
Preface

Machine learning develops intelligent computer systems that are able to generalize from previously seen examples. Traditionally, machine learning has mostly been involved with generalizing decisions from examples, as in the case of *classification*, or predicting a scalar number, as in the case of *regression*. Recent developments, however, have approached a wider domain where the prediction has to satisfy additional constraints, i.e., the output has *structure*.

Predicting structured data can be considered as one of the big challenges in machine learning, namely:

Learning functional dependencies between *arbitrary* input and output domains.

The current book aims at collecting and reviewing the state of the art in machine learning algorithms and machine learning theory which considers learning in this scenario. The material included in this collection covers applications as diverse as machine translation, document markup, computational biology, image restoration, and information extraction — to name just a few motivating examples.

The present book contains a number of papers based on talks presented at two Neural Information Processing Systems (NIPS) workshops on “Learning with Structured Outputs” and “Graphical Models and Kernels,” along with several invited articles describing recent progress made since the workshops have taken place. We believe that it provides a timely overview of this exciting field, covering a wide range of promising methods. Structured prediction being a rather novel field, this overview cannot be comprehensive nor anticipate all future trends, but we hope it will provide a good starting point for entering this exciting area.

We would like to thank everybody who contributed toward the success of this book, in particular Karin Bierig and Bob Prior for their continuing support, assistance, and patience.

Gökhan Bakır, Thomas Hofmann, Bernhard Schölkopf,
Alexander J. Smola, Ben Taskar, S.V.N. Vishwanathan

Tübingen, Zürich, Canberra, Berkeley; September 2006