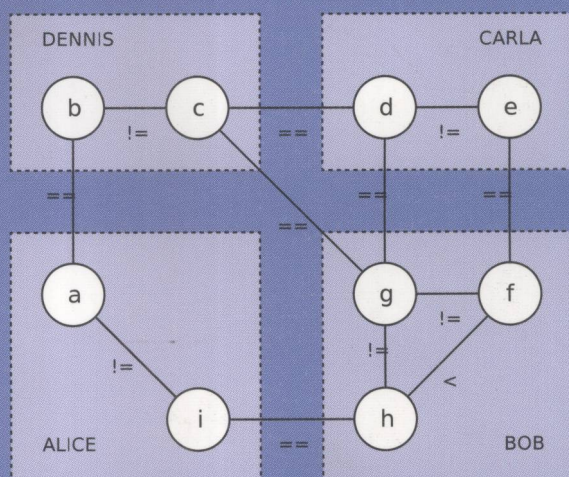


Nicolas Maudet
Simon Parsons
Iyad Rahwan (Eds.)

Argumentation in Multi-Agent Systems

Third International Workshop, ArgMAS 2006
Hakodate, Japan, May 2006
Revised Selected and Invited Papers



Springer

TP18-53

A694
2006

Nicolas Maudet Simon Parsons
Iyad Rahwan (Eds.)

Argumentation in Multi-Agent Systems

Third International Workshop, ArgMAS 2006
Hakodate, Japan, May 8, 2006
Revised Selected and Invited Papers



Springer



E2007003601

Series Editors

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

Volume Editors

Nicolas Maudet
Université Paris-Dauphine
LAMSADE
75775 Paris Cedex 16, France
E-mail: maudet@lamsade.dauphine.fr

Simon Parsons
City University of New York
Brooklyn College
Department of Computer & Information Science
2900 Bedford Avenue, Brooklyn, NY 11210, USA
E-mail: parsons@sci.brooklyn.cuny.edu

Iyad Rahwan
The British University in Dubai
Institute of Informatics
P.O. Box 502216, Dubai, UAE
E-mail: irahwan@acm.org

Library of Congress Control Number: 2007937617

CR Subject Classification (1998): I.2.11, I.2, C.2.4, H.5.2-3

LNCS Sublibrary: SL 7 – Artificial Intelligence

ISSN 0302-9743
ISBN-10 3-540-75525-X Springer Berlin Heidelberg New York
ISBN-13 978-3-540-75525-8 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 2007
Printed in Germany

Typesetting: Camera-ready by author, data conversion by Scientific Publishing Services, Chennai, India
Printed on acid-free paper SPIN: 12171521 06/3180 5 4 3 2 1 0

Lecture Notes in Artificial Intelligence

4766

Edited by J. G. Carbonell and J. Siekmann

Subseries of Lecture Notes in Computer Science

Preface

This volume presents the recent developments of the growing area of research taking place at the interface of argumentation theory and multiagent systems. Argumentation can be abstractly defined as the interaction of different arguments for and against some conclusion. Over the last few years, argumentation has been gaining increasing importance in multiagent systems, mainly as a vehicle for facilitating “rational interaction” (i.e., interaction which involves the giving and receiving of reasons). This is because argumentation provides tools for designing, implementing and analyzing sophisticated forms of interaction among rational agents. Argumentation has made solid contributions to the practice of multiagent dialogues. Application domains include: legal disputes, business negotiation, labor disputes, team formation, scientific inquiry, deliberative democracy, ontology reconciliation, risk analysis, scheduling, and logistics. A single agent may also use argumentation techniques to perform its individual reasoning because it needs to make decisions under complex preferences policies, in a highly dynamic environment.

Following the success of its two first editions, the International Workshop on Argumentation in Multiagent Systems (ArgMAS 2006) took place for the third time in May 2006 in Hakodate, Japan, as a satellite workshop of the Autonomous Agents and Multiagent Systems conference. The workshop series is concerned with the use of the concepts, theories, methodologies, and computational models of argumentation in building autonomous agents and multiagent systems. In particular, the workshop aims at bridging the gap between the vast amount of work on argumentation theory and the practical needs of multiagent systems research. While the revised contributions of ArgMAS 2006 indeed constitute the backbone of this volume, it also includes revised versions of papers presented in recent conferences: Autonomous Agents and Multiagent Systems (AAMAS 2006), and the European Conference on Artificial Intelligence (ECAI 2006). These additional contributions were selected on the basis of their scientific quality and relevance to the topics emphasized here. Our objective has been to offer a comprehensive and up-to-date overview of this rapidly evolving landscape, as we did in the previous volumes of this series (LNAI 3366, LNAI 4049).

