Preface

This book is based on my experience teaching introductory courses on machine learning to business school students and executive groups. The purpose of the material is not to convert the reader into a data scientist. Instead, it is to give the reader an understanding of the tools used by data scientists and how they can further the objectives of an organization. The third edition improves the presentation of material and includes a number of new case studies and examples. There is an extra chapter extending the discussion of different ways neural networks are designed and used.

Most students recognize that they need some knowledge of machine learning to survive in a world where jobs will be increasingly impacted by it. Today, all executives need to know how to use computers. Tomorrow, all executives will need to be comfortable managing large data sets and working with data science professionals to improve their productivity.

I have used no matrix or vector algebra and no calculus in this book. Although these areas of study can help specialists, it has been my experience that most business school students and most executives are not comfortable with them.

The book explains the most popular algorithms used by data scientists. This will enable the reader to assess their strengths and weaknesses for a particular situation and work productively with data science professionals. The algorithms are illustrated with a number of different data sets, which can be downloaded from my website:

www-2.rotman.utoronto.ca/~hull

Both Excel worksheets and Python code accompany the data sets. Virtually all my students are comfortable with Excel before taking my courses. I insist that all become comfortable with Python as well. This is

not a hard sell. Students recognize that coding skills have become a necessary prerequisite for many jobs in business.

Several hundred PowerPoint slides can be downloaded from my website. Instructors who choose to adopt the book are welcome to adapt the slides to meet their own needs.

A number of people have helped me move this book to a third edition. I would particularly like to thank Emilio Barone, Raymond Kan, and Jun Yuan, who made many suggestions for improving the material. I am grateful to Jay Cao, Jacky Chen, Jeff Li, and Niti Mishra who worked on some of the Python code that accompanies the book. I would also like to thank Rotman's FinHub center, the TD bank, and the Global Risk Institute in Financial Services for providing funding for the development of research and teaching materials in machine learning and financial innovation at the Rotman school. Peter Christoffersen (prior to his untimely death in 2018) and Andreas Park have been great colleagues at FinHub and provided much of the inspiration for the book.

I welcome comments on the book from readers. My email address is hull@rotman.utoronto.ca.

John Hull

About the Author

John Hull is a University Professor at the Joseph L. Rotman School of Management, University of Toronto. Prior to writing this book, he wrote three best-selling books in the derivatives and risk management area. His books have an applied focus and he is proud that they sell equally well in the practitioner and college markets. He is academic director of FinHub, Rotman's Financial Innovation Lab, which carries out research and develops educational material in all aspects of financial innovation. He has consulted for many companies throughout the world and has won many teaching awards, including University of Toronto's prestigious Northrop Frye award.