Position Estimation and Autonomous Travel by Mobile Robots in Natural Terrain

Eric Krotkov

Robotics Institute
Carnegie Mellon University

1. Introduction

In this chapter, we will address two senses of the term "mobile robot navigation:" (1) finding a position or a course; and (2) travelling from place to place. In simplistic language, these two meanings correspond to answers to the questions "Where am I?" and "How do I get there?"

Although the robotics community has devoted significant effort to unattended air and water vehicles, in this chapter we restrict attention to ground robots. One reason for this restriction is that the majority of research and development efforts have concentrated on land vehicles, so we have learned more about ground-travelling robots than about flying, gliding, swimming, or floating robots. Another reason for the restriction is that, in general, the density of obstacle fields is greater on the ground than in either air or water, and as a consequence, mobile robots on the ground face the more challenging navigation problems.

Mobile robotics is a young field, having started, by most accounts, in the late 1960s with work on Shakey at SRI and the Cart at Stanford. In the three decades since the field's inception, an estimated 100 to 1,000 mobile robots have been built and operated. In the process, many lessons have been learned and many problems have been solved, enabling mobile robots to achieve such feats as driving 140 km/h on the Autobahn, rappelling into an active volcano, and driving on highways from Pittsburgh to San Diego.

Despite these successes, no set of guiding principles for mobile robotics has been widely articulated or accepted. Particular designers and practitioners have developed effective approaches and distinguished dogma, and perhaps these will one day evolve into principles. However, the Handbook of Mobile Robot Development has not yet been written.

Given the immature state of the field, we have chosen for the purposes of this chapter to approach the two senses of the term "mobile robot navigation" through case studies. These case studies, distilled from the author's research, will illustrate how selected approaches have been developed and how they perform. We do not argue that the selected approaches are optimal or even superior to competing approaches. Instead, we aim first to articulate the approaches, and then to investigate whether they exhibit characteristics of being organized by coherent principles.

The case studies have been conducted as part of the Lunar Rover Initiative, whose goal is to launch and land a lunar mission within the decade. The initiative has targeted a mission to land two rovers on the surface of the Moon, where they will navigate semi-autonomously for 1,000 km, visiting historic sites and places of scientific interest (Krotkov 1994). One of its goals is to provide the general public an opportunity to drive a robot on the Moon. For additional information, see http://www.cs.cmu.edu/~epk.

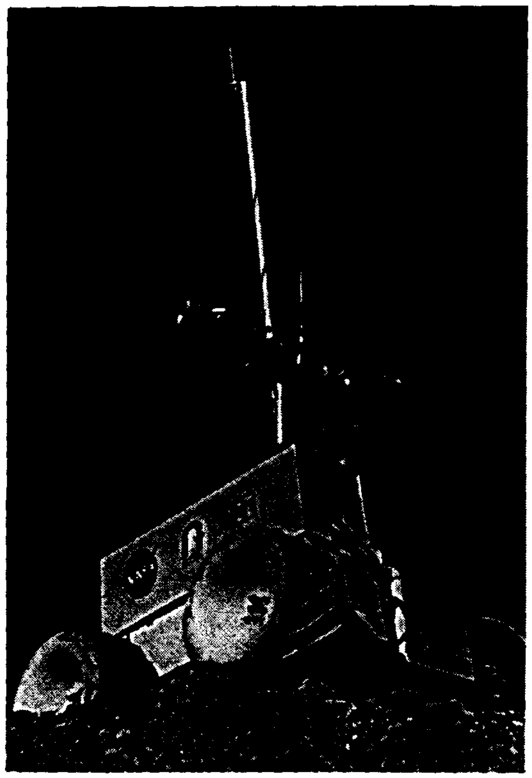

Figure 1 Ratler

This goal is ambitious, because remotely driving a lunar rover presents a number of technical challenges. One major challenge is position estimation, since there is no Global Positioning System (GPS) satellite constellation orbiting the Moon. Another major challenge is time-delay. The largest distance between Earth and the Moon is 407,000 kilometers, leading to a communication delay of about 2.5 seconds in each direction (1.4 seconds between Earth and the Moon, and an estimated 1.1 second communication delay on Earth). Thus, if the rover is about to hit an obstacle on the Moon, the teleoperator on Earth notices this situation 2.5 seconds later, and (assuming instant response) their evasive motion commands reach the rover another 2.5 seconds after that -- five seconds after hitting the obstacle, potentially crippling the rover and ending the mission.

In order to develop and validate the technology required to meet these two, and other, challenges, we have conducted increasingly challenging field trials in terrestrial settings. For these, we have used Ratler as the mobile vehicle (Purvis & Klarer 1992). This robot (Figure 1) was built by Sandia National Labs, and outfitted with sensors and control systems at Carnegie Mellon University (CMU). Ratler is about the size of a tractor mower (1.2 m by 1.2 m by 1.5 m). Its body is divided into two halves that rotate against each other to increase the robot's ability to traverse rugged terrain.

