Planning for Symbiotic Action (Doctoral Consortium)

Tathagata Chakraborti
Advisor: Subbarao Kambhampati
3^{rd} yr PhD Student, Computer Science Department
Arizona State University, Tempe AZ 85281
tchakra2@asu.edu

ABSTRACT

The field of Artificial Intelligence (AI) has become extremely prominent in recent times, with the integration of different intelligent components into devices and services we use in everyday life. As the capabilities of such systems become more and more complex, one branch of AI that becomes relevant is that of *automated planning* or *sequential decision making*, in order for these components to participate in diverse long term tasks. A key aspect of such systems is increased interaction with humans.

Challenges in Human-in-the-Loop Planning (HILP)

Classical planning has traditionally emphasized on the efficiency or accuracy of the plan generation process. However, in real world applications, especially involving humans, planners must deal with typical challenges including uncertainty and partial knowledge, and issues involving priorities and authority. Technologies that become crucial in this context involve abilities to dynamically predict, anticipate and adapt to changing needs while making task plans. My research focuses on such aspects of *"human-in-the-loop planning"*.

Modalities of HILP - My Research Focus

I have looked at two specific ways in which automated planners may interact with humans. First I will describe how planners can enable different types of autonomous behavior of robots sharing their workspace with humans - i.e. interacting with human *colleagues*. Then I will look at possible roles of automated planners in platforms that involve *collaboration* with or among human planners. The aim of my thesis is to provide planning technologies for and motivate well-informed and principled design of complex symbiotic man-machine systems of the future.

Humans as Colleagues

Many of today's robots built for tasks like household assistance or hotel/office service or security guards, do not operate in teams, i.e. they do not have common goals and commitments with humans sharing the environment, and interactions with such agents depart from traditional notions of proximal human-robot teams. Even though these scenarios require significantly different levels of autonomy from the robot, the underlying theme of autonomy in such settings involves the robot achieving some sense of independence of purpose in so much as its existence is not just defined by the goals of the humans around it but is rather contingent on tasks it is supposed to achieve on its own.

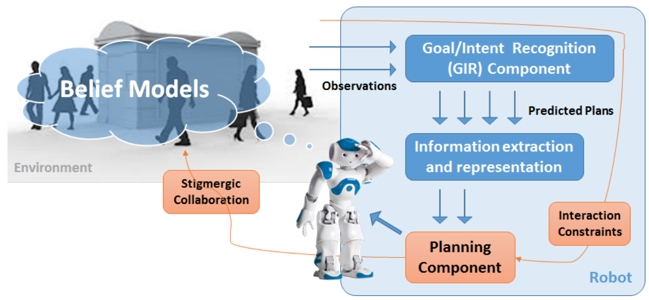

Thus the robots in a sense become colleagues rather than teammates. This becomes even more prominent when we consider *inter-team* and *intra-team* interactions between multiple independent teams in a human-robot cohabited environment. We postulate that interaction with the human cohabitants in such cases should be similar to how we interact with our human colleagues rather than teammates. This provides many interesting possibilities in modeling autonomous behavior in these scenarios - the agents must learn modes of passive or *stigmergic* collaboration. In [1, 2, 3, 4] I model such interactions at three levels of granularity - resources, plans and goals - and motivate the need for developing appropriate metrics for quantifying performance in such settings. Much of the challenge in developing such behaviors is in modeling the appropriate interaction constraints (I use integer programming formulations for this purpose). In the following discussion I will briefly introduce two such models.

Planning with Resource Conflicts

In [4], we look at how robots sharing their workspace with humans, and using shared resources, can plan to minimize conflicts on resource usage. We propose a planner that models the intentions of the human colleagues and produces plans that decouple its resource demands with that of the humans'. Note that there is no explicit or direct interaction here between the human and the robot, the interaction is successful inasmuch as the human's plan was successful.

Resource Profiles. We represent information from predicted plans in the form of *resource profiles*, that enables the robot to reason with how the environment will evolve with time, at different levels of abstraction. This way we compile the complexity of accounting for individual predictions into

Appears in: *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2016), J. Thangarajah, K. Tuyls, C. Jonker, S. Marsella (eds.), May 9–13, 2016, Singapore.*

the number of resources being modeled, which makes our approach independent of both the number of agents being modeled and the size of the hypothesis goal set.

Modes of Behavior. We show how our approach models different autonomous behaviors: (1) *compromise* - the robot opts for suboptimal plans to respect the human's plans; (2) *opportunism* - the robot anticipates favorable changes to the world and produces plans even better than the original optimal ones; and (3) *negotiation* - the robot negotiates desired times of use of the resources with the human.

