Hello? Is Someone in this Office Available to Help Me?

Proactively Seeking Help from Spatially-Situated Humans

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
mmv@cs.cmu.edu

Anind K. Dey
Human-Computer Interaction
Carnegie Mellon University
Pittsburgh PA USA
anind@cs.cmu.edu

ABSTRACT

Robots are increasingly autonomous in our environments, but they still must overcome limited sensing, reasoning, and actuating capabilities while completing services for humans. While some work has focused on robots that proactively request help from humans to reduce their limitations, the work often assumes that humans are always available to help. In this work, we propose a model for task-embedded robot navigation that includes the people who can help the robot and benefit from the robot's services - those who are assigned static locations in the environment, in particular offices. These occupants have different challenges compared to traditional helpers such as teleoperators in that they are not always available to help and they are spatially-situated and therefore physically cannot help in every location.

Categories and Subject Descriptors

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

General Terms

Experimentation

Keywords

Human-Robot Interaction, Asking for Help, Planning

1. INTRODUCTION

Robots are becoming increasingly autonomous in their ability to perform services for us in our environments. They can give visitors directions in malls [11] and tours in museums [12], and act as companions for individual users [9]. Despite these great strides, robots are still not ubiquitous as they have sensing and actuation limitations that can affect their task performance. For example, many robots have difficulty recognizing speech in loud or busy environments, recognizing objects or obstacles to avoid while navigating, and may not have the physical ability to manipulate objects in the environment.

To overcome these limitations, some work has focused on reasoning about the robot's current state and proactively requesting help from humans to correct predictions (*e.g.,* in speech) and direct the robot's action if necessary during tasks. However, this work has been limited to asking for only one kind of help at a time. Additionally, whether the human helper is a supervisor that typically assigned the robot its task and therefore have high incentive for it to perform [1, 6] or a passer-by willing to help a robot even when they are not receiving services from it [14], there is an assumption that humans are almost always available to help the robot and can help the robot anywhere in its environment.

We instead focus on asking for many types of help from the actual occupants of the environment and beneficiaries of the robots' services as the robot performs tasks for them (*e.g.,* [9]). We argue that robots that ask for help from occupants combine the benefits of asking passers-by with supervisors in the following ways, discussed next.

2. SPATIALLY- SITUATED OCCUPANTS

We define **occupants** of buildings as having predefined spatially-situated work spaces and conducting long-term work which requires that they be present over a period of time. While they have similarities to both supervisors and bystanders, they also have constraints which violate the assumptions of this previous work. In particular, they are spatially-situated in the environment and no single occupant can help the robot at every location in the environment. As a result, the robot will need to navigate to an occupant's office to request help. Additionally, they may not be available to help at their location, and the robot will need to learn and model this availability through long-term interaction.

2.1 Distributed Help

The idea of human computation (*e.g.,* [13]) and crowdsourcing on websites like Amazon.com's Mechanical Turk [8] have been used to gather help from a **distributed** set of people. In robot domains, bystanders and passers-by in busy environments have helped robots complete tasks in locations as varied as offices [2], conferences [7], and even on the street [14]. Because the number of people in these areas is so high, there is a very limited possibility that any particular person will be asked for help too frequently or will be asked to spend a lot of time with the robot.

Because there are many occupants in a building, the work to help the robot is distributed among them, significantly reducing the burden of help especially in the office environment in which occupants are busy and may not be able to answer frequently. It is unclear whether occupants are even willing to help the robot at all given that they have other work to do. Our work will address the following questions:

- Are office occupants willing and available to help a robot perform its tasks?
- Are they willing to provide some types of help more than others (*e.g.,* are they unwilling to help with tasks that require them to leave their office)?
- Are there available occupants distributed around the environment or only in one area of the building?