We have organized the chapter as follows. In Section 2, we identify several position estimation approaches common in the mobile robotics community, and describe a working multi-sensor position estimation system implemented by the author. In Section 3, we enumerate selected highlights from the last decade of natural terrain mobile robot research, and present an operational integrated guidance, navigation, and control system implemented by the author (along with a cast of thousands) on a lunar rover prototype. We conclude in Section 4 with perspectives on the future of mobile robot navigation, and reflections on the multiple meanings of the term "navigation."

2. Position Estimation in Natural Terrain

Position estimation has received attention in several domains, including marine, terrestrial, aeronautic, and space navigation. There are two essentially different approaches, which the following paragraphs will briefly review. Interested readers will find a comprehensive survey elsewhere (Borenstein 1996).

The first approach, *relative position estimation,* integrates measurements that are internal to the moving equipment, producing a position measurement relative to a previous position. This approach includes the method of deduced reckoning, more commonly called dead reckoning. Dead reckoning has been used on virtually every mobile robot for ground travel, whether wheeled, legged, or tracked. In the case of wheels, the underlying approach, called odometry, counts revolutions of the wheels, and uses knowledge of the wheel radius to calculate the distance travelled. In order to infer change in elevation, it is necessary to employ some kind of tilt sensor (for example, an inclinometer) in addition to odometry. A variation on the theme of odometry double integrates output from accelerometers to infer distance travelled, and determines attitude by integrating output from rate gyroscopes. The fundamental weakness of dead reckoning, however it is implemented, is that the error grows without bound over time.

The second approach, *absolute position estimation,* makes measurements of features that are external to the moving equipment, producing a position measurement relative to an external reference system. The advantage of absolute estimators is that their error does not grow over time, in contrast to the relative estimators. The disadvantages include considerable complexity, and a greater likelihood of failure. Two absolute position estimators have been implemented on many mobile robots operating in natural terrain.

The most popular absolute position estimators exploit the Global Positioning System (GPS), provided by a constellation of satellites in low Earth orbit. There are numerous advantages to this approach, for it does indeed solve a large class of position estimation problems at a moderate cost. However, the approach suffers from several disadvantages. First, unless operated with a differential receiver located away from the moving equipment, the accuracy and precision are not sufficient for many purposes. The requirement for a second receiver restricts the possible operating range of the robot system, which must retain a line of communication to the differential station. Second, GPS performs poorly in urban, jungle, and mountainous environments due to multipath effects, and completely fails indoors and underground where the satellite signal waveforms cannot be detected reliably.

Another absolute position estimator solves the landmark navigation problem of finding position and orientation with respect to landmarks, that is, easily distinguishable features in the environment. One version of this is called the Backpacker's Problem: Given a topographic map, and a feature-rich image taken from some location within the mapped area, determine the position at which the image was taken. Another version of this is well known to mariners: Given an almanac enumerating the position of the stars, and given a sextant, determine the current viewing position.

We now turn to case studies of one relative (Section 2.1) and one absolute (Section 2.2) estimators.

2.1 Dead Reckoning Case Study

Eight different sensor measurements are available inside Ratler: four encoders (one for each wheel), three inclinometers (roll, left pitch, and right pitch), and a compass. We have presented one approach to using these measurements for dead reckoning (Krotkov 1995). We have also presented an alternative approach using an inertial measurement unit (Fuke 1996). In this section, we present an evolution of the first approach, and the results of a rigorous test program.

The signals from the three inclinometers are quite stable and noiseless; on the other hand, the compass signal presents high levels of noise during motion. Further tests with the compass and similar devices from other robots available in our group pointed to the same general conclusion: the existence of large metallic structures causes sudden and unexpected changes in orientation measurements. Since we conducted our tests in the vicinity of large buildings, we were forced to abandon the compass altogether, and rely solely on encoder and inclinometer integration.

Take (x, y, z) for position, θ for yaw (orientation) and ϕ for average pitch from the filtered inclinometer signals. We computed the coordinates with respect to a global coordinate system:

$$x_{n+1} = x_n + \frac{e_r + e_l}{2}\cos\theta_n\cos\phi$$

$$y_{n+1} = y_n + \frac{e_r + e_l}{2}\sin\theta_n\sin\phi$$

$$z_{n+1} = z_n + \frac{e_r + e_l}{2}\sin\phi$$

$$\theta_{n+1} = \theta_n + \frac{e_r - e_l}{B}$$

where B is the wheelbase, and e_r and e_l are respectively the encoder displacements expressed in metric units.

To quantify the performance of this approach, we built a test site in a 10 x 20 meter area with rough but uniform terrain, in which we drew a grid with intervals of 0.25 meter. We attached a piece of chalk to the center of Ratler, so that motion of the robot produced a trail of chalk on the ground.

