A Few Issues on Human-Robot Interaction for Multiple Persistent Service Mobile Robots

Manuela Veloso,
with Joydeep Biswas, Brian Coltin, Max Korein, Juan Pablo Mendoza, Vittorio Perera.
Thanks to Cetin Mericli, Thomas Kollar, Yichao Sun, Stephanie Rosenthal
mmv@cs.cmu.edu
School of Computer Science, Carnegie Mellon University

Introduction

AI and robotics researchers aim at having robots in our environments coexisting with humans, as artificial creatures that will help humans and collaborate with humans to improve our societies. There will be more than one robot. Robots will not interact with some humans just once, or a few times, but many times. Humans will interact with robots to change their requests and to teach and correct their behaviors. This abstract briefly discusses a few issues for AI and HRI for such challenging repeated interactions in space and time between robots and humans. We have made different levels of research progress on these issues, as we discuss. Our presentation is motivated by our work with the CoBot mobile service robots, which have been running in our environments for the last three years, and for more than 500kms.

The CoBot Robots

The CoBot robots have several contributions to service robots (Veloso et al. 2012), including: robust real-time autonomous localization (Biswas, Coltin, and Veloso 2011), based on WIFI data (Biswas and Veloso 2010), and on depth information (Biswas and Veloso 2012); symbiotic autonomy in which the deployed robots can overcome their perceptual, cognitive, and actuation limitations by proactively asking for help from humans (Rosenthal, Biswas, and Veloso 2010), and from the web (Kollar, Samadi, and Veloso 2012; Samadi, Kollar, and Veloso 2012), and from other robots (Aguero and Veloso 2012); human-centered planning in which models of humans are explicitly used in robot task and path planning (Rosenthal, Veloso, and Dey 2011); semi-autonomous telepresence enabling the combination of rich remote visual and motion control with autonomous robot localization and navigation (Coltin et al. 2011); web-based user task selection and information interfaces; and creative multi-robot task scheduling and execution (Coltin and Veloso 2012). The robots, built by Mike Licitra and maintained by Joydeep Biswas, have operated over 500km for more than three years basically without hardware failures. Our robots include a modest variety of sensing and computing devices, including the Microsoft Kinect depth-camera, vision cameras for telepresence and interaction, a small Hokuyo LIDAR for obstacle avoidance and localization comparison studies, and a touch-screen and speech-enabled tablet, microphones and speakers. Figure 1 shows three of our CoBot robots.

Figure 1: Three CoBot serive mobile robots.

We have been experiencing the CoBot service robots for a long time at CMU navigating in our indoor environments fully robustly. The robots perform tasks of escorting people, transporting items, and just going to places. We raise four issues: (i) the robots currently have no model, no memory of any of the tasks performed or of the humans with whom they interacted; (ii) the robots need to be interruptible by humans and possibly correctable; (iii) the robots need to coordinate among the multiple robots, not only in terms of task scheduling, but also in terms of their interactions with humans; and (iv) the robots need to provide a safe motion.

Robots with Past, Present, and Future

CoBot can very effectively escort people in the building. It does it by itself and not followed by any developer. So for example, consider Manuela interacting with CoBot to request that CoBot escort a visitor from Manuela's office to the visitor's next appointment at 3pm. The dialog-based human-robot interaction is perfectly effective and the robot is able to ground its task, including the destination location from dialoging with the task requester. Manuela stays in her office and CoBot leaves escorting the visitor. Later that day, or on the following day, CoBot comes by Manuela's office again, for some other task, e.g., delivering a message or a package. Naturally, Manuela would like to ask "How did it go with the visitor this morning? Did you reach the destination well? Did you have any problem finding the place? Was the meeting host ready to receive our visitor? Did you apologize that the visitor was a bit late?" Although the robot logs most of its execution, we don't have yet a way to reference past experience to be able to engage in an appropriate human-robot interaction that refers to the filtered past of the robot relevant to a specific user. References to the past, and correlations with the present, and predictions of the future need to be part of the human-robot interaction with persistent robots, i.e., robots that are present everyday.

We have a long experience and interest on reasoning based on memory (Veloso 1994; Ros et al. 2009; Fernandez and Veloso 2012; Mericli, Veloso, and Ak?n 2012), where experience needs to be stored and indexed, retrieved based on some similarity metric, and effectively adapted. We hope that such research is of interest to persistent service robots.

