Robot Navigation

Panel Chairperson:
David P. Miller
Jet Propulsion Laboratory/California Institute of Technology
Pasadena, CA 91109 U.S.A.

Panel Members

Rod Brooks,
AI Lab, MIT
545 Technology Square
Cambridge, MA 02139

Raja Chatila
LAAS-CNRS
7, Ave. Du Colonel Roche
31077 Toulouse Cedex, France

Scott Harmon,
Robot Intelligence International
PO Box 7890
San Diego, CA 92107

Stan Rosenschein
Teleos Research
576 Middlefield Rd
Palo Alto, CA 94301

Chuck Thorpe,
The Robotics Institute, CMU
Pittsburgh, PA 15213

Chuck Weisbin
Oak Ridge National Lab
PO Box 20008, Bldg 6025, M.S. 6364
Oak Ridge, TN 37831

Abstract

Robot navigation is a problem that encompasses most of the major areas of AI research. Machine architecture, search, knowledge acquisition and representation, planning, scheduling, reaction, perception, and of course robotics, all can play integral roles in a mobile robot navigation system. For this reason, a large part of the AI community is now interested (and has been since STRIPS was used to guide the Shakey robot) in mobile robot navigation.

1. David P. Miller

There are two pieces of AI folk wisdom that show why mobile robot navigation is an important problem for researchers to address: 1) *Simulations are doomed to succeed* 2) *Real robots don't work; at least not the way you thought they were going to.* These bits of knowledge point out the reasons why robotics is becoming more important to AI and why mobile robot navigation is becoming one of *the* problems in this area.

For years researchers in AI have worked on narrow problems in the field, and many promising solutions to these problems have been found. Work in spatial reasoning, resource allocation, task planning, sensing, sensor fusion, etc. have resulted in a variety of very interesting work. The question arises: why, with all this previous work, are there so few functioning systems that make use of these modules? I believe the answer, at least in part, is that the simplifying assumptions made about the modules associated with the module actually being designed are too simple. Implementing a module in a real robot system makes it difficult to make oversimplifications.

Robot navigation covers most of the major areas of AI research. Issues, for example, in representation that used to be discussed in AI circles in terms of abstraction, elegance, and clarity, are discussed in mobile robot circles in terms of whether or not it works.

This panel is composed of people who have made systems that work. Interestingly there still remain unresolved issues. Some of those that are addressed below are:

° What aspects of the robot's environment need to be modeled internally by the robot?

° What kind of sensory input is appropriate/vital for navigation in natural terrains?

° Should different sensor modalities be fused or diffused?

° How much execution monitoring should be preplanned?

° Is general purpose route planning any easier a target than general purpose planning?

2. Rodney A. Brooks

The field of robot navigation has in many ways become abstracted away from reality and many groups and individual researchers have spent large amounts of time solving problems which I believe to be irrelevant to real world mobile robot navigation. Actual experiments with mobile robots, such as the ALV, have shown that good performance can be achieved by using an approach which is quite different from the traditional knowledge-based AI approach. In contrast, such traditional approaches have not shown anywhere near as impressive performance in the few cases that they have been implemented on real robots.

In particular I will argue that

o Explicit complete three dimensional world models are unnecessary. Rather, only task dependent aspects of the world need be modeled, such as clear paths ahead, dynamic obstacles which could cause navigation problems, and landmarks need be explicitly extracted from the sensors.

o Different modalities from the sensorium do not necessarily need to be fused. Rather they may be fed directly to appropriate task achieving behaviors as appropriate. For instance laser scanning data might be used to correlate with an a priori map in order to do localization if that happens to be important for some particular task, whereas sonar data might be used only for local obstacle detection and never get correlated with the map.

o Execution monitoring should not be preplanned. It should be built.

o The appropriate sensorium is not dependent just on the fact that the robot is to navigate over natural terrain. Rather it is critically dependent on both the reason for needing to navigate, on the relative scale of the vehicle and the features over which it must travel, and on the method of locomotion of the vehicle. For instance, a multi-legged machine that is to crawl over a terrain consisting of densely packed boulders of the same size, and whose goal is to move in some general direction on an exploratory mission, might work well with only force feedback on its legs, feelers (or whiskers) to predict obstacles just ahead, and a compass or sun sensor for general direction.

o General purpose route planning is a red herring. What is needed is not a route plan, but an ability to get to where you want to go. Route planning is an unnatural abstraction of this task, and the Hughes ALV work has clearly demonstrated the pitfalls of this abstraction.