We commanded Ratler to follow 10 different arcs, varying from a sharp right turn (radius of 3 meters) to a sharp left turn (same radius). We repeated each arc 3 times. We repeated straight motions approximately 15 times. Misalignment of one of the wheels to the right by a mere 2 degrees caused the robot to drift 0.6 meter every 20 meters (3% error). Skid steering contributed an additional 0.75 meter drift (4% error), on average, for commanded linear trajectories. We compensated these errors by introducing correction factors in the encoder readings. This substantially reduced the mean error. For straight lines, average error fell to 0.3 meter for every 20 meter trajectory (1.5% error). However, for abrupt turns performance is poor; for instance, a 3 meter radius turn leads to 10% mean error. This poor performance during turns appears to be an inevitable consequence of the slippage caused by skid steering.

2.2 Visual Landmark Navigation Case Study

A mobile robot performing a long traverse on the Moon (or other celestial bodies) poses a great challenge for position estimation technologies. The GPS infrastructure is not available, and the

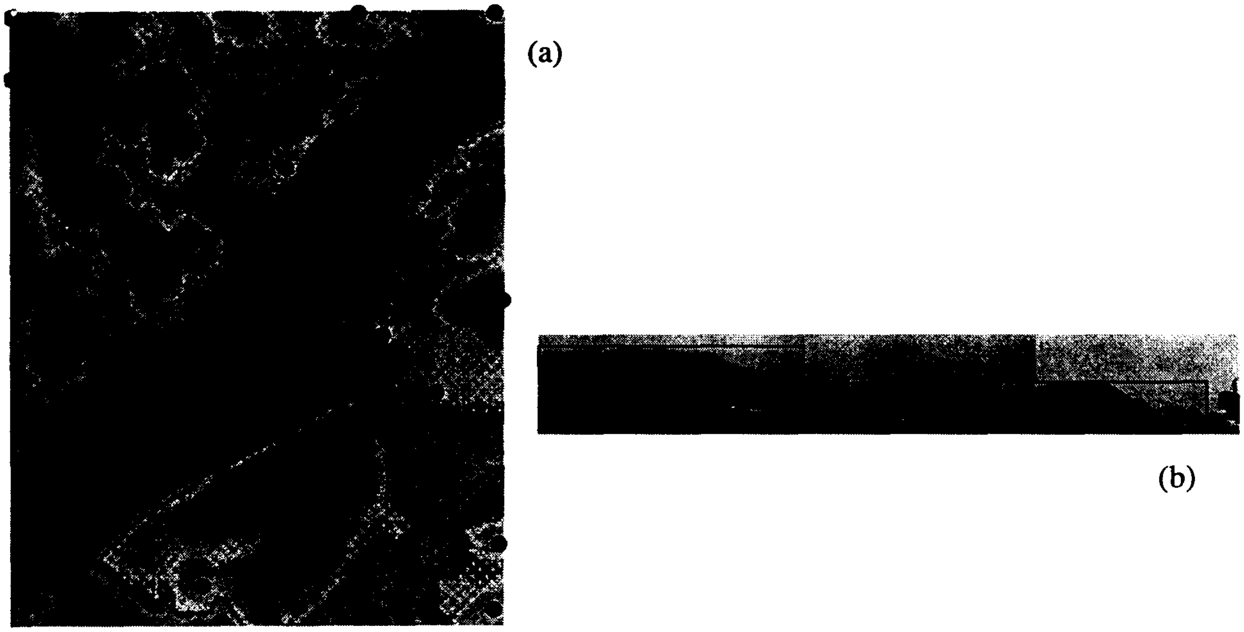

(a)

(b)

Figure 2 Hills in Pittsburgh: (a) Preprocessed digital elevation map, (b) Skyline features

magnetic field is insufficient for compass measurements. On celestial bodies, attitude measurements can be generated accurately from a star sensor (Wertz 1978), since the visibility of stars from bodies without atmosphere is excellent. In this situation, relative position and absolute orientation are available, but absolute position is not.

We have developed a vision-based method for self-localization of mobile robots travelling in outdoor terrain. The method is more accurate than any competing approach, by virtue of a probabilistic foundation.

The method assumes that the terrain has appreciable three-dimensional relief, and that a digital elevation map of the terrain exists. The approach will be valuable in environments where the Global Position System does not function, for example, on the Moon or planets.

2.2.1 Map Processing

Given digital elevation maps, such as those provided by the U.S. Geological Survey, we extract features likely to be identifiable in rover-acquired images. The processing searches for local elevation maxima in a 13x13 window, labelling them as topographic peaks. Figure 2a illustrates a map of the Pittsburgh area with peak features marked as black circles (the gray circle near the river marks the robot position).

2.2.2 Image Understanding

Our automatic mountain detection algorithm follows two steps: it first detects the skyline in a given image, and then searches that skyline for mountain peaks.