This book opens with a brief survey paper (“Argumentation in Multiagent Systems: Context and Recent Developments”) by the editors, which aims at presenting the broad framework of the volume. Light is shed more specifically on a couple of “hot topics.”

The rest of the book is then divided into two parts. The first one is dedicated to the exploration of the fundamentals and possible (and desirable in agent systems) extensions of argumentation-based reasoning (“Foundations and Explorations”). For instance, most argumentation frameworks do not really cater

for the dynamic aspects of multiagent systems since they assume fixed knowledge bases to start with. Two papers of this volume specifically address this issue. Fukumoto and Sawamura investigate how argumentation may result in a modification of agents' beliefs. They tackle this problem by introducing a new learning method based on argumentation, developed in line with the logic programming paradigm, but necessitating different extensions. In the context of argumentation-based joint deliberation, Ontañón and Plaza study how learning agents can make use of past examples to generate arguments and counter-arguments as to what course of action should be taken in a given situation. Using a specific bilateral protocol, they show that the overall performance of the system is improved because joint predictions resulting from this process are typically more accurate than individual agent prediction. One other well-known limitation of Dung's original abstract framework is that it does not allow for coalitions of arguments to attack other arguments. Nielsen and Parsons explore which semantics can be defined when such a possibility is taken into account. While all the aforementioned papers are concerned with epistemic reasoning, Rahwan and Amgoud present an approach that puts together the different pieces of an argumentation-based agent. Indeed, different argumentation frameworks can be integrated to manage not only beliefs, but also desires and plans intended to achieve these desires. This capacity to reason on the basis of different attitudes is a crucial component of autonomous deliberative agents, as witnessed and argued by BDI-like agenthood theories. Finally, Harvey, Chang and Ghose show how argumentation can be used to enhance some aspects of distributed constraint satisfaction algorithms. Agents (variables) argue about partial assignments (of variables), by exhibiting counter-examples and making counter-proposals. The technique proposed in this paper makes it possible to resolve the problem of cycles without relying on a total ordering of the agents. In the last paper of this part of the volume, Karunatillake and colleagues present an empirical study of the use of argumentation-based negotiation as a means to manage conflict involving "social influences" in societies of agents. This kind of conflict will typically occur in environments where not all roles and relationships (and obligations attached to them) can be assumed to be known in advance. They show that, in this context of study, argumentation-based interaction is an improvement both in terms of efficiency and effectiveness over non-argumentative approaches.

The second part of the book is dedicated to a more specific but highly challenging question (as witnessed by the number of contributions related to that topic during the workshop): how should agents select arguments when engaged in complex interactions ("Strategic Issues")? Amgoud and Hamerlain regard the strategy problem as a two-step decision process: first select the preferred speech act, then select the best content to instantiate this speech act. What is shown in this paper is that these two steps involve different types of beliefs and goals. As a consequence, the formal framework for defining strategies is composed of two different systems, both grounded on argumentation theory. One especially important parameter of the resulting decision problems is provided by agents' generic profiles (e.g., cautious or adventurous), that is, attitudes regarding argument-

based comparison of candidate decisions. Mbarki, Bentahar, and Moulin make a slightly different distinction: they distinguish (dynamic) strategies (which involves global planning of an agent communication, in terms of sub-goals to be achieved), and tactics (which amounts to selecting the best argument with respect to the selected strategy). Each tactic is attached to a sub-goal selected at the strategy level. This articulation, often overlooked by other approaches, is at the core of the formal framework they propose. Oren, Norman, and Preece investigate two specific heuristics for dialogue move selection: one simply consists in revealing as little information as necessary in a given context; the second one involves a more sophisticated computation to assess the utility cost induced by revealing a given piece of information. Such heuristics make sense in particular in domains where privacy concerns are important, hence the need to understand more precisely how they can affect dialogue outcomes. Another interesting specific negotiation strategy is explored by Ramchurn et al. in the context of repeated interactions (that is, when agents typically interact more than once). Here, arguments are seen as promises of rewards in future interactions. Their strategy, which is based on a reward generation algorithm, achieves better outcomes than standard negotiation algorithms. On a slightly different tone, in the last paper of this book, Pasquier and colleagues develop an approach which accounts for the generative aspects of argumentative communication. Departing from the mainstream dialectical line of research, they ground their proposal on the notion of cognitive coherence, a theory coming from behavioral cognitive science.

We conclude this preface by extending our gratitude to the members of the Steering Committee, members of the Program Committee, and the auxiliary reviewers, who together helped make the ArgMAS workshop a success. We also thank the authors for their enthusiasm in submitting papers to the workshop, and for revising their papers on time for inclusion in this book.

