

negotiation

machine learning

specification

argumentation

verification

agent

social choice

decision making

programming

second edition

Multiagent Systems

edited by

Gerhard Weiss

coalition formation

auctions

autonomy

organization

joint planning

communication

intelligence


cooperation

competition

Multiagent Systems

second edition

edited by Gerhard Weiss



The MIT Press
Cambridge, Massachusetts
London, England

©2013 Massachusetts Institute of Technology

All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from the publisher.

MIT Press books may be purchased at special quantity discounts for business or sales promotional use. For information, please email special_sales@mitpress.mit.edu or write to Special Sales Department, The MIT Press, 55 Hayward Street, Cambridge, MA 02142.

This book was set in Times Roman by the editor. Printed and bound in the United States of America.

Library of Congress Cataloging-in-Publication Data

Multiagent systems / edited by Gerhard Weiss. — Second edition.

p. cm.—(Intelligent robotics and autonomous agents series)

Includes bibliographical references and index.

ISBN 978-0-262-01889-0 (hardcover : alk. paper)

1. Intelligent agents (Computer software) 2. Distributed artificial intelligence.

I. Weiss, Gerhard, 1962– editor of compilation.

QA76.76.I58M85 2013

006.3—dc23

2012024600

10 9 8 7 6 5 4 3 2 1

Multiagent Systems

Intelligent Robotics and Autonomous Agents

Edited by Ronald C. Arkin

Dorigo, Marco, and Marco Colombetti, *Robot Shaping: An Experiment in Behavior Engineering*

Arkin, Ronald C., *Behavior-Based Robotics*

Stone, Peter, *Layered Learning in Multiagent Systems: A Winning Approach to Robotic Soccer*

Wooldridge, Michael, *Reasoning About Rational Agents*

Murphy, Robin R., *An Introduction to AI Robotics*

Mason, Matthew T., *Mechanics of Robotic Manipulation*

Kraus, Sarit, *Strategic Negotiation in Multiagent Environments*

Nolfi, Stefano, and Dario Floreano, *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*

Siegwart, Roland, and Illah R. Nourbakhsh, *Introduction to Autonomous Mobile Robots*

Breazeal, Cynthia L., *Designing Sociable Robots*

Bekey, George A., *Autonomous Robots: From Biological Inspiration to Implementation and Control*

Choset, Howie, Kevin M. Lynch, Seth Hutchinson, George Kantor, Wolfram Burgard, Lydia E. Kavraki, and Sebastian Thrun, *Principles of Robot Motion: Theory, Algorithms, and Implementations*

Thrun, Sebastian, Wolfram Burgard, and Dieter Fox, *Probabilistic Robotics*

Mataric, Maja J., *The Robotics Primer*

Wellman, Michael P., Amy Greenwald, and Peter Stone, *Autonomous Bidding Agents: Strategies and Lessons from the Trading Agent Competition*

Floreano, Dario and Claudio Mattiussi, *Bio-Inspired Artificial Intelligence: Theories, Methods, and Technologies*

Sterling, Leon S. and Kuldar Taveter, *The Art of Agent-Oriented Modeling*

Stoy, Kasper, David Brandt, and David J. Christensen, *An Introduction to Self-Reconfigurable Robots*

Lin, Patrick, Keith Abney, and George A. Bekey, editors, *Robot Ethics: The Ethical and Social Implications of Robotics*

Weiss, Gerhard, editor, *Multiagent Systems*, second edition

For Sofie, Alina, and Tina
—G. W.

Preface

Gerhard Weiss

The Subject of This Book

Multiagent systems are systems composed of multiple interacting intelligent agents. An agent is a computational entity such as a software program or a robot that is situated in some environment and that to some extent is able to act autonomously in order to achieve its design objectives. As interacting entities, agents do not simply exchange data but are actively engaged in cooperative and competitive scenarios; they may communicate on the basis of semantically rich languages, and they achieve agreements and make decisions on the basis of processes such as negotiation, argumentation, voting, auctioning, and coalition formation. As intelligent entities, agents act flexibly, that is, both reactively and deliberatively, in a variety of environmental circumstances on the basis of processes such as planning, learning, and constraint satisfaction. As autonomous entities, agents have far-reaching control over their behavior within the frame of their objectives, possess decision authority in a wide variety of circumstances, and are able to handle complex and unforeseen situations on their own and without the intervention of humans or other systems. And as entities situated in some environment, agents perceive their environment at least partially and act upon their environment without being in full control of it. Concrete multiagent systems and their environments have several relevant attributes in which they can differ; the table below (on page xxxvi) indicates the wide range of possible instantiations of these attributes.

