropping old samples and collecting unnecessary samples). For example, if the rover is to communicate an image of an objective `objective2`, all it needs to do is move to a waypoint (`waypoint3`) from where `objective2` is visible and perform the action –

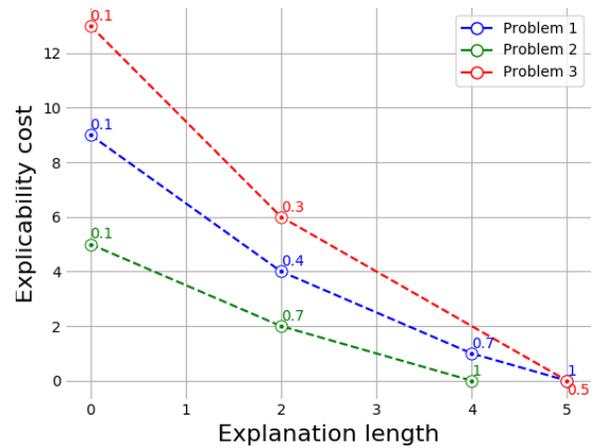
```
(take_image waypoint3 objective2 camera0 high_res)
```

If the rover was to produce a plan that better represents the Martian’s expectations, it would look like –

```
(sample_soil store waypoint3)
(communicate_soil_data general waypoint3 waypoint3 waypoint0)
(drop_off store)
(sample_rock store waypoint3)
(communicate_rock_data general waypoint3 waypoint3 waypoint0)
(take_image waypoint3 objective1 camera0 high_res)
```



(a) The Rover (Meets a Martian) Domain



(b) The Barman (in a Bar) Domain

Figure 2: Trade-off between explicability versus explanation cost for plans produced at different values of α .

Now if the rover chose to directly use an MCE it could end up explaining up to six different model differences based on the problem and the plan under execution. In some case, this may be acceptable, but in others, it may make more sense for the rover to bear the extra cost rather than laboriously walking through all the updates with an impatient Martian. MEGA lets us naturally model these scenarios through the use of the α parameter – the rover would choose to execute the Martian’s expected optimal plan when the α parameter is set to zero (which means the rover does not care about the extra cost it needs to incur to ensure that the plan makes sense to the Martian with the least explaining involved).

Figure 2 shows how the explicability cost and explanation cost varies for three typical problem instances in this domain. The algorithm starts converging to the smallest possible MCE, when α is set to one. For smaller α , MEGA chooses to save explanation costs by choosing more expensive (and explicable) plans.

2.1.2 The Barman (in a Bar) Domain. Here, the brand new two-handed Barman robot is wowing onlookers with its single-handed

skills, even as its admirers who may be unsure of its capabilities expect, much like in the original IPC domain, that it is required to have one hand free to perform actions like fill-shot, refill-shot, shake etc. This means that to make a single shot of a cocktail with two shots of the same ingredient with three shots and one shaker, the human expects the robot to execute the following plan –

```
(fill-shot shot2 ingredient2 left right dispenser2)
(pour-shot-to-used-shaker shot2 ingredient3 shaker1 left 11 12)
(refill-shot shot2 ingredient3 left right dispenser3)
(pour-shot-to-used-shaker shot2 ingredient3 shaker1 left 11 12)
(leave left shot2)
(grasp left shaker1)
```

The robot can, however, directly start by picking both the shot and the shaker and does not need to put either of them down while making the cocktail. Similar to the Rover domain, we again illustrate on three typical problems from the barman domain (Figure 2) how at lower values of α the robot choose to perform plans that require less explanation. As α increases the algorithm produces plans that require larger explanations with the explanations finally converging at the smallest MCE required for that problem.

2.2 Empirical Results: Computation Time

Contrary to classical notions of planning that occurs in state or plan space, we are now planning in the model space, i.e. every node in the search tree is a new planning problem. As seen in Table 1 this becomes quite time consuming with increasing number of model differences between the human and the robot, even as there are significant gains to be had in terms of minimality of explanations, and the reduction in cost of explicable plans as a result of it. This motivates the need for developing approximations and heuristics [10] for the search for multi-model explanations.

