Towards A Framework for Human-Robot Interaction

Sebastian Thrun

Computer Science Department

Stanford University

Stanford, CA 94305

Abstract

The goal of this article is to introduce the reader to the rich and vibrant field of robotics, in hopes of laying out an agenda for future research on human robot interaction. Robotics is a field in change; the meaning of the term "robot" today differs substantially from the term just a decade ago. The primary purpose of this article is to provide a comprehensive description of past and present-day robotics. It identifies the major epochs of robotic technology and systems—from industrial to service robotics—and characterizes the different styles of human robot interaction paradigmatic for each epoch. To set an agenda for research on human robot interaction, the article articulates some of the most pressing open questions pertaining to modern-day human robot interaction.

1 Introduction

The field of robotics is changing at an unprecedented pace. At present, most robots operate in industrial settings, where they perform tasks such as assembly and transportation. Equipped with minimal sensing and computing, robots are slaved to perform the same repetitive task over and over again. In the future, robots will provide services directly to us people, at our workplaces, and in our homes.

These development are sparked by a number of contributing factors. Chief among them is an enormous reduction in costs of devices that compute, sense, and actuate. Of no lesser importance are recent advances in robotic autonomy, which have critically enhanced the ability of robotic systems to perform in unstructured and uncertain environments (see [Thrun, 2002] for an overview). All these advances have made it possible to develop a new generation of service robots, posed to assist people at work, in their free time, and at home.

From a technological perspective, robotics integrates ideas from information technology with physical embodiment. Robots share with many other physical devices, such as household appliances or cars, the fact that they 'inhabit' the same physical spaces as people do in which they manipulate some of the very same objects. As a result, many forms of human robot interaction involve pointers to spaces or objects that are meaningful to both robots and people [Kortenkamp et al., 1996]. Moreover, many robots have to interact directly with people while performing their tasks. This raises the question as to what the right modes are for human robot interaction. What is technologically possible? And what is desirable?

Possibly the biggest difference between robots and other physical devices—such as household appliances—is autonomy. More than any other research discipline, the field of robotics has striven to empower robots with an ability to make their own decisions in broad ranges of situations. Today's most advanced robots can accommodate much broader circumstances than, for example, dishwashers can. Autonomy opens the door to much richer interactions with people: some researchers have gone as far as proclaiming their systems 'social' [Simons et al., 2003] or 'sociable' [Breazeal, 2003a]. Naturally, sociable interaction offers both opportunities and pitfalls: It offers the opportunity for the design of much improved interfaces by exploiting rules and conventions familiar to people from different contexts. However, it does so at the danger people reflecting capabilities into robotic technology that do not exist. For these and other reasons, it remains unclear if we ever want to interact with robots the same way we interact with our next door neighbor, our maid, our colleagues.

2 The Three Kinds of Robots

Robotics is a broad discipline. The broadness of the field becomes apparent by contrasting definitions of some of the most basic terminology. The Robot Institute of America (1979) defines a robot as "a reprogrammable, multifunctional manipulator designed to move materials, parts, tools, or specialized devices through various programmed motions for the performance of a variety of tasks." In contrast, the Webster dictionary defines a robot as "An automatic device that performs functions normally ascribed to humans or a machine in the form of a human." A technical introduction into robotic sensors, actuators, and algorithms can be found in [Thrun, 2002].

The United Nations (UN), in their most recent robotics survey [U.N. and I.F.R.R., 2002], groups robotics into three major categories. These categories—*industrial robotics*, *professional service robotics*, and *personal service robotics*—are primarily defined through their application domains. However, they also represent different technologies and correspond to different historic phases of robotic development and commercialization.

Industrial robots represent the earliest commercial success, with the most widespread distribution to date. An industrial robot has three essential elements: it manipulates its physical environment (e.g., by picking up a part and placing it somewhere else); it is computer-controlled, and it operates in industrial settings, such as on conveyor belts. The boundary between industrial robots and non-robotic manufacturing devices is somewhat fuzzy; the term robot is usually used to refer to systems with multiple actuated elements, often arranged in chains (robotic arm). Classical applications of industrial robotics include welding, machining, assembly, packaging, palletizing, transportation, and material handling. For example, Figure 1 shows an industrial welding robot in the left panel, next to a robotic vehicle for transporting containers on a loading dock in the right panel.

Industrial robotics started in the early 1960s, when the world's first commercial manipulator was sold by *Unimate*. In the early 1970s, Nissan Corp. automated an entire assembly line with robots, starting a revolution that continues to the present date. To date, the vast majority of industrial robots are installed in the automotive industry, where the ratio of human workers to robots is approximately ten to one [U.N. and I.F.R.R., 2002]. The outlook of industrial robots is prosperous. In 2001, the UN estimates the operational stock of industrial robots to be 780,600, a number that is expected to grow by just below 25% until 2005. According to [U.N. and I.F.R.R., 2002], the average cost of an industrial robots has decreased by 88.8% between 1990 and 2001. At the same time, US labor costs increased by 50.8%. These opposing trends continue to open up new opportunities for robotic devices to take over jobs previously reserved for human labor. However, industrial robots tend not to interact directly with people. Interface research in this field focuses on techniques for rapidly configuring and programming these robots.

