Cambridge University Press 978-0-521-87157-0 - Computational Principles of Mobile Robotics, Second Edition Gregory Dudek and Michael Jenkin Excerpt More information

1

Overview and motivation

"... let's start with the three fundamental Rules of Robotics – the three rules that are built most deeply into a robot's positronic brain." In the darkness, his gloved fingers ticked off each point.

"We have: one, a robot may not injure a human being, or through inaction, allow a human being to come to harm."

"Right!"

"Two," continued Powell, "a robot must obey the orders given it by human beings except where such orders would conflict with the First Law."

"Right!"

"And three, a robot must protect its own existence as long as such protection does not conflict with the First or Second Laws."¹

Powell and Donovan discuss the laws of robotics.

The ability to navigate purposefully is fundamental to most animals and to every intelligent organism. In this book we examine the computational issues specific to the creation of machines that move intelligently in their environment. From the earliest modern speculation regarding the creation of autonomous robots, it was recognized that regardless of the mechanisms used to move the robot around or the methods used to sense the environment, the computational principles that govern the robot are of paramount importance. As Powell and Donovan discovered in Isaac Asimov's story "Runaround," subtle definitions within the programs that control a robot can lead to significant changes in the robot's overall behavior or action. Moreover, the interactions among multiple complex components can lead to large-scale emergent behaviors that may be hard to predict.

Mobile robotics is a research area that deals with the control of autonomous and semi-autonomous vehicles. What sets **mobile robotics** apart from other research areas such as conventional manipulator robotics, artificial intelligence, and computer vision is the emphasis on problems related to the understanding of **large-scale space**, that is, regions of space substantially larger than those that can be observed from a single vantage point. While at first blush the distinction between sensing in large-scale space, with its requirement for mobility, and local sensing may appear obscure, it has far-reaching implications. To behave intelligently in a large-scale environment not only implies dealing with the incremental acquisition of knowledge, the estimation of positional error, the

¹ I. Asimov, "Runaround" [37]. Reprinted by permission of the Estate of Isaac Asimov c/o Ralph M. Vicinanza, Ltd.

Cambridge University Press 978-0-521-87157-0 - Computational Principles of Mobile Robotics, Second Edition Gregory Dudek and Michael Jenkin Excerpt More information

2 **Overview and Motivation**

ability to recognize important or familiar objects or places, and real-time response, but it also requires that all of these functionalities be exhibited in concert. This issue of extended space influences all of mobile robotics; the tasks of moving through space, sensing about space, and reasoning about space are fundamental problems within the study of mobile robotics. The study of mobile robots in general, and this volume in particular, can be decomposed into the study of these three subproblems.

Mobile robots are not only a collection of algorithms for sensing, reasoning, and moving about space, they are also physical embodiments of these algorithms and ideas that must cope with all of the vagaries of the real world. As such, mobile robots provide a reality check for theoretical concepts and algorithms. They are the point where literally the "rubber meets the road" for many algorithms in path planning, knowledge representation, sensing, and reasoning.

In the context of humanity's ongoing quest to construct more capable machines – machines that match or even surpass human capabilities – the development of systems that exhibit mobility is a key hurdle. The importance of spatial mobility can be appreciated by observing that there are very few sophisticated biological organisms that cannot move or accomplish spatially distributed tasks in their environment. Just as the development of the wheel (and hence wheeled vehicles) marked a turning point in the evolution of manually operated tools, the development of mobile robots is an important stepping stone in the development of sophisticated machines.

Many different terms have come to be applied to the field of autonomous systems or mobile robotics. The words **autonomous**, as in autonomous system, and **automaton** have their roots in the Greek for *self-willed* (*auto+matos:* $\alpha \upsilon \tau \sigma \mu \alpha \tau \sigma \varsigma$). The term **robot** itself was introduced by Karel Čapek in his 1923 play *R.U.R.* (R.U.R. stands for Rossum's Universal Robots). The word 'robot' is derived from the Czech or Polish words 'robota,' meaning 'labour,' and 'robotnik,' meaning 'workman.' It is interesting to note that the word *automaton* implies a degree of self-will that is not conveyed by the term *robot*, and that **autonomous robot** might be construed as self-contradictory.

