Department of Electrical and Computer Systems Engineering

Technical Report MECSE-17-2004

A Review of Outdoor Robotics Research

D. J. Spero

MONASH UNIVERSITY

A Review of Outdoor Robotics Research

Dorian J. Spero
Intelligent Robotics Research Centre
Monash University,
Clayton, Victoria 3800, Australia
Email: dorian@ieee.org

Abstract

This report provides an introductory review of the field of mobile robot navigation, with a domanial bias towards unstructured outdoor environments. The navigational process is reviewed in terms of the following fundamental elements: locomotion mechanism, control system, sensing, environmental mapping, path planning, and localisation. These elements are also described using several representative examples of experimental robotic systems, which have been developed for various complex tasks in outdoor environments. Some of these tasks include cross-country and interplanetary exploration, search and rescue, law enforcement, and military operations.

It is important to note that this field is still in its technological infancy. Compared to a robot's potential applicability, somewhat borne out of our own imagination, there exists a substantial amount of research yet to be done and, hence, there is ample room for aspiring roboticists to join the research efforts. Currently, there are numerous open problems that need to be solved before mobile robots can be released into the real-world for affordable, flexible and reliable service. This report discusses a few of the significant problems of late, and provides several avenues for future research.

Contents

1 Introduction

The field of robotics is a relatively new area, however, the original concept of automation can be traced back to ancient civilisations [Malone, 1978]. The Greek philosopher Aristotle (384-322 B.C.) once wrote:

> *If every instrument could accomplish its own work, obeying or anticipating the will of others ... if the shuttle could weave, and the pick touch the lyre, without a hand to guide them, chief workmen would not need servants, nor masters slaves**.

Aristotle's notion of having autonomous instruments that are able to perform complex tasks will soon be a reality. A substantial research effort is currently underway to develop intelligent† robots, which are machines that are capable of performing complex tasks, in a variety of settings, by exploiting perception, reasoning, and control. These robots, which are now in their third generation (see [Jarvis, 1992a] for the phases of robot evolution), have a close correlation to living organisms and so their potential uses are vast and varied.

There are two types of robots that have emerged: *robot manipulators* and *mobile robots*. The former involves manipulating objects on tables or conveyor belts from a fixed location, whereas the latter is based on robots that can navigate in large-scale space (that is, space that cannot be appropriately observed from a single vantage point).

In this report, a review of the research on mobile robot navigation is presented. The process of navigation is made up of the following fundamental elements: *locomotion mechanism*, *control system*, *sensing*, *environmental mapping*, *path planning*, and *localisation*. Each of these elements are surveyed in terms of their relevance to an unstructured outdoor environment, which is characteristically complex and uncertain. The author conjectures that research in this environment is more challenging than in indoor areas, as more emphasis is placed on a robot's robustness and generality. Also, common assumptions such as a planar surface, parallel walls, constant lighting conditions and consistent landmarks cannot be used as a crutch to simplify the navigation process.

Section 2 provides a review of each of the navigational elements. These elements are described in a combined form in section 3, by way of several examples of experimental robotic systems. The report concludes with a discussion in section 4, which suggests several ideas for future work.

*This quote is given in the context of ancient practices.

†The definition of "intelligence" is a highly debated topic, which will not be entertained herein.

2 Navigational Elements

A mobile robot is a combination of computational and physical elements. The computational system consists of a set of algorithms and models that control the robot's function. Conversely, the physical (hardware) system is the mechanical structure and sensory device that provides the operating base for the computational processes. A vital attribute of the physical system, is the locomotion mechanism that enables robot motion from place to place.

2.1 Locomotion

There are numerous terrestrial locomotion methods that have been developed, in support of robot movement on a solid surface [Dudek and Jenkin, 2000]. They can be categorised into four broad groups:

1. **Wheeled Mobile Robots** – This is the most commonly used locomotive method, which is based on rolling wheel motion. A variety of wheel arrangements, enabling translational and rotational movement, have been adopted such as differential steering, skid steering, synchronous drive, omni-directional drive, tricycle drive, and car drive (also known as Ackerman or kingpin steering).

2. **Tracked Vehicles** – These vehicles have a similar mechanism to differential steering, however, the wheels are extended into treads and rotational movement is performed through ground skid. Generally, they are suitable for rough terrain operation but they can have a destructive effect upon the ground whilst turning. When going straight, tracks tend to exert less ground pressure than that of wheels, assuming no slippage.

3. **Limbed Vehicles** – Legged motion is a biologically inspired locomotive method that is potentially very versatile. However, it has proven to be pragmatically complex, due to problems in stability maintenance, limb control and placement. A variety of limb arrangements and gaits have been investigated. The actuators which are used vary widely from servo motors to shape memory alloys.

4. **Miscellaneous Locomotive Strategies** – Includes uncommon strategies that do not cleanly fit within the first three groups (for example, the slithering locomotion of NASA's Snakebots).

For outdoor navigation, wheeled motion has been the preferred method. Wheeled mobile robots are commercially available from numerous manufactures (for example, iRobot, Nomadic, Robuter and Denning), eliminating the man-hours involved in designing and constructing a robot platform.

The building process can also be simplified by converting an outdoor vehicle into a mobile robot [Jarvis, 1993a]. When using this approach, it needs to be considered that the majority of outdoor vehicles have turning circle constraints; for instance, those with Ackerman steering. These constraints can severely effect a vehicle's maneuverability within a cluttered environment, making it difficult to navigate. Tracked vehicles are also highly suited to rough terrain work, however, they are usually more expensive.

Either way, regardless of how the hardware is developed, it is still only a dormant unit that requires a computational system for functionality.

2.2 Control Systems

In support of mobile robot navigation, the *control system* governs the following components: sensing, environmental mapping, path planning and localisation. Many different techniques and approaches for robotic control have been developed [Arkin, 1998]. The particular approach chosen determines the robot's intelligence, adaptability, predictability, speed of response, and computational complexity.

There are two extremities of the robot control spectrum: *deliberative reasoning* and *reactive control*. Usually its a matter of finding an appropriate point between the two extremities, to overcome the weaknesses of either extremity in isolation.

2.2.1 Deliberative Control

Deliberative control entails modeling the world and then using this knowledge to plan actions, predict outcomes and optimise performance. It is generally believed that a certain level of this type of control is required to solve complex problems. For instance, purposeful planning over an environmental model is required to find the shortest distance trajectory between two positions in a cluttered environment.

A problem with this control paradigm is that the models often rely on strong assumptions, and the more structured these assumptions, the greater the hazards of failure when they are not all met. Thus, assumptions that enable narrowly defined, high efficiency and accuracy often comes at the price of reduced robustness in the face of uncertainty outside these narrowly defined constraints.

Also, because the reasoning process takes a certain amount of time to make decisions, the operating environment may significantly change between data acquisitions. The resulting lag can cause errors in the model and, as a consequence, planned actions that are error prone. This problem can be combated by using reactive control, in a complementary fashion, to increase the system's response time and, thus, provide a safety net that minimises the risk of collision.

2.2.2 Reactive Control

In contrast to deliberative reasoning, reactive control exhibits a low-level of intelligence by forgoing the planning process. Reactive control was pioneered by Brooks ([Brooks, 1986]), Arkin ([Arkin, 1987]) and numerous other proponents. It is a technique that is based on tightly coupling perception and action, through an architecture that coordinates sensorimotor behaviours. The reactive paradigm can be characterised by its computational simplicity and real-time response. The robot reacts swiftly to sensory input, which is a property that has proliferated the use of reactive control for time-critical tasks (for example, collision avoidance).

