Towards Personal Service Robots for the Elderly

Nicholas Roy, Gregory Baltus, Dieter Fox, Francine Gemperle, Jennifer Goetz, Tad Hirsch
Dimitris Margaritis, Mike Montemerlo, Joelle Pineau, Jamie Schulte, Sebastian Thrun

Computer Science and Robotics
Carnegie Mellon University
http://www.cs.cmu.edu/∼nursebot

Abstract

This paper describes the state-of-the art of a large-scale project, aimed towards the development of personal service robots for the elderly population. Taking care of elderly and chronically ill people is one of the major challenges currently faced by society. Needs range from support in manipulation to assisting those with dementia and cognitive impairment. To respond to this challenge, we have developed a first prototype robot. Using natural language, the robot can provide information related to activities of daily living obtained from the Web. It also enables remote care-givers to establish a "tele-presence" in people's home, by relaying back video and audio stream through the Next Generation Internet. The paper describes this early prototype, and it lays out our research agenda towards building service robots for the elderly.

1 The Problem

At the turn of the millennium, the number of elderly in need of care is increasing dramatically. Today, more than 12.5% of the US population is over 65 years in age. As the baby-boomer generation approaches the retirement age, this number will increase significantly. By 2030, more than 15% of the population will be 65 and over. According to information summarized from the U.S. Bureau of Census, the population of people 85 and older is expected to increase 39.3 percent by the year 2000 and 33.2 percent between the years 2000 and 2010. Current living conditions for the majority of elderly people are already alarmingly unsatisfactory, and situation will worsen in the future. According to (US Department of Health and Human Services 1999), nearly 9 percent of non-institutionalized persons 70 years of age and over were unable to perform one or more activities of daily living such as bathing, dressing, using the toilet, and getting in and out of bed or chairs.

At the same time, the nation is facing an explosion of costs in the health-care sector. Current nursing home costs range between $30,000 and $60,000 annually. Over the last decade along, costs have more than doubled (in adjusted dollars). The dramatic increase of the elderly population along with the explosion of costs pose extreme challenges to society. The current practices of providing care for the elderly population is already insufficient. Undoubtedly, this problem will multiply over the next decade. Thus, as a society we need to find alternative ways of providing care to the elderly and chronically ill population. Such ways must not only lower the costs. They must also increase the comfort of living, and approach people with the level of dignity that our elderly deserve.

Robotic technology, at the same time, is going through major revolutions. Sparked by a dramatic increase of computation per dollar, and by substantial decreases in the costs of major sensor technologies (e.g., cameras), we are now closer than ever to the goal of intelligent service robots than can assist people in their daily living activities. In the last few years, service robots (Engelberger 1999) were successfully fielded in hospitals (King & Weiman 1990), museums (Burgard *et al.* 1999), and office buildings/department stores (Endres, Feiten, & Lawitzky 1998), where they perform janitorial services, deliver, educate, or entertain (Schraft & Schmierer 1998). Robots have also bee developed for guiding blind people (Lacey & Dawson-Howe 1998). Robotic aids for the elderly have been developed, but many of these robotic aids are mechanical aids (Song *et al.* 1998; Dario, Laschi, & Guglielmelli 1999; Takahashi *et al.* 1999) that put the cognitive load on the patient side. There has been little research to date in terms of assisting elderly people with cognitive tasks, such as remembering medication schedules. However, human-robot interaction, autonomous systems and planning have seen major developments recently. The time is ripe to leverage the various technologies into the lives of elderly people, where the need for personal assistance is larger than in any other age group.

This paper described initial results obtained by CMU's "Nursebot project." The goal of this project is the development of personal robotic aids that serves five primary functions:

- **Cognitive prosthesis.** A large fraction of the elderly population suffers from varying degrees of dementia. The inability to remember can have severe consequences. For example, subjects may forget to take medicine; they might forget to see the bathroom, etc. When conditions become too severe, patients need regular supervision in carrying out their daily activities, which often means moving into a nursing home. People also might use a robot for lesser purposes, such as finding out what's on TV. Reminding is an important (and time-consuming) activity in a health-care professional's life.
- **Safeguarding.** As elderly people become physically and

cognitively impaired, the private home can pose substantial risks. For example, accidents relating to falling can have, if undetected by others, severe consequences (up to a patient's death). loss of stability is a leading problem for independently living elderly people. By reducing such risks through systematic monitoring and safeguarding, the move into dependent living (e.g., nursing home) can be delayed.

