Task Behavior and Interaction Planning for a Mobile Service Robot that Occasionally Requires Help

Stephanie Rosenthal, Manuela Veloso, and Anind K. Dey
Carnegie Mellon University
Pittsburgh, PA USA 15213

Abstract

In our work, a robot can proactively ask for help when necessary, based on its awareness of its sensing and actuation limitations. Approaches in which humans provide help to robots do not necessarily reason about the human availability and accuracy. Instead, we model the availability of humans in the robot's environment and present a planning approach that uses such model to generate the robot navigational plans. In particular, we contribute two separate planners that allow a robot to distinguish actions that it cannot complete autonomously from ones that it can. In the first planner, the robot plans autonomous actions when possible and requests help to complete actions that it could not otherwise complete. Then for actions that it can perform autonomously, we use a POMDP policy that incorporates the human availability model to plan actions that reduce uncertainty or that increase the likelihood of the robot finding an available human to help it reduce its uncertainty. We have shown in prior work that asking people in the environment for help during tasks can reduce task completion time and increase the robot's ability to perform tasks.

Introduction

Robotic technology has had many advances, but mobile robot agents are still not universally present in our daily environments. While the ultimate goal is for robots to perform tasks autonomously, we realize that robots still have many limitations, at the perception, cognition, and execution levels. Interestingly, many of the limitations may not be limitations for humans. In particular, humans are capable of helping robots in two ways:

- *Increasing Capabilities:* performing physical tasks that a robot does not have the capability to perform autonomously, and
- *Reducing Uncertainty:* reducing a robot's uncertainty about its state or the effects of its actions.

Robots that can plan for these limitations during their tasks and request help to increase its capabilities or reduce uncertainty can increase their task performance (Rosenthal, Biswas, and Veloso 2010).

To complete tasks in uncertain environments, many robots have relied on human supervisors who are always 1) available to monitor their progress and 2) accurate to help them and tell them which action to take (*e.g.,* teleoperators (Dorais et al. 1999)). As more robots are deployed in our environments, it will be infeasible to employ supervisors for each robot. To reduce the dependence on supervisors during tasks, we are interested in robots that ask for help from people already located *in the environment* - particularly those in known static locations such as offices (Rosenthal, Veloso, and Dey to appear). Compared to traditional supervisors, humans in the environment also have limitations:

- the robot must travel to find them in their offices,
- the robot must interrupt them in their offices,
- they may not be in their offices and have limited availability to provide help, and
- they may not always be accurate.

As a robot plans to navigate in the environment, we argue that it must not only consider the distance and expected uncertainty on its many possible paths, but also who is available to help and where, the cost of interrupting and asking them, and whether they will provide an accurate response. A robot that relies on humans in the environment but does not model those humans may navigate along shorter paths with no humans available or with humans who provide inaccurate help. As a result, a robot may not be able to receive the help it needs and may fail to complete tasks.

In this work, we first present our robot CoBot that performs tasks autonomously but has localization uncertainty and manipulation limitations, and our environment with humans located offices that the robot performs tasks in. We then contribute a two-level framework for task-planning that includes asking humans in the environment to help which our robot uses to perform services. Our Behavior Interaction Planner first plans the high level goals of the robot such as navigating to the elevator or another location, speaking its goals to humans in the environment, and finding humans if it is incapable of performing an action autonomously (Rosenthal, Biswas, and Veloso 2010; Veloso et al. 2011). This high level plan ensures that the task can be completed. Then, the robot plans its motions through the environment using a Human Observation Provider POMDP (HOP-POMDP (Rosenthal, Veloso, and Dey 2011)) that takes into account human accuracy and availability to reduce localization uncertainty if it may need help navigating along its path. We have shown that humans in our environment are willing to help the robot increase its capabilities and reduce its uncertainty, and that this help can improve task performance.