A common strategy is to combine a high-level deliberative framework with low-level reactive control for maintaining vehicle safety. This hybrid strategy accommodates both complex problem solving and quick decision making, as the need arises.

2.3 Sensing

The robot needs to perceive various external stimuli, through the use of sensors, for tasks such as sensing environmental properties (for example, surface colour), ranging to objects, or ascertaining the robot's position within the environment. Each sensor has strengths and weaknesses, making sensor selection a tradeoff between pertinent criteria. For instance, the criteria involved in selecting a camera may include considerations like cost, resolution, frame rate, picture quality, field of view and versatility. A plethora of sensing techniques have been developed (see [Everett, 1995]), which can be broadly categorised into two types: *internal state* and *external state* sensors.

2.3.1 Proprioception

Internal state, or proprioceptive‡, sensors provide feedback on the robot's internal parameters, with no direct reference to the external world. These

‡A term borrowed from its use in a biological context.

sensors may include: odometry, gyroscopes, accelerometers, battery level indicators, and motor stall current detectors.

Odometry, as discussed in section 2.6.1, is the most widely used sensor for localisation. This sensor measures the angular rotation of a wheel through a device such as an *incremental optical encoder* or *synchro*, which is fitted to the motor shaft or the ground contact wheel. In the case of an incremental optical encoder, the number of pulses generated is directly proportional to shaft revolution; whereas with a synchro, a measure of magnetic coupling is used to indicate the absolute shaft orientation.

Inertial sensors[§] (that is, gyroscopes and accelerometers) are regarded as being internal state sensors with one provision: some of them reference the external world to counteract its interference with the measurements, as is the case with accelerometers which factor out the effects of the local gravity vector. Gyroscopes and accelerometers use Newtonian mechanics to measure the rate of rotation and acceleration, respectively. There are two commonly used types of gyroscopes: *mechanical gyroscopes* and *optical gyroscopes*. The former is based on the conservation of angular momentum of a spinning mass suspended in a gimbal, whereas the latter exploits the *Sagnac effect* (discovered by Sagnac in 1913 [Sagnac, 1913]) and has little or no moving parts. In the case of accelerometers, acceleration is measured from the displacement of a spring-mounted mass.

2.3.2 Exteroception

For navigatory purposes, a mobile robot needs to be able to observe the outside world. Such a task is performed using an array of external state, or exteroceptive, sensors. Some of the various kinds of external state sensors, include: tactile feelers (that is, touch), proximity sensors (for example, near infrared proximity detectors), magnetic compasses, global positioning systems (GPS), sonar, time-of-flight laser ranging, and machine vision. Since sensors are inherently noisy, gathering accurate environmental data is usually attained through redundancy or multi-sensor fusion (as described in [Kam *et al.*, 1997; Luo and Kay, 1989]). Sensor fusion is also a technique for combining sensors, in a complementary fashion, to overcome the limitations of using a particular sensor in isolation.

Vision can arguably provide the richest source of sensory data, in support of rudimentary tasks such as navigation and object manipulation. As a result, there is a large research contingent in the area of computer vision, investigating the many facets involved in the acquisition, processing and

[§] A class of sensors that measure derivatives of pose.

interpretation operations (see [Ballard and Brown, 1982] for the fundamentals of computer vision). Currently, camera images can be used, in a limited sense¶, for identifying objects, target tracking, or extracting surface shape and colour.

An ongoing problem with artificial vision is obtaining depth from a 2D image of an unstructured scene. Consequently, the field of rangefinding has grown to be an integral part of machine vision, in an endeavour to find complementary techniques that overcome the limitations of current camera technology. By fusing 2D vision with rangefinding sensors, as demonstrated in [Spero and Jarvis, 2002a; Jarvis, 1992b], a solution to 3D vision can be realised – circumventing the problem of inferring 3D from 2D.

Numerous rangefinding methods have been developed, each with strengths and weaknesses in terms of applicability constraints, accuracy, robustness, weight, power consumption, cost and safety [Jarvis, 1993d; Hebert, 2000]. These methods can be classified using three sets of dichotomies:

1. **Passive versus Active** – Passive methods rely on ambient lighting conditions, as they do not emit energy into the environment. Typical methods include lateral and temporal stereopsis, range from texture, range from focus, and range from attenuation. In contrast, active methods impose structured energy sources (for example, light, ultrasonic or microwave) upon the environment. Methods that fall into this category are striped lighting, sonar, radar and time-of-flight laser ranging.

2. **Image Based versus Direct** – Image based methods make range measurements by using image analysis, as opposed to direct methods that obviate such a need (for example, time-of-flight laser ranging).

3. **Monocular versus Multiple View** – Monocular methods extract range data from a single viewpoint, as in the case of range from texture or focus. Multiple view, or triangulation based, methods rely on identifying features in multiple images and determining range related disparities via correspondence matching. Triangulation based methods include lateral and temporal stereopsis – methods that are inherently subjected to missing parts/obscured edge problems.

Passive triangulation based methods are generally limited to short range applications and their performance is highly dependent on ambient environmental conditions. As a consequence, active methods are generally used for large-scale outdoor work (for instance, time-of-flight laser ranging as used in [Guivant *et al.*, 2000; Mayora *et al.*, 1998]).

¶Artificial vision has yet to reach the enormous expectations borne of our own facility.

2.3.3 Load-Bearing Surface Determination

A purely geometric description of the environment, obtained using rangefinders, is insufficient to characterise the traversability of a path in vegetated terrain. Due to the compressibility of terrain cover (grass, bushes, foliage), the *visible surface* does not always correspond to the *load-bearing surface* [Manduchi *et al.*, 2001]. The visible surface is the geometric surface as perceived by rangefinding sensors, while the load-bearing surface is the actual surface touched by the robot's wheels. As a representative example, a patch of tall grass may be easily traversable by a robot, without risk of damage. However, the grass would probably be incorrectly identified as an obstacle based solely on its geometry.

The extent of research towards determining the load-bearing surface is relatively small (several recent methods are surveyed in [Manduchi *et al.*, 2001]). A colour-based classification scheme is given in [Bellutta *et al.*, 2000], where images from a colour camera were used to identify green vegetation, dry vegetation, soil and rocks. This approach is highly dependent on ambient lighting conditions; so taken to extremes, it cannot be used at night. Another approach, proposed in [Macedo *et al.*, 2001; 2000], uses statistical analysis of laser rangefinder (ladar) data to differentiate between rocks (non-traversable) and grass (traversable). It is based on the premise that the distribution of range readings from rocks are not as scattered as those from grass. There are also other approaches that are currently being investigated, such as the analysis of ultrasonic or non-visible spectral signatures.

2.4 Environmental Mapping

To carry out complex missions in an unknown rough terrain environment, the robot must be able to incrementally generate and maintain a map of its environment. The robot builds a map by gathering sensor data as it moves through the environment. The map, as an internal representation of space, can then be used by the robot for deliberative reasoning so actions can be preplanned and optimised before execution.

Comprehensive surveys on environmental modeling techniques are given in [Thrun, 1998; Chatila and Laumond, 1985]. Two modeling paradigms have clearly emerged: *geometric models* and *topological models*.

