Mathematical Methods in Artificial Intelligence

Edward A. Bender

Contents

Preface

If AI is ever to become a respectably hard science, then a firm, formal basis is needed.

—Derek Partridge (1991)

Philosophy

Teaching an AI course presents a problem. The field is so broad that an attempt to cover most of it is bound to result in a fairly shallow survey course. Nevertheless, it is important to discuss some important tools of AI in some depth. The tools can be roughly divided into three types: tools for implementing a plan (e.g., Lisp, microprocessors), tools for designing a plan (e.g., algorithms), and tools for designing tools.

A hands-on approach based on implementing plans is often pursued in computer science. Unfortunately, toy AI problems are of limited pedagogical use while real AI problems are often on such a scale that programming only one of them is a major project. Moreover, a hands-on approach often gives students the ability to implement some plans without giving them the ability to understand or develop the tools on which such plans are based.

As a result, I believe it's critical to focus on tools for designing AI tools. Since AI is a young science, we must to some extent anticipate what these tools will be. Mathematics has been the major tool designing tool in the sciences; therefore, I am persuaded that AI will not be an exception.

Possible Courses

It's popular to say that a book does not require much formal mathematics, but does require some mathematical maturity. That's true here. Much of the material can, in theory, be read and understood with no more background than high school algebra. In practice, however, students need more than this or they will be overwhelmed by the need to think mathematically. Furthermore, after the first ten chapters, some background in calculus is needed.

This text can be used for an introductory course in AI for upper-division or graduate students who have had a standard lower-division calculus course. Many courses are possible, depending on the time available, the capabilities of the students, and the interests of the instructor. All courses should include at least Chapter 1 and most of Chapters 2, 3, 5, and 10. Possible supplements are (a) logic from Chapters 4 and/or 6; (b) neural nets from Chapter 11; (c) probability and its uses from Chapters 7, 8, and 9; and (d) material from Chapter 14. The more mathematical exercises and proofs can be emphasized or deemphasized as circumstances dictate.

This text can also be used for a second course in AI for students interested in AI research. Many monographs and research papers are inaccessible to such students because they assume a mathematical background not provided by standard AI courses. The mathematics in this text helps bridge that gap.

Instructors may obtain a TEX diskette containing solutions to many of the exercises from Computer Society Press.

The following diagram illustrates some dependencies among all chapters but the last, ranging from weak (dotted lines) to nearly complete dependency (solid lines). Dashed lines indicate that only some sections are essential. More details on dependencies are found at the end of each chapter introduction.

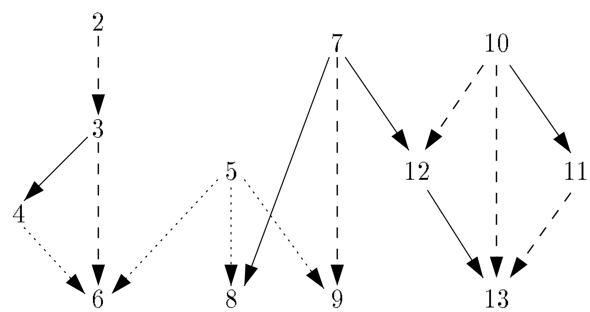

Acknowledgments

Various people have helped with this text. I'd particularly like to thank my students whose confusions and misunderstandings uncovered poorly written passages, my colleagues Frederic Bien, Te C. Hu, Fred Kochman, Alfred Manaster, Jeff Remmel, and Malcolm Williamson for their suggestions, my copy editor Emily Thompson for her ample, apposite use of red ink, and Computer Society Press for their helpful editorial assistance.

Dear Student

History teaches that new technology will require new mathematics. ... The question is: Which mathematics to use?

—Monique Pavel (1989)

Many introductory AI texts give the impression that AI is a collection of heuristic ideas and data structures implemented in Lisp and Prolog. The prognosis for such a discipline would be grim. Fortunately, AI researchers use mathematics and are developing new tools. Unfortunately, most of what you need is found in monographs and research articles—inappropriate material for a beginning course. This text is my attempt to fill the gap.

Since some of the mathematics used in AI is not part of a standard undergraduate curriculum, you'll be learning mathematics and seeing how it's used in AI at the same time. As with most mathematically oriented texts, this one isn't easy. I've written the next couple of pages to help you through it. Please read them.

Goals

In this text I hope to introduce you to some mathematical tools that have been important in AI and to some of their applications to the design of algorithms for AI. Since expert systems (broadly interpreted) comprise a large part of AI and have been the main focus of mathematically based tools, I have centered the book around the expert system idea.

