The State of the Art in Mobile Robotics

James L. Crowley
Institut National Polytechnique de Grenoble
Grenoble, France

Abstract

This paper presents a survey of the information processing aspects of mobile robotics, including architecture, position estimation, world modeling, sensor data integration, vision, and planning. For each topic, a definition is offered, followed by a discussion of the major issues. Points for which a consensus seems to have emerged within the scientific domain, are then presented, followed by a discussion of alternative approaches and currently unsolved problems. The paper finishes with a conclusion section in which some of the more important topics of scientific consensus are reviewed and a discussion is given of some applications areas where mobile robots are thought to be feasible with existing technology. We complete the article with an appendix which lists the names and address for many of the more well known research groups.

1 Introduction

A survey of the state of the art in any scientific domain is a dangerous task. Dangerous, because in order to cover the domain, one must organize the subject into a coherent presentation. No matter how objective one tries to be, the organization of the material reflects unconscious prejudices. Thus one can not escape offending some members of the scientific community. To those who find that this survey has ignored or mis-represented their point of view, I offer my apologies.

In preparing this survey, I have been greatly aided by participation in a NATO workshop on mobile robotics, May 10 to 15, 1987, at Estoril Portugal. At this workshop, a group of 40 of my colleagues met and systematically discussed many the issues concerning mobile robotics. This paper is my attempt to summarize some of the more important points from these discussions. The interesting ideas are provided by my colleagues. The biases are, alas, my own.

1.1 Organization of this Paper

The following sections present a breakdown of the field in a number of topic areas. For each topic, there is an attempt to define the topic and the problems it presents. This is followed by a discussion of issues or aspects of the problem and a discussion of any consensus viewpoints that seemed to have emerged. Following this is a discussion of unusual or alternative approaches to the problem. Finally, some sections finish with a discussion of currently unsolved problems. I have not, however, held rigidly to this schema. In some sections I have skipped some of these sub-section because nothing remained to be said.

In the conclusion section I attempt to extract some genuine conclusions from all this. In particular, I synthesize my opinions of which issues in mobile robotics are solved problems and to what application areas mobile robots may be introduced today.

1.2 Scope of the Paper

In order to limit the size of the paper (and the quantity of my work) I have taken a computer scientist's view of what is interesting in mobile robotics. This paper concerns the information processing aspects of a mobile robot. Thus it ignores issues such as power generation, and electrical distribution. Perhaps with less justification it also ignores the mechanical aspects of mobility (i.e. wheel configuration). This is a conscious decision: Mobility is simply too domain dependent to be adequately treated in a paper of such limited scope.

2 Architecture

Most of the issues addressed below are influenced by the information processing architecture of a mobile robot. Thus we begin by considering such architectures.

2.1 Definition of Architecture for a Mobile Robot

The architecture of a mobile robot is defined to be the decomposition and the inter-connection of the information processing sub-systems. Such sub-systems typically receive signals from sensors and process these to produce control signals for actuators. The mechanical and power sub-systems are deliberately not included in this definition, although admittedly they form an important aspect of any mobile robot.

We note that the information processing architecture is distinct from the computer architecture on which it is implemented. For a given decomposition of the perception and control systems there are often many ways that these systems may be mapped into processing hardware. It is our view that the form of the computer architecture should follow the function of the information processing architecture. Granted that for any individual mobile robot project, the information processing architecture is often constrained by the available computer hardware.

2.2 Issues or aspects of problem

The basic architectural question concerns how to decompose and interconnect the perception and control sub-systems. Two related issues are debated with in the mobile robot scientific community at present:

1) Should the architecture by organized as a set of layers distinguished by response time and level of abstraction of the information processed?

2) Is an explicit geometric model of the local environment necessary for local navigation and sensor perception?

2.3 Consensus Concerning Multiple Layers

With two notable exceptions described below, both the scientific and applications community have converged towards a consensus that the control and perception processes for a mobile robot <u>do decompose</u> naturally into a set of layers. Even more interestingly, the same set of layers seem to appear across a wide spectrum of projects, including projects involving tele-operation. For most of these layers, the issue that is discussed is the <u>Name</u> that should be applied to its existence.

An example of a four layer decomposition, first published in last year's symposium [Crowley 86] is presented in figure 1. This decomposition is dictated by both the response time of the feedback loops at each layer as well as the level of abstraction of the information. Such a decomposition groups perceptual and control together into feedback processes at each layer.

At the lowest level are the individual motor controllers as well as processes which take in and describe sensor signals. Both processes (ideally) operate with cycle times on the order of 10 to 100 milliseconds (100 to 10 cycles per second). Their is little disagreement over the need for such a layer. There are many groups working with a variety of approaches to produce hardware to make possible to perform the lowest level perception. Motor controllers which operate at such rates have long been commercially available. We repeat the caution that motor control at this level must accept commands asynchronously if the more abstract layers of the system are to function.

