

Louise F. Gunderson
James P. Gunderson

A decorative graphic consisting of a grid of vertical white lines of varying lengths, creating a barcode-like effect. On the left side, there are two solid orange rectangular blocks of different sizes, one above the other, partially overlapping the grid.

Robots, Reasoning, and Reification

 Springer

ROBOTS, REASONING, AND REIFICATION

ROBOTS, REASONING, AND REIFICATION

“Where is my Robot?”

L.F. Gunderson

Gamma Two, Inc.
Denver, Colorado, USA

J.P. Gunderson

Gamma Two, Inc.
Denver, Colorado, USA

 Springer

Authors

L.F. Gunderson
Gamma Two, Inc.
209 Kalamath Street
Unit 13
Denver, CO
80223, USA
lgunders@gamma-two.com

J.P. Gunderson
Gamma Two, Inc.
209 Kalamath Street
Unit 13
Denver, CO
80223, USA
jgunders@gamma-two.com

ISBN: 978-0-387-87487-6 e-ISBN: 978-0-387-87488-3

DOI: 10.1007/978-0-387-84816-7

Library of Congress Control Number: 2008934306

© 2009 Springer Science+Business Media, LLC

All rights reserved. This work may not be translated or copied in whole or in part without the written permission of the publisher (Springer Science+Business Media, LLC, 233 Spring Street, New York, NY 10013, USA), except for brief excerpts in connection with reviews or scholarly analysis. Use in connection with any form of information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed is forbidden.

The use in this publication of trade names, trademarks, service marks, and similar terms, even if they are not identified as such, is not to be taken as an expression of opinion as to whether or not they are subject to proprietary rights.

Printed on acid-free paper

springer.com

While in principle everything may be under strict control within the machine, the remote space-time surroundings are in the general case known to the system by extrapolation only, that is predicted with some uncertainty. As psychological functionalism, when actually carried out, has thus been found to be forced into probabilism, a cybernetics with ecological involvement must contain probabilistic elements. – Egon Brunswik, 1950

Preface

This work was created from the statement “But, all you have to do is make the robot recognize its surroundings. Salamanders do it, and how complex are they?” Little did we know what a long path was started with those simple words. This book is a small step on that path, which we hope leads to robots that can serve as true and useful assistants to humans. At the least, we hope for some help with the tasks that are described by the 3 d**** words (dull, dirty, or dangerous).

Fair warning, this work is a synthesis of ideas from many disciplines. As such, we have depended on the work of many other researchers and philosophers. The heart of this work, the lens model, comes from the work of Egon Brunswik. Even though he died in the 1950’s, his ideas are still strong enough to resonate into the 2000’s and into our robot. Another researcher who’s work has greatly influenced this work is Walter Freeman, Professor Emeritus of Neurobiology at the University of California, Berkeley. We have relied heavily on his work on preafference and attention to guide the development of our robot. In addition, we have used research from a myriad of different fields. Our huge thanks to all the researchers who’s work we used to synthesize this new theory.

Denver, CO
July 2008

Louise F. Gunderson
James P. Gunderson

Acknowledgements

This book would not have come into being without the support of a number of people. Much of what we are presenting here is the outgrowth of discussions with some really intelligent people at conferences and workshops, around the world. High on that list are the people who put together and orchestrate the Performance Metrics for Intelligent Systems (PerMIS) workshop, sponsored by the National Institutes of Standards and Technology. This includes an enormous amount of work by Elena Messina, Raj Madhava, Jim Albus, and Alex Meystel. We particularly want to thank Alex, whose encouragement has meant the world to both of us.

Many thanks to Christian Brown and Randi Himelgrin. Without your support, this book would never have happened. Chris put in untold hours of hard work, soldering circuit boards, chasing robots, and above all participating in endless design meetings, and asking the hard questions. Much of what we have developed in testing methodology was a direct result of long discussions with Randi about the state of testing in the real world.

Jim

Much of the chapter on surviving in a dynamic world resulted from a seemingly innocent question posed by my thesis advisor at the University of Virginia: Worthy Martin. He asked “What do you mean it can deal with the unexpected?” Thanks, Boss!

Louise

Thanks to Ellen Bass for introducing me to the work of Egon Brunswik and to Don Brown for being patient while I worked out the details of judgment analysis in my dissertation.