Since its inception in the late 1970s, the field of multiagent systems has evolved impressively and today it is an established and vibrant field in computer science. The field has a profound and broad conceptual and theoretical foundation, drawing on and bringing together results, techniques, and tools not only from computer science and artificial intelligence (which traditionally has dealt

	attribute	range
agents	number	from two upward
	uniformity	homogeneous ... heterogeneous
	goals	contradicting ... complementary
	flexibility	purely reactive ... purely deliberative
	abilities (sensors, effectors, cognition)	simple ... advanced
	autonomy	low ... high
interaction	frequency	low ... high
	persistence	short-term ... long-term
	level	signal-passing ... knowledge-intensive
	language	elementary ... semantically rich
	pattern (flow of data and control)	decentralized ... hierarchical
	variability	fixed ... changeable
	purpose	competitive ... cooperative
environment	predictability	forseeable ... unforeseeable
	accessibility and knowability	unlimited ... limited
	dynamics	fixed ... variable
	diversity	poor ... rich
	availability of resources	restricted ... ample

Variety of multiagent systems and their environments.

with single-agent systems) but also from mathematics, logics, game theory, and other areas. The multiagent systems field is multidisciplinary in nature. Examples of disciplines to which the field is related are cognitive psychology, sociology, organization science, economics, and philosophy.

A main reason for the vast interest and attention multiagent systems are receiving is that they are seen as an enabling technology for applications that rely on distributed and parallel processing of data, information, and knowledge in complex – networked, open, and large-scale – computing environments. With advancing technological progress in interconnectivity and interoperability of computers and software, such applications are becoming standard in a variety of domains such as e-commerce, logistics, supply chain management, telecommunication, health care, and manufacturing. More generally, such applications are characteristic of several widely recognized computing paradigms known as grid computing, peer-to-peer computing, pervasive computing, ubiquitous computing, autonomic com-

puting, service-oriented computing, and cloud computing. In a sense multiagent technology is complementary to and cuts across these paradigms. Another reason for the broad interest in multiagent systems is that these systems are seen as a technology and tool that helps to analyze and develop models and theories of interactivity in large-scale human-centered systems. Research motivated by this interest is mainly based on computer-based experimental analysis and simulation rather than mathematical studies. The general goal of this line of research is to gain a deeper understanding of complex organizational, societal, economic, and political phenomena (including their underlying mechanisms and observable dynamics) for which the individual abilities and traits of the involved humans are crucial and thus cannot be neglected.

Main Features of This Book

The book offers several features that support its use as a textbook and reference volume:

- *Scope* – It captures the state of the art in the field in breadth and depth.
- *Theory* – It conveys the theoretical foundations and underpinning of multi-agent systems.
- *Clarity* – It provides many illustrations and examples.
- *Practice* – It includes many exercises of varying degrees of difficulty.
- *Expertise* – Its chapters are written by leading experts and authorities in the field.

It is worth saying a few words about the last-mentioned feature. This feature ensures that the book is built on an outstanding, broad, and profound basis of knowledge and experience. As the readers will also see, tremendous effort has been invested in carefully coordinating the individual chapters and in ensuring overall coherence of the book. In a way, there is no approach to writing a book on *multiagent* systems that is more natural and obvious than the *multi-author* approach taken for this book. A list of the thirty-one contributing authors is provided on pages xliii–xlvi.

Readership and Prerequisites

The book is primarily intended for use in undergraduate, graduate, and postgraduate courses and is also suited for self-study. The main academic audiences are *stu-*

dents and teachers of artificial intelligence, computer science, information technology, and related fields. Because of the multidisciplinary nature of multiagent systems (both as a technology and a field), it can also serve as a course text for students and teachers from disciplines such as psychology, economics, sociology, and philosophy. Moreover, because of its breadth and depth, the book can serve as a basic reference volume for both *researchers* who want to branch out beyond their own subfields and *professionals from industry* who want to explore potential usages of multiagent technology in their application areas.

As far as possible, the chapters are written so they can be understood without advanced prior knowledge. The main prerequisite for making the most of the book and for understanding its contents in detail is familiarity with basic concepts of computer science (especially algorithms and programming) and mathematics (especially logics and game theory) at the freshman level. Some useful background in logics and game theory is supplied in Part VI of this book.