Sensor Integration and Vehicle Control: At the second layer, perceptual information is integrated and control assured in terms of the vehicle's behavior. Processes at this level should ideally operate with cycle times of 100 millisecond to 1 second. A majority of the community seems to agree that the integration of perceptual information from different sources and different requires a geometric model expressed in a common coordinate system. A disagreement occurs over whether this coordinate system should be in terms of the vehicle or a fixed external coordinate system. Proponents of vehicle coordinates argue that relative position information for local landmarks may be known with more precision than the vehicles location in the external world. Proponents of a local fixed coordinate frame maintain that a fixed local coordinate frame facilitates recalling and integrating expectations from a pre-learned model of the environment. Reflex responses may occur at both the level of signal processing and control, and the level of sensor integration and vehicle control.

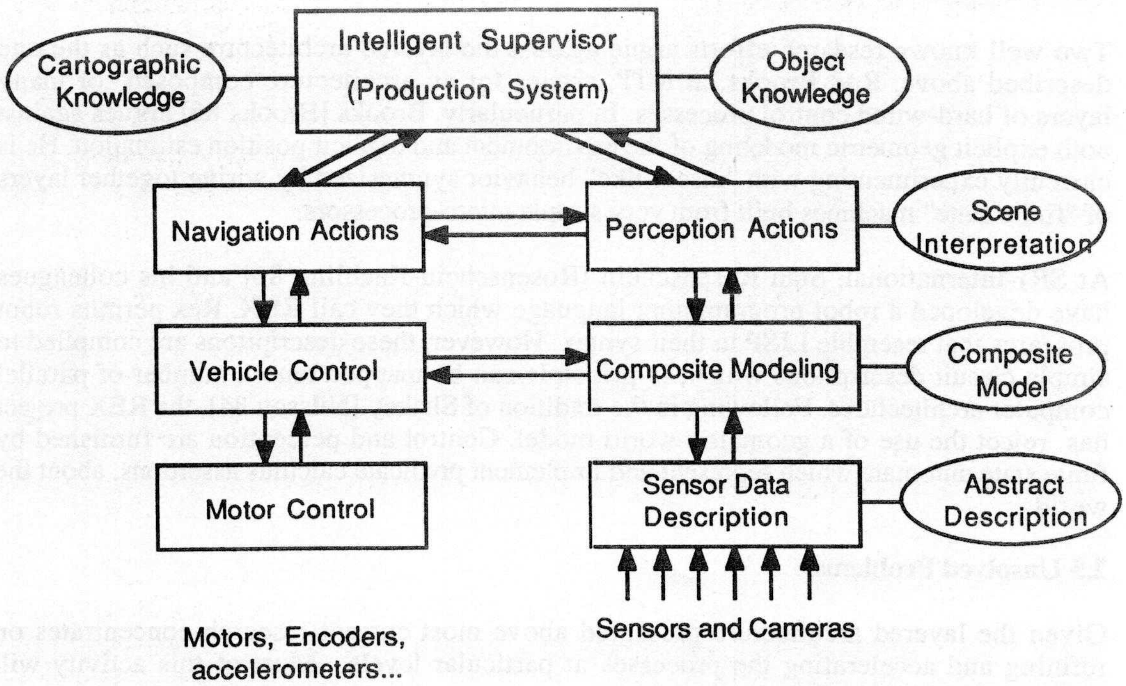

Figure 1. The architecture of a mobile robot is naturally decomposed into layers for 1) motor control and signal processing, 2) sensor integration and vehicle control, 3) perception and locomotion actions, 4) Intelligent Supervision.

Perceptual and Locomotive Actions: The third layer in the architecture integrates what we refer to as perception and locomotion "actions" [Crowley 87a]. These are algorithmic processes which might also be referred to as skills. Examples of perceptual actions are the functions Find-Object, Find-Path, Free-Path, or Find-Road. Locomotion actions might include Follow-Road, Follow-Wall or Go-Around-Obstacle. There does not seem to be much controversy over the need for such processes. The major issue concerns whether such processes should be hard-wired, programmed algorithmically or rule based. While certain members of the AI community have argued for a rule based approach, it seems that no one within the mobile robot community currently defends such an approach. The reasons seem to be problems of complexity and temporal response. Perception and locomotion tasks require cycle times ranging from 100 milliseconds for road following to 10 seconds for local path planning.