2.2 Incentives

Because a robot would know the office locations, it could provide services or **incentives** to the occupants to encourage occupants to help it. Like supervisors who must help the robot for tasks to be completed, if occupants want to continue to receive incentives from the robot (*e.g.,* mail delivery), they must also agree to help the robot at some times. In this work, we assume that the robot will ask for help more often than it will provide incentives (*e.g.,* it will require help navigating to deliver mail to other occupants more often than any occupant will receive mail themselves). We will answer the following questions relating to incentives:

- Are office occupants report that they are motivated to answer questions with incentives?
- Do they actually answer more frequently when offered an incentive?

2.3 Long-Term Interaction

Because supervision occurs over time, there are additional opportunities for robots to take advantage of the **long-term interactions**. Models of humans have been proposed to take into account the expertise of the human to determine the type of question that the robot should ask [3, 6], to ground or familiarize the helper with the robot's current state to increase the likelihood of accurate responses [4, 10], and to model the helper's interruptibility or availability to answer questions [5, 11].

We assume that the occupant would not be supervising a robot while conducting their work and often may not be available to help the robot if it needs it. Unlike bystanders who may only have to answer a particular robot's questions once, an office occupant will be asked questions more often allowing the robot to learn who is often available at certain times to avoid interruptions as well as keep track of the frequency of questions to avoid asking a single person too often. We address the following questions:

- Is there a novelty effect associated with willingness to help the robot and does willingness to help decrease over time?
- Can the availability of occupants be modeled to take into account who will be able to help?
- Does occupant availability change through the day?

2.4 Navigation

Finally, due to the lack of constant supervision, the robot must **navigate** to the spatially-situated occupants to determine their availability and to ask them for help if needed. Current navigational models take into account the uncertainty in the path and the path distance but not who is available to help the robot along a path. Intuitively, a robot should choose short paths that also have humans available to help it if necessary. We will answer the following question:

- Can a robot use availability information from its long-term interactions to determine its best path?

3. REFERENCES

[1] 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.

[2] 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 *15th International Joint Conference on Artificial Intelligence*, pages 880–885, 1997.

[3] D. J. Bruemmer, 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.

[4] H. Clark and D. Wilkes-Gibbs. Referring as a collaborative process. *Cognition*, Jan 1986.

[5] J. Fogarty, S. E. Hudson, C. G. Atkeson, D. Avrahami, J. Forlizzi, S. Kiesler, J. C. Lee, and J. Yang. Predicting human interruptibility with sensors. *ACM ToCHI*, 12(1):119–146, 2005.

[6] 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.

[7] M. Michalowski, S. Sabanovic, C. DiSalvo, D. Busquets, L. Hiatt, N. Melchior, and R. Simmons. Socially distributed perception: Grace plays social tag at aaai 2005. *Autonomous Robots*, 22(4):385–397, 2007.

[8] M.Turk. Amazon.com mechanical turk – artificial, artificial intelligence, 2010.

[9] S. Rosenthal, J. Biswas, and M. Veloso. An effective personal mobile robot agent through a symbiotic human-robot interaction. In *AAMAS '10: 9th International Joint Conference on Autonomous Agents and Multiagent Systems*, 2010.

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

[11] 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 *HRI '08: 3rd ACM/IEEE International Conference on Human Robot Interaction*, pages 303–310, 2008.

[12] S. Thrun, M. Beetz, M. Bennewitz, W. Burgard, A. Cremers, F. Dellaert, D. Fox, D. Hahnel, C. Rosenberg, N. Roy, et al. Probabilistic algorithms and the interactive museum tour-guide robot minerva. *The International Journal of Robotics Research*, 19(11):972–999, 2000.

[13] L. von Ahn and L. Dabbish. Labeling images with a computer game. acm conference on human factors in computing systems. In *CHI 2004*, pages 319–326, 2004.

[14] A. Weiss, J. Igelsböck, M. Tscheligi, A. Bauer, K. Kühnlenz, D. Wollherr, and M. Buss. Robots asking for directions: the willingness of passers-by to support robots. In *HRI '10: 5th ACM/IEEE International Conference on Human Robot Interaction*, pages 23–30, 2010.