

1 Introduction

How would you:

- How would you synthesize the sound of a violin?
- Analyze the sound of a violin?
- Model the traffic on a highway?
- Modify the traffic on a highway?
- Predict the bending of a beam?
- Design the bending of a beam?
- Simulate the weather?
- Forecast the weather?
- Recognize an image?
- Render an image?
- Animate a fish?
- Optimize a fish?

These questions do not have simple answers: all are active research areas. There cannot be a single recipe that covers this whole menu. There are many possible levels of description; choosing among them depends on your goals and on the available tools. This text is a tour through those spaces. For example, if you seek to make a mathematical model of a violin, you could use a numerical model based on a first-principles description. This lets you match your model parameters to measurements on a real instrument, and change parameters between a Stradivarius and a Guarneri. However, running it in real time will require a supercomputer, and the effort to find good parameters for the model is almost as much work as building a real violin. Alternatively, you could try to use an analytical (pencil-and-paper) solution to the governing equations; in return for some large approximations you may be able to find a useful explicit solution, but it might not sound very good. Finally, you could forget about the underlying governing equations entirely and experimentally try to find an effective description of how the player's actions are related to the sound made by the instrument (which is a reasonable thing to do because dissipation and symmetries in a system reduce the effective number of degrees of freedom [Temam, 1988]). These three approaches (analytical, numerical, and observational) comprise the three parts of this book.

To build a model there are many decisions that must be made, either explicitly or more often implicitly. Some of these are shown in Figure 1.1. Each of these is a continuum rather than a discrete choice. This list is not exhaustive, but it's important to keep

returning to it: many efforts fail because of an unintentional attempt to describe either too much or too little.



Figure 1.1. Some levels of description for mathematical model building.

These are *meta-modeling* questions. There are no rigorous ways to make these choices, but once they've been decided there are rigorous ways to use them. There's no single definition of a "best" model, although quasi-religious wars are fought over the question. One good attempt is the *Minimum Description Length* principle [Rissanen, 1986], essentially Occam's Razor: the best model is the one that is the smallest (including the information to specify both the form of the model and the values of the parameters). Unfortunately, this has two serious problems: finding the minimum description length for a given problem is an uncomputable task, and it says nothing about the error metric that will be used to judge the model. A stock trader, civil engineer, cardiologist, and video game designer have very different standards for success. They differ in the prior information they have about their problem, and the posterior criteria that they will use to evaluate and update their model. Ultimately, the strongest useful statement is that the best model is the one that works best for you.

Surprisingly little ambition is needed to exceed the performance of almost any available computer, and conversely computer hardware speeds have historically raced ahead of the development of software tools to use them effectively. Where computational speed is most important, the examples in this book will use efficient portable low-level tools. On the other hand, where algorithm clarity is most important, high-level environments will be used. The appendices provide brief introductions to these.

No single reference text covers the range of subjects in this book. To help access the literature, each chapter ends with a list of relevant general sources, and then cites the more specialized literature as needed throughout. Where important ideas are introduced without any references they are either so well known that they need no further citation, or are my own results that I have not published elsewhere (the context should make this distinction clear).

1.1 SELECTED REFERENCES

- [Press *et al.*, 2007] Press, William H., Teukolsky, Saul A., Vetterling, William T., & Flannery, Brian P. (2007). *Numerical Recipes in C: The Art of Scientific Computing*. 3rd edn. Cambridge: Cambridge University Press.

This is warmly recommended for almost any numerical problem. The numerical analysis literature is full of rigorous results that have little bearing on solving practical problems; *Numerical Recipes* gracefully merges theoretical insights with practical tricks for most useful algorithms. It's one of those rare books that's immediately useful by a beginner but that continues to hold new insights for an expert.

- [Pearson, 1990] Pearson, Carl E. (1990). *Handbook of Applied Mathematics: Selected Results and Methods*. 2nd edn. New York: Van Nostrand Reinhold.

This is a good example of one of a number of such large reference volumes that survey applied mathematics.