Related Work

Robots have limitations that affect their task performance. Human supervisors have traditionally been used to overcome a robot's limitations by monitoring robot progress and intervening when necessary, but such help can be very expensive in terms of monitoring time and cognitive load on the helper (Yanco, Drury, and Scholtz 2004). Much recent work has focused on different techniques to allow robots to reason about their own limitations and capabilities to proactively ask for help from supervisors, teachers, and passers-by or bystanders in the environment. These different types of help have been classified in a variety of ways (Goodrich and Schultz 2007; Hearst 1999; Parasuraman, Sheridan, and Wickens 2000; Scholtz 2002), but we differentiate our work from previous approaches that assume humans are always available to help and therefore do not plan when and who it is appropriate to ask.

Task Planning to Increase Capabilities

Much of the work on increasing robots' capabilities beyond their autonomous behavior has focused on teleoperation. Many different methods have been proposed for supervisors to provide help including providing assistance at different levels of granularity depending on the robot's capabilities (Dorais et al. 1999) and participating in mixed-initiative interactions to ensure the robot performs its task effectively. Teachers have been able to accurately label data both when the robot requests it and through corrective feedback after the action is performed in order for robots to learn policies for which actions to take (Argall et al. 2009). These techniques vary in the amount of understanding the robot has about its own capabilities. If it knows its capabilities, it can proactively plan to ask for a supervisor to take control when needed. However, the prior work all assumes that there is a human supervisor that is always available to help the robot.

Task Planning to Reduce Uncertainty

Supervisors have also been used to help reduce uncertainty that robots may have when interpreting their sensor data (Fong, Thorpe, and Baur 2003). For example, recent work has focused on navigation planning using "oracles" who are always available and accurate to help robots execute using POMDPs. Oracular POMDPs (OPOMDPs) have been proposed to plan for needing help to reduce uncertainty through its own actions or through asking humans for help without modeling the human in states explicitly (Armstrong-Crews and Veloso 2007). However, OPOMDPs assume that there is an always-available oracle that can be queried for observations from any of the robot's states at a cost of asking.

More recently, there has been an interest in distributing the burden of uncertainty help among crowds of bystanders (Asoh et al. 1997; Hüttenrauch and Eklundh 2006; Michalowski et al. 2007; Weiss et al. 2010)) Bystanders and passers-by in busy environments have helped robots navigate in locations as varied as offices (Asoh et al. 1997), conferences (Michalowski et al. 2007), and even on the street (Weiss et al. 2010). While humans are not actively in contact with the robot all the time, there is still an assumption that at least **one** human will help the robot shortly after it requests it. To the authors' knowledge, there has been little work on planning who to ask in an environment (Shahaf and Horvitz 2010). Because these robots do not have a way to model bystanders in the environment in terms of who will be available or where they are located, nor can they proactively plan to contact a known helper, they have little control over the help they receive and cannot plan to optimize their performance using that help.

Task Planning with Capabilities and Uncertainty

In this work, we are interested in planning tasks for robots that overcome both capability and uncertainty limitations. Beyond the aforementioned work on teleoperation and mixed-initiative robots which requires a human supervisor to monitor the robot at all times to help with any problems that may arise, there has been little work focused on planning for the need of several different types of help from several different people. Even our prior work (Rosenthal, Biswas, and Veloso 2010) which incorporated both types of help did not plan for *who* was available.

In this work, we propose a two-layered task planner to plan the autonomous actions and the interaction behaviors to request help (*e.g.,* (Simmons et al. 1997)). In the Behavior Interaction Planner, the robot plans actions to ensure that it can complete tasks. In other words, it plans autonomous actions when possible and when it lacks a capability it plans to find someone to ask for help. The HOP-POMDP planner plans to complete tasks under uncertainty while explicitly modeling the availability and accuracy of different people to determine who to ask and where to navigate. Next, we describe our mobile robot and academic environment, and then we describe its task planning within that environment.

Robot and the Environment

Our environment consists of one nine-floor of an academic building containing approximately 79 offices per floor. On one floor, for example, there are 35 individual offices for faculty and staff and 44 offices each shared by 2-3 graduate students. Our robot, CoBot (Fig. 1(a)), is capable of autonomous localization using the Kinect and WiFi and omnidirectional navigation in the environment as well as dialog with humans[1]. It has a laptop with the screen facing forwards, towards from the direction of movement, that occupants can use to interact with the robot while it performs tasks autonomously in our building for the occupants.