- **Systematic Data Collection.** A key problem currently faced in the health care sector is the inability to collect data for people living in private homes. Such data include statistics on medication (when did the person take what?), daily living activities, and factors related to the prediction of specific medical risks (blood sugar, leg diameter). Assuming that the necessary mechanisms are in place to guarantee the privacy of such information, such data can be of tremendous value for professional care givers in their diagnosis and selection of treatment.
- **Remote tele-medicine.** In the US, home visits by healthcare professionals (e.g., doctors) are extremely rare, due to the high costs involved. The idea of robotic telepresence is to use Internet technology to relay live video and audio stream from the doctor's office to the patient's living room, thereby enabling the doctor to establish a tele-presence in the patient's home. The ability to move (and manipulate) provides an enhanced degree of flexibility currently lacking in video-conferencing and other, competing alternatives.
- **Social interaction.** Finally, the vast majority of independently living elderly people is forced to live alone, and is deprived of social interaction. Social engagement can significantly delay the deterioration and health-related problems. While robots cannot replace humans, we seek to understand the degree at which robots can augment humans, either by directly interacting with the person, or by providing a communication interface between different people that is more usable than current alternatives.

To accommodate these needs, we are currently developing a first generation personal service robot specifically targeted at people with mild forms of dementia and other physical inabilities (e.g., low blood pressure). This paper describes the current system design, along with initial results obtained in a controlled experiment.

2 Hardware Design

The current prototype robot, called *Flo* (in honor of *Florence Nightingale*) is shown in Figure 1. Flo is built on top of a Nomad Scout differential drive mobile base, equipped with 16 ultrasonic range finders. The custom-made robot is equipped with a SICK PLS laser range finder, capable of measuring distances at an angular resolution of one degree and a spatial resolution of 5 cm, within a planar perceptual field that covers a 180 degree range. Flo is also equipped with two on-board PCs, connected to the Internet via a 2mbit/sec wireless Ethernet link manufactured by BreezeCom. A bright, touch-sensitive color display is mounted conveniently at approximate eye height for for sitting people. On top of that, FLO possesses an actuated face that

Figure 1: Side view of Flo, the robot. The robot is equipped with a touch-sensitive display, a laser range finder, an array of 16 sonar sensors, and two on-board PCs.

enables it to show different facial expressions by modifying the angle of its mouth and that of its eyebrows. The face is mounted on a 2D pan/tilt unit (by Directed Perception), capable of swiveling the face at high speeds. Additionally, the eyes are motorized, which enables them to saccade when tracking a person's face. Flo's eyes are white-balanced color CCD cameras with an approximate aperture angle of 100 degrees. The cameras are connected to a pair of frame grabbers and JPEG encoders for image processing and high-bandwidth communication. Flo is also equipped with a speaker system and a microphone, necessary for recording and synthesizing speech and other acoustic signals (e.g., music). Flo's battery lifetime is approximately 45 minutes. The robot currently lacks a mechanism for connecting itself to a battery charger, making it necessary that a human assists the robot in operation.

3 Software

At the current point, Flo consists of four major software subsystems, each of which is designed with a specific goal in mind for assisting the elderly.

3.1 Tele-Presence Interface

One of the most important goals of a robot assistant is not to supplant communication between the users and other people such as nurses and doctors, but to facilitate it. Rather than replacing a nurse, Flo needs to allow real nurses to monitor and interact with the user. Furthermore, while having family and relatives visit can often be difficult, Flo can allow for virtual visits using the tele-presence, in this way increasing the user's contact with the outside world at relatively little effort.

Figure 2: Flo interacts with a person.

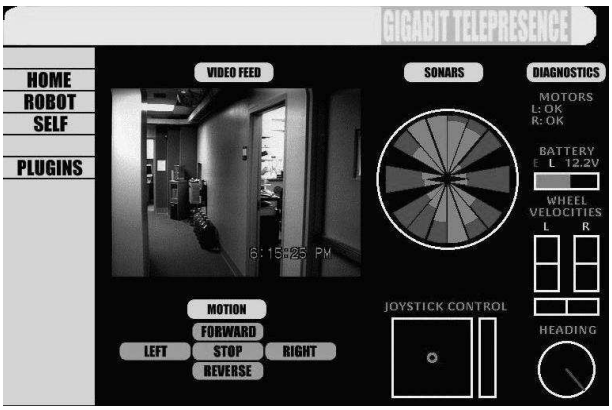

Figure 3: A picture of the tele-presence interface, at the remote console.