2.4.1 Geometric Models

Geometric models represent the physical location and configuration of environmental features, such as obstacles and landmarks. The models can be categorised as either *Cartesian based* or *tessellated*. A Cartesian model consists of discrete geometric primitives (for example, points, lines or polynomial functions), which are described in terms of a Cartesian coordinate system. This approach, as used in [Zhang and Ghosh, 2000; Leonard *et al.*, 1990; Ayache and Faugeras, 1989], is highly space-efficient as each stored object is represented by only a few numerical parameters. However, assumptions are made regarding object type and pose, reducing the model's generality. Tessellated models, on the other hand, do make these assumptions as they represent space itself – not the individual objects within it.

A tessellated, or grid-based, model has a simple lattice structure, with each element corresponding to a region in the environment. The commonly used approach, proposed by Moravec and Elfes ([Elfes, 1989; Moravec and Elfes, 1985]), is to spatially decompose the environment into an evenly-spaced grid. Each grid cell then indicates, in a probabilistic manner, the presence of an obstacle.

Due to this spatial sampling, a tessellated model is storage intensive; however, this is becoming less of a concern with the rapidly falling price of computer memory.

2.4.2 Topological Models

In contrast to relying on metric data that can be error prone, topological models represent an environment by graphs [Choset and Nagatani, 2001; Kuipers and Byun, 1991]. The graphs indicate the spatial relationships between environmental features, with a set of nodes corresponding to places and/or landmarks. An arc connects a pair of nodes, if there exists a direct path between them.

Topological models can be associated with the way humans communicate directions. For instance, to reach a particular research laboratory at Monash University, from an initial position, one may need to travel down a hallway and then turn left at the end; head down another hallway and then turn right at the second doorway. Such an approach does not require accurate metric information. However, it is possible that important places in the environment are difficult to detect. Also, ambiguities can result from multiple places that are detected but look the same.

To overcome the inherent weaknesses in geometric and topological models, a number of hybrid schemes have been proposed (for example, [Thrun and Bücken, 1996]). These schemes integrate metric and relational information, facilitating consistent modeling and timely problem solving, respectively.

2.5 Path Planning

Path planning is an area that has been extensively researched, with comprehensive surveys given in [Hwang and Ahuja, 1992; Latombe, 1991]. A mobile robot needs to be able to plan a collision-free trajectory, to efficiently move from an initial pose to a specified target pose within the environment. Most of the algorithms have been developed for known indoor workspaces (sometimes referred to as the piano movers problem), where the objective is to usually find paths that are of minimal distance (as in A* search [Nilsson, 1971]).

However, the quest for distance optimality alone is insufficient in *a priori* unknown rough terrain, as the complex and harsh nature of the environment cannot be disregarded. Consequently, path planning is at best a suboptimal arrangement of distance, time, energy consumption and safety factors. In practice, optimality is often sacrificed for feasibility so system objectives can be achieved in a rational manner. For instance, the path planner may purposefully bias the robot towards open spaces for safety; or away from open spaces, as is the case with surreptitious missions.

Several techniques have been proposed for navigating in unstructured environments, which can be categorised into two groups: *global* and *local* path planners. The groups are differentiated by the scope of environmental knowledge that is incorporated into the reasoning process.

2.5.1 Global Path Planners

Global path planners consider all environmental information and are characterised as planning complete paths from initial to target poses. The bug algorithm, proposed in [Lumelsky and Stepanov, 1987], is a global path planner that finds viable paths by switching between two reactive behaviours: moving directly towards the target position; and circumnavigating an obstacle. Apart from assuming a holonomic point robot with perfect odometry, this simple algorithm can generate paths that are significantly worse than the optimal path, with respect to distance.

The distance transform methodology [Jarvis, 1993b] is a simple and robust technique used for both finding optimal collision free paths and obstacle growing (to accommodate the physical extent of the robot). The distance transform is calculated over a grid structured spatial representation. In the case of path planning, distances are propagated throughout each grid cell in an outwards direction from specified goal points, to ultimately fill the entire free space. Optimal paths are then found by using a steepest descent trajectory from any point in free space, without risk of local entrapment. Even though the distance transform has a generality (conferred by the tessellated model) that is useful in rough terrain environments, it is computationally expensive and does not scale well to large navigational areas at high resolution.

A versatile path planning strategy was proposed by LaValle and Kuffner ([LaValle and Kuffner Jr., 1999]), termed the Rapidly-exploring Random Tree (RRT). An RRT is a randomised data structure that is used to compute collision-free kinodynamic trajectories for high degree-of-freedom problems. To account for a robot's dynamic constraints, a state space representation is used which includes both configuration and velocity parameters. A path plan is generated by defining two RRTs, one rooted at the start state and the other at the goal state. Both trees are grown by first selecting a random state from the state space and then searching each tree for the nearest neighbouring state. All viable control inputs are applied to the neighbouring states to generate possible successor states, based on whether they are collision-free, satisfy velocity bounds and minimise some chosen metric (for example, minimum distance) to the random state. The procedure is repeated until two states, one from each tree, are regarded as being sufficiently close in the state space to render a solution. The algorithmic pseudo code is listed in [LaValle and Kuffner Jr., 1999].

RRTs rapidly explore the state space and therefore scales well to large navigational areas, however, it is assumed that the environment is known *a priori* and is structured to assist with collision detection in real space. Also, the nearest-neighbour search is a computationally expensive step. These problems are addressed in [Spero and Jarvis, 2002b], by adapting the RRT approach to a tessellated model.

While numerous other global path planners exist, not many can handle the complexities of a rough terrain environment. In any case, a global map may not be available, making it necessary to use a local path planner.

2.5.2 Local Path Planners

Local path planners rely on local information about nearby obstacles for the purpose of moving towards the target pose, while avoiding these obstacles. As a representative example, the artificial potential field method [Khatib, 1986; Montano and Asensio, 1997] is a local path planner that is based on a physical analogy. It models the goal point as an attractant and the obstacles as repellents and then computes a continuous potential field to reactively guide the robot to the goal point. One of the problems with this path planner is that it commonly suffers from *cul de sacs* (concave obstacle configurations) that trap the robot, which is generally a problem incurred by all local path planners.

Since local methods can perform real-time obstacle detection, they are sometimes used to augment global path planners as a safety measure. Obstacles detected by the local path planner, during the robot's motion, can then be safely avoided; even if they are not present in the global map of the environment.

2.6 Localisation

A mobile robot needs to be able to determine its pose* within the environment, for the purposes of path planning and environmental mapping. The process of estimating the robot's pose, usually termed localisation, is considered by some researchers to be the most important process in robot autonomy, as stated in [Cox, 1991].

Localisation can be performed either on an ongoing or sporadic basis. Depending on the environmental representation, localisation can also be in qualitative or quantitative terms – the latter being the most commonly used in outdoor environments. Surveys of localisation methodology can be found in [Borenstein *et al.*, 1997; Jarvis, 1993c].

Compared to the simplicity, accuracy and reliability of localisation methods used in factory environments (for example, buried guide wires or painted stripes on the ground), localisation in rough terrain is an extremely complex problem with no robust solution as yet. However, there are currently several localisation methods and technologies used for outdoor navigation, in a limited capacity†. They can be categorised into two groups: *dead-reckoning* and *reference-based* systems.

