Proceedings of the 1997 IEEE
International Conference on Robotics and Automation
Albuquerque, New Mexico - April 1997

Navigational Support for Robotic Wheelchair Platforms: An Approach that Combines Vision and Range Sensors

P.E. Trahanias, M.I.A. Lourakis, A.A. Argyros and S.C. Orphanoudakis

Institute of Computer Science
Foundation for Research and Technology – Hellas
P.O.Box 1385, Heraklion, 711 10 Crete, Greece
and
Department of Computer Science, University of Crete
P.O.Box 1470, Heraklion, 714 09 Crete, Greece
e-mail:{trahania,lourakis,argyros,orphanou}@ics.forth.gr

Abstract

An approach towards providing advanced navigational support to robotic wheelchair platforms is presented in this paper. Contemporary methods that are employed in robotic wheelchairs are based on the information provided by range sensors and its appropriate exploitation by means of obstacle avoidance techniques. However, since range sensors cannot support a detailed environment representation, these methods fail to provide advanced navigational assistance, unless the environment is appropriately regulated (e.g. with the introduction of beacons). In order to avoid any modifications to the environment, we propose an alternative approach that employs computer vision techniques which facilitate space perception and navigation. Computer vision has not been introduced todate in rehabilitation robotics, since the former is not mature enough to meet the needs of this sensitive application. However, in the proposed approach, stable techniques are exploited that facilitate reliable, automatic navigation to any point in the visible environment. Preliminary results obtained from its implementation on a laboratory robotic platform indicate its usefulness and flexibility.

1 Introduction

Current advances in robotics have facilitated the introduction of related technologies in many application areas, such as surveillance systems, autonomous vehicles, delivery robots and cleaning machines [1]. A distinctive and very important application sector is that of rehabilitation robotics [2]. The latter has been greatly advanced through the introduction of flexible manipulators, mobile platforms, fixed robotic workstations and sensors. The navigational capabilities offered by such hardware components are typically based on range and/or proximity sensor measurements [3].

In this work we are interested in providing navigational assistance to robotic wheelchair users. More specifically, we aim at enhancing current robotic wheelchairs with the capability of automatic "targeted navigation" (move to *that* point). Therefore, a robust and effective navigation approach is required, that will work synergistically with motor impaired users, ultimately providing assistive navigation in uncontrolled environments. In the remainder of this section, navigational capabilities and related shortcomings of current robotic wheelchairs are shortly presented, followed by a brief overview of the proposed approach.

Most navigational approaches in current generation robotic wheelchairs are based on the measurements obtained by a ring of sonars [3]. Sonars provide measurements that can be readily interpreted for obstacle avoidance tasks [4]. There are, however, certain limits to what can be achieved by using only local range measurements. Although they support reactivity to local environment features, they are inadequate for autonomous navigation, since range information constitutes a very restricted environment representation.

To overcome the above problem, environment maps are usually employed [5, 6]. Maps are constructed either by exploiting *a priori* knowledge or by employing an initial *learning* phase and are constantly updated by the local sonar measurements. Motion planning methods [7, 8] are then employed to yield a path from an initial configuration to a desired one and path following is performed by some sort of odometric (dead-reckoning) technique. A non-trivial problem, however, is that path following is not error-free [9]. Another major flaw inherent in these approaches, which also shares responsibility in the above, is the lack of any intermediate localization information. Although there have been some efforts towards automatic localization (e.g. [10]), they are expected to perform poorly in environments cluttered with obstacles.

To deal with the localization problem, radio beacons [11] or easily recognized patterns [12] are usually introduced. Beacons, combined with planning techniques that take uncertainty into account [9, 13], facilitate accurate localization and, therefore, path following. Such environment modifications are, however, very restricting since they confine robot roving in a predefined space. This is more profound in the case of robotic wheelchair platforms, where the goal is to support user mobility in various environments.