Contents

| | | |
|----------|--|----|
| 1 | Introduction | 1 |
| 1.1 | Bridging the Gap | 2 |
| 1.1.1 | Bidirectional Mapping | 3 |
| 1.2 | Reification and Preference in Biological Entities | 4 |
| 1.3 | More Advanced Brains | 5 |
| 1.4 | What This Book Is and What It Is Not | 6 |
| 1.5 | Structure of the Book | 7 |
| 1.6 | A Note on Typefaces and Terminology | 7 |
| 1.6.1 | Anthropomorphization | 8 |
| 2 | Some background material on probability and biology | 9 |
| 2.1 | Layout | 9 |
| 2.2 | Probability in the Real World | 10 |
| 2.3 | Why a Biologically Principled Argument? | 11 |
| 2.3.1 | Biological Principles | 13 |
| 2.4 | What Is a Biologically Principled Argument? | 14 |
| 2.4.1 | Biology Is an Observational Science | 14 |
| 2.4.2 | Life Has Structure | 16 |
| 2.4.3 | The Theory of Evolution Explains the Observed Diversity of Life | 17 |
| 2.5 | So Why Is Our Model Biologically Principled? | 18 |
| 2.5.1 | Why Not Just Use Expected Value? | 19 |
| 3 | Using Cognition and Physiology to Build a Cognitive Model | 21 |
| 3.1 | Reification in Biological Entities | 21 |
| 3.1.1 | Recognition | 22 |
| 3.1.2 | Preference | 23 |
| 3.2 | Biological Storage | 26 |
| 3.2.1 | Explicit Memory | 27 |
| 3.3 | Emotion | 29 |
| 3.3.1 | Emotion as mediator | 29 |

- 4 Representation** 31
 - 4.1 Representing Features of the World 33
 - 4.2 Representing Goals 34
 - 4.3 Representing Actions in the World 35
 - 4.3.1 Enabling Conditions 35
 - 4.3.2 Outcomes 35
 - 4.3.3 Representing Likelihoods 36
 - 4.4 Exogenous Events 37

- 5 Perception/Action System** 39
 - 5.1 Robot as Perception/Action System 40
 - 5.1.1 Robot as Body 41
 - 5.1.2 Robot as Sensor 42
 - 5.1.3 Robot as Agent of Change 43
 - 5.1.4 Low Level Control Loop - Procedural Memory 45
 - 5.1.5 System Safety and Routine Actions 47
 - 5.2 Examples of Perception/Action Systems 47
 - 5.2.1 Fred - a simple test robot 47
 - 5.2.2 Basil 51
 - 5.3 Summary of Perception/Action Systems 54

- 6 Design of a Reification Engine** 57
 - 6.1 Model Selection Criteria 57
 - 6.2 Judgment Analysis 59
 - 6.3 Designing the Reification Engine 62

- 7 Bridging the Sensor to Symbol Gap** 65
 - 7.1 Supporting Bidirectional Mapping 66
 - 7.1.1 A Third Approach 68
 - 7.2 Reification Architecture 68
 - 7.3 PerCepts and Reification 71
 - 7.3.1 PerCept Data 73
 - 7.3.2 PerCept Function 76
 - 7.4 Mental Model 77
 - 7.5 Current World State 79
 - 7.6 Reification functionality 80
 - 7.6.1 Initialization 80
 - 7.6.2 Mapping the World onto its Model - Recognition 81
 - 7.6.3 Projecting the Model onto the World - Preference 82
 - 7.6.4 Updating the Current World State 84
 - 7.7 Wrapping Up Reification 84

