On the Many Interacting Flavors of Planning for Robotics

Kartik Talamadupula[†] and **Matthias Scheutz**[§] and **Gordon Briggs**[§] and **Subbarao Kambhampati**[†]

[†]**Dept. of Computer Science and Engg.**
Arizona State University
Tempe AZ
{krt,rao} @ asu.edu

[§]**HRI Laboratory**
Tufts University
Boston MA
{mscheutz,gbriggs} @ cs.tufts.edu

Abstract

Automated planners have been employed in a variety of robotics applications. However, there still remains a distinct divide between task planning, or high-level planning, and its counterparts in robotics. In particular, navigation and dialogue planning have emerged as important concerns in the quest to make realistic end-to-end robotic systems a reality. However, in the absence of a unifying problem to solve, collaborations between these three fields have been sparse and mostly narrow and project-driven. In this paper, we discuss Human-Robot Teaming as that unifying problem, and outline via a simple example the various sub-fields of the different kinds of planning that interact naturally to produce a solution to the overall problem. Our hope is to spur the various, fragmented planning communities into further collaboration by highlighting the rich potential of these interactions.

1 Introduction

Automated planning systems have come a long way since the days of the STRIPS planner (Fikes and Nilsson 1972) and Shakey the Robot. Specifically, the evolution of fast heuristics and various compilation methods has enabled the application of state-of-the-art planners to cutting edge research problems in robotics. Planners – in one form or the other – now regularly guide robotic systems that hitherto had to rely on painstakingly pre-programmed scripts in a robust and real-time manner. The idea that robotic agents need to be endowed with autonomy is not new – from depictions in popular culture to actual deployed agents, robots are assumed to be autonomous and independent in many crucial ways. However, it is the meaning of this autonomy that is constantly changing. Much remains to be done is defining a full taxonomy for the usage of the word *planning* when it comes to robotic systems. Such an effort would go a long way towards recognizing the various disciplines and subfields that have hitherto been treated as separate from each other, and spur research on further bridging the divide between the automated planning and robotics communities.

Apart from automated planning – alternatively known as *task* planning – robotic systems have also had to contend with navigation or *path* planning. Increasingly, *dialogue* planning is gaining importance as well, reflecting the key role that speech and dialogue play in interaction with humans. All three of these fields house thriving research communities that report progresses at highly rated venues year after year. However, collaborations between these fields are few and far in between, and usually only come about as a result of integrated systems that are task or project specific.

In this paper, we discuss a motivating problem that brings these different types of planning together, and outline a simple example task that demonstrates the need for different kinds of planning. We conclude by pointing out various connections between existing work in the field of automated planning, and important and outstanding problems in other fields associated with robotics.

2 Planning for Human-Robot Teaming

Human-Robot Teaming scenarios are defined as those that involve humans working with autonomous robotic agents to achieve high-level goals that are determined and specified by a human (Talamadupula et al. 2011). Alternatively, any general problem that considers symbiotic interaction between humans and robots (Rosenthal, Biswas, and Veloso 2010) can be used to illustrate the point that we wish to make, too. Here we present a simple example of an HRT task.

A Motivating Example Consider a robot, *Cindy*, that must deliver a medical kit to *Commander Z*. Cindy is told that there is such a kit in a room at the end of the hallway, what the kit looks like, and instructed to remain undetected by the enemy while performing this task. Just before entering the room, Cindy encounters *Commander Y*, who asks her to void her earlier, more important goal in order to follow him. Cindy declines while indicating urgency and interruption in her voice, and negotiates a commitment to meet Commander Y wherever he happens to be when she achieves her current goal. Arriving outside the room where Commander Z is located, she senses that the door is closed, thus triggering a further query to her handler. Cindy is instructed to try a new action – pushing the door open with her hand. She enters the room and delivers the kit to Commander Z, who reinforces the commitment that she must go meet Commander Y at his current location at once.

Even in this simple task, various planning modalities must interact and occur in parallel to enable the script.

1. **Task Planning**: Agents must be able to plan for changing or conditional goals like the medical kit (Talamadupula et al. 2010), elaboration of the goals associated with the task (Baral and Zhao 2008) as well as trajectory constraints like 'remain undetected' on the form of the plan (Mayer et al. 2007). Additionally, the task planner may have to deal with updates to the model that are either learned, or specified by humans (Cantrell et al. 2012).

