

Andrew C. Staugaard, Jr.

**ROBOTICS AND AI:
AN INTRODUCTION TO
APPLIED MACHINE INTELLIGENCE**



Andrew C. Staugaard, Jr.

*Department of Computer Science
The School of the Ozarks*

**ROBOTICS AND AI:
AN INTRODUCTION TO
APPLIED MACHINE INTELLIGENCE**

PRENTICE HALL, INC., Englewood Cliffs, N.J. 07632

STAUGAARD, ANDREW C.

Robotics and AI

Bibliography: p. 337

Includes index

1. Robotics. 2. Artificial intelligence. I. Title.
TJ211.S73 1987 006.3 86-25140

ISBN 0-13-782269-3

Editorial/production supervision and
interior design: Tom Aloisi

Cover design: 20/20 Services, Inc.

Manufacturing buyer: Carol Bystrom

To You, the Student: May This Book Enhance
Your Quest To Be All That You Can Be

© 1987 by Prentice-Hall, Inc.,

A Division of Simon & Schuster, Inc.

Englewood Cliffs, New Jersey 07632

All rights reserved. No part of this book may be
reproduced, in any form or by any means,
without permission in writing from the publisher.

Printed in the United States of America

10 9 8 7 6 5 4 3 2 1

ISBN: 0-13-782269-3 025

Prentice-Hall International (UK) Limited, London

Prentice-Hall of Australia Pty. Limited, Sydney

Prentice-Hall Canada Inc., Toronto

Prentice-Hall Hispanoamericana, S. A., Mexico

Prentice-Hall of India Private Limited, New Delhi

Prentice-Hall of Japan, Inc., Tokyo

Prentice-Hall of Southeast Asia Pte. Ltd., Singapore

Editora Prentice-Hall do Brasil, Ltda., Rio de Janeiro

Preface

Welcome to the fascinating world of intelligent machines! Imagine machines that might someday think, see, and communicate by voice very much the same as you and I do. Such machines will become a reality through the direct application of concepts from a branch of computer science called *artificial intelligence*, or *AI*. In this book, you will learn about the basic concepts of AI and how AI is applied to produce intelligent machines.

An area that is particularly suited to applied machine intelligence is robotics. An intelligent robot must be capable of sensing its surroundings and responding to a changing environment much the same as we do. In fact, the robot that can perform humanlike intelligent tasks is the ultimate intelligent machine. For this reason, this book focuses on the applications of artificial intelligence to robotic systems. You will learn about voice communication systems, vision systems, range and navigation systems, and tactile sensing systems.

In the field of artificial intelligence, it often seems as though the author is talking to his or her colleagues rather than **you**, the student. In robotics, important operational concepts are often left out of the discussion. This book is my attempt to separate the wheat from the chaff. I have strived to write a readable text that zeros in on the important concepts of artificial intelligence and how those concepts can be applied to produce intelligent machines. At the end of this book, you will not only understand the **what's**, but more important, the **why's** and **how to's**.

One prerequisite to reading this book is a fundamental knowledge of Boolean algebra, such as you might get in a digital electronics or introductory computer science course. In addition, you should have a basic understanding of how simple electronic components, such as resistors, capacitors, inductors, diodes, and transistors, operate.

This book is meant to be used at the introductory level as a first exposure to artificial intelligence and/or intelligent machines. As a result, the book is ideal for a one semester introductory course in undergraduate technology, engineering, and computer science programs. The basic concepts learned here will prepare you for more advanced study in the areas of artificial intelligence, vision systems, voice communication systems, tactile sensing systems, and intelligent machines in general.

The first chapter lays the foundation for the rest of the text. Chapters 2 and 3 are devoted entirely to a discussion of the basic concepts of artificial intelligence and the related field of knowledge representation. Chapters 4 through 8 discuss the application areas of speech synthesis, speech recognition and understanding, vision, range finding and navigation, and tactile sensing, respectively.

Several people have made valuable suggestions and contributions during the preparation of this book. I want to thank Vince Leonard and Bob Kochersberger of Jamestown Community College, Pete Ho of the University of Missouri-Rolla, and finally, Marvin Minsky of the Massachusetts Institute of Technology for giving me insight into the tremendous potential of AI.

Andrew C. Staugaard, Jr.

