CHAPTER 1

Introduction

Figure 1.1 Science Fiction Becomes Fact. Many of the author's generation were introduced to robotics and space travel at the same time when the first *Star Wars* movie appeared in 1977. Little did we suspect that real robots of our own design would drive around on Mars in 1997 for the Pathfinder Mission—only 20 short years later.

Although robot arms that spot weld our cars together have been around for some time, a new class of robots, the mobile robot, has been quietly growing in significance and ability. For several decades now, behind the scenes in research laboratories throughout the world, robots have been evolving to move automatically from place to place. Mobility enables a new capacity to interact with humans while relieving us from jobs we would rather not do anyway.

Mobile robots have recently entered the public consciousness as a result of the spectacular success of the Mars rovers, television shows such as *Battlebots,* and the increasingly robotic toys that are becoming popular at this time.

Mobility of a robot changes everything. The mobile robot faces a different local environment every time it moves. It has the capacity to influence, and be influenced by, a much larger neighborhood than a stationary robot. More important, the world is a dangerous place, and it often cannot be engineered to suit the limitations of the robot, so mobility raises the needed intelligence level. Successfully coping with the different demands and risks of each place and each situation is a significant challenge for even biological systems.

1.1 Applications of Mobile Robots

Every situation in which an animal, human, or vehicle does useful work today is a potential application for a mobile robot. Generally, some of the reasons why it may be a good idea to automate are:

- Better. Manufacturers can improve product quality – perhaps because results are more consistent, easier to measure, or easier to control.
- Faster. Automation can be more productive than alternatives either due to increased rates of production, reduced downtime, or reduced consumption of resources.
- Safer. Sometimes the risk to humans is simply not justified when machines are a viable alterative.
- Cheaper. Using robots can reduce overheads. Robot maintenance costs can be much lower than the equivalent for man-driven vehicles.
- Access. Sometimes, humans cannot even exist at the scales or in the environments in question.

1.2 Types of Mobile Robots

We can classify mobile robots based on such dimensions as their physical characteristics and abilities, the environments for which they are designed, or perhaps the job that they do. Following are some examples of a few different classes of mobile robots.

Figure 1.2 Tug AGV (JBT Corporation, Philadelphia, USA). These laser guided vehicles are used in factories to move materials from place to place.

1.2.1 Automated Guided Vehicles (AGVs)

AGVs are designed to move materials (an application known as *material handling*) in factories, warehouses, and shipping areas in both indoor, and outdoor settings. They may convey automotive parts in manufacturing settings, newsprint in publishing companies, or waste in nuclear power plants.

Early vehicles had guidance systems based on sensing wires embedded in the floor whereas contemporary systems use laser triangulation systems, or inertial systems augmented by occasional magnetic landmarks in the floor.

It is typical for contemporary systems to employ wireless communications to link all vehicles to a central computer responsible for controlling traffic flow. Vehicles are further classified based on whether they pull trailers filled with material (tug-AGV), pick and drop it with forks (forked-AGV) or convey it on an platform on the top of the vehicle (unit load AGV).

Figure 1.3 Straddle Carrier. Used to move containers to and from ships, automated versions of these vehicles are perhaps the largest AGVs in use today.

AGVs are perhaps the most developed market for mobile robots. Companies exist to sell components and controls to many competing vehicle manufacturers, and vehicle manufacturers sometimes compete with each other to sell to value-added systems integrators who assemble a solution for a particular application. In addition to moving material, the loading and unloading of trucks, trains, ships, and planes are potential applications for future generations of vehicles.

1.2.2 Service Robots

Service robots perform tasks that would be considered service industry jobs if they were performed by humans. Some service tasks, like the delivery of mail, food, and medications, are considered to be "light" material handling, and are similar to the job of AGVs. Many service tasks, however, are distinguished by higher levels of intimacy with humans, ranging from coping with crowds to answering questions.

Medical service robots can be used to deliver food, water, medications, reading material, and so on to patients. They can also move biological samples and waste, medical records, and administrative reports from place to place in a hospital.

Surveillance robots are like automated security guards. In some cases, the automated ability to move through an area competently and to simply sense for intruders is valuable. This application was one of early interest to mobile robot manufacturers.

Figure 1.4 Health Care and Surveillance Service Robots. (Left) The Aethon Corp. "Tug" Vehicle is used to move food, linens, records, specimens, and biological waste in hospital settings. (Right) A surveillance robot like Robart might scan a warehouse for unwanted intruders on its regular rounds.

1.2.3 Cleaning and Lawn Care Robots

Other service robots include machines for institutional and home floor cleaning and lawn care. Cleaning robots are used in airports, supermarkets, shopping malls, factories, and so on. They perform such operations as washing, sweeping, vacuuming, carpet shampooing, and trash pickup.

