

Robert Babuška
Frans C.A. Groen (Eds.)

Interactive Collaborative Information Systems

Interactive Collaborative Information Systems

Dr. Robert Babuška
Delft Center for Systems and Control
Delft University of Technology
Mekelweg 2
2628 CD Delft
The Netherlands
E-mail: r.babuska@tudelft.nl

Dr. Frans C.A. Groen
Faculty of Science
Informatics Institute
Science Park 107
1098 XG, Amsterdam
The Netherlands
E-mail: F.C.A.Groen@uva.nl

ISBN 978-3-642-11687-2

e-ISBN 978-3-642-11688-9

DOI 10.1007/978-3-642-11688-9

Studies in Computational Intelligence

ISSN 1860-949X

Library of Congress Control Number: 2010920828

© 2010 Springer-Verlag Berlin Heidelberg

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, reuse of illustrations, recitation, broadcasting, reproduction on microfilm 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.

The use of general descriptive names, registered names, trademarks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

Typeset & Cover Design: Scientific Publishing Services Pvt. Ltd., Chennai, India.

Printed on acid-free paper

9 8 7 6 5 4 3 2 1

springer.com

Robert Babuška and Frans C.A. Groen (Eds.)

Interactive Collaborative Information Systems

Preface

The increasing complexity of our world demands new perspectives on the role of technology in decision making. Human decision making has its limitations in terms of information-processing capacity. We need new technology to cope with the increasingly complex and information-rich nature of our modern society. This is particularly true for critical environments such as crisis management and traffic management, where humans need to engage in close collaborations with artificial systems to observe and understand the situation and respond in a sensible way. We believe that close collaborations between humans and artificial systems will become essential and that the importance of research into Interactive Collaborative Information Systems (ICIS) is self-evident.

Developments in information and communication technology have radically changed our working environments. The vast amount of information available nowadays and the wirelessly networked nature of our modern society open up new opportunities to handle difficult decision-making situations such as computer-supported situation assessment and distributed decision making. To make good use of these new possibilities, we need to update our traditional views on the role and capabilities of information systems.

The aim of the *Interactive Collaborative Information Systems* project is to develop techniques that support humans in complex information environments and that facilitate distributed decision-making capabilities. ICIS emphasizes the importance of building actor-agent communities: close collaborations between human and artificial actors that highlight their complementary capabilities, and in which task distribution is flexible and adaptive. To fulfill such a prospect, we need intelligent systems that observe their environment, interpret and fuse information, have learning and decision-making capabilities, and have the ability to work in teams. It also means that we need to study the interaction of humans with their artificial counterparts in such settings and how their information needs can be met. Research within the ICIS projects helps create such views. ICIS combines research from information technology, artificial intelligence and human sciences to obtain a

multidisciplinary foundation from which innovative actor-agent systems for critical environments can emerge.

This book focuses on the employment of innovative agent technology, advanced machine learning techniques, and cognition-based interface technology for the use in collaborative decision support systems. It consists of five parts: *Reinforcement learning*, *Collaborative decision making*, *Computer-human interaction modeling*, *Architectures for distributed agent-actor communities*, and *Case studies and applications*.

Reinforcement Learning is the main subject of the first part of the book. This type of learning plays an important role in developing intelligent systems. As the systems (agents) have to realize a common goal, the question is which actions an agent should take in its environment to contribute to that common goal. Reinforcement learning is an established way to implement this process. The core research questions within this subject are the representation of the environment's state, the representation of the actions that an agent can take in a given state, and the long-term reward representing the common goal. Finding the optimal sequence of actions often becomes intractable, especially when the number of agents increases. As a result, approximate solutions are needed. These are discussed in the chapter "*Approximate dynamic programming and reinforcement learning*." This chapter focuses on an approximate solution for reinforcement learning given a representation with continuous states. The next chapter "*Learning with whom to communicate using relational reinforcement learning*" exploits relational structures to come up with strategies for multi-agent systems. The last chapter "*Switching between representations in reinforcement learning*" investigates when to switch online between feature sets representing the states.

