Using Symbiotic Relationships with Humans to Help Robots Overcome Limitations

Stephanie Rosenthal
Computer Science Dept.
Carnegie Mellon University
Pittsburgh, PA, USA
srosenth@cs.cmu.edu

Manuela Veloso
Computer Science Dept.
Carnegie Mellon University
Pittsburgh, PA, USA
veloso@cs.cmu.edu

ABSTRACT

We are interested in task-driven robots in our environments that can communicate with humans. While today's robots often communicate with humans to overcome their limited perception and execution, the relationship between humans and robots is often one-sided in which the human is providing all the help to the robot without their own benefits. Instead, we propose a *symbiotic relationship* in which the robot performs tasks for humans and only ask for help to complete the task successfully. The symbiotic relationship is a more balanced one in which the robot and human mutually benefit each other through their actions and help. We introduce the Visitor-Companion Task for a robot to accompany a human visitor to meetings throughout the day as an example of the relationship and our robot, CoBot, that implements the task. We discuss both the planning requirements and benefits for a robot in a symbiotic relationship as well as the benefits and limitations of the humans.

Categories and Subject Descriptors

H.4 [**Information Systems Applications**]: Miscellaneous

General Terms

Experimentation

Keywords

Human-robot/agent interaction

1. INTRODUCTION

While we aim for robots to perform autonomous tasks in close proximity to humans (*e.g.,* robots in our offices [9] or malls [10]), currently these robots may not be capable of completing all actions successfully due to limitations in perception and execution. To overcome robot limitations and improve robot performance, we take advantage of human knowledge and expertise by requesting help from humans available in the environment. We propose *symbiotic relationships* in which robot perform autonomous tasks for humans and may ask for help from people in the environment. The symbiotic relationship is a more balanced relationship compared to previous human-robot relationships and requires that 1) the robot asks during its task to help humans, 2) the help that the robot requests is within human capability without special instruction and 3) the robot gives incentive for the robot to answer through the benefit of its own task.

Figure 1: The CoBot Visitor-Companion Robot, designed and built by Mike Licitra

To illustrate the symbiotic relationship, we contribute the Visitor-Companion Task (see [11] for complete details). Because the companion problem requires that a human be present near the robot for a majority of the time, it offers the flexibility of the robot proactively requesting assistance from the visitor or other humans when needed. The robot helps the visitor navigate to each meeting without getting lost and can perform other tasks for the visitor. The visitor in can answer questions to help the robot overcome its limitations (*e.g.,* tell the current location) or physically help the robot (*e.g.,* lift a coffee cup). The visitors actions satisfy the robots' subgoals which in turn satisfy the shared goals of both the human and robot. The help mutually benefits the robot, which can now complete the task, and the human when the request is accomplished or expectation is satisfied.

We have implemented this task on a robot, CoBot, that is both limited in its ability to sense its location in the environment or to perform other tasks the visitor might request (it has no arms to lift objects or open doors). We describe the symbiotic relationship in detail with examples from the Visitor-Companion Task and discuss the benefits and and requirements of the robot in order to implement the relationship with enough benefit to the humans from which it requests help using examples from our experiences.

2. SYMBIOTIC RELATIONSHIPS

Many robots require specific human supervisors to be continuously monitoring its progress to take control or direct the robot whenever an error occurs (*e.g.,* robots using collaborative control [7] and sliding autonomy [4][8]). This supervisory relationship is extensively explored in the context of Urban Search and Rescue (USAR) by Yanco *et. al* who find that this relationship can actually be detrimental to the success of the a task when the supervisor is unfamiliar with the interfaces or becomes disoriented and cannot give accurate commands [15]. Unlike systems in which the robot can seek assistance or confirmation from humans, symbiotic agents are autonomous and do not control or direct each others' actions in any way. All agents can take these actions to achieve the goals of the team, and coordinate through *synchronous* communication actions to request and provide help to team members.