Vision, being a more powerful sense, can be em-

*This work was partially supported by EC Contract No ERBFMRX-CT96-0049 under the TMR Programme and the General Secretariat for Research and Technology, Greece, under Grant No. 6060.

ployed towards this end. Indeed, rich information regarding the environment can be extracted from images. Moreover, it is fairly straightforward to assign semantic content to this information. On the other hand, computer vision techniques are not yet reliable enough for coping with the uncertainty and unpredictability of the real world. Thus, the proposed approach is semi-automatic, relying on the user for some decisions that are hard to make automatically, and trying to combine some of the advantages of range sensing and vision by fusing information acquired by a sonar ring and a camera. Computer vision techniques are involved for target tracking; sonar-based reactivity is employed for local, fine control of motion. More specifically, the camera "locks" on a user-selected target, while at the same time the sonars are checking for obstacles that may be in the wheelchair's course. Whenever the wheelchair completes a detour for avoiding an obstacle, the camera instructs it to move in the direction of the target and approach the desired destination. The visual capabilities introduced facilitate an environment representation that is quite appropriate for the task addressed. This is the crux of this work, which integrates existing, robust methods for achieving accurate (targeted) navigation in uncontrolled environments.

The rest of the paper is organized as follows. Section 2 describes the proposed navigation approach and section 3 focuses particularly on its adaptation to wheelchair platforms. A prototype implementation and experimental results are presented in section 4. Section 5 concludes the paper and gives directions for future work.

2 Platform Navigation

As already mentioned, current approaches to automatic navigation employ beacons and environment maps [7, 11], which limits their usability. To compensate for that, in our approach we completely avoid space charting. Instead, we introduce a deictic, visual representation of the target pattern [14] which the user wants to reach. More specifically, we do not employ any kind of workspace maps but use visual images to represent the viewed scene. Since this representation is at the lowest possible (image) level, it does not introduce any errors as would be the case with higher level representations. On the other side, however, it can not support *fully* automatic navigation since it is lacking a detailed environment model. To overcome this handicap, target selection is entrusted to the operator. In other words, our approach circumvents the issue of platform localization and lets the user pick-up a desired environment pattern from an image of the viewed scene; the selected pattern constitutes the target position of the system.

In order to reach the selected target, we employ a *visual fixation capability* and a *hierarchical motion-planner*. Visual fixation is triggered by the target selection and is maintained throughout the whole navigation process. The motion planner operates in two levels. At the higher level, the *global planner* consults the fixation module and commands the platform motion towards the direction pointed to by the vision system, i.e. the platform moves in a straight path towards the target. At a lower level, the *tactical planner* controls local, fine platform motion. To achieve this, it constantly checks for obstacles in the direction of motion, using the platform sonars. In case that an obstacle is encountered, the tactical planner takes over control temporarily from the global planner and performs a motion to avoid the obstacle. During this motion, the global planner is inhibited; visual fixation is, however, active by tracking the icon of the target. When obstacle detour is complete, the global planner resumes control of the platform and redirects its motion towards the target. This procedure is performed in a closed loop until the target is reached, yielding assistive navigational capabilities to the platform. The above two level scheme can also be considered as a realization of the subsumption architecture proposed by Brooks [15]. According to the latter, a hierarchy of concurrent processes are pursuing a common goal, with processes higher in the hierarchy having more general subgoals and lower priority processes inhibiting higher level ones whenever dynamic changes occur in the environment. Figure 1 presents an algorithm in pseudocode that implements our approach; technical issues involved in that are elaborated in the following sections.

```
0. Select Visual_Target
1. Continuously Fixate on Visual_Target
2. Invoke Global_Planner
   2a. Initiate Robot_Motion Towards Visual_Target
   2b. While VisualTargetNotReached do
         Acquire Sensor_Measurements
         Check For_Obstacles
         If ObstacleFound
            Inhibit Global_Planner
            Use Tactical_Planner to Avoid_Obstacle
            Resume Robot_Motion Towards Visual_Target
         EndIf
       EndWhile
3. End.
```

Figure 1: Algorithmic implementation of the proposed approach for assistive navigation.

2.1 Environment Representation

Motivated by recent works on *deictic representations* [14], the proposed navigation approach employs them as an effective model for the external world. A deictic representation is initiated by a higher level process; for the sake of reliability and robustness, this is performed by the user, who selects a target in the viewed scene using appropriate man-machine interfaces.

