Developing Intelligent Wheelchairs for the Handicapped

Takashi Gomi and Ann Griffith

Applied AI Systems, Inc. (AAI)
340 March Road, Suite 600
Kanata, Ontario, CANADA K2K 2E4
gomi@Applied-AI.com
griffith@Applied-AI.com
Home Page: http://fox.nstn.ca/~aai

Abstract

A brief survey of research in the development of autonomy in wheelchairs is presented and AAI's R&D to build a series of intelligent autonomous wheelchairs is discussed. A standardized autonomy management system that can be installed on readily available power chairs which have been well-engineered over the years has been developed and tested. A behavior-based approach was used to establish sufficient on-board autonomy at minimal cost and material usage, while achieving high efficiency, sufficient safety, transparency in appearance, and extendability. So far, the add-on system has been installed and tried on two common power wheelchair models. Initial results are highly encouraging.

1 Introduction

In recent years, the concept of applying behavior-based intelligent robots to service tasks [Gomi, 92} has been discussed. With the accelerated rate of aging of the population being reported in many post-industrial countries, demand for more robotic assistive systems for people with physical ailments or loss of mental control is expected to increase. This is a seemingly major application area of service robots in the near future. For the past six years, we have been developing a range of autonomous mobile robots and their software using the behavior-based approach [Brooks,86] [Maes, 92]. In our experience the behavior-based approach [Brooks, 86] [Brooks, 91a][Steels, 93] [Pfeifer & Scheier, 96] [Maes, 92] allows developers to generate robot motions which are more appropriate for use in assistive technology than traditional Cartesian intelligent robotic approaches [Gomi, 96a]. In Cartesian robotics, on which most conventional approaches to intelligent robotics are based, "recognition" of the environment, followed by planning for the generation of motion sequence and calculation of kinematics and dynamics for each planned motion, occupy the center of both theoretical interest and practice. By adopting a behavior-based approach wheelchairs can be built which can operate daily in complex real-world environments with increased performance in efficiency, safety, and flexibility, and greatly reduced computational requirements. In addition, improvements in the robustness and graceful degradation characteristics are expected from this approach.

In the summer of 1995, an autonomy management system for a commercially available Canadian-made power wheelchair was successfully designed and implemented by our development team. The system looks after both longitudinal (forward and backward) and angular (left and right) movements of the chair. In addition, we implemented on-board capability to carry out "recognition" of the environment followed by limited vocal interactions with the user. The results were exhibited in August 1995 at the Intelligent Wheelchair Event organized by David Miller at the International Joint Conference on Artificial Intelligence (IJCAI’95) held in Montreal. Despite a very short development period (33 days), the chair performed remarkably well at the exhibition.

Encouraged by the initial success, we developed a three year plan to build a highly autonomous power wheelchair for use by people with various types and degrees of handicap. The intelligent wheelchair project, now called the TAO Project, intends to establish a methodology to design, implement, and test an effective add-on autonomy management system for use in conjunction with most common commercially available power wheelchairs. In order to demonstrate the principle, the project will build, during its life, an autonomy management system for several well-established electric wheelchair models currently available on the market throughout North America and Japan.

In late 1995, a sister R&D company was established in Japan exclusively for the development of intelligent robotic technologies for the disabled and the aged. With the initiative of this new R&D group, the development of TAO-2 autonomous wheelchair using a commercially available Japanese wheelchair began in the spring of 1996.

Based on our experience, methods used and some issues related to the application of the behavior-based approach to realize an intelligent wheelchair and possibly other assistive technologies are discussed. A brief survey is also presented of other groups who are working in this area.

2 Brief survey of the field

Below is a description of research on intelligent wheelchairs that has been conducted and still ongoing at some institutions. The survey is not intended to be complete but to provide an idea of the different approaches used.

2.1 IBM T.J. Watson Research Center

Some of the earliest work in the development of intelligent wheelchairs was a system implemented by Connell and Viola, [Connell & Viola, 90] in which a chair is mounted on top of a robot to make it mobile. Mr. Ed, as the chair was called, could be controlled by the user using a joystick mounted on the arm of the chair and connected to the robot. The user could also delegate control to the system itself to perform certain functions such as avoid obstacles or follow other moving objects. In addition to the joystick, input to the robot comes from bumper switches at the front and rear of the robot, eight infrared proximity sensors for local navigation and two sonar sensors at the front of the robot for following objects. Control is passed from the user to the robot through a series of toggle switches.

A set of layered behaviors were used to control the chair’s movement. These were broken into competencies with each small set of rules becoming a toolbox to achieve a particular goal. These groups could be enabled or disabled by means of switches controlled by the operator. It worked as a partnership in which the machine took care of the routine work and the user decided what needed to be done.

2.2 KISS Institute for Practical Robotics

The KISS Institute for Practical Robotics (KIPR), located in Virginia is a non-profit educational corporation performing R&D on the integration of robotics in assistive technology, space robotics and autonomous underwater vehicles as well as education in robotics and related fields.

David Miller and Marc Slack at KISS Institute have developed TinMan I and II. In TinMan II shown in Figure 1, a supplementary wheelchair controller is installed between the joystick and the standard wheelchair motor controller. Along with sensors installed on the chair, the chair avoids obstacles and goes through openings with minimum input from the user. It has been tested with two power wheelchairs, Dynamics and Penny & Giles.

Figure 1 TinMan II from KISS Institute

2.3 CALL Centre, University of Edinburgh

CALL Centre at the University of Edinburgh has developed the CALL Centre Smart Wheelchair. It was originally developed as a motivating educational and therapeutic resource for severely disabled children. The chairs were designed to assist in the assessment and development of physical, cognitive, social and communicative skills. Thirteen chairs have been built and evaluated in three local school, one in a residential hospital and three others in pre-vocational establishments.

The chairs are adapted, computer-controlled power wheelchairs which can be driven by a number of methods such as switches, joysticks, laptop computers, and voice-output. The mechanical, electronic and software design are modular to simplify the addition of new functions, reduce the cost of individualized systems and create a modeless system. Since there are no modes and behaviors are combined transparent to the user, an explicit subsystem called the *Observer* was set up to report to the user what the system is doing. The *Observer* responds and reports its perceptions to the user via a speech synthesizer or input device.

The software runs on multiple 80C552 processors communicating via an I2C serial link monitoring the sensors and user commands. Objects or groups of objects form modules which encapsulate specific functional tasks. It is multitasking with each object defined as a separate task. The architecture of behaviors each performing a specific functional task is similar to Brooks' Subsumption Architecture.

2.4 University of Michigan

Simon Levine, Director of Physical Rehabilitation at the University of Michigan Hospital began development of NavChair in 1991 with a grant for a three year project from the Veteran's Administration [Bell et al, 94]. The Vector Field Histogram (VFH) method was previously developed for avoiding obstacles in autonomous robots and was ported to the wheelchair. However, this method was designed for fully autonomous robots and it was soon determined that there were sufficient differences in the power base between robots and wheelchairs and in the requirements of human-machine systems that significant modifications were required. This resulted in a new method, called Minimum VFH (MVFH) which gives greater and more variable control to the user in manipulating the power wheelchair.

The NavChair (shown in Figure 2) has a control system designed to avoid obstacles, follow walls, and travel safely in cluttered environments. It is equipped with twelve ultrasonic sensors and an on-board computer. This team uses a shared-control system in which the user plans the route, does some navigation and indicates direction and speed of travel. The system does automatic wall following and overrides unsafe maneuvers with autonomous obstacle avoidance. Since it is desirable that the system change the user's commands as little as possible, the system and user must cooperatively adapt to environmental or function conditions. A new method called " Stimulus Response Modelling" has been developed in which the system qualitatively monitors changes in the user's behavior and adapts in realtime. It is designed so that the adaptation is smooth and the change in modes intuitive to the user. By adjusting the degree of autonomy of obstacle avoidance the control modes of NavChair can be changed giving the user more or less control depending on the situation.

Figure 2 NavChair, University of Michigan

2.5 Nagasaki University and Ube Technical College

Existing ceiling lights in an indoor environment are used as landmarks for self-localization of a motorized wheelchair by [Wang et al, 97]. The chair is therefore restricted to use within one building, the layout of which is known in advance. An azimuth sensor is used to give the angle between a fixed point and a particular object and a vision sensor detects the ceiling lights. The ceiling lights are used as the landmarks but if the lights are missed then the azimuth sensor and the rotating angle of both wheels provide the information necessary to continue the navigation.

