Field Robots

Chuck Thorpe and Hugh Durrant-Whyte

Robotics Institute, Carnegie Mellon University, Pittsburgh USA; Australian Centre for Field Robotics, The University of Sydney, Sydney NSW 2006, Australia

Abstract

Field Robots are machines that work in unstructured environments, including under water, in mines, in forests and on farms, and in the air. These applications involve both advanced ideas in robotics and careful attention to engineering details. This paper discusses the overall challenges of field robotics and the current state of the art. The paper is not a thorough survey, nor a tutorial; but is instead intended to serve as a discussion starter for further design and development.

Introduction

Field Robots is concerned with the automation of vehicles and platforms operating in harsh, unstructured environments. Field robotics encompasses the automation of many land, sea and air platforms in applications such as mining, cargo handling, agriculture, underwater exploration and exploitation, highways, planetary exploration, coastal surveillance and rescue, for example. Field robotics is characterized by the application of the most advanced robotics principles in sensing, control, and reasoning in unstructured and unforgiving environments. The appeal of field robotics is that it is challenging science, involves the latest engineering and systems design principles, and offers the real prospect of robotic principles making a substantial economic and social contribution to many different application areas.

In general, field robots are mobile platforms that work outdoors, often producing forceful interactions with their environments, with no human supervision. Classic examples are backhoes that automatically load trucks, or automated bulldozers that clear fields. Other examples don't quite meet all the criteria, but are still considered field robots. Automated underground mining machines are not technically operating outdoors, and are operating in semistructured environments, but still share most of the technical challenges of other field robots. Robot cranes are not necessarily mobile in the usual sense, but still have to reach over wide areas, sense obstacles, deal with vagaries of outdoor illumination and wind gusts, and otherwise handle field robotics issues. Military scout robots are not (supposed to) perform forceful interaction, but do deal with highly irregular terrain.

A major impetus in field robotics is the comparative advantage provided by automation of large expensive manned machines (a mining haul vehicle or a container handling crane) over the automation of low cost domestic systems (a vacuum cleaner or product mover for example). A typical automation package may cost of the order of US$50-$100,000. While this is a substantial increase in cost for a typical product mover (AGV), it is quite modest in comparison to the $2M cost of a mining haul truck. Additionally, the relative pay-off in increased equipment utilization, reduced maintenance and fleet controllability is itself a major motivation for automation.

The last ten years have seen substantial real progress in development and implementation of field robotics systems. Indeed, in the next few years there is a very good chance that substantial commercially available field robotics systems will go into operation. Consequently, there is a high degree of excitement in the area about the future prospects and opportunities in field robotics.

In this paper, we try to take a broader view of the field than is offered by any single application. We structure our discussion around categories of difficulty: what technical problems are fairly well understood and can be used in applications? Which areas are partially understood, and are somewhat useable, but still require significant research? Finally, which areas are still difficult problems, where a solution can perhaps be kludged up but for which no solid engineering solutions are readily available? This paper is not meant to be a survey of the state of the art, but rather a discussion starter; we cannot hope to adequately cover the growing body of literature and increasing number of projects. (Major conferences that cover the topic include the Field and Service Robotics Conferences, as well as the International Symposium on Robotics Research). Instead, we concentrate on big emerging themes, and illustrate those areas with projects with which we are familiar.

Things we understand

Field robotics has made considerable progress in the past 10 years with a significant number of systems now deployed in both advanced research and near-commercial applications. Technically, field robotics has been able to make best progress in applications where the environment is cooperative, constrained, and, ideally, structured for automation. Typical of such applications are cargo container handling, haulage in mining, and high altitude unmanned air vehicles (UAVs).

Figure 1: Automated container hauler

Field robotics is as much about engineering as it is about developing basic technologies. While many of the methods employed derive from other robotics research areas, it is the application of these in large scale and demanding applications which distinguishes what can currently be achieved in field robotics. Broadly, field robotics technologies can be divided into three main areas; navigation and sensing, planning and control, and safety.

In the sensing and navigation area the problem of position determination (localization) has been the subject of considerable work. GPS has made the localization problem tractable in many field robotics applications. This is particularly true in environments where good views of the sky can be guaranteed (in surface mining, in UAV or autonomous ship systems for example). The development of low cost inertial sensing has also been a major advance for these applications. Combined GPS inertial systems are now in common use in many field robotics applications. In applications where GPS availability is poor, such as port cargo handling (the cranes get in the way!), or non-existent, such as underground mining or indoor vehicles, the localization problem is not so easily solved. In such situations there has been a common use of lasers and mm-wave radar to observe the relative location of artificial landmarks or beacons and, by referencing these to a map, to deduce position. The use of artificial landmarks together with appropriate position estimation algorithms is now well understood.