For skyline detection, the algorithm first segments the image into sky and ground regions using standard region growing operators. The algorithm then describes the sky/ground border as a set of line segments, using a chi-square test to terminate lines, and using a test based on the Akaike Information Criterion to merge lines.

Figure 3 Dromedary Peak, Utah

Figure 4 North and South Massifs, Apollo 17 Landing Site

For mountain peak detection, the algorithm applies an A*-type search to the skyline, looking for nearby lines that could form the image of a mountain. Figure 2b illustrates the results of the peak detection process, in this case, two broad peaks and one relatively narrow peak.

2.2.3 Position Estimation

Our three-step algorithm identifies correspondences between features extracted from maps and images. First, off-line, it creates a table containing all the peaks that are visible from every possible position. Second, it visits each possible position, and compares the peaks extracted by the image understanding routines with the stored peaks. The comparison is based on a probabilistic figure of merit, which is the product of a prior distribution and an error distribution derived from the image understanding process. Third, for the positions with significant posterior probability, it verifies that the actual image agrees with a synthetic image of what the scene should look like from that position.

2.2.4 Results

The skyline navigation approach has been tested on three data sets. For the Pittsburgh data set, whose peaks lie several kilometers away, the system estimated position with an error of 87 meters, so small that it lies inside the red dot in Figure 2a. In comparison, errors ranging from several hundred meters to 5,000 meters have been reported in previous research with real data (Sutherland 1993). The difference is in the scale of our features. If we constrain ourselves to a small number of very salient features, huge errors ensue. But if we allow many map features to be considered, we must impose some quantitative structure in order to manage the large number of competing interpretations that arise. The likelihood calculations in our algorithms provide this needed structure.

For the Utah data set (Figure 3), whose peaks lie tens of kilometers away, the error was 95 meters, which compares favorably to the error of 71,700 square meters obtained by other systems (Thompson 1993). For the Apollo 17 data set covering the North and South Massifs (Figure 4), whose peaks lie tens to hundreds of kilometers away, the error was 180 meters. Again, we benefit greatly from our reliance on all available features, large and small, present in the image.

2.3 Discussion

In this section, we presented two case studies of position estimators representative of the principles that mobile robots use to estimate their position when GPS not available or not sufficient. The first—dead reckoning—employs computationally simple sensor processing to estimate relative position, but the quality of the position estimate declines over time and with the roughness of the terrain traversed. The second—visual landmark navigation—uses computationally intensive sensor interpretation to estimate absolute position, but the quality of the position estimate is independent of time and the terrain traversed.

Because no single, general, powerful method has been invented, mobile robot developers typically combine two methods, one from each of the above categories.

3. Autonomous Travel over Natural Terrain

A significant number of systems have been fielded that are capable of self-reliant operation in outdoor, natural terrain. These include land vehicles such as the Navlab (Hebert 1997) and air vehicles such as the cruise missile. However, these vehicles operate under constraints (on mass, power, volume, etc) far different than those faced by planetary rovers.

Restricting attention to planetary rover class mobile robots, we know of only six systems other than our own that have achieved rough terrain traverses approaching or equalling 100 meters under control modes other than teleoperation.

- Robby is a six-wheeled rover developed at the Jet Propulsion Laboratory (JPL). Robby used stereo and a semi-autonomous control mode to traverse 100 m in a desert arroyo. For this traverse, operators specified navigational waypoints, between which the vehicle traveled autonomously.
- Rocky is a family of six-wheeled rovers developed at JPL. The Rocky rovers use a control architecture that creates task-oriented modules defined as "behaviors" and provides for programmable inter-module connections (Gat 1994). Rocky rovers have performed a number of traverses in outdoor terrain using short-range sensors such as laser light stripers. The Rocky concept served as the baseline design for the Mars Pathfinder mission, launched in 1996, which calls for travel tens of meters by the Sojourner rover (Stone 1993).
- Adam is a six-wheeled rover developed as part of the Iares project in Europe, whose goal is to build a ground demonstrator with capabilities to perform an ambitious (1,000 km traverse over 13 months) scientific mission to Mars by the end of the millennium (Giralt 1994). The approach has been tested in an experimental testbed called Eden (Chatila 1996). In Eden, Adam used vision for landmark recognition, a laser rangefinder for obstacle detection, a coarse-scale planner for sub-goal selection, and 2D and 3D motion planners for obstacle avoidance. Adam executed autonomously a "Go To(Landmark)" task, where the landmark is a known artificial object, in an unknown environment that is gradually discovered by the robot. Although the degree of autonomy is high, the system has yet to be tested in extreme terrain for long durations.
- Ambler is a six-legged walker developed at Carnegie Mellon University. Under autonomous control, Ambler traversed over 100 m in rugged terrain including meter-scale obstacles, and over 500 m in benign terrain including 15 degree slopes and cross-slopes

(Krotkov & Simmons 1995). The navigation system relied on accurate terrain maps built from laser rangefinder data.

