When Robots Meet People: Research Directions In Mobile Robotics

Sebastian Thrun
Computer Science Department and Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15213-3891

Recent research in the field of mobile robotics has led to significant progress along various dimensions. Applications such as robots that guide blind or mentally handicapped people, robots that clean large office buildings and department stores, robots that assist people in recreational activities, etc., are clearly in reach, and for many of those target domains prototypes are readily available. This recent, ongoing revolution has been triggered by advances along various dimensions. Robotic hardware has steadily become cheaper and more reliable. Robotic software has matured, reaching critical levels of reliability, robustness, and flexibility. Mobile robots have an enormous potential to change our everyday lives. It is worth noting that an increasing fraction of these robots rely on methods developed in artificial intelligence.

Together with researchers at the University of Bonn (W. Burgard, A.B. Cremers, D. Fox, D. Hähnel, G. Lakemeyer, D. Schultz and W. Steiner), and motivated in part by the work of Horswill [3], the author recently installed a mobile tourguide robot in a branch of the *Deutsches Museum* [1, 4]. The robot's task was to guide visitors through the museum, explaining to them a subset of the museum's exhibits. The major challenges that arose in this project can be grouped into two categories: navigation and human robot interaction.

- **Navigation.** Navigating safely was a primary concern, since collisions with obstacles—human and exhibits alike—had to be avoided at all costs. At the same time, the robot had to find its way at approximately walking speed. Difficulties arose from the fact that a variety of obstacles were practically "invisible" to the robot's sensors, despite the fact that the robot possessed five independent sensor systems (tactile, infrared, sonar, laser, cameras). Invisible hazards included staircases, glass cages, railings, and various strangely shaped exhibits. The fact that the crowds surrounding the robot sometimes deliberately tried to mislead the robot contributed to the difficulty of the general problem.

- **Human robot interaction.** Other challenges arose from the need to interact with people in "natural" and appealing ways. The very success of the robot depended on its ability to engage people in all age groups (2-80), to "delude" them into tours, and, if successful, to communicate both information and intent effectively. Apart from the museum's visitors, the robot also interacted with people all over the world through the Web, who could control the robot and watch its operation.

An additional challenge was the development of software that facilitates the rapid installation of robots in new and unknown environments. To this end, our approach did not require any modification of the museum itself. Strictly speaking, this requirement does not directly follow from the application scenario at hand, but it is essential for a large number of similar mobile robot applications, and therefore has been a primary scientific interest of ours.

Figure 1: Left: The robot. Center: The museum. Right: Map of the Carnegie Museum of Natural History, learned in less than 2 hours.

The robot was controlled by a modular, distributed software architecture that performed functions such as collision avoidance, mapping, localization, path planning, mission planning, and human interaction. Among the various principles underlying the software design, three stand out as the most notable ones:

1. **Probabilistic representations and learning.** Probabilistic representations enabled the robot to reason with multiple interpretations of its sensor data in a mathematically consistent way. They critically enhanced the robot's robustness and flexibility. In particular, they enabled the robot to learn aspects of its environment, such as the map shown in Figure 1 [5]. The robot learned both off-line and on-line, and both learning modes were essential for the robot's success.

2. **Resource adaptability.** To guarantee continuous operation under the large variety of situations faced by the robot, many modules were capable of adapting their computational and memory requirements to the available resources. In particular, several "any-time" algorithms [2] came to bear. These algorithms produce results when needed, but the quality of the results increase with the available resources.

3. **Distributed, asynchronous control.** A modular, distributed software architecture, consisting of up to 25 individual modules, proved effective in controlling the robot in the presence of various computational and networking bottlenecks. For example, the architecture successfully managed failures such as temporary downtime of the radio link, which connected on-board to off-board computers.

In a nutshell, the robot did its job and it excited the visitors. During a six-day installation period, the robot guided more than 1,500 visitors and 2,000 Web users, traversing more than 18.5km at an average speed of almost 40 cm/sec. Figure 1 shows pictures of the robot, the exhibition, and a map learned by the robot in a different museum. In fact, the robot raised the museum's attendance by more than 50%, suggesting that robotic applications in the field of entertainment and education might be commercially viable.

So what did we learn? First, we learned that research on mobile robot navigation has progressed to a level that robots can now navigate reliably even in densely populated spaces. In most aspects, such robots can now be installed in new sites within days, or even hours. We also learned that human robot interfaces are key if robots are to become part of people's everyday life, and that adaptive mechanisms are essential for operating robots in highly dynamic and unpredictable environments. Finally, we learned that entertainment is a highly promising application domain for mobile robotics. In most envisioned service robot applications, robots have to compete with human labor, whereas in the entertainment sector, robots may generate revenue by simply exploiting the fact that they differ.