These devices are concerned, not with getting somewhere, or carrying anything, but instead with getting everywhere at least once. They want to cover every part of a particular area of floor in order to clean it.

Figure 1.5 Floor and Lawn Care Service Robots. These kinds of mobile robots care about area coverage. They try to "visit" every place in some predefined area.

1.2.4 Social Robots

Social robots are service robots that are specifically designed to interact with humans and often their main purpose is to convey information or to entertain. Although stationary information kiosks convey information, social robots require mobility for one reason or another.

Figure 1.6 Entertainment and Tour Guide Robots. (Left, Center) The SONY QRIO and AIBO robots dance and play, respectively. (Right) The EPFL tour guide moves from one station to another, often surrounded by people, and it describes museum exhibits.

Some potential applications include answering product location questions in a retail store (grocery, hardware). A robot that delivers hamburgers to kids in a restaurant would be fun. Robot assistants for elderly and infirm individuals could help their owners see (robot seeing-eye dog), move, or remember their medication.

In recent years, SONY Corporation has produced and marketed some impressive robots intended to entertain their owners. The earliest such devices were packaged as "pets." Automated tour guides in museums and expositions can guide customers through a particular set of exhibits.

1.2.5 Field Robots

Field robots perform tasks in the highly challenging "field" conditions of outdoor natural terrain. Almost any type of vehicle that must move about and do useful work in an outdoor setting is a potential candidate for automation. Most things are harder to do outdoors. It's difficult to see in bad weather and it's difficult to decide how to move through complicated natural terrains. It's easy to get stuck, too.

Figure 1.7 Field Robots. Field robots must engage the world exactly as it exists. (Left) Semi-automated fellerbunchers similar to this one have been designed to gather trees. (Right) Automated excavators have been prototyped for mass excavation applications—where large amounts of dirt are loaded over short time periods.

Vehicles that do real work look the way they do for good reasons. Hence, field robots look a lot like their human-driven equivalents. Field robots are often of the form of arms and/or tools (called *implements* in general) mounted on a mobile base. As such, they exemplify a more general case of a mobile robot that not only goes somewhere but also that interacts physically with the environment in some useful way.

In agriculture, real and potential applications include planting, weeding, chemical (herbicide, pesticide, fertilizer) applications, pruning, harvesting and picking fruit and vegetables. In contrast to household grass mowing, large scale mowing is necessary in parks, and on golf courses, and highway medians. The specialized man-driven vehicles used in mowing are good candidates for automation. In forestry, tending nurseries and the harvesting of full grown trees are potential applications.

There are diverse applications in mining, and excavation. Above ground, excavators, loaders, and rock trucks have been automated in open pit mines. Underground, drills, bolting machines, continuous miners, and load-haul-dump (LHD) vehicles have been automated.

1.2.6 Inspection, Reconnaissance, Surveillance, and Exploration Robots

Inspection, reconnaissance, surveillance and explorations robots are field robots that deploy instruments from a mobile platform in order to inspect an area or find or detect something in an area. Often, the best justification for a robot is that the environment is too dangerous to risk using humans to do the job. Clear examples of such environments include areas subject to high radiation levels (deep inside nuclear power plants), certain military and police scenarios (reconnaissance, bomb disposal), and space exploration.

In the energy sector, robots have been deployed to inspect components of nuclear reactors including steam generators, calandria, and waste storage tanks. Robots to inspect high tension power lines, and gas and oil pipelines have been prototyped or deployed. Remotely piloted undersea vehicles are becoming increasingly more autonomous and they and have been used to inspect oil rigs, communications cables on the seabed and even to help find shipwrecks like that of the *Titanic*.

Research into developing the robotic soldier has become particularly intense in recent years. Robotic vehicles are being considered for such missions as reconnaissance and surveillance, troop resupply, minefield mapping and clearing, and ambulance services. Manufacturers of military vehicles are already working hard to get a diverse array of robotic technologies into their products. Bomb disposal is already an established niche market.

In space, several robotic vehicles have now driven autonomously for kilometers over the surface of Mars and the concept of vehicles that maneuver around space stations under thruster power has been on the drawing table for some time.

Figure 1.8 Exploration Robots. (Left) The military robot is intended to explore environments where it may be too unsafe for soldiers. (Right) The Mars Science Laboratory searches for signs of life in the hostile environment on Mars.

1.3 Mobile Robot Engineering

1.3.1 Mobile Robot Subsystems

There are a host of challenges that become instantly obvious when attempting to construct systems exhibiting autonomous mobility. At the lowest level in a conceptual hierarchy of abilities, robots require the capacity of automatic *control*. Doing this involves the sensing of the states of actuators such as steering, speed, or wheel velocities and the precision application of power to those actuators to cause them to exert the correct forces. However, moving around competently requires more than an accelerator, steering and engine that do what they are told – there needs to be a driver. This book is mostly about the construction of such a driver.