However, our goal is to make an agile, inexpensive robot platform and, as a result, CoBot has some limitations. CoBot has high localization uncertainty in large open spaces (Fig. 1(b) - darker grey areas indicate more uncertainty) and also

[1]Thanks to Mike Licitra and Joydeep Biswas for their work on the CoBot hardware and localization/navigation.

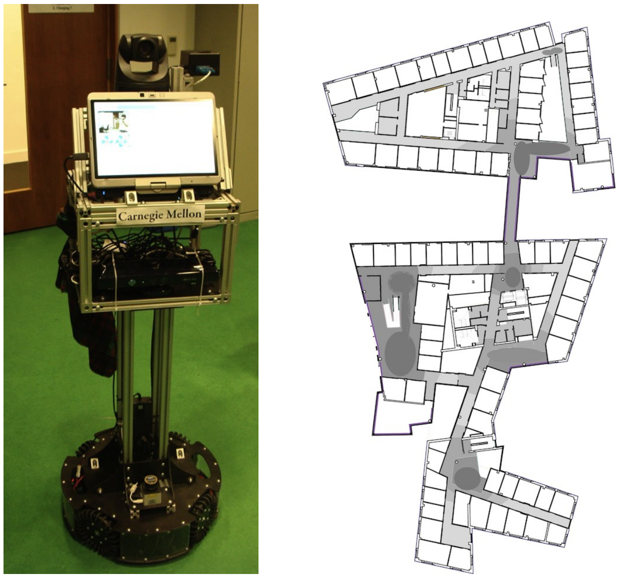

(a) The CoBot2 Robot (b) Areas of Uncertainty

Figure 1: (a) CoBot is capable of autonomous localization and navigation but has manipulation limitations without arms and (b) has localization uncertainty in the hallways of the building (the darker the grey the more uncertainty in that location).

has difficulty perceiving chairs found in common areas resulting in increased navigation time as it attempts to relocalize or avoid these areas entirely. Additionally, CoBot does not have arms or the ability to manipulate objects to push chairs out of the way, press elevator buttons to navigate between floors, or pick up the mail or other objects to give to the building occupants. While the robot can overcome some of these challenges autonomously, inevitably, it must ask for help from humans sometimes to resolve each of these limitations, particularly the physical ones.

In prior work, we have found that several of our building occupants are willing to help the robot with localization questions to reduce uncertainty, as well as help it increase its capabilities by moving chairs out of the way and writing notes to other building occupants (Rosenthal, Veloso, and Dey to appear). Given that people in the environment are willing to help, we would like to model their availability so that CoBot can proactively plan to visit offices when it needs help. First, we describe the Behavior Interaction Planner which plans autonomous actions and help to increase its capabilities with manipulation tasks. Then, we'll describe our HOP-POMDP navigational planner which plans paths where humans are likely to be available in case the robot needs help reducing uncertainty.

Behavior Interaction Planner

Typically, high-level task planners plan only the autonomous actions to complete a task and a separate dialog manager interacts with humans to receive the task requests such as transporting objects. Our goal is for CoBot to complete tasks as autonomously as possible, but we increase its capabilities through asking for help from humans in the environment. Additionally, as CoBot performs actions, humans may want to know what robot's goals are. Our Behavior Interaction Planner therefore reasons about a robot's incapabilities (Rosenthal, Biswas, and Veloso 2010) and human interest in the robot and plans for both human interactions in addition to the autonomous actions.

Modeling Capabilities, Actions and Interactions

We define actions and interactions that are required to complete a task along with their preconditions and effects and we model the CoBot's capabilities in terms of the actions it can perform. In particular, if CoBot has probability $p = 0$ of completing an action, it should ask for help; otherwise, it should attempt to complete the action autonomously (Rosenthal, Biswas, and Veloso 2010).

