

# Preface

I believe you will find this book has a refreshing new perspective on a subject that needs a new perspective. In particular, these are some of the unique viewpoints that are found in this book:

First, the book takes the viewpoint that symbol processing techniques are not working out very well and these methods alone will never be able to give the depth and breadth of capabilities found in human beings. I believe the solution is to introduce new foundational principles to AI and these new principles get a prominent place in this book. At the moment, the new principles that are under consideration are: connectionist/neural networking methods, case-based and memory-based methods, and picture processing. Even though there is a movement to include quantum mechanics, I do not cover this possibility except to mention it in passing. All this is not to say that symbol processing can be thrown out, only that a good solution to intelligence certainly consists of many more techniques than just symbol processing.

Second, that the methods of AI are best seen as different ways of doing pattern recognition. One way to do pattern recognition is simply to store lots of cases and when a new problem comes up you compare it to all your stored cases. At the other end of the spectrum, Classical Symbol Processing AI has taken the viewpoint that cases should be condensed or compressed down to a relatively small set of rules and then the cases should be thrown away. A system should work with only its condensed knowledge. In between these two extremes are neural networks, especially backprop type networks, where the networks develop rulelike capabilities but they can also store cases. As much as possible, the book compares these three basic methods using actual AI programs.

Third, the structure of the book starts at the bottom of human abilities with vision and other simple pattern recognition abilities and moves on to the higher levels of problem solving and game playing and finally to the level of natural language and understanding of the world. At the higher levels you need more than a simple pattern recognition approach, you really need some more complex computer architecture that includes methods for structuring thoughts.

Fourth, there is the matter of what constitutes AI. One of the early major goals for the subject was to produce systems that were as capable of dealing with the real world as human beings. The early researchers knew this would take time and that they had to start with much smaller systems and then scale them up as time went on. None of the systems have shown any promise of scaling up. The smaller systems can be highly useful, but for the most part they do not display all the characteristics of intelligence that people would demand of an intelligent system, although they typically show some characteristics of intelligence. Because of this, to a great extent AI has become “advanced algorithms” or “advanced programming techniques,” covering topics that deal with specialized areas of human achievement like game playing or natural language processing and where there may be little or no evidence of intelligent behavior on the part of the program. Unfortunately, in

many cases, in order to do a good job within these specialized areas, you need a system that has all the flexibility and knowledge of a whole human being. Therefore, to a certain extent I have chosen topics that consider what methods are necessary in a system that has all the flexibility of a human being. This also means that certain other traditional AI methods that would normally come up in an AI course have been ignored.

The main target groups for this book are juniors, seniors, and first year graduate students in computer science or psychology. Other people without computer science and psychology backgrounds should still be able to learn a lot about AI from the book. In the way of mathematics, a very small amount of linear algebra and calculus is called for. In linear algebra a knowledge of vectors, matrices, and matrix multiplication is sufficient. Calculus is only used for the SAINT integration program and for a derivation of the back-propagation algorithm in the appendix and even this derivation can be skipped if necessary. In the way of programming abilities, the ACM's CS-2 course that includes linked lists, trees, and recursion is adequate.

Some Prolog is introduced as a useful notation for describing the symbol processing algorithms. If you want to do more Prolog than this small amount, I have what used to be an additional chapter on Prolog available free on the Internet. If you prefer to have your students use Lisp instead, there is a Lisp chapter there as well. I do believe doing some symbolic programming is important to give students the flavor of the symbol processing approach. If you want them to do a lot of such programming, then instead of starting with Chapter 1, start with Section 4.2 and then go on to either the Internet Prolog or Lisp chapter and then start at Chapter 1.

The book is organized in a manner in which I think the reader will slowly get an intuitive feeling for the principles of AI. (I like to think that the book reads like a novel. A few reviewers have agreed with this assertion and one has disagreed.) Throughout the book, the applications of the basic principles are demonstrated by examining some classic AI programs in detail. To a large extent, people learn by seeing specific examples and then generalizing from them (a case-based approach).

I recommend that you read all the exercises even if you do not want to do them in a formal way. I feel this is worthwhile since many exercises raise issues that are not discussed in the text. Some of them are worth debating in class. Some exercises can be done without any programming, but many of them are programs. Some of these can reasonably be done in general purpose languages like Pascal, C, or Fortran. Instructors, take note that you cannot simply go to a chapter and assign exercises 1–5 one day, 6–10 the next, and so on. That would be altogether too much work for a student. Instead, you need to look at each exercise and choose only those that are appropriate for your type of student. Most exercises are as easy or hard as they look, except some long or hard ones that may seem short or easy have warnings attached to them.

More material is available online including neural networking software for Unix and DOS/MS-Windows and an outline and commentary on the topics. As time goes on there will be a list of frequently asked questions. If there is more material that deserves attention and fits in with the perspective of this book, please let me know and I will consider adding it online. For online information, see my Pattern Recognition Basis of AI page at:

<http://www.mcs.com/~drt/basisofai.html>

or send me an email note at:

[drt@mcs.com](mailto:drt@mcs.com).