Planning with Humans in the Loop

Tathagata Chakraborti
Advisor: Dr. Subbarao Kambhampati
Department of Computer Science
Arizona State University
Tempe, AZ 85281, USA

Abstract

Automated planners have traditionally looked at optimizing cost or duration of plans using complete domain information and user preferences, or improving efficiency of the planning process itself. However as artificial intelligence becomes more and more part of our daily lives, interactions between autonomous systems and humans throw up different challenges. With regards to automated planning, these not only involve reasoning with incomplete information and unknown preferences, but also the ability to account for different aspects of interacting with humans. In my doctoral work, I consider these issues in the context of planning for synergy in human-robot cohabitation and planning for collaborative systems involving human computation. In this brief abstract, I present the overarching goal of my research, the overall challenges addressed in my work till now and the state of the current work in progress.

Planning with humans in the loop involves not only the ability to reason with uncertainty or incompleteness in terms of both knowledge of the environment and the preferences of the human colleagues but also the ability to deal with the different aspects of interacting with humans. As we will see in course of this discussion, different modes of interacting with humans introduce different challenges into the planning process. In my present work I look at two modalities of planning with humans in the loop. In the first half of this abstract I will discuss the human-aware planning paradigm and its role in human-robot cohabitation. Particularly we will look at how we can achieve both active and passive coordination among agents sharing resources in the same environment - in case of active synergy we will look at how an autonomous agent can offer help without expectations or commitments from human colleagues, while in case of passive synergy we will see how the autonomous agent can avoid conflicts of interest while sharing resources. In the second half of the abstract I will describe the role of a planner in human computation tasks that involve planning and scheduling. Specifically, we will investigate how the role of a planner changes with the type (expert or non-expert) of the human collaborators and the richness of the domain information in the context of crowdsourced planning dealing with construction of tour plans and proactive decision systems dealing with disaster response management.

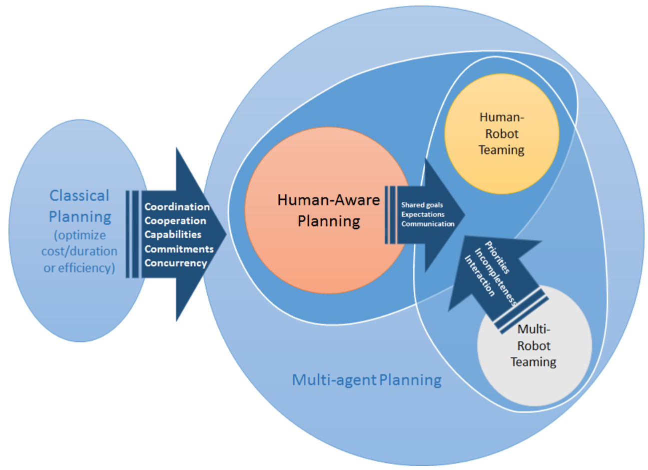

Figure 1: The scope of human-aware planning - as a subclass of multi-agent planning with and without some distinctive features with respect to multi-robot or human-robot teaming.

Synergy in Human-Robot Cohabitation

As robots become ubiquitous in our daily lives, it becomes important to model acceptable or desired behaviors of these autonomous agents that cohabit our environment. Indeed there has been a lot of work under the umbrella of "human-aware" planning, both in the context of path planning (Sisbot et al. 2007; Kuderer et al. 2012) and in task planning (Koeckemann, Pecora, and Karlsson 2014; Cirillo, Karlsson, and Saffiotti 2010), that aim to provide social skills to robots so as to make them produce plans conforming to desired behaviors when humans and robots operate in shared settings. Human-aware planning in fact holds a unique spot in the area of multi-agent planning as illustrated in Figure 1. As we move over to multi-agent planning from classical planning, we have to deal with challenges in coordination, capability and commitment modeling, handling concurrency, etc. But even within the multi-agent planning paradigm, introducing a human in the context of multi-robot teaming produces new challenges like model completeness, priorities and interaction issues. Further, the presence or absence of a team itself determines if we can assume shared goals and expectations. Thus, we can think of human-aware

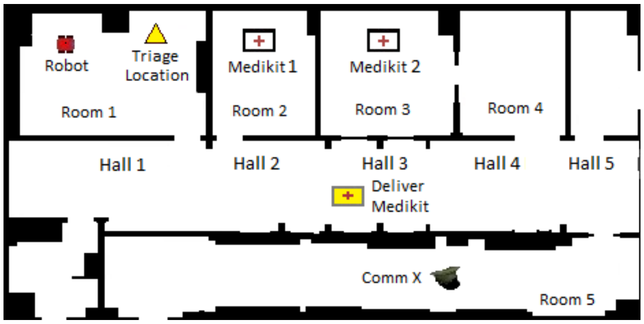

Figure 2: A mock-up of the Urban Search and Rescue (USAR) setting used as a running example throughout the work on human-robot cohabitation.