For `ask` interactions, there are no preconditions, so CoBot can ask for help at any time. In order to do so, the robot plans to `speak`s the defined question text and display the question with multiple choice answers on the screen, and the effect is the required human response (*e.g.* clicking a 'Done' button on CoBot's user interface when a task is completed). If the human clicks a button that is not the required response or there is no response for 30 seconds, the robot replans its task and will find another person to ask for help.

For `navigate` actions, the precondition is that the robot `speak` aloud its new goal to humans in the area, and at execution time the robot sends a desired location to the HOP-POMDP navigation planner to plan and execute a route that is likely to have available help if the robot needs to reduce uncertainty. The planning effect is that the robot is in the location that it should navigate to. Any other actions needed for a task can be defined similarly in terms of autonomous actions and requests for help.

Autonomous Planning and Execution

Given a new task, the robot plans the sequence of actions necessary to complete it. For example, in the $\mathrm{Transport(fromroom\#, toroom\#, object)}$ task, the Behavior Interaction Planner plans the following sequence of actions (conditional plan in Figure 2a, illustrated in Figure 3): `navigation` to the *pick-up* location $\mathrm{fromroom\#}$, `ask` for the object o, `navigation` to the *drop-off* location $\mathrm{toroom\#}$, and `ask` for task completion confirmation. It also does an initial check to ensure that there is enough battery power to get the robot to the location.

The Behavior Interaction Planner reasons about the requirements of $\mathrm{navigation}$ to plan for a robot's incapabilities. For example, if CoBot (with no arms) must navigate between different floors of the building, the robot must not only execute `navigateto` actions in $\mathrm{navigation}$, but also ask for help pressing elevator up/down buttons, pressing the correct floor button, and holding the door open for the robot. In these cases, the Behavior Interaction Planner plans not just to navigate to the goal location but also the human interactions and questions necessary to perform the task completely (conditional plan in Figure 2b:

- `navigateto` to elevator,
- `ask` for help pressing the up/down button,

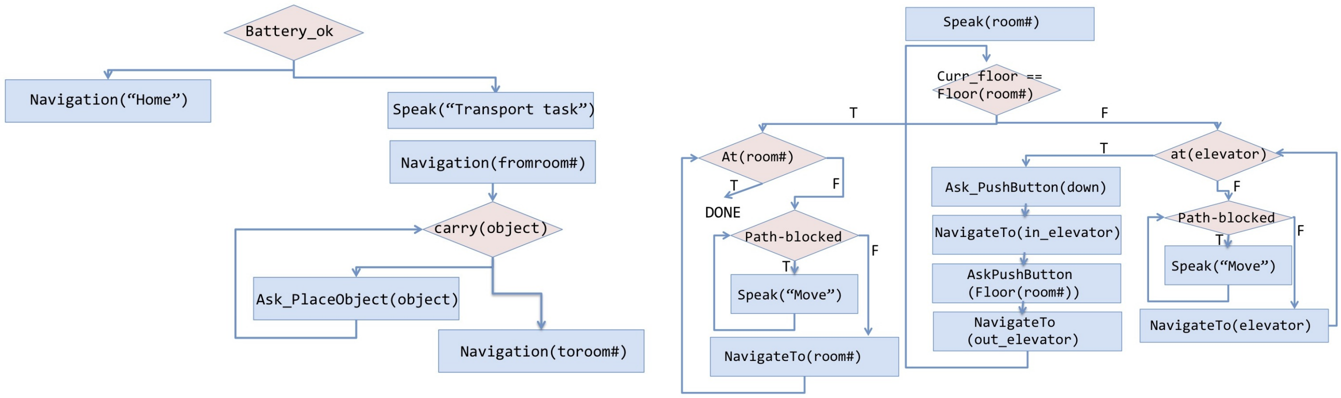

(a) Transport Task Plan

(b) `Navigation` Plan

Figure 2: (a) CoBot executes the conditional plan for the Transport task. (b) For each of the `Navigation` actions in the Transport task, it executes this plan to travel between floors and ask for help.