2.6.1 Dead-Reckoning

Dead-reckoning‡ refers to pose estimation based on the observation of internal parameters. In a biological context, dead-reckoning (known as idiothetic sensing) corresponds to navigating a workspace without one's facilities of sight, smell, hearing, etcetera. Since no reference is made to the external world, this localisation approach uses relative pose measurements and is subject to accumulative errors.

Odometry, as described in Section 2.3.1, is the most widely used localisation method. By integrating internal motion information over time, the robot's pose can be estimated in an open loop manner. This is a simple,

*The robot's "pose" refers to its position and orientation.

†Experiments are usually conducted, over short time durations, in agreeable outdoor environments.

‡Derived from "deduced reckoning".

low-cost solution that can provide adequate short-term accuracy, especially when the robot is traversing over smooth planar surfaces.

However, odometry is prone to two types of errors: *systematic* and *non-systematic* errors. The former results from kinematic imperfections (for example, irregular wheels or imprecise wheelbase calculations), while the latter is caused by the unpredictable interaction of the locomotion system with the environment (for example, wheel slippage). Large non-systematic errors can occur when the robot traverses over loose ground, such as soil or gravel surfaces, causing significant wheel slippage.

The problem with a momentary orientation error in odometry is that it causes a constantly growing lateral position error [Borenstein *et al.*, 1997]. As a consequence, gyroscopes are sometimes incorporated into the dead-reckoning process to increase the reliability of orientation measurements. A more robust solution that is commonly used, is to relegate dead-reckoning to short-term pose estimation and use a reference-based system for long-term pose fixes.

2.6.2 Reference-Based Systems

Reference-based systems make absolute pose measurements by referencing the external world. In doing so, they eliminate accumulative errors in the robot's position that result from the use of dead-reckoning. Reference-based systems are therefore needed for robust outdoor navigation, however, they tend to be relatively complex and often require costly site preparation. Some of the commonly used reference-based systems include: *active beacon systems*, *global positioning systems (GPS)* and *landmark-based localisation*.

Active beacon systems involve placing at least three[§] transmitters, or transceivers (transmitter-receiver units), within the environment and then estimating the robot's pose through either *triangulation* or *trilateration*. In the former case, the angles to three beacons are measured using a receiver on-board the robot; while, in the latter case, the distances to three beacons are measured instead, using time-of-flight information. By solving a series of constraint equations that describe the geometric relationships, accurate pose estimates can be made. This approach is only reliable within certain areas, as the beacons' transmitted radio frequency (RF) signals have a focused propagation pattern[¶] and are usually only visible via line-of-sight (LOS). System reliability is also dependent on the beacon configuration, with respect to the robot's pose.

[§]For 2D pose fixes.

[¶]Also called a communication footprint.

GPS is a trilateration based technology that was developed by the United States Department of Defence. The system uses 24 satellites (excluding redundancy) as active beacons, each transmitting encoded RF signals. A ground receiver needs LOS to three satellites to make a 2D fix. For national security reasons, the US government used to intentionally degrade the accuracy of this system through a safety protocol termed *selective availability* (SA). As a consequence, localisation accuracy was around 100 metres (2drms) best-case, which is generally too large for robotic operations. Selective availability is now no longer in use, however, the errors are still too large for navigating in some cluttered environments.

A viable alternative, based on the same system, is differential GPS (DGPS) which is more accurate. DGPS involves being in close proximity (that is, within several tens of kilometres) to another ground receiver that has a precisely surveyed position, and can therefore act as a reference beacon. Corrective signals are transmitted by this reference beacon to counteract the errors in the GPS signals.

There are a variety of DGPS schemes that exist, each with a different positional accuracy. For instance, a typical code-phase DGPS has an accuracy well under ten metres, whereas the more costly carrier-phase DGPS has an accuracy in the order of a few centimetres [Everett, 1995]. However, regardless of what type of GPS is used, the fact that LOS to three satellites is needed can present a problem in vegetated terrain. The natural surroundings can occlude the signals, causing a loss of tracking and possibly task failure.

Landmark-based localisation is a method that involves the detection and recognition of distinct features in the environment. Similar to active beacon systems, triangulation to three features is used to estimate the robot's pose. The features, or landmarks, can be either artificial or natural – a distinction based on whether objects are purposefully placed in the environment for robot navigation. Artificial landmarks, such as bar-coded reflectors and signposts, are designed to be highly visible through various kinds of markings, patterns, colours, sizes, shapes, and materials. They are strategically placed in the environment in an endeavour to maximise system reliability and positioning accuracy. However, relying on artificial landmarks for localisation, restricts the robot's configuration space to areas where the appropriate number of landmarks are visible.

Natural landmark based localisation is a generalised approach, which does not require cumbersome, and perhaps expensive, site preparation. This method has biological correlates that confer its potential applicability and corresponding value to the robotics field. Selecting natural landmarks is based on a desirable criteria, including: observability, uniqueness, temporal stability, geometric distribution, and lateral compactness [Lipton, 1996;

Clark and Dissanayake, 1999]. As opposed to the natural landmarks exploited in indoor workspaces (for example, corners and edges), definitive primitives cannot be reliably used in rough terrain, where free-form features dominate. Consequently, natural landmark based localisation is extremely complex; with no robust solution as yet.

2.6.3 Simultaneous Localisation and Mapping

To globally localise in an *a priori* unknown environment, the robot needs to perform the process of *simultaneous localisation and mapping* (SLAM)‖. SLAM involves incrementally building a feature based map of the environment, while using the map to globally localise the robot [Williams *et al.*, 2000]. The key problem in SLAM is coping with three forms of uncertainty [Leonard *et al.*, 2001], including: data association uncertainty, navigation error, and sensor noise.

As described in [Dissanayake *et al.*, 2001], there are currently three philosophical approaches used to address the SLAM problem: *estimation-theoretic (statistical)*, *numerical (non-statistical)*, and *qualitative* approaches. The estimation-theoretic approach, based on the Kalman Filter (KF), is the most popular approach due to its consistent theoretical framework and long history in other fields such as the aerospace and maritime sciences.

The KF solution, as devised in [Leonard and Durrant-Whyte, 1991]*, involves a recursive update procedure that comprises prediction, observation, and update steps. Statistical models are used in this procedure to estimate the uncertainty in the robot and landmark locations, along with their intercorrelations. It is these models that enable a thorough investigation into various SLAM properties such as the convergence to a solution and the evolution of positional uncertainties. However, they are also the source of practical vulnerabilities, as they are based on several underlying assumptions. Another issue to consider is that this approach tends to be computationally expensive, especially when a large number of landmarks are being tracked.

Other numerical techniques to solve SLAM have emerged, such as the probabilistic Bayesian (maximum likelihood) approach [Thrun *et al.*, 1998], the particle filter [Doucet *et al.*, 2001], and the set-theoretic approach [Di Marco *et al.*, 2000]. While some of these techniques are more robust than the KF based approach, they generally are computationally expensive and rely on the accuracy of models. A recent technique proposed in [Spero and Jarvis, 2004b; 2004a] moves away from these models and simplifies the process of

‖A phrase first coined in [Leonard and Durrant-Whyte, 1991].

*Based on key precursors: [Smith *et al.*, 1990; Moutarlier and Chatila, 1989] tracking positional uncertainties. It has shown promising results, however, a more in-depth investigation is required.

Several qualitative techniques have also been proposed (for example, [Choset and Nagatani, 2001]), that obviate the need for absolute pose estimates. Instead, they employ relational knowledge of the relative position of the robot and landmarks. They are generally computationally efficient, however, they incur the same problems inherent with topological modeling (described in Section 2.4.2), such as spatial ambiguity (aliasing).