Professional service robots constitute a much younger kind of robots. Service robotics is mostly in its infancy, but the field grows at a much faster pace than industrial robotics. Just like industrial robots, professional service robots manipulate and navigate their physical environments. However, professional service robots assist people in the pursuit of their professional goals, largely outside industrial settings. Some of these robots operate in environments inaccessible to people, such as robots that clean-up nuclear waste [Blackmon et al., 1999, Brady et al., 1998] or navigate abandoned mines [Thrun et al., 2003]. Others assist in hospitals, such as the Helpmate robot [King and Weiman, 1990] shown in Figure 2a which transports food and medication in hospitals, or the surgical robotic system shown in Figure 2b, used for assisting physicians in surgical procedures. Robot manipulators are also routinely used in chemical and biological labs, where they handle and manipulate substances (e.g. blood samples) with speeds and precisions that people cannot match; recent work has investigated the feasibility of inserting needles into human veins through robotic manipulators [Zivanovic and Davies, 2000]. Most professional service applications have emerged in the past decade. According to the UN [U.N. and I.F.R.R., 2002], 27% of all operational professional service robots operate underwater; 20% perform demolitions; 15% offer medical services and 6% serve people in agriculture (e.g., by milking cows [Reinemann and Smith, 2000]). Military applications, such as bomb diffusal, search and rescue [Casper, 2002], and support of SWAT teams [Jones et al., 2002], are of increasing relevance [U.N. and I.F.R.R., 2002]. According to the UN, the total operational stock of professional service robots in 2001 was 12,400, with a 100% growth expectation by 2005. The amount of direct interaction with people is much larger than in the industrial robotics field, since service robots often share the same physical space with people.

Personal service robots possess the highest expected growth rate. According to optimistic estimates [U.N. and I.F.R.R., 2002], the number of deployed personal service robots will grow from from 176,500 in 2001 to 2,021,000 in 2005—a stunning 1,145% increase. Personal service robots assist or entertain people in domestic settings or in recreational activities. Examples include robotic vacuum cleaners, lawn mowers, receptionists, robots assistants to elderly and handicapped people, wheelchairs, and toys. Figure 3 shows two example from the medical sector: a robotic assistant to the elderly [Montemerlo et al., 2002, Schraft et al., 1998] and a robotic walker [Dubowsky et al., 2002, Glover et al., 2003, Lacey and Dawson-Howe, 1998, Morris et al., 2003]. Figure 3c depicts a series of humanoid robots developed with an eye towards domestic use. Another example is shown in Figure 4c: a robotic toy, popular through its use for robotic soccer [Kitano, 1998]. Personal service robots are just beginning to emerge. The sales of robotic toys alone is projected to increase by a factor of ten in the next four years. Many of these robots interact with people who, in general, possess no special skills or training to operate a robot. Finding effective means of interaction is therefore more crucial in this new market segment of robotic technology than in industrial robotics or professional service robotics.

The widely acknowledged shift from industrial to service robotics, and the resulting increase of robots that operate in close proximity to people, raises a number of research and design challenges. Some of these challenges are outside the scope of this article, such as those pertaining to safety and cost. From the HRI perspective, the most important characteristic of these new target domains is that service robots share physical spaces with people. In some applications, these people will be professionals that may be trained to operate robots. In others, they may be children, elderly, or handicapped people whose ability to adapt to robotic technology may be limited. The design of the interface, while dependent on the specific target application, will require substantial consideration of the end user of the robotic device. Herein lies one of the great challenges that the field of robotics faces today.

3 Robotic Autonomy

Autonomy refers to a robot's ability to accommodate variations in its environment. Different robots exhibit different degrees of autonomy; the degree of autonomy is often measured by relating the degree at which the environment can be varied to the mean time between failures, and other factors indicative of robot performance. Human robot interaction cannot be studied without consideration of a robot's degree of autonomy, since it is a determining factor with regards to the tasks a robot can perform, and the level at which the interaction takes place.

The three kinds of robotics are characterized by different levels of autonomy, largely pertaining to the complexity of environments in which they operate. It should come at little surprise that industrial robots operate at the lowest level of autonomy. In industrial settings, the environment is usually highly engineered to enable robots to perform their tasks in an almost mechanical way. For example, pick-and-place robots are usually informed of the physical properties of the parts the be manipulated, along with the locations at which to expect parts and where to place them. Driverless transportation vehicles in industrial settings often follow fixed paths defined by guide wires or special paint on the floor. As these examples suggest, careful environment engineering indeed minimizes the amount of autonomy required—a key ingredient of the commercial success of industrial robotics.