Robots that are manufactured following the same general structure as humans are known as **anthropomorphic robots** or **humanoid robots**, and in fiction, robots that are indistinguishable from humans are sometimes known as **androids**.

Although androids are beyond today's technology, anthropomorphic robots and robots with anthropomorphic features are quite common. There are many reasons why researchers develop robots in an anthropomorphic mold. In addition to a desire to develop an agent in "one's own image," there are practical reasons for developing systems with anthropomorphic features. The operating environment for many mobile robots is the same environment that humans inhabit, and we have adapted our environment to suit our performance specifications. By mimicking human structures, at least at an operational or functional level, a robot may be better suited to operate in our environment. Human physiology, perception, and cognitive processes have been studied extensively. Thus by using locomotive, sensing, and reasoning systems based on biological models, roboticists can exploit the extensive literature that already exists in these fields. In addition, people seem to have a fascination with human-looking robots that goes beyond the pragmatic. That being said, mobile robots are not limited to mimicking existing biological systems, and there exist many other mechanisms, from infrared sensors to alternative drive mechanisms, that can be exploited in the design of a mobile robot.

3

The study of mobile robots is an intrinsically interdisciplinary research area that involves:

Mechanical engineering: vehicle design and in particular locomotive mechanisms.

Computer science: representations and sensing and planning algorithms.

Electrical engineering: system integration, sensors, and communications.

Cognitive psychology, perception, and neuroscience: insights into how biological organisms solve similar problems.

Mechatronics: the combination of mechanical engineering with computer science, computer engineering, and/or electrical engineering.

Although many classes of the mobile robot systems currently in operation are fundamentally research vehicles and are thus experimental in nature, a substantial number of mobile robot systems are deployed in domestic or industrial settings. Real applications in which current mobile robots have been deployed successfully are characterized by one or more of the following attributes: the absence of an on-site human operator (often due to inaccessibility), a potentially high cost, long duty cycles, and the need to tolerate environmental conditions that might not be acceptable to a human. As such, mobile robots are especially well suited for tasks that exhibit one or more of the following characteristics:

- An environment that is inhospitable, so that sending a human being is either very costly or very dangerous.
- An environment that is remote, so that sending a human operator is too difficult or takes too long. Extreme instances are domains that are completely inaccessible to humans, such as microscopic environments.
- A task with a very demanding duty cycle or a very high fatigue factor.
- A task that is highly disagreeable to a human.

Successful industrial applications for mobile robots typically involve more than one of these characteristics. Consider the application of mobile robotics to underground mining as an example. The environment is dangerous, in that the possibility of rock fall or environmental contamination due to the release of hazardous gas or dust is quite real. The environment is remote, in that humans operating in underground mines must travel considerable distances, typically many kilometers, in order to reach the rock face being worked. At the rock face, the miner is confronted with an operational environment that can be cramped, poorly illuminated, hot, and dangerous. Other 'ideal' robotic operational environments include nuclear, extraterrestrial, and underwater environments.

Mobile robots are feats of engineering. The actuators, processors, user interfaces, sensors, and communication mechanisms that permit a mobile robot to operate must be integrated so as to permit the entire system to function as a complete whole. The physical structure of a mobile robot is complex, requiring a considerable investment of both human and financial resources in order to keep it operating. "Robot wranglers"² are an essential component for the successful operation of any robotic system. Thus, one of the goals of this book, in addition to provoking new research, is to act as a reference of mobile robot

² Graduate students and technicians.

tools and techniques for those who would develop or maintain a mobile robot. Rather than concentrate strictly on the sensors required for a mobile robot [204] or on the physical design of small autonomous robots [310] or collect the seminal papers of the field [143], this volume considers the computational processes involved in making a robot sense, reason, and move through its environment.

1.1 From Mechanisms to Computation

Robots can be considered from several different perspectives. At a physical, hardware, or mechanistic level, robots can be decomposed into the following:

- A power source, typically based on batteries.
- A mechanism for making the robot move through its environment the physical organization of motors, belts, and gears that is necessary to make the robot move.
- A computer or collection of computers that controls the robot.
- A collection of sensors with which the robot gathers information concerning its environment.
- Communications hardware to enable the robot to communicate to an off-board operator and any externally based computers.