A laser range finder is used to detect obstacles in the chair's path. Two CCD cameras are used, one is used to detect the ceiling light landmarks and the other is used in conjunction with the laser range finder to detect objects. A slit-ray is emitted from the laser emitter and this is detected by the CCD camera. The image signal is processed by a logic circuit constructed with an FPGA which informs the controller if passage is clear or where obstacles exist. In twenty test runs in a room with ten ceiling lights the maximum position error was 0.35 meters and the maximum orientation error was 17 degrees.

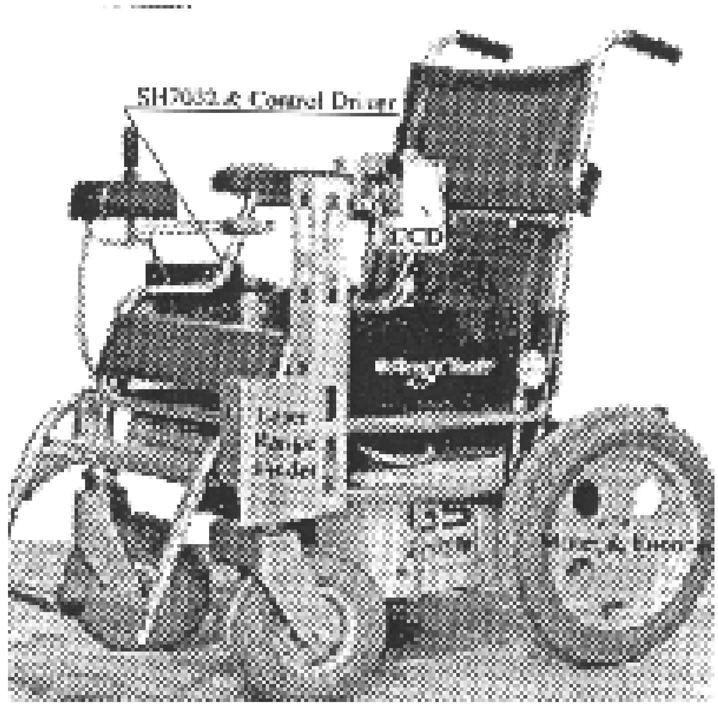

Figure 3 Chair used by Nagasaki University

2.6 TIDE Programme

Technology initiative for disabled and elderly people (TIDE) programme of the European Union began in 1991 as a pilot action with 21 development projects and a budget of ECU18 million. The SENARIO project (SENsor Aided intelligent wheelchair navigatIOn), one of the initial projects within TIDE, includes 6 member companies from Greece, Germany, the UK, and France to introduce intelligence to the navigation system of powered wheelchairs.

The system consists of five subsystems: risk avoidance, sensoring, positioning, control panel, and power control. The risk avoidance subsystem includes the central intelligence and inputs information from the sensoring and positioning subsystems. The sensoring subsystem includes ultrasonic, odometer, and inclinometer sensors. The positioning subsystem identifies the initial position of the chair by means of a laser range finder and allows the chair to be used in known environments. The control panel subsystem accepts user's instructions and the power control subsystem converts the system's instructions into vehicle movements.

The system has two modes of operation, the Teach mode and Run mode. In the Teach mode the user selects the desired paths from a topological diagram. In the Run mode (on a predefined path) the user selects a path and the system will follow it based on stored information obtained during the Teach mode. On a free route, the system takes instructions from the user and navigates semi-autonomously while monitoring safety and taking action or warning the user of the level of risk.

2.7 Wellesley College, MIT

Wheelesley is the name given to the chair used for experimental development by Holly Yanco, first at Wellesley College and now at MIT [Yanco et al, 95]. This chair has a Subsumption Architecture-like layered approach to its performance. By means of a graphical interface the user of the chair points to the direction in which the chair should head. The chair then goes in that direction while performing other tasks such as obstacle avoidance. The interface also allows the user to tell the chair when specific tasks such as going up a ramp are required and to have a record of a particular environment and important features of that environment.

The chair is designed in such a way that it can turn in place. It has 12 proximity sensors, 6 ultrasonic range sensors, 2 shaft encoders and a front bumper with sensors. A 68332 computer is onboard and the interface runs on a Macintosh Powerbook. Work is underway to incorporate information from the angle of the eyes of the user to control the computer as a replacement for the mouse.

Figure 4 Wheelesley Robot

2.8 Northeastern University

The long-term goal of Crisman and Cleary [Crisman & Cleary,96] is to develop a robot which can go to a destination, retrieve an object and return it to the operator. A teleoperated and autonomous approach each has its strength and weaknesses. Therefore, a shared control approach is suggested to divide the task between the user and the robot, taking advantage of the strengths of each. The user performs high-level functions such as object recognition and route planning while the robot performs safety and motion controls. Since the user points the objects out explicitly in a video image, the robot has been named "Deictic". The robot, after receiving instructions how to move relative to the object, performs the local motion and waits for further instruction. This means there is continuous interaction between the user and the robot with the user giving instructions to the robot every minute or so.

Commands are given to the robot by means of a button interface in which a verb description describes the desired motion of the robot and a noun describes the object relative to which the motion should be performed. The robot is able to navigate in almost any situation using its vision system to identify corners, edges, and polygonal patches.

The initial work was done in simulation followed by an implementation on an Invacare Arrow wheelchair. Motion controller cards, optical encoders, and a vision system were added to the wheelchair. New directional ultrasonic transducers were developed to detect obstacles at a wide angle in one direction and at a narrow angle in the opposite direction. This gave the robot the ability to detect objects not at standard height. A bumper with piezo-electric film embedded was installed to detect when the chair did bump an obstacle. A Puma 200 was used for the reaching experiments.

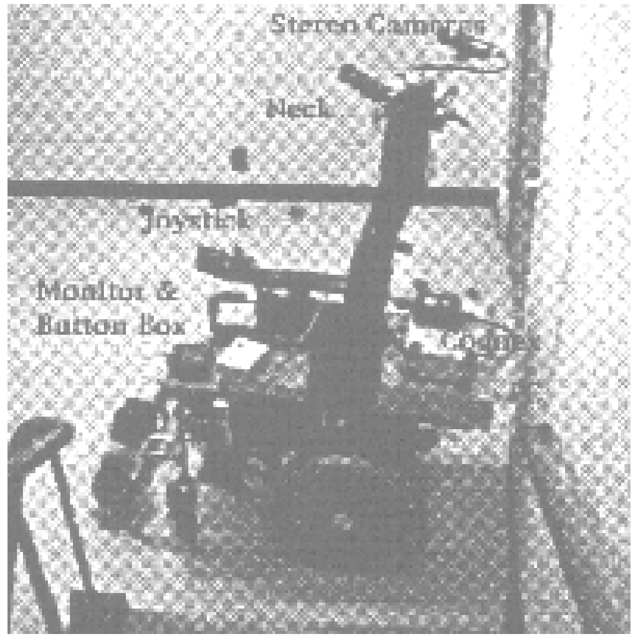

Figure 5 Deictic robot from University of Michigan

3 Desirable characteristics of robots for the handicapped

3.1 Background

Since around 1992, AAI began a number of exchanges with people with various handicaps and the individuals who assist them. This was preceded by a few years of on-going interactions with the handicapped community through marketing, installing, servicing, and training individuals on a speech-to-text voice interface system for computers. This device proved to be effective for people with several types of handicap, particularly for individuals who had lost arm/hand usage. Since late 1995, voluntary work has been attempted by members of AAI at two institutions for the mobility handicapped in Japan: a senior citizen’s hospice for severe physical/mental problems, and an institution for people with severe physical handicaps. A considerable amount of time practising physical assistive work has been carried out by members of the R&D team, including the designer involved in the conceptual design of the robots, engineers and a technician responsible for the construction of the robots, and the project manager and administrators of the robotics projects. In early 1995, an individual with a severe physical disability (a quadriplegic) joined AAI as a regular data entry/bookkeeping clerk and as a future tester of autonomous wheelchairs.

Based on these exposures, as well as earlier volunteer work, a preferable approach to robotics for service tasks [Gomi, 96b] and a tentative list of desirable characteristics for future robots built for the purpose of interacting directly with severely handicapped or fully disabled individuals has been compiled. Some of the desirable characteristics are discussed below.