Since SLAM is thought to be the most critical element in robot autonomy, a growing research contingent continues en masse to find that elusive solution. Once found, robotics will move one giant step closer to autonomous operation in the real-world.

3 Robotic Systems

In this section, the navigational elements are discussed in terms of complete robotic systems. Several representative examples of outdoor robots that have been developed, or are in development, were chosen to form the basis of this discussion. This is by no means an exhaustive survey. However, these examples do provide a sample of some of the work conducted at various robotic research centres around the world.

An important issue to note is that these robots are not yet autonomous; that is, human operators are always, to some extent, in the control loop. At the very least, a human operator is responsible for pressing a "stop" button as soon as the robot deviates from what is considered normal behaviour. However, usually there is more human involvement in a robot's operation by means of a *teleoperation interface* [Fong and Thorpe, 2001]. This interface may comprise a joystick, voice commands, human gestures, or Web-based control.

Currently, several application areas for outdoor robots are being investigated. Some of these areas include: *cross-country and interplanetary exploration*, *law enforcement*, and *search and rescue*.

3.1 Cross-Country and Interplanetary Exploration

The ability to explore, or discover, unknown outdoor areas is useful for such tasks as military reconnaissance, map construction (cartography), mining, and hazardous area inspection. There are numerous robots that have been developed for the purpose of exploration, some of which were designed for competition (for instance, the DARPA Grand Challenge over desert terrain [DARPA, 2004]).

For the purpose of cross-country exploration, the locomotion mechanism and chassis is often based on that of a commercially available vehicle which has been converted into a semi-autonomous robot. The conversion process involves the installation of actuators for the steering wheel, brake pedal and throttle (see [Bentivegna *et al.*, 1998]). Also a variety of sensors are fitted, along with a computational unit and, possibly, wireless communication to a remote control station.

As an example, the NAVLAB II robot is based on a modified four-wheel drive HMMWV (High-Mobility Multi-purpose Wheeled Vehicle) [Stentz and Hebert, 1995]. This robot has been instrumented with a scanning laser rangefinder for detecting obstacles in the environment. The obstacles are represented using an occupancy grid map, which evolves over time as the robot explores the environment. Global path planning is performed using an efficient version of A* search, called the D* algorithm. This path planner dynamically adjusts the global path based on a low-level reactive system that performs obstacle avoidance in real-time. For safety reasons, a person is stationed inside the vehicle so that the robot controller can be overridden, if required.

In the context of interplanetary exploration, the robots are built with respect to the rigours of the targeted planet and their intended application [Pedersen *et al.*, 2003]. For the purpose of Mars exploration, two well-known examples include the Russian built *Marsokhod* [Kemurdjian *et al.*, 1992] and the US built *Sojourner* [Stone, 1996]. The Marsokhod has never been deployed to Mars, but has facilitated important research contributions from its use as a testing platform. The Sojourner, however, was sent on a mission to Mars in 1997 and was the first to do so.

Both robots have a similar locomotion mechanism, which consists of a six-wheeled skid steering arrangement and an articulated chassis for overcoming difficult relief. In terms of control, these robots are teleoperated either in real-time or via command sequences that are uploaded by human operators. They do, however, have autonomous capabilities, but this type of control is restricted to avoiding obstacles and other physical hazards. Some of the sensors that were fitted to the Sojourner, on its mission to Mars, included: accelerometers, gyroscope, odometry, temperature sensors, touch (or bump) sensors, cameras, laser stripe projectors, real-time clock, telemetry, and an Alpha Proton X-Ray Spectrometer (for examining the composition of the Martian surface).

3.2 Law Enforcement

The most well-known application of robotics in law enforcement is explosive ordnance disposal. However, there are other possible applications that go far beyond this. Robots are currently being developed for such tasks as surveillance, reconnaissance, under-vehicle inspection, breaching doors, and launching less-lethal ammunition (for example, snare nets, tear gas, ballistic bags, and rubber balls).

Several of these robots are discussed in [Nguyen and Bott, 2000]. As a representative example, the SPIKE robot (produced by II-Tracker of Portland, OR) is a 425-pound tracked vehicle that can breach solid-core doors by simply driving through them. Apart from breaching doors, this robot can also be used for reconnaissance, bomb disposal, or even intimidation (as a psychological weapon). It has an 8-horsepower diesel engine, and a radio control link for teleoperation. There is little or no autonomous functionality, as every law enforcement operation presents a unique scenario that requires complex tactical planning and execution, which is best left to trained personnel. A variety of sensors and manipulators can be selected for use on this robot, depending on the specifics of the task.

The ODIS robot, described in [Smuda *et al.*, 2002], is one of the few robots that have been developed for inspecting the underside of vehicles at security checkpoints. This robot is a low profile omni-directional vehicle that is only 3.75-inches tall and highly maneuverable. It is able to navigate autonomously using a behaviour based control strategy, which invokes desired behaviours based on the environmental information gathered from sonar, infrared and laser sensors. ODIS is equipped with a colour camera, mounted on a pan-tilt mechanism, for the purpose of detecting bombs and other contraband. To enable the inspection of dark cavities, the camera is surrounded by an array of bright LEDs that provide active lighting of the viewable area.

The benefit of using ODIS is that security inspectors do not have to come in close contact with the vehicle under inspection to conduct the standard “mirror on a stick” examination. Also, a swarm of these robots can potentially be used in instances where a large group of vehicles need to be inspected, as was the case at the US/Canadian border crossings just after the September 11, 2001 terrorist attack. However, due to the critical nature of the inspection process, the autonomous aspect of this technology has not yet been accepted by the law enforcement community, especially when it lacks a proven track record.

3.3 Search and Rescue

The area of search and rescue is of utmost importance in adverse situations such as earthquakes, fires, bombings, and collapsed buildings. It is the perpetual occurrence of these situations, along with the RoboCupRescue competition [RoboCupRescue, 2001], that has stimulated a growing interest in the development of robots in this area. Robots can potentially assist rescue workers in carrying out a number of different tasks, some of which include: assessing danger zones (reconnaissance); locating structural weaknesses, gas leaks, or HAZMATs (hazardous materials); searching for and rescuing survivors trapped under ruble, lost at sea, or missing in remote bushland/mountains; patrolling the beaches for drowning swimmers**; extinguishing fires; and removing debris.

In the case of a collapsed building, there are many issues involved in a search and rescue operation and so a robot can provide support in a variety of different ways [Casper *et al.*, 2000]. As an example, there were several robots used in the World Trade Center disaster (on September 11, 2001) to find survivable voids within the rubble pile that could be rapidly excavated, providing access to the basement where there might be survivors [Murphy, 2004]. These robots were tracked vehicles that were mainly teleoperated through a tether that provided both communications and power. There was one robot, called SOLEM, that was teleoperated using wireless communication, however, the rubble caused intermittent wireless dropout. Also, SOLEM had to be connected to a safety rope, which restrained it in the same manner as a tether. In terms of sensors, most of the robots were fitted with a colour camera and two-way audio, for controlling the robot and communicating with victims, respectively.

