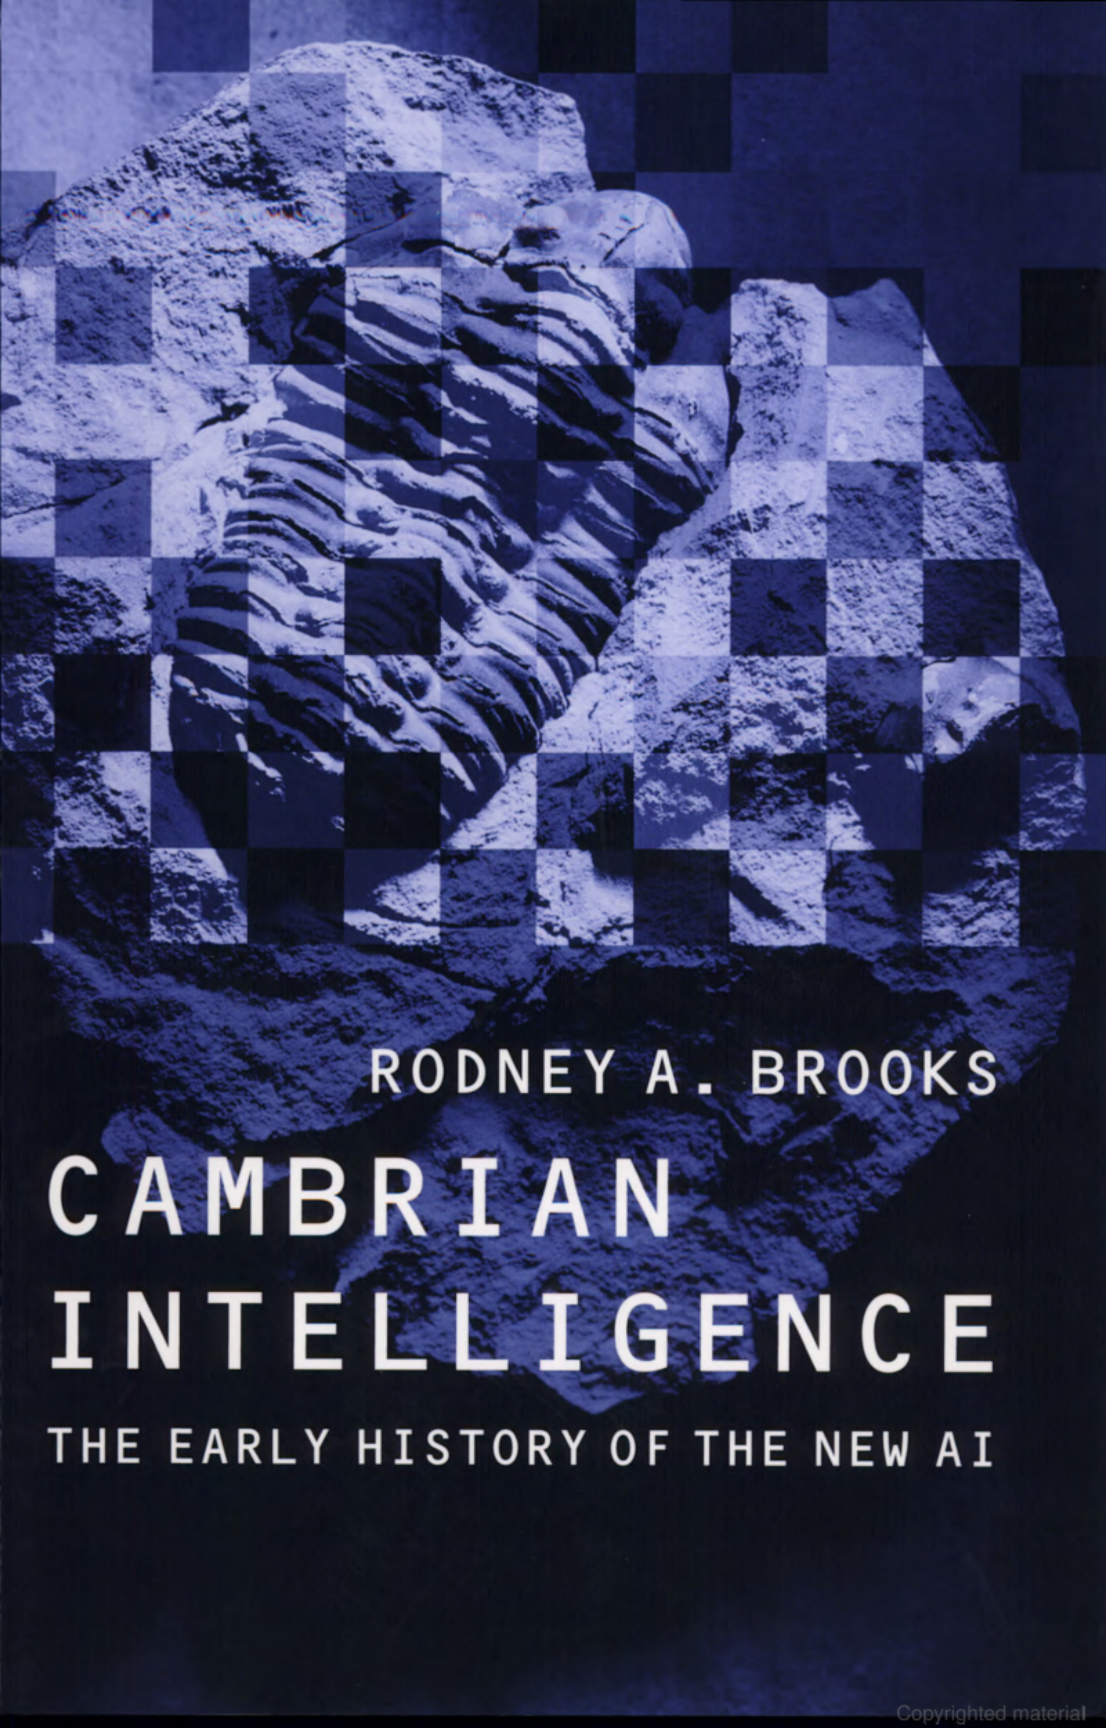

RODNEY A. BROOKS

CAMBRIAN INTELLIGENCE

THE EARLY HISTORY OF THE NEW AI

CAMBRIAN INTELLIGENCE

The Early History of the New AI

RODNEY A. BROOKS

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

in Computer Modern by the author using LaTeX
and bound in the United States of America.

Congress Cataloging-in-Publication Data

Allen.
the early history of the new AI /
Brooks.

Book."
references.
–0-262-52263-2

paper)
2. Artificial intelligence. I. Title.
1999

9 8 7 6 5 4 3 2 99-17220
CIP

Contents

Preface

This is a collection of scientific papers and essays that I wrote from 1985 to 1991, a period when the key ideas of behavior-based robotics were developed. Behavior-based robotics has now been widely adopted for mobile robots of all sorts on both Earth and Mars.

Mobile robotics was an almost non-existent curiosity fourteen years ago. There were many unmanned guided vehicles developed to carry parts in automated manufacturing plants, but they just had to follow guide wires in the ground. In contrast, autonomous mobile robots are conceived of as robots that can operate in worlds where the layout of obstacles are not known ahead of time; they are to operate in the sorts of worlds that people and animals can negotiate with ease. The few autonomous mobile robots that did exist fourteen years ago were programmed to build elaborate three-dimensional world models and to *plan* a long series of actions within those world models before they started to move. Out of necessity the handful of researchers in the area made a restriction that their robots could operate only in static worlds.