This picture is quite different in service robotics. While environment modifications are still commonplace—the satellite-based GPS system that helps outdoor robots determine their locations is such a modification—the complexity of service robot environments mandate higher degrees of autonomy than in industrial robotics. The importance of autonomy in service robotics becomes obvious in Figure 5a: This diagram depicts the trajectory of a museum tourguide robot [Burgard et al., 1999], as it toured a crowded museum. Had the museum been empty, the robot would have been able to blindly follow the same trajectory over and over again—just as industrial robots tend to repeatedly execute the exact same sequence of actions. The unpredictable behavior of the museum visitors, however, forced the robot to adopt detours. The ability to do so sets this robot apart from many industrial applications.

Autonomy-enabling technology has been a core focus of robotics research in the past decade. One branch of research is concerned with acquiring environmental models. An example is shown in Figure 5b, which depicts a 2-D map of a nursing home environment acquired by the robot in Figure 3b by way of its laser range finders. Such a 2-D map is only a projection of the true 3-D environment; nevertheless, paired with a planning system it is sufficiently rich to enable the robot to navigate in the absence of environmental modifications. Other research has focused on the capability to detect and accommodate people. In general, robots which operate in close proximity to people require a high degree of autonomy, partially because of safety concerns and partially because people are less predictable than most objects. It is common practice to endow service robots with sensors capable of detecting and tracking people [Schulz et al., 2001]. Some researchers have gone as far as devising techniques whereby robots learn about people's routine behavior and actively step out of the way when people approach [Bennewitz et al., 2003].

The type and degree of autonomy in service robotics varies more with the specific tasks a robots is asked to perform and the environment in which it operates. Personal robots tend to be etching at low-cost markets. As a result, endowing a personal robot with autonomy can be significantly more difficult than its more expensive professional relative. For example, the robotic dog shown in Figure 4 is equipped with a low-resolution CCD camera and an onboard computer whose processing power lags behind most professional service robots by orders of magnitude—which adds to the challenge of making it autonomous.

4 Human Robot Interfaces

Robots, like most other technological artifacts, require user interfaces for interacting with people. Interfaces for industrial robots tend to differ from interfaces for professional service robots, which in turn differ from that for personal service robots.

In industrial robotics, the opportunity for human robot interaction is limited, since industrial robots tend not interact directly with people. Instead, their operational space is usually strictly separated from that of human workers. Interface technology in industrial robotics is largely restricted to special purpose programming languages and graphical simulation tools [Nof, 1999], which have become indispensable for configuring robotic manipulators. Some researchers have developed techniques for programming robots through demonstration [Friedrich et al., 1996, Ikeuchi et al., 1996, Mataric, 1994, Schaal, 1997]. The idea here is that a human demonstrates a task (e.g., an assembly task) while being monitored by a robot. From that, the robotic device learns a strategy for performing the same tasks by itself.

Service robots require richer interfaces. Here we distinguish interfaces for indirect interaction with interfaces for direct interaction. For the sake of this article, I define indirect interaction to be the interaction that takes place when a person operates a robot: for example, a person may give a command which the robot then executes. A robotic surgeon interacts indirectly with a surgical robot, like the one shown in Figure 2b, in that the robot merely amplifies the surgeon's force. A nice way to distinguish indirect from direct interaction pertains to the flow of information and control: In indirect interaction, the operator controls the robot, which communicates back to the operator information about its environment and its task. In direct interaction, the information flow is bi-directional: information is communicated between the robot and people in both directions, and the robot and the person are interacting on "equal footing." An example is the robotic caregiver in Figure 3a, which talks to people and interacts with people in ways motivated by people's interaction with nurses. In particular, it asks questions, and it can also respond to questions asked by people. As a general rule of thumb, the interaction with professional service robots is usually indirect, whereas the interaction with personal service robots tends to be more of the direct kind. There are exceptions to this rule, such as the robotic vacuum cleaner in Figure 3c, which is a personal service robot whose interaction is entirely indirect.

There exists a range of interface technologies for indirect interaction. The classical interface is the master-slave interface, in which a robot duplicates the exact same physical motion of its operator. A recent implementation of this idea is given by Robonaut [Ambrose et al., 2001], a robot developed as a telepresence device on a space station. The goal of this project is to demonstrate that a robotic system can perform repairs and inspections of space flight hardware originally designed for human servicing. Some robots are operated remotely using interfaces familiar from RC cars [Casper, 2002]; others possess haptic displays and control interfaces [Ruspini et al., 1997].