Upon selection of a visual target, a *marker* M is *bound* to it [14]; this binding is kept permanent for the period of a navigation session. Markers, a central concept in deictic representations around which perception and action revolve, can be thought of as *pointers* to environment objects. When a marker is set to point to an object, it registers the object's features and initiates specific action(s). Formally, we define a marker M as

$$M(F, A)\ ,\ F \in S_F\ ,\ A \in S_A$$

where, S_F denotes the set of object feature values and S_A denotes the set of associated actions. The issue of feature values is deferred until the next subsection where target tracking is discussed. In our model, S_A has only *one* member; in other words, regardless of the target selection, the same action is always invoked. The latter consists of a two-step procedure, corresponding exactly to the two steps 1 and 2 of the algorithm presented in Fig 1. More specifically, S_A is defined as

$$S_A = \{Continuously\ Fixate \text{ on Visual\_Target };\ Invoke \text{ Global\_Planner}\}$$

Intuitively, the above definition associates with each user selection one, twofold system action: (a) fixation on the target which is active at all times during a navigation session, and (b) invocation of the global planner which commands platform motion towards the target.

The employed representation, although rudimentary, is quite adequate for the needs of the application addressed. Moreover, since it remains at the low, image level, it is robust and not susceptible to ambiguities introduced by object- or scene-based representations.

2.2 Visual Target Fixation and Tracking

The deictic representation of the target facilitates visual target fixation and tracking. Since the camera is moving due to the motion of the robotic platform, this is actually the case of a moving observer tracking a target that can be either stationary or moving. This problem has been studied by various researchers, using mainly predictive techniques [16]. However, these techniques have considerable computational requirements and are, therefore, inappropriate for our case where real time performance is required. Consequently, we have adopted a "template-based" approach, where a template is used to represent the target selected and tracking is achieved by a "template-pursuit" technique.

More specifically, user pick-up of a target triggers the construction of a template, which covers an *area of interest* around the target. A region growing technique is employed to delineate the area of interest. This starts with a small window which is used as a "seed" for subsequent growing; window expansion terminates according to a criterion based on the distribution of the color histogram. The contents of set S_F, introduced previously in section 2.1, can now be defined. For a particular object (template), they consist of a representation of the distribution of the object's color histogram. In other words, the object features used for pattern matching refer to the latter distribution. As shown in [17], color distributions can be used to index objects with a good accuracy and, they do not change significantly when computed in sub-parts of an object. Based on this observation, growing of the seed template is terminated when, at a certain point, the color histogram change, with respect to the initial one, is above a predefined threshold.

Following template construction, fixation of the vision system to the selected target is performed. For the case of an intrinsically calibrated camera, this is easily achieved by *pan* and *tilt* motions with angles ϕ_p and ϕ_t, respectively, given as

$$\tan\phi_p = \frac{x}{f} \quad , \quad \tan\phi_t = \frac{y}{f}$$

where, f is the camera focal length and (x, y) are the image coordinates of the template center. The above is illustrated in Fig 2. When f is not known, iterative fixation procedures can be employed for this task [18].

Target fixation is maintained during the course of platform motion by tracking the template across subsequent frames. Template tracking is accomplished by

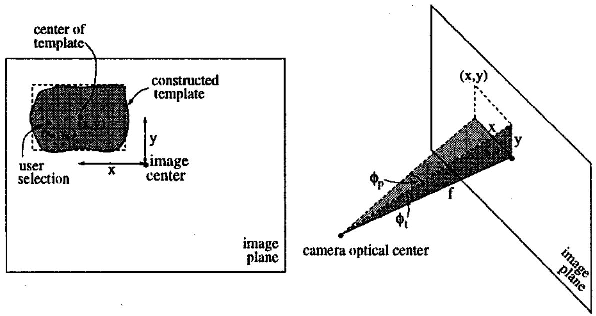

Figure 2: Angles ϕ_p and ϕ_t used in target fixation.

minimizing a sum of squared differences criterion. In order to avoid template comparisons between distant (in time) image frames, the current (stored) template is continuously updated. This involves replacement of the template with the one that constitutes the best match in the next frame.

Visual fixation on the target results in a direction D_v, pointed to by the vision system. D_v is completely determined by ϕ_p and ϕ_t. This is subsequently utilized for guiding the motion of the robotic platform which is performed on the horizontal, 2D plane. Therefore, D_v is simply computed on the horizontal plane, $D_v = \phi_p$.