Intelligent Supervisor: At the top of these layers is some kind of a supervisor level. While it is generally agreed that such a layer is needed, little work has yet been done on this level. We argue [Crowley 87b] that such a level requires reasoning with a large amount of poorly structured domain specific knowledge. As a result we use a forward-chaining production system, implemented in OPS-5 for this level. Herman [Herman-87] has recently organized similar knowledge is a set of states, described as a finite state machine. We note that the contexts of our production system are also expressed as a set of states. The principal difference would seem to be the implementation of reasoning within such each state. Control cycles at this level may operate from 1 to 100 seconds.

2.4 Alternate Architectures

Two well known research efforts argue against the layered architecture such as the one described above. Rod Brooks, at MIT, argues for an architecture composed of many layers of hard-wired control processes. In particularly, Brooks [Brooks 86] argues against both explicit geometric modeling of the environment and explicit position estimation. He is currently experimenting with "insect-like" behavior synthesized by wiring together layers of "finite state" machines built from very simple micro-processors.

At SRI-International, Stan Rosenschein [Rosenschein-Kaebling 86] and his colleagues have developed a robot programming language which they call REX. Rex permits robot programs that resemble LISP in their syntax. However, these descriptions are compiled to simple circuit descriptions which, in principle can be mapped onto a number of parallel computer architectures. Following in the tradition of Shakey [Nilsson 84], the REX project has reject the use of a geometric world model. Control and perception are furnished by finite state automata which represent and implement predicate calculus assertions. about the world.

2.5 Unsolved Problems

Given the layered architecture presented above most current research concentrates on refining and accelerating the processes at particular levels. Most of this activity will become evident in the following sections.

3 Position Estimation

Navigating requires knowing where the vehicle is with respect a goal. Thus, it is a bit surprising that the issue of position estimation remains a hotly debated topic. The most interesting new twist is the recognition of the utility of an explicit representation of uncertainty in the estimated position.

3.1 Definition

The estimated position of a mobile robot includes at least an estimate of the location and orientation of the vehicle with respect to some coordinate system. In many systems it is necessary to complement this information with the vehicle velocity and acceleration as well as the velocity and acceleration in orientation. This information is increasingly coupled with a covariance estimate of the uncertainty of these parameters. In view of this additional information, a better term might be estimation of the vehicle's "dynamic state".

3.2 Issues or Aspects of Problem

Among the majority who agree on the utility of estimating part or all of the robot vehicle's dynamic state, the following issues are discussed.

1) Should position and orientation be expressed in a relative or absolute coordinate system.
2) How should a dynamic model of the uncertainty in the vehicle state be measured and expressed.

3.3 Consensus

There is little argument over the utility of explicitly representing the uncertainty in the vehicle's position and orientation. Most groups also agree that such uncertainty may be represented by a covariance matrix and updated using the Kalman filter formulation [Smith-Cheeseman 87]. It is also generally accepted that the first and second derivatives of position and orientation are useful for modeling the vehicle's state. Most discussion concerns techniques for measuring these parameters and their uncertainty for specific vehicles. In general, measuring the uncertainty requires identifying and modeling sources of uncertainty in the vehicle mechanics. Most current commercially available systems navigate using beacons. The Kalman filter makes it possible to estimate a vehicles location from the observed angle to a beacon at a known location.

The point of view of representing the vehicle position and orientation relative to the immediate environment has been eloquently defended by Binford. He has argued that it is often possible to known the relative position of objects with much greater accuracy than their position in some absolute coordinate system. Chatila [Chatila 84] and Cheeseman and Smith [Smith-Cheeseman 87] have proposed specific techniques to handle such a problem.

An alternative view-point argues that for many applications, this problem does not present itself, where-as a fixed local coordinate system permits expectations of sensor data to be recalled, thus greatly facilitating both position estimation, and interpretation of noisy perceptions.

3.4 Alternative Approaches

Both Brooks at MIT and Rosenschein and Kaibling at SRI argue that an estimate of the vehicle position is not necessary with their approaches.

3.5 Unsolved Problems

Much research centers around modeling the sources of error in the estimated position and orientation for specific vehicles. There is also significant work in measuring the vehicle motion using visual motion techniques. Most such techniques search for a coherent displacement velocity from the 2-D image flow (tokens or spatio-temporal derivatives). Olivier Faugeras, of INRIA now argues for measurement of motion by matching of observed 3-D shapes to a dynamically maintained 3-D local model.

4 Modeling and Sensor Fusion

All known perceptual processes are noisy. Driven in part by the needs of mobile robots, sensor data integration (also known as sensor data fusion) is an area of great activity for both mobile robotics and vision.