3. Raja Chatila

Consider navigation (reaching a given goal) in a gradually discovered environment, with various conditions for perception and motion (lighting, kind of terrain, shape and behavior of objects: fixed, movable, moving, etc.). The robot has specific intrinsic constraints (sensors, computing capacities, physical features, locomotion structure), and a task to perform.

1. Representations of the geometry, topology and some semantics of the environment are necessary. Exact model contents are sensor, robot and task dependent. Some tasks may require no model, but one is needed in order to:

a. self-locate with respect to the environment. If this is impossible, reaching a given goal (that is not in sight) will be impossible in general.

b. plan and execute a collision-free path to reach a given goal minimizing a criterion (energy, distance, time,...) and satisfying some constraints.

2. Natural terrain is essentially unstructured. Thus navigation requires data on the terrain and its nature. A depth map (e.g. provided by a 3D laser range-finder) fused with color vision is very suitable for extracting geometric features with attributes on navigability. Dead-reckoning and sensors on the locomotion structure might be vital for detecting slippage, unstable postures, etc. Sensors providing a poor but quick response (e.g. sonar) might be useful for monitoring.

3. Each type of sensor has limitations and uncertainties. Since the environment and robot state change, a single sensor will not be sufficient in general to build consistent environment models. Therefore fusion (spatial and temporal) of sensory data is necessary for navigation but not systematically, for example, for obstacle avoidance, tracking, etc.

4. Execution monitoring ensures correct task execution, reactivity to events and error recovery. Reflex loops must exist at several levels, but provision should be made to make them context and task dependent. This implies that decisional capacities, for analyzing a situation and foreseeing some events to be monitored, are necessary. Thus, execution monitoring should be preplanned, with some flexibility given to the on-line system.

5. Path planning itself is complex in general, and cannot be solved by local methods or by "trial and error". But navigation is not only geometric reasoning since the robot has a task to perform that may put various constraints on it. Feasible trajectories should be planned, using a topological and semantic knowledge on the environment, according to the capabilities of the robot's locomotion system. The problem is thus of the same nature as task planning in the general case.

4. Scott Harmon

As a general rule, a robot's configuration depends entirely upon its task. Its perception, modeling, execution monitoring and planning capabilities are driven completely by its task goals and environment. Absolutely no more capability than is needed should be incorporated.

Following this philosophy, a robot should model only as much of the environment as it needs to reliably accomplish its task. This depends upon both the specifics of the robot and the task. If the robot is relatively insensitive to the environment's complexity then little of the environment must be modeled.

A robot's sensory input depends completely upon the environment's character and the task demands. Changes in any of these factors alters the sensor information requirements. The robot should sense the environment's geometrical nature at a scale which is significant to its mobility mechanisms. Any environmental elements which affect robot mobility must be sensed in order to guarantee reliable passage. The robot may need to characterize significant nongeometric conditions as well.

Sensor fusion is primarily useful for obtaining more information about a situation than is available from independent sensors. If independent sensor measurements provide all the information at the accuracy and certainty which is adequate for planning and successful execution then no sensor fusion is needed. However, this is seldom the case for complex mobility tasks because of environmental complexities and uncertainties. Sensor fusion enriches the understanding of the task state in geometric space, in time, in modality and in certainty.

All execution monitoring must be preplanned. Effective execution monitoring is simply impossible if a priori expectations of execution results are not available. However, many other conditions must also be met to guarantee monitoring effectiveness. The important components of the execution process must be sensed. Execution monitoring is extremely useful if the execution process is nondeterministic. The particular execution process must be important to the success of accomplishing the task. If the process is not very important then don't monitor it; if it is very important then monitor it closely. Finally, the robot must be able to take some corrective action if the execution process goes awry. Monitoring is only useful if something can be done when a problem is detected. Monitoring becomes more detailed as more corrective actions are available.

Finally, truly general purpose route planning for a complex robot in a complicated environment reduces to general purpose planning so they, in the limit, have the same level of difficulty.