Murphy, in [Murphy, 2000], suggests the use of shape-shifting robots for navigating through tight spaces in rubble that are too small or dangerous for humans and search dogs. These robots can alter their physical configuration to adapt to their immediate surroundings. For example, the Bujold robot (built by Inuktun Services of Canada) is a chemical-inspection microrobot that can change its tracked locomotion mechanism between three configurations. Each configuration changes the robot's point of view, along with its height, center of gravity, and the ease at which it can pivot on its tracks. While this robot has shown promising results, the high number of degrees of freedom and its limited view of the world make teleoperation a difficult exercise. As a result, a certain level of autonomy is warranted, but given the hostile environment, this is extremely complex to implement.

Several other innovative robots are proposed in [Hirose and Fukushima,

**This may be construed as an oxymoron.

2002]. Two examples of these robots include the ACM-R3 robot and the Genbu robot, which are both the subject of ongoing research. ACM-R3 is a mechanical snake with a smooth and slender articulated body, modeled after its biological counterpart. This robot can enter through small cracks and crevices in rubble, and follow tortuous tunnels. Because of this unique penetrating ability, it is currently being promoted for use in reconnaissance missions. The Genbu robot, by contrast, is a fire-fighting robot that has a wheeled locomotion mechanism and a jointed body for flexibility. It is connected to a high pressure water hose from a fire truck, and novelly uses the hydraulic energy in this hose to actuate itself.

In future, this and many other robots may prove to be extremely useful in carrying out real search and rescue operations. Note that the success of these operations is often based on one critical factor – **time**. Of course, the longer that victims are trapped or incapacitated, the more likely they are to succumb to serious threats such as hypothermia, hypovolemia, dehydration, or asphyxiation. Robots can therefore play an important role by decreasing rescue times, and providing special capabilities that are not otherwise available. However, substantially more research is required in this humanitarian domain for such benefits to eventuate.

4 Discussion and Future Work

This report has described the fundamental concepts and ideas in mobile robot navigation, and how it pertains to unstructured outdoor environments. Numerous key references have also been given, allowing the reader to delve deeper into more focused areas.

There are numerous open problems in robotics that remain unresolved. Many of them relate to a robot's fallibility when faced with the chaos of the real-world. Usually, these type of problems are dealt with by either solving them or, more commonly, confining the robot's world until they are eliminated. This leads to one irrefutable truth:

A human's freedom is a robot's poison.

A significant amount of research needs to be done to remove this anomaly between humans and robots. Perhaps only when robots have an affinity for the unknown, will they be truly autonomous. Until then, there are numerous avenues for future research in robotics, some of which are listed below (in no particular order):

- **Safety** – As stated by Asimov's laws [Asimov, 1942], a robot should not harm human beings. Conversely, human beings should not harm a robot either. Therefore the idea of self-preservation needs to be investigated so that techniques can be developed, within legal limits.

- **Unstructured and/or Dynamic Environments** – Complex environments demand that robots be robust and adaptable for survival. This invariably leads to more generalised solutions than those developed in known structured environments.

- **Power Supplies** – Mobile robots rely heavily on portable power supplies for their mobility. The search for better power supplies in terms of size, weight, cost, energy supply and recharging ability is ongoing.

- **Locomotion Mechanisms** – Robots with a reconfigurable or flexible locomotion mechanism are more able to adapt to their immediate surroundings than rigid designs. Consequently, more research needs to done in developing adaptable robots, along with their complex control.

- **Sensors and Sensor Fusion** – There is a continuous demand for better nonvisual and vision-based sensors, along with their fusion to overcome individual limitations. Currently, 3D vision is one of the many topics of interest.

- **Landmark Detection and Recognition** – Free-form landmarks in a natural environment are difficult to detect and recognise. However, these abilities are needed for natural landmark based localisation.

- **Load-Bearing Surface Determination** – The problem of differentiating between the visible surface and the load-bearing surface has not been thoroughly investigated. Therefore load-bearing surface determination, and its impact on the path planning process, is one of the possible avenues for future research.

- **SLAM** – Currently, most of the SLAM methods are highly dependent on models, and associated assumptions, which compromise their robustness in natural environments. Thus, the search for a robust and flexible solution continues.

- **Multi-Agent Robotics** – A team of simple robots that cooperate to accomplish some common goal can potentially outperform a single complex robot. However, there are numerous problems, specific to multi-agent systems, that first need to be resolved (see [Arkin and Hobbs, 1992; Dudek *et al.*, 1996]). These robots may in future prove useful for tasks such as interplanetary exploration, pushing heavy objects, and cleaning up toxic waste.

- **Artificial Intelligence** – The learning and reasoning processes of a robot's artificial brain have been well researched. However, significantly more research needs to be conducted before a robot can mentally evolve within its environment for survival.
- **System Integration** – The integration of the many different hardware and computational elements that make up a mobile robot is a challenging task. Consequently, there is a need for tools and techniques that can assist in this process.

References

[Arkin and Hobbs, 1992] R. C. Arkin and J. D. Hobbs. Dimensions of communication and social organization in multi-agent robotic systems. In *Proc. 2^{nd} Int. Conf. Simulation of Adaptive Behavior*, pages 486–493, Honolulu, HI, Dec. 1992.

[Arkin, 1987] R. C. Arkin. Motor schema based navigation for a mobile robot: An approach to programming by behavior. In *Proceedings of the IEEE Conference on Robotics and Automation*, pages 264–271, Raleigh, NC, 1987.

[Arkin, 1998] R. C. Arkin. *Behavior-Based Robotics*. MIT Press, Cambridge, MA, 1998.

[Asimov, 1942] I. Asimov. *Runaround*. Street and Smith Publications, Inc., New York, 1942.

[Ayache and Faugeras, 1989] N. Ayache and O. D. Faugeras. Maintaining representations of the environment of a mobile robot. *IEEE Transactions on Robotics and Automation*, 5(6):804–819, Dec. 1989.

[Ballard and Brown, 1982] D. H. Ballard and C. M. Brown. *Computer Vision*. Prentice Hall, Englewood Cliffs, NJ, 1982.

[Bellutta *et al.*, 2000] P. Bellutta, R. Manduchi, L. Matthies, K. Owens, and A. Rankin. Terrain perception for DEMO III. In *Proceedings of the IEEE Intelligent Vehicles Symposium 2000*, pages 326–331, Dearborn, MI, 3-5 Oct. 2000.

[Bentivegna *et al.*, 1998] D. C. Bentivegna, K. S. Ali, R. C. Arkin, and T. Balch. Design and implementation of a teleautonomous hummer. In *Mobile Robots XII*, volume 3210, pages 130–138, Pittsburgh, PA, 16 Oct. 1998. SPIE.

[Borenstein *et al.*, 1997] J. Borenstein, H. R. Everett, L. Feng, and D. Wehe. Mobile robot positioning: Sensors and techniques. *J. Robotic Systems*, 14(4):231–249, Apr. 1997.

[Brooks, 1986] R. Brooks. A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, RA-2(1):14–23, 1986.

[Casper *et al.*, 2000] J. Casper, M. Micire, and R. R. Murphy. Issues in intelligent robots for search and rescue. In *Unmanned Ground Vehicle Technology II*, volume 4024, pages 292–302, Orlando, FL, 24-25 Apr. 2000. SPIE.

[Chatila and Laumond, 1985] R. Chatila and J. Laumond. Position referencing and consistent world modeling for mobile robots. In *IEEE International Conference on Robotics and Automation*, pages 138–145, St. Louis, MO, 25-28 Mar. 1985.

