

LNAI 3898

Karl Tuyls  
Pieter Jan 't Hoen  
Katja Verbeeck  
Sandip Sen (Eds.)

# Learning and Adaption in Multi-Agent Systems

First International Workshop, LAMAS 2005  
Utrecht, The Netherlands, July 2005  
Revised Selected Papers



Springer

TP181-53

L217

2065

Karl Tuyls Pieter Jan 't Hoen  
Katja Verbeeck Sandip Sen (Eds.)

# Learning and Adaption in Multi-Agent Systems

First International Workshop, LAMAS 2005  
Utrecht, The Netherlands, July 25, 2005  
Revised Selected Papers



E200603509



Springer

## Series Editors

Jaime G. Carbonell, Carnegie Mellon University, Pittsburgh, PA, USA  
Jörg Siekmann, University of Saarland, Saarbrücken, Germany

## Volume Editors

Karl Tuyls  
Universiteit Maastricht  
Tongersestraat 6, Maastricht, The Netherlands  
E-mail: k.tuyls@cs.unimaas.nl

Pieter Jan 't Hoen  
Center for Mathematics and Computer Science (CWI)  
Kruislaan 413, P.O. Box 94079, 1090 GB Amsterdam, The Netherlands  
E-mail: hoen@cwi.nl

Katja Verbeeck  
Vrije Universiteit Brussel  
Faculty of Sciences (WE), Department of Computer Science  
Pleinlaan 2, 1050 Brussels, Belgium  
E-mail: kaverbee@vub.ac.be

Sandip Sen  
University of Tulsa  
Department of Mathematical and Computer Sciences  
600 S. College, Tulsa, OK 74104, USA  
E-mail: sandip-sen@utulsa.edu

Library of Congress Control Number: 2006923098

CR Subject Classification (1998): I.2.11, I.2, C.2.4

LNCS Sublibrary: SL 7 – Artificial Intelligence

ISSN 0302-9743  
ISBN-10 3-540-33053-4 Springer Berlin Heidelberg New York  
ISBN-13 978-3-540-33053-0 Springer Berlin Heidelberg New York

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, re-use of illustrations, recitation, broadcasting, reproduction on microfilms or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer. Violations are liable to prosecution under the German Copyright Law.

Springer is a part of Springer Science+Business Media  
springer.com

© Springer-Verlag Berlin Heidelberg 2006  
Printed in Germany

Typesetting: Camera-ready by author, data conversion by Scientific Publishing Services, Chennai, India  
Printed on acid-free paper SPIN: 11691839 06/3142 5 4 3 2 1 0

# Lecture Notes in Artificial Intelligence

3898

Edited by J. G. Carbonell and J. Siekmann

Subseries of Lecture Notes in Computer Science



## Preface

This book contains selected and revised papers of the International Workshop on Learning and Adaptation in Multi-Agent Systems (LAMAS 2005), held at the AAMAS 2005 Conference in Utrecht, The Netherlands, July 26.

An important aspect in multi-agent systems (MASs) is that the environment evolves over time, not only due to external environmental changes but also due to agent interactions. For this reason it is important that an agent can learn, based on experience, and adapt its knowledge to make rational decisions and act in this changing environment autonomously.

Machine learning techniques for single-agent frameworks are well established. Agents operate in uncertain environments and must be able to learn and act autonomously. This task is, however, more complex when the agent interacts with other agents that have potentially different capabilities and goals. The single-agent case is structurally different from the multi-agent case due to the added dimension of dynamic interactions between the adaptive agents.

Multi-agent learning, i.e., the ability of the agents to learn how to cooperate and compete, becomes crucial in many domains. Autonomous agents and multi-agent systems (AAMAS) is an emerging multi-disciplinary area encompassing computer science, software engineering, biology, as well as cognitive and social sciences. A theoretical framework, in which rationality of learning and interacting agents can be understood, is still under development in MASs, although there have been promising first results.

The goal of this workshop was to increase awareness and interest in adaptive agent research, encourage collaboration between machine learning (ML) experts and agent system experts, and give a representative overview of current research in the area of adaptive agents. The symposium served as an inclusive forum for the discussion of ongoing or completed work concerning both theoretical and practical issues. An important part of the workshop was to model MASs for different applications and to develop robust ML techniques.

Contributions in this book cover topics on how an agent can learn, using ML techniques, to act individually or to coordinate with one another towards individual or common goals, which is still an open issue in real-time, noisy, collaborative and adversarial environments. The book start with an extensive overview article on cooperative and competitive multi-agent learning, which also contains a description of the contributions of this book and places them in context with the state of the art. It is a good starting point for newcomers in the field, wishing to read a self-contained overview of the state of the art in multi-agent learning, and also a good introduction for experts wishing to explore the contributions of this book.