The behavior-based approach to mobile robots changed the way robots were programmed, at least at the lower levels of control. The eight chapters of this book are papers that I wrote (one with a student) introducing the behavior-based approach; four are technical papers describing robot systems and approaches, and four are philosophical papers describing the changes in intellectual point of view that enables this approach.

In an idealized scientific world one might think that the new intellectual point of view was developed first and then the technology and technicians caught up and implemented systems guided by those new intellectual points of view. But that is certainly not the way things developed for me. In all cases the technological implementations came first, and the philosophical realizations followed later. For this reason the technical papers are presented first and the philosophical papers second. However, either part of the book can be read first as all the papers were written to be read as single works. Philosophers may safely

model where cognition mediates between
actions.

without finding themselves ungrounded in the

led philosophy there was one key earlier
in nature, that was critical to the development
approach. In hindsight, it may seem to be the
that enabled and outlines the whole approach but I
new approach was not at all obvious at the time of

was that the so-called central systems of intelligence
has been referred to more recently — was perhaps an
and that all the power of intelligence arose from the
and actuation systems. This is the cornerstone
robotics, both mobile robots as they have developed
years and humanoid robots that have been developed

came about after spending a number of years
mostly vision-based perception, and motion
two-dimensional shapes and some for robot

that there would be some intelligent system
together doing so-called *high-level reasoning.*
generally accepted order of things, largely validated

by the early successes of the Shakey robot at the SRI Artificial Intelligence Center in Menlo Park, California, near Stanford University. The higher-level reasoning system of Shakey worked in predicate logic, and the perception system produced descriptions of the world in first-order predicate calculus. However, that work in the late sixties and very early seventies was carried out in a very restricted world; a well-lit area consisting of only matte-painted walls, clean floors, and large cubes and wedges with each face painted a different color. These restrictions made the vision problem fairly straightforward though not trivial.

By the early eighties it seemed to me that if there was to be any validity to this approach then computer vision should be able to deliver predicate calculus descriptions of much less constrained and much more complex worlds. Computer vision systems ought to be able to operate in the ordinary sorts of environments that people operated in, cluttered offices with things stuck on walls and disorderly piles of papers that partially obscured objects. A computer monitor, for instance, should be visually recognizable as such even if it were a new model with a shape and size slightly different from all those that the vision system had been told about or had seen before. Even more challenging, a computer vision system should be able to operate outdoors and pick out trees, hills, pathways, curbs, houses, cars, trucks, and everything else that a three-year-old child could name.

There were no vision systems around doing anything even remotely as sophisticated. Any such dreams had been put on hold by the computer vision community as they worked on very difficult but much simpler challenges. But even so, the validity of the approach was not questioned. Rather, the progress in computer vision was questioned. Alex (Sandy) Pentland, then a young researcher at SRI held a weekly seminar series to try to uncover the sources of the seeming stall in computer vision. The title of the seminar series was *From Pixels to Predicates*, setting the tone for the direction of inquiry. I was a very junior faculty member at Stanford at the time and was invited to give one of the weekly talks during the 1983–84 academic year.

I have kept my crude hand-drawn opening slides from that talk, and they are reproduced in figures 1 and 2. I did not have a cogent argument to make based on these slides. It may have been that I had recently discovered the joy of brashly telling everyone that their most implicit beliefs or assumptions were wrong, or at least open to question; the central systems people were giving us perception people a hard time that we just weren't producing the sorts of outputs that would enable their *intelligent* systems to do all the really hard stuff. Or it may have been that I was extraordinarily frustrated both with the difficulty of producing complete descriptions of the world from visual input, and

that had to be maintained in such models
uccessful motor actions within them. Or I
ever insight. I doubt the latter, and think
ıt two, perhaps dominated by the perverse