The use of natural landmarks for relative navigation is also reasonably well developed. Good examples include the use of vision to delineate crop lines in agricultural applications, the use of vision to maintain an automobile on a highway, or the use of lasers to determine approach vectors for vehicles in cargo and mining applications.

Figure 2: Automated Straddle Carrier

Over the past 10 years or so, we have become better at controlling the motion of large vehicles and of planning both small and large motions for these machines. At the beginning of the 1990's, speeds of a few meters per second were typical. Now, it is not uncommon for autonomous land vehicles to achieve speeds of 60-100Kmph. There are two basic reasons for this. First, processing of navigation, particularly visual, information has become much faster and can now easily achieve the required bandwidths necessary for high-speed vehicle control. Second, we have got much better at modeling the complexity of high-speed vehicle motion, particularly the effect of pneumatic tires or tracks, and the interaction of vehicles and rough ground. Many complex land vehicle control problems have now been solved. Planning motions of large vehicles has also been tackled with some success. At the simplest level, issues such as linear route planning and scheduling of vehicles has been implemented on a number of commercial field robotics systems. The planning of motions on rough 3D terrain has also received considerable attention although this is not yet a closed book. The planning of vehicle missions, in air or sub-sea vehicles for example, has also made progress although there is some way to go to make these methods generally applicable.

Safety is a major (and in some ways, unique) element of field robotics applications. Large vehicles traveling at high speed pose a significant health hazard. There are a number of elements to this problem. First is the issue of collision detection. Some quite advanced collision detection systems exist in advanced research or near commercial stages of development. These include mm-wave radar in automobile and mining applications, and lasers or radar in cargo handling or air-borne applications. A major requirement in field applications is long range (for stopping distance from high speed), all weather operation and reliability. It is the reliability issue which is hard to come to grips with. In a rough terrain environment it is not easy to obtain the sensitivity and discrimination necessary to distinguish a potential collision from a simple artifact of the environment. A second major safety issue is the problem of overall system integrity. The cost, size and speed of vehicles in many applications means that some guarantee must be made about the operation of the machine over its operating life. A 15 minute demo is not enough. The challenge is to ensure continuous operation 24 hours a day 365 days a year. For autonomous systems, the main issues are in sensor and software integrity. Safety is understood although there are still some very complex and challenging issues to overcome.

Given a reasonably structured, ideally open air environment that can be controlled to a degree, field robotics is now able to deliver operational systems.

Things we kind of understand

As the environment becomes less structured or the interactions between machine and environment become more complex, we understand less about what we are doing. Examples of such applications include autonomous reconnaissance in unstructured terrains, excavation in mining or construction, sub-sea navigation and low altitude UAVs.

Figure 3: Underwater robot for reef exploration

Probably the most complex and demanding research issue facing field robotics is the full 3D perception and understanding of typical unstructured environments. This may include the problem of land-vehicle terrain estimation from sensors such as vision and lasers, or equivalent problems with sonar underwater and with air-borne laser and radar systems. Such terrain models will be an essential component of functions such as navigation, path and mission planning in unstructured domains. While there is, and has been, a great deal of research currently being undertaken in these areas (outside and inside field robotics), the reliable construction and *understanding* of terrain models is still some way off. This is for two main reasons. First, we need a general representation of such environments capable of dealing with the intrinsic uncertainty (location, structure and data association) together with the broad (unpredictable, time varying, multi structured) types of terrain likely to be encountered by field robotics machines. A number of terrain construction methods have been employed with varying degrees of success. These include the use of grids, triangular tessellations, and probabilistic methods. Second, the problem of terrain understanding remains a major hole in current research effort. Terrain understanding really involves the manipulation, reasoning and dynamic use of terrain models in navigation and planning. While there has been some notable work (on grids for example), these fall a long way short of the necessary genericity required in typical field applications. To some degree the problem of terrain understanding is not even well posed until there is some better consensus on the methods required for terrain construction. While we understand what is needed, the necessary methods have yet to be developed.