2.3 Navigation Approach

Using the fixation capability described above, platform navigation is implemented in two levels. At the higher level, the global planner is responsible for achieving the navigation goal, i.e. reaching the selected target. Towards this end, it instructs a motion of the platform in a direction D_p, which coincides with D_v. Due to motor drifts and inaccuracies of the target tracking module, D_p and D_v may change slightly during the course of this motion. This is easily adjusted, however, by using a feedback loop between the output-input of the global planner, as shown in Fig 3(a).

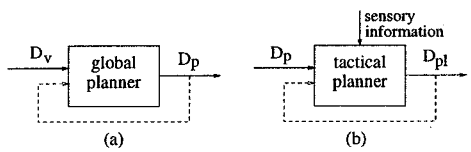

Figure 3: (a) Global planner, (b) tactical planner.

The tactical planner operates at a lower level, being the responsible module when fine, local motion control is required. Towards this, it constantly checks for environment obstacles using the platform sonars. In case that an obstacle is detected at a nearby distance, the tactical planner takes over control from the global planner. It uses D_p and the sonar information to render the *local platform direction* D_{pl}, as shown in Fig 3(b).

The latter is determined as illustrated in Fig 4(a). Let us denote with s_0 the sensor pointing in the direction of the platform motion, D_p. Then, $s_1, s_2, \ldots$, denote the sonars on the one side of s_0, whereas $s_{-1}, s_{-2}, \ldots$, denote the sonars on the other side of s_0. Let us also assume that all sonars between s_{-i} and s_j (included) indicate the presence of an obstacle. Then, the platform local direction of motion D_{pl} is set in the direction pointed to by s_{j+1}, if the angle between the latter direction and D_p is smaller than the angle between D_p and the direction pointed to by $s_{-(i+1)}$ (see Fig 4(a)). Otherwise, it is set in the direction pointed to by $s_{-(i+1)}$. In other words, the platform is set to avoid the obstacle in a way that its initial direction (D_p) is minimally modified.

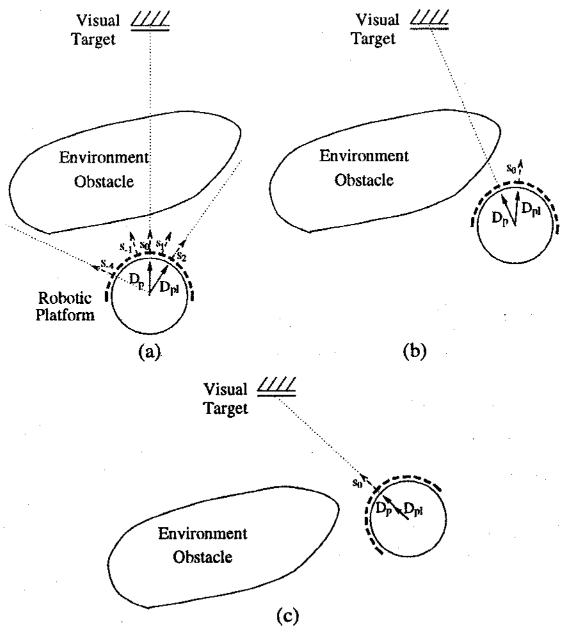

Figure 4: Obstacle avoidance with simultaneous target tracking; see text for explanation.

During this motion in the direction D_{pl}, the vision system maintains D_v, by virtue of the target fixation capability. Therefore, the global planner determines a direction of motion, $D_p = D_v$. The tactical planner tries to minimize the difference between D_p and D_{pl}, taking into account the sonar information. Figure 4(b) shows an intermediate snapshot of the motion commanded by the tactical planner, where this concept is illustrated. When matching between D_p and D_{pl} is made possible (Fig 4(c)), the tactical planner is deactivated and the global planner resumes control again; at this point the platform starts moving in a straight path towards the selected target. Upon reaching the desired target, the user either instructs the platform to stop or selects a new destination.

3 Application to Robotic Wheelchairs

In this section we focus particularly on robotic wheelchair platforms and consider the adaptation of the proposed navigation approach in this case.