- Dante II is an eight-legged walking and rappelling robot developed at Carnegie Mellon University (Krotkov 1994). Dante II descended 400 m over 35 degree slopes into the crater floor of Mt. Spurr, an active volcano in Alaska, where it successfully collected gas samples. The navigation system employed a number of control modes, ranging from direct teleoperation to full autonomy, using cameras for visual feedback and a laser rangefinder for terrain mapping.
- The Surrogate Semi-Autonomous Vehicle is a modified HMMWV developed at Lockheed-Martin under the Unmanned Ground Vehicle program. This vehicle traversed hundreds of meters of cross-country terrain using stereo vision for obstacle avoidance at a speed of 2 meters/sec (Hebert 1997). In a previous effort, the CARD-II system demonstrated a version of safeguarded teleoperation (Lescoe 1991).

The accomplishments of these navigation systems are remarkable and historic. One distinguishing feature of the present work is greater travel distance (43 km, to date). Another distinguishing feature is that the other systems, except Robby, rely on laser rangefinders or proximity sensors rather than on stereo vision. The main difference with the Robby system is our utilization of general-purpose computing hardware.

3.1 Case Study

For a lunar rover mission like the one described in Section 1, a spectrum of navigation modes are applicable. At one end of the spectrum lies pure teleoperation. Although this is feasible, as the Soviets showed in the 1960s with the Lunokhod program (Newman 1962, Gatland 1978), the time delay proves troublesome and contributes to significant stress and fatigue for operators. At the other end of the spectrum lies pure autonomy. Although research in autonomous vehicles has progressed dramatically, it has not yet produced techniques for safe, reliable operation of rovers using limited computing and sensing resources.

Our approach, called safeguarded teleoperation, occupies a position in the center of the spectrum. In this approach, the user specifies high-level goals such as desired direction of travel, and the vehicle autonomously decides how to execute the command in a way that optimizes a performance criterion. The principal virtue of the approach is that it enables safe and reliable operation, maintaining system integrity by avoiding terrain hazards and faulty uplink commands of the type that crippled the Phobos lander. The safeguarded teleoperation approach has been described in detail elsewhere (Krotkov, 1995).

3.2 Safeguarded Teleoperation

The safeguarding system consists of four primary processes: stereo vision, laser scanning, obstacle avoidance, and command arbitration. The safeguarding system runs entirely on-board the robot on a 586 processor. Figure 5 presents a block diagram of the overall navigation software architecture. The basic data flow is that the stereo process produces terrain elevation maps which are passed to the obstacle avoidance planner, which uses them to evaluate the efficacy of traveling along different paths. In safeguarded teleoperation mode, these recommendations are combined with the desires of a human operator to choose the best path to traverse (in autonomous mode, there is typically no

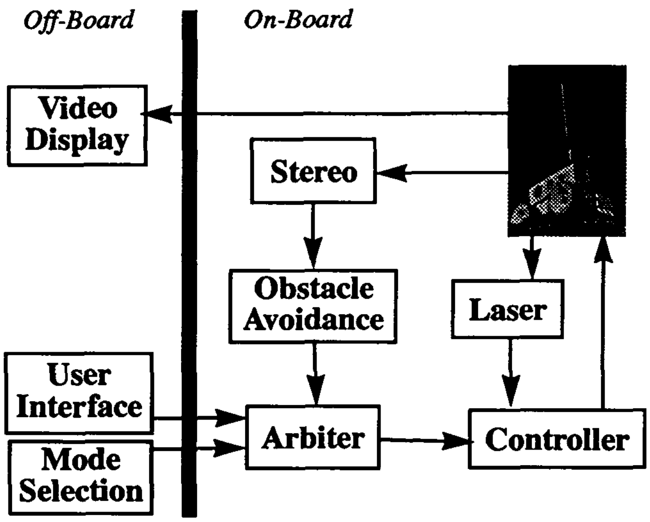

Figure 5 Safeguarded teleoperation system architecture

operator input). The command arbiter then forwards steering and velocity commands to on-board controller, which executes the commands and returns status and sensor information. All components operate concurrently and receive their inputs asynchronously.

This arbitration architecture facilitates combining the best features of human and machine control: The human operator provides high-level guidance as to where the rover should travel (minimizing the need for global path planning by the rover), and the rover provides safeguarding by detecting and avoiding imminent hazards (which increases overall safety, and can reduce the level of fatigue and frustration typically experienced by operators in time-delayed, remote driving tasks).

3.3 Stereo Perception

The stereo perception system consists of a stereo module that derives terrain information from camera images. The hardware consists of two CCD cameras, with auto-iris 8 mm lenses mounted on a motion-averaging mast, and a three-channel frame grabber.

The software takes as input a stereo pair and outputs arrays of the three coordinates X, Y, and Z of the image pixels in the camera frame of reference. Prior to stereo matching, the images are rectified in order to align their scanlines. The stereo system uses area-based correlation to establish correspondences between left and right images.