4.1 Definition

We define the problem of integrating sensor data as "the combining of independent and qualitatively different sensor data to reinforce consistent interpretations and disregard inconsistent interpretations". Sub-problems include integrating information from the same sensor taken from different view-points as well as integrating expectations with perceived data. Most researchers would agree in its richest form, this problem involves integrating information from more than one sensing modality. Different techniques for extracting 3-D shape from images, (shape from motion, shape from stereo, shape from texture, etc.) are considered as different sensing modalities, or "logical sensors" which are based on the same physical sensor (1 or more cameras).

4.2 Issues or aspects of problem

Debate over this issue usually concerns the following questions:

1) Is it useful to maintain a local model of the environment, or is it sufficient to react to immediate sensory perceptions?
2) How should the information in such a model be represented?
3) Is a 2-D model sufficient or is a 3-D model necessary?
4) How can imprecise and measurements be represented and combined?
5) How can the uncertainty in existence be represented and updated?

4.3 Consensus

A surprising and sudden consensus has taken form in the last few years. A number of scientists now argue that some form of dynamically maintained local model of the geometry of the environment is useful, both for the integration of noisy perceptual data and for the local control of vehicle and locomotion actions. Most researchers agree that the key to combining disparate sensor data lies in an explicit representation of uncertainty using covariances and some form of Kalman filter. Most researchers also agree that geometric information forms the basic framework for such a local model. Color, texture, temperature and other properties may be attached to the geometric primitives that make up such a model. Most researchers would also agree that when vehicle motion is in a 2-D plane, a 2-D description of the limits of free-space is the minimal information needed for vehicle control. Examples of systems which operate using such information include [Chatila 82], [Laumond 84], [Kuan 85], [Crowley 85], [Binford 87], [Arkin 87], [Weisbin 87] and [Elfes 85].

Many researchers feel that the integration of visual information requires a 3-D local model. [Faugeras, Lustman and Toscani 87], [Ayache and Faugeras 87] have recently demonstrated a system which models the environment using stereo. Based on their results, Faugeras has recently argued that measuring image flow is unnecessary, as the same information is available more reliably from matching 3-D measurements from 3-Camera stereo.

Most researchers address the sensor data fusion problem with geometric primitives for surfaces or volumes. Such primitives are typically represented by a set of parameters. The classic form of such primitives are generalized cylinders, developed by Binford, Agin and Nevatia [Agin 76]. An example of similar primitives for 3-D surfaces, contours and corners are discussed in [Crowley 87a]. Such primitives for 2-D line segments is discussed in [Crowley 87c]. Explicitly including the uncertainty for each parameter provides the means for integrating imprecise and varied perceptual data. This point has recently been championed by Durrant-Whyte [Durrant-Whyte 86].

4.4 Alternative Approaches

An interesting alternative to parametric geometric primitives is the occupancy grid. Various forms of such an approach have been reported by Moravec and Elfes [Elfes 85] as well as Weisbin and his associates [Weisbin-87]. The basic idea is to model the local environment as a 2-D or 3-D grid. Each grid cell may be marked as occupied, free or unknown. In the Moravec-Elfes system, occupancy is represented by a number in the range -1 (occupied) to 1 (free). A form of MYCIN like update rule is used to update the confidence occupancy based on projections of sonar data. An interesting aspect of this approach is that path planning with a potential field or velocity field may be performed directly within such an array by either graph search or various cellular automata processes. This raises the potential for a single "chip" which both integrates 2-D free space information and directly plans local paths that avoid obstacles.

5 Vision

For diverse reasons, most researchers agree that passive vision is the ideal perceptual method for a mobile robot. Systems that exploit active range measurement systems (i.e. laser ranging, infrared ranging, light stripe, and ultra-sonic ranging) usually do so to avoid the high computational cost and slow response time of current vision technology. However, with the steady decrease in the cost and size of computer hardware, vision is becoming a feasible alternative.

5.1 Definition

A conservative definition of a computer vision system is one which measures 2-D and 3-D information from 2-D images produced from visible and near infra-red light by a some form of television camera.

5.2 Issues or aspects of problem

It is possible to view activity in vision for mobile robots along the independent axes of sophistication and speed, as shown in figure 2. While most researchers would prefer both speed and sophistication, they are forced by existing hardware to use either fast simple vision or slow sophisticated vision. A trade-off curve exists between the extremes. Fortunately, the hardware trend is favorable, and the entire curve is moving up and to the left.

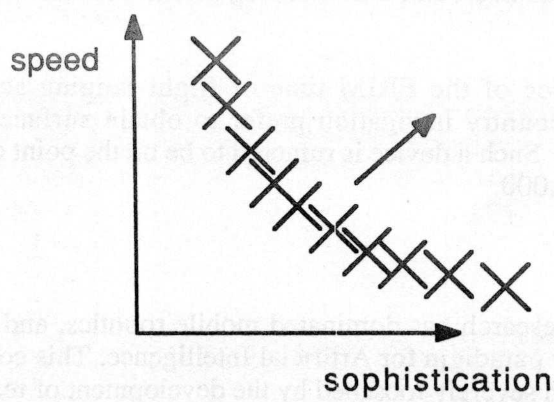

Figure 2. Idealization of the trade-off between speed and sophistication in machine vision.