- 8 Working Memory and the Construction of Personal Experiences** 87
 - 8.1 Transient Memory 88
 - 8.1.1 Working Memory and the Current World State 89
 - 8.1.2 Internal State 92
 - 8.2 Episodic Memory 92
 - 8.2.1 Emotive Tags 96
 - 8.3 Memory Services 97
 - 8.4 Providing Memory Services to the Reification Process 98
 - 8.5 Memory, What Was That Again? 98
- 9 Semantic Memory and the Personal Rough Ontology** 101
 - 9.1 Semantic Memory 101
 - 9.1.1 What is a Personal Rough Ontology? 102
 - 9.2 Building Semantic Memory 104
 - 9.2.1 Structure of the Ontology 106
 - 9.2.2 The nodes in the multi-graph 107
 - 9.2.3 Relationships, the Edges of the Graph 112
 - 9.2.4 A Note on Representing Probabilities 114
 - 9.3 Persistent Storage in the Personal Rough Ontology 115
 - 9.4 Transient versus Persistent Knowledge 115
 - 9.5 Extracting Problems for the Deliberative System 117
 - 9.6 Focusing Attention by Finding Sub-Ontologies 117
 - 9.6.1 Weighted Transitivity 118
- 10 Deliberative System** 121
 - 10.1 Deliberation 122
 - 10.2 Reasoning About the Present 124
 - 10.2.1 Sense-Symbols from the Reification Engine 125
 - 10.2.2 Symbols from the Ontology 125
 - 10.2.3 Internal State 126
 - 10.2.4 Reasoning with WorldSets 126
 - 10.3 Choosing the Future 128
 - 10.3.1 Planning as Search 129
 - 10.3.2 Adapting to Failure 134
 - 10.4 Plan Evaluation and Selection 135
 - 10.4.1 Acquiring Distributions 135
 - 10.4.2 Simulator Fidelity 136
 - 10.5 Summary 137
- 11 Putting it All Together** 139
 - 11.1 How it Fits Together 139
 - 11.2 Goals and Environment 141
 - 11.3 Knowledge Sources 143
 - 11.3.1 Ontological Knowledge 144
 - 11.3.2 Reification Knowledge 148

- 11.3.3 Perception/Action Knowledge 148
- 11.4 The process 149
 - 11.4.1 Perception/Action 149
 - 11.4.2 Reification 149
 - 11.4.3 Execution 149
 - 11.4.4 Deliberation 150
 - 11.4.5 Execution, Reification and Action 151
 - 11.4.6 Perception/Action - Reflex 152
 - 11.4.7 Execution Failure 152
 - 11.4.8 Back Up to Deliberation 152
 - 11.4.9 Procedural Memory and Localization 153
- 11.5 A Few Notes About the General Flow 154
- 12 Testing** 157
 - 12.1 Testing the Robot, or How Does One Test an Embedded System? .. 157
 - 12.2 eXtreme Programming 158
 - 12.3 Methodology for Testing Embodied Systems 159
 - 12.3.1 Benefits of Partitioning the Tests 161
 - 12.4 General Testing Guidelines 162
 - 12.4.1 General Partitioning Guidelines 162
 - 12.5 Testing in the lab 163
 - 12.5.1 Hardware 163
 - 12.5.2 Static Tests 164
 - 12.5.3 Dynamic tests 167
 - 12.6 Formal System Tests - Testing In The Real World 169
 - 12.6.1 Testing Recognition 169
 - 12.6.2 Testing Preference 174
 - 12.6.3 Testing Self-Localization 179
 - 12.7 Summary 181
- 13 Where do we go from here** 183
 - 13.1 A Stopping Point 183
 - 13.2 Next Steps 185
 - 13.2.1 Adding Learning to the Model 185
 - 13.2.2 Adding Additional Data Sources 185
 - 13.2.3 Porting the Brain into New Bodies 186
- Glossary** 187
 - References 190
- Index** 199

Acronyms

List of abbreviations, symbols, and acronyms

| | |
|---------|---|
| ATV | All Terrain Vehicle |
| CNS | Central Nervous System |
| CWS | Current World State |
| DARPA | Defense Advanced Research Projects Agency |
| EET | Estimated Execution Time |
| fMRI | functional magnetic resonance imaging |
| GPS | Global Positioning System |
| I/O | Input / Output |
| IR | Infra-Red, typically referring to sensors |
| LIDAR | Light Detection and Ranging |
| MCMC | Markov Chain Monte Carlo [simulation] |
| MRM | MultiResolution Modeling |
| NASA | National Aeronautics and Space Administration |
| OODA | Observe - Orient - Decide - Act [loop] |
| POC | Probability of Occurrence |
| ProPlan | Probability–Aware Planning and Execution System |
| PWM | Pulse Width Modulation |
| RAM | Random Access Memory |
| RDF | Resource Description Framework |
| ROF | Retry On Failure |
| SAM | Surface to Air Missile |
| UART | Universal Asynchronous Transmitter Receiver |
| UAV | Unmanned Aerial Vehicle |
| UML | Unified Modeling Language |
| XML | eXtensible Markup Language |
| XP | eXtreme Programming |

Chapter 1

Introduction

Where is my robot?