We hope that our readers will enjoy reading the efforts of the researchers. A special word of gratitude also goes to our invited speakers, Peter Stone and Ann Nowé.

The first invited talk, “Multi-Robot Learning for Continuous Area Sweeping”, by Peter Stone, University of Texas at Austin, USA, has been shaped into the paper: “Multi-Robot Learning for Continuous Area Sweeping”, by Mazda Ahmadi and Peter Stone.

In their paper they study the problem of multi-agent continuous area sweeping. In this problem agents are situated in a particular environment in which they have to repeatedly visit every part of it such that they can detect events of interest for their global task and coordinate to minimize the total cost. Events are not uniformly distributed, such that agents need to visit locations non-uniformly. The authors formalize this problem and present an initial algorithm to solve it. Moreover they nicely illustrate their approach with a set of experiments in a routine surveillance task.

The second invited talk of the workshop was by Ann Nowé, professor of computer sciences at the university of Brussels, Belgium, resulting in the paper: “Learning Automata as a Basis for Multiagent Reinforcement Learning”, by Ann Nowé, Katja Verbeeck and Maarten Peeters. In their work they start with an overview on important theoretical results from the theory of learning automata in terms of game theoretic concepts and consider them as a policy iterator in the domain of reinforcement learning problems. Doing so they gradually move from the variable structure automaton, mapping to the single-stage single-agent case, over learning automata games, mapping to the single-stage multi-agent case, to interconnected learning automata, considering multi-stage multi-agent problems. The authors also show the most interesting connection with the field of ant colony optimization.

## Acknowledgements

When organizing a scientific event like LAMAS, a word of gratitude is always in place. This book would not have been produced without the help of many persons. First of all, the organizers would like to thank the members of the PC, who guaranteed a scientifically strong and interesting LNCS volume. Secondly, we would like to express our appreciation to the invited speakers, Ann Nowé and Peter Stone, for their distinguished contribution to the workshop program. Finally, we also would like to thank the authors of all contributions for submitting their scientific work to the LAMAS workshop!

December 2005

Karl Tuyls  
Pieter Jan 't Hoen  
Katja Verbeeck  
Sandip Sen  
Tulsa

# Organization

## Organizing Committee

Co-chairs:

Karl Tuyls  
Pieter Jan 't Hoen  
Katja Verbeeck  
Sandip Sen

## Program Committee

Stephane Airiau  
Bikramjit Banerjee  
Ana Lucia Bazzan  
Sander Bohté  
Michael Goodrich  
Daniel Kudenko

Han La Poutré  
Michael Littman  
Peter McBurney  
Ann Nowé  
Simon Parsons  
Steve Phelps