The agents in symbiotic relationships benefit each other by requesting and receiving help on actions they could not have performed alone due to lack of *capabilities*, coordinating their actions only when they need help. The help can come in two forms:

- an agent performs an action for another (*e.g.,* socially embedded learning [3] in which the human escorts the robot to the desired location)
- an agent increases another's capability to complete the action either through learning or explaining state information(*e.g.,* learning by demonstration [2] in which a human tells the robot which state they are in or which action to take)

While the robot could learn to perform actions for which it has or can learn capabilities, we do not expect any robot to be able to complete all actions. For example, a robot without arms cannot ever lift a cup of coffee. Using our formalism for planning around these capabilities, if a robot does eventually acquire arms it can simply stop asking for help by updating the robot's capabilities.

Finally, because there may be many possible plans that achieve the same goals, the agents assign costs to their state (*expectations*) which all the agents can use when evaluating the best actions. When the agents take actions that affect each other, they take actions to minimize cost of each others' state while achieving the goal, further benefitting the group. This relationship is in contrast to those in which the human or the robot is responsible for helping the other without benefit in return.

3. STATES AND ACTIONS

Symbiotic agents do not control or direct each others' actions in any way. Instead, all agents can take actions to achieve the goals, and coordinate through *synchronous* communication actions to request and provide help to other agents in the environment. In the Visitor-Companion Task, although both the robot and human have the same goal to attend all meetings, they are not performing joint actions together. When possible, the robot acts autonomously, allowing the human to follow it to their meetings, and performs actions to satisfy both the visitor's and its own goals.

While the robot maintains state mostly about itself, it also maintains some state about the visitor in order to evaluate the visitor's *expectations* when deciding which of its actions it should take. We divide the actions into categories - asynchronous (Execute, Inform, Ask, Request) and synchronous (Respond, Notify) (see [11] for full details). While the asynchronous actions can happen whenever the preconditions are met, the synchronous actions require another communication action be performed before they can be invoked and affect the state of the visitor. Both humans and the robot can perform both asynchronous and synchronous actions, asking/requesting and offering help to benefit each other.

Asynchronously, the robot can inform visitor about different locations such as labs that might be of particular interest. The robot can `move` past these different locations around the building using the `nav-target` state to maintain knowledge of where it is going. This autonomous action as well as `open-door` and `put-coffee` include a *capability* or probability of success based on the robot's uncertainty (discussed later), which can result in either a success or failure. Based on the failure (*e.g.,* localization error from `move`), the robot can ask a human nearby for help. When the visitor responds to a location question, the robot processes the response and updates its location information to continue moving. Otherwise, the robot waits for the action to be taken, updates its state, and continues with its plan.

4. CAPABILITIES

The agents in symbiotic relationships benefit each other by requesting and receiving help on actions they could not have performed alone due to lack of *capabilities*, coordinating their actions only when they need help. The robot may have limitations, either due to state uncertainty or physical limitations which may cause some of its actions to fail. The visitors do not know their way around the building and do not know how the robot works to help it all the time.

4.1 Robot Capabilities

Unlike methods which require training before the robot can be deployed (*e.g.,* learning by demonstration [2] and socially embedded learning [3]), the robot asks for help while it is performing a task autonomously and either does not have any ability to perform an action or cannot determine which action to take. A robot can ask a human for help when it cannot determine its own state with certainty, and therefore cannot determine what action to take. Specifically, the robot asks a clarification question to determine which state it is in. Because it acts autonomously, when its state is known, it can plan (or replan) its actions and execute them without further human intervention. Unlike learning by demonstration in which a robot does not know what action to take, the robot is given a policy a priori but does not know which state it is in to take the action.

While the robot could learn to distinguish states through asking questions, we do not expect any robot to be able to complete all actions. In symbiotic relationships, a robot can ask for human help when its plan to complete a task requires actions it cannot perform itself. For example, a robot without arms cannot ever lift a cup of coffee. If the robot's task requires it to bring coffee to someone, it must ask a human in the environment to pour the coffee and put it on the robot.