You know - the one that acts like the ones in the movies; the one that I just tell what to do, and it goes out and does it. If it has problems, it overcomes them; if something in the world changes, it deals with the changes. The robot that we can trust to do the dirty, dangerous jobs out in the real world - where is that robot? What is preventing us from building and deploying robots like this? While there are a number of non-trivial and necessary hardware issues, the critical problem does not seem to be hardware related. We have many examples of small, simple systems that will (more or less) vacuum a floor, or mow a lawn, or pick up discarded soda cans in an office. But these systems have a hard time dealing with new situations, like a t-shirt tossed on the floor, or the neighbor's cat sunning itself in the yard. We also have lots of teleoperated systems, from Predator aircraft, to deep sea submersibles, to bomb disposal robots, to remote controlled inspection systems. These systems can deal with changes to the world and significant obstacles provided that one or more humans are in the loop to tell the robot what to do.

So, what happens when a person takes over the joystick, and looks through the low-resolution, narrow field of view camera of a perimeter-patrol security robot? Suddenly, where the robot was confounded by simple obstacles and easy to fix situations, the teleoperated system is able to achieve its goals and complete its mission. This is despite the fact that in place of a tight sensor-effector loop, we now have a long delay between taking an action and seeing the results (very long in the case of NASA's Mars rovers). We have the same sensor data, we have the same effector capabilities, we have added a massive delay yet the system performs better. Of course, it is easy to say that the human is just more intelligent (whatever that means), but that does not really answer the question. What is it that the human operator brings to the system?

We believe that a major component of the answer is the ability to reify: the ability to turn sensory data into symbolic information, which can be used to reason about the situation, and then to turn a symbolic solution back into sensor/effector actions that achieve a goal. This bridging process from sensor to symbol and back is the focus of this book. Since it is the addition of a human to the system that seems to

enable success, we draw heavily from current research into what biological systems (primarily vertebrates) do to succeed the world, and how they do what they do. We look at some research into cognition on a symbolic level, and research into the physiology of biological entities on a physical (sensor/effector) level. From these investigations we derive a computational model of reification, and an infrastructure to support the mechanism. Finally, we detail the architecture that we have developed to add a reification to existing robotic systems.

1.1 Bridging the Gap

There has long been a gulf between artificial intelligence researchers who focus on deliberative symbol manipulation and those who focus on embedding control systems into robots. Much of this gulf has been ascribed to the different approaches, working from the symbolic down versus working from the control system up. The general consensus has been that as the two ends work toward the middle, the gulf will narrow and narrow until it disappears. Underlying both these beliefs is the assumption that once the core research is addressed, it will just be a matter of pushing the research frontier toward the opposing viewpoint until they meet. If one continues the bottom up (or top down) approach long enough, eventually one gets to the top (or bottom) and the complete problem is solved. However, recent research has suggested that the gulf may not be bridgeable by work from either side, rather it may require a specific research approach that is different from either the sensor-based or the symbolic domains.

From the viewpoint of the embedded systems approach, the critical task is the recognition of physical and perceptual cues, while mapping those cues onto a symbol system is outside the scope of the research. From the point of view of the deliberative approach, a symbol manipulation system is developed, and it is outside the scope of the symbol system to recognize the physical and perceptual characteristics that define the thing referred to by the symbol. A purely deliberative system might be manipulating abstract strings such as 'block' and 'red'. These abstract symbols have no meaning other than the allowed manipulations in the symbol system. However, if these symbols are meant to refer to real-world objects or characteristics (e.g., if the things referred to have concrete or material existence) then the symbols must correspond to objects in the real world to be effectively used. In recent research the terms *symbol grounding* and *symbol anchoring* have been used to describe the process as well.

In a recent paper by Coradeschi and Saffiotti[38], the argument is made that the Symbol Grounding problem, as presented by Harnad[93], has features in common with Pattern Recognition. Coradeschi and Saffiotti argue that these two problems have an area of overlap (See Figure 1.1A), which also overlapped with the anchoring problem. However, it is more likely that there is in fact no such area of overlap, and that the process of anchoring or reification spans the gap between these two domains, as in Figure 1.1B.

The term reification is taken from philosophy and is defined[134] as “the process of regarding or treating an abstraction or idea as if it had concrete or material existence.” Reification is a two way process, because there are two primary information flows that must be maintained to effectively connect symbols to objects: one is the flow from objects in the physical world onto the symbols, the second is from the symbols onto the objects. This problem is compounded by the fact that a symbol system typically does not have direct access to the objects in the physical world except via the mediation of the perceptual system.