3.2 Softness and flexibility

Establishment of rapport between the handicapped person and the caregiver is essential for the care to be successful. So much so, there will be a great deal of anxiety in those treated by future robotized arms, support boards, and wheels. The need for softness realized between the physical interface of the end effectors of such a robot and the human body surface or limbs does not stop at simple padding of otherwise solid effector surfaces, or use of softer materials, or passive or active compliance of effectors. The softness must also be architectural in that the entire physical support structure must be able to alter, reconfigure, and even completely restructure moment to moment reactions and responses to accommodate, whenever necessary, changes in not only the physical but also the perceived psychological situation of the user.

The flexibility of the system as a whole, as well as that of the end effectors, must essentially come from this "structural softness". The flexibility must be founded on the openness of the design of motions the system can generate so that it does not rely on fixed modes of operation or rigid scenarios defined *a priori*. In most circumstances humans in general behave without a prepared set of motion patterns, and since we are dealing with such an existence, a man-made system itself must not act with a fixed set of motions which are algorithmically describable. This places the appropriateness of most existing system control methods in doubt as a tool to seriously deal with many types of physically handicapped people.

Learning has often been hailed as a scheme with which a system can be made more adaptable. We would also have to question this relishable notion as a candidate that would sufficiently increase adaptability of systems such as service robots dealing directly with humans. Learning schemes, particularly those so far studied to the greatest extent and depth in the symbolic AI community, have failed to make significant contributions to robotic systems operating in highly dynamic application areas. In general, learning research has focussed on methods to improve the chosen performance index of systems but variables involved in the scheme are most often not grounded through sensors or actuators.

3.3 Fail safe and robust

A robot arm holding a fragile human body must not drop the person when a bug is hit for the first time. The concept of fail safe implies readiness of a system against possible failure. In traditional system engineering disciplines, such as Fault Tolerant Computer Systems (FTCS) research and practice, this typically translates into the preparation of additional capabilities in the form of a standby in computer hardware and software. The concepts of hot-standby and cold-standby are commonly employed in system design. Since it is impossible to prepare for every possible failure, the provision of readiness should exist, however, more in the form of capabilities spread across the system in atomic form and meshed fine grain with the competence structure which also functions in the normal execution of tasks. This is analogous to the way readiness to failure is implemented in life forms found in nature. If a small animal or an insect temporarily loses the use of a limb, it tries to adjust to the situation by immediately enlisting the use of other limbs or even other portions of the body. The additional capability readied in this form would be quickly organized and mobilized the moment a fault is detected.

3.4 Graceful degradation

A cousin to the concept of fail safe, graceful degradation is more important in systems that physically interface with humans than in systems that deal with materials and artifacts. A control system designed as a monolith or components with relatively larger granularity would have less chance of realizing the concept fully. When one loses a limb, the resulting transition is not smooth, causing great suffering to the individual. However, every day we lose a large number of brain cells that we know won't reproduce, but we do not deteriorate or lose capabilities as drastic as loosing a limb. Systems composed of finer grain active units seem to offer more desirable results.

3.5 Evolvability

Another reason for the failure of learning in symbolic AI would be the relatively short time the methods have typically tried to achieve the "result". In fact, we probably do not know what desirable results are as much as we think we do. Both shortcomings, this and the lack of grounding, are due mostly to the very nature of being symbolic rather than pragmatic.

In evolution, changes occur along a much longer time scale. In *situated* and *embodied* systems, such as life forms in nature and well-built autonomous robots, a search through a very high dimensional space of the real world for adaptation demands "experiments" on a vast number of combinations of dimensional parameters, if such dimensionalization or parameterization makes sense at all. Evolutionary Robotics (ER) is an emerging field of science and technology [Harvey, 92], where physical or virtual robots' autonomy structures are evolved to achieve collective trans-generational learning. ER seems to be a scheme that could well be applied to robots operating to tend and care for humans because of the open nature of human autonomy and ER's basic principle that can provide long term learning. Here, the concept of learning should probably be replaced by a more comprehensive concept of evolution, which implies perpetual adaptation of an autonomous system to a constantly changing operational environment rather than optimization of one or more performance indices of such a system.

3.6 The development plan

The development of autonomous wheelchairs at AAI is carried out in the following four phases. Some of the phases overlap in their execution.

(1) The basic safety phase,
(2) The mobility phase,
(3) The human interface phase, and
(4) The exploration phase.

Currently, we are in the second phase of the project which began on April 1, 1996. Prior to the start of the project on July 20, 1995, a study was conducted to identify various requirements by potential users of the autonomous wheelchair both in Canada and Japan through interactions with people with various types of handicap. Causes of the handicaps we came across included gradual mobility loss by aging, recent sudden loss of body control due to brain damage, and prolonged motion limitations and bodily contortion due to stroke suffered at a young age. The project continues to enjoy cooperation from institutions for the handicapped and individuals with disabilities. The TAO project is scheduled to end in the summer of 1998. For a description of the development plan, please refer to [Gomi & Ide, 96].

4 Implementation of the first prototype, TAO-1

A regular battery powered wheelchair (a motorized chair) produced and marketed in Canada (FORTRESS Model 760V) was used as the base of the first implementation of the concept. A set of sensors, a computerized autonomy management unit, and necessary harnesses were built and added to TAO-1 (Figure 6) through the summer of 1995.

Figure 6 Autonomous wheelchair TAO-1. Cover is removed to show autonomy unit.

4.1 Planned functions of the chair

The selection of functions to be implemented on TAO-1 was somewhat influenced by the rules set out for the IJCAI'95 robotics contest. However, later demonstrations of our prototype and observations made at an institution for the aged confirmed that the guideline was in fact appropriate. Of the following functions which we now follow, only the first two were attempted at our IJCAI'95 entry. However, all five of them are currently pursued.

(a) Basic collision avoidance

This is achieved by behaviors which monitor and respond to inputs from on-board CCD cameras or those which respond to active infrared (IR) sensors. When the chair encounters an obstacle, it first reduces its speed, and then depending on the situation it faces, stops or turns away from the obstacle to avoid hitting it. The obstacle can be inanimate (e.g., a column in a hallway, a light pole on the sidewalk, a desk, a standing human) or animate (a passerby, a suddenly opened door in its path, an approaching wheelchair). Encountering a moving obstacle, the chair first tries to steer around it. If it cannot, it stops and backs off if the speed of the advancing obstacle is slow enough (e.g., 20 centimeters per second). Otherwise, it stays put until the obstacle passes away. Thus, if the chair encounters another wheelchair, both chairs can pass each other smoothly as long as there is enough space in the passage for two chairs. A fast paced human usually does not affect the chair's progress and at most causes the chair to temporarily slow down or steer away.

(b) Passage through a narrow corridor

When surrounded by walls on each side of the path, as in a hallway, the chair travels autonomously from one end to the other parallel to the walls.

(c) Entry through a narrow doorway

The chair automatically reduces its speed and cautiously passes through a narrow doorway which may leave only a few centimeters of space on each side of the chair. Some types of ailment such as Parkinson's disease or polio often deprive a human of the ability to adjust the joystick of a power wheelchair through such a tight passage.

(d) Maneuver in a tight corner

Similarly, when the chair is surrounded by obstacles (e.g., walls, doors, humans), it is often difficult to handle the situation manually. The autonomous chair should try to find a break in the surroundings and escape the confinement by itself unless instructed otherwise by the user.

(e) Landmark-based navigation

Two CCD color cameras on-board the chair are used for functions explained in (a), (b), and (c) above. They constantly detect the depth and size of free space ahead of the chair. The cameras are also used to identify landmarks in the environment so that the chair can travel from its present location to a given destination by tracing them. An on-board topological map is used to describe the system of landmarks.

4.2 Hardware structure

As a standard powered wheelchair, model 760V has two differentially driven wheels and two free front casters. Although they are designed to rotate freely around their vertical and horizontal axis, these casters typically give fluctuations in delicate maneuvers due to mechanical hysteresis that exists in them because of design constraints (the rotating vertical shaft of the support structure of the caster cannot be at the horizontal center of the caster). This sometimes causes the chair to wiggle particularly when its orientation needs to be adjusted finely. Such fine adjustments are necessary typically when a wheelchair tries to enter a narrow opening such as a doorway.

The entire mechanical and electrical structure, the electronics, and the control circuitry of the original power wheelchair were used without modification. The prototype autonomy management system still allows the chair to operate as a standard manually controlled electric wheelchair using the joystick. The joystick can be used anytime to seamlessly override the control whenever the user wishes even in autonomy mode.