3.1 Platform/Workspace Configuration

The proposed approach for assistive navigation of robotic wheelchairs presupposes a certain platform configuration. This is shown schematically in Fig 5.

The platform is assumed to be equipped with a vision system, consisting of a camera mounted on an active head. The head supports at least two degrees of freedom: pan and tilt. The whole system is placed at

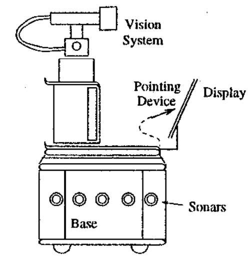

Figure 5: Robotic platform layout.

an appropriate height on top of the platform, so that it overlooks the workspace, without any parts of the scene being occluded by other parts of the platform or by the user. It is also assumed that the head provides a 360^o pan capability. Finally, an image display and a pointing device are assumed to be available on the platform. The display is simply connected to the output of the camera, whereas the pointing device facilitates target selection in the viewed scene. At a relatively low height, a ring of sonars is attached on the platform. The whole configuration is completed with on-board processing power and mechanical motors that execute the platform-motion commands.

Typical navigation sessions are assumed to take place indoors. Although this is by no means an inherent limitation of the approach, current technology and the strict requirements of this sensitive application area suggest indoor environments as an ideal workspace.

3.2 Autonomy vs. Reliability

Various approaches for robotic platform navigation range considerably in the degree of autonomy they support. In the one end, one may regard "manually-controlled" platform motion, whereas, in the other end, fully autonomous navigation in unstructured environments is considered. For the case of robotic wheelchairs, the former approach has already led to market products. Regarding autonomous navigation approaches, they are currently far from being reliable in order to be introduced in this sector. However, there is a clear demand for technology that would increase the independence of people with special needs [19, 2].

Towards this end, the proposed approach for assistive navigation presents a very good compromise regarding the autonomy/reliability trade-off. It relieves the user from the continuous operation of the wheelchair, involving him/her only in the target selection process. On the other side, it does not support navigation in non-visible areas, neither recognition of the target objects. It is, however, of utmost importance that the enhanced navigation capabilities, compared to contemporary ones, are offered without compromising robustness and reliability in platform operation. This is due to the fact that no higher-level, cognitive procedures are involved in any of the steps employed. Moreover, performance reliability has been verified experimentally, as will be presented in the next section.

4 Experimental Results

A prototype of the proposed navigation approach has been implemented and tested in a laboratory environment. Preliminary results are reported here, which demonstrate its feasibility in real scenarios.

4.1 Implementation

The mobile robotic platform available at the Computer Vision and Robotics Laboratory (CVRL) of ICS-FORTH, namely TALOS, has been used as a testbed in all our experiments. TALOS includes:

- A mobile platform (equipped with a 486 and a PENTIUM processors running Linux, wireless communications, sonar, infrared and tactile sensors).
- A binocular, active vision head (independent control of pan, tilt, left and right vergence).

The system is configured so that the PENTIUM is responsible for vision processing and control of the head, while the 486 controls the motion of the robot as well as its sensors. Communication between the two processors is facilitated by the TCX package [20]. The prototype developed on TALOS does not include any advanced user interface for target selection. Currently, this is done interactively by the system operator.

4.2 Laboratory Experimentation

Several experiments have been conducted to test the effectiveness of the proposed navigation approach. In these experiments we have considered various workspace environments and target objects. Moreover, obstacles have been artificially placed to obstruct the initial, straight path to the target. These experiments have verified the appropriateness of this approach for assistive navigation. Results from a sample experiment are presented here for demonstration purposes.

The workspace for this experiment is shown in Fig 6(a). It consists of a room with various "objects" placed in it. The platform's initial position was at the one end of the room. More specifically, Fig 6(b) shows a top view of the workspace; the robotic platform is denoted with the filled circle, whereas its initial orientation is indicated with the corresponding arrow. Some objects (chair, table, box, cart with an amplifier on top of it) are placed in various spots in the room. An obstacle has been intentionally placed as indicated in Fig 6(b).