The second part of the book addresses **Collaborative Decision Making**. Decision-theoretic models are given to describe cooperation between multiple agents under uncertainty. In these cases the state of the environment given the agent's observations is uncertain and can be described by, for instance, a probability distribution over the states. Approximate solution methods for these cases are presented in the chapter "*A decision-theoretic approach to collaboration: principal description methods and efficient heuristic approximations*". The next chapter "Efficient Methods for Near-Optimal Sequential Decision Making Under Uncertainty" discusses both Bayesian and distribution-free algorithms for sequential decision making when costs are known. A completely different approach is presented in the chapter on "*Ant colony learning algorithm for optimal control*". In this case an optimization heuristic is used and a novel algorithm is introduced in which the artificial agents (ants) work together to collectively learn optimal control policies. The chapter on "*Map-based support for effective collaboration in micro-mobile virtual teams*" presents collaboration in geo-spatial support systems where the maps aid the distributed decision-making process.

The topic of the third part of the book is **Computer-Human Interaction Modeling**. The first chapter of this part “*Affective dialogue management using factored POMDPs*” shows that partially observable Markov decision processes (POMDPs) are appropriate for this purpose and presents a novel approach to develop an affective dialogue model. The next chapter, “*Context-aware multimodal human computer interaction*,” presents multimodal interaction techniques including speech, lip movement, facial expression, and text and visual communication. The chapter “*Design issues for pen-centric interactive maps*” focuses on pen-input recognition systems, in particular on new features for classifying iconic gestures. The response of users to intelligent systems showing adaptive behavior is studied in the chapter “*Interacting with adaptive systems*”. Human poses are important for both interaction and situational awareness in human-inhabited environments. The last chapter, “*Example-based human pose recovery under predicted partial occlusions*,” deals with that topic also in the case when partial occlusion occurs.

The architecture of intelligent decision-making system is an important glue that lets all the parts work together. This is the topic of the fourth part of the book, **Architectures for Distributed Agent-Actor Communities**. Chapter “*Agility and adaptive autonomy in networked organizations*” addresses the tradeoff in a multi-actor environment between global coordination of activities and respecting the autonomy of the actors involved. The underlying principles of multi-agents organizations that are not only hierarchical, but that can also adapt their structure are the topic of the chapter “*Adaptive hierarchical multi-agent organizations*”. In the chapter “*Method for designing networking adaptive interactive hybrid systems*” the various architectures for this type of systems are given and a top-down design methodology is introduced. The tradeoff in a multi-actor environment between global coordination of activities and respecting the autonomy of the actors involved forms the topic studied in the chapter “*Agility and adaptive autonomy in networked organizations*”.

Case Studies and Applications are discussed in the last part of the book. Crisis management is the topic of the first two chapters. In the first chapter, “*A call for sensemaking support systems in crisis management*”, the information challenges in crisis management are explored and three case studies are investigated. In the next chapter, “*A distributed approach to gas detection and source localization using heterogeneous information*”, a system for early detection of gaseous substances and coarse source estimation is presented by using heterogeneous sensor measurements and human reports. Mobility is the application discussed in the next two chapters. The chapter “*Traffic light control by multiagent reinforcement learning systems*” discusses extensions to improve a basic setting of multiple local controllers (agents), each responsible for the optimization of traffic lights around a single junction using reinforcement learning. The chapter “*Fusing heterogeneous and unreliable data from traffic sensors*” deals with traffic data fusion from a variety of traffic sensors using conservation laws and Poisson statistics. The last

chapter, “*Bayesian networks for expert systems, theory and practical applications*”, shows the strength of Bayesian modeling approaches in three different applications: medical diagnosis support, petrochemical decision support and victim identification.

Acknowledgement. The work described in this book is part of the ICIS project – Interactive Collaborative Information Systems, see <http://www.icis.decis.nl/>. This project is one of nine ICT research projects funded by the BSIK program (SenterNovem, grant no. BSIK03024) of the Dutch government. The ICIS project started in the fall of 2003 and finished in 2010. The ICIS project is hosted by the D-CIS Lab, an open research partnership of Thales Nederland, the Delft University of Technology, the University of Amsterdam and the Netherlands Foundation of Applied Scientific Research (TNO). The ICIS consortium involves leading IT-driven industries, academic research institutions, technology institutes and high-tech small and medium enterprises. The financial support of SenterNovem is gratefully acknowledged.