The tele-presence interface consists of a camera and microphone on-board Flo, that transmit the video and audio signal to a remote station. Both camera and microphone are mounted inside Flo's head, to provide as robot-centric a representation as possible. The video feed is compressed into a JPEG feed on board, and then both signals are transmitted to a local base-station over the wireless Ethernet, and then to a remote station. At the remote station the JPEG is decompressed and synchronized with the audio before they are played back on the remote computer's screen and speakers. JPEG cards are used instead of the higher compression afforded by MPEG cards, because MPEG compression causes a significant latency in the signal transmission, usually around .25sec and enough to be annoying [1].

At the moment, the wireless bandwidth does not support bi-directional tele-presence, although the recent availability of 11 Mbps wireless Ethernet will eliminate this problem, making wireless bi-directional tele-presence a reality. Currently, an audio signal can be transmitted from the remote station to on-board the robot where it can be played out the Flo's speakers. Thus, it is possible to have a live conversation with someone's presence "embedded" in Flo.

In addition to the communication modes, the tele-presence offers control of the robot to the remote user. Using a joystick, a health care giver, friend or relative can drive the robot around the user's rooms, and also direct the robot's gaze by controlling the head configuration. The safety of the robot is guaranteed by the robot's navigation software, which limits the robot's velocity so as to avoid collisions with obstacles. Figure 3 shows the graphical interface, which is run inside a Web browser. The interface displays the video steam, along with the robot's sensor readings (sonar only). It offers and easy-to-use joystick interface for remote robot operation.

[1]MPEG compression uses information from previous and future frames, and therefore needs to acquire several frames before it can begin compression.

3.2 Speech Interface

One of the major goals that has dictated the design process has been to develop a robot that allows the most natural interaction between the users and the robot. Elderly people often have difficulties interacting through unfamiliar means, such as keyboards and computer screens. It is therefore of great importance that the robot communicates in ways familiar to elderly people. To that end, spoken interaction with the robot is absolutely essential.

Flo possesses a real-time speech interface. The speech recognition system is based on CMU's SPHINX II system (Lee 1989; Ravishankar 1996). This system has the principle virtues of being speaker-independent, and requiring no pre-training by any user. Furthermore, most existing speech recognition systems require very small distances between the speaker and the microphone for reasonable recognition rates, however, our experience has shown that we can control the robot across a lab. SPHINX is capable of handling vocabularies of thousands of words, but the command-and-control tasks that Flo is predicted to perform do not require large vocabularies. At the moment, Flo's vocabulary consists of approximately one hundred words, enabling it to understand a variety of questions relating daily living activities such as inquiries for the television program and the weather forecast.

The speech recognition system is controlled by a dialog manager that generates the appropriate response, based on the observed utterance from the user. Since much of the speech around the robot is assumed not to be directed to the robot, it is necessary to signal the attention of the dialog manager by using the phrase "Flo" in each utterance, much as one might call a nurse over before talking to her. The dialog manager is currently based on keyword-spotting over the utterance strings, although more sophisticated techniques are in development using Markov Decision Process algorithms. One of the advantages of using probabilistic dialog managers is the ability to handle varying recognition rates, and to increase the managers ability to detect out-of-vocabulary words. Natural Language Processing can be used for dialog management, however, our experience has

Domain	# of possible responses
Weather	6
TV Schedule (ABC, NBC, CBS)	4
Appointment Calendar	4
MP3 Player	9
Time, Date & Location	3
Miscellaneous	4

Table 1: The domains of the dialog manager.

shown that NLP is not necessary for the domains we are addressing. Table 1 lists the information domains that the dialog manager is capable of processing.

The speech synthesis system is Festival, from the University of Edinburgh (Black, Taylor, & Caley 1999) system allows for producing a waveform for any English-language text, in a variety of voices, both male and female, and a variety of accents. When Flo speaks, she produces output to the screen that closely resembles her spoken output, for clarity and for users who may have hearing loss.