May 2007

Nicolas Maudet
Simon Parsons
Iyad Rahwan

Organization

Program Chairs

Nicolas Maudet	Université Paris-Dauphine, France
Simon Parsons	City University of New York, USA
Iyad Rahwan	British University in Dubai, UAE (Fellow) University of Edinburgh, UK

ArgMAS Steering Committee

Antonis Kakas	University of Cyprus, Cyprus
Nicolas Maudet	Université Paris-Dauphine, France
Peter McBurney	University of Liverpool, UK
Pavlos Moraitis	Université René Descartes-Paris 5, France
Simon Parsons	City University of New York, USA
Iyad Rahwan	British University in Dubai, UAE (Fellow) University of Edinburgh, UK
Chris Reed	University of Dundee, UK

Program Committee

Leila Amgoud	IRIT, Toulouse, France
Katie Atkinson	University of Liverpool, UK
Jamal Bentahar	Laval University, Canada
Carlos Chesnevar	Universitat de Lleida, Spain
Frank Dignum	Utrecht University, The Netherlands
Rogier van Eijk	Utrecht University, The Netherlands
Anthony Hunter	University College London, UK
Antonis Kakas	University of Cyprus, Cyprus
Nikos Karacapilidis	University of Patras, Greece
Nicolas Maudet	Université Paris-Dauphine, France
Peter McBurney	University of Liverpool, UK
Jarred McGinnis	University of Edinburgh, UK
Pavlos Moraitis	Université René Descartes-Paris 5, France
Xavier Parent	King's College, UK
Simon Parsons	City University of New York, USA
Philippe Pasquier	University of Melbourne, Australia
Henry Prakken	Utrecht University, The Netherlands
	University of Groningen, The Netherlands
Iyad Rahwan	British University in Dubai, UAE (Fellow) University of Edinburgh, UK

Chris Reed	University of Dundee, UK
Carles Sierra	IIIA, Spain
Guillermo Simari	Universidad Nacional del Sur, Argentina
Katia Sycara	Carnegie Mellon University, USA
Francesca Toni	Imperial College, London, UK
Paolo Torroni	Università di Bologna, Italy
Bart Verheij	Maastricht University, The Netherlands
Gerard Vreeswijk	Utrecht University, The Netherlands
Mike Wooldridge	University of Liverpool, UK

Auxiliary Referees

P.M. Dung	Asian Institute of Technology, Thailand
-----------	---

Lecture Notes in Artificial Intelligence (LNAI)