2.3 Demonstration: The USAR Domain

We first demonstrate MEGA on a robot performing an Urban Search And Reconnaissance (USAR) task – here a remote robot is put into disaster response operation often controlled partly or fully by an external human commander. This is a typical USAR setup [5], where the robot’s job is to infiltrate areas that may be otherwise harmful to humans, and report on its surroundings as and when required / instructed by the external. The external usually has a map of the environment, but this map is no longer accurate in a disaster setting - e.g. new paths may have opened up, or older paths may no longer be available, due to rubble from collapsed structures like walls and doors. The robot (internal) however may not need to inform the external of all these changes so as not to cause information overload of the commander who may be otherwise engaged in orchestrating the entire operation. This calls for an instantiation of the MEGA algorithm where the model differences are contributed to by changes in the map, i.e. the initial state of the planning problem (the human model has the original unaffected model of the world).

Figure 3 shows a relevant section of the map of the environment where this whole scenario plays out. The orange marks indicate rubble that has blocked a passage, while the green marks indicate collapsed walls. The robot (Fetch), currently located at the position marked with a blue **O**, is tasked with taking a picture at the location marked with an orange **O** in the figure. The external commander’s

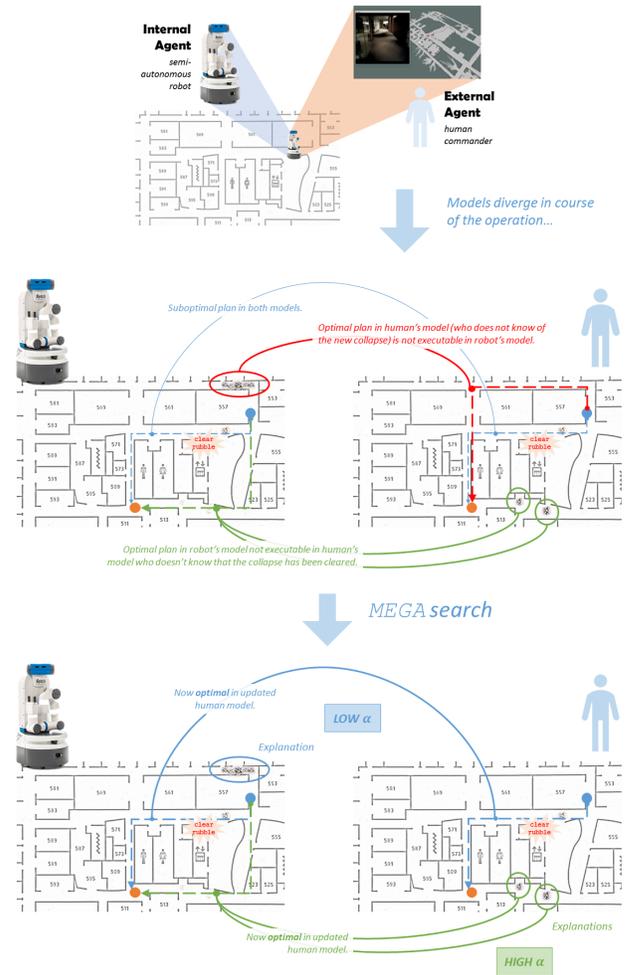


Figure 3: A typical search and reconnaissance scenario with an internal semi-autonomous agent (robot) and an external supervisor (human) – a video demonstration can be accessed at https://www.youtube.com/watch?v=u_t1TQotzo4.

expects the robot to take the path shown in red, which is no longer possible. The robot armed with MEGA has two choices – it can either follow the green path and explain the revealed passageway due to the collapse, or compromise on its optimal path, clear the rubble and proceed along the blue path. A video demonstration of the scenario can be viewed at https://www.youtube.com/watch?v=u_t1TQotzo4. The first part of the video demonstrates the plan generated by MEGA for low α values. As expected, it chooses the blue path that requires the least amount of explanation, and is thus the most explicable plan. In fact, the robot only needs to explain a single initial state change to make this plan optimal, namely –

Explanation >> remove-has-initial-state-clear_path p1 p8

This is also an instance where the plan closest to the human expectation, i.e. the most explicable plan, still requires an explanation,

which previous approaches in the literature cannot provide. Moreover, in order to follow this plan, the robot must perform the costly `clear_passage p2 p3` action to traverse the corridor between `p2` and `p3`, which it could have avoided in its optimal plan (shown in green on the map). Indeed, MEGA switches to the robot's optimal plan for higher values of α along with the following explanation –

```
Explanation >> add-has-initial-state-clear_path p6 p7
Explanation >> add-has-initial-state-clear_path p7 p5
Explanation >> remove-has-initial-state-clear_path p1 p8
```

By providing this explanation, the robot is able to convey to the human the optimality of the current plan as well as the infeasibility of the human's expected plan (shown in red).