5. Stanley J, Rosenschein

The navigation task, like other underspecified tasks, can be made complex enough to require modeling of virtually any fact about the environment and detection of almost any imaginable condition. Navigation means traversing a continuous path that terminates where the robot needs to be. Because the reasons for taking one path rather than another (e.g., goals of seeing or being seen, not falling into elephant traps, avoiding poison ivy, etc.) are so diverse, and because conditions to be maintained along the way may be qualitative and not merely positional (e.g., maintain ability to see the road ahead), the primary research challenge is in integrating navigation and other perceptual and reasoning processes. Practical, general-purpose navigation appears to be as hard as general-purpose planning.

On the question of sensor fusion, although some have argued for autonomous, task-specific sensing modalities, we feel that, whenever practical, sensor modalities should be fused, for two fairly obvious reasons: (1) so that more conditions can be recognized, making possible more sophisticated robot control and more fluent transfer between tasks, and (2) so that the same conditions can be recognized more robustly through redundant evidence.

6. Chuck Thorpe

1. What aspects need to be modeled?

Answer: As many as possible, and as explicitly as possible. There are always models of the environment, even if no more than the distance to the nearest sensed object. Robots without models cannot accomplish interesting tasks, at least not efficiently. It is possible to build an extremely simple robot with no models, only reflexes, that does some task such as wanders a table top and turns when it encounters an edge. Such robots don't need models, computers, or even electronics; everything is done mechanically. But a robot to follow a road, or clean the floors, or perform any other really useful task, must have a map of its environment. Any intelligent coupling of perception to action, any planning at a higher level than a reflex, needs a model of the world.

In short, robots without models may make great protozoans, or even with much cleverness may produce artificial insects; but I wouldn't trust one to be a chauffeur.

2. What kind of sensors is needed for natural terrains?

Answer: Maybe not as much as we have thought. For our AMBLER project), the primary sensors will be range data and foot forces. In an environment such as Mars, the entire problem is geometry and on surface strength.

I think a much richer set of sensing (sensors and processing) may be needed for travel in semi-natural terrains, where it is important to recognize flowers and roads, and to properly interpret stop signs and moving cars and pedestrians and police cars and ...

3. Fusion or fission?

Answer: Both, of course. There is interesting evidence of two different biological perception mechanisms, one for locomotion and one for recognition. In practical robot vision systems, there may be several different activities going on in parallel. Each of these activities may need information from several sensors, such as Anita Flynn's work on fusing sonar and IR range data to build maps. So at least some of the data needs to be spread over at least two different processes.

4. Preplan execution monitoring?

Answer: Definite maybe. I greatly prefer to keep my planners so simple that they can be rerun at almost every step, or at least any time something goes wrong with the original plan. So rather than monitor execution, I prefer to update the world map / model, and replan. In the robots I've known and loved, perception has always been much slower than planning (and much less reliable), so the extra planning wins, especially if it can produce a plan which does less sensing.

5. Route planning vs. planning.

Answer: Route planning is much easier, except in really contrived cases. It is of course possible to fold any problem into route planning, e.g. "plan the route that would be necessary to bake a cake" could incorporate all the steps in cake-baking-planning. It is of course also obvious that optimal route planning can be tough, e.g. traveling salesman. But practical path planners, looking for "sufficing" rather than optimal paths, can be much simpler than full-blown planners.

7. C.R. Weisbin

I.The aspects of the robot's environment which need to be modeled seem to me to depend critically on the task (e.g., navigation in a known world, exploration and goal finding, ability to deal with unexpected events, etc.). We would provide an "a priori" model appropriate for the task domain, and confirm that model dynamically based upon sensing feedback.

2. Sensory input for navigation in natural terrain seems again to depend on the task and the available processing time. Sensors for autonomous navigation of an automobile moving at 60mph, could be different than sensors required for delicately defusing a bomb.

3. If by fused or diffused, you mean that independent sensor information is consistently combined (fusion) versus independent sensor data used for different task functions, we would tend to support multisensory fusion where time and resources were available, and independent sensor readings for different task functions where speed or cost effectiveness dominated.

4.1 do not understand how execution monitoring can be anything but preplanned. Do you not need to anticipate the types of events to monitor before implementation of monitoring?

5. General purpose route planning can be envisaged and extended to cover many, if not all, of the most general purpose planning paradigms; but I do not think this question (as it is posed) is very important.