1.1.1 Bidirectional Mapping

To be effective the reification bridge must be capable of answering two fundamental questions:

1. How will symbols appear in the my sensors; and
2. How will this sensor pattern correspond to a symbol?

These correspond to the two functions that a reification system must provide (See Figure 1.2). If the deliberative system has a reachable goal to achieve and a collection of operators that it can apply to modify the world, it can (with sufficient time and computational resources) find a sequence of actions or set of behaviors to achieve that goal. This has been a solved problem since the earliest days of artifi-

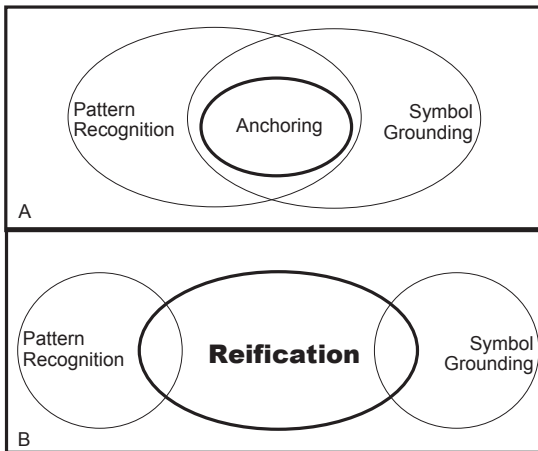


Fig. 1.1 Possible relationships of pattern recognition, symbol grounding, and reification. In A, the problem of anchoring symbols to sensor/action patterns should be approachable by either top-down or bottom-up improvements. However, in B the problem cannot be solved by either top-down or bottom-up approaches, since there is no area of overlap. Rather, a third approach is required, one that solves the reification problem first, which then provides the bridge between symbol and sensors.

cial intelligence research. However, to achieve this goal in the real world the system must be capable of finding the things in the real world needed to achieve each of the actions. It is one thing to produce the step “Pick up the red block from on top of the blue table,” it is quite another know what the sensor pattern that corresponds to the symbols will appear like to its sensors; to find the red block in the real world and grasp it. To be effective, the system needs the ability to build a sensor map that corresponds to the expected state of the things in the real world. This map must correspond to the symbols in the internal model. This is the process of determining how symbols will appear in sensor data, and is one necessary function. The biologic equivalent of this ability is prefference which will be discussed in detail in Chapter 3.1.2.

The second necessary function is the ability to create symbols out of the sensor data. If one has a robot tasked to deliver mail around the office, it needs to be capable of noticing the stairs as stairs, not as a series of parallel lines on a level floor. Failure to correctly put the sensor data into a semantic context can result in the robot tumbling down the stairs, when it thought it was simply crossing a decoration on the floor. Without this ability, it is not possible for the perceptual system to recognize exogenous changes to the world, which must be recognized to either take opportunistic advantage of conditions or to avoid problems which crop up after the plan has been put into effect. This is the symbol grounding problem, which we call recognition, and will also be discussed in more detail later. These two basic functions seem to be features common to almost all vertebrate brains. So it seems reasonable to begin by looking at the research into primitive vertebrate cognition.

1.2 Reification and Prefference in Biological Entities

For any species to survive, the members of that species must be able to sense and manipulate their environment so as to find food, avoid predators, and reproduce. In the case of vertebrate species, these survival mechanisms require the ability to map sensory data onto a neuronal representation, and to take the resulting behavior choices and map those onto motor actions. They must perform this bidirectional




Fig. 1.2 Reification provides a bidirectional mapping between the symbol system used by the deliberative system and the sensor based system.

mapping between the sensory-motor systems and the (potentially primitive) deliberative system. Discussing only the problem of finding food, they must be able to discover how their perceptions of the environment relate to the presence of food. For extremely simple, non-vertebrate species (e.g., amoebas) this might be a purely reactive mapping between chemical sensors on the surface and a gradient ascent behavior. However, for more complex (e.g., vertebrate) species, there is a mapping between the perception of sensory information, and some neuronal representation that is manipulated to assure survival. This is the process of recognition. Conversely, this vertebrate organism must be able, after sensing hunger, to know what features of the environment to use in the search for food. Current research indicates that this is done by priming the sensory cortex with the sensations to expect after taking goal directed action. This is the process of prefference. The combination of these two processes is called *reification*. Both of these processes are discussed in more detail later.