2.4 Human Factors Evaluations

Finally, we will now use the above search and reconnaissance domain to analyze how humans respond to the explicability versus explanations trade-off. This is done by exposing the external commander's interface to participants who get to analyze plans in a mock USAR scenario. The participants were incentivized to make sure that the explanation does indeed help them understand the optimality of the plans in question by formulating the interaction in the form of a game. This is to make sure that participants were sufficiently invested in the outcome as well as mimic the high-stakes nature of USAR settings to accurately evaluate the explanations.

Figure 4 shows a screenshot of the interface which displays to each participant an initial map (which they are told may differ from the robot's actual map), the starting point and the goal. A plan is illustrated in the form of a series of paths through various waypoints highlighted on the map. The participant has to identify if the plan shown is optimal. If the player is unsure, they can ask for an explanation. The explanation is provided to the participant in the form of a set of model changes in the player's map.

The scoring scheme for the game is as follows. Each player is awarded 50 points for correctly identifying the plan as either optimal or satisfying. Incorrectly identification costs them 20 points. Every request for explanation further costs them 5 points, while skipping a map does not result in any penalty. The participants were additionally told that selecting an inexecutable plan as either feasible or optimal would result in a penalty of 400 points. Even though there were no actual incorrect plans in the dataset, this information was provided to deter participants from taking chances with plans they did not understand well.

Each participant was paid \$10 dollars and received additional bonuses based on the following payment scheme –

- Scores higher than or equal to 540 were paid \$10.
- Scores higher than 540 and 440 were paid \$7.
- Scores higher than 440 and 340 were paid \$5.
- Scores higher than 340 and 240 were paid \$3.
- Scores below 240 received no bonuses.

The scoring systems for the game was designed to make sure –

- Participants should only ask for an explanation when they are unsure about the quality of the plan (due to small negative points on explanations).

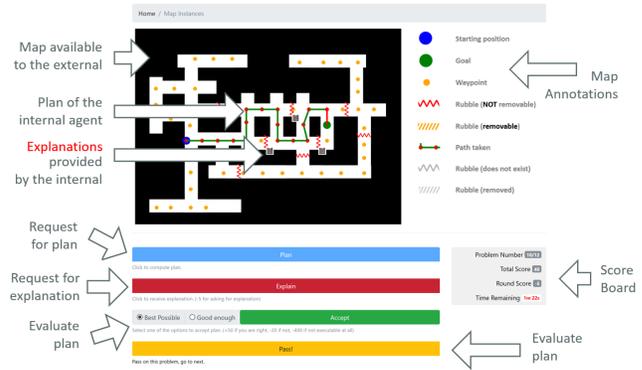


Figure 4: Interface to the external commander in a mock search and reconnaissance study.

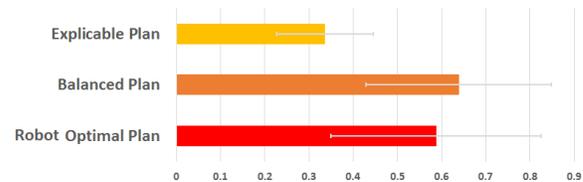


Figure 5: Responses to explicable plans versus balanced or robot optimal plans with explanations.

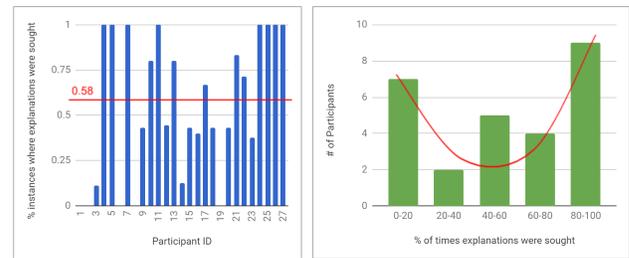


Figure 6: Click-through rates for explanations.

- Participants are incentivized to identify the feasibility and optimality of the given plan correctly (large reward and penalty on doing this wrongly).