The dialog manager has a connection to a number of external sources of information, such as the World Wide Web, and thus is able to answer questions on a number of topics. For example, Flo is able to warn the user of impending bad weather, and can serve as a rudimentary TV guide. Flo also can consult an electronic datebook, reminding the user to take their medication, or that it is time to visit the doctor. Future plans for Flo's dialog include allowing the robot to answer the phone, control the TV and VCR using infrared transmitters, and to control many of the living areas lights and appliances using wireless X10 technology.

The touch-sensitive display is also integrated with the dialog manager; the touch-sensitive display provides input to the dialog manager as if it were also the speech recognizer. When the user presses the "medication" button, this is considered to be equivalent to a user's verbal request to take their medication. The principle difference between information provided by the touch screen and speech recognizer is the very low error rate on the information from the touch-screen. Each button on the screen corresponds to some verbal request that the user can also make.

3.3 Face Finding and Tracking

As mentioned above, Flo's face is equipped with two color cameras actuated by two independent servo motors (one per camera).

Flo uses a neural network approach for face detection, developed by Henry Rowley and colleagues (Rowley, Baluja, & Kanade 1998). This algorithm scans the image using a neural network trained to detect faces in camera images. It reliably finds faces under a wide range of viewing and lighting conditions. Unfortunately, the face finding algorithm requires approximately four seconds per image on Flo's on-board computers. Once a face has been found, a fast color-based tracking algorithm tracks the face at a rate of 15 frame/sec. This rate is sufficient to track people's faces even when they are moving rapidly. While tracking a person's face, the cameras are continually adjusted to keep the person centered in the camera image. Whenever the angle of the cameras surpasses a certain threshold (30 degrees), the whole head is rotated so that the cameras (eyes) can move back to its canonical position.

The ability to find and track faces is important for several reasons. It enables to direct the robot's sensors (in particular, its microphone and its cameras) into the direction of the person. This is important for the speech interface, whose recognition accuracy depends crucially on the strength of the speech signal. While the typical user of speech technology has no difficulties speaking into a microphone at the appropriate distance, cognitive impairment and technology barriers often make it difficult for elderly people to speak into microphones. Thus, the ability to direct the microphone automatically is important.

The ability to track faces is also important for the telepresence interface, specifically when interacting with people that move (or through a moving robot). The current joy-stick interface does not enable the user to *simultaneously* control the robot motion and the camera direction; in fact, the cognitive load of controlling these devices simultaneously is probably too high for health-case professionals without excessive training. Thus, the ability to track the face is essential for being able to interact with a person while the robot is in motion.

Finally, the current face tracking mechanism gives people a feeling of awareness. To an observer whose face is tracked visually, the coordinated motion of "eyes" and the "neck" resembles that of a person. The emotional ramifications are important, since one of our goals is to understand to which extent a robot can become a "social" companion for elderly people.

3.4 Navigation

Flo inherits its navigation system from a series of mobile robots, previously developed the Robot Learning Laboratories at CMU and the University of Bonn, Germany. These include the tour-guide robots Minerva (Thrun *et al.* 1999) and Rhino (Burgard *et al.* 1999), which successfully exhibited reliable navigation through crowds in unmodified public places (e.g., museums). The description of the navigation system is available elsewhere; thus, we will not describe it here in any depth.

Functionally, the navigation system enables the robot to navigate safely to arbitrary target locations in indoor environments. It does this by first learning a map of the environment, which is represented through a 2D occupancy grid map (Elfes 1989) (see (Thrun, Burgard, & Fox 1999) for a recent extension to 3D building models). Figure 4 shows a learned map of a large-scale indoor environment (left diagram). This specific map stems from a museum in San Jose. On the right side, an architectural drawing of the building is shown for comparison. Maps can be augmented by hand to attach symbolic names to specific places, and to constrain the motion in regions that might cause hazard to the robot or the person (such as: open staircases that the robot currently cannot detect through its sensors). Mapping environments of the size shown here takes in the order of 15 minutes; all software runs in real-time. The ability to learn accurate maps in real-time makes it extremely easy to "install" a robot in a

Figure 4: Left: Map of the environment shown graphically on the right. This map covers a large open area in the Tech Museum in San Jose, CA. The floor plan was developed by the building's designers; not every line therein corresponds to an actual obstacle (and vice versa). As can be seen, the map is accurate.

new home. Homes do not have to be modified in any way to facilitate robot navigation.