While it is clear that humans can reify, it has been argued that more primitive biological entities are simply “hardwired” reactive systems - they simply respond to a stimulus without any cognition. However, it can be postulated that, in a changing environment, an organism that relies only on an inherited reactive system will be at a disadvantage to one that can reify. If this is true, one would expect to see reification in very primitive organisms. This leads to the question “How complex does a brain have to be before it can reify?”

Salamanders have been used for decades by scientists researching brain function. While the nervous systems of all vertebrates have a common structural plan, the salamanders and their allied species have preserved a type of brain structure which closely resembles that of the most primitive amphibians[97]. These brains have most of the critical functional areas that are shared by all vertebrates, yet their brains are simple enough to allow clear research results. For example, amphibians do have specialized, hardwired prey recognition cells, which allow for the recognition of an object as potential prey[172]. This would suggest that they have the structure of a hardwired reactive system. However, at least one amphibian, the tiger salamander, can be trained to recognize a new scent, which implies that they are capable of reification of new sensory input[51]. Reification occurs at a very low level in the vertebrate brain. The reification methodology described in this book is guided by the example of these very primitive brains. It will be described in more detail in Chapter 3.

1.3 More Advanced Brains

Of course, it might be claimed that these simple creatures use this primitive process, but we humans are more sophisticated and rely on a more advanced mechanism to do the same thing. One of the reasons that the tiger salamander brain was chosen as a model, is that the core functions of all mammalian brains (including ours) have the same structural components as this primitive brain. It is clear that humans have

some sort of a reification mechanism. Artists have long known that we interpret visual images into familiar (if distorted) representations. One practice to overcome this mapping from the distal image to a distorted proximal image is to inverting the images, and then draw the upside down image. This allows the artist to duplicate what is actually there, rather than the interpreted image. Psychologists and philosophers have addressed this non-conscious automatic mechanism for much longer than artificial intelligence has been a discipline:

We do not see patches of color, but trees and houses; we hear, not indescribable sound, but voices and violins[119].

It is clear that, in humans, the conscious mind deals not with low-level sensor data, but with symbols. It is also clear that when we look for things in the environment we do not look for “three orthogonal rectilinear surfaces of similar dimension, with a reflective electromagnetic signal with a wavelength of approximately 650 nanometers.” Instead we look for the red block, and some non-conscious mechanism translates this into the sensory/perceptual indicators that can be used to recognize the block when we see it.

1.4 What This Book Is and What It Is Not

In this book, we construct a framework that can be used for the construction of a biologically principled cybernetic brain. We use a mathematical model from cognitive science to construct the *Reification Engine*. Freeman, among others, has proposed that only a true working neuronal model of a brain can extract semantics from sensory information[70]. While such a neuronally based model might be necessary to build a human level intelligence, we believe that for simpler intelligences, this level of fidelity is not required. Therefore, this book does not contain an attempt to build a working neuronal model. Readers who are interested in that type of work should look at work by Kosma[117] or Edelman[61], among others.

However, we do not believe that it is sufficient to simply describe an architecture or a framework that might achieve a gain in intelligence. There is an enormous gulf between the design and the reality, and the discipline of engineering is based on bridging that gulf. Except where we specifically call out otherwise, all of the theories and designs we present have been encoded and tested on an actual robot. We have found that the practice of embodying the architecture has exposed problems that can cause the design to fail. Among the aspects of the design that we have not yet implemented is the learning side of the overall loop.

In addition, in order to be able to proceed with confidence and make claims about the ability of the system, it is necessary to have complete confidence in the underlying system. During the development of the software and hardware we have made extensive use of automated testing. In addition we have done testing in the robot’s ecosystem. If the test case requires the robot to travel across the room, and return, we must wait, patiently (or not), for the robot to trundle there and back again.

We feel that this level of testing is required to demonstrate that reification gives a robot enough semantic intelligence to reason effectively and achieve goals in the real world.