- Vol. 4790: N. Dershowitz, A. Voronkov (Eds.), *Logic for Programming, Artificial Intelligence, and Reasoning*. XIII, 562 pages. 2007.
- Vol. 4772: H. Prade, V.S. Subrahmanian (Eds.), *Scalable Uncertainty Management*. X, 277 pages. 2007.
- Vol. 4766: N. Maudet, S. Parsons, I. Rahwan (Eds.), *Argumentation in Multi-Agent Systems*. XII, 211 pages. 2007.
- Vol. 4755: V. Corruble, M. Takeda, E. Suzuki (Eds.), *Discovery Science*. XI, 298 pages. 2007.
- Vol. 4754: M. Hutter, R.A. Servidio, E. Takimoto (Eds.), *Algorithmic Learning Theory*. XI, 403 pages. 2007.
- Vol. 4737: B. Berendt, A. Hotho, D. Mladenic, G. Semeraro (Eds.), *From Web to Social Web: Discovering and Deploying User and Content Profiles*. XI, 161 pages. 2007.
- Vol. 4733: R. Basili, M.T. Pazienza (Eds.), *AI*IA 2007: Artificial Intelligence and Human-Oriented Computing*. XVII, 858 pages. 2007.
- Vol. 4724: K. Mellouli (Ed.), *Symbolic and Quantitative Approaches to Reasoning with Uncertainty*. XV, 914 pages. 2007.
- Vol. 4722: C. Pelachaud, J.-C. Martin, E. André, G. Chollet, K. Karpouzis, D. Pelé (Eds.), *Intelligent Virtual Agents*. XV, 425 pages. 2007.
- Vol. 4720: B. Konev, F. Wolter (Eds.), *Frontiers of Combining Systems*. X, 283 pages. 2007.
- Vol. 4702: J.N. Kok, J. Koronacki, R. Lopez de Mantaras, S. Matwin, D. Mladenič, A. Skowron (Eds.), *Knowledge Discovery in Databases: PKDD 2007*. XXIV, 640 pages. 2007.
- Vol. 4701: J.N. Kok, J. Koronacki, R. Lopez de Mantaras, S. Matwin, D. Mladenič, A. Skowron (Eds.), *Machine Learning: ECML 2007*. XXII, 809 pages. 2007.
- Vol. 4696: H.-D. Burkhard, G. Lindemann, R. Verbrugge, L.Z. Varga (Eds.), *Multi-Agent Systems and Applications V*. XIII, 350 pages. 2007.
- Vol. 4694: B. Apolloni, R.J. Howlett, L. Jain (Eds.), *Knowledge-Based Intelligent Information and Engineering Systems, Part III*. XXIX, 1126 pages. 2007.
- Vol. 4693: B. Apolloni, R.J. Howlett, L. Jain (Eds.), *Knowledge-Based Intelligent Information and Engineering Systems, Part II*. XXXII, 1380 pages. 2007.
- Vol. 4692: B. Apolloni, R.J. Howlett, L. Jain (Eds.), *Knowledge-Based Intelligent Information and Engineering Systems, Part I*. LV, 882 pages. 2007.
- Vol. 4687: P. Petta, J.P. Müller, M. Klusch, M. Georgeff (Eds.), *Multiagent System Technologies*. X, 207 pages. 2007.
- Vol. 4682: D.-S. Huang, L. Heutte, M. Loog (Eds.), *Advanced Intelligent Computing Theories and Applications*. XXVII, 1373 pages. 2007.
- Vol. 4676: M. Klusch, K.V. Hindriks, M.P. Papazoglou, L. Sterling (Eds.), *Cooperative Information Agents XI*. XI, 361 pages. 2007.
- Vol. 4667: J. Hertzberg, M. Beetz, R. Englert (Eds.), *KI 2007: Advances in Artificial Intelligence*. IX, 516 pages. 2007.
- Vol. 4660: S. Džeroski, L. Todorovski (Eds.), *Computational Discovery of Scientific Knowledge*. X, 327 pages. 2007.
- Vol. 4659: V. Mařík, V. Vyatkin, A.W. Colombo (Eds.), *Holonic and Multi-Agent Systems for Manufacturing*. VIII, 456 pages. 2007.
- Vol. 4651: F. Azevedo, P. Barahona, F. Fages, F. Rossi (Eds.), *Recent Advances in Constraints*. VIII, 185 pages. 2007.
- Vol. 4648: F. Almeida e Costa, L.M. Rocha, E. Costa, I. Harvey, A. Coutinho (Eds.), *Advances in Artificial Life*. XVIII, 1215 pages. 2007.
- Vol. 4635: B. Kokinov, D.C. Richardson, T.R. Roth-Berghofer, L. Vieu (Eds.), *Modeling and Using Context*. XIV, 574 pages. 2007.
- Vol. 4632: R. Alhajj, H. Gao, X. Li, J. Li, O.R. Zaiane (Eds.), *Advanced Data Mining and Applications*. XV, 634 pages. 2007.
- Vol. 4629: V. Matoušek, P. Mautner (Eds.), *Text, Speech and Dialogue*. XVII, 663 pages. 2007.
- Vol. 4626: R.O. Weber, M.M. Richter (Eds.), *Case-Based Reasoning Research and Development*. XIII, 534 pages. 2007.
- Vol. 4617: V. Torra, Y. Narukawa, Y. Yoshida (Eds.), *Modeling Decisions for Artificial Intelligence*. XII, 502 pages. 2007.
- Vol. 4612: I. Miguel, W. Ruml (Eds.), *Abstraction, Reformulation, and Approximation*. XI, 418 pages. 2007.
- Vol. 4604: U. Priss, S. Polovina, R. Hill (Eds.), *Conceptual Structures: Knowledge Architectures for Smart Applications*. XII, 514 pages. 2007.
- Vol. 4603: F. Pfenning (Ed.), *Automated Deduction – CADE-21*. XII, 522 pages. 2007.
- Vol. 4597: P. Perner (Ed.), *Advances in Data Mining*. XI, 353 pages. 2007.
- Vol. 4594: R. Bellazzi, A. Abu-Hanna, J. Hunter (Eds.), *Artificial Intelligence in Medicine*. XVI, 509 pages. 2007.

