INTRODUCTION TO

AI ROBOTICS

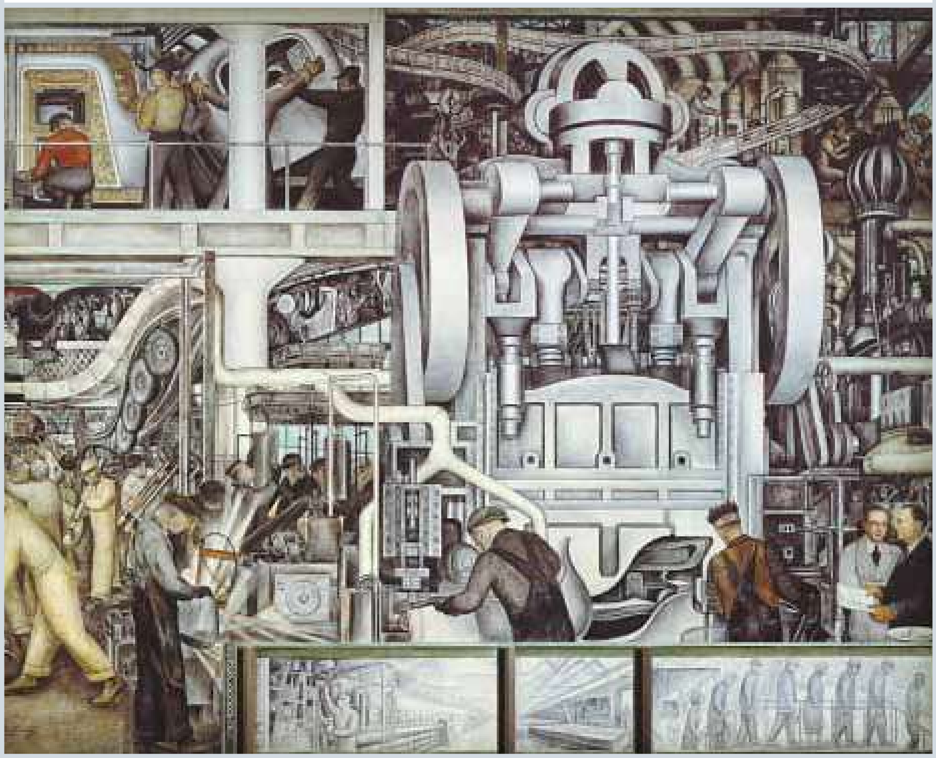

ROBIN R. MURPHY

Introduction to AI Robotics

Robin R. Murphy

A Bradford Book
The MIT Press
Cambridge, Massachusetts
London, England

PART I

Robotic Paradigms

Contents:

Overview

The eight chapters in this part are devoted to describing what is *AI robotics* and the three major paradigms for achieving it. These paradigms characterize the ways in which intelligence is organized in robots. This part of the book also covers architectures that provide exemplars of how to transfer the principles of the paradigm into a coherent, reusable implementation on a single robot or teams of robots.

What Are Robots?

One of the first questions most people have about robotics is "what is a robot?" followed immediately by "what can they do?"

In popular culture, the term "robot" generally connotes some anthropomorphic (human-like) appearance; consider robot "arms" for welding. The tendency to think about robots as having a human-like appearance may stem from the origins of the term "robot." The word "robot" came into the popular consciousness on January 25, 1921, in Prague with the first performance of Karel Capek's play, R.U.R. (Rossum's Universal Robots).[37] In R.U.R., an unseen inventor, Rossum, has created a race of workers made from a vat of biological parts, smart enough to replace a human in any job (hence "universal"). Capek described the workers as robots, a term derived from the Czech word "robota" which is loosely translated as menial laborer. Robot workers implied that the artificial creatures were strictly meant to be servants to free "real" people from any type of labor, but were too lowly to merit respect. This attitude towards robots has disastrous consequences, and the moral of the rather socialist story is that work defines a person.

The shift from robots as human-like servants constructed from biological parts to human-like servants made up of mechanical parts was probably due to science fiction. Three classic films, *Metropolis* (1926), *The Day the Earth Stood Still* (1951), and *Forbidden Planet* (1956), cemented the connotation that robots were mechanical in origin, ignoring the biological origins in Capek's play. Meanwhile, computers were becoming commonplace in industry and accounting, gaining a perception of being literal minded. Industrial automation confirmed this suspicion as robot arms were installed which would go through the motions of assembling parts, even if there were no parts. Eventually, the term robot took on nuances of factory automation: mindlessness and good only for well-defined repetitious types of work. The notion of anthropomorphic, mechanical, and literal-minded robots complemented the viewpoint taken in many of the short stories in Isaac Asimov's perennial favorite collection, *I, Robot*.[15] Many (but not all) of these stories involve either a "robopsychologist," Dr. Susan Calvin, or two erstwhile trouble shooters, Powell and Donovan, diagnosing robots who behaved logically but did the wrong thing.

