Preface

This book distills the experiences and lessons learned from the international Trading Agent Competition (TAC) series. Motivated by TAC, a community of academic and industry researchers has been inventing and polishing techniques for autonomous bidding by software agents. We, the authors, have been both organizers of TAC and successful participants. As such, we have tackled the problems posed by TAC with our own independent efforts, and we have closely observed the evolution of approaches developed by the community as a whole.

TAC is a stylized setting exemplary of the rapidly advancing domain of electronic marketplaces. It is also a benchmark, motivating researchers to apply innovative approaches to a common task. A key feature of TAC is that it provides an academic forum for open comparison of agent bidding strategies in a complex scenario, as opposed to, for example, automated trading in real-world securities markets, in which practitioners are less inclined to share their technologies. As the product of sustained focus and cross-fertilization of ideas over time, TAC provides a unique case study of the current capabilities and limitations of autonomous bidding agents.

Throughout the text, we balance the contextual reporting of results from the specific TAC scenario with the desire to generalize to the broader problem of autonomous bidding. To ground the discussion, we include substantial data from controlled TAC experiments and TAC tournaments, methods employed by particular TAC agents, and anecdotes from TAC events. To generalize these lessons and techniques, we develop a generic trading agent architecture unifying the approaches observed, define abstract versions of trading agent subproblems, and highlight important properties of these problems and proposed solutions through theoretical and experimental analysis.

We consider this dual approach—intensive design focused on a concrete scenario, interleaved with abstraction and analysis aimed at drawing general lessons—essential for deriving principled trading agent designs. Real-world markets are too complex to rely solely on abstract modeling, and specific markets are too idiosyncratic to admit direct transfer of techniques. By testing general ideas in particular scenarios, we are forced to work through operational details that tend not to arise in more abstract models. Through careful evaluation of proposed designs, we can achieve some confidence in their viability, and gather evidence about their limitations. Lifting the methods back up to more generic market scenarios enables adaptation to similar environments, and even transfer across qualitatively different market domains.

The main contributions of this book are (i) the story of the development and evolution of the TAC research initiative, including anecdotal accounts of

TAC agent interactions over the years; (ii) detailed analyses of specific TAC agent designs and bidding techniques; and (iii) development of some general engineering foundations for trading agent design. Our intended audience includes individuals interested in developing TAC agents, but we expect most readers are primarily motivated by other trading domains. By inviting all to immerse yourselves in the TAC domain, we lead you on the very path we have taken in developing our understanding of the current best practices for designing autonomous bidding agents.

Acknowledgments

By its very nature, the Trading Agent Competition is a collective enterprise, and this book would not have been possible without the contributions of scores of developers who have implemented TAC agents over the years. We would especially like to thank the people central to the design and operation of the TAC Travel game servers: Peter Wurman and Kevin O'Malley (University of Michigan) and Sverker Janson, Joakim Eriksson, and Niclas Finne (Swedish Institute of Computer Science). Individuals serving as GameMaster over the years are listed in Appendix A. Other volunteers who have served important roles enabling the TAC Travel tournaments include Eric Aurell, Maria Fasli, Nick Jennings, David Parkes, Norman Sadeh, and Shihomi Wada.

Many of the results presented herein were produced collaboratively with our colleagues and students. In most cases, specific attribution is provided by citations throughout the book to articles previously published in various journals and conference proceedings. The contributions of these coauthors and others listed below are to a great extent inseparable from our own.

ATTac was developed by Peter Stone in collaboration with colleagues at AT&T Labs—Research: Michael Littman (ATTac-00 and ATTac-01); Michael Kearns and Satinder Singh (ATTac-00); and Jánosz Csirik, David McAllester, and Robert Schapire (ATTac-01). Many of the ideas described in connection with ATTac, particularly in Chapter 6 would not have come about without their cooperation. Stone's research related to this book was supported in part by the National Science Foundation under Grant No. IIS-0237699, and by an Alfred P. Sloan Foundation research fellowship.

RoxyBot was developed by Amy Greenwald, initially in collaboration with Justin Boyan (2000–2002), and later in collaboration with her students at Brown University, specifically Jesse Funaro (2003), Jonathan Bankard (2004),

Bryan Guillemette (2005), Seong Jae Lee (2006), and Victor Naroditskiy (2003–2006). Victor deserves our sincerest gratitude for his substantial contributions to the theoretical development in Chapter 3 (Theorems 3.2 and 3.4) and for creating the series of examples of the stochastic bidding problem presented in Chapter 5. Seong Jae also deserves special thanks for conducting the experiments reported in Chapter 5. Greenwald's research related to this book was supported in part by the National Science Foundation under Grant No. IIS-0133689, and by an Alfred P. Sloan Foundation research fellowship.

The Walverine team, supervised by Michael Wellman, has engaged many graduate and undergraduate students at the University of Michigan. Participants from 2002–present included Shih-Fen Cheng, Evan Leung, Kevin Lochner, Daniel Reeves, Julian Schvartzman, and Yevgeniy Vorobeychik. Kevin O'Malley, Christopher Kiekintveld, Daniel Reeves, and William Walsh were instrumental in the operation of the first two TAC events. Daniel Reeves, Kevin Lochner, and Rahul Suri provided indispensable testbed and analysis tools. Wellman's research related to this book was supported in part by the National Science Foundation under Grant Nos. IIS-9988715, IIS-0205435, and IIS-0414710.