ıally three transparencies. An underlying
ve rise to the two different figures. The
the words PERCEPTION, WORLD, and AC-
wo arrows, one from action to the world,
:ception, indicating the causality chain of
;ave the image reproduced in figure 1. In
g, the computational box, that was used
o order up actions. This was very much
ıew of how to build an intelligent system,
ch implicitly promoted by Sandy's semi-
ːrception goes on by itself, autonomously
that are fed to a cognition box that does
ıtiates the real *intelligence* of the system.
ıe action box what to do, in some sort of
ınguage.
the overlay to give the image in figure 2.
approach to intelligence upside down. It
that is devoted to cognitive tasks. Instead
ion and action subsystems do all the work
ıl observer that has anything to do with
g cognitive abilities to a system that works
xplicit place where cognition is done.
:ursive problem of deciding what is in the
ıding yet another little homunculus inside
ɔwer of reason is reduced to a rather dumb
:e operating in such highly abstract terms
truly intelligent action to be generated.
up and made more real as each level of

is whole problem by denying its existence.
idea at the time how to blend the percep-
'ork together and to achieve all the desired
tainly no hint of it in the actual paper I
on the series that Sandy produced.
and I have spent the last fourteen years
gure, by building ever more complex arti-
igent behavior using the second model.

PART I

TECHNOLOGY

CHAPTER 1

A ROBUST LAYERED CONTROL SYSTEM FOR A MOBILE ROBOT

This is by far the most referenced paper that I have written. It both introduced the notion of behavior-based robotics, although that name only came much later, and described the first instance of a robot programmed in this manner. I first gave a talk on this work in Gouvieux-Chantilly, France, in late 1985. Only many years later did I learn that in the back row senior robotics people were shaking their heads asking each other why I was throwing my career away. The content of this paper was shocking because it argued for simplicity rather than for mathematical complexity of analysis and implementation. To this day many people still find the contents of this paper entirely disreputable because it is not filled with mathematical equations. Having spent six years in a mathematics department I am not afraid of mathematics, indeed I revere beautiful mathematics. But I am afraid that many people have severe cases of physics envy *and feel that their work is not complete if it does not have pages of equations, independently of whether those equations shed any light at all on the deep questions. To my mind there has to date been no decent mathematical analysis of the ideas presented in this paper and further developed by many, many researchers during the rise of behavior-based robotics. That is not to say there should not or can not be such an analysis. But I think it will require some very deep insights and can not rely on surface level equation generation.*

Abstract. We describe a new architecture for controlling mobile robots. Layers of control system are built to let the robot operate at increasing levels of competence. Layers are made up of asynchronous modules which communicate over low bandwidth channels. Each module is an instance of a fairly simple computational machine. Higher level layers can subsume the roles of lower levels by suppressing their outputs. However,

---

Originally appeared in the *IEEE Journal of Robotics and Automation, RA-2* April, 1986, pp. 14–23.

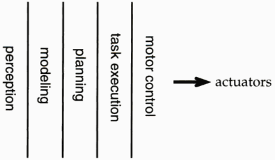

decomposition of a mobile robot control system

to function as higher levels are added. The result
robot control system. The system has been used
robot wandering around unconstrained laboratory
rooms. Eventually it is intended to control
the office areas of our laboratory, building maps
using an onboard arm to perform simple tasks.

for a completely autonomous mobile robot must per-
information processing tasks in real time. It op-
where the boundary conditions (viewing the
control problem in a classical control theory formulation)
rapidly. In fact the determination of those boundary con-
over very noisy channels since there is no straightforward
sensors (e.g., TV cameras) and the form required of
conditions.

approach to building control systems for such robots is
problem into a series (roughly) of *functional units* as
a series of vertical slices in Figure 1. After analyzing the
requirements for a mobile robot we have decided to use
behaviors as our primary decomposition of the problem.
by a series of horizontal slices in Figure 2. As with
we implement each slice explicitly then tie
to a robot control system. Our new decomposi-
different architecture for mobile robot control
different implementation strategies plausible at

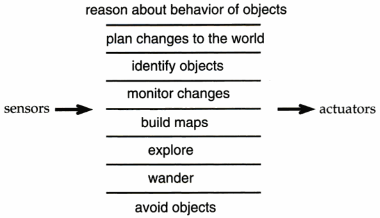

Figure 2: A decomposition of a mobile robot control system based on task achieving behaviors.

the hardware level, and with a large number of advantages concerning robustness, buildability and testability.