The shift from human-like mechanical creatures to whatever shape gets the job done is due to reality. While robots are mechanical, they don't have to be anthropomorphic or even animal-like. Consider robot vacuum cleaners; they look like vacuum cleaners, not janitors. And the HelpMate Robotics, Inc., robot which delivers hospital meals to patients to permit nurses more time with patients, looks like a cart, not a nurse.

INTELLIGENT ROBOT

It should be clear from Fig. I.1 that appearance does not form a useful definition of a robot. Therefore, the definition that will be used in this book is *an intelligent robot is a mechanical creature which can function autonomously*. "Intelligent" implies that the robot does not do things in a mindless, repetitive way; it is the opposite of the connotation from factory automation. The "mechanical creature" portion of the definition is an acknowledgment of the fact that our scientific technology uses mechanical building blocks, not biological components (although with recent advances in cloning, this may change). It also emphasizes that a robot is not the same as a computer. A robot may use a computer as a building block, equivalent to a nervous system or brain, but the robot is able to interact with its world: move around, change

a. b.

Figure I.1 Two views of robots: a) the humanoid robot from the 1926 movie Metropolis (image courtesty Fr. Doug Quinn and the Metropolis Home Page), and b) a HMMWV military vehicle capable of driving on roads and open terrains. (Photograph courtesy of the National Institute for Standards and Technology.)

it, etc. A computer doesn't move around under its own power. "Function autonomously" indicates that the robot can operate, self-contained, under all reasonable conditions without requiring recourse to a human operator. Autonomy means that a robot can adapt to changes in its environment (the lights get turned off) or itself (a part breaks) and continue to reach its goal.

Perhaps the best example of an intelligent mechanical creature which can function autonomously is the Terminator from the 1984 movie of the same name. Even after losing one camera (eye) and having all external coverings (skin, flesh) burned off, it continued to pursue its target (Sarah Connor). Extreme adaptability and autonomy in an extremely scary robot! A more practical (and real) example is *Marvin*, the mail cart robot, for the Baltimore FBI office, described in a Nov. 9, 1996, article in the Denver Post. Marvin is able to accomplish its goal of stopping and delivering mail while adapting to people getting in its way at unpredictable times and locations.

What are Robotic Paradigms?

PARADIGM

A *paradigm is a philosophy or set of assumptions and/or techniques which characterize an approach to a class of problems*. It is both a way of looking at the world and an implied set of tools for solving problems. No one paradigm is right; rather, some problems seem better suited for different approaches. For example, consider calculus problems. There are problems that could be solved by differentiating in cartesian (X, Y, Z) coordinates, but are much easier to solve if polar coordinates (r, θ) are used. In the domain of calculus problems, Cartesian and polar coordinates represent two different paradigms for viewing and manipulating a problem. Both produce the correct answer, but one takes less work for certain problems.

Applying the right paradigm makes problem solving easier. Therefore, knowing the paradigms of AI robotics is one key to being able to successfully program a robot for a particular application. It is also interesting from a historical perspective to work through the different paradigms, and to examine the issues that spawned the shift from one paradigm to another.

ROBOTIC PARADIGMS

There are currently three paradigms for organizing intelligence in robots: hierarchical, reactive, and hybrid deliberative/reactive. The paradigms are described in two ways.

ROBOT PARADIGM PRIMITIVES

1. **By the relationship between the three commonly accepted primitives of robotics: SENSE, PLAN, ACT.** The functions of a robot can be divided into three very general categories. If a function is taking in information from the robot's sensors and producing an output useful by other functions, then that function falls in the SENSE category. If the function is taking in information (either from sensors or its own knowledge about how the world works) and producing one or more tasks for the robot to perform (go down the hall, turn left, proceed 3 meters and stop), that function is in the PLAN category. Functions which produce output commands to motor actuators fall into ACT (turn $98°$, clockwise, with a turning velocity of 0.2mps). Fig. I.2 attempts to define these three primitives in terms of inputs and outputs; this figure will appear throughout the chapters in Part I.