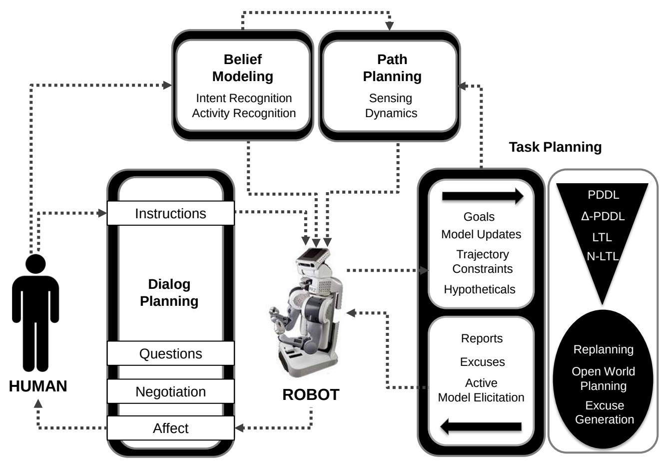

Figure 1: *A schematic of the various interactions present in a simple Human-Robot Teaming task.*

2. **Path Planning**: Autonomous robots must be endowed with capabilities of planning their paths. These may include planning with goal-oriented actions like looking for the medical kit (Simmons and Koenig 1995), finding the shortest path to the room that holds the kit (Koenig, Likhachev, and Furcy 2004), obeying constraints on the trajectories of the path (Saffiotti, Konolige, and Ruspini 1995) or planning for agents that exhibit different dynamics, like UAVs and AUVs (McGann et al. 2008).
3. **Dialogue Planning**: Robots need to skilled at both recognizing and producing subtle human behaviors vis-a-vis dialogue (Briggs and Scheutz 2013) – for example, in the above scenario, Cindy needs to both understand the superiority in Commander Y's voice when requesting a new task, as well as inflect her own response with urgency in order to indicate that the task at hand cannot be interrupted. Negotiation is another possibility, for which the robot needs to be informed by the task planner regarding excuses (Göbelbecker et al. 2010) and other hypotheticals.
4. **Mental Modeling**: The agent must be in a position to model the beliefs and mental state of other agents that are part of the scenario (Briggs and Scheutz 2012); in this case, Cindy may want to model Commander Y's mental state to determine his location at the end of the first task.
5. **Intent and Activity Recognition**: Closely tied in to both dialogue and mental modeling is the problem of recognizing the intents of, and activities performed by, other agents (Vail, Veloso, and Lafferty 2007). Humans are endowed with these capabilities to a very sophisticated degree, and agents that interact and team with humans must possess them as well.
6. **Architecture**: Finally, the integrated architecture that all these processes execute in plays a big role in determining the planning capabilities of the autonomous system. A good control structure must display programmability, adaptability, reactivity, consistent behavior, robustness, and extensibility (Alami et al. 1998). By dint of having to interact with humans, it must also fulfill the notions of attending and following, advice-taking, and tasking (Konolige et al. 1997). Finally, it must be able to detect and recover from failure, and tide all the other planning components over that failure.

3 Planning, and More Planning

As the above list shows, even a simple task requires various kinds of planning components in order to present humans with a seamless teaming experience. However, the mere presence of these components is not enough – they must interact in order to process data that comes in both from the human team-member as well as the world, so that the scenario objectives may be furthered. Figure 1 presents an outline of the various components described previously, and the interactions among them. Here, we look at three of the most important ones:

Task and Motion Planning The interaction between these two kinds of planning is well-established; in the above scenario, task planning sets the waypoints that must be visited in order to fulfill the higher level goals, and these waypoints are then passed on as goals themselves to the motion planning process. In return, the motion planning provides updates to the task planner on new objects in the world, and the current location of the robot.

Task and Dialogue Planning Real-world robotics applications are seeing an increase in the use of dialogue managers as more and more systems try to interact with and engage humans gainfully. On the one hand, dialogue systems provide more information and context to task planners in the form of instructions (goals), actions models and user preferences. In the reverse direction, task planners can be used to inform robotic agents about relevant questions to ask in order to elicit more information from humans, and to help determine the tone and affect of the interaction.

Mental Modeling and Task Planning Many human-robot teaming applications consider just the presence of one agent in the world; however, the real world is more closely approximated by multi-agent scenarios. Although a given agent will not interact with *all* of the agents in the world with the same degree of closeness, sometimes it is useful to model those other agents. In the example scenario described in Section 2, for instance, Cindy the robot must model Commander Y's mental state to some level in order to know his possible location when she achieves her prior goal, so that she may rendezvous with him. Similarly, given a model as well as the goals of Commander Y, a task planner can be used to simulate a mental model (in lieu of a real mental or belief model).

The challenge – and opportunity – for the automated planning community is to develop algorithms and systems that take all these interactions into account explicitly, in order to support more real-world robotics applications.