Changes from the First Edition

The first edition appeared in 1999, and since then the field of multiagent systems has developed considerably. Some topics and themes that were characteristic of the field some twelve years ago play only a minor role today, and some of today's core topics played no or only a marginal role at that time. This second edition captures all these changes and shows the current state of the art in the field. Much more work was necessary in creating the new edition other than just rewriting or restructuring some of the first-edition chapters. Only one of the seventeen chapters was already included in the first edition (some passages of the chapter have been rewritten and some material has been added); the other sixteen chapters are entirely new.

What remained *unchanged* is the unique conception and vision behind the book: to have a high-quality course book and reference volume on multiagent systems whose parts are all written by acknowledged authorities in the field.

Structure and Chapters

This book is divided into six parts:

- Part I introduces basic concepts and principles of computational agency and covers key issues of both individual agents (Chapter 1) and agent organizations (Chapter 2).

- Part II focuses on communication among agents and discusses agent communication languages (Chapter 3) as well as two forms of agent-agent interaction – negotiation and bargaining (Chapter 4) and argumentation (Chapter 5) – for which communication is particularly crucial.
- Part III focuses on coordination among agents from different perspectives, including social choice (Chapter 6), mechanism design and auctions (Chapter 7), coalition formation (Chapter 8), and trust and reputation (Chapter 9).
- Part IV focuses on distributed cognition in multiagent systems and deals with several basic cognitive abilities, namely, learning (Chapter 10), planning and decision making (Chapter 11), and constraint handling and optimization (Chapter 12).
- Part V focuses on the development and engineering of multiagent systems and deals with programming (Chapter 13), specification and verification (Chapter 14), and agent-oriented software engineering (Chapter 15).
- Part VI provides relevant and useful background knowledge in logics (Chapter 16) and game theory (Chapter 17).

Each chapter starts with a motivating introduction, then delves into its topic, and concludes with considerations on the current state of the art, open challenges, and latest developments. The chapters provide a number of pointers to relevant literature and put particular emphasis on providing examples and illustrations. Moreover, each chapter comes with various exercises.

The Exercises

At the end of each chapter, exercises of varying difficulty are provided, which concern relevant theoretical and practical aspects of multiagent systems. The following four levels of difficulty are distinguished to roughly indicate the amount of effort required for solving the exercises:

- **Level 1** Simple test of comprehension or slightly more subtle problem, solvable within a few hours or days. (Appropriate for Bachelor education)
- **Level 2** Much harder problem (e.g., requires writing a non-trivial program); solving it could take several days or weeks. (Bachelor)
- **Level 3** Even harder problem, typically related to a “hot” topic of current research; solving it could take weeks or months. (Bachelor/Master)

- **Level 4** An open research question; solution cannot be expected within a few months or could even be a topic of a PhD. (Master/PhD)

I recommend addressing as many Level-1 and Level-2 exercises as possible and dealing with at least a few of the Level-3 and Level-4 exercises. Most of the Level-1 and Level-2 exercises can be solved with the knowledge provided in this book, whereas Level-3 and Level-4 exercises typically require additional literature studies and extensive theoretical and/or experimental studies. Carefully working through Level-1 and Level-2 exercises will reward a reader with a real understanding of the material treated in the chapters, and solving Level-3 and Level-4 exercises will turn a reader into a real expert!

How to Use This Book

The book can be used for teaching as well as self-study. The chapters and thus the overall book are designed to be self-contained and understandable without additional material. Of course, there are many relationships between the chapters, but in principle they can be treated independently and read in any sequence. In general, I recommend starting off with Part I (i.e., Chapters 1 and 2) in order to set up a proper contextual understanding, especially if the reader is new to this field.

All chapters together can easily fill two one-semester courses. There are several ways to use the book. One possibility is to work linearly through the book from front to back, thereby covering each of the individual chapters fully or just partially. Another possibility, resulting from the self-containment of the chapters, is to tailor a course to specific needs and interests by using only some selected chapters in any preferable order while dropping the others. The book can also be employed as a *complementary text for courses on “classical” (single-agent) AI* in order to selectively cover elements of multi-agency. For instance, if such a course deals with machine learning, then an obvious option is to include also one or another multiagent learning algorithm. Last but not least, the chapters in this book can be used as *supportive material for specialized courses* on topics such as electronic auctions, smart devices, cooperative information systems, and autonomous systems engineering.