- Vol. 4585: M. Kryszkiewicz, J.F. Peters, H. Rybinski, A. Skowron (Eds.), *Rough Sets and Intelligent Systems Paradigms*. XIX, 836 pages. 2007.
- Vol. 4578: F. Masulli, S. Mitra, G. Pasi (Eds.), *Applications of Fuzzy Sets Theory*. XVIII, 693 pages. 2007.
- Vol. 4573: M. Kauers, M. Kerber, R. Miner, W. Windsteiger (Eds.), *Towards Mechanized Mathematical Assistants*. XIII, 407 pages. 2007.
- Vol. 4571: P. Perner (Ed.), *Machine Learning and Data Mining in Pattern Recognition*. XIV, 913 pages. 2007.
- Vol. 4570: H.G. Okuno, M. Ali (Eds.), *New Trends in Applied Artificial Intelligence*. XXI, 1194 pages. 2007.
- Vol. 4565: D.D. Schmorow, L.M. Reeves (Eds.), *Foundations of Augmented Cognition*. XIX, 450 pages. 2007.
- Vol. 4562: D. Harris (Ed.), *Engineering Psychology and Cognitive Ergonomics*. XXIII, 879 pages. 2007.
- Vol. 4548: N. Olivetti (Ed.), *Automated Reasoning with Analytic Tableaux and Related Methods*. X, 245 pages. 2007.
- Vol. 4539: N.H. Bshouty, C. Gentile (Eds.), *Learning Theory*. XII, 634 pages. 2007.
- Vol. 4529: P. Melin, O. Castillo, L.T. Aguilar, J. Kacprzyk, W. Pedrycz (Eds.), *Foundations of Fuzzy Logic and Soft Computing*. XIX, 830 pages. 2007.
- Vol. 4520: M.V. Butz, O. Sigaud, G. Pezzulo, G. Baldassarre (Eds.), *Anticipatory Behavior in Adaptive Learning Systems*. X, 379 pages. 2007.
- Vol. 4511: C. Conati, K. McCoy, G. Paliouras (Eds.), *User Modeling 2007*. XVI, 487 pages. 2007.
- Vol. 4509: Z. Kobti, D. Wu (Eds.), *Advances in Artificial Intelligence*. XII, 552 pages. 2007.
- Vol. 4496: N.T. Nguyen, A. Grzech, R.J. Howlett, L.C. Jain (Eds.), *Agent and Multi-Agent Systems: Technologies and Applications*. XXI, 1046 pages. 2007.
- Vol. 4483: C. Baral, G. Brewka, J. Schlipf (Eds.), *Logic Programming and Nonmonotonic Reasoning*. IX, 327 pages. 2007.
- Vol. 4482: A. An, J. Stefanowski, S. Ramanna, C.J. Butz, W. Pedrycz, G. Wang (Eds.), *Rough Sets, Fuzzy Sets, Data Mining and Granular Computing*. XIV, 585 pages. 2007.
- Vol. 4481: J. Yao, P. Lingras, W.-Z. Wu, M. Szczuka, N.J. Cercone, D. Ślęzak (Eds.), *Rough Sets and Knowledge Technology*. XIV, 576 pages. 2007.
- Vol. 4476: V. Gorodetsky, C. Zhang, V.A. Skormin, L. Cao (Eds.), *Autonomous Intelligent Systems: Multi-Agents and Data Mining*. XIII, 323 pages. 2007.
- Vol. 4457: G.M.P. O'Hare, A. Ricci, M.J. O'Grady, O. Dikenelli (Eds.), *Engineering Societies in the Agents World VII*. XI, 401 pages. 2007.
- Vol. 4456: Y. Wang, Y.-m. Cheung, H. Liu (Eds.), *Computational Intelligence and Security*. XXIII, 1118 pages. 2007.
- Vol. 4455: S. Muggleton, R. Otero, A. Tamaddon-Nezhad (Eds.), *Inductive Logic Programming*. XII, 456 pages. 2007.
- Vol. 4452: M. Fasli, O. Shehory (Eds.), *Agent-Mediated Electronic Commerce*. VIII, 249 pages. 2007.
- Vol. 4451: T.S. Huang, A. Nijholt, M. Pantic, A. Pentland (Eds.), *Artificial Intelligence for Human Computing*. XVI, 359 pages. 2007.
- Vol. 4441: C. Müller (Ed.), *Speaker Classification*. X, 309 pages. 2007.
- Vol. 4438: L. Maicher, A. Sigel, L.M. Garshol (Eds.), *Leveraging the Semantics of Topic Maps*. X, 257 pages. 2007.
- Vol. 4434: G. Lakemeyer, E. Sklar, D.G. Sorrenti, T. Takahashi (Eds.), *RoboCup 2006: Robot Soccer World Cup X*. XIII, 566 pages. 2007.
- Vol. 4429: R. Lu, J.H. Siekmann, C. Ullrich (Eds.), *Cognitive Systems*. X, 161 pages. 2007.
- Vol. 4428: S. Edelkamp, A. Lomuscio (Eds.), *Model Checking and Artificial Intelligence*. IX, 185 pages. 2007.
- Vol. 4426: Z.-H. Zhou, H. Li, Q. Yang (Eds.), *Advances in Knowledge Discovery and Data Mining*. XXV, 1161 pages. 2007.
- Vol. 4411: R.H. Bordini, M. Dastani, J. Dix, A.E.F. Seghrouchni (Eds.), *Programming Multi-Agent Systems*. XIV, 249 pages. 2007.
- Vol. 4410: A. Branco (Ed.), *Anaphora: Analysis, Algorithms and Applications*. X, 191 pages. 2007.
- Vol. 4399: T. Kovacs, X. Llorà, K. Takadama, P.L. Lanzi, W. Stolzmann, S.W. Wilson (Eds.), *Learning Classifier Systems*. XII, 345 pages. 2007.
- Vol. 4390: S.O. Kuznetsov, S. Schmidt (Eds.), *Formal Concept Analysis*. X, 329 pages. 2007.
- Vol. 4389: D. Weyns, H. Van Dyke Parunak, F. Michel (Eds.), *Environments for Multi-Agent Systems III*. X, 273 pages. 2007.
- Vol. 4386: P. Noriega, J. Vázquez-Salceda, G. Boella, O. Boissier, V. Dignum, N. Fornara, E. Matson (Eds.), *Coordination, Organizations, Institutions, and Norms in Agent Systems II*. XI, 373 pages. 2007.
- Vol. 4384: T. Washio, K. Satoh, H. Takeda, A. Inokuchi (Eds.), *New Frontiers in Artificial Intelligence*. IX, 401 pages. 2007.
- Vol. 4371: K. Inoue, K. Satoh, F. Toni (Eds.), *Computational Logic in Multi-Agent Systems*. X, 315 pages. 2007.
- Vol. 4369: M. Umeda, A. Wolf, O. Bartenstein, U. Geske, D. Seipel, O. Takata (Eds.), *Declarative Programming for Knowledge Management*. X, 229 pages. 2006.
- Vol. 4363: B.D. ten Cate, H.W. Zeevat (Eds.), *Logic, Language, and Computation*. XII, 281 pages. 2007.
- Vol. 4343: C. Müller (Ed.), *Speaker Classification I*. X, 355 pages. 2007.
- Vol. 4342: H. de Swart, E. Orlowska, G. Schmidt, M. Roubens (Eds.), *Theory and Applications of Relational Structures as Knowledge Instruments II*. X, 373 pages. 2006.
- Vol. 4335: S.A. Brueckner, S. Hassas, M. Jelasity, D. Yamins (Eds.), *Engineering Self-Organising Systems*. XII, 212 pages. 2007.
- Vol. 4334: B. Beckert, R. Hähnle, P.H. Schmitt (Eds.), *Verification of Object-Oriented Software*. XXIX, 658 pages. 2007.