Physical additions to the chair were also kept to a minimum. AI components added to the chair were made visually as transparent as possible. Two processor boxes, one for vision-based behavior generation and the other for non-vision behavior generation are tacked neatly under the chair's seat, hidden completely by the wheelchair's original plastic cover. Sensors are hidden under the footrests, inside the battery case, and on other supporting structures. Only the two CCD cameras are a little more visible: they are attached to the front end of the two armrests for a good line of sight. A small keypad and miniature television set are installed temporarily over the left armrest to enter instructions and for monitoring.

The non-vision behavior generator is based on a Motorola 68332 32-bit micro controller. A multi-tasking, real-time operating system was developed and installed as the software framework. This combination gave the system the capability to receive real-time signals from a large number of sensors and to send drive outputs to the two motors which govern the wheels. The chair currently has several bump sensors and 12 active infrared (IR) sensors which detect obstacles in close vicinity (less than 1 meter) of the chair. Signals from the cameras are processed by a vision-based behavior generation unit based on a DSP board developed by a group at MIT. Vision processing is discussed in Section 6.6 below.

4.3 Software structure

The over-all behavior structure of TAO-1 is shown in Figure 7. Smaller behaviors are lumped up to save space on the diagram. Software for the vision system is also built according to behavior-based principles. The major difference between this and conventional image processing is that it consists of behaviors, each of which generates actual behavior output to the motors. It can presently detect depth and size of free space, vanishing point, indoor landmarks, and simple motions up to 10 meters ahead in its path. Indoor landmarks are a segment of ordinary office scenery that naturally comes in view of the cameras. No special markings are placed in the environment for navigation.

There are also a large number of behaviors invoked by IRs and bumpers which collectively generate finer interactions with the environment. Vision-based and non-vision behaviors jointly allow the chair to proceed cautiously but efficiently through complex office spaces. Note that there is no main program to coordinate behaviors.

Currently, the autonomy program occupies about 35 KBytes for all of the vision related processing and 32 KBytes for other behavior generation and miscellaneous computation. Of the 35 KBytes for vision related processing, only about 10 KBytes are directly related to behavior generation. The rest are involved in various forms of signal preprocessing: generation of depth map, calculation of the size of free space, estimation of the vanishing point, and detection of specific obstacles in the immediate front of the chair.

Of the remaining 25 KBytes, approximately 20 KBytes are used in the neural network system for detecting landmarks and referencing a topological map. The current implementation of the landmark system consumes only 256 Bytes per landmark, although this figure may change in the future as more sophisticated landmark description might become necessary. The current system has space for up to 64 landmarks but this can also be adjusted in future versions.

Of the 32 KBytes of non-vision processing (i.e., processing of inputs from IR's , bump sensors, voice I/O, etc.), again no more than several KBytes are spent for generating behaviors. Altogether, there are some 150 behaviors in the current version of TAO-1. A considerable amount of code has been written to deal with trivial periphery, such as keypad interface, voice I/O, and LCD display. The comparable inefficiency of coding is because these non-behavioral processing had to be described in more conventional algorithms.

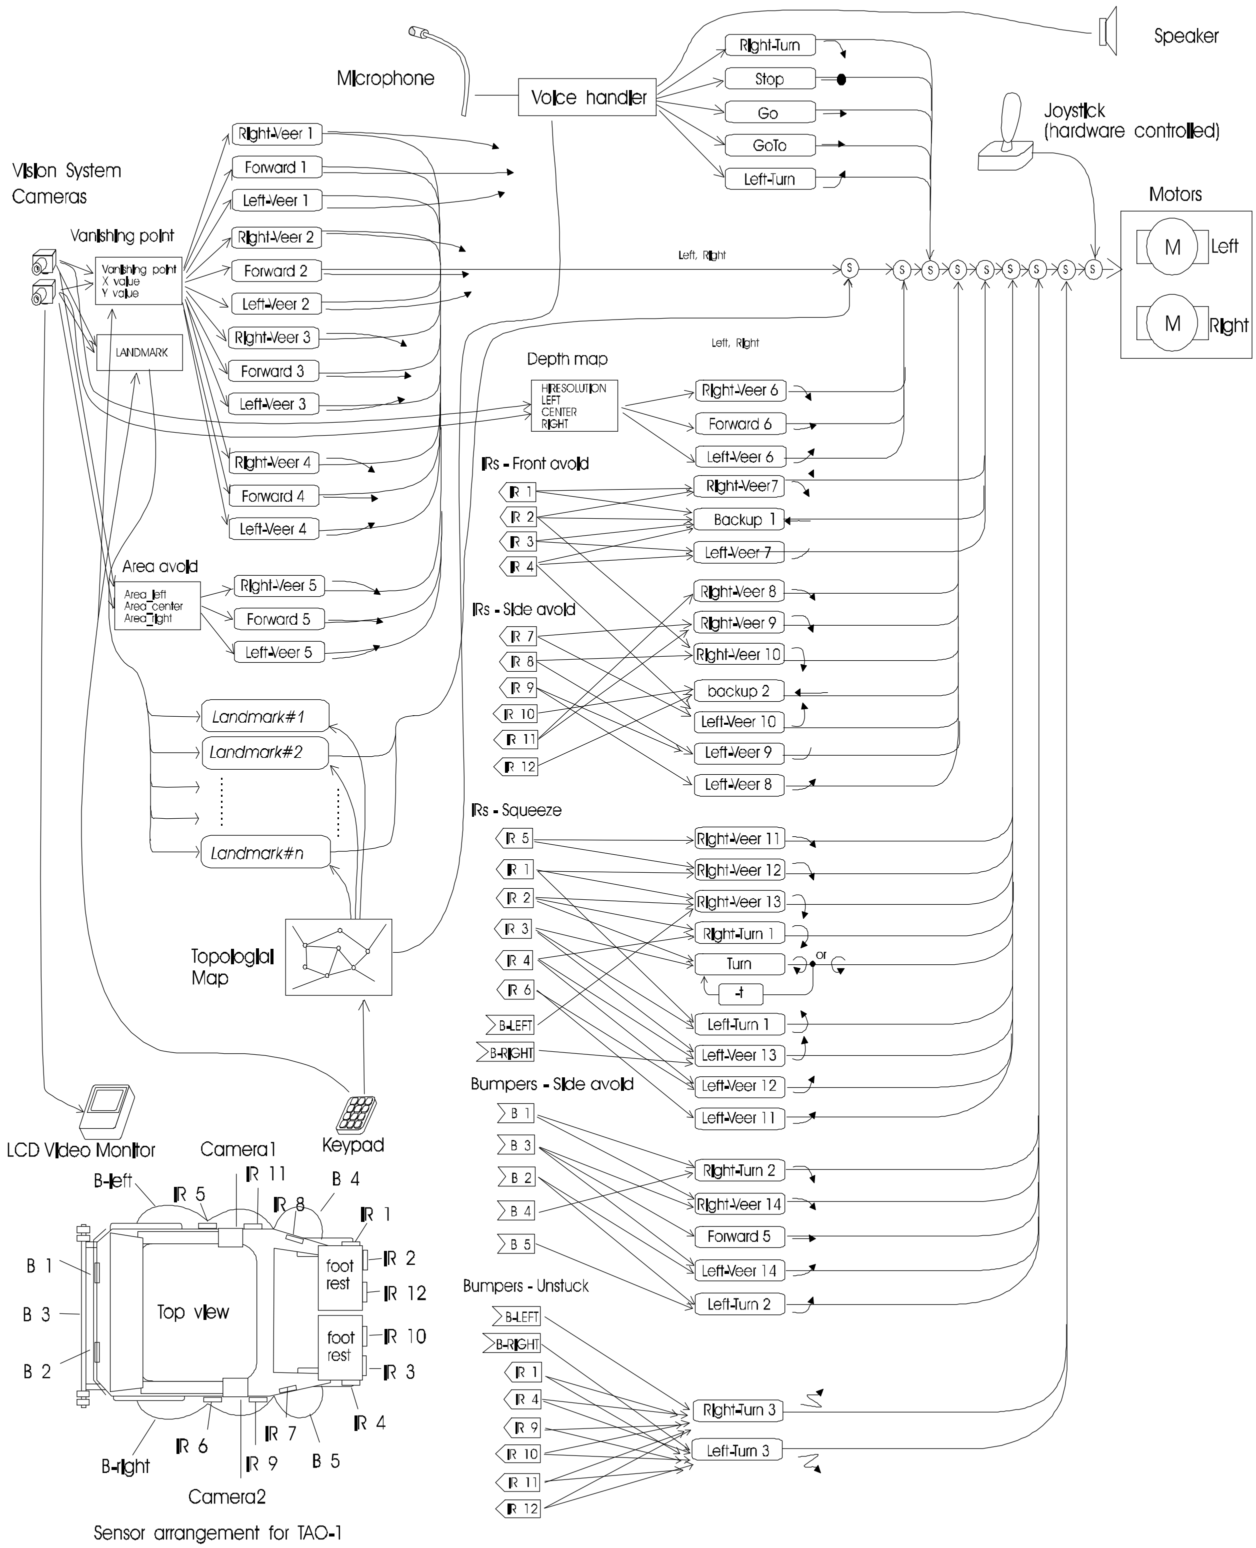

Figure 7 TAO-1 behavior structure (Not all behaviors are shown)