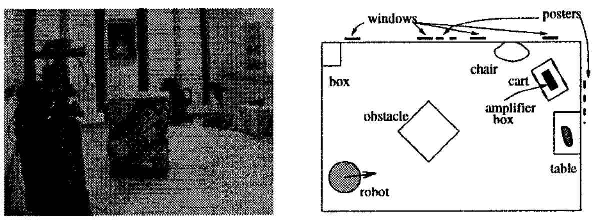

(a) (b)

Figure 6: (a) Workspace for navigation experiments, (b) top view of the workspace.

In this experiment we have simulated a navigation session to reach the amplifier box, lying on the cart at the far end corner of the room. The scene, as viewed by the robot's camera is shown in Fig 7. The target selection, performed by the user, is depicted in Fig 7(a) with an arrow. As can be observed, this selection corresponds to the amplifier box. The template constructed for this user selection is shown in Fig 7(b), as a window superimposed on the actual image.

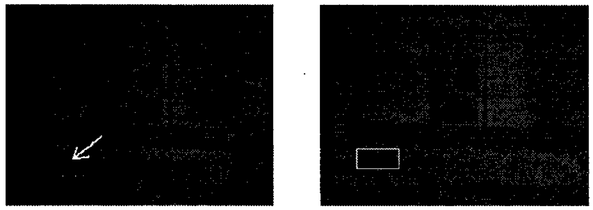

(a) (b)

Figure 7: Scene viewed by the camera mounted on TALOS; (a) user selection, (b) constructed template.

Following that, target fixation has been performed and platform navigation has been initiated in a straight line towards the selected target. This is shown in the first few images of Fig 8. Actually, Fig 8 shows a sequence of snapshots of the whole navigation session, until the target was reached. After a short straight motion, the obstacle, cutting the robot's way, has been detected. The images in the second and third row of Fig 8 illustrate the detour performed by TALOS in order to avoid it. During this, fixation of the vision system to the target has been maintained. After completion of the obstacle avoidance motion, TALOS has again initiated a motion in a straight line towards the target. This is shown in the images in the last row of Fig 8.

The result presented here, although conducted in a laboratory workspace, demonstrates clearly our approach for assistive navigation and serves as an indication of its performance in indoor environments.

Figure 8: Snapshots from a navigation session.

5 Conclusions and Future Work

In this paper we have presented an approach towards introducing assistive navigational capabilities in robotic wheelchair platforms. Contemporary approaches usually fail to support autonomous navigation and provide some low level functionalities; however, the motion of the platform is user controlled. This can be attributed to the fact that these approaches lack some kind of *environment representation* that would facilitate execution of navigation goals.

In our approach, we overcome this limitation by introducing visual representations of the selected target. This is coupled with sonar-based obstacle avoidance techniques. The resulting approach exhibits a navigational behavior that may be useful for robotic wheelchairs. To the best of our knowledge, computer vision has not been employed before in this sensitive application area. This is mainly due to the fact that computer vision is still not reliable enough to be employed in cases where safety and robustness are at a premium. By excluding, however, high level cognitive tasks from the vision system and relying on the user for performing them, we have been able to achieve robust system performance.

The proposed approach concerns only user-selected targets. Towards assisting the user in the target selection process, vision techniques that will suggest regions of interest can be employed. Such techniques will extract precategorical visual information that corresponds to potentially interesting features (e.g. color, symmetry) or dynamic events of the environment. The case of dynamic events is of particular importance, since they signal changes in the environment. For example, recent work on motion perception [21, 22, 23] and tracking of moving objects [24, 18] can be used for enabling a robotic wheelchair to follow a person that moves in the static environment.

Since the proposed approach does not make any limiting assumptions about the "robotic platform", it can effectively be exploited by other robotic actuators, e.g. flexible manipulators, in tasks such as object manipulation. User selection of an object triggers fixation on that, which can then be easily manipulated (e.g. picked up). Such functionalities can be effectively integrated with the navigational functionalities presented above, resulting in robotic platforms with advanced navigational and manipulation capabilities.

References

[1] R.C. Dorf. *Concise International Encyclopedia of Robotics: Applications and Automation.* Wiley, New York, NY, 1990.

[2] E. Donna, C. Bacon, T. Rahman, and W.S. Harwin. *Fourth Intl. Conf. on Rehabilitation Robotics: Conference Proceedings.* Applied Sci. & Eng. Lab., Univ. of Delaware/A.I. duPont Inst., Wilmington, Delaware, USA, 1994.