Jan Ramon  
Sandip Sen  
Peter Stone  
Kagan Tumer  
Danny Weyns  
David Wolpert

## Additional Referees

Jacob Crandall  
Alexander Helleboogh  
Koen Mertens

# Lecture Notes in Artificial Intelligence (LNAI)

- Vol. 3918: W.K. Ng, M. Kitsuregawa, J. Li, K. Chang (Eds.), *Advances in Knowledge Discovery and Data Mining*. XXIV, 879 pages. 2006.
- Vol. 3910: S.A. Brueckner, G.D.M. Serugendo, D. Hales, F. Zambonelli (Eds.), *Engineering Self-Organising Systems*. XII, 245 pages. 2006.
- Vol. 3904: M. Baldoni, U. Endriss, A. Omicini, P. Torroni (Eds.), *Declarative Agent Languages and Technologies III*. XII, 245 pages. 2006.
- Vol. 3899: S. Frintrop, *VOCUS: A Visual Attention System for Object Detection and Goal-Directed Search*. XIV, 216 pages. 2006.
- Vol. 3898: K. Tuyls, P.J. 't Hoen, K. Verbeeck, S. Sen (Eds.), *Learning and Adaption in Multi-Agent Systems*. X, 217 pages. 2006.
- Vol. 3891: J.S. Sichman, L. Antunes (Eds.), *Multi-Agent-Based Simulation VI*. X, 191 pages. 2006.
- Vol. 3890: S.G. Thompson, R. Ghanea-Hercock (Eds.), *Defence Applications of Multi-Agent Systems*. XII, 141 pages. 2006.
- Vol. 3885: V. Torra, Y. Narukawa, A. Valls, J. Domingo-Ferrer (Eds.), *Modeling Decisions for Artificial Intelligence*. XII, 374 pages. 2006.
- Vol. 3881: S. Gibet, N. Courty, J.-F. Kamp (Eds.), *Gesture in Human-Computer Interaction and Simulation*. XIII, 344 pages. 2006.
- Vol. 3874: R. Missaoui, J. Schmidt (Eds.), *Formal Concept Analysis*. X, 309 pages. 2006.
- Vol. 3873: L. Maicher, J. Park (Eds.), *Charting the Topic Maps Research and Applications Landscape*. VIII, 281 pages. 2006.
- Vol. 3863: M. Kohlhase (Ed.), *Mathematical Knowledge Management*. XI, 405 pages. 2006.
- Vol. 3862: R.H. Bordini, M. Dastani, J. Dix, A.E.F. Seghrouchni (Eds.), *Programming Multi-Agent Systems*. XIV, 267 pages. 2006.
- Vol. 3849: I. Bloch, A. Petrosino, A.G.B. Tettamanzi (Eds.), *Fuzzy Logic and Applications*. XIV, 438 pages. 2006.
- Vol. 3848: J.-F. Boulicaut, L. De Raedt, H. Mannila (Eds.), *Constraint-Based Mining and Inductive Databases*. X, 401 pages. 2006.
- Vol. 3847: K.P. Jantke, A. Lunzer, N. Spyrtos, Y. Tanaka (Eds.), *Federation over the Web*. X, 215 pages. 2006.
- Vol. 3835: G. Sutcliffe, A. Voronkov (Eds.), *Logic for Programming, Artificial Intelligence, and Reasoning*. XIV, 744 pages. 2005.
- Vol. 3830: D. Weyns, H. V.D. Parunak, F. Michel (Eds.), *Environments for Multi-Agent Systems II*. VIII, 291 pages. 2006.
- Vol. 3817: M. Faundez-Zanuy, L. Janer, A. Esposito, A. Satue-Villar, J. Roure, V. Espinosa-Duro (Eds.), *Nonlinear Analyses and Algorithms for Speech Processing*. XII, 380 pages. 2006.
- Vol. 3814: M. Maybury, O. Stock, W. Wahlster (Eds.), *Intelligent Technologies for Interactive Entertainment*. XV, 342 pages. 2005.
- Vol. 3809: S. Zhang, R. Jarvis (Eds.), *AI 2005: Advances in Artificial Intelligence*. XXVII, 1344 pages. 2005.
- Vol. 3808: C. Bento, A. Cardoso, G. Dias (Eds.), *Progress in Artificial Intelligence*. XVIII, 704 pages. 2005.
- Vol. 3802: Y. Hao, J. Liu, Y.-P. Wang, Y.-m. Cheung, H. Yin, L. Jiao, J. Ma, Y.-C. Jiao (Eds.), *Computational Intelligence and Security, Part II*. XLII, 1166 pages. 2005.
- Vol. 3801: Y. Hao, J. Liu, Y.-P. Wang, Y.-m. Cheung, H. Yin, L. Jiao, J. Ma, Y.-C. Jiao (Eds.), *Computational Intelligence and Security, Part I*. XLI, 1122 pages. 2005.
- Vol. 3789: A. Gelbukh, Á. de Albornoz, H. Terashima-Marín (Eds.), *MICAI 2005: Advances in Artificial Intelligence*. XXVI, 1198 pages. 2005.
- Vol. 3782: K.-D. Althoff, A. Dengel, R. Bergmann, M. Nick, T.R. Roth-Berghofer (Eds.), *Professional Knowledge Management*. XXIII, 739 pages. 2005.
- Vol. 3763: H. Hong, D. Wang (Eds.), *Automated Deduction in Geometry*. X, 213 pages. 2006.
- Vol. 3755: G.J. Williams, S.J. Simoff (Eds.), *Data Mining*. XI, 331 pages. 2006.
- Vol. 3735: A. Hoffmann, H. Motoda, T. Scheffer (Eds.), *Discovery Science*. XVI, 400 pages. 2005.
- Vol. 3734: S. Jain, H.U. Simon, E. Tomita (Eds.), *Algorithmic Learning Theory*. XII, 490 pages. 2005.
- Vol. 3721: A.M. Jorge, L. Torgo, P.B. Brazdil, R. Camacho, J. Gama (Eds.), *Knowledge Discovery in Databases: PKDD 2005*. XXIII, 719 pages. 2005.
- Vol. 3720: J. Gama, R. Camacho, P.B. Brazdil, A.M. Jorge, L. Torgo (Eds.), *Machine Learning: ECML 2005*. XXIII, 769 pages. 2005.
- Vol. 3717: B. Gramlich (Ed.), *Frontiers of Combining Systems*. X, 321 pages. 2005.
- Vol. 3702: B. Beckert (Ed.), *Automated Reasoning with Analytic Tableaux and Related Methods*. XIII, 343 pages. 2005.
- Vol. 3698: U. Furbach (Ed.), *KI 2005: Advances in Artificial Intelligence*. XIII, 409 pages. 2005.
- Vol. 3690: M. Pěchouček, P. Petta, L.Z. Varga (Eds.), *Multi-Agent Systems and Applications IV*. XVII, 667 pages. 2005.