In order to model these limitations, some actions have both success and failure effects that happen according to capabilities - the probability of success p. If there is no chance of completing an action, $p = 0$. For example, if the robot does not have arms, there is no change it could perform `open-door` or `put-coffee` itself. These actions will always result in failure and the robot will always request help from a human near the coffee maker with action `(ask` giveCoffee). When $p > 0$, the robot may not complete an action successfully due to the uncertainty in the robot models. For example, in the move action, the robot may be uncertain of its location which contributes its successful completion.

The ability of a robot to ask these kinds of questions significantly extends its abilities. While a robot's tasks have been previously limited by its physical structure, it can now accomplish many more by asking for help. Additionally, a robot can more successfully complete tasks by reducing its state uncertainty. For example, CoBot asks for help when it cannot determine its location with high certainty (more details in [11]). While CoBot can navigate autonomously, with help from a visitor it is able to navigate to goals nearly 10% faster than autonomous navigation. Because the questions reduce its uncertainty, it misses fewer turns and stops fewer times to replan its path to the goal.

4.2 Human Capabilities

While a robot can ask for help when it lacks capabilities, it should also recognize human capabilities and expectations. While human supervisors are very accurate when helping the robot, but they cannot participate in the environment during monitoring and are very time-expensive to require as we scale up the number of robots in the environment. We would like any human in the environment who does not know about how the robot works to answer as accurately as possible under these conditions. However, humans in the environment may not always be accurate when answering questions. They may not understand the question or they may not know the answer. They may be busy and unable to spend a lot of time helping the robot. The robot may be asking a question related to something that happened in the past and the human might not have been aware of the environment at the time. Recently, much research has focused both on modifying interfaces to make it easier for supervisors to understand the feedback that is required (*e.g.,* [1][13]) and by taking into account inaccuracies in feedback to make the robots more robust (*e.g.,* [5]).

In order for humans to answer questions as accurately as possible, a robot will need to *ground* them in its frame of reference before asking for help (*e.g.,* [6][14]). We performed experiments with two different robot tasks to understand what kinds of information a robot might need to tell a human when it asks for help in determining its state - a shape recognition task on a Wizard of Oz'd robot and the localization talk on CoBot performing the Visitor-Companion Task. In these studies, we tested whether each combination of four kinds of robot state information affected the accuracy of the responses humans gave: a description of the context of the robot, its prediction of its state, its uncertainty in that prediction, and an additional question asking humans to further describe the state (described further in [12]).

The results of our initial shape recognition task show that providing users all of this information together (context, prediction, uncertainty, and asking an extra question) improved the accurate of their responses to the robot the most. However, because the shape recognition robot was Wizard of Oz'd, we tested users' responses when CoBot asked where it was (*e.g.,* "Can you point to where we are on this map of the building") while taking a tour of the building. Participants in this study were randomly assigned to one of five conditions: 1) no state information, 2) uncertainty and context, 3) uncertainty and prediction, 4) uncertainty, context, and prediction, and 5) uncertainty, context, prediction, and extra question (the more accurate from the first study). Though the experimenter remote-controlled the robot to each location in the building and triggered the appropriate responses, the participants believed the robot was moving autonomously. After participants completed the 15-minute tour containing 13 questions, they were given a survey about their experiences with the robot.

We collected the clicks on the map for each participant and calculated the Euclidean distance from the clicks to the actual robot location. Because the distribution of these distances was not a Normal Distribution, we performed a log transformation to normalize the data. We, then, analyzed the results of the localization test of log distances with a mixed model with participant ID as a random effect and the question condition as a fixed effect analyzed using the F statistic. Our results show there are statistically significant differences between the five question conditions ($F(4,38.53) = 3.93$, $p < 0.001$). We used Contrasts to analyze the whether there were statistically significant differences between our guidelines from the shape recognition study and the other four conditions tested. Running four contrasts means significance is measured at $.05/4 = 0.0125$.

