Preface

This book considers supervised learning in a class of artificial neural networks called multilayer perceptrons (MLP). This covers just a small part of the field of neural networks, but it is a significant part worth considering in detail. Interested readers are, of course, encouraged to consult other sources for information on the broader field.

The book is oriented to the practical reader whose goal is to train neural networks for a specific task. This may include students and practitioners in many fields. The typical reader may already have some familiarity with neural networks. This is a multidisciplinary field though and readers are likely to have a variety of backgrounds so we start with basic properties of single-layer networks and build from there. A mathematical background including college calculus and basic statistics is assumed.

The book surveys MLP training algorithms and includes practical hints for obtaining networks that train in reasonable amounts of time and generalize well. It is not a high-level debate about the fundamental capabilities of neural networks or their possible role as models of human intelligence. The goal is to describe selected techniques in enough detail to allow implementation by the motivated reader. Where possible, we attempt to explain how and why methods work (or don't) and the conditions that affect their success.

Part of our intent is to suggest ideas and give pointers for further study. We attempt to summarize theory where it is available and point out the major implications without going into rigorous derivations. In most cases the reader is referred to other sources for more detailed development. A warning: Some of the ideas are rather speculative and have been criticized as ad hoc. Exploration often precedes theoretical explanation, but ideas must be tested. Where possible, we try to provide references to empirical evaluations.