Each participant was shown a total of 12 maps. For 6 of the 12 maps, the participant was assigned the optimal robot plan, and when they asked for an explanation, they were randomly shown different types of explanations as introduced in [10]. For the rest of the maps, in place of the robot's optimal plan, participants could potentially be assigned a plan that is optimal in the human model (i.e. an explicable plan) with no explanation or somewhere in between (i.e. the balanced plan) with a shorter explanation. Note that out of the 6 maps, only 3 had both balanced plans as well as explicable plans, the other 3 either had a balanced plan or the optimal human plan. In total, we had 27 participants for the study, including 4 female and 22 male participants between the age range of 19-31 (1 participant did not reveal their demographic).

Optimal Plan		Balanced Plan		Explicable Plan	
$ \mathcal{E} $	$C(\pi, M^R)$	$ \mathcal{E} $	$C(\pi, M^R)$	$ \mathcal{E} $	$C(\pi, M^R)$
2.5	5.5	1	8.5	-	16

Table 2: Statistics of explicability versus explanation trade-off with respect to explanation length and plan cost.

Figure 5 shows how people responded to the different kinds of explanations / plans. These results are from 382 problem instances that required explanations, and 25 and 40 instances that contained balanced and explicable plans respectively. From the perspective of the human, the balanced plan and the robot optimal plan do not make any difference since both of them appear suboptimal. This is evident from the fact that the click-through rate for explanations in these two conditions are similar. However, the rate of explanations is significantly less in case of explicable plans as desired.

Table 2 shows the statistics of the explanations / plans. These results are from 124 problem instances that required *minimal* explanations as per [10], and 25 and 40 instances that contained balanced and explicable plans respectively, as before. As desired, the robot gains in length of explanations but loses out in cost of plans produced as it progresses along the spectrum of optimal to explicable plans. Thus, while Table 2 demonstrates the cost of explanation versus explicability trade-off from the robot’s point of view, Figure 5 shows how this trade-off is perceived from the human’s perspective.

It is interesting to see that in Figure 5 about a third of the time participants still asked for explanations even when the plan was explicable, and thus optimal in their map. This is an artifact of the risk-averse behavior incentivized by the gamification of the explanation process and indicative of the cognitive burden on the humans who are not (cost) optimal planners. Thus, going forward, the objective function should incorporate the cost or difficulty of analyzing the plans and explanations from the point of view of the human in addition to the current costs in equation MEGA(4) and Table 2 modeled from the perspective of the robot model.

Finally, in Figure 6, we show how the participants responded to inexplicable plans, in terms of their click-through rate on the explanation request button. Such information can be used to model the α parameter to situate the explicability versus explanation trade-off according to preferences of individual users. It is interesting to see that the distribution of participants (right inset) seem to be bimodal indicating that there are people who are particularly skewed towards risk-averse behavior and others who are not, rather than a normal distribution of response to the explanation-explicability trade-off. This further motivates the need for learning α interactively with the particular human in the loop.

3 DISCUSSION AND FUTURE WORK

In the following section, we will elaborate on some of the exciting avenues of future research borne out of this work.

3.1 Model learning and picking the right α

We assumed that the hyper-parameter α is set by the designer in determining how much to trade-off the costs of explicability versus explanations on the part of the autonomous agent. However, the

design of α itself can be more adaptive and “human-aware” in the sense that the parameter can be *learned* in course of interactions with the human in the loop to determine what kind of plans are preferred (as seen in Figure 6) and how much information can be transmitted. This is also relevant in cases where the human mental model is not known precisely or if there is uncertainty towards what the new model is after an update or explanation. This is a topic of future work; existing literature on iterative model learning [18, 29] can provide useful guidance towards the same. Authors in [9] discuss a few useful representations for learning such models for the purposes of task planning at various levels of granularity. Note that search with uncertainty over a learned human (mental) model can often times be compiled to the same planning process as described in [31] by using annotated models, so the same techniques as introduced in this paper still apply.

3.2 Cost of explanations and cognitive load

Currently, we only considered the cost of explanations and explicability from the point of view of the robot. However, there might be additional (cognitive) burden on the human – measured in terms of the complexity of interpreting an explanation and how far away the final plan is from the optimal plan in the human’s mental model. This again ties back to the assumptions on the cognitive abilities (i.e. optimality) of the human in the loop, and needs calibration [18, 29] based on repeated interactions (as seen in Figure 5).