At the device level, the hardware details can be abstracted, and a robot can be considered as follows:

- A software-level abstraction of the motors, encoders, and motor driver boards that allow the robot to move. Most mobile robot hardware manufacturers provide support for the underlying hardware at this level rather than force the user to deal with the details of actually turning motors.
- Software-level mechanisms or libraries to provide access to the robot's sensors, for example, the current image obtained by a video camera as an array of intensities.
- A standard communications mechanism, such as a serial interface or network access to the outside world.

From a still more abstract perspective, we can consider mobile robots at a purely computational level such that the sensors, communications, and locomotive systems are seen simply as software modules that enable the robot to interact with its environment. Typical components in a software architecture include the following:

- A motion control subsystem,
- A sensor control subsystem,
- A sensor interpretation subsystem.
- A mission control subsystem.

Even higher levels of abstraction exist. The term **cognitive robotics** is used to refer to the use of Artificial Intelligence (AI) techniques within a mobile robot and often assumes the existence of an idealized computational abstraction of the robot.

Historical Context

5

1.2 Historical Context

Autonomous Robots in Fiction

Thou shalt not make a machine in the likeness of a human mind.³

Autonomous devices have a long and checkered past in legend and literature. From ancient legends to modern films and literature, many different robots and robot-like devices have been constructed to extend the will of their creator or owner. Much of the fictional literature on autonomous systems is cautionary in nature: the 'robot' may follow its instructions too literally, or it may grow to have a will of its own and not follow its instructions at all. For example, in Isaac Asimov's story "Runaway" a robot is told to "get lost," which of course it does, while "Robots of Empire" and "Robots of Dawn," also by Asimov, describe the process of robots evolving their own rules of operation. Given their supposed infallibility, fictional robots have also been proposed as final arbitrators of judgment. In the 1951 film *The Day the Earth Stood Still*, Gort is a universal policeman who enforces the law without being influenced by sentiment.

Perhaps the earliest reference to a 'robot' in literature can be found in Greek mythology. According to ancient Greek or Cretan mythology, Talos was an 'animated' giant man made of bronze who guarded the island of Crete. Talos guarded the island and enforced the law. One of Talos' flaws was that he was too literal minded in the interpretation of his directives, so that he became a burden. Even in this legend, problem specification and representation was an issue! This notion of the robot as protector also appears in Jewish folklore. According to legend, in sixteenth-century Prague, the Jewish population turned to a Golem to protect them from the gentiles who wanted to kill them. A rabbi fashioned the Golem out of clay and breathed life into it.

Clay and bronze are not the only potential building materials for fictional 'robots.' In works of fiction, autonomous agents are also constructed out of biological components. In 1818, Mary Shelley wrote *Frankenstein*, which tells the story of Dr. Frankenstein and his efforts to animate dead tissue. As one inspired job advertisement put it, "Dr. Frankenstein was more than just a scientist – he was an electrical engineer with the creative capability for bringing extraordinary ideas to life." Nor are all fictional accounts of robots based on anthropomorphic designs. In his 1880 *The Demon of Cawnpore*, Jules Verne describes a steam-powered elephant,⁴ whereas more recently the film *Blade Runner* (based on Philip K. Dick's novel *Do Androids Dream of Electric Sheep?*) [166] describes a world in which animals are almost extinct and robotic pets are popular.

Isaac Asimov is often regarded as a key contributor to the genesis of robotics due to his copious science fiction writings on the topic and, most notably, his introduction of the "three laws of robotics." Introduced in 1942 in "Runaround" and reprinted at the beginning of this chapter, they are as follows:

1. A robot may not injure a human being, or, through inaction, allow a human being to come to harm.

³ F. Herbert, *Dune* [267].

⁴ The Demon of Cawnpore was also published as The End of Nana Sahib.

- 2. A robot must obey the orders given it by human beings except when such orders would conflict with the First Law.
- 3. A robot must protect its own existence as long as such protection does not conflict with the First or Second Law.

In later works, Asimov added a zeroth law that required a robot not to injure humanity. Many of Asimov's stories center around robot (and human) attempts to find new definitions or loopholes in these laws. Although the relevance of these laws to real robotics research is questionable, they nevertheless have proven to be both inspirational and provocative.