The book contains a number of exercises that allow readers to test and further deepen their knowledge. Course instructors may find them helpful and inspiring in view of formulating review questions, in-class exercises, homework problems, or course exams. Some exercises are fairly simple and are intended to make sure that basic material provided in the chapters is mastered. Others are more difficult and challenging and may serve as subjects of class discussion or advanced team work.

Throughout the book numerous references to relevant literature are provided. They enable interested students to further explore specific aspects, and they support teachers in choosing additional course material.

Slides and More – The Website of the Book

This book is accompanied by the website accessible via

- <http://mitpress.mit.edu/multiagentsystems>

The site is intended to provide useful teaching material for students and teachers. The website starts with lecture slides for the chapters (prepared by the respective chapter authors) and some other resources such as copies of the figures and a list of exercises in the book. I hope to extend the supplementary material once this book is in use.

Teachers, students, and industrial professionals are invited and encouraged to contribute additional resources based on the contents of the book. Examples of such additional resources are

- exercises, lab projects, and exam papers,
- alternative slides and lecture videos,
- descriptions/syllabi of courses employing this book, and
- errors slipped in the book (errata).

Teachers using the book are asked to notify me of their courses (with URLs). I will maintain an online list of these courses. The material I receive will be made available on the website so that all readers and the multiagent community as a whole can benefit from it. Additional resources and related questions can be mailed to gerhard.weiss@maastrichtuniversity.nl

Acknowledgments

This book has been developing for nearly two years (taking into account that the 17 chapters had been prepared in parallel by different authors, this amounts to a total development period of 34 years invested in this book!). During that time many people have contributed to the book project. I cannot thank all of them, and so below I mention those to whom I am particularly indebted. Please also see the acknowledgments at the end of the individual chapters.

I am most grateful to the contributing authors for their engagement and enthusiasm. They not only provided the chapters and chapter slides, but also gave many

useful comments and suggestions on how the overall quality of the book could be further improved. In particular, they all put tremendous effort in carefully coordinating the contents of their chapters. Developing a textbook like this, with 31 leading experts in the field being involved as the chapter authors, is an endeavor that is thrilling and appealing on the one hand but can easily fail for many reasons on the other. Obviously this book project did not fail, and in large part this is due to the professionalism and commitment of the involved authors. I enjoyed a lot working with them on the new edition. Here is my advice to all who think about editing a comparable multi-author textbook: do it, but only if you have such a strong author team behind you as I had.

My special thanks also goes to the authors of the first edition: their excellent chapters laid the foundation for this second edition.

I want to thank Tuomas Sandholm for valuable discussions in the early phase of this book project on possibilities for covering certain key topics in the best possible way.

At MIT Press, I am grateful to Marc Lowenthal, James DeWolf and Virginia Crossman for providing professional guidance, assistance, and support during this book project whenever necessary. Ada Brunstein provided great support in the start-up phase of the project. Thanks also to the anonymous reviewers provided by MIT Press for their useful remarks and comments.

I “misused” numerous family-time evenings and weekends to work on this book project – my warmest thanks go to Tina, Alina, and Sofie for their patience and understanding.

Contributing Authors

Rafael H. Bordini

Faculty of Informatics
PUCRS - Pontifical Catholic University of Rio Grande do Sul
Porto Alegre, RS, Brazil

⇒ *Chapter 13*

Felix Brandt

Institut für Informatik
Technische Universität München
Munich, Germany

⇒ *Chapter 6*

Amit K. Chopra

Department of Information Engineering and Computer Science
University of Trento
Trento, Italy

⇒ *Chapter 3*

Vincent Conitzer

Department of Computer Science
Duke University
Durham, NC, USA

⇒ *Chapter 6*

Virginia Dignum

Faculty of Technology, Policy and Management
Delft University of Technology
Delft, The Netherlands

⇒ *Chapter 2*

Jürgen Dix

Department of Computer Science
Clausthal University of Technology
Clausthal, Germany

⇒ *Chapters 13 and 14*

Ed Durfee

Computer Science and Engineering Division
University of Michigan
Ann Arbor, MI, USA

⇒ *Chapter 11*

Edith Elkind

Division of Mathematical Sciences
School of Physical and Mathematical Sciences
Nanyang Technological University
Singapore

⇒ *Chapters 8 and 17*

Ulle Endriss

Institute for Logic, Language and Computation
University of Amsterdam
Amsterdam, The Netherlands

⇒ *Chapter 6*

Alessandro Farinelli

Computer Science Department
University of Verona
Verona, Italy

⇒ *Chapter 12*