Trading Agents

Synthesis Lectures on Artificial Intelligence and Machine Learning

Action Programming Languages
Michael Thielscher
2008

Representation Discovery using Harmonic Analysis
Sridhar Mahadevan
2008

Essentials of Game Theory: A Concise Multidisciplinary Introduction
Kevin Leyton-Brown and Yoav Shoham
2008

A Concise Introduction to Multiagent Systems and Distributed Artificial Intelligence
Nikos Vlassis
2007

Intelligent Autonomous Robotics: A Robot Soccer Case Study
Peter Stone
2007

Trading Agents

Michael P. Wellman

ISBN: 978-3-031-00426-1 paperback
ISBN: 978-3-031-01554-0 ebook

DOI 10.1007/978-3-031-01554-0

A Publication in the Springer series
SYNTHESIS LECTURES ON ARTIFICIAL INTELLIGENCE AND MACHINE LEARNING

Lecture #12
Series Editor: Ronald J. Brachman, *Yahoo Research*
William W. Cohen, *Carnegie Mellon University*
Thomas Dietterich, *Oregon State University*
Series ISSN
Synthesis Lectures on Artificial Intelligence and Machine Learning
Print 1939-4608 Electronic 1939-4616

Trading Agents

Michael P. Wellman
University of Michigan

SYNTHESIS LECTURES ON ARTIFICIAL INTELLIGENCE AND MACHINE LEARNING #12

ABSTRACT

Automated trading in electronic markets is one of the most common and consequential applications of autonomous software agents. Design of effective trading strategies requires thorough understanding of how market mechanisms operate, and appreciation of strategic issues that commonly manifest in trading scenarios. Drawing on research in auction theory and artificial intelligence, this book presents core principles of strategic reasoning that apply to market situations. The author illustrates trading strategy choices through examples of concrete market environments, such as eBay, as well as abstract market models defined by configurations of auctions and traders. Techniques for addressing these choices constitute essential building blocks for the design of trading strategies for rich market applications. The lecture assumes no prior background in game theory or auction theory, or artificial Intelligence.

KEYWORDS

trading agent, bidding agent, algorithmic trading, auction, computational game theory, eBay, multiagent systems

Contents

Preface

Delegating decisions to automated software agents is both compelling and fraught with risk. Computers can process massive volumes of information, respond in microseconds to events, and reliably follow instructions despite their tedium. But automation requires that we provide in advance the right instructions for the situations that will arise in the complex environments in which we set the agents loose. By including in the agent adaptive components that are shaped by the agent's own experience, we may improve its flexibility and potential performance, but at the same time increase the difficulty of predicting its behavior.

When the decisions are about trading in markets, the stakes are significant on both sides of the automation dilemma. The vast amount of data and complexity of calculations bearing on the trading decision—valuing goods and services, predicting future market prices—necessitate algorithmic assistance. Rapidity of response required by dynamics of the market may preclude direct human involvement. On the other hand, trading decisions inherently produce economic consequences. Delegating authority to trading agents means trusting to software the disposition of real resources—money, commitments to provide goods or services, or whatever is being exchanged in the market at hand.

This lecture is designed for those who would design trading agents. Whereas it cannot address the general problem of automating complex decisions, the lecture instead focuses on issues specific to market environments. By presenting key principles drawn from research in economics and artificial intelligence, it aims to equip the designer with some basic tools for strategic reasoning about market interactions. By "tools" I refer to patterns of thinking that support strategic analysis of well-specified market scenarios. We derive these patterns through study of canonical market forms, and identification of pivotal strategic questions that come up in market situations. Although the markets and situations we explore are inevitably idealized simplifications, the concepts and insights we take as lessons from these studies form the building blocks for more comprehensive agent design.

Much of our study of canonical market models lies within the realm of *auction theory*. There are many existing surveys and even textbooks covering this ground, and numerous references to same are provided within. The vantage taken in this lecture differs from standard treatments in two broad respects, which affect selection of topics as well as the technical level and manner in which they are treated.

1. *Trading agent focus*. Many discourses on auction theory emphasize the problem of the auction designer, highlighting such issues as relative revenue properties of alternative market mechanisms. The perspective here is exclusively that of the trading agent designer, whose problem is to participate in an existing market regardless of how or why it came to be. We engage in

comparative mechanism analysis to highlight implications of rule differences for trading agent strategy, but avoid distracting attention to questions about market design per se.

2. *Practical strategic reasoning*. Where other introductions to auction theory cater to those who aspire to extend the theory in new directions, the aim here is to generate lessons for constructing strategies in analogous situations. Accordingly, the exposition relies heavily on worked-out numeric examples, as opposed to formally stating and demonstrating abstract propositions. Our premise is that often the best way to understand the import of a strategic issue is to work it out in a concrete instance, notwithstanding the simplifications invoked. As will become apparent to the reader, even simplified auction scenarios can prove fairly complicated to work through in detail.

The lecture presupposes no prior background in game theory or auction theory, or artificial intelligence, although no doubt those with some prior exposure to these fields will find some of this material easier to get through. I make no attempt to systematically cover game theory, particularly as that topic is nicely treated by another *Synthesis Lecture* [Leyton-Brown and Shoham, 2008]. Game-theoretic concepts are rather introduced as needed, and applied going forward. The mathematics employed in this lecture lie mainly within algebra and probability, and a little bit of calculus. The reader with less mathematical inclination or energy can gloss over most derivations with measured loss in insight. However, I would urge those with fortitude to work through the example scenarios, as there is no substitute for actively solving a detailed problem for making a concept really sink in.

Except for an early chapter that motivates strategic issues through an extended case study of bidding on eBay, the lecture emphasizes generic market structures rather than specific markets or trading domains. In particular, the lecture contains no specific treatment of electronic trading in financial markets, despite the preponderance of trading activity devoted to financial securities. Similarly, those looking for dedicated case studies on electricity trading or Internet advertising or other vertical market will not find them here. Nevertheless, agent designers for any of these domains will find that many of the more abstract strategic lessons are directly applicable to their problems. For example, the double auction mechanism studied in Chapter 4 is an abstract version of the market institution employed almost universally by commodity and stock exchanges. Emphasis of issues such as bidding in simultaneous interdependent markets (Chapter 5) is fundamentally motivated by the relevance of these problems to real trading domains like electricity and keyword search advertising. No doubt specialized knowledge would inform trading agent design for particular applications, but limitations in space and time and expertise preclude detailed treatments here.

Writing this lecture was certainly a great learning experience for its author. The content draws significantly on collaborative research experiences with students and colleagues, especially Amy Greenwald, Chris Kiekintveld, Jeff MacKie-Mason, Anya Osepayshvili, Dan Reeves, and Julian Schvartzman. Along with other members of the Strategic Reasoning Group at the University of Michigan Artificial Intelligence Laboratory over the years, they deserve much credit for shaping the

perspectives and ideas I attempt to convey here. I am particularly grateful to individuals who provided constructive feedback on earlier drafts, namely Ben Cassell, Quang Duong, Amy Greenwald, Erika Homann, Kevin Leyton-Brown, Peter Stone, Prateek Tandon, and Bryce Wiedenbeck. You are fortunate to have been spared the errors and confusions they uncovered, if not those I kept hidden until after their reviews.

This book is dedicated to Julian and Clara, my favorite autonomous agents.

Michael P. Wellman
Ann Arbor, Michigan
July 2011