- Vol. 3684: R. Khosla, R.J. Howlett, L.C. Jain (Eds.), Knowledge-Based Intelligent Information and Engineering Systems, Part IV. LXXIX, 933 pages. 2005.
- Vol. 3683: R. Khosla, R.J. Howlett, L.C. Jain (Eds.), Knowledge-Based Intelligent Information and Engineering Systems, Part III. LXXX, 1397 pages. 2005.
- Vol. 3682: R. Khosla, R.J. Howlett, L.C. Jain (Eds.), Knowledge-Based Intelligent Information and Engineering Systems, Part II. LXXIX, 1371 pages. 2005.
- Vol. 3681: R. Khosla, R.J. Howlett, L.C. Jain (Eds.), Knowledge-Based Intelligent Information and Engineering Systems, Part I. LXXX, 1319 pages. 2005.
- Vol. 3673: S. Bandini, S. Manzoni (Eds.), AI\*IA 2005: Advances in Artificial Intelligence. XIV, 614 pages. 2005.
- Vol. 3662: C. Baral, G. Greco, N. Leone, G. Terracina (Eds.), Logic Programming and Nonmonotonic Reasoning. XIII, 454 pages. 2005.
- Vol. 3661: T. Panayiotopoulos, J. Gratch, R.S. Aylett, D. Ballin, P. Olivier, T. Rist (Eds.), Intelligent Virtual Agents. XIII, 506 pages. 2005.
- Vol. 3658: V. Matoušek, P. Mautner, T. Pavelka (Eds.), Text, Speech and Dialogue. XV, 460 pages. 2005.
- Vol. 3651: R. Dale, K.-F. Wong, J. Su, O.Y. Kwong (Eds.), Natural Language Processing – IJCNLP 2005. XXI, 1031 pages. 2005.
- Vol. 3642: D. Ślęzak, J. Yao, J.F. Peters, W. Ziarko, X. Hu (Eds.), Rough Sets, Fuzzy Sets, Data Mining, and Granular Computing, Part II. XXIII, 738 pages. 2005.
- Vol. 3641: D. Ślęzak, G. Wang, M. Szczuka, I. Düntsch, Y. Yao (Eds.), Rough Sets, Fuzzy Sets, Data Mining, and Granular Computing, Part I. XXIV, 742 pages. 2005.
- Vol. 3635: J.R. Winkler, M. Niranjan, N.D. Lawrence (Eds.), Deterministic and Statistical Methods in Machine Learning. VIII, 341 pages. 2005.
- Vol. 3632: R. Nieuwenhuis (Ed.), Automated Deduction – CADE-20. XIII, 459 pages. 2005.
- Vol. 3630: M.S. Capcarrère, A.A. Freitas, P.J. Bentley, C.G. Johnson, J. Timmis (Eds.), Advances in Artificial Life. XIX, 949 pages. 2005.
- Vol. 3626: B. Ganter, G. Stumme, R. Wille (Eds.), Formal Concept Analysis. X, 349 pages. 2005.
- Vol. 3625: S. Kramer, B. Pfahringer (Eds.), Inductive Logic Programming. XIII, 427 pages. 2005.
- Vol. 3620: H. Muñoz-Ávila, F. Ricci (Eds.), Case-Based Reasoning Research and Development. XV, 654 pages. 2005.
- Vol. 3614: L. Wang, Y. Jin (Eds.), Fuzzy Systems and Knowledge Discovery, Part II. XLI, 1314 pages. 2005.
- Vol. 3613: L. Wang, Y. Jin (Eds.), Fuzzy Systems and Knowledge Discovery, Part I. XLI, 1334 pages. 2005.
- Vol. 3607: J.-D. Zucker, L. Saitta (Eds.), Abstraction, Reformulation and Approximation. XII, 376 pages. 2005.
- Vol. 3601: G. Moro, S. Bergamaschi, K. Aberer (Eds.), Agents and Peer-to-Peer Computing. XII, 245 pages. 2005.
- Vol. 3600: F. Wiedijk (Ed.), The Seventeen Provers of the World. XVI, 159 pages. 2006.
- Vol. 3596: F. Dau, M.-L. Mugnier, G. Stumme (Eds.), Conceptual Structures: Common Semantics for Sharing Knowledge. XI, 467 pages. 2005.
- Vol. 3593: V. Mařík, R. W. Brennan, M. Pěchouček (Eds.), Holonic and Multi-Agent Systems for Manufacturing. XI, 269 pages. 2005.
- Vol. 3587: P. Perner, A. Imiya (Eds.), Machine Learning and Data Mining in Pattern Recognition. XVII, 695 pages. 2005.
- Vol. 3584: X. Li, S. Wang, Z.Y. Dong (Eds.), Advanced Data Mining and Applications. XIX, 835 pages. 2005.
- Vol. 3581: S. Miksch, J. Hunter, E.T. Keravnou (Eds.), Artificial Intelligence in Medicine. XVII, 547 pages. 2005.
- Vol. 3577: R. Falcone, S. Barber, J. Sabater-Mir, M.P. Singh (Eds.), Trusting Agents for Trusting Electronic Societies. VIII, 235 pages. 2005.
- Vol. 3575: S. Wernter, G. Palm, M. Elshaw (Eds.), Biomimetic Neural Learning for Intelligent Robots. IX, 383 pages. 2005.
- Vol. 3571: L. Godo (Ed.), Symbolic and Quantitative Approaches to Reasoning with Uncertainty. XVI, 1028 pages. 2005.
- Vol. 3559: P. Auer, R. Meir (Eds.), Learning Theory. XI, 692 pages. 2005.
- Vol. 3558: V. Torra, Y. Narukawa, S. Miyamoto (Eds.), Modeling Decisions for Artificial Intelligence. XII, 470 pages. 2005.
- Vol. 3554: A.K. Dey, B. Kokinov, D.B. Leake, R. Turner (Eds.), Modeling and Using Context. XIV, 572 pages. 2005.
- Vol. 3550: T. Eymann, F. Klügl, W. Lamersdorf, M. Klusch, M.N. Huhns (Eds.), Multiagent System Technologies. XI, 246 pages. 2005.
- Vol. 3539: K. Morik, J.-F. Boulicaut, A. Siebes (Eds.), Local Pattern Detection. XI, 233 pages. 2005.
- Vol. 3538: L. Ardissono, P. Brna, A. Mitrović (Eds.), User Modeling 2005. XVI, 533 pages. 2005.
- Vol. 3533: M. Ali, F. Esposito (Eds.), Innovations in Applied Artificial Intelligence. XX, 858 pages. 2005.
- Vol. 3528: P.S. Szczepaniak, J. Kacprzyk, A. Niewiadomski (Eds.), Advances in Web Intelligence. XVII, 513 pages. 2005.
- Vol. 3518: T.-B. Ho, D. Cheung, H. Liu (Eds.), Advances in Knowledge Discovery and Data Mining. XXI, 864 pages. 2005.
- Vol. 3508: P. Bresciani, P. Giorgini, B. Henderson-Sellers, G. Low, M. Winikoff (Eds.), Agent-Oriented Information Systems II. X, 227 pages. 2005.
- Vol. 3505: V. Gorodetsky, J. Liu, V.A. Skormin (Eds.), Autonomous Intelligent Systems: Agents and Data Mining. XIII, 303 pages. 2005.
- Vol. 3501: B. Kégl, G. Lapalme (Eds.), Advances in Artificial Intelligence. XV, 458 pages. 2005.
- Vol. 3492: P. Blache, E.P. Stabler, J.V. Busquets, R. Moot (Eds.), Logical Aspects of Computational Linguistics. X, 363 pages. 2005.
- Vol. 3490: L. Bolc, Z. Michalewicz, T. Nishida (Eds.), Intelligent Media Technology for Communicative Intelligence. X, 259 pages. 2005.