2. **By the way sensory data is processed and distributed through the system**. How much a person or robot or animal is influenced by what it senses. So it is often difficult to adequately describe a paradigm with just a box labeled SENSE. In some paradigms, sensor information is restricted to being used in a specific, or dedicated, way for each function of a robot;

ROBOT PRIMITIVES	INPUT	OUTPUT
SENSE	Sensor data	Sensed information
PLAN	Information (sensed and/or cognitive)	Directives
ACT	Sensed information or directives	Actuator commands

Figure I.2 Robot primitives defined in terms of inputs and outputs.

SENSING ORGANIZATION IN ROBOT PARADIGMS

in that case processing is *local* to each function. Other paradigms expect all sensor information to be first processed into one *global* world model and then subsets of the model distributed to other functions as needed.

Overview of the Three Paradigms

In order to set the stage for learning details, it may be helpful to begin with a general overview of the robot paradigms. Fig. I.3 shows the differences between the three paradigms in terms of the SENSE, PLAN, ACT primitives.

HIERARCHICAL PARADIGM

The *Hierarchical Paradigm* is the oldest paradigm, and was prevalent from 1967–1990. Under it, the robot operates in a top-down fashion, heavy on planning (see Fig. I.3). This was based on an introspective view of how people think. "I see a door, I decide to head toward it, and I plot a course around the chairs." (Unfortunately, as many cognitive psychologists now know, introspection is not always a good way of getting an accurate assessment of a thought process. We now suspect no one actually plans how they get out of a room; they have default schemas or behaviors.) Under the Hierarchical Paradigm, the robot senses the world, plans the next action, and then acts (SENSE, PLAN, ACT). Then it senses the world, plans, acts. At each step, the robot explicitly plans the next move. The other distinguishing feature of the Hierarchical paradigm is that all the sensing data tends to be gathered into one global world model, a single representation that the planner can use and can be routed to the actions. Constructing generic global world models

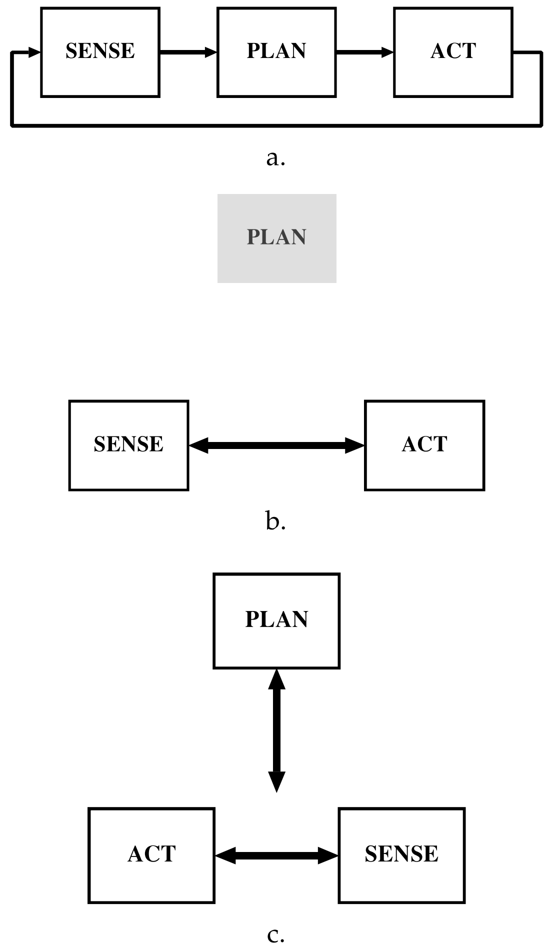

Figure I.3 Three paradigms: a.) Hierarchical, b.) Reactive, and c.) Hybrid deliberative/reactive.

turns out to be very hard and brittle due to the *frame problem* and the need for a *closed world assumption*.

Fig. I.4 shows how the Hierarchical Paradigm can be thought of as a transitive, or Z-like, flow of events through the primitives given in Fig. I.4. Unfortunately, the flow of events ignored biological evidence that sensed information can be directly coupled to an action, which is why the sensed information input is blacked out.