- `navigateto` into the elevators,
- `ask` for help pressing the floor number and recognizing that floor,
- `navigateto` out of the elevator,
- `navigateto` to goal

To summarize the $\text{Transport}(\text{fromroom}\#, \text{toroom}\#, \text{object})$ task within the Behavior Interaction Planner (Fig. 2), the robot first executes the `Navigation` plan to travel to the $\text{fromroom}\#$. Upon arriving at the location, the robot asks for help picking up the object. Then, it executes the `Navigation` plan to travel to the $\text{toroom}\#$.

Next, we describe the HOP-POMDP navigation planner which we use for `navigateto` actions generated from the Behavior Interaction Planner to increase the likelihood of finding available people to help the robot along its paths.

Humans Observation Provider POMDPs

While the robot can plan generally to ask for help, it must navigate to proactively find people who are available at execution time. We will model locations, availability, cost of interruption and asking, and accuracy of humans so that robots can consider the benefits of asking different humans in addition to the distance to its goal and determine who to ask and where to navigate. Without a model of humans, the robot may choose a path that has no help available or one where humans often provide inaccurate help, and the robot may fail to complete its task if it actually needs help along these paths. We introduce the HOP-POMDP as a planning framework to determine which path to take in the environment (Rosenthal, Veloso, and Dey 2011).

Humans as Observation Providers

Unlike oracles modeled in OPOMDPs (Armstrong-Crews and Veloso 2007), humans in the environment are not always available or interruptible (Shiomi et al. 2008), may not be accurate (Rosenthal, Dey, and Veloso 2009), and they may have a high cost of asking or interruption (Rosenthal, Biswas, and Veloso 2010). We formalize these limitations within the POMDP framework. In particular, we will model the probability of a robot receiving an observation from a human in terms of the human's availability and their accuracy to reduce the uncertainty of the robot. A similar formulation can be achieved for increasing capabilities.

Location We assume that humans are located in a particular known location in the environment, and can only help the robot from that location. When the robot is in state s it can only ask for help from the human h_s in the same state. As a result of taking the ask action `ask`, the robot receives an observation o from the human.

Availability The *availability* of a human in the environment is related to both their presence and their interruptibility (Fogarty et al. 2005). We define availability α_s as the probability that a human provides a non-null observation o in a particular state s:

$$0 \leq \alpha_s \leq 1 \tag{1}$$

If there is no human available in particular state, $\alpha_s = 0$. A human provides observations with probability

$$\sum_{o \neq o_{null}} p(o|s, \texttt{ask}) = \alpha_s \tag{2}$$

and would provide no observation o_{null} otherwise

$$p(o_{null}|s, \texttt{ask}) = 1 - \alpha_s \tag{3}$$

Receiving the o_{null} is equivalent to receiving no observation or timing out waiting for an answer. This is to ensure that $\sum_o p(o|s, \texttt{ask}) = 1$.

Accuracy The non-null observation o that the human provides when they are available depends on their *accuracy* η. The more accurate the human h_s, the more likely they are to

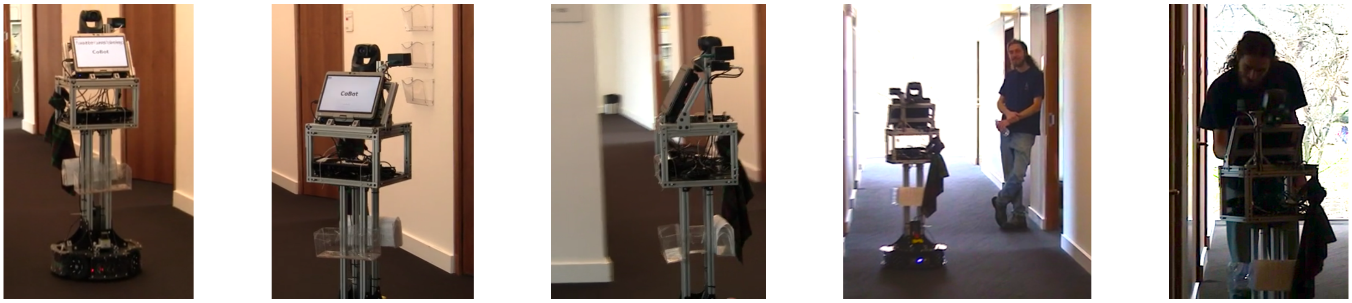

(a) (b) (c) (d) (e)