Interruptible, Correctable, Own Autonomy

We have spent the last few years getting our mobile robots to autonousmly localize, navigate, and perform tasks in the environment. It was a long path to get to this point, where everything is working, namely: hardware, perception, sensing, localization, path planning, navigation, task requests, task-based interaction with humans, and symbiotic autonomy enabling the robots to proactively ask for help from humans and the web, when they find needed. We now realize that the robots seem to be "too," autonomous, i.e., they get tasks, they generate plans to achieve their tasks, and they execute them with the sole goal of executing their plans. We realize that humans may want *to interrupt* the robot on its navigational path, "CoBot, please wait!" a user can say. CoBot is executing its plan, can now be interrupted. We have researched on how to represent tasks, and their constraints and execution status.

The human-robot interaction for allowing a human to interrupt a robot is complex, as many situational features and constraints need to be consider, including task priorities, humans requesters, interruption frequency, and timings.

In addition to interruptions of a task for creating a new task, of particular interest are interruptions related to corrections of the robot behavior. We have previously worked on language-based teaching and correction of new tasks to our CoBots (Mericli et al. 2014), as well as robots demonstrating and correcting robot's behaviors in general (Chernova and Veloso 2009; Argall, Browning, and Veloso 2012; Mericli, Veloso, and Akin 2012), to capture new tasks and revise behaviors, in particular also to be socially appropriate.

Furthermore, persistent robots need to generate their own actions to match any metric to value their own autonomy. They can explore on their own to learn task-related features of the environment that can improve their user-task success rate (Korein, Coltin, and Veloso 2014). Due to their experience, they can also act to anticipate requests, based on learned personal task preferences and routines.

Safe Execution

The CoBot robots move by themselves, and are not chaperoned by any person. Such autonomous execution behavior in human environments raises important safety objectives. The robots are already equipped with a variety of mechanisms to detect possible failures, in particular sensing failures, and to respond to such failures appropriately. However the general problem of guaranteeing safe robot execution in their mobile interaction with humans, is a complex research problem, which needs to be addressed for complete autonomous navigation in the environments and interaction with people.

One particular aspect is to enable our robots to recognize internal and external factors that may interfere with successful motion execution. Even when our robots are equipped with appropriate obstacle avoidance algorithms, collisions and other forms of motion interference might be inevitable: there may be obstacles in the environment that are invisible to the robot's sensors, or there may be people who could interfere with the robot's motion. We have already started working on the general safe execution focused on motion interference detection (Mendoza, Veloso, and Simmons 2012). We have introduced a Hidden Markov Model-based model for detecting such events in mobile robots that do not include special sensors for specific motion interference. We have identified the robot observable sensory data and model the states of the robot. Our algorithm is motivated and implemented on our omnidirectional CoBot mobile service robot equipped with a depth-camera. Our experiments have shown that our algorithm can detect over 90% of motion interference events while avoiding false positive detections.

Complete guaranteed safe autonomy of a mobile robot is our overall target, but we are well aware of the challenges of such goal, due to the inevitable uncertainty of the robot interaction with its environment, in particular humans.

Multiple Robots

There are multiple opportunities for multi-robot coordination, in terms of distributed and shared perception, planning, execution, and sharing of learned personalized information about humans and the environment.

For our multiple service robots, we have introduced a task planning and scheduling for the robots to fulfill all requests as an instance of the Vehicle Routing Problem (VRP). We have introduced an interesting extension of the VRP, namely VRP with Transfers, i.e., which allows for the robots to *transfer* items between each other. By having one robot pick up an object and transfer it to a different robot (or a series of robots) for delivery, both time and the battery life of the robots can be conserved. Transferring items makes the problem significantly harder because this creates exponentially more possible schedules.

Multiple robots can share all sorts of information that may be able to enhance their interaction with humans in general and specific humans in particular. With multiple robots, we care not only about human-robot interaction, but also about robot-robot interaction still in the presence of humans. We need to consider multiple robots with multiple online task requests, with communication-limited distributed scenarios, and with a variety of capacity and priority constraints. Furthermore multiple robots can share learned knowledge, and they can tune their task scheduling and planning to the possible robots' capabilities and contexts.

In conclusion, if the future will include multiple autonomous robots capable of asking for help from humans, and servicing humans in a persistent manner, we will develop novel human-robot interaction approaches involving past, present, and future time handling, interruptibility, correction, own task generation from learning, multiple robots, and safe mobile execution.