5 The second prototype, TAO-2

Encouraged by the success of TAO-1, in late 1995 a sister company of AAI (AAI Japan, Inc.) was established in northern Japan. AAI Japan is dedicated to the development of advanced intelligent robotics to aid people with various handicaps. In May 1996, AAI Japan purchased a new power wheelchair (Suzuki MC-13P), which is a model widely used in Japan. MC-13P has a form of power steering in which the two front casters alter their orientation in synchrony with the drive wheels when a turn is indicated by the joystick. The servo controller also halts the inside turn wheel of the two drive wheels while the chair is making a tight turn. This is a significant departure from the way the FORTRESS model makes a turn. The latter simply turns the two differentially driven main wheels in opposite directions, allowing the chair to turn in place. The intent of providing a power steering feature on the Suzuki chair is obviously for ease of use, and the user is freed from the wiggly caster problem described above. However, this prevented the chair from making turns in a tight turn circle. The feature was felt undesirable for an autonomous chair.

Figure 8 TAO-2 autonomous wheelchair

Immediately following the purchase of the Suzuki chair, the development team began building an autonomy management system for TAO-2; a new prototype autonomous chair based on MC-13P. The over-all computer hardware and software structures as well as sensors are almost identical to those for TAO-1, except for a few changes listed below to accommodate the above mentioned and other minor differences in characteristics.

(1) The behaviors responsible for turning TAO-2 needed their parameters adjusted.
(2) The locations of touch sensors made up of thin piano wires needed to be moved forward in order to compensate for a larger turn circle.
(3) The back bumper was not activated since it was hardly used. The difference in turning characteristics reduced the chance of the Suzuki chair performing frequent switch backs.
(4) Two prominent side bumpers were added to protect bystanders when the chair makes a turn in their direction. This was necessitated by the lack of structure on which to mount sensors.

TAO-2 is shown in Figure 8. It was fitted with the autonomy management system at AAI in Canada in the span of one week. After two days of testing, it was shipped back to Japan in time for a public demonstration in the town of Kosaka, Akita Prefecture.

6 Evaluation of the Prototypes

6.1 Demonstrations

When TAO-1 was demonstrated at IJCAI'95 in Montreal on the 22nd of August, it was the 33rd day of the development of the first prototype. Everything from the motherboard, vision system, sensor arrangements and their harnessing, operating system (based on an earlier prototype), a large number of behaviors (some 60 by that time) were all developed and tested in that period. The chair could perform functions (a) and (b) in Section 4.1 well and functions (c) and (d) moderately well, although they were not initially targeted. Function (e) was not yet implemented. In all, it performed as well as other chairs at the exhibition most of which took much longer time to develop. All five functions are now implemented on TAO-1 and are undergoing continuous improvement.

TAO-2 was demonstrated on June 4, 1996 at a gymnasium of a local school in Kosaka, Japan. The chair ran smoothly throughout the 1 hour demonstration persistently avoiding by-bystanders, other obstacles and the walls. Unsolicited, a severely handicapped spectator who could not even reach the joystick volunteered to test ride the chair. The chair performed to her satisfaction and excitement as it went through the gymnasium among a large number of spectators.

The success of the two prototypes suggests that our intention to build a standardized add-on autonomy unit is a valid one. The concept has at least been proven in two power wheelchair types which come from drastically different backgrounds. The divergence in design philosophy and practical variances in implementation, some fairly significant, of a base power wheelchair can be absorbed by relatively minor hardware and software alterations made on the standardized add-on unit. TAO-2 also showed that the installation, testing, and adjustment of a separately built autonomy unit can be made in a very short period of time. In both TAO-1 and TAO-2, no cooperation from the manufactures was sought. In each case, characteristics of the joystick were studied and a seamless interface was designed around it.

6.2 TAO-2 experiments

After successfully testing the basic navigation functions of TAO-2 at our laboratory in Canada, it was transported to AAI Japan’s facility in Akita Prefecture, Japan in May, 1996 for additional tests and adjustments. Two types of experiments were conducted with TAO-2: indoor experiments and running of the autonomous chair outdoors in snow. The indoor experiments included unassisted navigation of the chair in a circular corridor and the gymnasium of a local primary school, and in corridors of an institution for physically handicapped adults. At the school, the chair navigated smoothly both in the circular corridor and the gymnasium except when it hit a glass door separating the corridor and one of classrooms next to the corridor. The incident was due to the fact that the chair bases its collision avoidance on vision (incapable when faced with a planer glass surface under rare lighting conditions) and active infrared (IR) sensors (IR emission is transparent through most glass surfaces). This, however, does not mean the present sensors are inferior. On the contrary, combined they are vastly more efficient and capable than other sensors such as laser range finders and ultrasonic sensors. Nevertheless, the addition of local ultrasonic sensors is being considered to cover this imperfection.

In the gymnasium which was populated by several dozen spectators, some of whom were surrounding the chair, TAO-2 constantly found a break in the crowd and escaped from the human wall without touching anyone. A female spectator with severe bodily contortion volunteered to try the chair. Her condition was such that she was not even capable of extending her arm to reach the joystick. As in TAO-1, the control structure of the original power wheelchair (Suzuki MC-13R model) was left intact when the autonomy management system was added. The intelligent chair is designed to allow the user to take over the entire control system by touching the joystick. It then simply acts as a standard motorized chair. Despite the total absence of input from the user, the chair navigated smoothly, always successfully avoiding walls and spectators. When completely surrounded by the spectators, it stopped until a break which was approximately 50% wider than the width of the chair developed roughly in front of it. It then moved out of the circle through the opening. The ability to locate a break anywhere in a circle regardless of its orientation when surrounded by people has been implemented and tested in other behavior-based robots.

When tested at a local institution for the severely physically handicapped, the chair managed to travel along corridors in most cases. Interest in an autonomous wheelchair that can take individuals to a desired destination was strong, and the experiment had to be conducted amid many spectators who were themselves in a chair. TAO-2 encountered some difficulties when surrounded by other wheelchairs in close proximity. This difficulty includes at its core a common problem for both TAO chairs: the autonomy management system still requires better processes to detect thin pipes or tubes in the environment. Such processes will likely depend on inputs from the vision system as it provides the widest communication path between the chair and the environment and is amenable to the addition of new processes to deal with specific problems such as detection of vertical and horizontal thin pipes in the path of the autonomous chair. Landmark navigation was not attempted in these experiments due to the shortage of time and manpower necessary to prepare an on-board topological map. In all, TAO-2 at this stage appeared to have basic navigational capacity in populated indoor space.

In February 1997, TAO-2 was tested outdoors on the snow covered pavement and sidewalks of Kosaka, Japan. No particular modifications were made to the basic functioning of the indoor version of the chair except for minor adjustments to the vision system, the active IR sensors and the software. The outdoor temperature was around -10 degrees Celsius when the chair was tested. First, TAO-2's ability to interpret signals obtained through the vision system and other sensors (IR’s and bumpers) when navigating through the mostly white surrounding snow-scape was checked. The chair successfully navigated through a narrow corridor sided by walls of snow. Most of the time the chair depended on both the vision system and IR sensors to position itself roughly in the middle of the narrow (changing from 1.2 to 1.5 meters) corridor. The surface of the floor of the corridor was mostly covered by snow with some foot prints. The height of the snow walls on both sides of the corridor was about one meter. The sunlight which was shining through a thin layer of clouds at an angle from behind the chair caused one of the walls to appear quite dark and the other slightly brighter, while the floor was yet another tone. Such a contrast was good enough for the vision system to distinguish the geometry and guide TAO-2 roughly in the middle of the snow corridor. Whenever the chair’s course noticeable deviated from the center of the corridor, mostly due to friction and slippage caused by the uneven surface of the snow covered floor, the IRs on either side would detect the deviation and associated processes were invoked to cancel the deviation.