Since the 1940s, mobile robots have become a common feature of science fiction literature and film. Famous fictional robots include Robbie (*Forbidden Planet*), Gort (*The Day the Earth Stood Still*), Rosie (*The Jetsons*), Robot (*Lost in Space*), Floyd (*Stationfall* and *Planetfall*), R2D2 and C3PO (*Star Wars*), Data and the partly biological Borg (*Star Trek*), HAL (2001 and 2010), Bender (*Futurama*), the Terminator (*Terminator*), and of course Marvin, the paranoid android (*The Hitchhiker's Guide to the Galaxy*). More details on the evolution of robots in literature can be found in [34]. (See also [257].) It is interesting to note that fictional robots usually do not suffer from the computational, sensing, power, or locomotive problems that plague real robots. How the Daleks (see Figure 1.1) from the long-running BBC television series *Doctor Who* managed to conquer most of the galaxy without having to navigate a set of stairs was only finally addressed in a recent episode. On the other hand, fiction serves not only to predict the future, but also to inspire those who might create it. Stork [585] provides some insights into the differences between a specific fictional autonomous system – HAL from 2001 – and the current state of the art in terms of real systems.

Early Autonomous Robots

Various robotic or robotic-like systems can be found scattered throughout history. Mechanical robotic systems can be traced back to Greek and Roman times. The Roman historian Appian reported a mechanical simulation of Julius Caesar. In the fifteenth century, Leonardo da Vinci developed a number of robotic systems, perhaps the most famous of which was an anthropomorphic device with controllable arms and legs [543]. Less well



Figure 1.1. A Dalek, a half-robot/half-biological creature from the BBC TV series *Doctor Who*. Copyright Barry Angel. Used with permission.

Historical Context



(a) Tesla's robot

(b) Walter's robot

Figure 1.2. Analog robots. (a) Reprinted from M. Cheney, *Tesla: Man Out of Time*, Prentice-Hall [118]. Used with permission. (b) Copyright Owen Holland. Used with permission.

known perhaps is Leonardo's robotic car, or automata. This device resembled a differential drive robot and used springs and cams to program the robot to follow different trajectories as it moved. (See [543] for an in-depth review of the design of this vehicle). The late eighteenth-century autonomous Japanese tea carrier Karakuri is described in [541].

Autonomous vehicles built by Nikola Tesla in the 1890s are probably the earliest electrical mobile robots. In the 1890s, Tesla built wireless, radio-controlled vehicles [118]. One of his remote-controlled aquatic vehicles is shown in Figure 1.2a. The first steps towards modern electronic robotic systems were made during the early to mid 1940s. Norbert Wiener is considered the inventor of **cybernetics** and hence modern robotics. A mathematician, Wiener studied regulatory systems and their application to control. During World War II, he was involved in a project to develop a controlling device for an anti-aircraft gun. The development of such a device, which integrates sensory information (radar) via processing (simple control laws executing on an analog computer) into action (directing and firing the anti-aircraft gun), resulted in one of the first robotic systems. As Wiener mused in his January 1949 article in *Electronics*,

It has long been clear to me that the modern ultra-rapid computing machine was in principle an ideal central nervous system to an apparatus for automatic control; and its input and output need not be in the form of numbers or diagrams, but might very well be, respectively, the readings of artificial sensors such as photoelectric cells or thermometers, and the performance of motors or solenoids.⁵

⁵ Electronics, January 1949. Reprinted in [402].



Figure 1.3. V1 flying bomb. Currently hanging in the Imperial War Museum, London.

At the same time that Wiener was developing an automatic anti-aircraft gun, work in Germany on the V1 and V2 – autonomous aircraft and self-guided rocketry – was establishing the basis for autonomous vehicle design (see Figure 1.3). The V1 and V2 were known as **Vergeltungswaffen** (reprisal weapons). The V1 was equipped with simple sensors to measure distance travelled (a propeller in the nose), altitude, and heading. This was sufficient to permit the device to be launched from France and Holland and to strike at population centers in England. Roughly 8000 were launched [526].

W. Grey Walter built one of the earliest fully autonomous vehicles. Described in a series of articles published in 1950 and 1951 in *Scientific American* and in his book *The Living Brain* [635], Walter's electronic turtle (see Figure 1.2b) had phototube eyes, microphone ears, contact-switch feelers, and capacitors used as memory devices to perform associations. Walter named the robot Tortoise after the creature in *Alice in Wonderland*. The Tortoise performed tasks such as heading towards well-lit regions, locating the recharging hutch, and wandering without mishap.