4 CONCLUSION

We saw how an agent can achieve human-aware behavior while at the same time keeping in mind the cost of departure from its own optimality which could otherwise have been explained away if given the opportunity. This raises several intriguing challenges in the plan generation process, most notably in finding better heuristics in speeding up the model space search process as well as dealing with model uncertainty and identifying the sweet spot of the algorithm in explicability-explanations trade-off. Indeed, the revised human-aware planning paradigm opens up exciting new avenues of research such as learning human mental models, providing explanations at different levels of abstractions, and so on.

REFERENCES

- [1] Mitchell Ai-Chang, John Bresina, Len Charest, Adam Chase, JC-J Hsu, Ari Jonsson, Bob Kanefsky, Paul Morris, Kanna Rajan, Jeffrey Yglesias, et al. 2004. MAPGEN: Mixed-Initiative Planning and Scheduling for the Mars Exploration Rover Mission. *IEEE Intelligent Systems* (2004).
- [2] Rachid Alami, Aurélie Clodic, Vincent Montreuil, Emrah Akin Sisbot, and Raja Chatila. 2006. Toward Human-Aware Robot Task Planning. In *AAAI Spring Symposium: To Boldly Go Where No Human-Robot Team Has Gone Before*.
- [3] Rachid Alami, Mamoun Gharbi, Benjamin Vadant, Raphaël Lallemand, and Adolfo Suarez. 2014. On human-aware task and motion planning abilities for a teammate robot. In *Human-Robot Collaboration for Industrial Manufacturing Workshop, RSS*.
- [4] James F Allen. 1994. Mixed initiative planning: Position paper. In *ARPA/Rome Labs Planning Initiative Workshop*.
- [5] Cade Earl Bartlett. 2015. Communication between Teammates in Urban Search and Rescue. *Thesis* (2015).
- [6] Alexandros Belesiotis, Michael Rovatsos, and Iyad Rahwan. 2010. Agreeing on plans through iterated disputes. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: volume 1-Volume 1*. International Foundation for Autonomous Agents and Multiagent Systems, 765–772.
- [7] Rodney Brooks. 1986. A robust layered control system for a mobile robot. *IEEE journal on robotics and automation* 2, 1 (1986), 14–23.
- [8] T. Chakraborti, G. Briggs, K. Talamadupula, Yu Zhang, M. Scheutz, D. Smith, and S. Kambhampati. 2015. Planning for serendipity. In *IROS*.