In service robotics, the utility of direct interaction is much less established than that of indirect interaction. To study direct interaction, numerous research prototypes have been equipped with speech synthesizers and recognizers or sound-synthesizing devices. Some robots only generate speech but do not understand spoken language [Thrun et al., 2000]; others also understand spoken language [Asoh et al., 1997, Bischoff and Graefe, 2003] or use keyboard interfaces to bypass speech recognition altogether [Torrance, 1994]. Speech as output modality is easy to control and can be quite effective. Several researchers have reported excellent results for office robots, in which speakers were instructed with regards to vocabulary the robot was able to understand [Asoh et al., 1997]. Encouraging results have also been reported for a museum tour-guide robot that understands spoken commands in multiple languages [Bischoff and Graefe, 2003]; although to my knowledge none of these systems have been evaluated systematically with regards to the effectiveness of the speech interface. Experiments with a service robot in an elderly care facility [Roy et al., 2000], where subjects were not instructed about the robot's vocabulary, resulted in the finding that the speech interface was difficult to use: Only about 10% of the words used by the target group were in the robot's vocabulary. Further, anecdotal evidence suggests that the ability to talk can create a false perception of human-level intelligence [Nourbakhsh et al., 1999, Schulte et al., 1999].

A number of robots carry graphical screens capable of displaying information to the user [Nourbakhsh et al., 1999, Simons et al., 2003]. Researchers have used both regular and touch-sensitive displays [Nourbakhsh et al., 1999]. The information on the display may be organized in menus similar to the ones found on typical information kiosks. Some researchers have used the display to project an animated face [Simons et al., 2003], such as the one shown in Figure 4a. Gesture recognition [Kahn et al., 1996, Kortenkamp et al., 1996, Perzanowski et al., 2000, Waldherr et al., 2000] has also been investigated, as has gaze tracking [Heinzmann and Zelinsky, 1998, Zelinsky and Heinzmann, 1996] as an interface to refer to physical objects in a robot's workspace. Such interfaces have shown moderate promise for robots assisting severely handicapped people. A recent study has investigated modalities as diverse as head motion, breath expulsion, and electro-oculographic signals (eye motion recorded by measuring electrical activity in the vicinity of the eye) as alternative interfaces for handicapped people [Mazo et al., 2000].

There exist also interface technologies that are unique to robotics, in that they require physical embodiment. A classical example is that of a mechatronic face [Breazeal, 2003a, Breazeal, 2003b] shown in Figure 4b. This face is capable of exhibiting different facial expressing, such as smile, frown, surprise. It does this by moving actuated elements into position reminiscent of human muscular movement when expressing certain emotions. In the past decade, dozens of such faces have been developed for service robot applications, with varying degrees of dexterity and expressiveness. Many robotic faces are able to change the expression of the mouth and the eyes, emulating basic expressions such as smiling and frowning. The face shown in Figure 4b possesses fifteen independently actuated elements and consequently can express quite a range of different postures.

Some researchers have begun to explore the social aspects of service robotics. Humanoid robots, by virtue of their appearance and behavior, appeal to people differently than other technological artifacts, as a recent survey suggests [Fong et al., 2003]. Most research on sociable robots has focused on humanoid robots and robots with humanoid elements (see Figure 6a). A recent study [Kiesler and Goetz, 2002] found that the shape of the robot, and specifically the presence and absence of humanoid features, influences people's behavior towards the machine and their assumptions about its capabilities. Somewhat complementary is the work by Scassellati [Scassellati, 2000], who has investigated the use of humanoid robots to aid human development. His prototype robot uses its ability to track where a toddler is looking to create a model of gaze movement and focus of attention. Robots like his may be used in the future to help train autistic children to behavior more socially; however, we do not know yet if children will interact with a robot as they will with a person. More generally, it remains unclear whether we, the people who will ultimately interact with service robots on a daily basis, will seek social-style interactions with robots that parallel our interactions with other people.

5 Open Questions

Human robot interaction is a field in change. The technical developments in robotics have been so fast that the types of interactions that can be performed today differ substantially from those that were possible even a decade ago. Interaction can only be studied relative to the available technology. Since robotics is still in its infancy, so is the field human robot interaction.

What follows is a number of open questions that arose while writing this article:

- How effective is a mechatronic face over a simulated one? How effective is a humanoid torso over a non-humanoid one? Does the physical shape of a robot affect the way people interact with them, and are those changes beneficial or detrimental to the robot's task performance?
- What happens when robots work with or interact with groups of people? How do interfaces scale up to multiple "operators?"
- To what extent is progress in human robot interaction tight to progress in other core disciplines of robotics, such as autonomy? In which way will future advances in

robotic autonomy change the way we interact with robots?

- To what extent does robotics benefit from direct interactions, as opposed to the well-established indirect interaction? Assuming that interaction is merely a means to an end, does direct interaction improve the on-task performance, relative to indirect interaction? Does it make robots amenable to a broader range of problem domains (e.g., elderly people unfamiliar with robotic technology)? How will these interfaces be affected when service robots are as familiar to people as personal computers are today?
- Much of the recent commercial success of PDAs cam be attributed to a specialized alphabet—graffiti—designed to maximize the recognition accuracy. Will there be a graffiti for robotics? If so, what will is look like, and what will it be good for?
- When referencing an object, it is better to recognize gestures of a person pointing towards the object, or is it more effective to have a person use a touch-sensitive display of a camera image containing the object?