Planning for Serendipity

In [1], we look at how robots can provide assistance to their human colleagues without the latter asking for help, i.e. the robot plans to produce *serendipitous interventions* (which appear as positive exogenous events) during execution of the human's plans. Contrary to the discussion in the previous section, here the robot's actions directly modify the human's plans rather than just effect it indirectly by avoiding potential conflicts. Note that the human, being unaware, cannot plan to exploit these interventions in advance, which imposes several constraints on the plan generation process.

Plan Interruptibility. In order to produce serendipitous interventions, the robot computes parts of the predicted human plan that can be replaced by a (cheaper) plan involving both the human and the robot.

Plan Preservation. But of course all removable sub-plans do not lead to serendipitous interventions, the robot has to ensure additional constraints that involve preserving the prefix of the original human plan and the world state he originally intended to be in after the intervention. A demonstration can be viewed at `http://bit.ly/1Q6tOBW`.

Humans as Collaborators

In this section we will look at the role of automated oversight in human computation tasks - specifically we investigate how an automated planner can contribute to a plan generation process carried out by humans. In order to collaborate with human planners, we postulate that automated planners must be able to perform two fundamental tasks -

Interpretation. In order for the planner to make useful contributions, it must first comprehend the current state of the planning process, as well as the intentions and (often implicit) preferences of the human planners.

Steering. The planner must iteratively refine, validate and critique the planning process, while respecting the authority of the human planners. The goal of this work is to augment the vast domain knowledge of humans with the superiority of automated constraint checking and plan validation.

Crowdsourced Planning - `AI-MIX`

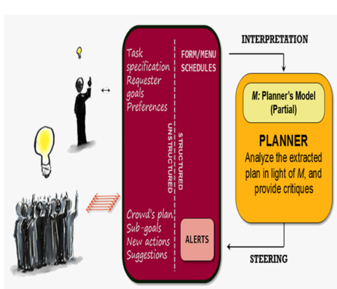

In [5] we investigate how planners can work with a crowd, using extremely shallow models (in tour planning domain). We show how a PDDL domain may be used to support and critique the plan generation process, and how declarative languages like ASP can be used to compile all the constraints to generate the final plan. Our system is the proud winner of the People's Choice Best System Demonstration Award at ICAPS'14. A demonstration of the system is available at `http://bit.ly/1o22ajr`.

Proactive Decision Support - `RADAR`

We are currently working on a system that can use more complete domain models to collaborate with experts on specialized tasks (like disaster response). An aspirational video is available `http://bit.ly/1KEJZKj`. I am more involved in the planning aspect of the project. Planning applications include plan recognition for preemptive assistance with both the plan generation and the proactive context-based information integration process, and plan validation and refinement using techniques like landmark generation and plan robustness. A preliminary demonstration of these capabilities can be found at `http://bit.ly/1o22ajr`.

Acknowledgments

I would like to take this opportunity to thank the contributions of several people who have made this work possible, including Kartik Talamadupula (IBM T.J. Watson Research Center) and Gordon Briggs (Tufts University) for their early work on coordination in the USAR environment, and David Smith (NASA Ames Research Center) and Matthias Scheutz (Tufts University) for their valuable inputs and support.

Also special thanks to all the funding agencies that have supported my work till now. My research has been funded in part by the ARO grant W911NF-13-1-0023, and the ONR grants N00014-13-1-0176, N00014-13-1-0519, N00014-15-1-2027 and N00014-15-1-2344.

REFERENCES

[1] T. Chakraborti, G. Briggs, K. Talamadupula, Y. Zhang, M. Scheutz, D. Smith, and S. Kambhampati. Planning for serendipity. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2015.

[2] T. Chakraborti, V. V. Meduri, V. Dondeti, , and S. Kambhampati. A game-theoretic approach to ad-hoc coalition formation in human-robot societies. In *AAAI 2016 Workshop on Multiagent Interaction without Prior Coordination (MIPC)*, 2016.

[3] T. Chakraborti, K. Talamadupula, Y. Zhang, and S. Kambhampati. A formal framework for studying interaction in human-robot societies. In *AAAI 2016 Workshop on Symbiotic Cognitive Systems (SCS)*, 2016.

[4] T. Chakraborti, Y. Zhang, D. Smith, and S. Kambhampati. Planning with resource conflicts in human-robot cohabitation. In *Proc. of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2016.

[5] L. Manikonda, T. Chakraborti, S. De, K. Talamadupula, and S. Kambhampati. AI-MIX: How a Planner Can Help Guide Humans Towards a Better Crowdsourced Plan. In *Innovative Applications of Artificial Intelligence (IAAI) 2014.Winner of People's Choice Best Demonstration Award at ICAPS 2014!*