Contents

PREFACE ix

1 INTELLIGENT ROBOTICS 1

- 1-1 Robot Technology 1
- 1-2 End-effectors 5
- 1-3 Sensory Perception 6
- 1-4 Robot Control and Intelligence 11
- Summary 16
- Problems 17

2 BASIC CONCEPTS OF ARTIFICIAL INTELLIGENCE 19

- 2-1 Intelligence 19
- 2-2 Problem Representation in Artificial Intelligence 25
- 2-3 Problem-Solution Techniques Used in Artificial Intelligence 40
- Summary 59
- Problems 60

3 ELEMENTS OF KNOWLEDGE REPRESENTATION 63

- 3-1 Logic 64
- 3-2 Production Systems 80
- 3-3 Semantic Networks 90
- 3-4 Frames 96
- 3-5 Expert Systems 112
- 3-6 Programming Languages 122
- Summary 123
- Problems 124

4 SPEECH SYNTHESIS 128

- 4-1 Elements of Speech 128
- 4-2 Time-domain Analysis/Synthesis and Waveform Digitization 138
- 4-3 Frequency-domain Analysis/Synthesis 145
- 4-4 Phoneme Speech Synthesis 150
- Summary 166
- Problems 167

5 SPEECH RECOGNITION AND UNDERSTANDING 169

- 5-1 Types of Speech-recognition Systems 170
- 5-2 Isolated-word Recognition 172
- 5-3 Connected-speech Understanding 184
- Summary 193
- Problems 194

6 VISION 196

- 6-1 Image Transformation 197
- 6-2 Image Analysis 208
- 6-3 Image Understanding 230
- 6-4 Industrial Vision Systems 254
- Summary 262
- Problems 263

7 RANGE FINDING AND NAVIGATION 266

- 7-1 Triangulation 267
- 7-2 Time-of-Flight Ranging 276
- 7-3 Robot Position and Proximity Sensing 283
 - Summary 292
 - Problems 293

8 TACTILE SENSING 295

- 8-1 Tactile-sensing System 296
- 8-2 Sensing Joint Forces 298
- 8-3 Sensing Touch 311
- 8-4 Sensing Slip 328
 - Summary 333
 - Problems 335

REFERENCES 337

ANSWERS TO CHAPTER PROBLEMS 340

INDEX 368

Intelligent Robotics

The field of robotics is in a state of rapid development. Early robots were nothing more than clever mechanical devices with cams and stops that performed simple pick-and-place operations. As computer technology improved, robots became more sophisticated. Computer control allows robots to perform more precise industrial operations, such as welding, spray painting, and simple parts assembly. However, such operations do not really require the robot to "think." The robots are simply programmed to perform a series of repetitive tasks. If anything interferes with the preprogrammed task, the robot must be stopped, since it is not capable of sensing its external environment and thinking its way out of a problem. For robots to become more efficient, maintenance free, and productive, they must be capable of sensing external conditions and thinking very much like you do. Such abilities require the direct application of *artificial intelligence* and *sensory perception*.

This brief chapter lays the foundation for the entire text. The chapter begins with an overview of robot technology leading to a discussion of industrial versus personal robots, the limitations of robots, and the future of robotics. Elements of intelligent robotics, such as programmability, flexibility, input/output, and sensory feedback are introduced in preparation for subsequent chapters.

1-1 ROBOT TECHNOLOGY

There has been much discussion as to what type of apparatus constitutes a robot. For instance, a pilot flying an airliner operates controls that activate various air control surfaces to allow the plane to fly a course from point A to point B. Clearly, this is not a robotic type of action, but rather the effect of direct mechanical

linkages between the pilot and the control surfaces. But suppose that the airliner is equipped with a computer-based navigation control system. Such systems are capable of literally "flying" the airplane from take-off to landing without any interference from the pilot. A passenger cannot tell who is flying the airplane, the pilot or the computer. Does the plane become a robot when being flown by the computer? Many contend that a machine becomes a robot when it can perform physical tasks without human intervention. Does this make the computer-flown airplane, or for that matter your automatic dishwasher, a robot?

Webster defines a robot as "an automatic apparatus or device that performs functions ordinarily ascribed to human beings or operates with what appears to be almost human intelligence." The Robot Institute of America says, "a robot is a reprogrammable, multifunctional manipulator designed to move material, parts, tools, or specialized devices through variable programmed motions for the performance of a variety of tasks." Do either of these definitions rule out a computer-flown airliner or your automatic dishwasher from being classified as a robot?

Most people think of the devices pictured in Figure 1-1 as robots. But what makes these devices different from your automatic dishwasher? Is it the fact that they are computer controlled? Many automatic dishwashers on the market today are microprocessor based, and thus computer controlled.