Many researchers in robotics now recognize that human robot interaction plays a pivotal role in personal service robotics [Rogers and Murphy, 2001]. What constitutes an appropriate interface, however, is subject to much debate. Whatever the answer, it is likely to change in the future. Some of these changes will arise from the continuing stream of advances in robotic autonomy, which inevitably shifts the focus of the interaction to increasingly higher levels. Some changes will arise as a result of changed expectations: as we get used to robots in our environments, we might develop interaction skills that are succinctly different from our interactions with other technological artifacts, or with people. Regardless of the specifics of these developments, I conjecture that human robot interaction will soon become a primary concern in the majority of robotic applications.

References

[Ambrose et al., 2001] Ambrose, R., Askew, R., Bluethmann, W., Diftler, M., Goza, S., Magruder, D., and Rehnmark, F. (2001). The development of the robonaut system for space operations. In *Proceedings of ICAR, Invited Session on Space Robotics*, Budapest.

[Asoh et al., 1997] Asoh, H., Hayamizu, S., Isao, H., Motomura, Y., Akaho, S., and Matsui, T. (1997). Socially embedded learning of office-conversant robot jijo-2. In *Proceedings of IJCAI-97*. IJCAI, Inc.

[Bennewitz et al., 2003] Bennewitz, M., Burgard, W., and Thrun, S. (2003). Adapting navigation strategies using motion patterns of people. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*.

[Bischoff and Graefe, 2003] Bischoff, R. and Graefe, V. (2003). Hermes: an intelligent humanoid robot, designed and tested for dependability. In *Proceedings of the International Symposium on Experimental Robotics (ISER)*, Sant'Angelo d'Ischia, Italy. Springer Tracts in Advanced Robotics (STAR).

[Blackmon et al., 1999] Blackmon, T., Thayer, S., Teza, J., Broz, V., Osborn, J., and Hebert, M. (1999). Virtual reality mapping system for chernobyl accident site assessment. In *Proceedings of the SPIE, Vol. #3644*, pages 338–345.

[Brady et al., 1998] Brady, K., Tarn, T., Xi, N., Love, L., Lloyd, P., Davis, H., and Burks, B. (1998). Remote systems for waste retrieval from the oak ridge national laboratory gunite tanks. *International Journal of Robotics Research*, 17(4):450–460.

[Breazeal, 2003a] Breazeal, C. (2003a). Emotion and sociable humanoid robots. *International Journal of Human Computer Interaction*. Forthcoming.

[Breazeal, 2003b] Breazeal, C. (2003b). Towards sociable robots. *Robotics and Autonomous Systems,*, 42(3–4):167–175.

[Burgard et al., 1999] Burgard, W., Cremers, A., Fox, D., Hähnel, D., Lakemeyer, G., Schulz, D., Steiner, W., and Thrun, S. (1999). Experiences with an interactive museum tour-guide robot. *Artificial Intelligence*, 114(1-2):3–55.

[Casper, 2002] Casper, J. (2002). Human-robot interactions during the robot-assisted urban search and rescue response at the world trade center. Master's thesis, Computer Science and Engineering, University of South Florida, Tampa, FL.

[Dubowsky et al., 2002] Dubowsky, S., Genot, F., and Godding, S. (2002). PAMM: A robotic aid to the elderly for mobility assistance and monitoring: A ”helping-hand” for the elderly. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, San Francisco, CA. ICRA.

[Durrant-Whyte, 1996] Durrant-Whyte, H. (1996). Autonomous guided vehicle for cargo handling applications. *International Journal of Robotics Research*, 15(5).

[Fong et al., 2003] Fong, T., Nourbakhsh, I., and Dautenhahn, K. (2003). A survey of socially interactive robots. *Robotics and Autonomous Systems*, 42(3–4):143–166.

[Friedrich et al., 1996] Friedrich, H., Münch, S., Dillman, R., Bocionek, S., and Sassin, M. (1996). Robot programming by demonstration (rpd): Supporting the induction by human interaction. *Machine Learning*, 23(2/3):5–46.

[Glover et al., 2003] Glover, J., Holstius, D., Manojlovich, M., Montgomery, K., Powers, A., Wu, J., Kiesler, S., Matthews, J., and Thrun, S. (2003). A robotically-augmented walker for older adults. Technical Report CMU-CS-03-170, Carnegie Mellon University, Computer Science Department, Pittsburgh, PA.

[Heinzmann and Zelinsky, 1998] Heinzmann, J. and Zelinsky, A. (1998). 3-D facial pose and gaze point estimation using a robust real-time tracking paradigm. In *Proceedings of the IEEE International Conference on Automatic Face and Gesture Recognition*, pages 142–147.

[Ikeuchi et al., 1996] Ikeuchi, K., Suehiro, T., and Kang, S. (1996). Assembly plan from observation. In Ikeuchi, K. and Veloso, M., editors, *Symbolic Visual Learning*, pages 193–224. Oxford University Press.