¥359.00元

# Table of Contents

An Overview of Cooperative and Competitive Multiagent Learning <i>Pieter Jan 't Hoen, Karl Tuyls, Liviu Panait, Sean Luke, J.A. La Poutré</i> .....	1
Multi-robot Learning for Continuous Area Sweeping <i>Mazda Ahmadi, Peter Stone</i> .....	47
Learning Automata as a Basis for Multi Agent Reinforcement Learning <i>Ann Nowé, Katja Verbeeck, Maarten Peeters</i> .....	71
Learning Pareto-optimal Solutions in 2x2 Conflict Games <i>Stéphane Airiau, Sandip Sen</i> .....	86
Unifying Convergence and No-Regret in Multiagent Learning <i>Bikramjit Banerjee, Jing Peng</i> .....	100
Implicit Coordination in a Network of Social Drivers: The Role of Information in a Commuting Scenario <i>Ana L.C. Bazzan, Manuel Fehler, Franziska Klügl</i> .....	115
Multiagent Traffic Management: Opportunities for Multiagent Learning <i>Kurt Dresner, Peter Stone</i> .....	129
Dealing with Errors in a Cooperative Multi-agent Learning System <i>Constança Oliveira e Sousa, Luis Custódio</i> .....	139
The Success and Failure of Tag-Mediated Evolution of Cooperation <i>Austin McDonald, Sandip Sen</i> .....	155
An Adaptive Approach for the Exploration-Exploitation Dilemma and Its Application to Economic Systems <i>Lilia Rejeb, Zahia Guessoum, Rym M'Hallah</i> .....	165
Efficient Reward Functions for Adaptive Multi-rover Systems <i>Kagan Tumer, Adrian Agogino</i> .....	177
Multi-agent Relational Reinforcement Learning <i>Tom Croonenborghs, Karl Tuyls, Jan Ramon, Maurice Bruynooghe</i> .....	192