Table of Contents

Argumentation in Multi-Agent Systems: Context and Recent Developments	1
<i>Nicolas Maudet, Simon Parsons, and Iyad Rahwan</i>	
Part I: Foundations and Explorations	
Argumentation-Based Learning	17
<i>Taro Fukumoto and Hajime Sawamura</i>	
Arguments and Counterexamples in Case-Based Joint Deliberation	36
<i>Santiago Ontañón and Enric Plaza</i>	
A Generalization of Dung's Abstract Framework for Argumentation: Arguing with Sets of Attacking Arguments	54
<i>Søren Holbech Nielsen and Simon Parsons</i>	
An Argumentation-Based Approach for Practical Reasoning	74
<i>Iyad Rahwan and Leila Amgoud</i>	
Support-Based Distributed Search: A New Approach for Multiagent Constraint Processing	91
<i>Peter Harvey, Chee Fon Chang, and Aditya Ghose</i>	
Managing Social Influences Through Argumentation-Based Negotiation	107
<i>Nishan C. Karunatillake, Nicholas R. Jennings, Iyad Rahwan, and Sarvapali D. Ramchurn</i>	
Part II: Strategic Issues	
An Argumentation-Based Approach for Dialog Move Selection	128
<i>Leila Amgoud and Nabil Hameurlain</i>	
Specification and Complexity of Strategic-Based Reasoning Using Argumentation	142
<i>Mohamed Mbarki, Jamal Bentahar, and Bernard Moulin</i>	
Information Based Argumentation Heuristics	161
<i>Nir Oren, Timothy J. Norman, and Alun Preece</i>	
Negotiating Using Rewards	175
<i>Sarvapali D. Ramchurn, Carles Sierra, Lluís Godo, and Nicholas R. Jennings</i>	

