Contents

Preface

This is a book about the nature of mathematical modeling, and about the kinds of techniques that are useful for modeling systems (both natural and otherwise). It is oriented towards simple efficient implementations on computers. Digital systems are routinely used for modeling purposes ranging from characterizing and transmitting realities to experimenting with possibilities to realizing fantasies; whether the goal is to reproduce our world or to create new worlds, there is a recurring need for compact, rich descriptions of how systems behave.

This text, like its companion The Physics of Information Technology [Gershenfeld, 1999a], grew out of many questions from students that indicated that the pressure for specialization within disciplines often leads to a lack of awareness of alternative approaches to a problem. Just as a typical physicist learns about variational principles, but not about how they can be used to derive finite element algorithms, an engineer who learns about finite elements might not know about techniques for exact or approximate analytical solutions of variational problems. Many people learn about how continuum equations arise from microscopic dynamics, and about how finite difference approximations of them can be used for numerical solutions, but not about how it is possible to stop before the continuum description by using lattice gases. And few outside of specialized research communities encounter emerging ideas such as state-space reconstruction for nonlinear systems, or methods for managing complexity in machine learning.

Each of these topics can (and should) be the subject of a semester-long course, and so to cover this range I necessarily have had to sacrifice some depth. Although this brisk pace risks missing important points, the alternative for many people may be no exposure at all to this material. I've tried to strike a balance by introducing the basic ideas, giving simple but complete and useful algorithms, and then providing pointers into the literature. Other than assuming some calculus, linear algebra, and programming skills, the book is selfcontained. My hope is that the inexperienced reader will find this collection of topics to be a useful introduction to what is known, readers with more experience will be able to find usable answers to specific questions, and that advanced readers will find this to be a helpful platform from which to see further [Merton, 1993].

Included here are a number of "old-fashioned" subjects such as ordinary differential equations, and "new-fashioned" ones such as neural networks. One of the major goals is to make clear how the latter relate to the former. The modern profusion of fashionable techniques has led to equally exaggerated claims of success and failure. There really are some important underlying new insights that were not part of traditional practice, but these are best viewed in that richer context. The study of neural networks, for example, has helped lead to more generally applicable lessons about functional approximation and search in high-dimensional spaces. I hope that by demystifying the hype surrounding some of these ideas I can help bring out their real magic.

The study of modeling is inseparable from the practice of modeling. I've tried to keep this text concise by focusing on the important underlying ideas, and left the refinements and applications to the problem sets as much as possible. I've included the problem solutions because I use problem sets not as a test to grade correct answers, but rather as a vehicle to teach problem solving skills and supporting material. My students are cautioned that although the solutions are freely available, their grade is based on how they approach the problems in class rather than on what answers they get, and so reciting my solutions back is both obvious and a waste of everyone's time. This organization helps them, and me, focus attention where it should be: on developing solutions instead of looking them up.

In this text I usually start with the assumption that the governing equations or experimental observations to be modeled are already known, and leave the introduction of specific applications to the relevant literature (see, for example, [Gershenfeld, 1999a] for details, and [Gershenfeld, 1999b] for context). Therefore, an important component of the corresponding course at MIT is a semester modeling project that provides a broader chance for the students to develop practical modeling skills in particular domains. This is the best way I've found to communicate the matters of taste and judgement that are needed to turn the raw material in this text into useful results. These projects are always a pleasure to supervise, ranging from studying the spread of gossip to the spread of fire. Some of the most valuable lessons have come from students sharing their experience building models. To help extend that interchange to all of the readers of this text there is a Web page at http:/www.cup.cam.ac.uk/online/nmm that provides pointers to explore, comment on, and contribute modeling projects.

I am a great admirer of real mathematicians, who I define to be people who always are aware of the difference between what they know and what they think they know; I certainly am not one. I hope that my presumption in reducing whole disciplines to ten or so pages apiece of essential ideas is exceeded by the value of such a compact presentation. For my inevitable sins of omission, comission, and everything in between, I welcome your feedback at nmm@media.mit.edu.

It has been a great pleasure to watch my sketchy lecture notes grow up into this text under the guidance of the many students and colleagues who have helped shape it by their thoughtful, challenging, exasperating, inspiring questions and comments (with a particular thanks to F. Joseph Pompei). And I am grateful to the Media Lab and its community of sponsors and collaborators who have created an environment that lets dreams be chased wherever they lead.

Cambridge, MA Neil Gershenfeld

April, 1998