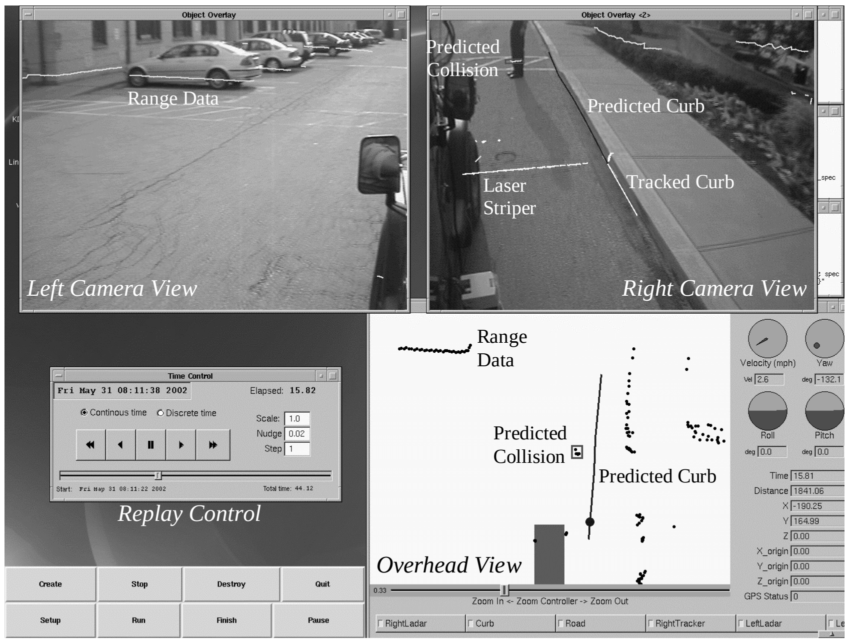

Replay of the integrated surround sensing system

Figure 3: Safe driving in cluttered environments using multi-sensor fusion

A second major issue that we have yet to come to grips with is how to plan and make robust decisions in environments with little structure or model. Fundamentally, the problem is concerned with making decisions in the presence of structural uncertainty or error; placing a foot on terrain of unknown geometry and composition, planning a path through an area with incomplete geometric or structural knowledge. This is the essence of working in field robotics where the environment can not be known and will not conveniently fall into some standard model. While there are well-paraded mathematical methods for formulating such problems it has been hard to see how these can reasonably be applied in even simple field environments. A demonstration in which a machine can autonomously make decisions such as these, and be equally effective in a variety of domains, is a long way from our current achievements. It may be that less traditional approaches (controller compositions, qualitative modeling and reasoning methods) will be a better way of dealing with this type of structural uncertainty. Again, there is a general understanding of the problem, but a concerted effort to formulate a method is still required. As always, substantive implementations are also necessary.

There are several other significant issues which the field robotics community has been addressing, but which still lack a basic understanding. Notable, is the general problem of contact control and environment modification as typified by autonomous excavation. Here, a large (very flexible, low bandwidth) manipulator comes into contact with a poorly known environment (variable geometry and material composition), then interacts and changes that environment in a manner which is unknown and can not often be predicted with any degree of accuracy. It is not at all clear how to deal with this problem in general. However, it is worth pointing out (maybe a lesson to be learned?) that the simple solution, so far, seems to work; just dig until certain forces and motions are detected, and carry on regardless.

The problem with lack of structure is at the heart of many field robotics applications. Once we have a grip on how to describe and model this lack of structure and once we can make decisions robust to these models, many new field applications will become solvable.

Things we mostly don't understand

The "grand challenge" problem for field robotics is:

> Create smart, reliable, mobile machines, capable of moving more capably than equivalent manned machines in a wide variety of unstructured environments.

Current field robots all fall well short of that goal, on at least one axis. Most field robots operate at a much lower performance than would be dictated by the raw hardware; usually, this is driven by limitations in perception and judgment rather than limitations in control. Humans are still much better than robots in driving over rough terrain, maneuvering through heavy traffic, detecting small obstacles by texture or color cues, and general visual recognition. Humans excel in "common sense"; a combination of general-purpose perception, reasoning about the behavior of other entities (specifically other people), planning in the face of uncertainty, and generating reliable and redundant systems.

A typical example is the electronic chauffeur of the future. Autonomous driving research has focused on perception of the road and of other vehicles, and accurate vehicle controls. The resulting vehicles have a wide variety of capabilities, for road following, headway maintenance, lane changing, and so forth. But moving from an impressive demo system to a truly functional vehicle will require significant further advances, notably in reliable perception and in modeling the behavior of other drivers.