[Jones et al., 2002] Jones, H., Rock, S., Burns, D., and Morris, S. (2002). Autonomous robots in SWAT applications: Research, design, and operations challenges. In *Proceedings of the 2002 Symposium for the Association of Unmanned Vehicle Systems International (AUVSI '02)*, Orlando, FL. AUVSI.

[Kahn et al., 1996] Kahn, R., Swain, M., Prokopowicz, P., and Firby, R. (1996). Gesture recognition using the perseus architecture. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 734–741, San Francisco, CA.

[Kiesler and Goetz, 2002] Kiesler, S. and Goetz, J. (2002). Mental models of robotic assistants. In *Proceedings of the Conference on Human Factors in Computing Systems (CHI 2002)*, Minneapolis, Minnesota. ACM.

[King and Weiman, 1990] King, S. and Weiman, C. (1990). Helpmate autonomous mobile robot navigation system. In *Proceedings of the SPIE Conference on Mobile Robots*, pages 190–198, Boston, MA. Volume 2352.

[Kitano, 1998] Kitano, H., editor (1998). *Proceedings of RoboCup-97: The First Robot World Cup Soccer Games and Conferences*, Berlin. Springer Verlag.

[Kortenkamp et al., 1996] Kortenkamp, D., Huber, E., and Bonasso, P. (1996). Recognizing and interpreting gestures on a mobile robot. In *Proceedings of AAAI-96*, pages 915–921. AAAI Press/The MIT Press.

[Lacey and Dawson-Howe, 1998] Lacey, G. and Dawson-Howe, K. (1998). The application of robotics to a mobility aid for the elderly blind. *Robotics and Autonomous Systems*, 23:245–252.

[Mataric, 1994] Mataric, M. (1994). Learning motor skills by imitation. In *Proceedings of the 1998 AAAI Spring Symposium*.

[Mazo et al., 2000] Mazo, M., , Garcia, J. C., Rodriguez, F. J., Urena, J., Lazaro, J. L., and Espinosa, F. (2000). Integral system for assisted mobility. *IEEE Robotics and Automation Magazine*, 129(1-4):1–15.

[Montemerlo et al., 2002] Montemerlo, M., Pineau, J., Roy, N., Thrun, S., and Verma, V. (2002). Experiences with a mobile robotic guide for the elderly. In *Proceedings of the AAAI National Conference on Artificial Intelligence*, Edmonton, Canada. AAAI.

[Morris et al., 2003] Morris, A., Donamukkala, R., Kapuria, A., Steinfeld, A., Matthews, J., Dunbar-Jacobs, J., and Thrun, S. (2003). A robotic walker that provides guidance.

In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*.

[Nof, 1999] Nof, S., editor (1999). *Handbook of Industrial Robotics, 2nd Edition*. John Wiley & Sons.

[Nourbakhsh et al., 1999] Nourbakhsh, I., Bobenage, J., Grange, S., Lutz, R., Meyer, R., and Soto, A. (1999). An affective mobile robot with a full-time job. *Artificial Intelligence*, 114(1–2):95–124.

[Perzanowski et al., 2000] Perzanowski, D., Adams, W., Schultz, A., , and Marsh, E. (2000). Towards seamless integration in a multi-modal interface. In *Proceedings of the Workshop on Interactive Robotics and Entertainment (WIRE)*, pages 3–9, Pittsburgh, PA. Carnegie Mellon University. Also available as NCARAI Report AIC-00-003.

[Reinemann and Smith, 2000] Reinemann, D. and Smith, D. (2000). Evaluation of automatic milking systems for the united states. In Hogeveen, H. and Meijering, A., editors, *Robotic Milking, Proceedings of the international symposium*, Lelystad, The Netherlands.

[Rogers and Murphy, 2001] Rogers, E. and Murphy, R. (2001). Final report for DARPA/NSF study on human-robot interaction. On the Web at http://www.aic.nrl.navy.mil/hri/nsfdarpa/index.html.

[Roy et al., 2000] Roy, N., Pineau, J., and Thrun, S. (2000). Spoken dialogue management using probabilistic reasoning. In *Proceedings of the 38th Annual Meeting of the Association for Computational Linguistics (ACL-2000)*, Hong Kong.

[Ruspini et al., 1997] Ruspini, D., Kolarov, K., and Khatib, O. (1997). The haptic display of complex graphical environments. In *Computer Graphics (SIGGRAPH 97 Conference Proceedings)*, pages 345–352. ACM SIGGRAPH.

[Scassellati, 2000] Scassellati, B. (2000). How robotics and developmental psychology complement each other. In *Proceedings of the NSF/DARPA Workshop on Development and Learning*, Lansing, MI. Michigan State University.

[Schaal, 1997] Schaal, S. (1997). Learning from demonstration. In *Advances in Neural Information Processing Systems 8*, Cambridge, MA. MIT Press.