References

Aguero, C., and Veloso, M. 2012. Transparent Multi-Robot Communication Exchange for Executing Robot Behaviors. In *10th Int. Conf. on Applications of Agents and Multi-Agent Systems*.

Argall, B.; Browning, B.; and Veloso, M. 2012. Policy Feedback for the Refinement of Learned Motion Control on a Mobile Robot. *International Journal of Social Robotics* 4:383–395.

Biswas, J., and Veloso, M. 2010. Wifi localization and navigation for autonomous indoor mobile robots. In *IEEE International Conference on Robotics and Automation (ICRA)*.

Biswas, J., and Veloso, M. 2012. Depth camera based indoor mobile robot localization and navigation. In *Proceedings of the IEEE International Conference on Robotics and Automation*.

Biswas, J.; Coltin, B.; and Veloso, M. 2011. Corrective gradient refinement for mobile robot localization. In *Intelligent Robots and Systems (IROS), IEEE International Conference on*. IEEE.

Chernova, S., and Veloso, M. 2009. Interactive Policy Learning through Confidence-Based Autonomy. *Journal of Artificial Intelligence Research* 34.

Coltin, B., and Veloso, M. 2012. Optimizing for transfers in a multi-vehicle collection and delivery problem. In *Proceedings of the 11th International Symposium on Distributed Autonomous Robotic Systems (DARS'12)*. Best Paper Award.

Coltin, B.; Biswas, J.; Pomerleau, D.; and Veloso, M. 2011. Effective semi-autonomous telepresence. *Proceedings of the RoboCup Symposium* 289–300.

Fernandez, F., and Veloso, M. 2012. Learning Domain Structure through Probabilistic Policy Reuse in Reinforcement Learning. *Progress in Artificial Intelligence*. DOI:10.1007/s13748-012-0026-6.

Kollar, T.; Samadi, M.; and Veloso, M. 2012. Enabling robots to find and fetch objects by querying the web. In *Proceedings of AAMAS'12, the Eleventh International Joint Conference on Autonomous Agents and Multi-Agent Systems*.

Korein, M.; Coltin, B.; and Veloso, M. 2014. Constrained Scheduling of Exploration Tasks for Service Robots to Learn about their Environment. In *Proceedings of AAMAS'14, the Thirteenth International Joint Conference on Autonomous Agents and Multi-Agent Systems*.

Mendoza, J. P.; Veloso, M.; and Simmons, R. 2012. Motion Interference Detection in Mobile Robots. In *Proceedings of IROS'12, the IEEE/RSJ International Conference on Intelligent Robots and Systems*.

Mericli, C.; Klee, S. D.; Paparian, J.; and Veloso, M. 2014. An Interactive Approach for Situated Task Specification through Verbal Instructions. In *Proceedings of AAMAS'14, the Thirteenth International Joint Conference on Autonomous Agents and Multi-Agent Systems*.

Mericli, C.; Veloso, M.; and Akin, H. L. 2012. Improving Biped Walk Stability with Complementary Corrective Demonstration. *Autonomous Robots* 32(4):419–432.

Mericli, T.; Veloso, M.; and Ak?n, H. L. 2012. Experience guided achievable push plan generation for passive mobile objects. In *Workshop at IROS'12 on "Beyond Robot Grasping - Modern Approaches for Dynamic Manipulation"*. Extended abstract.

Ros, R.; Arcos, J. L.; de Mantaras, R. L.; and Veloso, M. 2009. A Case-Based Approach for Coordinated Action Selection in Robot Soccer. *Artificial Intelligence* 173:1014–1039.

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. International Foundation for Autonomous Agents and Multiagent Systems.

Rosenthal, S.; Veloso, M.; and Dey, A. K. 2011. Is someone in this office available to help me? *Journal of Intelligent & Robotic Systems* 1–17.

Samadi, M.; Kollar, T.; and Veloso, M. 2012. Using the Web to Interactively Learn to Find Objects. In *Proceedings of the Twenty-Sixth Conference on Artificial Intelligence (AAAI-12)*.

Veloso, M.; Biswas, J.; Coltin, B.; Rosenthal, S.; Kollar, T.; Mericli, C.; Samadi, M.; Brandao, S.; and Ventura, R. 2012. CoBots: Collaborative Robots Servicing Multi-Floor Buildings. In *Proceedings of IROS'12, the IEEE/RSJ International Conference on Intelligent Robots and Systems*. Video and Extended Abstract.

Veloso, M. M. 1994. *Planning and Learning by Analogical Reasoning*. Springer Verlag.