The Increasingly Fascinating Opportunity for Human-Robot-AI Interaction: The CoBot Mobile Service Robots

MANUELA M. VELOSO, Carnegie Mellon University

CCS Concepts: • **Computer systems organization** → *Embedded and cyber-physical systems*; *Robotics*; • **Computing methodologies** → **Artificial intelligence**;

Additional Key Words and Phrases: Autonomous mobile service robots, indoor localization and navigation, symbiotic autonomy, verbalization

ACM Reference format:
Manuela M. Veloso. 2018. The Increasingly Fascinating Opportunity for Human-Robot-AI Interaction: The CoBot Mobile Service Robots. *ACM Trans. Hum.-Robot Interact.* 7, 1, Article 5 (May 2018), 3 pages.
https://doi.org/10.1145/3209541

Human-robot interaction (HRI) and autonomous robotics remain worlds apart but need to come together for both fields to substantively advance. Seeing this need, my group and I have pursued the path of truly bringing mobile service robots to our daily human environments. This direction was a major leap of faith and a departure from our work in autonomous robotics, especially for the prospering robot soccer movement. As opposed to the autonomous soccer teams – in which robots are not allowed to get input or interact in any other way with humans while playing a game – service robots, by nature, open the opportunity and the need for interaction with humans. The result is several CoBot robots moving indoors in buildings. Through the CoBot project, we have learned many lessons about the joys and pitfalls of research that advances both HRI and autonomous robotics. The multifaceted role of humans in their interaction with highly capable robots has brought us to many unexpected and fascinating problems at the intersection of HRI and autonomous robotics, which we hope to share.

The basic goal of the CoBot robots is twofold: to receive task requests from users and to autonomously plan and execute such tasks in our buildings. Tasks are of a pick-up and delivery nature, which require the robot to go to some location to greet a visitor or to pick up an object, and then go to some destination location for the visitor or for the object. Initially, we focused on the three core directions in parallel: (*i*) platform issues for hardware, sensors, sensory processing, motor control, route planning, safe navigation, and accurate localization; (*ii*) coordination for representing and scheduling of user requests for multiple CoBots; and (*iii*) interaction with human users on a specific task, such as a visitor companion.

Real robots require real work to get the essentials in place. Our work proceeded tirelessly until we were able to complete the basic loop of going all the way from a task request to the autonomous

Author's address: M. M. Veloso, School of Computer Science, Carnegie Mellon University, Pittsburgh PA 15213, USA.

2573-9522/2018/05-ART5
https://doi.org/10.1145/3209541

execution of such task. The CoBot robot could move reliably using WiFi or depth-based sensory information in parts of our building. New challenges arose, however, that posed technical impasses. The core robot localization was brittle in parts of the building with varying objects, such as tables and chairs, and the robot could not perform many actions due to its intrinsic limitations, for example, no arm manipulation capabilities.

Our major breakthrough came in the form of *symbiotic autonomy*. Essentially, we realized and accepted that the CoBot or any other robots will have many limitations now and continue to face various forms of limitations for the foreseeable future. We introduced the concept of symbiotic autonomy, in which the robots explicitly include actions to ask for help from humans in their behavior policies. Actions for which the robot asks for help include pressing the elevator button and putting objects into their basket, as well as removing objects from their baskets.

The first new direction materialize from symbiotic autonomy involved language, such as speaking with the robot, retrieving information from the web, and learning from the responses. Soon after, another new direction appeared: while the robots move and have an accurate measure of their location, they can produce maps of any vital statistic of the building captured from sensors, such as temperature, noise level, and WiFi signal strength. Autonomous navigation in the building also enables the learning of models of aspects of the building and its users. The robot can learn models of task requests and can create its own navigation and task goals to learn such models. At this point, our robots became extremely robust in their localization and navigation and could move autonomously in any building.

As the robot reliably performs its tasks, however, we realized that the robot is really not very "intelligent." More important, who knows what intelligence really is? The CoBot robots are far from being "present" and "aware" of the humans and environments around them. If there is someone lying on the floor or asking for help, or there is a flood, or there are new blooming daffodils, a CoBot does not notice anything and blindly and robustly executes its navigation to its goal location.

Thus, one more research direction is essential – programming a robot to be alert and aware of its surroundings – so that one day we will enable a mobile service robot that wanders around executing tasks to call 911 if needed.

There will continue to be more directions, more questions. It does not seem that the HRI for any single new goal or new task will necessitate a person coding everything from scratch for every single new task. Learning from examples, demonstration, instruction, and correction plays a major role in HRI.

In particular, autonomous navigation leads robots to disappear from our sight and even from our hearing. A human can see a robot going down a corridor, but when it turns at the end of the corridor, it is no longer visible. Shortly after disappearing from view, the noise of the robot motors cannot even be heard.

We experience, then, an autonomous mobile robot that "disappears" from human control. When the robot later returns to one's office autonomously, maybe guiding a visitor, a new research direction emerges for "verbalization." The robot must be able to explain in natural language to a human what happened while away in its autonomous experience. Also, we want to have the human ask questions to the robot about its experience while it was away. Our introduction of "verbalization" offered an approach to enable the robot to translate its navigation and perceptual experience into natural language. The robot can describe its task execution in language such as "I went on the corridor 7100, and then turned to the bridge, and passed through the kitchen, until I reached my destination 7002."

From our verbalization work, we learned how natural-language queries of users about the robot experience would match different levels of detail for the generated verbalization. These findings naturally led us to having the robot not only respond to questions about its autonomy but also to

have the robot autonomously offer information about its own experience. We allowed the CoBot to engage the human user's attention to possible unusual facts, such as "today there were a lot of people in the kitchen area" and "I waited at the elevator today very little time before someone helped, calling the elevator for me."

Through the CoBot, we saw this increasingly complex HRI as moving toward general human-AI interaction, in which the artificial artifact is increasingly intelligent and the human and robot are in an increasingly intertwined complementary interaction. There are so many directions still to address to move forward: making sure that the robots perform useful tasks, are safe, transparent, trustable, reliable, instructable, correctable, and are of use to improve our human lives.

Our field of human-AI or HRI includes many questions and steps along many directions. All such steps and directions can be appreciated to collectively advance our science. We can study how humans react to the technology, but we can also develop the actual robots and can be driven by the increasingly fascinating and challenging research directions that open up as robots become more available. AI and HRI is not a one-shot research project but instead is a fascinating long-term incremental process. All steps contribute to advance our science. We are still in the infancy of having real autonomous interacting, learning, responsible, useful robots in our environments. Good luck to all of us!

Received April 2018; accepted April 2018