Preface

Regression and classification methods based on similarity of the input to stored examples have been part of the arsenal in statistics and computer science for decades. Despite consistently good performance in a number of domains, these methods have not been employed as widely in applications where very large sets of high-dimensional data are involved. Two of the main reasons for this are the computational complexity of similarity search in high-dimensional spaces, often seen as prohibitive, and the sensitivity of the exemplar-based methods to the choice of distance measure. The main focus of this book is on advances in computational geometry and machine learning that may alleviate these problems, and on emerging applications in the field of computer vision in which the benefit of these advances is often dramatic.

The book contains contributions by participants in the workshop on nearest-neighbor methods in learning and vision, held in Whistler, British Columbia, as part of the annual conference on Neural Information Processing Systems (NIPS) in December 2003. The workshop brought together researchers from theory of computation, machine learning, and computer vision. Its goal was to bridge the ostensible gaps between these disciplines and explore the state of the art in nearest-neighbor search methods on the one hand, and the emerging applications of these methods in learning and vision on the other hand. The chapters, organized into three corresponding parts, are representative of the ideas presented and discussed at the workshop.

We hope that this book will be of interest to the students, researchers, and practitioners of machine learning and computer vision, to whom it may provide inspiration and ideas as well as useful solutions to specific problems. In addition, we expect that the book will be of interest to researchers in computational geometry and algorithms, for whom it presents intersting application domains in need of new efficient algorithms.

We would like to aknowledge Paul Viola, who co-organized the NIPS workshop with us. We would also like to thank all those who attended the workshop and participated in the discussions, and the authors of the chapters for their excellent contributions to the workshop and now to this book. Finally, we are grateful to the workshop chairs of NIPS 2003, to the editors of the series, and to the staff of MIT Press, in particular Bob Prior

and Katherine Almeida, without all of whom this project would not have succeeded.

Gregory Shakhnarovich
Brown University

Trevor Darrell, Piotr Indyk
Massachusetts Institute of Technology, Cambridge