Often, the objective is one of *navigation*, to move somewhere in particular or to follow a particular path. Accomplishing that requires a vehicle *state estimation* system that knows where the vehicle is at any time along the way. Navigation and control give a robot the capacity to drive (albeit blindly) from place to place, provided there is nothing in the way – but what if there is something in the way? The need for a capacity to understand the immediate surroundings, known as *perception*, arises in many different contexts ranging from following the road in view, to dodging a fallen tree, to recognizing the object for which the robot has been searching.

Although such an understanding of the environment is one aspect of intelligence, *planning* is another. Planning involves a capacity to predict the consequences of alternative possible courses of action, and a capacity to select the most appropriate one for the present situation. Planning often requires models of both the environment and the vehicle in order to predict how they interact. The former models are called *maps* and they are used to help decide where to go.

Maps may be produced externally to the robot or they may be produced by an onboard *mapping* system that uses the navigation system and the perception system to record salient aspects of what has been seen in its correct relative place. Maps can have other uses besides support of planning. If maps record both what is seen and where it is seen, it becomes possible for the robot to determine its position in the map if it can match what it does see to what it would see if it was in a particular hypothesized position.

1.3.2 Overview of the Text

This book will discuss many aspects of each of these subsystems in the order in which they might be developed, integrated, and tested during the construction of a prototype robot.

Preliminary material in mathematics (Chapter 2), numerical methods (Chapter 3), and physical models (Chapter 4) is presented first and the rest of the text will rely on it heavily. Thereafter incremental aspects of mathematics, models, and methods will be introduced as needed. Chapter 5 presents certain aspects of probability and then a more advanced topic that has special importance in mobile robots – optimal estimation. The problems of knowing where you are and of knowing what is out there are solved more effectively using optimal estimation techniques. State estimation (Chapter 6) is presented next because it produces the most basic feedback necessary for control of mobility. It is also a good introduction to some of the issues that are most peculiar to mobile robots. As we will see, moving is often not so hard compared to the problem of knowing precisely how you have moved or where you are.

The topic of control system design and analysis (Chapter 7) is presented next. Here are the basic techniques that are used to make things move in a controlled and deliberate fashion. Control makes it possible to move at a particular speed, with a particular curvature, or to a particular goal position and heading. It also becomes possible to move any articulations to point sensors, to dig a hole, or to pick up something. When Perception (Chapter 8) is added to this basic moving platform, things get very interesting quickly. Suddenly the mobile system potentially becomes able to look for things, recognize and follow or avoid people, and generally move much more competently in the local area. Such a system can also construct maps on small scales, and plan its motions in those maps based on the results of state estimation.

The topics of Localization and Mapping (Chapter 9) are two sides of the problem of locating objects relative to the robot or locating the robot relative to a map of objects. Once the robot can see objects, it can remember where they are and the associated map can be very useful the next time this or any other robot inhabits the area. The final chapter presents Motion Planning (Chapter 10). When maps are available on a large scale, whether produced by a robot or a human, it becomes possible to do some very intelligent things. Given an accurate map, robots can rapidly compute the best way to get anywhere and they can update that map rapidly as new information is gathered on the move.

1.3.2.1 Layers of Autonomy

For mobile robots, and for us, problems that require deep thought cannot be solved in an instant. There are other problems like avoiding the jaywalking pedestrian, which must be solved in an instant. Yet, it is difficult to be both smart and fast so robots tend to use a hierarchy of perceive-think-act loops in order to allocate resources optimally (Figure 1.9). Higher levels tend to be more abstract and deliberative, whereas lower levels tend to be more quantitative and reactive.