Argumentation and Persuasion in the Cognitive Coherence Theory:
Preliminary Report 193
 Philippe Pasquier, Iyad Rahwan, Frank Dignum, and Liz Sonenberg

Author Index 211

Argumentation in Multi-Agent Systems: Context and Recent Developments

Nicolas Maudet¹, Simon Parsons², and Iyad Rahwan³

¹ LAMSADE, Université Paris-Dauphine
Paris 75775 Cedex 16, France
maudet@lamsade.dauphine.fr

² Department of Computer and Information Science, Brooklyn College
City University of New York, 2900 Bedford Avenue, Brooklyn, 11210 NY, USA
parsons@sci.brooklyn.cuny.edu

³ Institute of Informatics, The British University in Dubai
P.O.Box 502216, Dubai, UAE
(Fellow) School of Informatics, University of Edinburgh, UK
irahwan@acm.org

Abstract. This chapter provides a brief survey of argumentation in multi-agent systems. It is not only brief, but rather idiosyncratic, and focuses on the areas of research that most interest the authors, and those which seem to be the most active at the time of writing.

1 Introduction

The theory of argumentation [81] is a rich, interdisciplinary area of research lying across philosophy, communication studies, linguistics, and psychology. Its techniques and results have found a wide range of applications in both theoretical and practical branches of artificial intelligence and computer science [14,74]. These applications range from specifying semantics for logic programs [20], to natural language text generation [21], to supporting legal reasoning [9], to decision-support for multi-party human decision-making [31] and conflict resolution [80].

In recent years, argumentation theory has been gaining increasing interest in the multi-agent systems (MAS) research community. On one hand, argumentation-based techniques can be used to specify *autonomous agent reasoning*, such as belief revision and decision-making under uncertainty and non-standard preference policies. On the other hand, argumentation can also be used as a vehicle for facilitating *multi-agent interaction*, because argumentation naturally provides tools for designing, implementing and analysing sophisticated forms of interaction among rational agents. Argumentation has made solid contributions to the theory and practice of multi-agent dialogues.

In this short survey, we review the most significant and recent advances in the field, with no intention of being exhaustive. Thus, we ignore recent work that extends the basic mechanisms of argumentation with new semantics [12], bipolar arguments [13], and the ability to handle sets of arguments [49]. Indeed, we have

very little to say about *how to argue* and, instead, deal with *what one can argue about*, dealing with the uses of argumentation rather than the mechanisms by which it may be carried out¹, and restricting even that view to coincide with the topics of the other papers in this volume. In particular, this chapter first recalls some of the key notions in argumentation theory, and then outlines work on two major applications of argumentation in multi-agent systems, namely in the reasoning carried out by autonomous agents (Section 3) and in multi-agent communication (Section 4).

2 What Is Argumentation Good for?

According to a recent authoritative reference on argumentation theory, argumentation can be defined as follows:

Argumentation is a verbal and social activity of reason aimed at increasing (or decreasing) the acceptability of a controversial standpoint for the listener or reader, by putting forward a constellation of propositions intended to justify (or refute) the standpoint before a rational judge. [81, page 5]

Let us decompose the elements of this definition that are most relevant to our discussion. First, the ultimate goal of argumentation is to resolve a “*controversial*” standpoint; controversial in the sense that it is subject to both “*justification*” or “*refutation*” depending on the information available. This distinguishes argumentation from the classical deductive reasoning viewpoint, in which proofs for propositions cannot be contested. Moreover, the nature of the “*standpoint*” can vary. While the classical study of argumentation has focused mainly on propositional standpoints — i.e. things that are believed or known — there is no reason why the standpoint is confined to be propositional. A standpoint can, in principle, range from a proposition to believe, to a goal to try to achieve, to a value to try to promote. That is, argumentation can be used for theoretical reasoning (about what to believe) as well as practical reasoning (about what to do).

Secondly, argumentation is an “*activity of reason*”, emphasising that a particular process is to be followed in order to influence the acceptability of the controversial standpoint. This activity and the propositions put forward are to be evaluated by a “*rational judge*”: a system that defines the reasonableness of these propositions according to some criteria. An important objective of argumentation theory is to identify such system of criteria.