With the development of digital computers came the potential for more complex mobile robots. Between 1966 and 1972, Nils Nilssen, Charles Rosen, and other researchers at the Stanford Research Institute developed Shakey, the first mobile robot to be operated using artificial intelligence techniques [461]. The 5-foot-tall robot used two stepper motors in a differential drive arrangement to provide locomotion and was equipped with touch-sensitive bumpers. An optical rangefinder and vidicon television camera with controllable focus and iris were mounted on a tilt platform for sensing. Off-board communication was provided via two radio channels – one for video and the other providing command and control. Shakey is shown in Figure 1.4a.

Work on Shakey concentrated on automated reasoning and planning, which used logicbased problem solving based on **STRIPS** – the STanford Research Institute Problem Solver. The control of movement and the interpretation of sensory data were secondary to this logic-based component. Simple video processing was used to obtain local information about empty floor space, and Shakey constructed a global map of its environment based on this information. A typical mission for Shakey was to find a box of a given size, shape, and color in one of a specified number of rooms and then to move it to a designated position.

Historical Context



(a) Shakey

(b) CMU Rover



(c) Stanford Cart

Figure 1.4. Early wheeled digital robots. (a) Reprinted from D. Stork (Ed.), *HAL's Legacy* [585], MIT Press, 1997. Used with permission. (b) and (c) Reprinted from H. P. Moravec, The Stanford Cart and the CMU Rover [433]. Used with permission.

Being able to accomplish these tasks depended on a simplified environment containing simple wooden blocks in carefully constrained shapes. Shakey had to cope with obstacles and plan actions using a total of 192K of memory (eventually upgraded to 1.35 MB).

The **Stanford Cart** [431–433] (see Figure 1.4c) was developed at SAIL (the Stanford Artificial Intelligence Laboratory) between 1973 and 1979 and moved to Carnegie Mellon University (CMU) in 1980. Throughout this period it underwent major modifications and served as the initial test device upon which solutions to a number of classic robot problems

CAMBRIDGE

10 **Overview and Motivation**



Figure 1.5. The GE Quadruped. Reprinted from Song and Waldron, *Machines That Walk* [575], MIT Press, 1989. Used with permission.

were developed. The Stanford Cart relied on stereo vision in order to locate objects and planned paths to avoid sensed obstacles using a world model based on stereo data. The stereopsis algorithm was based on a single camera that was mounted on a sliding track that was perpendicular to the camera's optical axis. A single 'view' of the environment was based on nine images taken at different positions along this track. A comparison of the images over time was used to determine the motion of the cart, whereas comparisons of the images from a single position were used to build an environmental model. The robot was controlled by an off-board computer program that drove the cart through cluttered spaces. The cart moved roughly 1 m every 10 to 15 min.

The kinematic structure of the Stanford Cart introduced a number of limitations in the robot. Recognizing these limitations, the CMU Rover project (started in 1980) developed a robot that relied on a synchronous drive-like assembly rather than the car-like steering of the Stanford Cart. The Rover added infrared and sonar proximity sensors to the robot and modified the camera mount for the video sensor so that it could pan and tilt as well as slide, which were not possible with the Stanford Cart.

Another early robot system, the **Hilare** project and robot family developed at Laboratoire & Analyse et d'Architecture des Systèmes (LAAS) in France [93, 243], also represented a milestone in performance. Hilare I, developed in 1977, was an indoor mobile robot based on a differential drive system (two powered wheels and one free wheel for balance) and included a laser rangefinder. Hilare's perceptual system relied on sonar units, a video camera, and a laserrange finder. The laser and camera were mounted on a pan-and-tilt station in order to direct the sensor in different directions.

In parallel with these early wheeled mobile robots, legged robotic systems began to appear in the 1960s. The first legged or walking robots appeared in a patent for a mechanical horse in 1893, but it was not until the early 1960s that an operational walking vehicle was constructed. Perhaps the most famous of the early legged vehicles is the General Electric Quadruped (see Figure 1.5) [373, 437]. Each of the four legs of this vehicle had three simple joints; the knee joint was composed of a single joint, and the hip joint used two. As can be seen in Figure 1.5, the GE Quadruped was controlled by