1.1 Requirements

We can identify a number of requirements of a control system for an intelligent autonomous mobile robot. They each put constraints on possible control systems we might build and employ.

Multiple Goals: Often the robot will have multiple goals, some conflicting, which it is trying to achieve. It may be trying to reach a certain point ahead of it while avoiding local obstacles. It may be trying to reach a certain place in minimal time while conserving power reserves. Often the relative importance of goals will be context dependent. Getting off the railroad tracks when a train is heard becomes much more important than inspecting the last 10 track ties of the current track section. The control system must be responsive to high priority goals, while still servicing necessary "low level" goals (e.g., in getting off the railroad tracks it is still important that the robot maintains its balance so it doesn't fall down).

Multiple Sensors: The robot will most likely have multiple sensors (e.g., TV cameras, encoders on steering and drive mechanisms, and perhaps infrared beacon detectors, an inertial navigation system, acoustic rangefinders, infrared rangefinders, access to a global positioning satellite system, etc.). All sensors have an error component in their readings.

to build cheap robots which can wander around human
with no human intervention, advice or control and
useful work. Map making is therefore of crucial
when idealized blue prints of an environment are

world is three dimensional; it is not just a two di-
The robot must model the world as three
if it is to be allowed to continue cohabitation with

coordinate systems for a robot are the source of large
errors. Relational maps are more useful to a mobile
the design space for perception systems.

where mobile robots will do useful work are not con-
exact simple polyhedra. While polyhedra may be useful
realistic world, it is a mistake to build a special world
models can be exact. For this reason we will build
environment for our robot.

while easy to collect, does not by itself lead to rich
of the world useful for truly intelligent interactions.
is much better for that purpose. Sonar data may be
for low level interactions such as real time obstacle avoid-

For robustness sake the robot must be able to perform when one or more of its sensors fails or starts giving erroneous readings. Recovery should be quick. This implies that built-in self calibration must be occurring at all times. If it is good enough to achieve our goals then it will necessarily be good enough to eliminate the need for external calibration steps. To force the issue we do not incorporate any explicit calibration steps for our robot. Rather we try to make all processing steps self calibrating.

in building *artificial beings*—robots that can sur-
days, weeks and months, without human assistance, in a
complex environment. Such robots must be self sustaining.

Levels and Layers

possible approaches to building an autonomous intelli-
As with most engineering problems they all start by

decomposing the problem into pieces, solving the subproblems for each piece, and then composing the solutions. We think we have done the first of these three steps differently to other groups. The second and third steps also differ as a consequence.

2.1 Levels of Competence

Typically mobile robot builders (e.g., (Nilsson 1984), (Moravec 1983), (Giralt et al. 1984), (Kanayama 1983), (Tsuji 1985), (Crowley 1985)) have sliced the problem into some subset of:

- sensing,
- mapping sensor data into a world representation,
- planning,
- task execution, and
- motor control.

This decomposition can be regarded as a horizontal decomposition of the problem into vertical slices. The slices form a chain through which information flows from the robot's environment, via sensing, through the robot and back to the environment, via action, closing the feedback loop (of course most implementations of the above subproblems include internal feedback loops also). An instance of each piece must be built in order to run the robot at all. Later changes to a particular piece (to improve it or extend its functionality) must either be done in such a way that the interfaces to adjacent pieces do not change, or the effects of the change must be propagated to neighboring pieces, changing their functionality too.

We have chosen instead to decompose the problem vertically as our primary way of slicing up the problem. Rather than slice the problem on the basis of internal workings of the solution we slice the problem on the basis of desired external manifestations of the robot control system.