A classic example of a simple fast system is the road following hardware developed by Graephe and Dickmanns at the Universität der Bundeswehr München in Nuebiberg, FRG. [Dickmanns 86]. Their system uses specially designed hardware modules to correlate an ideal road edge at video rates with windows selected from a road image. By measuring the left and right road edges at two distances, the system has successfully steered a truck at 100 km/hour on the auto-bahn. The system is optimized for the German Autobahn and would not handle many situations that occur in non-highway roads. At the other end of the spectrum, one could consider the 3-Camera stereo system of Faugeras and Lustman, which currently takes several minutes to produces very accurate and complete 3-D scene descriptions. Between these extremes one may place the variety of systems developed at C-MU using color and texture to detect the road surface [Stentz-Goto 87], [Wallace 87], as well as such systems as the 2-D stereo system for mapping a hallway developed by Binford, Kreigman and Treindl [Binford 87].

5.3 Consensus

There is very little consensus on how to accomplish vision mobile robots in real time. Major issues include:

1) Fast 2-D image measurements vs. 3-D scene description.
2) The use of special purpose vs. general purpose vision hardware.
3) The usefulness of measuring camera motion from 2-D image flow.
4) Edge detection vs. segmentation of images.

Perhaps the only point on which everyone agrees is that processors based on the standard serial or "Von Neuman" architecture will never be sufficient to perform the low level aspects of vision in real time.

5.4 Alternative Approaches

The major issue in computer vision is how to get algorithms to deliver useful information at real time rates. The major controversy concerns the approach to this goal. One side, typified by the C-MU ALV project, argues for fast, programmable, parallel vision architectures which will support a large variety of techniques. The alternative, typified by Faugeras, argues for two step process in which working algorithms are developed on general purpose serial machines, and then these algorithms are cast in special purpose hard-wired processors.

Spurred by the existence of the ERIM time of flight ranging scanner, many groups concerned with cross country navigation prefer to obtain surface information from a scanning ranging device. Such a device is rumored to be on the point of entering the market at a price of under $100,000.

6 Planning

Historically, planning research has dominated mobile robotics, and planning for mobile robots provided a major paradigm for Artificial Intelligence. This connection between the AI and robotics has been severely loosened by the development of real time mobile robots. A definite trend now exists toward replacement of planning with real time control.

6.1 Different Planning Problems for Mobile Robots

Planning could be defined as the decomposition of a goal into a set of actions. With regard to mobile robots, It is important to distinguish 3 types of planning:

1) Task level or Mission Planning,
2) Global Path Planning, and
3) Local Trajectory Planning.

Of the three kinds of planning listed above, Task level or Mission Planning bears the closest resemblance to activities called planning by the artificial intelligence community. The main goal of such planning is

1) decomposition of a mission statement into tasks which can be performed by the robot, and
2) Verification of the feasibility of the mission.

The work that has been performed for mobile robots in this area tends to recognize that such planning requires manipulation of a large amount of poorly structured knowledge. Thus the AI tools developed for expert systems are well suited for such planning. We cite our own work in mission planning for a surveillance robot as representative of such systems [Crowley 87b].

Global path planning involves constructing a sequence of locomotion actions in order to arrive at a stated goal location. The essential information for such planning is <u>topographic</u>, i.e. the connectivity between known places. A strong consensus exists that if the environment is known, than global path planning is a solved problem. The accepted technique is to express the environment as a network of decision points and then to search the network using either the A* algorithm [Nilsson 80] or Dykstra's algorithm [Aho, Hopcroft and Ullman 83]. A major disagreement exists over whether such graph search is better accomplished by a set of production rules or an algorithmic implementation.

Local trajectory planning refers to the construction of a local vehicle trajectory. Such process typically operate from sensory perception or some form of composite local model. Typical applications involve maneuvering among obstacles, following walls, going through doors, and turning corners. A consensus has been emerging that trajectories should be planned and controlled using the dynamic state of the vehicle. Typically such state is the position and heading and their first 2 derivatives. Interestingly, such motion corresponds to a set of curves known as "clothoid" curves [Kanayama 85]. Binford refers to this family as "Euler Spirals", remarking that Euler is the first to have reported their existence. Most trajectory planning systems plan a set of straight line movements, joined by clothoid corners using the composite local model. Path execution is then assured by a vehicle level controller that generates commands to the motor controllers. Interesting work has been done by Laumond at LAAS in the problem of "parallel parking" vehicles with wheel configurations that resemble those of a car.

