Layout of the Section

This section is divided into eight chapters, one to define robotics and the other seven to intertwine both the theory and practice associated with each paradigm. Ch. 2 describes the Hierarchical Paradigm and two representative architectures. Ch. 3 sets the stage for understanding the Reactive Paradigm by reviewing the key concepts from biology and ethology that served to motivate the shift from Hierarchical to Reactive systems. Ch. 4 describes the Reactive Paradigm and the architectures that originally popularized this approach. It also offers definitions of primitive robot behaviors. Ch. 5 provides guidelines and case studies on designing robot behaviors. It also introduces issues in coordinating and controlling multiple behaviors and the common techniques for resolving these issues. At this point, the reader should be almost able to design and implement a reactive robot system, either in simulation or on a real robot. However, the success of a reactive system depends on the sensing. Ch. 6 discusses simple sonar and computer vision processing techniques that are commonly used in inexpensive robots. Ch. 7 describes the Hybrid Deliberative-Reactive Paradigm, concentrating on architectural trends. Up until this point, the emphasis is towards programming a single robot. Ch. 8 concludes the section by discussing how the principles of the three paradigms have been transferred to teams of robots.

End Note

Robot paradigm primitives.
While the SENSE, PLAN, ACT primitives are generally accepted, some researchers are suggesting that a fourth primitive be added, LEARN. There are no formal architectures at this time which include this, so a true paradigm shift has not yet occurred.

1 *From Teleoperation To Autonomy*

Chapter Objectives:

- Define *intelligent robot.*
- Be able to describe at least two differences between AI and Engineering approaches to robotics.
- Be able to describe the difference between *telepresence* and *semi-autonomous control.*
- Have some feel for the history and societal impact of robotics.

1.1 Overview

This book concentrates on the role of artificial intelligence for robots. At first, that may appear redundant; aren't robots intelligent? The short answer is "no," most robots currently in industry are not intelligent by any definition. This chapter attempts to distinguish an intelligent robot from a non-intelligent robot.

The chapter begins with an overview of artificial intelligence and the social implications of robotics. This is followed with a brief historical perspective on the evolution of robots towards intelligence, as shown in Fig. 1.1. One way of viewing robots is that early on in the 1960's there was a fork in the evolutionary path. Robots for manufacturing took a fork that has focused on engineering robot arms for manufacturing applications. The key to success in industry was precision and repeatability on the assembly line for mass production, in effect, industrial engineers wanted to automate the workplace. Once a robot arm was programmed, it should be able to operate for weeks and months with only minor maintenance. As a result, the emphasis was

The term AI is controversial, and has sparked ongoing philosophical de-
bates on whether a machine can ever be intelligent. As Roger Penrose notes
in his book, The Emperor's New Mind: "Nevertheless, it would be fair to
say that, although many clever things have indeed been done, the simula-
tion of anything that could pass for genuine intelligence is yet a long way
off."[115] Engineers often dismiss AI as wild speculation. As a result of such
vehement criticisms, many researchers often label their work as "intelligent
systems" or "knowledge-based systems" in an attempt to avoid the contro-
versy surrounding the term "AI."

A single, precise definition of AI is not necessary to study AI robotics. AI
robotics is the application of AI techniques to robots. More specifically, AI
robotics is the consideration of issues traditional covered by AI for applica-
tion to robotics: learning, planning, reasoning, problem solving, knowledge
representation, and computer vision. An article in the May 5, 1997 issue
of Newsweek, "Actually, Chess is Easy," discusses why robot applications
are more demanding for AI than playing chess. Indeed, the concepts of the
reactive paradigm, covered in Chapter 4, influenced major advances in tra-
ditional, non-robotic areas of AI, especially planning. So by studying AI ro-
botics, a reader interested in AI is getting exposure to the general issues in
AI.

1.3 What Can Robots Be Used For?

Now that a working definition of a robot and artificial intelligence has been
established, an attempt can be made to answer the question: what can intel-
ligent robots be used for? The short answer is that robots can be used for just
about any application that can be thought of. The long answer is that robots
are well suited for applications where 1) a human is at significant risk (nu-
clear, space, military), 2) the economics or menial nature of the application
result in inefficient use of human workers (service industry, agriculture), and
3) for humanitarian uses where there is great risk (demining an area of land
mines, urban search and rescue). Or as the well-worn joke among roboticists
THE 3 D'S goes, robots are good for *the 3 D's*: jobs that are dirty, dull, or dangerous.

Historically, the military and industry invested in robotics in order to build
nuclear weapons and power plants; now, the emphasis is on using robots for
environmental remediation and restoration of irradiated and polluted sites.
Many of the same technologies developed for the nuclear industry for pro-
adapted for the pharmaceutical indus-

try; processing immune suppressant drugs may expose workers to highly toxic chemicals.

Another example of a task that poses significant risk to a human is space exploration. People can be protected in space from the hard vacuum, solar radiation, etc., but only at great economic expense. Furthermore, space suits are so bulky that they severely limit an astronaut's ability to perform simple tasks, such as unscrewing and removing an electronics panel on a satellite. Worse yet, having people in space necessitates more people in space. Solar radiation embrittlement of metals suggests that astronauts building a large space station would have to spend as much time repairing previously built portions as adding new components. Even more people would have to be sent into space, requiring a larger structure. the problem escalates. A study by Dr. Jon Erickson's research group at NASA Johnson Space Center argued that a manned mission to Mars was not feasible without robot drones capable of constantly working outside of the vehicle to repair problems introduced by deadly solar radiation.[51] (Interestingly enough, a team of three robots which did just this were featured in the 1971 film, *Silent Running*, as well as by a young R2D2 in *The Phantom Menace*.)