To this end we have defined a number of *levels of competence* for an autonomous mobile robot. A level of competence is an informal specification of a desired class of behaviors for a robot over all environments it will encounter. A higher level of competence implies a more specific desired class of behaviors.

We have used the following levels of competence (an earlier version of these was reported in Brooks (1984*a*)) as a guide in our work:

0. Avoid contact with objects (whether the objects move or are stationary).

1. Wander aimlessly around without hitting things.

"Explore" the world by seeing places in the distance which look reachable and heading for them.

Build a map of the environment and plan routes from one place to another.

in the "static" environment.

the world in terms of identifiable objects and per-
related to certain objects.

and execute plans which involve changing the state of
some desirable way.

about the behavior of objects in the world and modify

each level of competence includes as a subset each ear-
Since a level of competence defines a class of
can be seen that higher levels of competence provide
on that class.

Layers of Control

idea of levels of competence is that we can build layers of a
corresponding to each level of competence and simply
layer to an existing set to move to the next higher level of

building a complete robot control system which achieves
It is debugged thoroughly. We never alter that sys-
zeroth level control system, Next we build a another
layer, we call the first level control system. It is able to
data from the level *0* system and is also permitted to inject
the internal interfaces of level *0* suppressing the normal data
layer, with the aid of the zeroth, achieves level *1* competence.
layer continues to run unaware of the layer above it which
interferes with its data paths.
same process is repeated to achieve higher levels of competence.
3.
this architecture a *subsumption architecture*.
we have a working control system for the robot very
soon as we have built the first layer. Additional
later, and the initial working system need never be

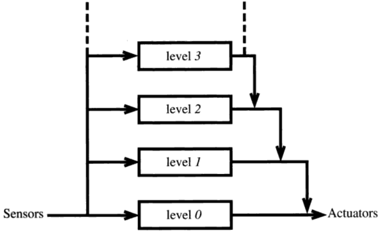

Figure 3: Control is layered with higher level layers subsuming the roles of lower level layers when they wish to take control. The system can be partitioned at any level, and the layers below form a complete operational control system.

We claim that this architecture naturally lends itself to solving the problems for mobile robots delineated in section 1.1.

Multiple Goals: Individual layers can be working on individual goals concurrently. The suppression mechanism then mediates the actions that are taken. The advantage here is that there is no need to make an early decision on which goal should be pursued. The results of pursuing all of them to some level of conclusion can be used for the ultimate decision.

Multiple Sensors: In part we can ignore the sensor fusion problem as stated earlier using a subsumption architecture. Not all sensors need to feed into a central representation. Indeed certain readings of all sensors need not feed into central representations—only those which perception processing identifies as extremely reliable might be eligible to enter such a central representation. At the same time however the sensor values may still be being used by the robot. Other layers may be processing them in some fashion and using the results to achieve their own goals, independent of how other layers may be scrutinizing them.

Robustness: Multiple sensors clearly add to the robustness of a system when their results can be used intelligently. There is another source of robustness in a subsumption architecture. Lower levels which have

n when higher levels are added. Since
he outputs of lower levels by actively
in the cases that it can not produce
r levels will still produce results which
. of competence.

o handle extensibility is to make each
sor. We will see below that this is
airly low bandwidth requirements on
layers. In addition we will see that
e spread over many loosely coupled

lividual layer? Don't we need to de-
itional manner? This is true to some
hat we don't need to account for all
; and generated behaviors in a single
different decompositions for different

from a set of small processors which

te machine with the ability to hold
end messages over connecting "wires".
wledgement of messages. The proces-
ly, monitoring their input wires, and
wires. It is possible for messages to
often. There is no other form of com-
ı particular there is no shared global

to as modules) are created equal in
e is no central control. Each module
n.

ressed and outputs can be inhibited
modules. This is the mechanism by
the role of lower levels.

pecification Language

onents of our layered control architec-
of the modules, and the second is the
n this section we flesh out the details

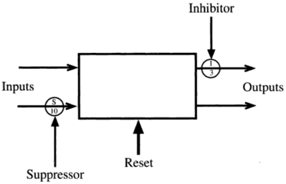

Figure 4: A module has input and output lines. Input signals can be suppressed and replaced with the suppressing signal. Output signals can be inhibited. A module can also be reset to state NIL.