[3] H.R. Everett. *Sensors for Mobile Robots: Theory and Application.* A K Peters, Ltd., Wellesley, MA, 1995.

[4] J. Borenstein and Y. Koren. The vector field histogram - fast obstacle avoidance for mobile robots. *IEEE Trans. on Robotics and Autom.*, 7(3):278–288, Jun. 1991.

[5] J. Borenstein and Y. Koren. Histogramic inmotion mapping for mobile robot obstacle avoidance. *IEEE Trans. on Robotics and Autom.*, 7(4):535–539, Aug. 1991.

[6] M. J. Mataric. Integration of representation into goal-driven behavior-based robots. *IEEE Trans. on Robotics and Autom.*, 8(3):304–312, Jun. 1992.

[7] J. C. Latombe. *Robot Motion Planning.* Kluver Academic Publishers, Boston, MA, 1991.

[8] Yong K. Hwang and Narendra Ahuja. Gross motion planning – a survey. *ACM Computing Surveys*, 24(3):221–291, Sept. 1992.

[9] H. Takeda, C. Facchinetti, and J.-C. Latombe. Planning the motions of a mobile robot in a sensory uncertainty field. *IEEE Trans. Pattern Anal. Machine Intell.*, 16:1002–1015, 1994.

[10] M. Drumheller. Mobile robot localization using sonar. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, PAMI-9(2):325–332, Mar. 1987.

[11] R.R. Rathbone, R.A. Valley, and P.J. Kindlmann. Beacon-referenced dead reckoning: A versatile guidance system. *Robotics Engineering*, Dec. 1986.

[12] M. Magee and J.K. Aggarwal. Robot self-location using visual reasoning relative to a single target object. *Pattern Recognition*, 28(2):125–134, 1995.

[13] P.E. Trahanias and Y. Komninos. Robot motion planning: Multi-sensory uncertainty fields enhanced with obstacle avoidance. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems, IROS'96*, Osaka, Japan, Nov. 4-8 1996.

[14] S.D. Whitehead and D.H. Ballard. Learning to perceive and act by trial and error. *Machine Learning*, 7:45–83, 1991.

[15] R. A. Brooks. A robust layered control system for a mobile robot. *IEEE J. Robotics Auromat.*, RA-2(7):14–23, Apr. 1986.

[16] A. Blake and A.L. Yuille, editors. *Active Vision – Chapter 1: Tracking.* The MIT Press, Cambridge, MA, 1992.

[17] M. Swain and D. Ballard. Color indexing. *Intl. Journal of Computer Vision*, 7(1):11–32, 1991.

[18] D. Murray and A. Basu. Motion Tracking with an Active Camera. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 16(5):449–459, May 1994.

[19] R. Foulds. Interactive robotic aids - one option for independent living: An international perspective. *Monograph 37, World Rehabilitation Fund, New York*, pages 7–17, 1986.

[20] C. Fedor. *TCX: An Interprocess Communication System for Building Robotic Architectures, Programmer's Guide to version 10.XX*, January 1994.

[21] A.A. Argyros, M.I.A. Lourakis, P.E. Trahanias, and S.C. Orphanoudakis. Independent 3d motion detection through robust regression in depth layers. In *British Machine Vision Conference (BMVC '96)*, Edinburgh, UK, Sep. 1996.

[22] A.A. Argyros, M.I.A. Lourakis, P.E. Trahanias, and S.C. Orphanoudakis. Qualitative detection of 3d motion discontinuities. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems, IROS'96*, Osaka, Japan, Nov. 4-8 1996.

[23] R. Sharma and Y. Aloimonos. Early detection of independent motion from active control of normal image flow patterns. *IEEE Trans. Syst. Man, Cybern.*, 26(1):42–53, Feb. 1996.

[24] N.P. Papanikolopoulos, P.K. Khosla, and T. Kanade. Visual Tracking of a Moving Target by a Camera Mounted on a Robot: A Combination of Control and Vision. *IEEE Trans. on Robot. and Automation*, 9(1):14–25, Fe. 1993.