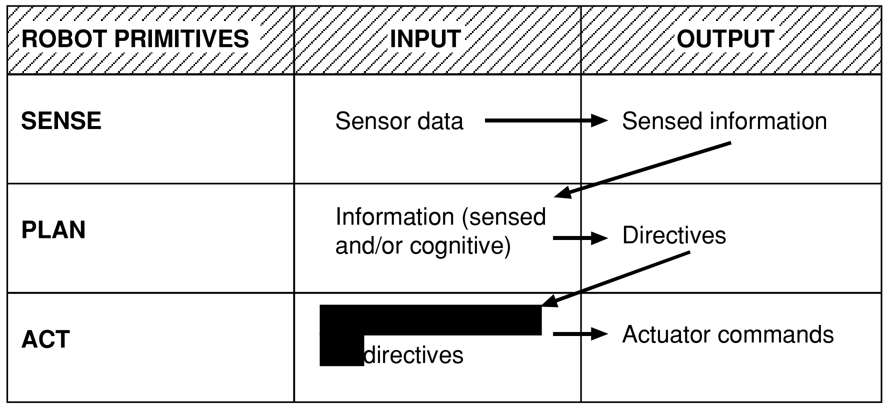

Figure I.4 Another view of the Hierarchical Paradigm.

REACTIVE PARADIGM

The *Reactive Paradigm* was a reaction to the Hierarchical Paradigm, and led to exciting advances in robotics. It was heavily used in robotics starting in 1988 and continuing through 1992. It is still used, but since 1992 there has been a tendency toward hybrid architectures. The Reactive Paradigm was made possible by two trends. One was a popular movement among AI researchers to investigate biology and cognitive psychology in order to examine living exemplars of intelligence. Another was the rapidly decreasing cost of computer hardware coupled with the increase in computing power. As a result, researchers could emulate frog and insect behavior with robots costing less than $500 versus the $100,000s Shakey, the first mobile robot, cost.

The Reactive Paradigm threw out planning all together (see Figs. I.3b and I.5). It is a SENSE-ACT (S-A) type of organization. Whereas the Hierarchical Paradigm assumes that the input to a ACT will always be the result of a PLAN, the Reactive Paradigm assumes that the input to an ACT will always be the direct output of a sensor, SENSE.

If the sensor is directly connected to the action, why isn't a robot running under the Reactive Paradigm limited to doing just one thing? The robot has multiple instances of SENSE-ACT couplings, discussed in Ch. 4. These couplings are concurrent processes, called behaviors, which take local sensing data and compute the best action to take independently of what the other processes are doing. One behavior can direct the robot to "move forward 5 meters" (ACT on drive motors) to reach a goal (SENSE the goal), while another behavior can say "turn 90°" (ACT on steer motors) to avoid a collision

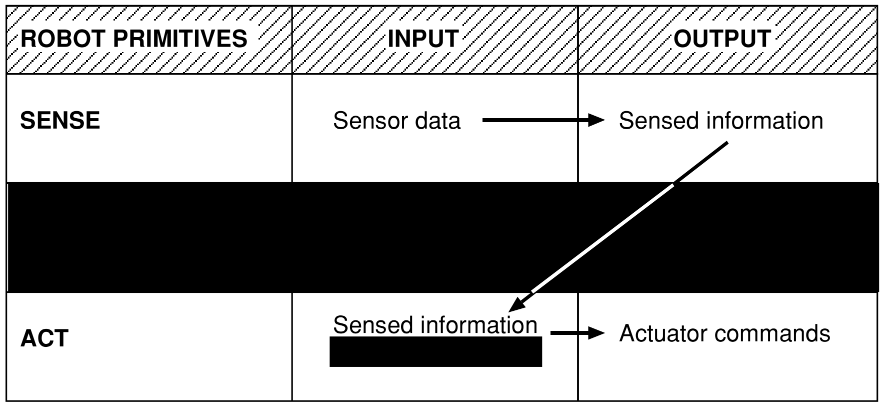

Figure I.5 The reactive paradigm.

with an object dead ahead (SENSE obstacles). The robot will do a combination of both behaviors, swerving off course temporarily at a 45° angle to avoid the collision. Note that neither behavior directed the robot to ACT with a 45° turn; the final ACT emerged from the combination of the two behaviors.

While the Reactive Paradigm produced exciting results and clever robot insect demonstrations, it quickly became clear that throwing away planning was too extreme for general purpose robots. In some regards, the Reactive Paradigm reflected the work of Harvard psychologist B. F. Skinner in stimulus-response training with animals. It explained how some animals accomplished tasks, but was a dead end in explaining the entire range of human intelligence.

