Human-Robot Communication in Automated Planning

Aleck MacNally
amacnally@student.unimelb.edu.au
University of Melbourne
Parkville VIC 3010

Abstract

With the continued integration of artificial intelligence and robotics into human society the need for human-aware and human-in-the-loop planning becomes ever more prominent. So far research has been conducted which addresses acting in a human world to achieve joint or independent goals or to assist humans in completing their own goals. However, little research has been published that pertains to human-robot communication whilst planning. A robot's ability to communicate efficiently and effectively with humans will allow the robot to be more useful to the human, in that the human may extract necessary information, or understand the actions that the robot may perform, as well as making the robot more efficient at achieving its own goals. With these benefits in mind this doctoral research will address when, why and how information is conveyed to humans by robots during planning s well as integrating into the planning process an ongoing negotiation between a human and the planner (a robot)

Introduction

Human-robot interaction presents unique challenges in the area of Automated Planning. Research into such problems as producing safe and predictable plans, inferring a human's goals from their actions and robot-human communication is on-going.

It is important when dealing with any form of human-robot interaction and human-aware planning that we understand the scenario between the robot and the human (Shah et al. 2011; Talamadupula et al. 2014; Chakraborti et al. 2015a; Levine and Williams 2014). These scenarios can be categorized as follows:

- The human is omnipotent
- The robot is omnipotent
- The robot and human plan separately but for a joint goal
- The robot and human plan separately for independent goals but must interact due simply to being in the same environment.
- The robot may adjust their goals to assist humans

The manner in which the challenge of communicating between humans and robots is addressed, will differ depending on which scenario the system is in. For instance when the human is omnipotent no information needs to go from the robot to the human. Whereas when humans and robots must work together to achieve a joint goal the communication may include negotiations with the human to produce joint plans which are less costly than working independently.

This doctoral research will address communication in two different areas of human-robot interaction. The first area is the communication of plans and goals to a human. This is important both when humans and robots work independently as well as when they share a joint goal. The second area of research is in negotiations between humans and robots, which is important when robots and humans work toward a joint goal.

Related Work

Shah et al. implemented the CHASKi executive which focused on a joint goal in a human-robot team. The research took inspiration from human-human interaction and used communication in a limited way to indicate when each agent was committing to a subtask or had completed one. CHASKI also allowed humans to issue requests to the robot.

Chakraborti et al. implemented a system which used the idea of resource profiles which temporally tracked the usage of specific resources. They constrained their resulting plan to minimize the overlap between these profiles. In this framework agents were able to negotiate for resources when there was no feasible plan. Similarly, negotiations were addressed by Karpas et al. who introduced temporal uncertainty to the PIKE executive (Levine and Williams 2014) by negotiating with the user to relax temporal constraints.

Unhelkar and Shah put forward CONTACT, a system handling multi-agent settings that decides whether an agent should communicate information based upon a reward system and an approximate idea of what is known by all other agents in the system.

Implicit human-robot communication has been studied by Zhang, Zhuo, and Kambhampati who used Machine Learning techniques to label the degree of "social acceptability" of plans

Human-Robot Communication

In this extended abstract two aspects of human-robot communication will be addressed. First, the communication of a

---

robots plans to the humans involved in the system. Without loss of generality, goals shared between humans and robots can be deemed to be independent. Second, the use of communication to negotiate between humans and robots, which will directly relate to humans and robots completing joint goals.

Communicating Plans

As robots become more prevalent in human society it will become necessary to ensure that they act in a manner which is understandable and predictable to the humans with whom they are cohabiting (Alterovitz, Koenig, and Likhachev 2014). To this end it is important that a plan which may seem strange to a human is explained properly by the robot agent. Stubbs, Wettergreen, and Hinds underlined this need when they found that humans working with robots were less productive when the autonomy of the robots was increased and their 'common ground' (shared knowledge) decreased. This introduces three challenges: which plans require explanation, when is an agent allowed to communicate their plan and what does an agent communicate when communicating a plan?

The first question follows from the assumption that a robot communicating all their plans to a human would become a nuisance. Therefore it is necessary that only plans which require explanation be explained. Once the robot has selected a plan to explain, this explanation must be heard and must not overlap any communication currently occurring and must also not happen when the human is completing a task which requires concentration such as driving. This is covered by the second question. The last question addresses how much information is necessary to convey to a human agents once a plan is to be explained. It is assumed that the optimal amount of information required to fully explain the plan to the human is the minimal amount, relating to the assumption from the first question.

Determining whether to communicate a goal It is important that a robot's goals are not ambiguous. This means that a human should be able to infer a robot's goals from its actions. To that end it is considered preferable to have a robot which acts obviously toward a single goal.

In many cases the nature of the environment and model of the robot result in plans which are optimal yet ambiguous. This ambiguity is formalised in the concept of the worst case distinctiveness (wcd). The wcd represents the size of the largest shared prefix of two or more optimal plans leading to different goals (Keren, Gal, and Karpas 2014). The larger the wcd, the larger the ambiguity inherent in the environment and robot's model. When the wcd of the model is too large it is necessary for the robot to communicate their goal to the human.