planning as multi-agent planning which includes the flavors of human-robot interaction, but mostly excludes the assumptions often made in explicit teaming scenarios. We will investigate these concepts using a running example of a typical search and rescue scenario. Figure 2 shows such an USAR setting, unfolding inside a building with interconnected rooms and hallways, with a human commander `CommX` and a robot `R`. The commander has capabilities to move and conduct triage at specified locations, and he can also meet with other agents, as well as pickup, dropoff or handover medkits to accomplish their task. The robot can similarly move about, search rooms, or handover or deliver the medkits. It can thus have its own goals (maybe from being directly assigned by the commander himself or due to long term task specifications), but can also help the commander in accomplishing his goals by fetching the medkits for him. In (Talamadupula et al. 2014) we introduced a scheme for the robot to represent first order beliefs of other agents in its environment. The question then is can the robot use such predictive models to make smarter decisions in its future?

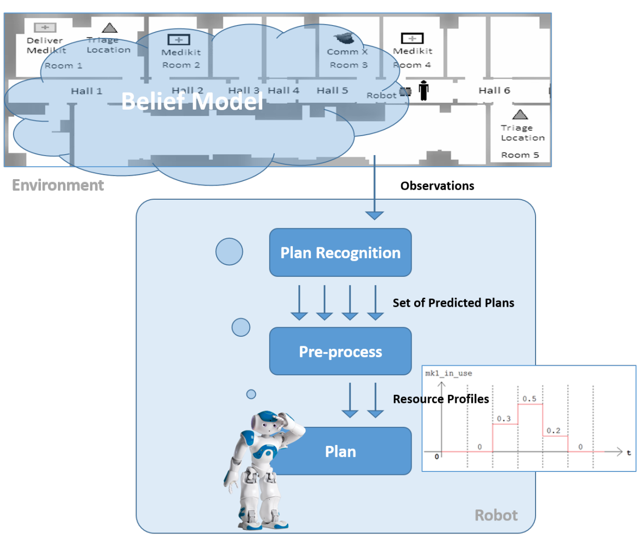

Figure 3: Architecture diagram showing different components of planning with resource profiles

Planning with Resource Profiles

The first form of synergy that we address is passive synergy, where the robot tries to avoid conflict of resource usages as it tries to use the medkits - specifically, in the USAR environment discussed before, the medkit is a constrained resource, that is demanded by both the agents. We approach the problem by using the distribution of goals of the human and the known model of the human to predict and approximate resource usage probabilities over time. Thus the problem of decoupling the human's plan and the robot's plan essentially boils down to minimizing the overlap between the resource usages induced by each of these plans. We convert the planning problem to an integer programming instance in order to minimize over this overlap elegantly. Recall here again that we do not have a team setting, and that the human and robot are not coordinating here to avoid a conflict, so that the onus is now on the robot to find a fail-safe plan that suits both of them (given the evidence).

Further we show that by modeling resources at various levels of abstractions we are able to model different types of interactions, and by varying the parameters of the IP we can produce different type of behaviors of the robot like opportunism, compromise and negotiation (when limited forms of communication is allowed). The representation of conflicts in terms of such profiles have an inherent advantage that they are no longer dependent on the number of agents the robot has to look out for but only the number of resources it has to reason with. For further details please refer to (Chakraborti, Zhang, and Kambhampati 2015).

Planning for Serendipity

Here we look at a case of active synergy - the robot now actively tries to seek out opportunities to help the human. Recall that this is not a team setting and hence the human does not expect the robot to help. So any helpful change the robot makes to the world will appear as positive exogenous events to the human while he is executing his plan. We refer to such exceptions as serendipitous moments for the human, and since the robot is trying to produce such moments, we name this planning paradigm as "planning for serendipity".

Note, however, that the absence of any expectations or commitments means that all seemingly helpful interventions do not turn out to be helpful when the plans of the individual agents are actually executed, which brings us to the concept of "plan interruptibility" and "plan preservation".

Plan Interruptibility. This outlines the structure of the resultant joint human-robot plan such that it conforms to the notion of serendipitous interventions.

Preservation Constraints. These outline a set of constraints on the structure of an interruptible plan that determines whether it is possible for the robot to produce a serendipitous intervention or not.