The implemented four-camera system features an 80 degree field of view, and processes terrain data 4 to 8 meters in front of the vehicle at approximately 1 Hz, with a typical ground resolution of 8 cm.

3.4 Laser-Based Hazard Detection

To complement the stereo system and provide a redundant sensor, we have developed a laser-based hazard detection system. We point the laser at the area some 75 to 125 cm in front of the vehicle, so that it can detect troublesome depressions such as craters. In addition to responding to terrain depressions, the laser system contributes to the overall safety by sensing areas not covered by stereo vision when performing a point turn, and by detecting small objects missed due to the coarse resolution of the vision system.

The hazard detection system consists of a data acquisition module and a hazard detection module. The acquisition module retrieves data from an Acuity 3000-LIR laser ranger operating in the near-infrared band, and validates the range data by performing integrity checks of the laser system and correcting minor problems if necessary. The detection module evaluates the elevation profile of the terrain looking for evidence of three different hazards: step, ditch, and belly.

The implemented system provides data with ground resolution better than 1 cm at rates in excess of 4 Hz.

3.5 Obstacle Avoidance Planner

To decide where it is safe to drive, we have developed an obstacle avoidance planner. The basic idea is to evaluate the hazards along a discrete number of paths (corresponding to a set of steering commands) that the rover could possibly follow in the next few seconds of travel. The evaluation produces a set of "votes" for each path/steering angle, including "vetoes" for paths that are deemed too hazardous to traverse. In this way, the rover steers itself away from obstacles, such as craters or mounds, that it cannot cross or surmount.

To evaluate the potential steering commands, the planner uses a detailed model of the vehicle's kinematics and dynamics to project the expected path of the rover forward in time over the terrain. This produces a set of paths, one for each potential steering direction. By examining the terrain map elevations beneath the wheels at each point along a path, the planner can estimate the expected rolls and pitches along the path. If the maximum (absolute) roll or pitch of a body segment exceeds an allowable threshold, or if the path crosses a significant area that has not yet been imaged by the cameras, then the path is vetoed. Otherwise, the "value" of that path is a weighted sum of the roll angle, the pitch angles, and the percentage of underlying terrain that is known.

The local obstacle avoidance algorithm dynamically chooses which portion of the image should be processed, based on the current vehicle speed, stopping distance, and expected cycle time of the perception/planning/control loop. Typically, only about 2% of the image is used for planning purposes.

3.6 Command Arbiter

The command arbiter asynchronously accepts path evaluations from various sources and chooses the best steering angle based on those evaluations. Each path evaluation consists of a steering angle, value, and speed. If the value is "veto" then that steering angle is eliminated from consideration. Otherwise, recommendations for that steering angle from all sources are combined using a weighted sum. This produces a new set of evaluations that combines all the sources of input.

The arbiter then finds, within this combined set of evaluations, the largest contiguous set of steering angles whose values are all within 90% of the maximum value, and chooses the midpoint of that set as the commanded steering angle. The speed chosen is the minimum speed recommended by all the sources. The idea is to prefer wide, easily traversable areas over directions that might be a bit more traversable, but have less leeway for error if the rover fails to track the path precisely.

The path evaluations are also tagged with a robot pose. If the tagged pose differs significantly from the rover's current dead-reckoned pose, then those path evaluations are invalidated. If the