1.5 Structure of the Book

The first two chapters examine the clues provided by the brain structure of living systems. In Chapter 2, we discuss some background material in biology and probability that is needed for the next chapter, in which we discuss the brain structure of land vertebrates. Chapters 4 – 5 describe the computational framework needed to support the `Reification Engine` and look at the embodiment of sample robots. A cognitively based mathematical model, taken from the work of Egon Brunswik, is described in Chapter 6. This is used to produce the design of the `Reification Engine` which will be used to bridge the gap from the ‘world-as-perceived’ to the ‘world-as-modeled’. We merge current research by cognitive scientists, neurophysiologists, and researchers into artificial intelligence with this model in Chapter 7 to complete the design for the `Reification Engine`. In Chapters 8 – 10, the remaining cognitive structures required for a cybernetic brain (memory and a deliberative system) are discussed. In Chapter 11, the construction of the cybernetic brain from its constituent parts is discussed. In Chapter 12, we discuss the unit testing used in the construction of the brain and the specific robotic testing done to validate the claims of this book. Finally, in Chapter 13, we draw conclusions and discuss future work. This future work includes the need for the robotic system to be able to learn from and adapt to changes in the world, including the ability to add new types of knowledge to its model of the world, and to be able to recognize and reason about new objects, tasks, and goals.

1.6 A Note on Typefaces and Terminology

The construction of the reification system draws on research from many different disciplines, and each of these has its own terminology. Regardless of the background of the reader, there will almost certainly be terms of art used in this book that are unfamiliar. We have tried to compile a glossary of the less well known terms, and when a term of art is introduced, we have generally called it out by using *emphasis*. If the term is unfamiliar, please take a quick look at the glossary (located just before the reference section), to make sure that we are using it in the way you expect.

1.6.1 Anthropomorphization

Anthropomorphization is defined as “The attribution of human motivation, characteristics, or behavior to nonhuman organisms or inanimate objects.” As you read this you will see that we often refer to the robots as ‘he’ or ‘she.’ This is due to a number of things, but high on the list is the fact that humans ascribe human characteristics to many of the inanimate objects in their environments. If we treat our cars as human-like, how much more should we anthropomorphize the human-like robots that we are attempting to create. This process has one significant effect: it defines our expectations of the object. Since we are designing an intelligent, autonomous robot, we have similar expectations of the behavior of the robot and the behavior of a servant. We will try to keep it to a minimum, but I am sure we will miss a few references.

Trying to keep track of the many aspects of the biologically inspired design, the implementation, and the concepts can be difficult. This is especially true when we may be referring to a term from neuroanatomy, one from psychology, and a similar one from the actual software that we built to embody the mechanism, all in the same sentence. We have tried to be consistent with the use of different typefaces to call out the various aspects. In general:

- Normal text is used for the body of the work, and for most psychological or neuroanatomical terms, once they have been introduced;
- *Italics* are used to introduce a new term, or to set off the *concept* from the thing;
- `Typewriter face` is used when we are referring to a software component; and finally,
- ‘single quotes’ are used to indicate a conceptual entity as opposed to the physical thing it refers to, and to set off one term from another, when the context is so complex that we need an additional mechanism (we have tried to keep this to a minimum, really).

Chapter 2

Some background material on probability and biology

In this book, we build a cognitive model that can deal with an uncertain and constantly changing universe. We have ample evidence that living organisms have this ability, and rely on it for daily survival. Rather than reinventing biology whenever possible, we use living creatures, such as salamanders or other primitive land vertebrates, to act as the design guide for the cognitive modules that must be present to create a successful autonomous robot. However, in order to take advantage of these biological examples, some of the basic assumptions that underly biology must also be discussed. The focus of this chapter is summarized by the following questions:

- How do living systems deal with a probabilistic universe?
- How can we discuss these models in a principled way?
- If we are going to use living systems as our guide, why use salamanders and rats instead of humans?

2.1 Layout

The general layout of this chapter is this: In Section 2.2, the features of the real world that make a probability-aware system important are discussed. Since all natural organisms live in this probabilistic environment, it makes sense to look at how these systems achieve the kind of performance that we desire in our robots. However, if we are going to derive our design from these biological entities, we need to explore the concept of a biologically principled argument, rather than the engineering approach of a mathematically principled argument. The need for a biologically principled argument is discussed in Section 2.3. The way in which an argument can be constructed to make it biologically principled is presented in Section 2.4. Finally, The reasons that we believe our model to be biologically principled are discussed. This section also includes a discussion of the conservation of the traits that are important to the success of a species in a dynamic and uncertain environment, and we will take a brief look at why we have chosen a biologically principled path, rather than using a mathematically principled mechanism..