Specific challenges for continued research in field robotics include the following, with an example task given for each challenge:

- General recognition: Perception systems are becoming adequate for well-defined tasks and well-defined objects. More general recognition is still difficult. Parts of the problem could be solved by better sensors or innovative combinations of sensing. Other parts of the problem sit at the intersection between sensing and world knowledge; the classical problem of recognizing a chair requires both shape measurement and reasoning about object use and social conventions. Example task: an automated car distinguishing between radar returns from another car vs. from a Mylar balloon blowing into the highway.
- Dealing with terrain: Perception works adequately for some indoor mobility tasks. Existing sensing systems are inadequate for driving over outdoor terrain. New sensors or sensing strategies need to be developed for detecting surface conditions (mud, ice, loose gravel), inferring partially-observable surfaces (rocks and ditches covered by grass), and estimating other surface properties (bearing strength of dirt, slip angle of sand). Example task: automated dune buggy deciding how to climb a steep hill.
- Building human interfaces: While some robots will be autonomous, many will work cooperatively with a human supervisor. There is a long and deep literature on pure teleoperation, where a human controls a robot directly for example via joystick. There is an emerging field of graphical interfaces, with good remote representations of current vehicle state. But there is very little work done on high-level interfaces between robots and humans. Users would like to generate commands at a variety of levels, including explicit joint commands; movement commands; safeguarded moves; subtasks; and task sequences. The robot in turn should be able to ask for clarification; report on problems and suggest alternatives; and understand what level of deviation is allowed from the commands. This work will require new results in dialogue modeling and planning systems, as well as more traditional robotic, graphics, and user interface studies. Example task: working with a lunar exploration robot over a link with time delays, in order to accomplish a prospecting task while safeguarding the robot.

- Dealing with humans: Robots operating in proximity to people need to be able to sense people and guarantee safe actions. Parts of this task involve sensing: which objects in the environment are people? How are they moving? Where are they looking? Parts of the task involve social conventions: will this person move to the right or the left when she meets me on the sidewalk? How likely is a driver to use a turn signal when changing lanes in front of an automated car? Example task: robot cleaning machine in a crowded grocery store, deciding when to go around people, when to wait for them to move, and when to politely ask them to clear a space.
- Legged systems: Rough terrain and narrow passages are best traversed by legged systems. Truly efficient and reliable legged locomotion requires fundamental advances in mechanisms, control, gait and footfall planning, and perception. Inspiration can come from biomimesis, but actual mechanisms may be different from biology. Example task: robot mule.
- Learning: Most mobile systems are brittle: if conditions change, or if they are asked to perform a task out of their design envelope, they are unable to cope. Systems of the future will need to use learning at many levels. Map learning is already being addressed. Example areas that need additional work include learning controls for extreme maneuvers, or changing controls as the system ages; learning behavior patterns of other agents; and learning new sensing strategies and encapsulated behaviors. Example task: adapting helicopter control as it picks up a load swaying in the wind.
- Cooperation: Until recently, most laboratories were fortunate to have a single functioning robot, so multi-robot research was infeasible. There is a large upsurge in multi-robot cooperative planning research, most visibly in the context of robot soccer. While some interesting performance has been achieved, there is still a great deal of work to be done on the fundamental principles of multi-agent cooperation. Subproblems include deducing another agent's plans from partial observation; problem decomposition without hierarchies or without communication; cooperative sensing; and cooperative mobility. Example task: multiple robot explorers in unknown terrain with difficult mobility, cooperating to examine all of the terrain in minimal time.

Conclusion

Field robotics is the most difficult, but most exciting, area of robotics. In terms of economic impact, field robots have arrived at a stage where they will soon be productive members of society, beginning to handle dirty and dangerous jobs and to be real contributors to the economy. In terms of scientific inquiry, robots that move through the real outdoor world, with all its unpredictability, offer a direct experimental testbed for developing truly intelligent machines. And in terms of public excitement (and our own motivation), it is fascinating to build machines that can extend our senses under the ocean, on the surface of Mars, or deep underground.

There are immediate applications, both in terms of complete systems and in terms of robotic components used in other applications. But there is also a rich set of open problems, in fundamental research as well as in applications, that will keep us all busy well into the future.

References

Third International Conference on Field and Service Robotics (FSR '01), A. Halme, ed., Helsinki University of Technology, Helsinki Finland, June 2001.

Ninth International Symposium on Robotics Research, October 1999, Snowbird Utah. Hollerbach and Koditschek, eds. Springer Verlag.

Second International Conference on Field and Service Robotics (FSR '99), J. Bares, ed., Carnegie Mellon University, August, 1999

Eighth International Symposium on Robotics Research, October 1997, Kamakura, Japan. Shirai and Hirose, eds. Springer Verlag.

First International Conference on Field and Service Robotics, A. Zelinsky, ed, Australian Robot Association, December, 1997

Seventh International Symposium on Robotics Research, October 1995, Herrsching, Germany. Giralt and Hirzinger, eds. Springer Verlag.