As you can see, it is rather difficult to nail down a precise definition for a






Figure 1-1 (a) Ideal for complex or heavy-duty materials-handling tasks, the T³586 Robot from Cincinnati Milacron has infinitely variable six-axis positioning and a payload capacity up to 225 pounds. In addition to handling hefty loads at high speeds and performing in hazardous or difficult-to-reach places, it can load and unload machines, palletize, inspect, sort, weld, and assemble. The T³586 features a hydraulic power system, a simple hand-held teach pendant for the operator to lead it through its moves, and an advanced computer control. (b) Cincinnati



Milacron's all electric, computer-controlled, T³726 Industrial Robot for welding, plasma cutting, parts handling, or assembly applications. The T³726 Robot features a payload capacity of 14 pounds, six axes of motion, and Milacron's unique three-roll wrist for true application flexibility. Each axis is powered by its own dc motor. (c) Cincinnati



Milacron has introduced a low-cost new robot specifically developed for materials-handling tasks. The new T³363 is the first in a series Milacron designed to bring the flexibility of robots to materials-handling applications at an affordable price. At an average cost of under \$30,000, the all-electric T³363 will move payloads up to 110 pounds and perform most materials-handling tasks, including palletizing, parts handling, package handling, and machine loading and unloading. It has three standard axes of motion, with the option of a fourth, and an easy-to-use microprocessor-based control capable of running a work cell. (Courtesy Cincinnati Milacron)

robot. However, most people think "robot" when they see a machine performing much like a human being, exhibiting a certain degree of "intelligence." Thus, a robot must possess a certain degree of machine, or artificial, intelligence. For now, I will leave the precise definition of a robot to Webster and the Robot Institute of America. However, I will attempt to define machine, or artificial, intelligence in a subsequent chapter.

Robots have developed along two paths: *industrial* and *domestic*. Industrial robots, such as those pictured in Figure 1-1, have been developed to perform a variety of manufacturing tasks, such as machine loading/unloading, welding, spray painting, and simple product assembly. However, most industrial robots have very limited sensing ability. If assembly parts are not presented to the robot in a precise, repetitive fashion, the robot cannot perform its task. If an object enters the work area of a robot, the robot and the object will likely collide, resulting in damage to both. In other words, most industrial robots are unintelligent, and cannot hear, see, or feel. Imagine trying to teach a person with all these handicaps to perform precision assembly operations. To become more efficient and productive, an industrial robot must be capable of *sensing* its surroundings and possess enough *intelligence* to respond to a changing environment much the same as we do. Thus, elements of sight, touch, and corrective action (intelligence) are essential to the evolution of industrial robotics.

Domestic or personal robots have been developed primarily for the home hobbyist market. Most of these devices are capable of voice synthesis (speaking), sensing light levels, detecting motion, and moving about with simple programmed instructions and sonar-type navigation systems. However, existing domestic robots, like their industrial counterparts, have limited intelligence, thereby limiting their applications.

Suppose you require a domestic robot that will perform a simple everyday task like cleaning off the kitchen table and washing the dishes after each meal. Imagine the sensing abilities and intelligence that such a robot must possess! First, it must be able to distinguish the dishes and eating utensils from other objects, like the morning paper, that might be on the table. Then it must have the ability to pick up and manipulate the delicate dishes, perform the washing and drying task, and place the dishes and utensils in their proper storage location. In addition, it must be capable of sweeping up and disposing of any broken dishes, as well as compensating for any change in its environment during the cleaning task. And what happens if the pet cat gets in its way during the washing operation? Of course, the robot must also be capable of receiving voice commands to warn it of any impending disaster or to alter its operation.

Now imagine the additional sensing abilities and intelligence that the same domestic robot must have to perform other household tasks, such as washing clothes, running the vacuum, cleaning the bathroom, taking out the garbage, and mowing the grass. Such an advanced robot must possess *sensing* and *intelligence* abilities similar to those of a human.

From the preceding discussion, it is clear that the two major requirements of

an advanced robot (industrial or domestic) are that it must sense its surroundings and be intelligent enough to compensate for changes in its environment. In fact, you could say that *sensory perception* and *intelligence* are the common denominators of any advanced robot.

In most cases, intelligent robots will eventually incorporate elements of vision, touch, and speech recognition, all of which require the direct application of artificial intelligence. Such a robot must also possess adequate manipulators and end-effectors to perform the given task. Thus, the three key ingredients of an intelligent robot must be *sensory perception*, *intelligence*, and *adequate end-effectors* to perform a variety of applications tasks. Let's take a closer look at each of these three fundamental components of an intelligent robot in preparation for the material that follows in later chapters.