The idea of our approach is that a human is watching a robot, but not attempting to achieve any goals of its own. The human uses goal recognition as set out by Ramirez and Geffner in order to infer the goal of the robot. The robot uses the wcd of the system to determine if the human will infer its goals correctly or if it will need to communicate its goals. We formalize this idea below.

The robot is given a planning problem $\Pi =< F, I, A, G >$, and has a set $\mathcal{G}$ of the possible goals, where $G \in \mathcal{G}$. The human in this case has knowledge of $\mathcal{G}$ as well as Π except for G. The human will not plan itself but will reason using goal recognition. Using the framework published by Keren, Gal, and Karpas it is possible to find the worst-case distinctiveness (wcd) of the problem. It is put forward that if the $wcd > \alpha$, then for any plan π the robot must communicate G to the human, otherwise G is deemed obvious enough to the human from π. The quantity α is yet to be established.

A more computationally intensive solution to this problem is that as we know the current goal of the robot, G, we can calculate:

$$wcd_G = \max_{G' \in \mathcal{G}, G' \neq G} wcd(\Pi, \{G, G'\})$$

where $wcd(\Pi, \{G, G'\})$ refers to the wcd of the problem with the possible goal set consisting of just G and G'. The wcd_G represents the worst case distinctiveness between each goal and G. The wcd_G has wcd as an upper-bound, so this method will never force the robot to communicate more than in the above method. Using wcd_G instead of wcd will reduce the amount of times that a robot will have to communicate with the human, however, it will require the robot to perform $O(|\mathcal{G}|)$ calculations of a wcd.

To improve upon this idea it would be advantageous to model the human belief in the goal of the robot as well as the human's own goals. Doing this would allow one to reason more accurately as to when the human might need the knowledge about what goal the robot is pursuing in order to pursue their own goals.

Determining whether to communicate a plan Given a situation in which the human has knowledge of the goals of the robot, it is preferable that a plan which is to be executed is predictable and familiar to the human. If it is not predictable or familiar, the plan must be communicated. Ensuring this will make robots more trustworthy to humans. The degree to which a given plan is surprising will be determined by combining in some suitable way the wcd measures for every sub-goal achieved by the plan. This will be addressed in future research.

Determining when an agent should communicate Once a robot has found it necessary to communicate a plan due to either the ambiguity in the robots goals or in its plan, the problem becomes determining when to communicate the plan. In particular in problems which relate to communicating over a channel which is being used by multiple agents and is affected by external sources. Assuming an axiomatization of communication actions over a channel as well as the ability to observe the state of the channel, this problem can be modeled as an MDP in which the state of the channel is modeled as observations and an agent can choose to communicate with respect to these observations. this idea will be explored in further research.

The second problem in this section is making sure that a communication action does not interrupt a critical human task. Such task are those which would jeopardize the safety of humans if they were interrupted such as driving or surgery. A simple solution to this problem is to have a human action which locks communication and a human action which unlocks communication. A second solution is to use a temporal multi-agent planner with mutual exclusion between actions which require concentration and robot communication actions. This topic will also be addressed in future research.

Determining the minimum amount of information to communicate The last question which will be addressed in this section is "how much information in a robots plan needs to be communicated?". A robot's goals will consist of a conjunction of logical statements. When these goals are created by a human they can be named easily, and then this name can be produced when it is determined that the goal is to be communicated. In the case where a robots goals are automatically generated such as in (Chakraborti et al. 2015a), the information within these goals must be summarized in a way which is human understandable. This is equally true for a plan which needs to be communicated.

A possible scheme for determining which actions of a plan need to be communicated to a human is described next. Given that a robot has an optimal plan π for a goal G, where the human has knowledge of G, we suggest that the optimal amount of information that needs to be communicated is the minimal subsequence $\pi_m \subset \pi$ (that is minimal in length not cost) which will render $P(G|\pi_m) = 1$. It is assumed that if from this π_m the human can calculate that the goal is G than, even if the plan is not human-intuitive, it is understood by the human. This will be addressed in future research.

Negotiating

Negotiations between humans and robots has already been attempted in various ways such as in Talamadupula et al. where negotiations were made over resources and in Karpas et al. where time was negotiated with. In this section we utilize the idea of negotiation to allow a problem to be solved by a joint human-robot team where the model of the robot is known but the model of the human is not. Using this idea the robot constructs a partial model of the human throughout the negotiations.

The idea of this approach is that an agent must solve a problem optimally and use the assistance of a human with a fixed communication penalty rate. To this end the agent assumes that the human may achieve any fact and after each iteration of communication, the robot updates the model of the human and then re-plans until there is a satisfying plan that both the human and robot agree is possible.