Although we analyzed the log distances, we report the true distances in meters for clarity (Figure 2). Participants who received no state information clicked further away from the robotÕs true location (4.5 meters) compared to those two received all four kinds of state information (1.65 meters) ($F(1,38.45) = 22.17$, $p < 0.001$). Participants who received only uncertainty and context or uncertainty and predictions clicked 2.76 and 2.74 meters respectively from the true location, a marginally significant difference compared to the four kinds of state information ($F(1,38.9) = 3.18$, $p = 0.082$) ($F(1.,38.7) = 3.78$, $p = 0.059$). While the participants with all four kinds of state information show a 1 meter improve-

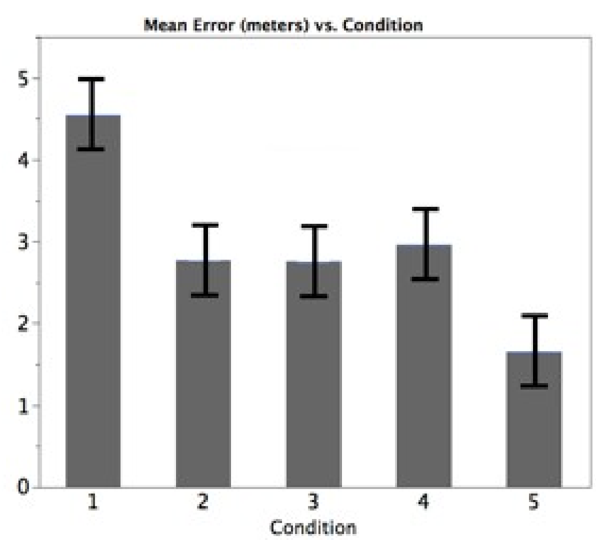

Figure 2: Our results show that providing context, prediction, uncertainty, and an additional question about context caused participants' responses to be most accurate compared to other combinations of the state information.

ment to these two conditions, there was a larger range of clicks leading to only marginal significance. Finally, participants who received context, uncertainty, and prediction clicked significantly further from the true location (2.94 meters) than those with all four kinds ($F(1,37.6) = 8.17$, $p < 0.001$).

While the robot can increase the accuracy capabilities of humans in the environment using our findings, it does not take into account the expectations of the human about the behaviors of the robot.

5. EXPECTATIONS

In symbiotic relationships, because there may be many possible plans that achieve the same goals, both humans and robots assign costs to their state (*expectations*) which all the agents can use when evaluating the best actions. When the agents take actions that affect each other, they take actions to minimize cost of each others' state while achieving the goal, further benefitting the group. This relationship is in contrast to those in which the human or the robot is responsible for helping the other without benefit in return. For example, humans in the environment would stop answering the robot if it asked too many questions without providing the human any benefit.

In order to understand how the symbiotic relationship (specifically, the robots' questions) affects the visitors in the Visitor-Companion Task, we invited five participants to participate in a four-meeting schedule. The participants were true visitors and had never been in the building before. Participants were told that the CoBot could assist them in the following ways on the way to their meetings:

- bring drinks to meetings
- providing additional information about meeting hosts (by displaying the host's website)

For each participant, CoBot gave the same information in each variation of the condition and asked the roughly the same number of questions about where the robot was in the building. After the visitors finished their schedule, we asked them to rate the robots' usefulness in its abilities as well as the number of questions the robot asked as it navigated (too many to too few).

While participants mostly felt the robot could have asked fewer questions, they had different opinions about how many were too many - reflecting different costs associated with the questions. When we combine the robot's abilities and the questions into a complete experience, we found that four out of five participants said they benefitted from the navigation guidance and other assistance and would use CoBot again, even thought they felt the robot asked them for help too many times. The one participant who would not use it again placed high cost on asking for help and said he would use it again if it asked fewer questions.

While the robot only necessarily requires state and actions in order to complete a task, we find that the robot should maintain an understanding of human expectations in order to choose which actions to take. Using these expectations, the robot can choose the best plan, the best action, or the best time to take an action that minimizes the cost to the visitor. For example, the visitor may not want to be late to any meeting and thus the robot would incur a high cost if it takes actions that result in the visitor being late. Additionally, while the robot may not always be able to avoid asking for help, it can ask raise the threshold of how uncertain it is to avoid asking questions if it may be able to perform the action itself. We believe that acting based on expectations could increase the usability of the robots and the willingness of humans to help them over long periods of time.