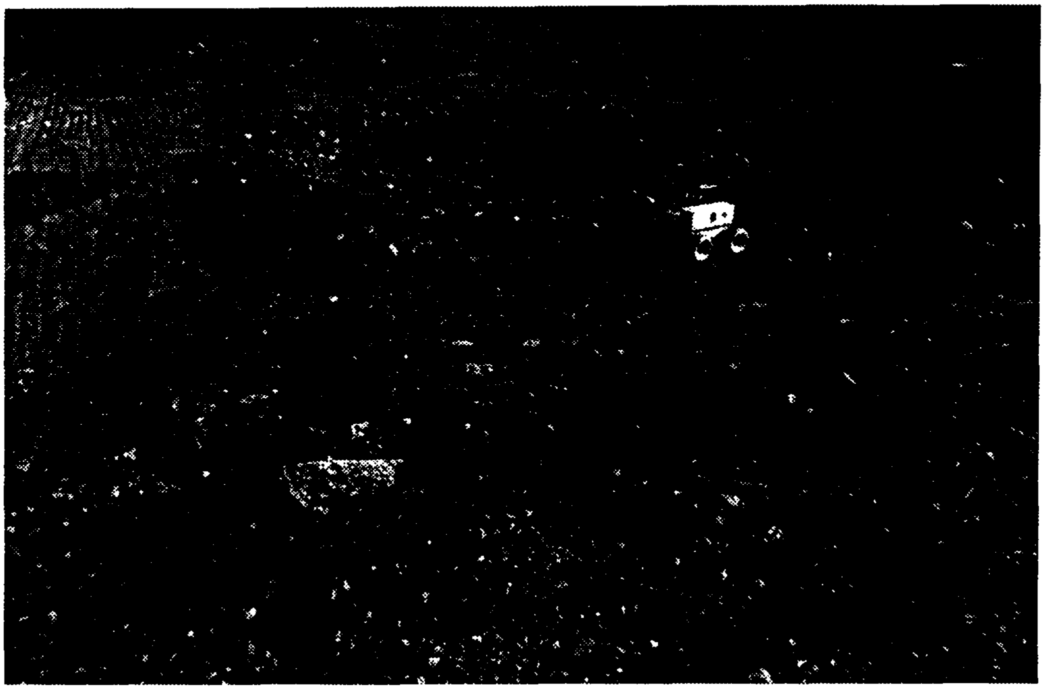

Figure 6 Ratler traversing the Moonyard

evaluations from all the sources are invalidated, then the arbiter issues a command to halt the rover. In this way, the arbiter safeguards against other modules crashing, or otherwise failing to provide timely inputs. Similarly, we guard against failure of the arbiter itself by having it send commands of the form "steer in this direction for X meters." The controller stops the robot if it does not receive a new command before that distance has been traversed.

3.7 Safeguarded Teleoperation Field Trials

We designed a test area called the MoonYard (Figure 6), a 2.5 acre site in Pittsburgh. The MoonYard is a lunar analogue containing crater patterns and topographic relief similar to the Apollo 11 landing site, the Apollo 17 landing site, and Mare Tranquilitatis. The MoonYard features 12 craters; the largest is 11 meters in diameter and 1.5 meters deep.

We tested the rover by commanding it to traverse the MoonYard hundreds of times, covering 12.5 km under a wide variety of conditions ranging in lighting from dawn to dusk, in weather from rain to shine, in wind from still to violent, and in ground moisture from baked dry to deep puddles. The safeguarding system proved effective in maneuvering around discrete obstacles, avoiding craters, and steep slopes. We also tested the safeguarded teleoperation system indoors during rainstorms, and outdoors at a Pittsburgh slag heap. The system performed well at all these sites, traversing a total of 43 km. Although the typical travel speed under safeguarded teleoperation control was 50 cm/sec, the rover successfully travelled at its maximum speed of 70 cm/sec while safeguarding.

During the field trials, the operator noted each occasion at which the safeguarding system did not adequately protect the rover from danger. Examples of failure include colliding with a rock while turning (which occurred numerous times) and colliding with a rock head-on (which occurred several times). We observed 2.4 failures per kilometer (16 failures over 6.7 km) using stereo-based safeguarding alone, and 0.8 failures per kilometer (36 failures over 42.7 km) using multi-sensor safeguarding with stereo and laser. These results indicate that the multi-sensor approach is 3 times more effective than the stereo-based approach. It is worth noting that the vehicle never tipped over, despite numerous opportunities.

3.8 Discussion

In this section, we have presented a case study of a system representative of the general principles that mobile robots use to travel autonomously in natural terrain. Three principles emerge from the case study as deserving specific mention:

- Acquire three-dimensional information about the environment. In the case study, stereo vision and a laser rangefinder supply this information; many other approaches are possible and desirable.
- Employ multiple sensing systems in order to simplify the processing that each must do. In the case study, the stereo process could execute more rapidly by virtue of relying on the laser system to recover from mistakes committed by the stereo system.
- Exploit multiple sensing systems in order to provide redundancy and system reliability. In the case study, the robot could continue driving (albeit more slowly) when the laser system or the stereo system were not operational. This redundancy enables graceful degradation of performance rather than catastrophic failure.

4. Discussion

In this chapter, we have addressed two senses of the term "mobile robot navigation:" (1) finding a position or a course; and (2) travelling from place to place.

As summarized in Section 2.3, mobile robot researchers rely on at least two principles, deduced reckoning and landmark navigation, to enable robotic position estimation in natural terrain. Many schemes have been developed to exploit those principles, and each scheme admits numerous variations. In practice, no single scheme or combination of schemes has yet been developed to estimate position with satisfactory accuracy or precision under a wide variety of natural terrain conditions.

As summarized in Section 3.8, mobile robot researchers rely on at least three guiding principles to enable autonomous travel by mobile robots in natural terrain: avoid collisions using three-dimensional perception, decrease cycle time by using multiple complementary sensors, and increase system reliability by using multiple sensor systems. In practice, no single approach to exploiting these principles has been recognized as superior, and prospects for the emergence of such an approach appear uncertain for the near future.

Presently, the vast majority of mobile robots operating in natural terrain do not exhibit position estimation or autonomous travel capabilities that rival those of humans, primates, or even many insects. However, research has succeeded in creating new mobile robotic capabilities, and there appears to be cause for optimism that this progress will continue. Further, it appears that roboticists an accelerate such advance by taking inspiration from the navigational principles that have evolved in the animal kingdom.