Multi-type ACO for Light Path Protection  
    *Peter Vrancx, Ann Nowé, Kris Steenhaut* ..... 207

**Author Index** ..... 217



# An Overview of Cooperative and Competitive Multiagent Learning

Pieter Jan 't Hoen<sup>1</sup>, Karl Tuyls<sup>2</sup>, Liviu Panait<sup>3</sup>,  
Sean Luke<sup>3</sup>, and J.A. La Poutré<sup>1,4</sup>

<sup>1</sup> Center for Mathematics and Computer Science (CWI),  
P.O. Box 94079, Amsterdam 1090 GB, The Netherlands

<sup>2</sup> Computer Science Department (IKAT), Tongersestraat 6,  
University of Maastricht, The Netherlands

<sup>3</sup> George Mason University, Fairfax, VA 22030

<sup>4</sup> TU Eindhoven, De Lismortel 2, Eindhoven 5600 MB, The Netherlands  
hoen@cwi.nl, k.tuyls@cs.unimaas.nl,  
{lpanait, sean}@cs.gmu.edu, hlp@cwi.nl

**Abstract.** *Multi-agent systems* (MASs) is an area of distributed artificial intelligence that emphasizes the joint behaviors of agents with some degree of autonomy and the complexities arising from their interactions. The research on MASs is intensifying, as supported by a growing number of conferences, workshops, and journal papers. In this survey we give an overview of multi-agent learning research in a spectrum of areas, including reinforcement learning, evolutionary computation, game theory, complex systems, agent modeling, and robotics.

MASs range in their description from cooperative to being competitive in nature. To muddle the waters, competitive systems can show apparent cooperative behavior, and vice versa. In practice, agents can show a wide range of behaviors in a system, that may either fit the label of cooperative or competitive, depending on the circumstances. In this survey, we discuss current work on cooperative and competitive MASs and aim to make the distinctions and overlap between the two approaches more explicit.

Lastly, this paper summarizes the papers of the first International workshop on Learning and Adaptation in MAS (LAMAS) hosted at the fourth International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS'05) and places the work in the above survey.

## 1 Introduction

*Multi-agent systems* (MASs) is an area of distributed artificial intelligence that emphasizes the joint behaviors of agents with some degree of autonomy and the complexities arising from their interactions. The research on MASs is intensifying, as supported by a growing number of conferences, workshops, and journal papers. This book of the first International workshop on Learning and Adaptation in MAS (LAMAS), hosted at the fourth International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS'05), is a continuation of this trend.

The goal of the LAMAS workshop was to increase awareness and interest in adaptive agent research, encourage collaboration between Machine Learning (ML) experts and agent system experts, and give a representative overview of current research in the area of adaptive agents. The workshop served as an inclusive forum for the discussion of ongoing or completed work concerning both theoretical and practical issues. More precisely, researchers from the multi-agent learning community presented recent work and discussed their newest ideas for a first time with their peers. An important part of the workshop was dedicated to model MASs for different applications and to develop robust ML techniques. Contributions cover on how an agent can learn using ML techniques to act individually or to coordinate with one another towards individual or common goals. This is an open issue in real-time, noisy, collaborative and possibly adversarial environments.

This introductory article has a twofold goal. The first is to give a broad overview of current MASs research. We present our overview of MASs research from the two main perspectives to be found in the literature; the cooperative and competitive perspective. Secondly, we briefly present an overview of the included papers and invited contributions and place them in the global context of ongoing research.

In cooperative systems, as suggested by the label, the agents pursue a common goal. Such systems are characterized by the fact that the designers of the MAS are free in their design of the agents. The agents can be built and learn with extensive knowledge of the system and the agents can expect benevolent intentions from other agents. Note that we do not claim that it is easy to design a cooperative MAS to have good emergent behavior, on the contrary!

In contrast to cooperative MASs, agents in a competitive MAS setting have non-aligned goals, and individual agents seek only to maximize their own gains. Recent work in competitive MASs has aimed at moving Reinforcement Learning (RL) techniques from the domain of single-agent to multi-agent settings. There is a growing body of work, algorithms and evaluation criteria, which we cover in the second part of our survey. Furthermore, this section also covers a growing body of work on non-cooperative agents [189] for economical and societal settings that have received increasing interest only in recent years. Such agents have their own, possibly conflicting goals and aim for local optimization. Their owners can e.g. be competing companies or autonomous departments within a bigger organization, where the multi-agent systems should facilitate trading, allocation, or planning between these owners, e.g. by means of negotiation or auctioning.