Once a map has been acquired, it enables the robot to maintain an accurate sense of its location in x-y space, along with its heading direction θ. This enables the robot to move to arbitrary target location, and to direct its camera towards items of interest (e.g., stove, the bathroom door). By differencing sensor readings and the learned map, the robot can detect changes in the environment, and in particularly the location of people. The ability to find people is essential for the envisioned care scenario in multiple ways: It facilitates the interaction with the person, for which knowledge of the location of a person is essential. It also makes it possible to analyze people's behavior to detect abnormal situations that require the attention of a care-giver (e.g., a person fell and is unable to get up by himself), and to systematically analyze peoples daily routines. The latter ability to gather data on people's behavior is important. Currently, health care professionals often lack critical data to assess the health status and potential risks for independently living elderly people. For example, people that suffer from low blood pressure are often advised to pause and catch breath after strenuous activities such as raising from a chair; however, whether or not a person adheres to these rules is typically unknown to the health care professional. Knowledge about such could be essential to improve the quality of life and risk assessment for elderly people.

4 Current Project Status

The first hardware prototype shown in Figure 1 has been developed and evaluated. The purpose of this prototype was to establish an "existence proof" to our project partners from the University of Pittsburgh's School of Nursing and CMU's Design Department, and to obtain their feedback on the design of robots for nursing. A secondary goal was to integrate the software described above into a running system. While the robot has not yet been tested with elderly people, feedback by our partners has identify a range of opportunities for improvement.

We also have integrated the software components described above using TCX (Fedor 1993) for inter-process communication. Table 2 shows a "typical" dialogue between the user and the robot. To answer the user's questions, Flo queries the NBC's and CNN's web sites for TV and weather information. Since phrases must begin with the word "Flo," Flo almost never responds to language tokens not directed at the robot. The current repertoire is sufficiently limited to guarantee high recognition accuracy, even if speakers deviate from the pre-programmed syntax. However, while our tests have included non-native speakers, no actual experiments with elderly have been conducted.

Figure 2 shows a person interacting with the robot in one of CMU's corridor. In our own experiments tele-operating the robot through the hallways, we found it extremely easy to navigate the robot and engage it in interactions with people. However, no experiments with health care professionals have been conducted.

5 User Feedback

Flo's "user interface" was recently evaluated in a systematic study involving 10 individuals (robotics graduate students). The students were chosen so that their prior exposure to this project was minimal. They were asked to communicate with the dialog manager without any instruction, although they

User:	*Flo, hello.*
Flo:	*Hello. What can I do for you.*
User:	*Flo, what's the weather?*
Flo:	*Do you mean the today's weather, tomorrow's weather, or the current weather?*
User:	*Flo, the current weather.*
Flo:	*It is currently raining with a temperature of 54 degrees.*
User:	*Flo, thanks.*
User:	*Flo, what's on NBC today?*
Flo:	*NBC is currently airing the Today Show.*
User:	*Flo, good-bye.*
Flo:	*Good-bye.*

Table 2: Example dialogue Flo's texts are also displayed on the screen. When answering the question of what's on TV, the robot displays the program for the next hours.

were informed before the experiment began what were the general areas of information that Flo contained. Although the speech recognition system was running, the dialog manager also contained a hidden human operator in a "wizard-of-oz" scenario, in case of dramatic failure of the speech recognition system. This precaution was in general not necessary.

The subjects were able to extract useful information, although a number of subjects reported uncertainty as to when the robot was "thinking" (e.g., retrieving information from the web), as opposed to merely waiting to be addressed, indicating that more subtle forms of feedback need to be added. Furthermore, some of the users addressed the robot in unexpected ways, indicating a need for a richer vocabulary and dialog manager. Certain simple functionality also needs to be added, such as asking the robot to repeat itself.