Acknowledgments

We gratefully acknowledge the assistance of Ben R. Brown, H. Ben Brown, Lonnie Chrisman, Cassandra Coombs, Fabio Cozman, Keith Duwel, Marie Elm, Richard Goodwin, Martial Hebert, Paul Klarer, James Martin, Patrick Muir, Mesta Electronics, David Pahnos, Yoshiakazu Shinoda, Gregory Whelan, and Red Whittaker. This work was supported in part by NASA Grants NAGW-1175 and NAGW-3863 from the NASA Telerobotics Program.

References

Angle, C., & Brooks, R. (1990). Small planetary rovers. In Proc. IEEE Intl. Workshop on Intelligent Robots and Systems, Tsuchiura, Japan, 383-388.

Borenstein, J., Everett, H. R., & Feng, L. (1996). Navigating Mobile Robots: Systems and Techniques.

Chatila, R., Lacroix, S., Betge-Brezetz, S., Devy, M., & Simeon, T. (1996). Autonomous mobile robot navigation for planet exploration: The EDEN project. In Proc. Workshop on Planetary Rover Technology and Systems, IEEE Intl. Conf. Robotics and Automation, Minneapolis, MN.

Cozman, F., & Krotkov, E. (1995). Mobile robot localization using a computer vision sextant. In Proc. IEEE Intl. Conf. Robotics and Automation, Nagoya, Japan, 106- 111.

Cozman, F., & Krotkov, E. (1996). Position estimation from outdoor visual landmarks for teleoperation of lunar rovers. In Proc. IEEE Workshop on Applications of Computer Vision, Sarasota, Florida, 156-161.

Fuke, Y., & Krotkov, E. (1996). Dead reckoning for a lunar rover on uneven terrain. In Proc. IEEE Intl. Conf. Robotics and Automation, Minneapolis, MN, 411-416.

Gat, E., Desai, R., Ivlev, R., Loch, J., & Miller, D. (1994). Behavior control for robotic exploration of planetary surfaces. IEEE Trans. Robotics and Automation, 10(4). 490-503.

Gatland, K. (1978). Robot Explorers: The Encyclopedia of Spaceflight.

Giralt, G. (1994). Technical thresholds between robotic, man-tended and permanently manned exploration and exploitation of the Moon. In Proc. ESA Intl. Lunar Workshop, Beatenberg, Switzerland, 103-121.

Hebert, M., Stentz, A., & Thorpe, C., editors (1997). Intelligent Unmanned Ground Vehicles: Autonomous Navigation Research at Carnegie Mellon.

Krotkov, E., Bares, J., Katragadda, L., Simmons, R., & Whittaker, R. (1994). Lunar rover technology demonstrations with Dante and Ratler. In Proc. Intl. Symposium on Artificial Intelligence, Robotics and Automation for Space, Jet Propulsion Laboratory, Pasadena, California.

Krotkov, E., Hebert, M., & Simmons, R. (1995). Stereo perception and dead reckoning for a prototype lunar rover. Autonomous Robots, 2(4), 313-331.

Krotkov, E., & Simmons, R. (1996). Perception, planning, and control for autonomous walking with the Ambler planetary rover. Intl. J. Robotics Research, 15(2), 155-180.

Lescoe, P., Lavery, D., & Bedard, R. (1991). Navigation of military and space unmanned ground vehicles in unstructured terrains. In Proc. Conf. Military Robotic Applications: Military Robotic Vehicles.

Newman, R. (1962). Time lag considerations in operator control of lunar vehicles from Earth. In Cummings, C., & Lawrence, H., editors, Technology of Lunar Exploration.

Purvis, J., & Klarer, P. (1992). Ratler: Robotic all terrain lunar exploration rover. In Proc. Sixth Annual Space Operations, Applications and Research Symposium, Johnson Space Center, Houston, TX.

Stone, H. (1993). Design and control of the MESUR/Pathfinder microrover. In Proc. Intl. Conf. Advanced Robotics, Tokyo, Japan, 263-270.

Sutherland, K., & Thompson, W. (1994). Localizing in unstructured environments: Dealing with the errors. IEEE Trans. on Robotics and Automation, 10(6), 740-754.

Thompson, W.B., Henderson, T.C., Colvin, T.L., Dick, L.B., & Valiquette, C.M. (1993). Vision-Based Localization. Image Understanding Workshop, 491-498.

Volpe, R., Balaram, J., Ohm, T., & Ivlev, R. (1996). The Rocky 7 Mars rover prototype. In Proc. Workshop on Planetary Rover Technology and Systems, IEEE Intl. Conf. Robotics and Automation, Minneapolis, MN.

Wertz, J.R. (1978). Spacecraft Attitude Determination and Control. Astrophysics and space science library, 73.