We believe that the museum tour-guide is a prototypical example of a new generation of mobile service robots. Many of the challenges in the museum tour-guide domain apply to a wide variety of mobile robot applications: the necessity to navigate through highly dynamic and unstructured environments; the necessity to interact with people; and the necessity to operate in unknown environments that cannot be modified. It is quite likely that robots similar to the one described here will soon be deployed in shopping malls, amusement centers, technology fairs, etc., where they will be receptionists, information kiosks, waiters, guides, but most of all: magnets that attract people. Similar robots may soon perform janitorial services, operate at sites too dangerous for humans, or assist people in various aspects of their lives.

The following is a highly subjective list of some of the most promising research directions:

- **Long living robots.** How can we build software that enables robots to operate reliably for many years? Such robots must be able to adapt to changes in their environments, changes in their physical properties, and they must also be able to perform new tasks. To what extent can we build robots that survive sensor and actuator failures? Most of today's mobile service robots are only deployed for short periods of time and are only capable of performing a single task. A science of software mechanisms that enable robots to operate for extended stretches of time, and to perform many tasks, would almost certainly lead to revolutionary changes in mobile robotics.
- **Human robot interaction/cooperation.** How can we build robots that can be installed and instructed by non-expert users? How can we design robots that can interact with people, and support them in their everyday activities? Most of today's mobile robots are installed and operated by experts, and they interactive capabilities are still poorly developed. To bring robots into the everyday life of people, advances in various aspects of human robot interaction are urgently needed.
- **Taskable robot teams.** How can we build large teams of mobile robots to collaboratively perform interesting tasks? Such software must be able to deal with individual robot failures and limited communication channels. We specifically lack methods that makes teams of robots taskable, i.e., that enable teams of robots to perform diverse families of user-specified tasks.
- **Mobile manipulation.** The integration of mobility and manipulation is essential for a wide range of service tasks. However, mobile robotics and robotic manipulations have largely remained separate fields, and integrations of both are often brittle at best. How can we build highly articulated robots that can both manipulate their environment and navigate therein?
- **Flexible software architectures.** While recent research has led to a large corpus of isolated component technologies, we still lack effective methods for their integration. How can we build software architectures that facilitate the assembly of large-scale robotics software? Learning, in particular, appears to be promising for providing the "glue" between different components that otherwise may not fit together.

While this list is highly subjective, it highlights some of the most promising new research directions in mobile robotics. Progress along any of those dimensions would almost certainly lead to interesting new science, along with practical algorithms with high societal impact.

Acknowledgment

The author gratefully acknowledges the contributions of W. Burgard, A.B. Cremers, D. Fox, D. Hähnel, G. Lakemeyer, D. Schultz and W. Steiner from the University of Bonn, with whom the museum tour-guide application was developed. He also acknowledges the contributions of S. Gutmann, B. Kuipers,

D. Margaritis, T. Mitchell, G. More, R. Romero, N. Roy, R. Simmons, T. Sawyer, and S. Waldherr, who contributed to various aspects of the system described here and influenced his opinion.

References

[1] W. Burgard, D. Fox, G. Lakemeyer, D. Hähnel, D. Schulz, W. Steiner, S. Thrun, and A.B. Cremers. Real robots for the real world — the RHINO museum tour-guide project. (submitted for publication), October 1997.

[2] T. L. Dean and M. Boddy. An analysis of time-dependent planning. In *Proceeding of Seventh National Conference on Artificial Intelligence AAAI-92*, pages 49–54, Menlo Park, CA, 1988. AAAI, AAAI Press/The MIT Press.

[3] I. Horswill. Specialization of perceptual processes. Technical Report AI TR-1511, MIT, AI Lab, Cambridge, MA, September 1994.

[4] S. Thrun, A. Bücken, W. Burgard, D. Fox, T. Fröhlinghaus, D. Hennig, T. Hofmann, M. Krell, and T. Schimdt. Map learning and high-speed navigation in RHINO. In D. Kortenkamp, R.P. Bonasso, and R. Murphy, editors, *AI-based Mobile Robots: Case studies of successful robot systems*. MIT Press, Cambridge, MA, to appear.

[5] S. Thrun, D. Fox, and W. Burgard. A probabilistic approach to concurrent mapping and localization for mobile robots. *Machine Learning and Autonomous Robots (joint issue)*, 1998. to appear.