Nuclear physics and space exploration are activities which are often far removed from everyday life, and applications where robots figure more prominently in the future than in current times.

The most obvious use of robots is manufacturing, where repetitious activities in unpleasant surroundings make human workers inefficient or expensive to retain. For example, robot "arms" have been used for welding cars on assembly lines. One reason that welding is now largely robotic is that it is an unpleasant job for a human (hot, sweaty, tedious work) with a low tolerance for inaccuracy. Other applications for robots share similar motivation: to automate menial, unpleasant tasks—usually in the service industry. One such activity is janitorial work, especially maintaining public rest rooms, which has a high turnover in personnel regardless of payscale. The janitorial problem is so severe in some areas of the US, that the Postal Service offered contracts to companies to research and develop robots capable of autonomously cleaning a bathroom (the bathroom could be designed to accommodate a robot).

Agriculture is another area where robots have been explored as an economical alternative to hard to get menial labor. Utah State University has been working with automated harvesters, using GPS (global positioning satellite system) to traverse the field while adapting the speed of harvesting to the rate of food being picked, much like a well-adapted insect. The De-

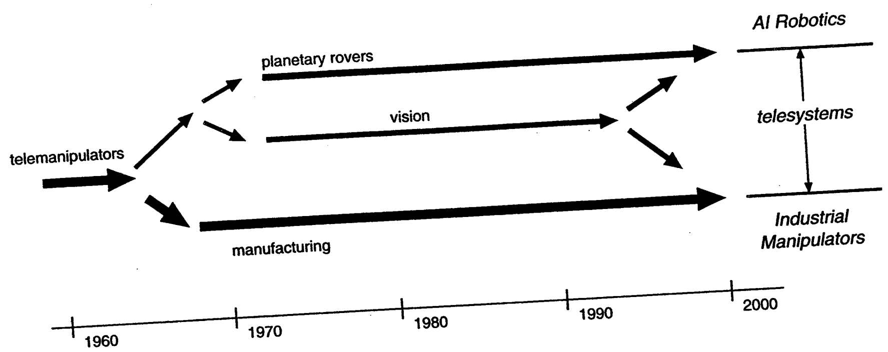

Figure 1.1 A timeline showing forks in development of robots.

placed on the mechanical aspects of the robot to ensure precision and repeatability and methods to make sure the robot could move precisely and repeatable, quickly enough to make a profit. Because assembly lines were engineered to mass produce a certain product, the robot didn't have to be able to notice any problems. The standards for mass production would make it more economical to devise mechanisms that would ensure parts would be in the correct place. A robot for automation could essentially be blind and senseless.

Robotics for the space program took a different fork, concentrating instead on highly specialized, one-of-a-kind planetary rovers. Unlike a highly automated manufacturing plant, a planetary rover operating on the dark side of the moon (no radio communication) might run into unexpected situations. Consider that on Apollo 17, astronaut and geologist Harrison Schmitt found an orange rock on the moon; an orange rock was totally unexpected. Ideally, a robot would be able to notice something unusual, stop what it was doing (as long as it didn't endanger itself) and investigate. Since it couldn't be pre-programmed to handle all possible contingencies, it had to be able to notice its environment and handle any problems that might occur. At a minimum, a planetary rover had to have some source of sensory inputs, some way of interpreting those inputs, and a way of modifying its actions to respond to a changing world. And the need to sense and adapt to a partially unknown environment is the need for intelligence.

The fork toward AI robots has not reached a termination point of truly autonomous, intelligent robots. In fact, as will be seen in Ch. 2 and 4, it wasn't until the late 1980's that any visible progress toward that end was made. So what happened when someone had an application for a robot which needed

real-time adaptability before 1990? In general, the lack of machine intelligence was compensated by the development of mechanisms which allow a human to control all, or parts, of the robot remotely. These mechanisms are generally referred to under the umbrella term: teleoperation. Teleoperation can be viewed as the "stuff" in the middle of the two forks. In practice, intelligent robots such as the Mars Sojourner are controlled with some form of teleoperation. This chapter will cover the flavors of teleoperation, given their importance as a stepping stone towards truly intelligent robots.

The chapter concludes by visiting the issues in AI, and argues that AI is imperative for many robotic applications. Teleoperation is simply not sufficient or desirable as a long term solution. However, it has served as a reasonable patch.

It is interesting to note that the two forks, manufacturing and AI, currently appear to be merging. Manufacturing is now shifting to a "mass customization" phase, where companies which can economically make short runs of special order goods are thriving. The pressure is on for industrial robots, more correctly referred to as industrial manipulators, to be rapidly reprogrammed and more forgiving if a part isn't placed exactly as expected in its workspace. As a result, AI techniques are migrating to industrial manipulators.

1.2 How Can a Machine Be Intelligent?

ARTIFICIAL INTELLIGENCE

The science of making machines act intelligently is usually referred to as *artificial intelligence*, or AI for short. Artificial Intelligence has no commonly accepted definitions. One of the first textbooks on AI defined it as "the study of ideas that enable computers to be intelligent,"[143] which seemed to beg the question. A later textbook was more specific, "AI is the attempt to get the computer to do things that, for the moment, people are better at."[120] This definition is interesting because it implies that once a task is performed successfully by a computer, then the technique that made it possible is no longer AI, but something mundane. That definition is fairly important to a person researching AI methods for robots, because it explains why certain topics suddenly seem to disappear from the AI literature: it was perceived as being solved! Perhaps the most amusing of all AI definitions was the slogan for the now defunct computer company, Thinking Machines, Inc., "... making machines that will be proud of us."