[Choset and Nagatani, 2001] H. Choset and K. Nagatani. Topological simultaneous localization and mapping (SLAM): Toward exact localization without explicit localization. *IEEE Transactions on Robotics and Automation*, 17(2):125–137, Apr. 2001.

[Clark and Dissanayake, 1999] S. Clark and G. Dissanayake. Simultaneous localisation and map building using millimetre wave radar to extract natural features. In *Proceedings of the 1999 IEEE International Conference on Robotics and Automation*, volume 2, pages 1316–1321, Detroit, MI, May 1999.

[Cox, 1991] I. J. Cox. Blanche - an experiment in guidance and navigation of an autonomous robot vehicle. *IEEE Transactions on Robotics and Automation*, 7(2):193–204, Apr. 1991.

[DARPA, 2004] DARPA. DARPA Grand Challenge. http://www.darpa.mil/grandchallenge/, 2004+. Visited on 3/11/2004.

[Di Marco *et al.*, 2000] M. Di Marco, A. Garulli, S. Lacroix, and A. Vicino. A set theoretic approach to the simultaneous localization and map building problem. In *Proceedings of the 39th IEEE Conference on Decision and Control*, volume 1, pages 833–838, Sydney, Australia, Dec. 2000.

[Dissanayake *et al.*, 2001] M. W. M. G. Dissanayake, P. Newman, S. Clark, H. F. Durrant-Whyte, and M. Csorba. A solution to the simultaneous localization and map building (slam) problem. *IEEE Transactions on Robotics and Automation*, 17(3):229–241, Jun. 2001.

[Doucet *et al.*, 2001] A. Doucet, N. de Freitas, and N. Gordon, editors. *Sequential Monte Carlo Methods in Practice*. Springer-Verlag, New York, 2001.

[Dudek and Jenkin, 2000] G. Dudek and M. Jenkin. *Computational Principles of Mobile Robotics*. Cambridge University Press, Cambridge, 2000.

[Dudek *et al.*, 1996] G. Dudek, M. R. M. Jenkin, E. Milios, and D. Wilkes. A taxonomy for multi-agent robotics. *Autonomous Robots*, 3(4):375–397, Dec. 1996.

[Elfes, 1989] A. Elfes. *Occupancy Grids: A Probabilistic Framework for Robot Perception and Navigation*. PhD thesis, Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA, 1989.

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

[Fong and Thorpe, 2001] T. Fong and C. Thorpe. Vehicle teleoperation interfaces. *Autonomous Robots*, 11(1):9–18, Jul. 2001.

[Guivant *et al.*, 2000] J. Guivant, E. Nebot, and S. Baiker. High accuracy navigation using laser range sensors in outdoor applications. In *Proceedings of the 2000 IEEE International Conference on Robotics and Automation*, volume 4, pages 3817–3822, San Francisco, CA, 24-28 Apr. 2000.

[Hebert, 2000] M. Hebert. Active and passive range sensing for robotics. In *Proc. IEEE Int. Conf. Robotics and Automation*, pages 102–110, San Francisco, CA, 24-28 Apr. 2000.

[Hirose and Fukushima, 2002] S. Hirose and E. F. Fukushima. Development of mobile robots for rescue operations. *Advanced Robotics*, 16(6):509–512, 2002.

[Hwang and Ahuja, 1992] Y. K. Hwang and N. Ahuja. Gross motion planning - a survey. *ACM Computing Surveys*, 24(3):219–291, Sep. 1992.

[Jarvis, 1992a] R. Jarvis. Computer vision and robotics - hand/eye coordination and autonomous navigation. In *Proceedings of the Australian Computer Science Conference 15*, pages 399–410, Hobart, Australia, Jan. 1992.

[Jarvis, 1992b] R. A. Jarvis. 3D shape and surface colour sensor fusion for robot vision. *Robotica*, 10:389–396, 1992.

[Jarvis, 1993a] R. Jarvis. Converting outdoor vehicles into autonomous mobile robots. In *Proceedings of the 1993 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 591–597, Yokohama, Japan, 26-30 Jul. 1993.

[Jarvis, 1993b] R. Jarvis. *Recent Trends in Mobile Robots*, volume 2, chapter Distance Transform Based Path Planning for Robot Navigation, pages 3–31. World Scientific, NJ, 1993.

[Jarvis, 1993c] R. Jarvis. A selective survey of localisation methodology for autonomous mobile robot navigation. In *Australian Robot Association's International Conference on Robots for Competitive Industry*, pages 310–321, Brisbane, Australia, 14-16 Jul. 1993.

[Jarvis, 1993d] R. Jarvis. *Three-Dimensional Object Recognition Systems*, chapter Range Sensing for Computer Vision. Elsevier, Singapore, 1993.

[Kam *et al.*, 1997] M. Kam, Z. Xiaoxun, and P. Kalata. Sensor fusion for mobile robot navigation. *Proceedings of the IEEE*, 85(1):110–119, Jan. 1997.

[Kemurdjian *et al.*, 1992] A. Kemurdjian, V. Gromov, V. Mishkinyuk, V. Kucherenko, and P. Sologub. Small Marsokhod configuration. In *Proc. IEEE Int. Conf. Robotics and Automation*, pages 165–168, Nice, France, 12-14 May 1992.

[Khatib, 1986] O. Khatib. Real-time obstacle avoidance for manipulators and mobile robots. *Int. J. Robotics Research*, 5(1):90–98, 1986.

[Kuipers and Byun, 1991] B. Kuipers and Y. T. Byun. A robot exploration and mapping strategy based on a semantic hierarchy of spatial representations. *Robotics and Autonomous Systems*, 8(1-2):47–63, Nov. 1991.

[Latombe, 1991] J. C. Latombe. *Robot Motion Planning*. Kluwer Academic Publishers, Boston, 1991.

[LaValle and Kuffner Jr., 1999] S. M. LaValle and J. J. Kuffner Jr. Randomized kinodynamic planning. In *Proceedings of the 1999 IEEE International Conference on Robotics and Automation*, volume 1, pages 473–479, Detroit, MI, May 1999.

[Leonard and Durrant-Whyte, 1991] J. J. Leonard and H. F. Durrant-Whyte. Simultaneous map building and localization for an autonomous mobile robot. In *IEEE/RSJ Int. Workshop on Intelligent Robots and Systems*, volume 3, pages 1442–1447, Osaka, Japan, 3-5 Nov. 1991.

[Leonard *et al.*, 1990] J. Leonard, H. Durrant-Whyte, and I. J. Cox. Dynamic map building for an autonomous mobile robot. In *IEEE International Workshop on Intelligent Robots and Systems '90*, volume 1, pages 89–96, Ibaraki, Japan, Jul. 1990.

[Leonard *et al.*, 2001] J. J. Leonard, P. M. Newman, R. J. Rikoski, J. Neira, and J. D. Tardós. Towards robust data association and feature modeling for concurrent mapping and localization. In *10th International Symposium of Robotics Research - ISRR 2001*, pages 5–11, Lorne, Victoria, Australia, Nov. 2001.

[Lipton, 1996] A. J. Lipton. *Vision Based Mobile Robot Navigation in Unknown Environments Using Natural Landmarks*. PhD thesis, Monash University, Melbourne, Australia, 1996.

[Lumelsky and Stepanov, 1987] V. J. Lumelsky and A. A. Stepanov. Path-planning strategies for a point mobile automaton moving amidst unknown obstacles of arbitrary shape. *Algorithmica*, 2(4):403–430, 1987.