[Schraft et al., 1998] Schraft, R. D., Schaeffer, C., and May, T. (1998). Care-o-bot: The concept of a system for assisting elderly or disabled persons in home environments. In *IECON: Proceedings of the IEEE 24th Annual Conference (Volume 4)*, pages 2476–2481.

[Schulte et al., 1999] Schulte, J., Rosenberg, C., and Thrun, S. (1999). Spontaneous short-term interaction with mobile robots in public places. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*.

[Schulz et al., 2001] Schulz, D., Burgard, W., Fox, D., and Cremers, A. (2001). Tracking multiple moving targets with a mobile robot using particles filters and statistical data association. In *Proceedings of the IEEE International Conference on Robotics and Automation*, Seoul, Korea.

[Simons et al., 2003] Simons, R., Goldberg, D., Goode, A., Montemerlo, M., Roy, N., Sellner, B., Urmson, C., Schulz, A., Abramson, M., Adams, W., Atrash, A., Bugajska, M., Coblenz, M., MacMahon, M., Perzanowski, D., Horswill, I., Zubek, R., Kortenkamp, D., Wolfe, B., Milam, T., and Maxwell, B. (2003). GRACE: An autonomous robot for the AAAI robot challenge. *AI Magazine*.

[Thrun, 2002] Thrun, S. (2002). Robotics. In *Chapter 25 of "Artificial Intelligence: A Modern Approach (second edition)" by S. Russell and P. Norvig*. Prentice Hall, Englewood Cliffs, NJ.

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

[Thrun et al., 2003] Thrun, S., Hähnel, D., Ferguson, D., Montemerlo, M., Triebel, R., Burgard, W., Baker, C., Omohundro, Z., Thayer, S., and Whittaker, W. (2003). A system for volumetric robotic mapping of abandoned mines. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*.

[Torrance, 1994] Torrance, M. C. (1994). Natural communication with robots. Master's thesis, MIT Department of Electrical Engineering and Computer Science, Cambridge, MA.

[U.N. and I.F.R.R., 2002] U.N. and I.F.R.R. (2002). *United Nations and The International Federation of Robotics: World Robotics 2002*. United Nations, New York and Geneva.

[Waldherr et al., 2000] Waldherr, S., Thrun, S., and Romero, R. (2000). A gesture-based interface for human-robot interaction. *Autonomous Robots*, 9(2):151–173.

[Zelinsky and Heinzmann, 1996] Zelinsky, A. and Heinzmann, J. (1996). Human-robot interaction using facial gesture recognition. In *Proceedings of the International Workshop on Robot and Human Communication*, pages 256–261, Silver Spring, MD. IEEE.

[Zivanovic and Davies, 2000] Zivanovic, A. and Davies, B. (2000). A robotic system for blood sampling. *IEEE Transactions on Information Technology in Biomedicine*, 4(1).

(a) Industrial welding robot (b) Robot for transporting containers

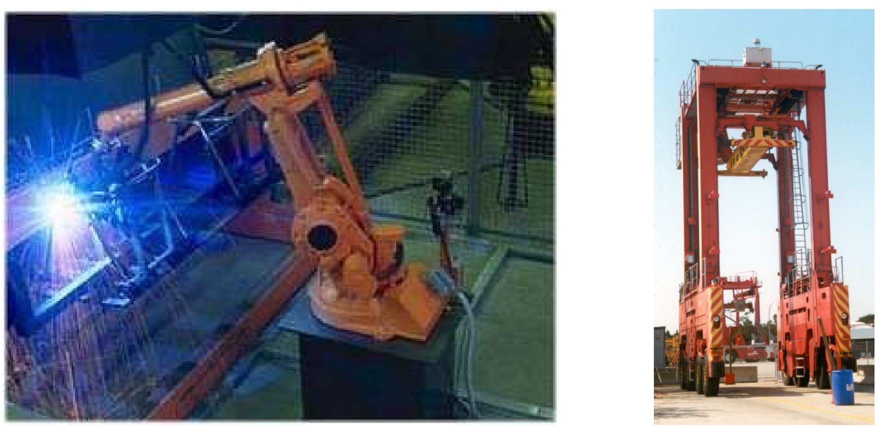

Figure 1: Industrial Robots: (a) a typical welding robot. (b) Autonomous robot for transporting containers on a loading deck [Durrant-Whyte, 1996].