With the help of these two guiding principles we convert the planning problem into an integer programming instance and show how the robot can preempt help both with and without scope of explicit communication. Please refer to (Chakraborti et al. 2015) for a detailed analysis.

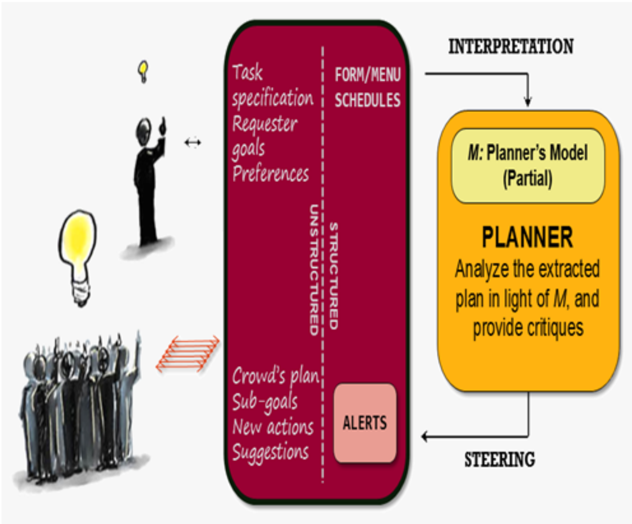

Figure 4: Crowdsourced Planning - architecture diagram showing the collaborative blackboard and the interaction between the planner and the crowd.

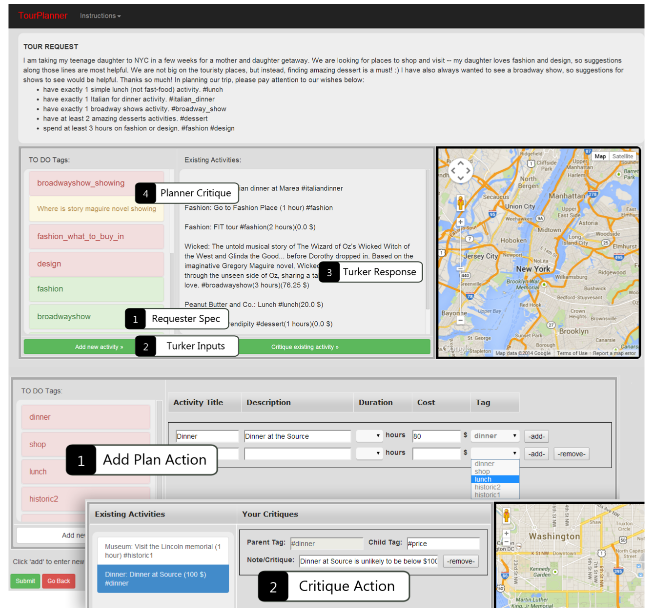

Figure 5: Snapshots showing different components of the award winning AI-MIX interface.

Humans as Collaborators

Traditionally mixed initiative planning has looked at scenarios where human experts can bring in their complex domain knowledge by critiquing automated planners during the planning process. A prime example of this is MAPGEN (Ai-Chang et al. 2004) - the mixed-initiative planning system for NASA's Mars Rover exploration mission. Interestingly, a large subclass of human computation applications are those directed at tasks that involve planning (e.g. tour planning) and scheduling (e.g. conference scheduling), where it is the humans that take part in the actual planning process. This can be looked upon as a *reverse mixed initiative planning* scenario where the humans can range from a small group of experts (e.g. commanders making battle plans) to a large crowd (as in crowdsourced tasks on MTurk).

Interestingly, work (Zhang et al. 2012) on such systems shows that even simple forms of automated oversight on the human contributors helps in significantly improving the effectiveness of the crowd. In this work, we argue that the automated oversight used in these systems can be viewed as a primitive automated planner, and explore several opportunities for more sophisticated automated planning in effectively steering the crowd. Straightforward adaptation of current planning technology is, however, hampered by the mismatch between the capabilities of human workers and automated planners. We identify and address two important challenges that need to be overcome before such adaptation of planning technology can occur: (i) **interpreting** inputs of the human workers (and the requester) and (ii) **steering** or critiquing plans produced by the human workers, armed only with incomplete domain and preference models. To these ends, we have built a tour plan generation system that uses automated checks and alerts to improve the quality of plans created by human workers (Manikonda et al. 2014; Talamadupula and Kambhampati 2013).