3.2 Communication

Figure 4 shows the best way to think about these finite state modules for the purposes of communications. They have some input lines and some output lines. An output line from one module is connected to input lines of one or more other modules. One can think of these lines as wires, each with sources and a destination.

Additionally outputs may be inhibited, and inputs may be suppressed.

An extra wire can terminate (i.e., have its destination) at an output site of a module. If *any* signal travels along this wire it *inhibits* any output message from the module along that line for some pre-determined time. Any messages sent by the module to that output during that time period is lost.

Similarly an extra wire can terminate at an input site of a module. Its action is very similar to that of inhibition, but additionally, the signal on this wire, besides inhibiting signals along the usual path, actually gets fed through as the input to the module. Thus it *suppresses* the usual input and provides a replacement. If more than one suppressing wire is present they are essentially 'or'-ed together.

For both suppression and inhibition we write the time constants inside the circle.

In our specification language we write wires as a source (i.e., an output line) followed by a number of destinations (i.e., input lines). For instance the connection to the force input of the **avoid** module defined be the wire defined as:

```
(feelforce force) (avoid force))
```

of the **feelforce** module to the input of one control system.
can also be described with a small exten-
we see the suppression of the command
by a signal from the level *1*

```
(turn command) 20.0)))
```

a signal can be connected to the reset input of

Instance

mobile robot control system to achieve levels *0*
as above, and have started implementation of
which exercises the fundamental subsump-
. We need more work on an early vision algorithm
2.

Zeroth Level

layer of control makes sure that the robot does not come
other objects. It thus achieves level *0* competence.
something approaches the robot it will move away.
of moving itself it is about to collide with an object
these two tactics are sufficient for the robot to
obstacles, perhaps requiring many motions, without
The combination of the tactics allows
with with very coarsely calibrated sonars and a
force functions. Theoretically, the robot is not
course, a sufficiently fast moving object, or a very
might result in a collision. Over the course of
hours of autonomous operation, our physical robot (see
collided with either a moving or fixed obstacle. The
have, however, been careful to move slowly.

and **forward** modules communicate with the actual
have extra communication mechanisms, allowing

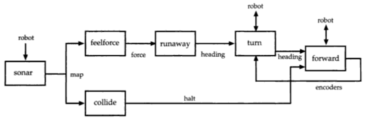

Figure 5: The level *0* control system.

them to send and receive commands to and from the physical robot directly. The **turn** module receives a heading specifying an inplace turn angle followed by a forward motion of a specified magnitude. It commands the robot to turn (and at the same time sends a busy message on an additional output channel illustrate in Figure 7) and on completion passes on the heading to the **forward** module (and also reports the shaft encoder readings on another output line shown in Figure 7) then goes into a wait state ignoring all incoming messages. The **forward** module commands the robot to move forward, but halts it if it receives and message on its halt input line during the motion. As soon as the robot is idle its sends out the shaft encoder readings–the message acts as a reset for the **turn** module, which is then once again ready to accept a new motion command. Notice the any heading commands sent to the **turn** module during transit are lost.

- The **sonar** module takes a vector of sonar readings, filters them for invalid readings, and effectively produces a robot centered map of obstacles in polar coordinates.

- The **collide** module monitors the sonar map and if it detects objects dead ahead it sends a signal on the halt line to the **motor** module. The **collide** module does not know or care whether the robot is moving. Halt messages sent while the robot is stationary are essentially lost.

- The **feelforce** module sums the results of considering each detected object as a repulsive force, generating a single resultant force.