(f) (g) (h) (i)

Figure 3: (a,b,c) After CoBot-2 receives a *Transport* task request, it autonomously navigates to the location l_p to pick up a bottle of water. (d,e) Upon arriving to l_p, CoBot-2 asks the user to place the bottle of water and press 'Done'. (f,g) Then, CoBot-2 navigates to location l_d to deliver the bottle of water. (h,i) When the user presses 'Done' CoBot-2 navigates back to its home location. (The complete example sequence is submitted as a video with this paper.)

provide a true observation o_s. Otherwise, h_s provides observations $o_{s'}$ where $s' \neq s$.

Formally, we define the accuracy η_s of h_s as the probability of providing o_s compared to the probability they provide any non-null observation $o \neq o_{null}$ (their availability α_s).

$$\eta_s = \frac{p(o_s|s, \texttt{ask})}{\sum_{o \neq o_{null}} p(o|s, \texttt{ask})} = \frac{p(o_s|s, \texttt{ask})}{\alpha_s} \qquad (4)$$

Cost of Asking It is generally assumed that supervisors are willing to answer an unlimited number of questions as long as their responses help the robot. However, in active learning, there is a *cost of asking* in terms of the time it takes for them to answer the question and the cost of interrupting them to limit the number of questions asked.

Let λ_s denote the cost of asking for help from h_s. These costs vary for each person, but are assumed to be known before planning. The cost for querying the human if they answer with a non-null observation $o \neq o_{null}$ is

$$R(s, \texttt{ask}, s, o_s) = -\lambda_s \qquad (5)$$

However, if the person is not available to hear the question or provide a response, there is no expected cost.

$$R(s, \texttt{ask}, s, o_{null}) = 0 \qquad (6)$$

Our reward structure has consequences that affect policy solutions. In particular, the robot does not receive negative reward when it tries unsuccessfully to ask someone for observations so it can afford to be riskier in who it tries to ask rather than incurring a higher cost of asking someone who is more available.

HOP-POMDP Formalization

To briefly review, POMDPs are represented as the tuple $\{\mathcal{S}, \mathcal{A}, \mathcal{O}, \Omega, T, R\}$ of states $\mathcal{S}$, actions $\mathcal{A}$, observations $\mathcal{O}$ and the functions:

- $\Omega(o, s, a) : \mathcal{O} \times \mathcal{S} \times \mathcal{A}$ - observation function, likelihood of observation o in state s after taking action a
- $T(s, a, s') : \mathcal{S} \times \mathcal{A} \times \mathcal{S}$ - transition function, likelihood of transition from state s with action a to new state s'
- $R(s, a, s', o) : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \times \mathcal{O}$ - reward function, reward received for transitioning from s to s' with action a and observation o

We define the HOP-POMDP as a POMDP for a robot moving in the environment with humans, and then discuss differences between humans as observation providers and noisy sensors.

Let HOP-POMDP be $\{\Lambda, \mathcal{S}, \alpha, \eta, \mathcal{A}, \mathcal{O}, \Omega, T, R\}$. where

- Λ - cost of asking each human
- α - availability for each human
- η - accuracy for each human
- $\mathcal{A} = A \cup \{\texttt{ask}\}$ - autonomous actions and a query action
- $\mathcal{O} = O \cup \{\forall s, o_s\} \cup o_{null}$ - autonomous observations, an observation per state, and a null observation

- $T(s, a_{ask}, s) = 1$ - self-transition for asking actions

Specifically, let h_s be the human in state s with availability α_s, accuracy η_s, and cost of asking λ_s. Our observation function Ω and reward function R reflect the limitations of humans defined in Equations 1-6. The remaining rewards, observations, and transitions are defined as with any other POMDP.