When TAO-2 travelled through the entire length of the corridor and reached the open pavement which was mostly covered by snow with some tire marks and sporadic black exposed surfaces of asphalt, it navigated among these ground marks just as humans would try to make sense of the orientation of the hidden roadway underneath the largely snow-covered pavement.

The TAO-2 chair was also tested on a sidewalk under similar climatic condition (snow on the ground, cloudy day with sufficient light, -10 degrees Celsius). However, the surface of the sidewalk was clear of snow because of a snow removal system that warms up the underside of the surface of the sidewalk using well-water. The system very successfully maintains a snow-free strip about 90 centimeters wide in the middle of a 1.2 meter wide sidewalk up until a certain temperature and rate of snowing. This optical contrast created an ideal condition for the vision system. Because of the high contrast between the wet surface of the sidewalk made up of dark brown bricks of the sidewalk and the white snow covered edges of the sidewalk, the vision system could easily follow the track. Light standards are erected at regular intervals on the edge of the sidewalk creating a particularly narrow passage. When passing by the light standards, the chair slowed down to negotiate past them but did not have particular difficulties to clear them. In general, the performance of TAO-2 in snowy outdoors was much better than expected. It became clear that the chair can cover the basic navigational requirements through a snow-covered town where a distinctive sidewalk system with snow removal is available.

6.3 Development time

The extremely short development time required for the initial prototype for both TAO-1 and TAO-2 can largely be attributed to the behavior-based approach. To achieve the demonstrated level of mobility and flexibility would normally have required another several months to a few years in conventional AI-based mobile robotics. In behavior-based robotics, the operational characteristics of the sensors need not be as precisely uniform as in conventional mobile robotics. For example, emission strength and angular coverage of the emitter, and the sensitivity and shape of the reception cone of the receptor of on-board IR sensors need not be homogeneous across all sensors, allowing the use of inexpensive sensors and simpler testing.

All sensors, including the CCD cameras, need not be installed at precise translational and angular coordinates. They also do not need calibration. They were placed on the chair in a relatively *ad hoc* manner at first, and continually moved around for better results as the development went on. In fact, the cameras and some of the sensors are attached to the chair by velcro detachable tape, so that their location and orientation can be adjusted easily. Such loose treatment of sensors is not common in conventional robotics where the robot's motions are derived after high-precision measurements of the relationships between its extremities and the environment. The large tolerance for signal fluctuation is due also to flexibility of processing and greater adaptability inherent in Subsumption Architecture [Brooks, 86].

With the absence of "sensor fusion", sensor inputs are directly linked to motor output only with a simple signal transformation and amplification (e.g., from sensor output voltage to motor drive current). The developer only needs to adjust the appropriateness of the definition and performance of the sensor-action pair or behavior in terms of its output without a detailed and precise analysis of input signal characteristics and elaborate planning and computation of output signal generation. Readers not familiar with the theoretical basis of behavior-based AI are encouraged to read [Brooks, 91b}. These theories are fully put into practice in our development.

6.4 Software structure

During development, sensor-actuator pairs or behaviors are simply "stacked up". They are added to the system one by one without much consideration for the design of the over-all software structure. Our operating system provided an adequate framework for the incremental development process allowing for shorter development time.

Thus, software development went totally incrementally side by side with finer adjustment of the sensors. Only general functions needed to be assigned to each sensor-actuator pair type first. For example, depth map - motor pairs are excellent for dealing with obstacles that suddenly appear in the path of the chair a few meters away. But the same sensor-actuator pair type is not at all effective for the management of the situation in which the chair has actually made physical contact with an obstacle.

Sometimes, competition or contradiction occurs between two or more behaviors. Such contradicting definitions of behaviors are in most cases easily observable and corrected quickly. An example of more complex contradiction occurs when two IR collision-detection sensors placed on the left and right front sides of the chair detect an approaching doorway in quick succession. Since the doorway is normally quite narrow, the reflection of infrared signals received by these sensors is usually strong enough to cause the chair’s immediate evasive action. As both sensors react alternatingly, the chair can get into an oscillatory motion, commonly known as "Braitenberg’s oscillation" after [Braitenberg, 84]. In this specific situation, other frontally-mounted IR sensors take in "just go ahead" signals that invoke behaviors which can break the tie.

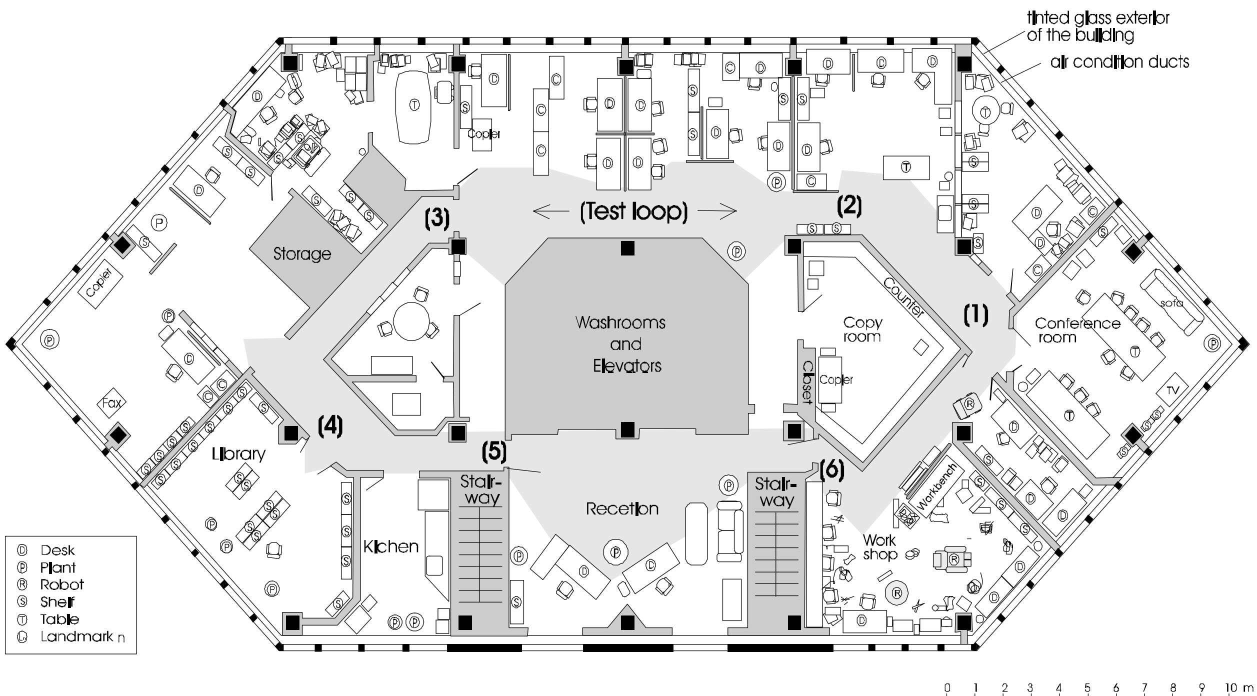

Figure 9 The office space which contains the test loop

6.5 Priority scheme

The priority arrangement is shown in the top right corner of Figure 7, where several output lines to the motors are joined by Ⓢ nodes or suppression nodes. Input from the left of the node is suppressed and replaced by one coming in vertically whenever it is active. Inputs from the joystick take the highest priority in deciding which action the chair should take. The electronically and mechanically seamless interface between the joystick controller and the autonomy management system allows the chair to run as a standard power wheelchair simply by operating the joystick. The second highest priority is given to behaviors which take in signals from left and right bumpers and some key frontal IR sensors. Behaviors are bundled up in Figure 7 with implied logical relationships among input lines to simplify the diagram. There are several groups of behaviors that mostly depend on signals from IR sensors for their invocation. These are followed by behaviors invoked by signals from the voice input system, followed by vision-driven behaviors as the lowest priority behavior groups. They are, in descending order of priority, depth map, vanishing point, and free area.