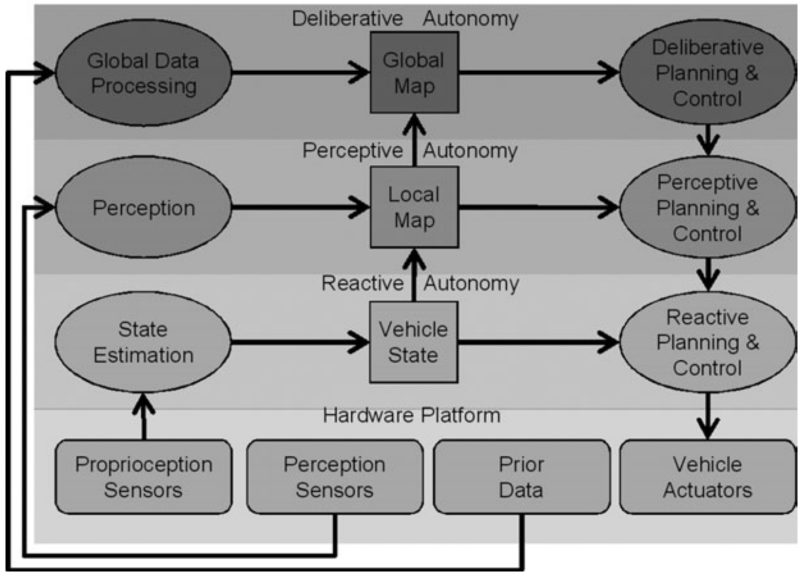

Figure 1.9 Layers of Autonomy. The entire mobile system can be described in terms of three nested perceive-think-act loops.

1.3.2.1.1 Reactive Autonomy. This layer is responsible for controlling the motion of the vehicle with respect to the environment and any required articulations. It typically requires feedback only of the articulation and motion state (position, heading, attitude, velocity) of the vehicle. The content of the book up to the middle of the Control chapter fits in this layer.

1.3.2.1.2 Perceptive Autonomy. This layer is responsible for responding to the immediately perceivable environment. It typically requires feedback of the state of the environment, which is derived from Perception. This layer requires estimates only of short-term relative motion and it tends to use environment models that are valid only locally. Prediction is limited to a few seconds into the future. The content of the book up to the Perception chapter fits in this layer.

1.3.2.1.3 Deliberative Autonomy. This layer is responsible for achieving longer term goals, sometimes called the “mission.” This layer requires earth-fixed position estimates and it tends to use environment models that extend over large areas. Prediction may extend arbitrarily far into the future. The content of the book up to Motion Planning chapter fits in this layer.

1.3.3 Fundamentals of Wheeled Mobile Robots

At this point in the history of the field there is no agreed-on list of fundamental concepts that constitute the core a of mobile robotics curriculum. Whereas homogeneous transforms take center stage in any text on robot manipulators, mobility changes everything – even the governing mathematics. This text takes the perspective that the most basic model of a wheeled mobile robot is a kinematic differential equation that is constrained and nonlinear. Most of the content of the text follows from that perspective.

The easiest way to formulate wheeled mobile robot (WMR) models is write the velocity kinematics and to do so in the body frame because the actuators, as well as the sensors, move with the robot. Fluency with ideas of moving coordinate systems makes the velocity kinematics very straightforward to derive in the general case. Once these ideas are mastered, they also play a role in mechanical dynamics, inertial navigation and stability control. However, in moving from the plane to three dimensions, the rate kinematics also get more complicated. In many cases, the basic WMR model requires explicit constraints for treatment of both rolling without slipping and terrain following. At this point, our model has become a differential algebraic system and we will cover both velocity and force driven models.

However, although the rate kinematics of WMRs thus produced are straightforward, the rotation matrix that appears to convert coordinates to the world frame makes the integrand nonlinear, and in a very bad way. The trig functions of orientation states make the problem of driving to a particular place – arguably the most basic problem of all – more formidable than the famous Fresnel integrals, and therefore not solvable in closed form. The elegance of manipulator inverse kinematics is simply not achievable for mobile robots and a host of related problems cannot be solved in closed form – even in the plane. Thus, any time we want to write a relationship between state (pose) and inputs, we have to be content to write an integral, so the text has a lot of integrals in it.

Then computers come to the rescue. The book covers numerical methods because a large number of important problems succumb to a short list of basic numerical methods. Some effort has been expended here to reduce every robotics problem into its associated canonical "method." Nonlinear least squares, for example, is fundamental to both estimation and control and the Kalman filter itself is derived here from weighted least squares. Likewise, the Newton iteration of constrained optimization applies to problems as diverse as terrain following, trajectory generation, consistent map building, dynamic simulation, and optimal control. As for that most basic problem of all? It is a numerical rootfinding problem that Newton's method handles readily once the problem is parameterized. Thus, although we cannot write elegant formulae for fundamental reasons, we can still get the job done in a computer.

Once the basic WMR model is in place we will integrate it, parameterize it, perturb it, take its squared expected value, invert it etc. in multiple ways in order to supply control and estimation systems with the models they need. They need them to provide the feedback and perform the actions necessary for a robot to know where it is and get where it is going. All of this is illustrated using essentially the same basic model of a robot actuated in linear and angular velocity.

Along the way the tools of linear systems theory prove to be powerful. In addition to providing the means to control the robot, these tools allow us to understand the propagation of systematic and stochastic error and therefore to understand the behavior of dead reckoning and to even calibrate the system while it is in operation.