

Preface

John Hull was the sole author of the first three editions of this book. He is fortunate to have had three colleagues (Jacky Chen, Zissis Poulos, and Jun Yuan) to work with him on this, the fourth edition. Similarly to John Hull, the three new co-authors have taught courses on machine learning to undergraduate and graduate students at the Joseph L. Rotman School of Management, University of Toronto. The world of machine learning is changing fast. We feel that our collaboration has helped to keep the material in this book (and in our courses) up to date.

The purpose of the book 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, how they can further the objectives of an organization, and potential pitfalls. Communicating effectively with data scientists is likely to become an increasingly important skill for all employees in the future.

The fourth edition improves the presentation of material and includes a number of new case studies and examples. There are new chapters discussing recent innovations in areas such as natural language processing and large language models.

The material in this book is not a “hard sell” to business school students. Most of these students recognize that they will 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.

We have used no matrix or vector algebra and no calculus in this book. Although these areas of study can help specialists, it has been our experience that many business school students and 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:

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

Both Excel worksheets and Python code accompany the data sets. Virtually all students are comfortable with Excel before taking our courses. We insist that they become comfortable with Python as well. This is also not a “hard sell.” Students recognize that rudimentary coding skills have become a necessary prerequisite for many jobs in business.

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

A number of people have contributed to the book at various times. We would particularly like to thank Emilio Barone, Raymond Kan, Jay Cao, Jeff Li, Niti Mishra, and Zeyu Wang. We would also like to thank Rotman’s FinHub center, the Global Association of Risk Professionals, 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.

We welcome comments on the book from readers. John Hull’s email address is hull@rotman.utoronto.ca.

About the Authors

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 was until 2024 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.

Dr. Jun Yuan is the Managing Director of Group Risk Management at the Royal Bank of Canada (RBC), an Adjunct Professor and Senior Fellow at the Rotman School of Management and Munk School of Global Affairs & Public Policy respectively at the University of Toronto. A successful financial and climate risk management expert, he leads a data analytics and machine learning team at RBC, focusing on optimizing supervised and unsupervised learning and natural language processing for risk management and financial innovation applications. Recognized as a thought leader in machine learning, Dr. Yuan partners with the Bank of Canada and the University of Toronto on machine learning research initiatives and has published insightful discoveries. He also advises leaders at financial regulatory bodies and institutions on policy reform of financial systems, with his recommendations adopted by Basel IV. Dr. Yuan is a sought-after speaker at international finance conferences.

Jacky Chen is the Managing Director of Completion Portfolio Strategies at OPTrust, one of Canada's largest defined benefit pension funds, and an Adjunct Professor at the University of Toronto's Rotman School of Management. A seasoned practitioner of machine learning and AI in the financial industry, he has extensive experience in developing quantitative investment strategies and leading corporate AI initiatives. At Rotman, Jacky is actively involved in both teaching and research. Affiliated with the Rotman Financial Innovation Lab, his research focuses on applying machine learning techniques to portfolio hedging, derivatives pricing, and risk management. His work has been published in leading academic journals.

Zissis Poulos is an Assistant Professor of Financial Technologies at the School of Information Technology, York University. He is also affiliated faculty at Rotman's Financial Innovation Lab. His research focuses on applications of machine learning in derivatives hedging, risk management, volatility modeling, and using natural language processing to analyze financial soft information. From 2017 to 2019, Zissis served as project coordinator for NSERC COHESA, a Canada-wide strategic network focused on promoting research at the intersection of hardware design and AI. He has been an active member of organizing committees for several IEEE conferences. Beyond his academic pursuits, Zissis serves as a technical advisor for AI startups in the United States and consults for asset management firms in Canada. In 2019, he co-founded a Toronto-based startup that designs compute engines for energy-efficient deep learning. The company successfully exited in 2021.