Plan Execution

The best HOP-POMDP policy is one in which the robot takes actions that result in low uncertainty or takes actions that leave it in states with a high possibility of a human reducing its uncertainty. As a result, the robot may plan longer paths to navigate in the hallways, but the robot is more likely to navigate with low uncertainty. With lower uncertainty, the robot will navigate faster to its goal locations (Rosenthal, Biswas, and Veloso 2010). Additionally, if the robot is taking paths with a high likelihood of human availability, it can ask these same people to help increasing its capabilities (*e.g.,* pressing elevator buttons).

Conclusions and Future Work

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 supervising the robot and always available to help. We relax these assumptions by asking for help from humans in the robot's environment and by planning to proactively ask for help. Using the two-layer planning, the robot can plan to complete tasks by first determining when it *must* ask for help and planning for those human interactions. Then, for other actions, it can use our HOP-POMDP framework to plan autonomous policies that limit uncertainty when possible and ask for help to reduce uncertainty otherwise. For future work, we will compare our two-layer planner to other planning algorithms on our real deployed robot, CoBot.

References

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

Armstrong-Crews, N., and Veloso, M. 2007. Oracular pomdps: A very special case. In *ICRA '07*, 2477–2482.

Asoh, H.; Hayamizu, S.; Hara, I.; Motomura, Y.; Akaho, S.; and Matsui, T. 1997. Socially embedded learning of the office-conversant mobile robot jijo-2. In *IJCAI-97*, 880–885.

Dorais, G. A.; Bonasso, R. P.; Kortenkamp, D.; Pell, B.; and Schreckenghost, D. 1999. Adjustable autonomy for human-centered autonomous systems. In *IJCAI Workshop on Adjustable Autonomy Systems*, 16–35.

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

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

Goodrich, M. A., and Schultz, A. C. 2007. Human-robot interaction: a survey. *Found. Trends Human-Computer Interaction* 1(3):203–275.

Hearst, M. A. 1999. Mixed-initiative interaction: Trends and controversies. *IEEE Intelligent Systems* 14–23.

Hüttenrauch, H., and Eklundh, S. 2006. To help or not to help a service robot: Bystander intervention as a resource in human-robot collaboration. *Interaction Studies* 7(3):455–477.

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

Parasuraman, R.; Sheridan, T.; and Wickens, C. 2000. A model for types and levels of human interaction with automation. *IEEE Transactions on Systems, Man and Cybernetics* 30(3):286–297.

Rosenthal, S.; Biswas, J.; and Veloso, M. 2010. An effective personal mobile robot agent through a symbiotic human-robot interaction. In *AAMAS '10*, 915–922.

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

Rosenthal, S.; Veloso, M.; and Dey, A. K. 2011. Learning accuracy and availability of humans who help mobile robots. In *AAAI '11*.

Rosenthal, S.; Veloso, M.; and Dey, A. K. to appear. Is someone in this office available to help me? proactively seeking help from spatially-situated humans. *Journal of Intelligent and Robotic Systems*.

Scholtz, J. 2002. Evaluation methods for human-system performance of intelligent systems. Technical Report 990, NIST Special Publication.

Shahaf, D., and Horvitz, E. 2010. Generalized task markets for human and machine computation. In *AAAI*.

Shiomi, M.; Sakamoto, D.; Takayuki, K.; Ishi, C. T.; Ishiguro, H.; and Hagita, N. 2008. A semi-autonomous communication robot: a field trial at a train station. In *HRI '08*, 303–310.

Simmons, R. G.; Goodwin, R.; Haigh, K. Z.; Koenig, S.; O'Sullivan, J.; and Veloso, M. M. 1997. Xavier: experience with a layered robot architecture. *SIGART Bulletin* 8:22–33.

Veloso, M.; Ventura, R.; Rosenthal, S.; Coltin, B.; and Biswas, J. 2011. Deploying robots in an office environment: Web-based task requests by users and autonomous robot navigation. In *submission*.

Weiss, A.; Igelsböck, J.; Tscheligi, M.; Bauer, A.; Kühnlenz, K.; Wollherr, D.; and Buss, M. 2010. Robots asking for directions: the willingness of passers-by to support robots. In *HRI '10*, 23–30.

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