HYBRID DELIBERATIVE/REACTIVE PARADIGM

But the Reactive Paradigm has many desirable properties, especially the fast execution time that came from eliminating any planning. As a result, the Reactive Paradigm serves as the basis for the *Hybrid Deliberative/Reactive Paradigm*, shown in Fig.I.3c. The Hybrid Paradigm emerged in the 1990's and continues to be the current area of research. Under the Hybrid Paradigm, the robot first plans (deliberates) how to best decompose a task into subtasks (also called "mission planning") and then what are the suitable behaviors to accomplish each subtask, etc. Then the behaviors start executing as per the Reactive Paradigm. This type of organization is PLAN, SENSE-ACT (P, S-A), where the comma indicates that planning is done at one step, then sensing and acting are done together. Sensing organization in the Hybrid Paradigm is also a mixture of Hierarchical and Reactive styles. Sensor data gets routed to each behavior that needs that sensor, but is also available to the planner

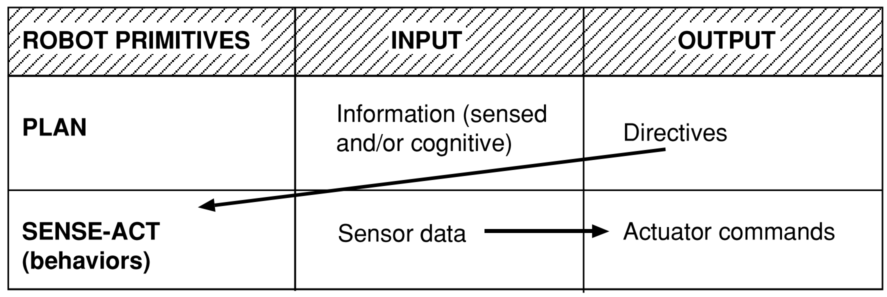

Figure I.6 The hybrid deliberative/reactive paradigm.

for construction of a task-oriented global world model. The planner may also "eavesdrop" on the sensing done by each behavior (i.e., the behavior identifies obstacles that could then be put into a map of the world by the planner). Each function performs computations at its own rate; deliberative planning, which is generally computationally expensive may update every 5 seconds, while the reactive behaviors often execute at 1/60 second. Many robots run at 80 centimeters per second.

Architectures

Determining that a particular paradigm is well suited for an application is certainly the first step in constructing the AI component of a robot. But that step is quickly followed with the need to use the tools associated with that paradigm. In order to visualize how to apply these paradigms to real-world applications, it is helpful to examine representative architectures. These architectures provide templates for an implementation, as well as examples of what each paradigm really means.

What is an architecture? Arkin offers several definitions in his book, *Behavior-Based Robots*.[10] Two of the definitions he cites from other researchers capture how the term will be used in this book. Following Mataric,[89] an architecture provides a principled way of organizing a control system. However, in addition to providing structure, it imposes constraints on the way the control problem can be solved. Following Dean and Wellman,[43] an architecture describes a set of architectural components and how they interact. This book is interested in the components common in robot architectures; these are the basic building blocks for programming a robot. It also is interested in the principles and rules of thumb for connecting these components together.

ARCHITECTURE

To see the importance of an architecture, consider building a house or a car. There is no "right" design for a house, although most houses share the same components (kitchens, bathrooms, walls, floors, doors, etc.). Likewise with designing robots, there can be multiple ways of organizing the components, even if all the designs follow the same paradigm. This is similar to cars designed by different manufacturers. All internal combustion engine types of cars have the same basic components, but the cars look different (BMWs and Jaguars look quite different than Hondas and Fords). The internal combustion (IC) engine car is a paradigm (as contrasted to the paradigm of an electric car). Within the IC engine car community, the car manufacturers each have their own architecture. The car manufacturers may make slight modifications to the architecture for sedans, convertibles, sport-utility vehicles, etc., to throw out unnecessary options, but each style of car is a particular instance of the architecture. The point is: by studying representative robot architectures and the instances where they were used for a robot application, we can learn the different ways that the components and tools associated with a paradigm can be used to build an artificially intelligent robot.

Since a major objective in robotics is to learn how to build them, an important skill to develop is evaluating whether or not a previously developed architecture (or large chunks of it) will suit the current application. This skill will save both time spent on re-inventing the wheel and avoid subtle problems that other people have encountered and solved. Evaluation requires a set of criteria. The set that will be used in this book is adapted from *Behavior-Based Robotics*:[10]

MODULARITY

1. **Support for modularity:** does it show good software engineering principles?

NICHE TARGETABILITY

2. **Niche targetability:** how well does it work for the intended application?

PORTABILITY

3. **Ease of portability to other domains:** how well would it work for other applications or other robots?

ROBUSTNESS

4. **Robustness:** where is the system vulnerable, and how does it try to reduce that vulnerability?

Note that niche targetability and ease of portability are often at odds with each other. Most of the architectures described in this book were intended to be generic, therefore emphasizing portability. The generic structures, however, often introduce undesirable computational and storage overhead, so in practice the designer must make trade-offs.

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.