1-2 END-EFFECTORS

Robot end-effectors, or hands, can take on many forms. In general, end-effectors are the mechanical devices with which the robot manipulates the real world. The devices required to operate the robot's sensors must also be included. If a robot uses a TV vision system, it must be able to turn and focus the camera in the direction it wishes to see. For industrial robots, end-effectors also include the tools required to perform a given task.

Much of the research on robot end-effectors has been centered on improving the flexibility of the robot's hands and arms. Currently, most industrial end-effectors are single degree of freedom devices specialized to a given task. Examples include welding torches, paint sprayers, two-fingered grippers, and vacuum cups. Several of these specialized industrial end-effectors are illustrated in Figure 1-2. However, such designs are not well suited for general-purpose operations. As a result, many researchers are attempting to develop multifingered grippers similar to the human hand. This is not an easy task since the human hand has 22 individual joints, or axis of motion. The best attempts have only resulted in three-fingered grippers.

The design of end-effectors is a science in itself. Aside from the actual mechanical design, any associated control software must be developed. This is complicated by the fact that the end-effector must incorporate various sensing devices. The sensing devices are used to measure and feed back process variable data so that control decisions can be made. This is referred to as *sensory feedback*. The controller software must then coordinate any movement of the end-effector with the sensory feedback data. The mechanical design technology required for sophisticated end-effectors will not be covered in this book. Rather, my emphasis in subsequent chapters will be on the fundamental concepts of artificial intelligence and sensory feedback that must be incorporated into any sophisticated end effector, or manipulator, design.




Figure 1-2 Most industrial end-effectors are single degree of freedom devices specialized to a given task: (a) standard gripper; (b) arc welder; (c) spot welder; (d) paint sprayer.

1.3 SENSORY PERCEPTION

Sensory perception, or feedback, must be incorporated into robot designs if robots are to become more than just mechanical manipulators. The categories of sensory perception include vision, tactile sensing, range finding, and voice communication. In most robot applications, the preceding order of categories represents the relative application enhancement priority. In other words, the possible range of applications of a robot is enhanced the most by vision. Tactile sensing, or touch, runs a close second, especially for industrial assembly operations. Many existing product assembly tasks employ some form of tactile sensing in lieu of vision. Range finding is required for three-dimensional (3-D) vision and also for position and proximity, or nearness, sensing. Range-finding techniques are also important for robot navigation, whether a robot is stationary and navigates its gripper or mobile and navigates its body. Finally, voice communication is less important, but makes the robot more user friendly and permits verbal person-to-robot communication for programming and control purposes. Let's take a short look at each of these sensory perception categories.

Vision

Vision is probably the single most important sensing ability that an intelligent robot can possess. An industrial robot that can "see" is capable of parts recognition, parts sorting, and precision assembly operations. Likewise, a domestic robot requires a sense of vision to perform everyday household tasks and to navigate from room to room. Any robot would be more intelligent if it could acquire information about its environment through its own vision system rather than being limited to a knowledge base provided by its programmer. Thus, a robot vision system can actually be used to build the knowledge base of the robot.

The science of computer vision is called *visual machine perception*. Visual perception is easy for us since once we become familiar with an object it is easily recognized. We seem to just "know" that a dog is a dog and a cat is a cat. However, visual perception is very difficult for a computer. An enormous amount of artificial intelligence is required for a computer, or robot, to distinguish between a dog and a cat.

Simple computer vision systems detect different levels of light. These systems typically utilize a light-dependent resistor (LDR) or phototransistor circuit to translate light levels to a proportional analog voltage level. The analog voltage is then converted to a digital value with an analog-to-digital (A/D) converter and read by the computer. Such a system is illustrated in Figure 1-3. A robot using this system can be readily programmed to awaken you at sunup or to detect when an object is in a given position. However, it cannot really recognize an object or perceive its surroundings. Actually, you could say that the system described is a light-detection system rather than a vision system.

True computer vision involves the transformation, analysis, and understanding of light images. As a result, the science of computer vision can be reduced to three fundamental tasks: *image transformation*, *image analysis*, and *image understanding*.