In summary, argumentation can be seen as the principled interaction of different, potentially conflicting arguments, for the sake of arriving at a consistent conclusion. Perhaps the most crucial aspect of argumentation is the interaction between arguments. Argumentation can give us means for allowing an agent to

¹ Not least because one can potentially make use of any mechanism for argumentation in the service of any of the applications of argumentation.

reconcile conflicting information within itself, for reconciling its informational state with new perceptions from the environment, and for reconciling conflicting information between multiple agents through communication. It is for these reasons that argumentation has begun to receive great interest in the multi-agent systems community. In particular, argumentation lends itself naturally to two main sorts of problems encountered in MAS:

- **Forming and revising beliefs and decisions:** Argumentation provides means for forming beliefs and decisions on the basis of incomplete, conflicting or uncertain information. This is because argumentation provides a systematic means for resolving conflicts among different arguments and arriving at consistent, well-supported standpoints;
- **Rational interaction:** Argumentation provides means for structuring dialogue between participants that have potentially conflicting viewpoints. In particular, argumentation provides a framework for ensuring that interaction respects certain principles (e.g. consistency of each participant’s statements).

In the next sections, we will discuss these applications in more detail and refer to some relevant literature. In particular, Section 3 deals with the topics of revising beliefs and making decisions, aspects that we can think of as being the concern of individual autonomous agents, while Section 4 deals with topics related to inter-agent communication and rational action, all aspects of argumentation that are decidedly multi-agent.

3 Argumentation for Reasoning in Autonomous Agents

Argumentation is a general process for reasoning. An autonomous agent that has to reason about could weigh arguments for and against different options in order to arrive at a well-supported stance. In this section, we discuss two main applications of argumentation to autonomous agent reasoning.

3.1 Argumentation for Belief Revision

One of the main challenges in specifying autonomous agents is the maintenance and updating of its beliefs in a dynamic environment. An agent may receive perceptual information that is inconsistent with its view of the world, in which case the agent needs to update its beliefs in order to maintain consistency. The major challenge of nonmonotonic reasoning formalisms [11] is to specify efficient ways to update beliefs. At the normative level, the AGM paradigm [29] specifies the rationality postulates that must be satisfied by an idealistic process of belief revision. On the operational level, formalisms for mechanising nonmonotonic reasoning include truth maintenance systems (TMS) [19], default logic [75] and circumscription [48].

Argumentation provides an alternative way to mechanise nonmonotonic reasoning. Argument-based frameworks view the problem of nonmonotonic reasoning as a process in which arguments for and against certain conclusions

are constructed and compared. Nonmonotonicity arises from the fact that new premises may enable the construction of new arguments to support new beliefs, or stronger counterarguments against existing beliefs. As the number of premises grows, the set of arguments that can be constructed from those premises grows monotonically. However, because new arguments may overturn existing beliefs, the set of beliefs is nonmonotonic. Various argument-based frameworks for non-monotonic reasoning have been proposed in the last 20 or so years. Some of the most notable are the following [42,60,79,41,22,27,67]².

While the above-mentioned frameworks have developed into a solid and mature sub-field of AI, their incorporation into *situated* autonomous agent reasoning remains an opportunity to be pursued. In order to do so, an adequate representation of the environment is needed, and a mechanism for integrating perceptual information into the belief-update mechanism is also required. Moreover, situated agents are required to update their beliefs in a timely fashion in order to take appropriate action accordingly.

3.2 Argumentation for Deliberation and Means-Ends Reasoning

An autonomous agent does not only maintain a mental picture of its environment. The agent is faced with two additional tasks: the task of *deliberation* in which it decides what state of the world it wishes to achieve — namely its goal — and the task of *means-ends reasoning* in which it forms a plan to achieve this goal. Argumentation is also potentially useful for tackling both these challenges.

Recently, argumentation has been applied to deliberation. For example, argumentation has been used as a means for choosing among a set of conflicting desires [1] and as a means for choosing between goals [3]. Another argument-based framework for deliberation has been presented by Kakas and Moraitis [39]. In this approach, arguments and preferences among them are used in order to generate goals based on a changing context. In addition, argumentation can be used to support standard BDI [73] models, as in [56].

More generally, as shown by Fox in his work since [26]³, argumentation provides a framework for making decisions. Just as one makes arguments and counterarguments for beliefs, one can make arguments and counterarguments for actions. While such a framework sounds as though it must be at odds with approaches based on decision theory [34], Fox and Parsons [28] provide an argumentation framework that reconciles the two approaches. In this system, argumentation is used to reason about the expected value of possible actions. In particular, one argument system is used to arrive at a stance on beliefs, while another argument system identifies the *outcomes* of possible actions. Together, arguments over beliefs and the results of actions can be combined to create arguments about the *expected value* of possible actions. This approach was later refined in [53].

² For comprehensive surveys on argument-based approaches to nonmonotonic reasoning, see [14,68].

³ Though this line of work, summarised in [52], did not explicitly use the term “argumentation” until [27], with hindsight it is clear that argumentation is exactly what Fox and his colleagues were using.