Figure 10a shows IR signals from a test run in which TAO-1 went around the test loop in our office floor shown in Figure 9 (shaded area). Note that signals from only 6 of the 12 IR sensors are plotted here. The x axis in Figures 10a through 10d shows the passage of time and its length corresponds to the time required to complete the loop from the workshop and back there counter-clockwise. Note that checkpoints (1) through (6) shown in Figure 9 are also marked on the diagrams. When there is no reflection from an obstacle, output of an IR is kept at 255. Depending on the strength of the reflected signal, a receptor may report lower values, 0 being the lowest. When the value becomes less than a threshold, the sensor would have "detected an obstacle." The threshold

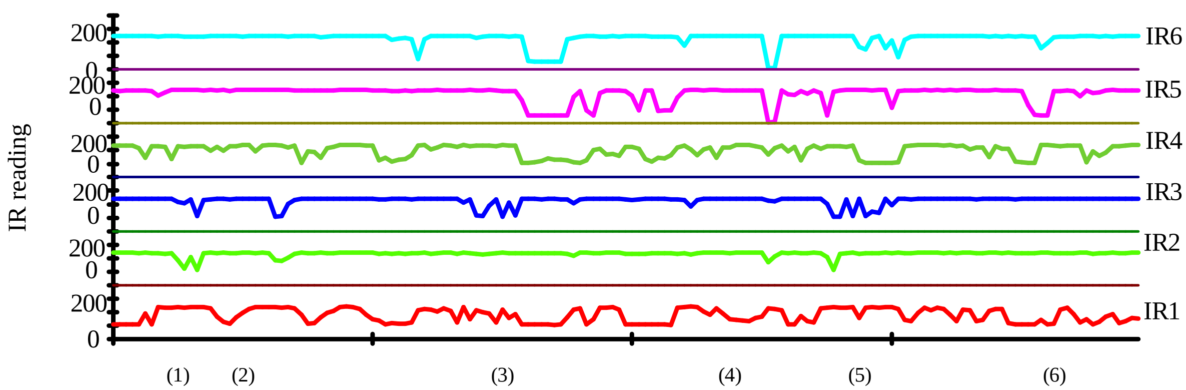

Figure 10a Output of active infrared (IR) sensors

is set as a function of the speed of the chair, and in this specific test is set at 210, 180, and 150, for when the chair is running, at fast, medium, and slow speed, respectively. In another mode of obstacle detection using an IR, changes in value are monitored for several sensor cycles. If the change is sufficiently large, detection of an obstacle is reported. The IR sensors take in values at 64Hz and several consecutive values are compared. Once invoked, a behavior corresponding to a specific IR sensor generates a predetermined reactive motion, altering the speed and orientation of the chair.

6.6 Vision processing

Inputs from 2 CCD cameras are alternatively processed through a single frame grabber into two primary vision planes of 256 x 128 pixels each at about 8 frame sets per second. Images in these primary vision buffers are averaged down to 64 x 32 pixel secondary vision plane by combining the left and right vision inputs after dividing each primary plane into left, center, and right. All vision processing described below occurs using image data in this secondary visual plane.

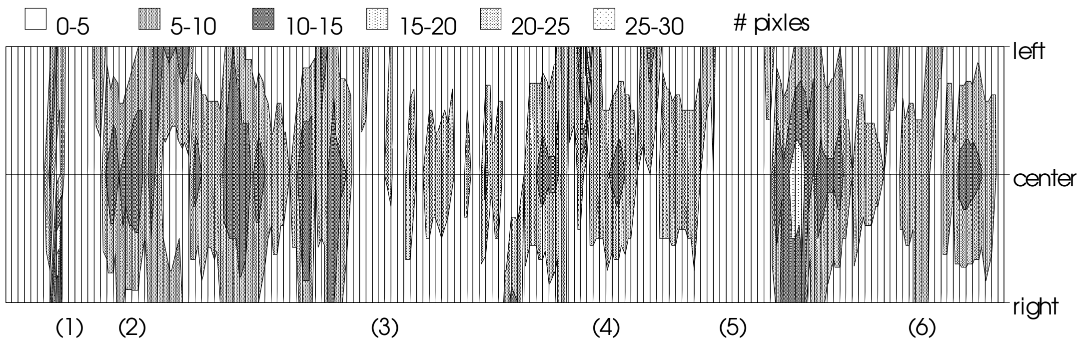

Figure 10b Depth map parameters from the vision subprocess

Figure 10b plots three depth values (left, center, and right) in terms of the number of pixels in the secondary visual plane determined according to Horswill's habitat constraint vision processing [Horswill, 92]. In the absence of active bumper and IR invoked behaviors, the parameter set directly dictates the orientation and speed of the wheels.

Output from the vanishing point detector of the vision system is shown in Figure 10c. The

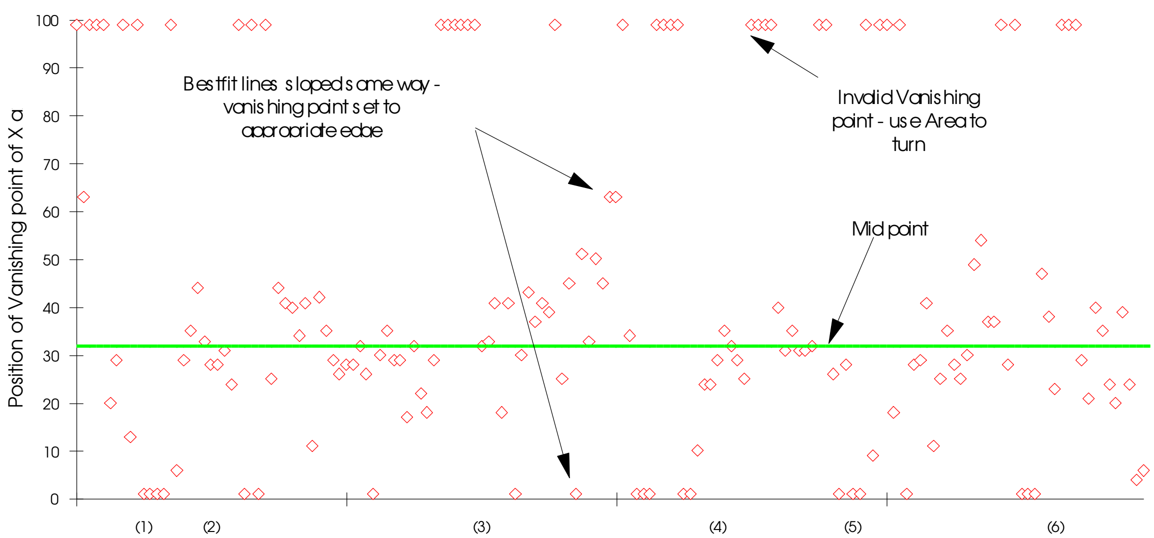

Figure 10c Output of vanishing point detector

detector attempts to find a vanishing point in the secondary visual plane and outputs its x axis value when it finds one. The value 0 corresponds to the left-most angle in the visual plane, and 63 to the right-most. When it fails to come up with a vanishing point, value 99 is output. The combined horizontal viewing angle of the left and the right cameras is approximately 100 degrees.

Figure 10d depicts output from the area detector. The number of pixels representing free space in the left, center and right visual fields are calculated by the detector. Steering and speed of the chair are determined by the size of available space as in depth map processing. The behaviors associated with area detection are invoked only when all other behaviors are not invoked.

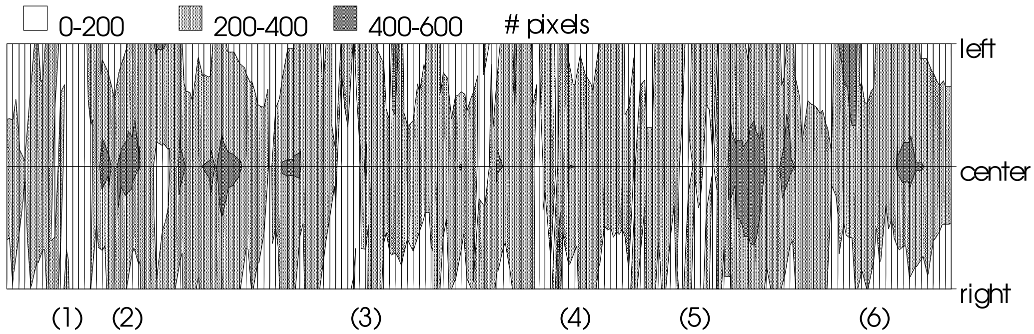

Figure 10d Output of the area detector

As the project proceeds the vision system will be enhanced to detect more objects and events such as outdoor landmarks, indoor landmarks that change in time, more complex and dynamic obstacles, and traffic signals in the path.