- [9] Tathagata Chakraborti, Subbarao Kambhampati, Matthias Scheutz, and Yu Zhang. 2017. AI Challenges in Human-Robot Cognitive Teaming. *arXiv preprint arXiv:1707.04775* (2017).
- [10] Tathagata Chakraborti, Sarath Sreedharan, Yu Zhang, and Subbarao Kambhampati. 2017. Plan Explanations as Model Reconciliation: Moving Beyond Explanation as Soliloquy. In *IJCAI*.
- [11] Tathagata Chakraborti, Yu Zhang, David Smith, and Subbarao Kambhampati. 2016. Planning with Resource Conflicts in Human-Robot Cohabitation. In *AAMAS*.
- [12] Marcello Cirillo. 2010. *Planning in inhabited environments: human-aware task planning and activity recognition*. Ph.D. Dissertation. Örebro university.
- [13] Marcello Cirillo, Lars Karlsson, and Alessandro Saffiotti. 2010. Human-aware Task Planning: An Application to Mobile Robots. *ACM Transactions on Intelligent Systems and Technology* (2010).
- [14] Anca Dragan, Kenton Lee, and Siddhartha Srinivasa. 2013. Legibility and Predictability of Robot Motion. In *Human-Robot Interaction*.
- [15] Chukwuemeka D Emele, Timothy J Norman, and Simon Parsons. 2011. Argumentation strategies for plan resourcing. In *The 10th International Conference on Autonomous Agents and Multiagent Systems-Volume 3*. International Foundation for Autonomous Agents and Multiagent Systems, 913–920.
- [16] George Ferguson, James F Allen, Bradford W Miller, et al. 1996. TRAINS-95: Towards a Mixed-Initiative Planning Assistant. In *AIPS*.
- [17] Maria Fox, Derek Long, and Daniele Magazzeni. 2017. Explainable Planning. In *IJCAI XAI Workshop*.
- [18] Dylan Hadfield-Menell, Stuart J Russell, Pieter Abbeel, and Anca Dragan. 2016. Cooperative inverse reinforcement learning. In *NIPS*.
- [19] Marc Hanheide, Moritz Göbelbecker, Graham S Horn, Andrzej Pronobis, Kristoffer Sjö, Alper Aydemir, Patric Jensfelt, Charles Gretton, Richard Dearden, Miroslav Janicek, et al. 2015. Robot task planning and explanation in open and uncertain worlds. *Artificial Intelligence* (2015).
- [20] International Planning Competition. 2011. IPC Competition Domains. <https://goo.gl/i35bxc>. (2011).
- [21] Subbarao Kambhampati and Kartik Talamadupula. 2014. Human-in-the-Loop Planning and Decision Support. <http://rakaposhi.eas.asu.edu/hilp-tutorial/>. (2014).
- [22] Uwe Koeckemann, Federico Pecora, and Lars Karlsson. 2014. Grandpa Hates Robots - Interaction Constraints for Planning in Inhabited Environments. In *AAAI*.
- [23] Anagha Kulkarni, Tathagata Chakraborti, Yantian Zha, Satya Gautam Vadlamudi, Yu Zhang, and Subbarao Kambhampati. 2016. Explicable Robot Planning as Minimizing Distance from Expected Behavior. *CoRR* abs/1611.05497 (2016).
- [24] Pat Langley, Ben Meadows, Mohan Sridharan, and Dongkyu Choi. 2017. Explainable Agency for Intelligent Autonomous Systems. In *AAAI/LAAL*.
- [25] Lydia Manikonda, Tathagata Chakraborti, Kartik Talamadupula, and Subbarao Kambhampati. 2017. Herding the Crowd: Using Automated Planning for Better Crowdsourced Planning. *Journal of Human Computation* (2017).
- [26] Hugo Mercier and Dan Sperber. 2010. Why Do Humans Reason? Arguments for an Argumentative Theory. *Behavioral and Brain Sciences* (2010).
- [27] Tim Miller, Jens Pfau, Liz Sonenberg, and Yoshihisa Kashima. 2017. Logics of common ground. *Journal of Artificial Intelligence Research* (2017).
- [28] Christian J Muise, Paolo Felli, Tim Miller, Adrian R Pearce, and Liz Sonenberg. 2016. Planning for a Single Agent in a Multi-Agent Environment Using FOND. In *IJCAI*.
- [29] Stefanos Nikolaidis, Przemyslaw Lasota, Ramya Ramakrishnan, and Julie Shah. 2015. Improved human-robot team performance through cross-training, an approach inspired by human team training practices. *International Journal of Robotics Research* (2015).
- [30] Sailik Sengupta, Tathagata Chakraborti, Sarath Sreedharan, and Subbarao Kambhampati. 2017. RADAR - A Proactive Decision Support System for Human-in-the-Loop Planning. In *AAAI Fall Symposium on Human-Agent Groups*.
- [31] S. Sreedharan, T. Chakraborti, and S. Kambhampati. 2017. Explanations as Model Reconciliation - A Multi-Agent Perspective. In *AAAI Fall Symposium on Human-Agent Groups*.
- [32] Kartik Talamadupula, Gordon Briggs, Tathagata Chakraborti, Matthias Scheutz, and Subbarao Kambhampati. 2014. Coordination in human-robot teams using mental modeling and plan recognition. In *Intelligent Robots and Systems (IROS 2014), 2014 IEEE/RSJ International Conference on. IEEE*, 2957–2962.
- [33] Stevan Tomic, Federico Pecora, and Alessandro Saffiotti. 2014. Too Cool for School??? Adding Social Constraints in Human Aware Planning. In *Workshop on Cognitive Robotics (CogRob)*.
- [34] Yu Zhang, Sarath Sreedharan, Anagha Kulkarni, Tathagata Chakraborti, Hankui Hankui Zhuo, and Subbarao Kambhampati. 2016. Plan Explicability for Robot Task Planning. In *RSS Workshop on Planning for Human-Robot Interaction*.
- [35] Yu Zhang, Sarath Sreedharan, Anagha Kulkarni, Tathagata Chakraborti, Hankui Hankui Zhuo, and Subbarao Kambhampati. 2017. Plan Explicability and Predictability for Robot Task Planning. In *ICRA*.