The rest of this document is structured as follows. Section 2 first informally introduces agents playing simple matrix games. We use this section to initially introduce the concepts of play, and whether the agents can be labeled as cooperative, competitive, or as something in between. Section 3 presents our overview of cooperative MASs. Section 4 continues with our overview of competitive MASs. Sections 3 and 4 are intended to be largely self contained, although there are cross-links between the sections. Section 5 presents the papers of this LAMAS proceedings and places this work in the context of the survey of MASs work, Sections 3 and 4. Lastly, Section 6 concludes with an agenda of future research opportunities for MASs. Appendix A includes some basic Game Theory (GT)

concepts universal to the domain of cooperative and competitive MASs as a general background for readers not familiar with the subject.

The next section continues with a discussion on the labels of cooperative and competitive as applied to MASs.

## 2 Agents Classified as Cooperative or Competitive

Multi-Agent Systems range in their description from cooperative to being competitive in nature. To muddle the waters, competitive systems can show apparent cooperative behavior, and vice versa. In practice, agents in a system, depending on the circumstances, can show a wide range of behaviors that may either fit the label of cooperative or competitive.

The fundamental distinction between systems labeled as cooperative or competitive is that for the former the agents are designed with as goal the maximization of a group utility. Competitive agents are solely focused on maximizing their own utility. We, in this section, label the agents as either utilitarian or selfish to stress more their intention, i.e. their design goal, than their actual behavior. For example, a competitive/selfish agent may cooperate with other agents in a temporary coalition. The selfish intentions of the agent are met due to a larger expected reward from cooperation. On the other hand, a cooperative/utilitarian agent may seem competitive if it accidentally hogs a resource to the detriment of other agents in its group. In complex cooperative systems, agents can easily hinder the other agents as the complexity of the interactions increase. The label utilitarian or selfish stresses more the intentional stance of the agent (and of its designer), as opposed to its apparent behavior.

The utilitarian stance for cooperative systems, as already mentioned in the introduction, is also reflected in the design of the agents. Commonly, a cooperative system is designed by one party (be that one designer or a team) to achieve a set of agreed upon goals. The behavior, or the algorithm that learns the behavior of the agents, is largely under the control of the designers of the system. This allows for possible intricate coordination to be a priori implemented in the system and many interactions in the system can be anticipated. An agent can essentially expect good intentions from other agents in the system. This is not the case for the competitive setting. Each agent is created by separate designers that all aim to achieve their own goals. This makes cooperation between selfish agents, even if this is rational, a more difficult and risky task. The designer of a competitive agent must also expend effort in considering the types of exploitive behavior that will be encountered. This distinction in design of agents for a cooperative or competitive setting must be kept in mind when choosing the range of strategies the agents can choose from.

### 2.1 Setting

In the following, we give a sample of the type of interactions that can be observed between agents. We discuss how these are a consequence of the utilitarian or selfish intentional stance.

We restrict our discussion to the well known two-agent, two-action matrix games. For a complete taxonomy we refer the reader to [132]. Of importance is that the listed games give an exhaustive overview of the types of settings that the agents can encounter. This gives a sound basis to inspect how agents can handle these types of games, both from the utilitarian and the selfish stance. We can then classify the agent behavior as either (apparent) cooperative, (apparent) competitive, or indistinguishable.

[132] classifies games from the perspective of selfish agents; the agents focus on maximizing their own gain, i.e. their private utility. Game theoretical notions prevail in the discussion of the choice of strategies of the agents. We take a slightly broader view and also focus on utilitarian agents and how they would play in the selected games. Utilitarian agents focus on achieving the highest possible group utility, i.e. the sum of their individual rewards.

Note that we only consider play between two selfish agents or between two utilitarian agents. We consider either a system of agents where all agents are intended to achieve a common goal, or a system of agents where all agents expect the worst. We do not cover the intricacies of a cooperative system that has to deal with selfish agents. For a more complete discussion of this topic, we refer the reader to [106] and Section A for a discussion on Evolutionary Stable Strategies.

The agents in the games know the complete payoff matrices. They know their own reward and that of their opponents for all joint actions<sup>1</sup>. They simultaneously must choose an action and receive their part of the reward based on the picked joint action. What they may not know is how the other agent, be that a malicious opponent or a benevolent agent, will play.

Note that we here as yet restrict ourselves to the single play of the presented matrix games. Agents may also have to learn these payoffs during repeated play of the game. We will give examples of this, along with a more formal treatment, in Section 4. After this initial exposition, we discuss how the choice of strategies can change due to repeated play.