7 Lessons learned so far from the chair project

Although the experience is still very limited, we can state that there is a strong expectation among the population for the development of an autonomous wheelchair for assisting and eventually fully taking care of the handicapped person's domestic transportation needs. We have demonstrated that the chair can travel at reasonable speeds through a standard North American office space with its peculiar attributes such as average width of passage ways, nature and volume of human traffic, size and orientation of doorways, etc.

In April 1996, TAO-1 was brought to a local shopping mall in Ottawa to freely roam around for an hour or so. TAO-1 skilfully skirted all internal structures of the mall such as escalators, flower pots, benches, signs, and showcases, as well as afternoon shoppers. TAO-1 and its rider visited stores as if he was window shopping or just strolling the mall. Virtually all fellow shoppers failed to notice that it was not driven manually. It tended to swerve downward when a sidewalk at the shopping center was slanted. This problem could be corrected in a few ways, and infact, when we encountered the same problem with TAO-2 on a sidewalk in Japan, we successfully implemented one of the methods. This made us feel that with proper engineering to increase the chair's dependability, it can already serve as an autonomous chair for the severely handicapped in limited application areas, such as strolling or window shopping. Usability of the chairs in more constrained places such as smaller homes and limited office spaces would require further testing and revisions.

In the United States in the early 20th Century when automobiles began hitting humans on the street killing or injuring them, many cities and towns passed by-laws mandating each driver to have a "battler" running and waiving a flag (or a lantern after dark) in front of the car. This practically limited the maximum speed of automobiles to about 10 miles per hour. Of course, the practical application and enforcement of these bylaws met strong resistance from the reality, and the isue was replaced with other arguments or simply forgotten in many cases. Some of the by-laws are said to be still in effect. The episode tells a lot about human nature and what will likely happen to the fate of intelligent wheelchairs and similar "intelligent" machines that are meant to assist and help would-be human users in need. After the modest demonstration of TAO-2 in Japan, which was reported in local television news and several local newspapers, we have received inquiries for the chair's availability. Needless to say, it will be at least a few more years before even a modestly autonomous chair can be released for use by the handicapped at large and put into daily use only with affordable amount of support.

Maintenance would be another issue if we proceed, not to mention various public liability issues that, unfortunately but undoubtedly, will follow. The public liability issue is potentially a problem in introducing an autonomous or semi-autonomous wheelchair to the general public and this can become a hindrance to the effort to bring these technologies to the handicapped.

We are not at all optimistic about the efforts required to establish an infrastructure for physical and moral support that encompasses all these and other yet to be found issues. Nevertheless, we can foresee that we will be able to answer, in the near future, some of the sincere wishes that already come from people who would be most benefitted by the technology.

Getting into technical issues, the list of things yet to be done is still quite long. Landmark detection, for example, requires a lot more work. Although we have succeeded in navigating the chair to go through a series of landmarks arbitrarily chosen in the chair's present operational environment, this is still a far cry from being able to state that it can run freely in any environment traversable by a wheelchair by detecting landmarks.

Apart from these and other shortcomings, we feel the technology as it is, is already useful in real world applications by individuals with certain types of handicap. Persons with bodily contortions such as those who suffered polio in earlier life, or individuals with involuntary hand/arm movements such as patients of Parkinson's disease, now could travel through confined and narrow spaces such as corridors and doorways without assistance. Other interface mechanisms such as neck control and a voice recognizer would also make the introduction of the autonomous chair easier. Less handicapped users can use the chair as a manual power wheelchair whenever desired, while autonomy management can assist in mundane situations and emergencies.

Everybody with whom we have interfaced so far, from a passer-by at the shopping center where TAO-1 was tested, to fellow robotics researchers, several handicapped people and caregivers who heard about the project and came to see and even volunteered to try an early prototype, willing investors, and journalists all gave us positive feedback. They agree in principle that mobility should be provided as much as and as soon as possible to those who otherwise are not capable of going to places by themselves. Although the development is still far from complete, TAO-1 and 2 have so far been covered by several TV programs and a few dozen newspaper and magazine articles in Europe, Japan, USA, and Canada, indicating the keen level of interest the public has on this subject.

8 Conclusions

Two prototype autonomous wheelchairs based on commercially available motorized wheelchairs have been built using behavior-based AI. The initial prototyping went very rapidly and the size of the software is significantly smaller than control programs for similar vehicles operating in the real world environment implemented using conventional AI and robotics methodologies. One of the chairs is now capable of travelling to its indoor destinations using landmark-based navigation. The performance of the prototypes indicates there is a cautious possibility today to build a functional intelligent wheelchair that is practical and helpful to people with certain types and degrees of handicap.

9 Acknowledgement

Koichi Ide is responsible for the detailed design and most of the implementation of both TAO-1 and TAO-2. Reuben Martin and Richard Edwards assisted Ide in implementation and testing of both chairs.

References

[Bell et al, 94]
D. A. Bell, S. P Levine, Y. Koren,, L. A Jaros, J. Borenstein,, "Design Criteria For Obstacle Avoidance in a Shared-Control System", *RESNA'94*, Nashville, 1994.

[Braitenberg, 84]
V. Braitenberg, "Vehicles: Experiments in Synthetic Psychology", MIT Press, Cambridge, MA. 1984.

[Brooks, 86]
R. A. Brooks, "A Robust Layered Control System For a Mobile Robot", *IEEE Journal of Robotics and Automation*, RA-2, 1986.

[Brooks, 91a]
R. A. Brooks, "New Approaches to Robotics", *Science 253*, 1991, pp 1227-1232.

[Brooks, 91b]
R.A. Brooks, "Intelligence Without Reason". In Proc. *IJCAI-91*, Sydney, Australia, August 1991.

[Connell & Viola, 90]}
Connell, J. and Viola, P., "Cooperative Control of a Semi-Autonomous Mobile Robot", *Robotics and Automation Conference*, 1990, pp 1118-1121.

[Crisman & Cleary, 96]
J.D. Crisman, and M.E Cleary, "Progress on the Deictic Controlled Wheelchair", AAAI'96 *Fall Symposium Series, Developing Assistive Technology for People with Disabilities*, MIT, 1996, pp12-18.

[Gomi, 96a]
Takashi Gomi, "Non-Cartesian Robotics", *Robotics and Autonomous Systems 18*, 1996, Elsevier, pp 169-184.

[Gomi, 96b]
T. Gomi, "Aspects of Non-Cartesian Robotics", *In Proceedings, Artificial Life and Robotics (AROB'96)*, Japan, February 1996.

[Gomi, 92]
T. Gomi, "Subsumption Robots and the Application of Intelligent Robots to the Service Industry", Applied AI Systems, Inc. Internal Report, 1992.

[Gomi & Ide, 96]
T. Gomi, and K. Ide, "The Development of an Intelligent Wheelchair". *In Proceedings IVS'96*, Tokyo, Japan, September 1996.

[Harvey, 92]
I. Harvey, "Evolutionary Robotics and SAGA: the case for hill crawling and tournament selection", In *Artificial Life 3 Proceedings,* C. Langton, editor, Santa Fe Institute Studies in the Sciences of Complexity, Proc. Vol XVI, 1992, pp 299-326.

[Horswill, 92]
I. Horswill, "A simple, cheap, and robust visual navigation system", *Second International Conference on Simulation of Adaptive Behavior (SAB'92)*, Honolulu, Hawaii , 1992.

[Maes, 92]
Pattie Maes, "Behavior-based Artificial Intelligence", *Second International Conference on Simulation of Adaptive Behavior (SAB'92)*, 1992.

[Pfeifer & Scheier, 96]
R. Pfeifer, C. Scheier, Introduction to "New Artificial Intelligence", Institut für Informatik der Universität Zürich, July 1996.

[Steels, 93]
Luc Steels, "Building agents out of autonomous behavior systems", *The Biology and Technology of Intelligent Autonomous Agents*, NATO Advanced Study Institute, Trento, 1993, Lecture Notes.

[Wang et al, 96]
H. Wang, C-U. Kang, T. Ishimatsu, T. Ochiai, "Auto Navigation on the Wheel Chair", In Proceedings, Artificial Life and Robotics, (AROB'96), Beppu, Japan, February 1996.

[Yanco et al, 95]
H.A. Yanco, A. Hazel, A. Peacock, S. Smith, S., and H. Wintermute, "Initial Report on Wheelesley: A Robotic Wheelchair System", *In Proceedings of the Workshop on Developing AI Applications for the Disabled, held at the International Joint Conference on Artificial Intelligence*, Montreal, Canada, August 1995.