As a result of studying this text, you should be in a much better position to read the technical literature in AI and should be able to easily fill in gaps in coverage by reading one of the more broadly based survey texts.

Reading Mathematics

Many people learn mathematics the way I learned history in high school. The exams contained two columns and the goal was to match each date in column A with one of the persons, places, and events in column B. Being lazy, I learned the “whens” of history but never the “whys.” I missed a whole world of ideas.

When mathematics is taught and learned by rote, students miss a world of ideas. Mathematics should be learned as an aid to thinking, not as a replacement for it. Learning mathematics is a skill that's seldom taught. If, like many students, you haven't mastered it, the following comments should be helpful.

The key is to work on understanding—not on memorization. How can you do this?

Let's begin with definitions. Whenever you meet a new concept, develop an understanding of it by relating it to ideas you already know and by looking at what it means in specific cases. For instance, when learning what a polynomial is, look at specific polynomials; when learning what continuity is, see what it means for a specific function like x^2. *The importance of understanding the general through the specific cannot be overemphasized*—even by using italics. The discussions and examples that immediately precede and follow definitions are often designed to foster understanding. If a definition refers to an earlier, unclear concept, stop! If you proceed, you may end up wandering aimlessly in a foggy landscape filled with shadowy concepts and mirages. Go back and improve your understanding of the earlier concepts so that they're practically solid objects that you can touch and manipulate. Finally, ask yourself why a definition has been introduced: What is the important or useful concept behind it? You may not be able to answer that question until you've read further in the text, but you can prepare your mind to recognize the answer when you see it.

What about theorems? The comments for definitions apply here, too: Look at specific examples, try to relate the theorem to other things you know, ask why it's important. Be sure you're clear on what the theorem *claims* and on what its words *mean*. In addition, attempt to see why the result seems reasonable before you read the proof. Reading and understanding the proof is the last step. If the proof is long, it may be helpful to make an outline of it. But don't mistake the ability to reproduce a proof for understanding. That's like expecting a photograph to understand a scene. There are better tests of understanding: Do you see where all of the assumptions are used? Can you think of a stronger conclusion than that in the theorem? If so, can you see why the stronger conclusion is not true, or at least why the proof is insufficient to establish the stronger conclusion?

Examples play a key role in mathematics. In practically every mathematics text, they fall into three categories.

- The type that aren't in the text: They're the ones you create by following the preceding advice.
- The obvious ones that are labeled "example" in the text. They're usually illustrations of definitions, algorithms, or theorems. Sometimes they develop related ideas.
- The type that comes from homework problems: These examples are the *solutions* to the problems you do yourself, not the problems themselves.

If you neglect any of these three types of examples, your mathematical text will be most useful to you as a doorstop.

Navigation Aids

Here's some information to help you navigate this text.

- A chapter introduction usually tells what's in the chapter, why it's there, and how the chapter is laid out. The overview it provides will help you organize the chapter in your mind.
- Numerous quotations highlight ideas and controversies, offer insight, provoke thought, and perhaps provide comic relief.
- Starred material either is more difficult than the text in which it is embedded or is peripheral.
- A remark that's somewhat off the track may appear as an Aside set in smaller type. Asides can be skipped without losing the thread of the discussion.
- There are four types of exercises. Here they are in order of difficulty.
  - Some exercises are lettered and some numbered; for example, 2.4.A versus 2.4.1. The purpose of lettered exercises is to make sure you absorbed the basic ideas. Their solutions can be found by rereading the preceding material. *You should do all lettered exercises.* It's often necessary to know the answers to these exercises before reading further.
  - A few numbered exercises are there to be sure you've picked up basic ideas that are needed soon. The answer to such an exercise is given immediately after the exercise section. *You should do all these exercises, then read the solutions.* If you've made an error, study the

section further or ask for help. It's important to understand how to do these exercises before reading further.

- The solutions to most exercises are neither very simple nor very difficult. Many you'll be able to do. Ask for help on those that baffle you.
- Starred exercises are ones that I consider difficult or that refer to starred material.

A few exercises don't have just one right answer. They may ask for your opinion or they may ask for you to construct an example of something. If an exercise asks for a proof, use full sentences. Read your proof aloud—it'll help you catch mistakes and incoherent thinking.

Enjoy your exploration of AI and its mathematical foundations.

Sincerely,
Ed Bender