(a) Hospital delivery robot (b) Surgical robot (c) Milking robot

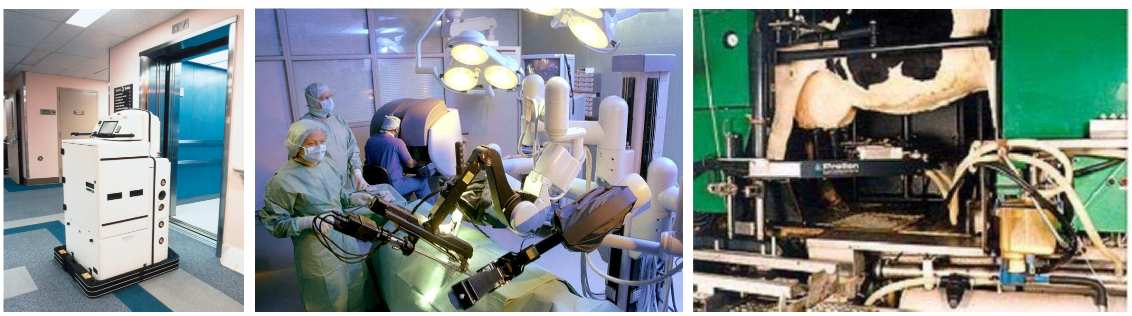

Figure 2: Professional service robots: (a) The Helpmate hospital delivery robot; (b) a surgical robot by Intuitive Surgical; and a (c) milking robot [Reinemann and Smith, 2000].

(a) Robotic caregiver (Nursebot)

(b) Robotic walker

(c) Robotic Vacuum Cleaner

Figure 3: Personal service robots. (a) The Nursebot, a prototype personal assistant fo rthe elderly; (b) robotic walker; (c) robotic vacuum cleaner "Roomba" (by iRobot, Inc.).

(a) Synthetic face on robot display (b) Actuated robotic face (c) Robotic animal

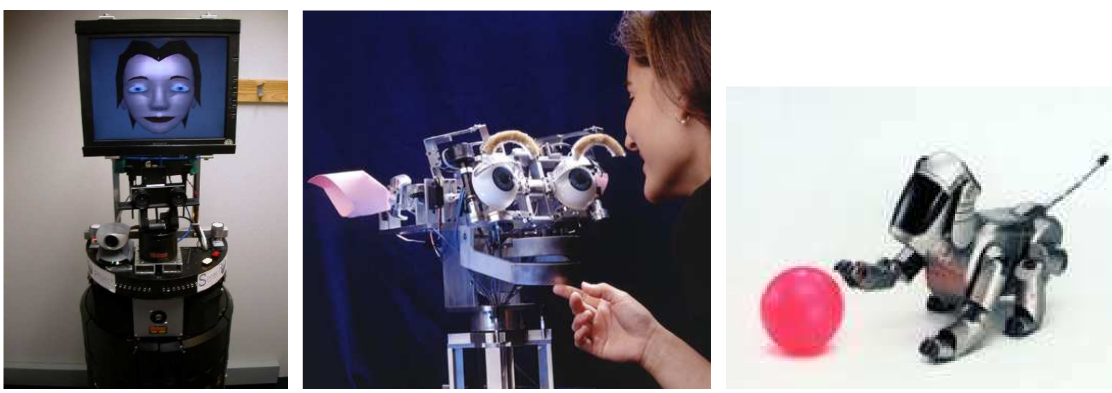

Figure 4: Faces in robotics: (a) animated face on the robot Grace [Simons et al., 2003]; (b) mechatronic face of KISMET [Breazeal, 2003a]; and (c) The Sony AIBO robotic dog.

(a) (b)

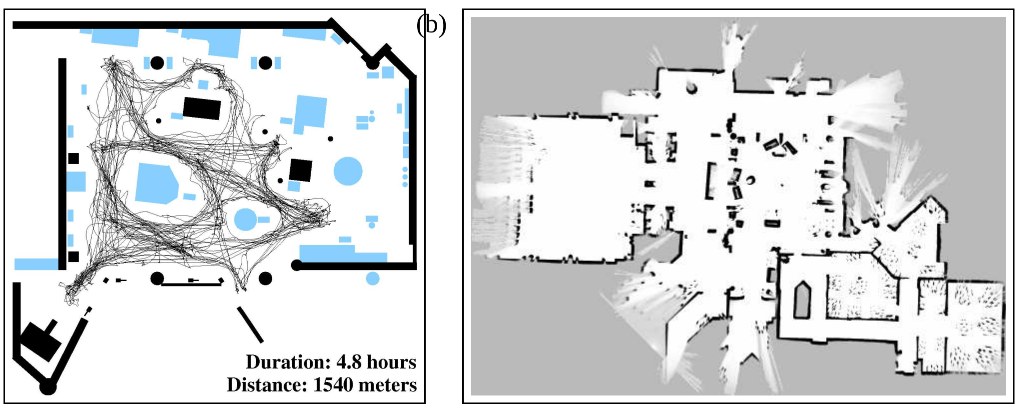

Figure 5: Left: Path taken by an autonomous service robot; (b) 2-D Map "learned" by a robot.

(a) Humanoid robots by Honda (b) Robot with robotic face

Figure 6: (a) Humanoid Robots Asimo and P3 by Honda. (b) Tourguide robot Minerva, with an actuated humanoid face but a non-humanoid torso.