Crowdsourced Planning - AI-MIX

Our work on crowdsourced planning uses the tour planning scenario to demonstrate the effectiveness of involving an automated planner in a human planning task. Note that this is a setting where the domain knowledge is extremely shallow, but we show that even though with such shallow models we cannot generate plans, we can still produce meaningful critiques of the suggested actions in the plan being constructed. The workflow has three major components -

The Request. This is the task description that provides a high level idea of the goals that need to be achieved and the preferences that need to be satisfied. The systems parses and interprets this information in order to produce subgoals that the crowd needs to address. The crowd can choose to do two things in response as follows.

Adding New Activities. The crowd can add a new activity or action in the plan to address outstanding subgoals or violated constraints.

Critiquing Existing Activities. The crowd can chose to critique existing activities if they feel these are invalid given the constraints or there are better options available.

Throughout this process, the system checks for consistencies in the background and pops up with suggestions to improve the plan quality, and computes the final plan using answer set programming. The work flow is shown in Figure 4 illustrating this iterative process on a shared platform we call the collaborative or distributed blackboard. Figure 5 shows a snapshot of the AI-MIX platform that won the People's Choice Best System Demonstration Award in ICAPS 2014. The system was used to build tour plans in the host city of Portsmouth for interested conference attendees.

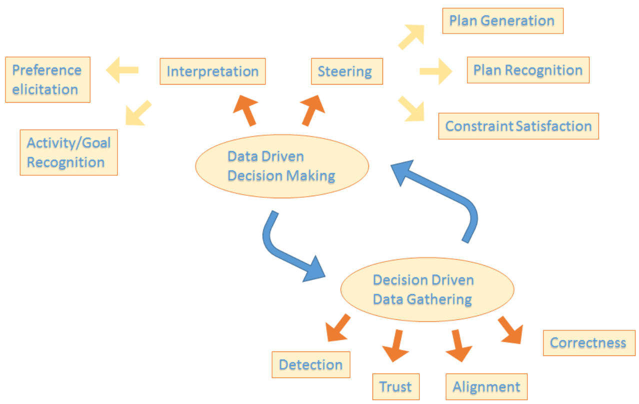

Figure 6: Proactive Decision Support (PDS) - an intricate interplay between data driven decision making and decision driven data gathering.

Proactive Decision Support - RADAR

Now we look at the other end of the spectrum - collaborative planning where both the domain is almost completely known, and the human collaborators are also experts in the domain. Note that this does not mean that the planner can preclude the humans and produce the optimal plan by itself because we do not know a priori the exact objective function the human is trying to optimize. Thus, as in the previous case of crowdsourced planning, any help from the automated system would be in terms of suggestions that improve and optimize the planning process of the humans. However, with a richer model, it turns out we can do much more. We thus divide the workflow into two self-perpetuating components as shown in Figure 6 -

Data Driven Decision Making During planning in complicated domains in the real world, there is a stream of information about the world state that needs to be taken into account during the planning process. In this age of information, it is easy for humans to get overwhelmed by the sheer amount of data coming from both structured (databases) and unstructured (Twitter) sources. Thus the planning component, while being able to generate and recognize plans with the detailed domain knowledge, must also has to be data-centric in its approach.

Decision Driven Data Gathering The planning process itself generates a demand on more data in the future (which can thus be proactively anticipated, pre-fetched and preprocessed) as the world unfolds - thus this completes the cycle of the data being used for planning and the data being generated and queried thereof. This process thus involves detection and alignment of events and verification of trust and correctness of data from heterogeneous sources.

Figure 7 shows a snapshot of the RADAR prototype currently under development. The use case we adopt is the real-time construction of an emergency action strategy (jointly built by the chiefs of police, health, transport and fire) in response to a major fire breakout.

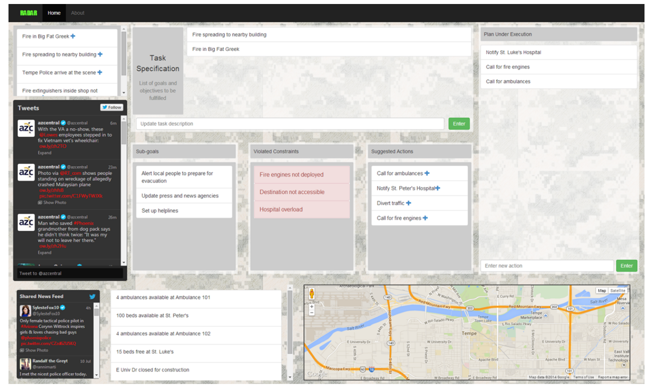

Figure 7: Snapshot of the RADAR prototype currently under development addressing a fire response strategy.