For a classical planning problem $\Pi =< F, I, A, G >$, where A is the robot's action set, we compile the model of the human into the robots model by augmenting it as follows. Let us consider the initial planning problem for the first iteration: $\Pi' =< F, I, A_0, G >$, where:

$$A_0 = A \cup \{a_h^f, a_h^{\neg f} | f \in F\}$$

where a_h^f is an action with:

- $prec(a_h^f) = \{\neg f\}$
- $add(a_h^f) = \{f\}$
- $del(a_h^f) = \varnothing$
- $cost = \lambda \times \sum_{a \in A} cost(a)$

Where $a_h^{\neg f}$ is defined similarly. a_h^f encodes the assumption that the human h can make f true. With these additional actions the robot can achieve all facts in the problem. The λ is a constant which encodes a fixed penalty for asking for help. A high λ would produce and scenario where the robot is expected to work on its own and to only seek help when it could not find a satisfying plan. If it is intended that the robot and human coordinate more than λ should be set lower.

Following the initial stage of planning the robot presents the human with the actions that she must complete (if there are any). She may then do one of two things: accept the plan as it is or negotiate with the robot. If she accepts then the problem is solved. If she chooses to negotiate, she presents the robot with additional knowledge about her model. For instance she may say that she cannot complete a_h^f but can complete $a_h^f\prime$ where $a_h^f\prime$ is equal to a_h^f with a different cost. She could also provide the robot with her entire set of actions if she chose to. With this framework the human may give as little or as much information to the robot as the human wishes. When the human has decided what information she is going to impart to the robot, she places the actions she cannot complete in A_j^{imposs}, this would include a_h^f from the example above, and all the actions she can complete in A_j^{poss}, which would include $a_h^f\prime$, where j is the current iteration of negotiation. The robot's action set is updated as follows:

$$A_j = \{A_{j-1} \setminus A_j^{imposs}\} \cup A_j^{poss}$$

The robot will then re-plan with the new set of actions A_j and negotiate until the human accepts the plan, or the robot cannot find a plan.

It is important to note that with this approach the model of what the human can do is not required. The robot assumes that the human can do anything and each turn the robot updates its understanding of the human's model with partial information given to it by the human

This approach will produce an optimal plan for the robot and for what the robot believes the human can do, however without a complete model for what the human can do an optimal plan for both agents cannot be found.

Conclusion

This extended abstract addressed two areas of communication between humans and robots. The first was the problem of communicating plans and goals to humans either in the same environment or to those completing the same goals. This involved three distinct challenges which were which plans and goals should a robot communicate, when should a robot communicate them and how much information was required to communicate a goal or plan? A solution to the first question was put forward which involved the computation of the worst case distinctiveness of the system (Keren, Gal, and Karpas 2014). The second area of research in this paper addresses negotiation. A scheme for negotiations between a robot and human was put forward which involved compiling a partial human action model and a planning problem into a classical planning problem which was then solved. The solution was presented to the human and then the partial human model was update. This process continues until a solution is found.

References

Alterovitz, R.; Koenig, S.; and Likhachev, M. 2014. Robot planning in the real world: research challenges and opportunities. *National Science Foundation*.

Chakraborti, T.; Briggs, G.; Talamadupula, K.; Zhang, Y.; Scheutz, M.; Smith, D.; and Kambhampati, S. 2015a. Planning for serendipity. *Intelligent Robots and Systems (IROS)* 5300–5306.

Chakraborti, T.; Zhang, Y.; Smith, D.; and Kambhampati, S. 2015b. Planning with stochastic resource profiles: An application to human-robot cohabitation. *International Conference on Automated Planning and Scheduling (ICAPS)*.

Karpas, E.; Levine, S. J.; Yu, P.; and Williams, B. C. 2015. Robust execution of plans for human-robot teams. *International Conference on Automated Planning and Scheduling (ICAPS)* 342–346.

Keren, S.; Gal, A.; and Karpas, E. 2014. Goal recognition design. *International Conference on Automated Planning and Scheduling (ICAPS)*.

Levine, S. J., and Williams, B. C. 2014. Concurrent plan recognition and execution for human-robot teams. *International Conference on Automated Planning and Scheduling (ICAPS)*.

Ramirez, M., and Geffner, H. 2009. Plan recognition as planning. *International Joint Conference on Artificial Intelligence (IJCAI)* 1778–1783.

Shah, J.; Wiken, J.; Williams, B.; and Breazeal, C. 2011. Improved human-robot team performance using chaski, a human-inspired plan execution system. *Proceedings of the 6th international conference on Human-robot interaction* 29–36.

Stubbs, K.; Wettergreen, D.; and Hinds, P. H. 2007. Autonomy and common ground in human-robot interaction: A field study. *Intelligent Systems, IEEE* 22(2):42–50.

Talamadupula, K.; Briggs, G.; Chakraborti, T.; Scheutz, M.; and Kambhampati, S. 2014. Coordination in human-robot teams using mental modeling and plan recognition. *Intelligent Robots and Systems (IROS)* 2957–2962.

Unhelkar, V. V., and Shah, J. A. 2016. Contact: Deciding to communicate during time-critical collaborative tasks in unknown, deterministic domains. *Association for the Advancement of Artificial Intelligence (AAAI)*.

Zhang, Y.; Zhuo, H. H.; and Kambhampati, S. 2015. Plan explainability and predictability for cobots. *arXiv preprint arXiv:1511.08158*.