6. CONCLUSION

In this work, we contribute a robot agent capable of being in a symbiotic relationship with humans. We introduce the Visitor-Companion Task as an example where a human visitor and a companion robot have joint goals and coordinate as a team, but interact asynchronously. The robot can ask for help from the human to overcome some of its limitations (*e.g.,* possible location confusion or need to open a door). The symbiotic relationship is a more balanced one in which the robot is expected to help the humans and provide incentive for answering its questions. We discuss both the benefits for a robot in a symbiotic relationship as well as the benefits and limitations of the humans.

7. REFERENCES

[1] B. Argall, B. Browning, and M. Veloso. Learning by demonstration with critique from a human teacher. In *International Conference on Human Robot Interaction (HRI'07)*, pages 57–64, 2007.

[2] B. Argall, S. Chernova, M. Veloso, and B. Browning. A survey of robot learning from demonstration. *Robotics and Autonomous Systems*, 57(5):469–483, 2009.

[3] H. Asoh, S. Hayamizu, I. Hara, Y. Motomura, S. Akaho, and T. Matsui. Socially embedded learning of the office-conversant mobile robot jijo-2. In *International Joint Conference on Artificial Intelligence*, volume 15, pages 880–887. Citeseer, 1997.

[4] D. J. Breummer, D. A. Few, R. L. Boring, J. L. Marble, M. C. Walton, and C. W. Nielsen. Shared understanding for collaborative control. *IEEE Transactions on Systems, Man, and Cybernetics*, 35(4):494–504, 2005.

[5] S. Chernova and M. Veloso. Learning equivalent action choices from demonstration. In *International Conference on Intelligent Robots and Systems (IROS 2008)*, pages 1216–1221, 2008.

[6] H. Clark. Talking as if. In *Proceedings of the 3rd ACM/IEEE international conference on Human robot interaction*, pages 393–394, 2008.

[7] T. W. Fong, C. Thorpe, and C. Baur. Robot, asker of questions. In *Robotics and Autonomous Systems*, volume 42, No. 3-4, pages 235–243, 2003.

[8] F. Heger, L. Hiatt, B. Sellner, R. Simmons, and S. Singh. Results in sliding autonomy for multi-robot spatial assembly. In *8th International Symposium on Artificial Intelligence, Robotics and Automation in Space (iSAIRAS)*, pages 5–8. Citeseer, 2005.

[9] J. Hong, Y. Song, and S. Cho. A hierarchical bayesian network for mixed-initiative human-robot interaction. In *IEEE International Conference on Robotics and Automation*, volume 4, pages 3808–3813, 2005.

[10] T. Kanda, M. Shiomi, Z. Miyashita, H. Ishiguro, and N. Hagita. An affective guide robot in a shopping mall. In *Proceedings of the 4th ACM/IEEE international conference on Human robot interaction*, pages 173–180. ACM, 2009.

[11] S. Rosenthal, J. Biswas, and M. Veloso. An effective personal mobile robot agent through a symbiotic human-robot interaction. In *Proceedings of the 9th international joint conference on Autonomous agents and multiagent systems*, 2010.

[12] S. Rosenthal, A. K. Dey, and M. Veloso. How robots' questions affect the accuracy of the human responses. In *Proceedings of The International Symposium on Robot - Human Interactive Communication*, pages 1137–1142, 2009.

[13] M. Shiomi, D. Sakamoto, K. Takayuki, C. T. Ishi, H. Ishiguro, and N. Hagita. A semi-autonomous communication robot: a field trial at a train station. In *International conference on Human robot interaction (HRI'08)*, pages 303–310, 2008.

[14] L. Steels. Evolving grounded communication for robots. *Trends in Cognitive Sciences*, 7(7):308–312, 2003.

[15] H. Yanco, J. L. Drury, and J. Scholtz. Beyond usability evaluation: analysis of human-robot interaction at a major robotics competition. *Human-Computer Interaction*, 19(1):117–149, 2004.