Conclusions and Future Work

So far we have seen the different ways in which planning with humans in the loop differ from classical planning. However, much of my work till now has made certain restrictive assumptions that might be of considerable interest to expand upon in future. Specifically, with regards to planning for synergy in human-robot cohabitation, we have so far assumed a single level belief model. In reality however, humans are seldom completely aloof of other agents in its environment, which means handling nested beliefs of agents, which can have significant impact on the outcome of plans. Further, in all settings involving such human-robot collaboration, researchers have largely overlooked the mismatch of agent models and its impact on the interaction itself. How do we know that humans will recognize a help when there is one? How can the robot act as the human expects it to during collaboration so as not to unpleasantly surprise him? In as much as HRI studies like (Narayanan et al. 2015) will provide clues as to how to address such issues, a general framework of planning to conform to behavioral expectations will be a fascinating direction to pursue in future.

Acknowledgments

This research is supported in part by the ARO grant W911NF-13-1-0023, and the ONR grants N00014-13-10176, N00014-13-1-0519 and N00014-15-1-2027, and a Google Research Grant. I would also like to thank Kartik Talamadupula (ASU, IBM), Lydia Manikonda (ASU), Yu Zhang (ASU) and Gordon Briggs (Tufts University) for their valuable contributions in my work.

References

[Ai-Chang et al. 2004] Ai-Chang, M.; Bresina, J. L.; Charest, L.; Chase, A.; jung Hsu, J. C.; Jnsson, A. K.; Kanefsky, B.; Morris, P. H.; Rajan, K.; Yglesias, J.; Chafin, B. G.; Dias, W. C.; and Maldague, P. F. 2004. Mapgen: Mixed-initiative planning and scheduling for the mars exploration rover mission. *IEEE Intelligent Systems* 19:8–12.

[Chakraborti et al. 2015] Chakraborti, T.; Briggs, G.; Talamadupula, K.; Scheutz, M.; Smith, D.; and Kambhampati, S. 2015. Planning for serendipity - altruism in human-robot cohabitation. In *ICAPS Workshop on Planning and Robotics (PlanRob)*.

[Chakraborti, Zhang, and Kambhampati 2015] Chakraborti, T.; Zhang, Y.; and Kambhampati, S. 2015. Planning with stochastic resource proles: An application to human-robot co-habitation. In *ICAPS Workshop on Planning and Robotics (PlanRob)*.

[Cirillo, Karlsson, and Saffiotti 2010] Cirillo, M.; Karlsson, L.; and Saffiotti, A. 2010. Human-aware task planning: An application to mobile robots. *ACM Trans. Intell. Syst. Technol.* 1(2):15:1–15:26.

[Koeckemann, Pecora, and Karlsson 2014] Koeckemann, U.; Pecora, F.; and Karlsson, L. 2014. Grandpa hates robots - interaction constraints for planning in inhabited environments. In *Proc. AAAI-2010*.

[Kuderer et al. 2012] Kuderer, M.; Kretzschmar, H.; Sprunk, C.; and Burgard, W. 2012. Feature-based prediction of trajectories for socially compliant navigation. In *Proceedings of Robotics: Science and Systems*.

[Manikonda et al. 2014] Manikonda, L.; Chakraborty, T.; De, S.; Talamadupula, K.; and Kambhampati, S. 2014. AI-MIX: How a Planner Can Help Guide Humans Towards a Better Crowdsourced Plan. *Innovative Applications of Artificial Intelligence (IAAI-14)*.

[Narayanan et al. 2015] Narayanan, V.; Zhang, Y.; Mendoza, N.; and Kambhampati, S. 2015. Automated planning for peer-to-peer teaming and its evaluation in remote human-robot interaction. In *Extended Abstract in ACM/IEEE International Conference on Human Robot Interaction (HRI)*.

[Sisbot et al. 2007] Sisbot, E.; Marin-Urias, L.; Alami, R.; and Simeon, T. 2007. A human aware mobile robot motion planner. *Robotics, IEEE Transactions on* 23(5):874–883.

[Talamadupula and Kambhampati 2013] Talamadupula, K., and Kambhampati, S. 2013. Herding the crowd: Automated planning for crowdsourced planning. *CoRR* abs/1307.7720.

[Talamadupula et al. 2014] Talamadupula, K.; Briggs, G.; Chakraborti, T.; Scheutz, M.; and Kambhampati, S. 2014. Coordination in human-robot teams using mental modeling and plan recognition. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2957–2962.

[Zhang et al. 2012] Zhang, H.; Law, E.; Miller, R.; Gajos, K.; Parkes, D.; and Horvitz, E. 2012. Human computation tasks with global constraints. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '12, 217–226. ACM.