We also performed quantitative analysis of the speech recognition and dialog management systems, over the sample dialogs acquired during the user testing. These results are summarized in Table 3. The word recognition rate was approximately 50%, largely because the vocabulary of the system was not large enough. The dialog manager performance was approximately the same, performing the correct action 50% of the time. Of the 87 errors, about half were as a result of incorrect recognition, and about half of the errors were a result of the users making a request that the dialog manager could not fill (i.e., asking for information the dialog manager did not have). In these cases, the dialog manager performed a best-effort action to fill the request (i.e, returning the information that best matched the request). The nature of the keyword spotting dialog management is that it is difficult to recognize when a request cannot be filled, often because the request lies outside the domains of expertise (and hence the outside the vocabulary) of the system.

Further analysis showed the effects of superior speech recognition. In a second experiment, the text of all the speech samples was compiled and used to build a new speech model with a larger corpus, with approximately twice the vocabulary. The speech samples were then re-processed by the speech recognition system and dialog manager offline; the results are summarized in Table 4. The word recognition rate increased to 83.7%, however, the dialog manager performance remained at about 50%.

Speech Recognition
17/172 (9.9%) exact sentence matches
593/1182 (50.2%) word recognition rate
Dialog manager
85/172 (49.4%) correct actions
41/172 (23.8%) recognized wrong request
46/172 (26.7%) performed best-effort action

Table 3: Results of dialog management on user samples, using basic corpus.

Most importantly, the increase in the speech recognition performance resulted in a dramatic decrease in actions that did not satisfy the user's request. In the majority of cases in both analyses, such incorrect recognition resulted in inappropriate responses. This was a source of substantial annoyance to the users, and therefore it is re-assuring that the level of usability of the system can be boosted simply by enlarging the corpus. It is worth noting that such improvements will only take the system so far, as indicated by the relatively minor increase in the overall system performance. The dialog manager clearly needs to be able to handle a wider range of requests than was originally anticipated, and also (perhaps most importantly) the dialog manager has to be able to recognize when it cannot fulfill a request. These two requirements are driving further development of the dialog manager.

6 Current and Future Research

Based in the initial feedback from our project partners and our first user study, we have begun developing a second, improved robot platform. In particular, the next generation will be equipped with a removable basket at its front. We also are integrating a handle that provides support for people with stability problems. This handle is *not* meant as a walking aid; instead, it will be equipped with a touch sensor that will stop the robot as soon as a person holds onto it. Finally, we plan to add an additional rotational degree of freedom to increase the robot's maneuverability in tight spaces. This robot is currently being developed in collaboration with CMU's Design Department.

In collaboration with the School of Nursing of the University of Pittsburgh, we are currently developing a detailed "script," laying out in detail modes of interaction between nursing robots and people. Finally, we are at the verge of integrating the University of Pittsburgh's system for intelligent scheduling and planning, with the goal of developing an intelligent aid that intelligent management support of daily living activities; in particular intelligent reminding and scheduling.

7 Discussion

This paper reported the initial design and results of a mobile robot aimed at the elderly population. Recognizing the importance of providing care for elderly, we are currently developing a mobile robot that will provide a range of services to non-institutionalized elderly people.

Speech Recognition
83/172 (48.3%) exact sentence matches
989/1182 (83.7%) word recognition rate
Dialog manager
91/172 (52.9%) correct actions
69/172 (40.1%) recognized wrong request
12/172 (6.9%) performed best-effort action

Table 4: Using the revised corpus, results of dialog management on user samples.

We believe that mobile robots are particularly well-suited to the task of providing assistance to cognitively-impaired elderly people. Research has shown that non-mobile devices such as PDAs are less effective for assisting elderly people because the user has to remember to carry the device at all times. Furthermore, small devices are limited in the range of sensing modalities they can support. Evidence indicates that hand-held devices increase compliance with prescribed medication only 5% – the total cost of non-compliance with medication is approximately $1 billion US per year (Dunbar-Jacob 2000).

The system design integrates existing systems from previous mobile robots (Thrun *et al.* 1999) and other research areas (Ravishankar 1996), however, several problems remain unsolved in the current robot. The dialog management problem has been addressed partially, however, we would like the system to tailor its dialogs to individual user preferences automatically. The POMDP-based approach has eliminated the need for full-scale natural language processing (NLP) in the current dialog domains, however, the problem of how to scale the POMDP approach to larger and more complex domains is unsolved. Finally, the user interface design has many open problems with respect to the specific user population the robot will be facing.