Acknowledgments

This research is supported in part by the ARO grant W911NF-13-1-0023, the ONR grants N00014-13-1-0176, N00014-09-1-0017, N00014-07-1-1049 and N00014-13-1-0519, and the NSF grant IIS201330813.

References

[Alami et al. 1998] Alami, R.; Chatila, R.; Fleury, S.; Ghallab, M.; and Ingrand, F. 1998. An architecture for autonomy. *The International Journal of Robotics Research* 17(4):315–337.

[Baral and Zhao 2008] Baral, C., and Zhao, J. 2008. Non-monotonic temporal logics that facilitate elaboration tolerant revision of goals. In *Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence, AAAI*, 13–17.

[Briggs and Scheutz 2012] Briggs, G., and Scheutz, M. 2012. Multi-modal belief updates in multi-robot human-robot dialogue interaction. In *Proceedings of 2012 Symposium on Linguistic and Cognitive Approaches to Dialogue Agents*.

[Briggs and Scheutz 2013] Briggs, G., and Scheutz, M. 2013. A hybrid architectural approach to understanding and appropriately generating indirect speech acts. In *Proceedings of the 27th AAAI Conference on Artificial Intelligence*, (forthcoming).

[Cantrell et al. 2012] Cantrell, R.; Talamadupula, K.; Schermerhorn, P.; Benton, J.; Kambhampati, S.; and Scheutz, M. 2012. Tell me when and why to do it!: Run-time planner model updates via natural language instruction. In *Human-Robot Interaction (HRI), 2012 7th ACM/IEEE International Conference on*, 471–478. IEEE.

[Fikes and Nilsson 1972] Fikes, R., and Nilsson, N. 1972. STRIPS: A new approach to the application of theorem proving to problem solving. *Artificial Intelligence* 2(3):189–208.

[Göbelbecker et al. 2010] Göbelbecker, M.; Keller, T.; Eyerich, P.; Brenner, M.; and Nebel, B. 2010. Coming up With Good Excuses: What to do When no Plan Can be Found. In *Proc. of ICAPS 2010*.

[Koenig, Likhachev, and Furcy 2004] Koenig, S.; Likhachev, M.; and Furcy, D. 2004. Lifelong Planning A*. *Artificial Intelligence* 155(1):93–146.

[Konolige et al. 1997] Konolige, K.; Myers, K.; Ruspini, E.; and Saffiotti, A. 1997. The saphira architecture: A design for autonomy. *Journal of experimental & theoretical artificial intelligence* 9(2-3):215–235.

[Mayer et al. 2007] Mayer, M. C.; Limongelli, C.; Orlandini, A.; and Poggioni, V. 2007. Linear temporal logic as an executable semantics for planning languages. *Journal of Logic, Language and Information* 16(1):63–89.

[McGann et al. 2008] McGann, C.; Py, F.; Rajan, K.; Thomas, H.; Henthorn, R.; and McEwen, R. 2008. A deliberative architecture for AUV control. In *Robotics and Automation, 2008. ICRA 2008. IEEE International Conference on*, 1049–1054.

[Rosenthal, Biswas, and Veloso 2010] Rosenthal, S.; Biswas, J.; and Veloso, M. 2010. An effective personal mobile robot agent through symbiotic human-robot interaction. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: volume 1-Volume 1*, 915–922.

[Saffiotti, Konolige, and Ruspini 1995] Saffiotti, A.; Konolige, K.; and Ruspini, E. H. 1995. A multivalued logic approach to integrating planning and control. *Artificial intelligence* 76(1):481–526.

[Simmons and Koenig 1995] Simmons, R., and Koenig, S. 1995. Probabilistic robot navigation in partially observable environments. In *International Joint Conference on Artificial Intelligence*, volume 14, 1080–1087.

[Talamadupula et al. 2010] Talamadupula, K.; Benton, J.; Kambhampati, S.; Schermerhorn, P.; and Scheutz, M. 2010. Planning for Human-Robot Teaming in Open Worlds. *ACM Transactions on Intelligent Systems and Technology (TIST)* 1(2):14.

[Talamadupula et al. 2011] Talamadupula, K.; Kambhampati, S.; Schermerhorn, P.; Benton, J.; and Scheutz, M. 2011. Planning for Human-Robot Teaming. In *ICAPS 2011 Workshop on Scheduling and Planning Applications (SPARK)*.

[Vail, Veloso, and Lafferty 2007] Vail, D. L.; Veloso, M. M.; and Lafferty, J. D. 2007. Conditional random fields for activity recognition. In *Proceedings of the 6th international joint conference on Autonomous agents and multiagent systems*, 235. ACM.