## 2.2 Types of Games

From the viewpoint of selfish agents, [132] broadly classifies the matrix games as either trivial, games of no conflict, games of complete opposition, or as games of partial conflict. The latter is also called a mixed motive game. We discuss each of the categories below. For each category, we sketch the game, give an example, and discuss how selfish and utilitarian agents would cope with the game.

*Trivial games:* In trivial games (TG), the expected reward of an agent does not depend on the choice of action of the other agent. In Table 1, we show such a trivial game. The Row player can choose either action  $A1$  or  $B1$  while the Column player can choose from actions  $A2$  or  $B2$ . The items in the table show the rewards for the Row player and Column player respectively for choice of action  $Ai$  or  $Bi$  respectively. For this game, the rewards of one player are not

---

<sup>1</sup> Agents in most Game Theory literature know the payoff matrix before play.

**Table 1.** A trivial game

TG	A2	B2
A1	2,2	2,2
B1	2,2	2,2

influenced by the choice of actions of the opponent. Such a game is therefore not of great interest in terms of formulating a best strategy. This strategy is based on what they think they should play given the logical action chosen by the opponent, a non-issue in this case.

Due to the simple nature of this game, there is no intrinsic difference in play between utilitarian and selfish agents.

*No conflict games:* In no-conflict games (NCG), both players benefit from choosing one, unambiguous joint action. Neither player benefits, in terms of individual rewards, by deviating from this logical choice. Consider the game in Table 2:

**Table 2.** A no-conflict game

NCG	A2	B2
A1	4,4	2,3
B1	3,2	2,2

Both the Row and Column player prefer the joint action  $A1A2$  (we give first the Row, and the Column player action) as this gives the most individual reward. Neither player has an incentive to choose another action when the sole goal is maximizing the private utility for selfish agents.  $A1A2$  is also the logical choice of action for the utilitarian players. We stress that in both cases, the Row and Column player individually choose  $A1$  and  $A2$  respectively without prior negotiations; the players base their individual choice solely on their own strategic reasoning.

Note that the Row player may prefer to play  $B1A2$  when the Row player aims to maximize the relative utility of play; the Row player wants to have more utility than the Column player. This aspect is not an issue for utilitarian players.

As for trivial games, there is little difference in play between utilitarian and selfish agents. One distinction that can be made is that a utilitarian row player will not pick action  $B1$  as such a player is not interested in achieving a higher reward than the other player. More importantly, this choice of action will lower the utility of the group and should be avoided.

*Games of Complete opposition (also known as zerosum games):* In games of complete opposition (CO), the gain of one agent is a loss for the other agent. The Table 3 shows a typical zerosum game (rewards for one joint action sum to 0). These games are characterized by fierce competition. On average, an agent can expect to have zero reward.

For selfish players, games of complete opposition are a difficult scenario. The best strategy for an unknown opponent, from a game theoretical viewpoint, is

Table 3. A game of complete opposition

CO	A2	B2
A1	0,0	2,-2
B1	1,-1	-3,3

to play a random strategy; all actions are equally probable. More technically, this is a mixed strategy. See Section A for a more formal definition. For two utilitarian agents, the game is also problematical as coordination of joint actions, by definition of a zerosum game, will not lead to a higher aggregated reward.

*Games of Partial Conflict Mixed Motive Games:* Games of partial conflict (PC) allow for both agents to choose profitable actions, but the agents prefer different joint actions. The latter point is the distinction between the no-conflict games and the mixed-motive games. We give an example in Table 4.

Table 4. A partial conflict game

PC1	A2	B2
A1	2,7	-1,-10
B1	1,-5	10,1

The Row agent prefers joint action *B1B2*. The Column agent prefers joint action *A1A2*. Blindly choosing *B1* by the Row player and *A2* by the Column player results in joint action *B1A2* that is preferred by neither player.

Games of partial conflict are difficult for selfish agents. Optimal play is achieved through a mixed strategy that maximizes expected utility. This aspect is handled in more detail in Appendix A.

Utilitarian agents that have as goal to maximize the group utility have a more clearcut strategy; choose the joint action that maximizes the total utility. For Table 4, joint action *B1B2* is the clear choice. For Table 5, the utilitarian agents are however faced with the choice of playing joint action *A1A2* or *B1B2*. The agents must however make their choices individually, with no a priori information of the action that will be played by the other agent.

Table 5. A second no-conflict game

PC2	A2	B2
A1	3,3	1,1
B1	1,1	3,3

2.3 Repeated Play

The above section has presented play for agents for single shot play of a selection of typical matrix games. We now focus on how the game can change if two agents repeatedly play the same game. Repeated play opens opportunities, especially to selfish agents, not available in single shot play of the game.