Image transformation involves the conversion of light images into electrical signals that can be used by a computer. Many existing computer vision systems

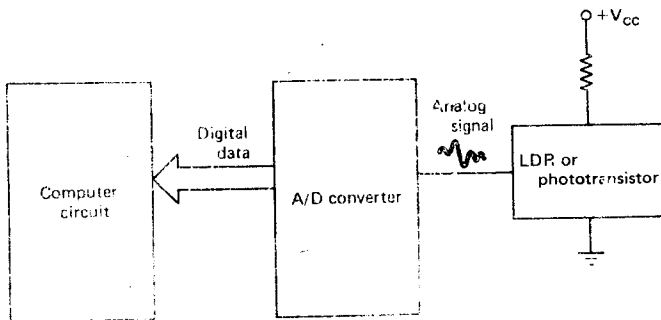


Figure 1-3 Simple robot light-detection systems only detect light levels and are not capable of vision.

utilize a TV camera, photodiode array camera, or charge-coupled device (CCD) camera for the *imaging device*, or eye, as illustrated in Figure 1-4. As the camera scans a scene, its output is converted to a digital code by an A/D converter and stored in memory as a digital image. The computer must then analyze the digital images and apply some degree of artificial intelligence to understand the scene.

Once a light image is transformed into an electronic image, it must be analyzed to extract such image information as object edges, regions, boundaries, color, and texture. In some systems, the digital images are compared to *image templates* to classify and recognize objects in the scene. The image templates have been previously stored in the computer memory for the comparison task. Using the template technique, the computer can recognize distinct, well-defined patterns. However, the problem is in storing enough memory templates to cover all possible scenes in a given environment. More intelligent systems use edge and region image analysis methods. In these systems, the edges and regions of objects in the scene are analyzed in order to generate information for the image-understanding process.

For a computer to truly see, it must be capable of understanding a given scene and use the knowledge gained from the scene for future problem-solving tasks. This is by far the most difficult of all computer vision tasks and requires the direct application of artificial intelligence. Image transformation, analysis, and understanding are covered in Chapter 6.




Figure 1-4 Many existing vision systems utilize a TV camera as the visual sensing device.

Tactile Sensing

Next to vision, tactile sensing, or touch, most enhances a robot's abilities. Like a blind person, a blind robot can be extremely effective in performing an assembly task using only a sense of touch. In fact, many times tactile sensing is more critical to a precision assembly operation than is vision.

During a precision assembly task, an industrial robot must be capable of

sensing any problems encountered during assembly from the interaction of parts and tools. Such interference and fit problems are created by part tolerance, misalignment of parts, tool wear, and the like. For example, suppose an industrial robot is making a gear assembly. One gear is to be placed on an axle, followed by a second gear on an adjacent axle. The teeth of the second gear must mesh with the teeth of the first gear. Using its tactile sensing abilities, the robot must sense any interference between the two gears and rotate the second gear until its teeth mesh with those of the first gear. If the robot cannot see what it is doing, this would be an impossible assembly task without tactile sensing. Imagine trying to perform the same operation blindfolded and without any sense of touch—impossible!

The term **compliance** describes the allowed movement between mating parts for the purpose of alignment during assembly tasks. In other words, a **compliant robot** “complies” with external forces by modifying its motions in order to minimize those forces. A robot that could successfully perform the preceding gear assembly operation by compensating for interference between the gears would be called a compliant robot. There is both **active** and **passive compliance**. Active compliance employs sensory feedback, such as tactile sensing, whereas passive compliance does not incorporate any sensory feedback.

As you will discover in Chapter 8, simple tactile sensing can be accomplished by placing microswitches, strain gauges, pressure transducers, and optical sensors in the end-effector of the robot. Magnetic Hall-effect devices and sonar sensors are also sometimes used. However, the most important parameter that must be measured to achieve tactile sensing is force. To accomplish this, strain gauge and pressure transducers are commonly placed on the robot’s arm, wrist, and fingers. In addition, artificial skin pads that sense pressure are placed on the robot’s gripper. These **tactile arrays** provide a sense of feel for the robot and can be used to determine the position and orientation of an object and to aid in the identification of unknown objects.

Information generated by these various force-sensing devices can be transmitted to the central robot computer or analyzed by **sensing cells** located within the manipulator. Sensing cells are single-chip microcomputers dedicated to the tactile sensing task. This subject is dealt with further in Chapter 8.

Range Finding and Navigation

Range, or distance, data must be obtained in order for a robot to create the 3-D information necessary for real-world navigation. Such information is required whether the robot is stationary and navigating its gripper or mobile and navigating its body. Vision systems can employ stereo cameras used like your eyes to determine depth through triangulation. However, vision systems are relatively sophisticated and expensive. Many simple robot navigation problems do not require such a sophisticated solution.

A simpler solution to range finding is found in time-of-flight ranging systems. These systems measure the amount of time it takes a radio signal, sound, or light to