

Preface

In the last few years machine learning has made its way into the areas of administration, commerce, and industry, in an impressive way. Data mining is perhaps the most widely known demonstration of this phenomenon, complemented by less publicized applications of machine learning, such as adaptive systems in various industrial settings, financial prediction, medical diagnosis, and the construction of user profiles for WWW-browsers. This transfer of machine learning from the research labs to the “real world” has caused increased interest in learning techniques, dictating further effort in informing people from other disciplines about the state of the art in machine learning and its uses.

The objective of this book is to provide the reader with sufficient information about the current capabilities of machine learning methods, as well as ideas about how one could make use of these methods to solve real-world problems. The book is based primarily on the material that was presented in the Advanced Course in Artificial Intelligence (ACAI '99), which took place in Chania, Greece and was attended by research students, professionals, and researchers. However, the book goes beyond the material covered in the course, in that it contains several position papers on open research issues of machine learning.

The book is structured in a way that reflects its objective of educating the reader on how machine learning works, what the open issues are, and how it can be used. It is divided into two parts: methods and applications.

The first part consists of 10 chapters covering to a large extent the field of machine learning, from symbolic concept learning and conceptual clustering to case-based reasoning, neural networks, and genetic algorithms. The research issues addressed include the relationship of machine learning to knowledge discovery in databases, the handling of noisy data, and the modification of the learning problem through function decomposition. This part of the book concludes with two chapters examining the basic principles of learning methods. The first of the two chapters examines the approaches to selecting the appropriate method for a particular problem or modifying the problem representation to suit a learning method. In contrast, the last chapter of the section reviews the approaches to integrating different machine learning methods, in order to handle difficult learning tasks.

The second part of the book exposes the reader to innovative applications of machine learning. This part consists of 11 chapters, presenting a range of applications, from data mining in finance, marketing, and economics to learning in human language technology and user modeling. Most of these chapters are based on the work that was presented during the workshops of ACAI '99. Each such chapter describes the types of problem that have been approached with the use of machine learning in a particular domain and gives an overview of the work in this area, as presented at the relevant workshop.

We hope that the combination of theoretical and empirical knowledge in this book will be of use to the reader who is interested in entering this exciting research field and using mature machine learning techniques to solve real-world problems. The editors of the book would like to thank the distinguished authors for their willingness and cooperation in making this special volume a reality.

May 2001

Georgios Paliouras
Vangelis Karkaletsis
Constantine D. Spyropoulos