An alternative approach to local trajectory planning is provided by variations of field theory. Originally proposed by Khatib [Khatib 87] for controlling arms, both potential fields and velocity fields have been widely explored for local trajectory generation for mobile robots [Thorpe 84], [Krogh 84], [Arkin 87].

7 Conclusions:

A number of conclusions can be drawn from this survey. This section will attempt to synthesize conclusions concerning both the scientific issues, and what applications are do-able with the state of the art today.

7.1 Scientific Issues

The following two paragraphs are the authors opinion on the state of the art in mobile robotics as of May 1987. It is highly opinionated and subject to debate.

Sufficient techniques now exist for a mobile robot to function in a known domain. The path planning problem can be divided into a global path planning system and a local trajectory generation or control system. The essential information for global path planning is topological: the interconnection of known decision points. The essential information for local trajectory planning is geometric in 2-D: the limits to free space, as well as the vehicle dynamics. A control theory approach to vehicle control is increasingly accepted as necessary. The vehicle state and its derivatives are seen as the relevant parameters. An explicit model of the uncertainty in vehicle state can be very useful. Position estimation is accepted as a problem of complementing odometry and inertial sensing with perception and beacon detection.

In the area of perception, direct range sensing and stereo vision continue to compete, although neither technique is yet fast enough on commercially available hardware. However, the hardware support is advancing rapidly; laboratory prototypes will increasingly appear over the next 5 years, with commercial products 2 to 3 years behind. The use of visual motion for motion estimation and structure is increasingly regarded as noisy and noise sensitive. The most realistic sensors for the moment remain the large foot-print ultrasonic sensors, although it is widely recognized that as much as 2/3 the measured data can be useless in some environments do to specularities. A geometric model of the environment is generally recognized as useful for, integrating sensor data and expectations, for planning and executing actions and for position estimation.

7.2 Applications of Mobile Robots

The most interesting conclusions concern what is do-able today with mobile robots. Given a properly engineered vehicle, a number of applications may be performed by mobile robots using technology that has been demonstrated under laboratory conditions. Applications which are accessible include the following.

Floor Cleaning: Principle problems include position estimation, environmental perception, and communicating a model of the area to be cleaned. Position estimation may use odometry and inertial sensing complemented with passive or active beacons. Perception of the limits of free space may be done with ultra-sound or infra-red sensors. The model of the area to be cleaned may be communicated with a bit mapped display and a mouse. Prototypes of such robots are under test at at least 4 laboratories world-wide.

Painting in construction: The technologies needed to build a robot to paint the ceiling, walls and floor of a newly constructed building are very similar to those of a floor-cleaning robot. This is primarily a question of proper engineering.

Parts Handling: Automatic guided vehicles (AGV's) already exist. The trend is toward systems which do not require stripes of wires on the floor. The local navigational needs are somewhat similar to the needs for floor-cleaning. A parts-handling robot also needs a "map" which states the topography of the environment (a network of places). Provisions are also needed to communicate goals and query the status of such a robot. Prototypes of such robots are known to be undergoing testing.

Surveillance Robot: The local and global navigational needs are similar to those of a parts-handling robot. A considerable amount of "payload engineering" is necessary. A surveillance robot which navigates with beacons is commercially available from Denning Mobile Robotics. Other prototypes are known to exist.

Military Applications on Land: This year has seen successful examples of real time road following at the Université Bundeswerh, Carnegie-Mellon University, Martin Marrietta, and FMC Corporation. Successful cross terrain navigation of a known environment has been demonstrated by FMC. Navigation across a poorly known terrain is now of intense interest to all of the above groups as well as the Jet Propulsion Laboratory. Numerous military applications are likely to follow.

Bibliography

[Agin 76] Agin G. J. and T. O. Binford, "Computer Description of Curved Objects", IEEE Trans. on Computers, C-25(4), April 1976.

[Aho, Hopcroft and Ullman 83] Aho, A. V. , J. E. Hopcroft, and J. D. Ullman, "Data Structures and Algorithms", Addison Wesley Publishing Co., 1983.

[Arkin 87] Arkin R. C. , "Motor Schema Based Navigation for a Mobile Robot", 1987 IEEE International Conference on Robotics and Automation, Raleigh, March 1987.

[Ayache and Faugeras 87] Ayache, N. and Faugeras, O. D, "Building, Registering and Fusing Noisy Visual Maps", International Conference on Computer Vision, June 87.

[Binford 87], Kreigman, J. K., E. Treindl and T. O. Binford T. , "A Mobile Robot: Sensing, Planning and Locomotion", Fourth IEEE COnference on Robotics and Automation, Raleigh N. C., March 1987.