- The **runaway** module monitors the 'force' produced by the sonar detected obstacles and sends commands to the **turn** module if it ever becomes significant.

robot

encoders

level *0* control system augmented with the level *1* system.

a complete description of how the modules are con-

when combined with the zeroth, imbues
to wander around aimlessly without hitting
was earlier as level *1* competence. This control
large degree on the zeroth level's aversion to hitting
it uses a simple heuristic to plan ahead a little in
collisions which would need to be handled by

module generates a new heading for the robot every

described in more detail in section 3, takes the
computation from the zeroth level, and com-
heading to produce a modified heading
in roughly the right direction, but is per-
obvious obstacles. This computation implic-
computations of the **runaway** module, in the
is also a heading to consider. In fact the output
suppresses the output from the **runaway**
it the **motor** module.

a complete description of how the modules are con-
Note that it is simply Figure 5 with some more modules

4.3 Second Level

Level two is meant to add an exploratory mode of behavior to the robot, using visual observations to select interesting places to visit. A vision module finds corridors of free space. Additional modules provide a means of position servoing the robot to along the corridor despite the presence of local obstacles on its path (as detected with the sonar sensing system). The wiring diagram is shown in Figure 7. Note that it is simply Figure 6 with some more modules and wires added.

- The **status** module monitors the **turn** and **forward** modules. It maintains one status output which sends either *hi* or *lo* messages to indicate whether the robot is busy. In addition, at the completion of every turn and roll forward combination it sends out a combined set of shaft encoder readings.

- The **whenlook** module monitors the busy line from the **status** module, and whenever the robot has been sitting idle for a few seconds it decides its time to look for a corridor to traverse. It inhibits wandering so it can take some pictures and process them without wandering away from its current location, and resets the **pathplan** and **integrate** modules–this latter action ensures that it will know how far it has moved from its observation point should any **runaway** impulses perturb it.

- The **look** module initiates the vision processing, and waits for a candidate freeway. It filters out poor candidates and passes any acceptable one to the **pathplan** module.

- The **stereo** module is supposed to use stereo TV images (Grimson 1985), obtained by the robot, to find a corridor of free space. At the time of writing final version of this module had not been implemented. Instead, both in simulation and on the physical robot, we have replaced it with a sonar-base corridor finder.

- The **integrate** module accumulates reports of motions from the **status** module and always sends its most recent result out on its integral line. It gets restarted by application of a signal to its reset input.

- The **pathplan** module takes a goal specification (in terms of an angle to turn, a distance to travel) and attempts to reach that goal. To do this it sends headings to the **avoid** module, which may perturb them to avoid local obstacles, and monitors its integral input which is an integration of actual motions. The messages to

CAMBRIAN INTELLIGENCE

THE EARLY HISTORY OF THE NEW AI

RODNEY A. BROOKS

Until the mid-1980s, AI researchers assumed that an intelligent system doing high-level reasoning was necessary for the coupling of perception and action. In this traditional model, cognition mediates between perception and plans of action. Realizing that this core AI, as it was known, was illusory, Rodney A. Brooks turned the field of AI on its head by introducing the behavior-based approach to robotics. The cornerstone of behavior-based robotics is the realization that the coupling of perception and action gives rise to all the power of intelligence and that cognition is only in the eye of an observer. Behavior-based robotics has been the basis of successful applications in entertainment, service industries, agriculture, mining, and the home. It has given rise to both autonomous mobile robots and more recent humanoid robots such as Brooks's Cog.

This book represents Brooks's initial formulation of and contributions to the development of the behavior-based approach to robotics. It presents all of the key philosophical and technical ideas that put this "bottom-up" approach at the forefront of current research in not only AI but all of cognitive science.

Rodney A. Brooks is the Fujitsu Professor of Computer Science and Engineering and Director of the Artificial Intelligence Laboratory at the Massachusetts Institute of Technology.

A Bradford Book

The MIT Press
Massachusetts Institute of Technology
Cambridge, Massachusetts 02142
http://mitpress.mit.edu

9 780262 522632 90000

BROBP 0-262-52263-2