[Luo and Kay, 1989] R. C. Luo and M. G. Kay. Multisensor integration and fusion in intelligent systems. *IEEE Transactions on Systems, Man, and Cybernetics*, 19(5):901–931, Sep.-Oct. 1989.

[Macedo *et al.*, 2000] J. Macedo, R. Manduchi, and L. Matthies. Ladar-based discrimination of grass from obstacles for autonomous navigation. In *Proceedings of the Seventh International Symposium on Experimental Robotics (ISER'00)*, Waikiki Beach, Dec. 2000.

[Macedo *et al.*, 2001] J. Macedo, R. Manduchi, and L. Matthies. Ladar-based terrain cover classification. In *Unmanned Ground Vehicle Technology III*, volume 4364, pages 274–280, Orlando, FL, 16-17 Apr. 2001. SPIE.

[Malone, 1978] R. Malone. *The Robot Book*. Push Pin Press, New York, 1978.

[Manduchi *et al.*, 2001] R. Manduchi, L. Matthies, and F. Pollara. From cross-country autonomous navigation to intelligent deep space communications: Visual sensor processing at JPL. In *11th International Conference on Image Analysis and Processing*, pages 472–477, Palermo, Italy, Sep. 2001.

[Mayora *et al.*, 1998] K. Mayora, I. Moreno, and G. Obieta. Perception system for navigation in a 3D outdoor environment. In *Proceedings of IAV'98 3rd IFAC Intelligent Autonomous Vehicles*, pages 381–385, Madrid, Spain, 25-27 Mar. 1998.

[Montano and Asensio, 1997] L. Montano and J. R. Asensio. Real-time robot navigation in unstructured environments using a 3D laser rangefinder. In *Proceedings of the 1997 IEEE/RSJ International Conference on Intelligent Robot and Systems*, pages 526–532, Grenoble, France, 7-11 Sep. 1997.

[Moravec and Elfes, 1985] H. P. Moravec and A. Elfes. High resolution maps from wide angle sonar. In *Proceedings of the IEEE Conference on Robotics and Automation*, pages 116–121, St. Louis, 1985.

[Moutarlier and Chatila, 1989] P. Moutarlier and R. Chatila. Stochastic multisensory data fusion for mobile robot location and environment modelling. In *Fifth International Symposium on Robotics Research*, pages 85–94, Tokyo, Japan, Aug. 1989.

[Murphy, 2000] R. R. Murphy. Marsupial and shape-shifting robots for urban search and rescue. *IEEE Intelligent Systems*, 15(2):14–19, Mar.-Apr. 2000.

[Murphy, 2004] R. R. Murphy. Trial by fire. *IEEE Robotics & Automation Magazine*, 11(3):50–61, Sep. 2004.

[Nguyen and Bott, 2000] H. G. Nguyen and J. P. Bott. Robotics for law enforcement: Applications beyond explosive ordnance disposal. In *Enabling Technologies for Law Enforcement and Security*, volume 4232, pages 433–454, Boston, MA, 5-8 Nov. 2000. SPIE.

[Nilsson, 1971] N. J. Nilsson. *Problem-Solving Methods in Artificial Intelligence.* McGraw-Hill, New York, 1971.

[Pedersen *et al.*, 2003] L. Pedersen, D. Kortenkamp, D. Wettergreen, and I. Nourbakshk. A survey of space robotics. In *Proc. 7th Int. Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS-03)*, Nara, Japan, 19-23 May 2003.

[RoboCupRescue, 2001] RoboCupRescue. Official website. http://www.rescuesystem.org/robocuprescue/, 2001+. Visited on 3/11/2004.

[Sagnac, 1913] G. M. Sagnac. L'ether lumineux demontre par l'effet du vent relatif d'ether dans un interferometre en rotation uniforme. *C.R. Academy of Science*, 95:708–710, 1913.

[Smith *et al.*, 1990] R. Smith, M. Self, and P. Cheeseman. Estimating uncertain spatial relationships in robotics. In I. J. Cox and G. T. Wilfong, editors, *Autonomous Robot Vehicles*, pages 167–193. Springer-Verlag, New York, 1990.

[Smuda *et al.*, 2002] W. Smuda, P. L. Muench, G. R. Gerhart, and K. L. Moore. Autonomy and manual operation in a small robotic system for under-vehicle inspections at security checkpoints. In *Unmanned Ground Vehicle Technology IV*, volume 4715, pages 1–12, Orlando, FL, 2-3 Apr. 2002. SPIE.

[Spero and Jarvis, 2002a] D. J. Spero and R. A. Jarvis. 3D vision for large-scale outdoor environments. In *Australasian Conf. Robotics and Automation*, pages 228–233, Auckland, New Zealand, 27-28 Nov. 2002.

[Spero and Jarvis, 2002b] D. J. Spero and R. A. Jarvis. Path planning for a mobile robot in a rough terrain environment. In *Third International Workshop on Robot Motion and Control*, pages 417–422, Bukowy Dworek, Poland, 9-11 Nov. 2002. IEEE.

[Spero and Jarvis, 2004a] D. J. Spero and R. A. Jarvis. On localising an unknown mobile robot. In 2^{nd} *Int. Conf. Autonomous Robots and Agents*, Palmerston North, New Zealand, 13-15 Dec. 2004. in press.

[Spero and Jarvis, 2004b] D. J. Spero and R. A. Jarvis. Towards exteroceptive based localisation. In *IEEE Conf. Robotics, Automation and Mechatronics*, Singapore, 1-3 Dec. 2004. in press.

[Stentz and Hebert, 1995] A. Stentz and M. Hebert. A complete navigation system for goal acquisition in unknown environments. In *Proc. IEEE/RSJ Int. Conf. Intelligent Robots and Systems*, pages 425–432, Pittsburgh, PA, 5-9 Aug. 1995.

[Stone, 1996] H. W. Stone. Mars pathfinder microrover: A low-cost, low-power spacecraft. In *Proc. 1996 AIAA Forum on Advanced Developments in Space Robotics*, Madison, WI, Aug. 1996.

[Thrun and Bücken, 1996] S. Thrun and A. Bücken. Integrating grid-based and topological maps for mobile robot navigation. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence (AAAI-96)*, volume 2, pages 944–950, Portland, OR, 4-8 Aug. 1996.

[Thrun *et al.*, 1998] S. Thrun, W. Burgard, and D. Fox. A probabilistic approach to concurrent mapping and localization for mobile robots. *Machine Learning*, 31(1-3):29–53, 1998.

[Thrun, 1998] S. Thrun. Learning metric-topological maps for indoor mobile robot navigation. *Artificial Intelligence*, 99(1):21–71, 1998.

[Williams *et al.*, 2000] S. B. Williams, P. Newman, G. Dissanayake, and H. Durrant-Whyte. Autonomous underwater simultaneous localisation and map building. In *Proceedings of the 2000 IEEE International Conference on Robotics and Automation*, volume 2, pages 1793–1798, San Francisco, CA, Apr. 2000.

[Zhang and Ghosh, 2000] L. Zhang and B. K. Ghosh. Geometric feature based $2\frac{1}{2}$D map building and planning with laser, sonar and tactile sensors. In *Proceedings of the 2000 IEEE/RSJ International Conference on Intelligent Robots and Systems*, volume 1, pages 115–120, Takamatsu, Japan, Oct. 2000.