[Brooks 86] Brooks, R., "A Robust Layered Control System for a Mobile Robot", IEEE Journal of Robotics and Automation, RA-2, no. 1 1986.

[Chatila 82] Chatila, R. "Path Planning and Environmental Learning in a Mobile Robot System." In ECAI. Orsay, France, August, 1982.

[Chatila 85] Chatila, R. and J.-P. Laumond , "Position Referencing and Consistent World Modeling for Mobile Robots", 1985 IEEE Conference on Robotics and Automation, St. Louis, March 1985.

[Crowley 85] Crowley, J. L. "Navigation for an Intelligent Mobile Robot", IEEE Journal on Robotics and Automation, 1 (1), March 1985.

[Crowley-Coutaz 86] Crowley, J. L. et J. C. Coutaz, "Navigation et Modèlisation pour un Robot Mobile", Technique et Science Informatique, Octobre, 1986.

[Crowley 86] Crowley, J. L. "Navigation and World Modeling for a Mobile Robot: A Progress Report", 3rd International Symposium on Robotics in Construction, Marseilles, June 1986.

[Crowley 87a] Crowley, J. L. "Using the Composite Surface Model for Perceptual Tasks", 1987 IEEE International Conference on Robotics and Automation, Raleigh, March 1987.

[Crowley 87b] Crowley, J. L., "Coordination of Action and Perception for a Surveillance Robot, Proc. of Tenth IJCAI, Milan, Aug, 1987.

[Crowley 87c] Crowley, J. L., "Mathematical Tools for Representing Uncertainty in Perception", To Appear.

[Dickmans 86] Dickmanns, E. D. and A. Zapp, "A Curvature-based Scheme for Improving Road Vehicle Guidance by Computer Vision", SPIE Conference on Mobile Robots, Cambridge Mass., USA, Oct. 1986.

[Durrant-Whyte 86] Durrant-Whyte, H. F., "Consistent Integration and Propagation of Disparate Sensor Observations", Third IEEE Conference on Robotics and Automation, San Francisco, April 1986. (To appear in Journal of Robotics Research).

[Faugeras and Hebert 83] Faugeras, O. D. and M. Hebert, "A 3-D Recognition and Positioning Algorithm using Geometrical Matching between Primitive Surfaces", Proc. of Eighth IJCAI, Aug, 1983.

[Faugeras, Lustman and Toscani 87] Faugeras, O. D, F. Lustman and G. Toscani, "Motion and Structure from Motion from Point and Line Matches", International Conference on Computer Vision, Brighton, June 87.

[Goto-Stentz 87], Goto, Y. and A. Stentz,"The CMU System for Mobile Robot Navigation", 1987 IEEE International Conference on Robotics and Automation, Raleigh, March 1987.

[Kaebling 1986] Kaebling, L. P. " An Architecture for Intelligent Reactive Systems", SRI Technical Note 400, October 1986.

[Kanayama 85] Kanayama, Y. "Trajectory Generation for Mobile Robots.", Third Int. Symp. on Robotics Research, ISRR-3, Paris, Oct. 1985.

[Khatib 85] Khatib, O., "Real Time Obstacle Avoidance for Manipulators adn Mobile Robots", 1985 IEEE Conference on Robotics and Automation, St. Louis, March, 1985.

[Krogh 84] Krogh, B., "A Generalized Potential Field Aproach to Obstacle Avoidance Control", in Robotics Research, the next 5 years and Beyond, SME, One SME Drive, Dearborn, MIch., 1984.

[Kuan 85], Kuan D., R. A. Brooks and J. C. Zamiska, "Natural Decomposition of Freespace for Path Planning", 1985 IEEE Conference on Robotics and Automation, St.

Louis, March, 1985.

[Herman 87] Herman, M., Personal Communication, 1987.

[Laumond 83] Laumond, J. P., "Model Structuring and Concept Recognition: Two Aspects of Learning for Mobile Robot." In Proc. of the Eighth IJCAI, William Kaufman Inc., August, 1983.

[Laumond 87] Laumond, J.-P., Finding Collision Free Smooth Trajectories for a Non-Holonomic Mobile Robot, Proc. of Tenth IJCAI, Milan, Aug, 1987.

[Elfes 85], Elfes R., "A Sonar Based Mapping and Navigation System", 1985 IEEE Conference on Robotics and Automation, St. Louis, March 1985.

[Nilsson 80], Nilsson, N. "Artificial Intelligence", Tioga Press, Los Altos Ca. 1980.

[Nilsson 84], Nilsson, N. "Shakey the Robot", Technical Note 323, SRI International Artificial Intelligence Center, Menlo Park, Cal. 1984.