A secondary goal of this paper is to make robotics researchers aware of a unique opportunity to develop personal service robots with high societal impact. We firmly believe that our current research only scratches the surface of this enormous challenge: Using personalized robotic technology to assist elderly and chronically ill people.

Acknowledgments

We are grateful to the following individuals, who all have contributed invaluably to this research: Tsuhan Chen, Jie Yang, Martha Pollack, Jacquelin Dunbar-Jacob, Joan Rogers, Sandy Engberg, Mike Vande Weghe, Richard Buchanan, Greg Armstrong, and the remaining members of the Nursebot team.

This project is supported through a Seed Funds grant from CMU's Provost's office, which is gratefully acknowledged. This research is also sponsored by the National Science Foundation (and CAREER grant number IIS-9876136 and regular grant number IIS-9877033), and by DARPA-ATO via TACOM (contract number DAAE07-98-C-L032) and DARPA-ISO via Rome Labs (contract number F30602-98-2-0137), which is also gratefully acknowledged. The views and conclusions contained in this document are those of the author and should not be interpreted as necessarily representing official policies or endorsements, either expressed or implied, of the United States Government or any of the sponsoring institutions. Nicholas Roy's research is supported by Le Fonds pour la Formation de Chercheurs et l'Aide à la Recherche (Fonds FCAR). Michael Montemerlo's research is supported in part by the Fannie and John Hertz Foundation.

References

Black, A.; Taylor, P.; and Caley, R. 1999. *The Festival Speech Synthesis System*. University of Edinburgh, Centre for Speech Technology Research, edition 1.4, for festival version 1.4.0.

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*.

Dario, P.; Laschi, C.; and Guglielmelli, E. 1999. Design and experiments on a personal robotic assistant. *Advanced Robotics* 13(2):153–69.

Dunbar-Jacob, J. 2000. Personal communication.

Elfes, A. 1989. *Occupancy Grids: A Probabilistic Framework for Robot Perception and Navigation*. Ph.D. Dissertation, Electrical and Computer Engineering, Carnegie Mellon University.

Endres, H.; Feiten, W.; and Lawitzky, G. 1998. Field test of a navigation system: Autonomous cleaning in supermarkets. In *Proc. of the 1998 IEEE International Conference on Robotics & Automation (ICRA 98)*.

Engelberger, G. 1999. Services. In Nof, S. Y., ed., *Handbook of Industrial Robotics*. John Wiley and Sons, 2nd edition. chapter 64., 1201–1212.

Fedor, C. 1993. *TCX. An interprocess communication system for building robotic architectures. Programmer's guide to version 10.xx.* Carnegie Mellon University, Pittsburgh, PA 15213.

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

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.

Lee, K.-F. 1989. *Automatic Speech Recognition: The Development of the SPHINX System*. Boston, MA: Kluwer Publishers.

Ravishankar, M. 1996. *Efficient Algorithms for Speech Recognition*. Ph.D. Dissertation, Carnegie Mellon.

Rowley, H.; Baluja, S.; and Kanade, T. 1998. Neural network-based face detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 20(1):23–38.

Schraft, R., and Schmierer, G. 1998. *Serviceroboter*. Springer verlag. In German.

Song, W.-K.; Lee, H.-Y.; Kim, J.-S.; Yoon, Y.-S.; and Bien, Z. 1998. Kares: intelligent rehabilitation robotic systme for the disabled and the elderly. In *Proc. 20th Inter. Conf. IEEE Engineering in Medicine and Biology Society*, volume 5, 2682–2685.

Takahashi, Y.; Kikuchi, Y.; T.Ibaraki; and Ogawa, S. 1999. Robotic food feeder. In *Proc. 38th Inter. Conf. SICE*, 979–982.

Thrun, S.; Bennewitz, M.; Burgard, W.; Cremers, A.; Dellaert, F.; Fox, D.; Hähnel, D.; Rosenberg, C.; Roy, N.; Schulte, J.; and Schulz, D. 1999. MINERVA: A second generation mobile tour-guide robot. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*.

Thrun, S.; Burgard, W.; and Fox, D. 1999. A real-time algorithm for mobile robot mapping with applications to multi-robot and 3D mapping. Submitted for publication.

US Department of Health and Human Services. 1999. Health, United states, 1999. Health and aging chartbook.