[Parodi-84], Parodi, A. "A Route Planning System for an Autonomous Vehicle", the First IEEE Conference on Artificial Intelligence Applications, Denver, December 1984.

[Rosenschein-Kaebling 86] Rosenschein, S. J, and L. P. Kaebling, "The Synthesis of Digital Machines with Provable Epistemic Properties", Proc. Conference of Theoretical Aspects of Reasoning about Knowledge, Asilomar, 1986.

[Smith-Cheeseman 87], Smith, R. C. and P. C. Cheeseman, "On the Estimation and Representation of Spatial Uncertainty", International Journal of Robotics Research 5 (4), Winter, 1987.

[Thorpe 84], Thorpe C. E. "Path Relaxation: Path Planning for a Mobile Robot.", Proceedings of the 1984 AAAI, Kaufman Press, 1984.

[Wallace et. al 85] Wallace, R. W. et. al., "First Results in Road Following", IJCAI 85, Los Angeles, August 1985.

[Wallace 87], Wallace, R. W., "Robot Road Following by Adaptive Color Classification and Shape Tracking", 1987 IEEE International Conference on Robotics and Automation, Raleigh, March 1987.

[Weisbin 87] M. Goldstein, F. G. Pin, G. de Saussure and C. Weisbin, "3-D World Modeling Based on Combinatorial Geometry for Autonomous Robot Navigation", 1987 IEEE Conference on Robotics and Automation, Raleigh N. C., March 1987.

Appendix. A Partial List of Research Efforts in Mobile Robotics

The following is an annotated list of active projects in mobile robotics for which sufficient information is known to the author. We are sure that we have omitted some major projects, and apologize in advance to those research teams which are not included. In particular we are certain that there are many fine projects in Japan of which we are ignorant.

LAAS-CNRS, T. Ave. Colonel Roche, 31077 Toulouse, France. Project Hilare, Principle investigators: George Giralt, Raja Chatila, J-P. Laumond. Perhaps one of the longest running projects. Recent applications include a cleaning robot for the Paris Metro as well as a ware-house robot.

SRI-International, AI Group, 333 Ravenswood Ave., Menlo Park, California, 94025 USA. Principle Investigators: Stan Rosenschein and Leslie Kaibling. Robot hardware: Flakey the robot, designed and built by Stan Reifle.

C-MU Robotics Institute, Schenley Park, Pittsburgh Pa, 15213 USA. Principle investigators: Takeo Kanade, Chuck Thorpe, Hans Moravec, Red Whittaker. At least 5 distinct projects have been performed at C-MU in the last 5 years. Major projects include The ALV NavLab, the Terragator, the Denning surveillance robot, the IMP, and Moravec's Pluto, Neptune and Uranus.

MIT AI Lab: 545 Technology Square, Cambridge MA, 02139, Principle Investigator: Rod Brooks. Robots: Allen, Tom, Jerry, Sydney, Seymour.

INRIA: Domaine de Voluceau, Rocquencourt, BP 105, 78153 Le Chesnay, France. Principal Investigators: Olivier Faugeras, Nicholas Ayache, Francis Lustman. A commercial copy of the INRIA robot is sold by RobotSoft SARL, of Asniers France.

LIFIA: INPG, 46 ave Fèlix Viallet, 38031 Grenoble, France. Principal Investigator: Jim Crowley. Developing a Surveillance Robot for project EUREKA - Mithra. Currently using a Denning robot named Lurch.

GM Research, Warren Michigan, USA: Principle Investigators: Steve Holland, Bob Tilove

Univ. of Amsterdam, Kruislaan 4090, The Netherlands. Principle Investigator: Dr. Willem Duinker.

Stanford University: AI Laboratory, Computer Science Dept, Stanford University, Stanford, Ca., 94305 USA. Principal Investigator: Thomas Binford.

ORNL: P.O. Box X, Oak Ridge Tenn, 37831 USA. Principle Investigator: Charles Weisbin.

FMC Corp. 1205 Coleman Ave., Santa Clara, Ca. 95052. Principle Investigator: Andrew Chang Military applications of mobile robots.

NBS: Industrial Systems Division, National Bureau of Standards, Building 220/B124, Gaithersberg, MD. 20899. Prinicipal Investigator: Marty Herman.

Insitüt der Bundeswehr Mûnchen, Inst. fur Messtechnik, 8014 Neubiberg, W. Germany. Principle Investigator: Volker Graefe. Real time road following system integrating simple real time vision with control theory.

Jet Propulsion Laboratory, 4800 Oak Grove Drive, Pasadena, Ca. 91109 USA. Principle Investigator: Brian Wilcox.