Amsterdam, Delft,
November 2009

Frans C.A. Groen
Robert Babuška

List of Contributors

Jelmer Marinus van Ast

Delft Center for Systems and
Control, Delft University of
Technology, Mekelweg 2,
2628 CD Delft,
The Netherlands
j.m.vanast@tudelft.nl

Robert Babuška

Delft Center for Systems and
Control, Delft University of
Technology, Mekelweg 2, 2628 CD
Delft, The Netherlands
r.babuska@tudelft.nl

Bram Bakker

Informatics Institute, University
of Amsterdam, Science Park 107,
1098 XG Amsterdam,
The Netherlands
p.b.bakker@uva.nl

Guido te Brake

TNO Defense, Security and
Safety, PO Box 21, 3769 ZG,
Soesterberg, The Netherlands
guido.tebrake@tno.nl

Lucian Buşoniu

Delft Center for Systems and
Control, Delft University
of Technology, Mekelweg 2,
2628 CD Delft, The Netherlands
i.l.busoniutudelft.nl

Trung H. Bui

Center for the Study of Language
and Information, Stanford
University, 210 Panama St, Stanford,
CA 94305, USA
thbui@stanford.edu

Willem Burgers

SNN Adaptive Intelligence, Geert
Grooteplein 21, 6525 EZ Nijmegen,
The Netherlands
w.burgers@science.ru.nl

Alin G. Chîtu

Man Machine Interaction, Delft
University of Technology, Mekel-
weg 4 2628 CD, The Netherlands
a.g.chitutudelft.nl

Henriette Cramer

Human-Computer Studies,
University of Amsterdam,
Science Park 107, 1098 XG
Amsterdam, The Netherlands
hrcramer@science.uva.nl

Tom Croonenborghs

KH Kempen University College,
Kleinhoefstraat 4, 2440 Geel,
Belgium
tom.croonenborghs@khk.be

Dragoş Datcu

Man Machine Interaction, Delft
University of Technology, Mekelweg
4 2628 CD, The Netherlands
d.datcu@tudelft.nl

Bart De Schutter

Delft Center for Systems and
Control & Marine and Transport
Technology Department, Delft
University of Technology, Mekelweg
2, 2628 CD Delft, The Netherlands
b@deschutter.info

Christos Dimitrakakis

Informatics Institute, University
of Amsterdam, Science Park 107,
1098 XG Amsterdam,
The Netherlands
christos.dimitrakakis@gmail.com

Kurt Driessens

DTAI, Katholieke
Universiteit Leuven,
Celestijnenlaan 200A,
B-3001 Heverlee, Belgium
kurt.driessens@cs.kuleuven.be

Vanessa Evers

Human-Computer Studies,
University of Amsterdam,
Science Park 107, 1098 XG
Amsterdam, The Netherlands
v.evers@uva.nl

Siska Fitrianie

Man Machine Interaction, Delft
University of Technology, Mekelweg
4, 2628 CD, The Netherlands
s.fitrianie@tudelft.nl

Marcel van Gerve

Institute for Computing &
Information Sciences, Radboud
University Nijmegen,
The Netherlands
marcelge@cs.ru.nl

Mattijs Ghijsen

Human Computer Studies Group,
Informatics Institute, Universiteit
van Amsterdam, Science Park 107,
1098 XG Amsterdam,
The Netherlands
m.ghijsen@uva.nl

Rianne Gouman

D-CIS Lab / Thales Research &
Technology Netherlands, P.O. Box
90, 2600 AB, Delft, The Netherlands
rianne.gouman@icis.decis.nl

Frans C.A. Groen

Informatics Institute, University
of Amsterdam, Science Park 107,
1098 XG Amsterdam,
The Netherlands
f.c.a.groen@uva.nl

Jaap van den Herik

Tilburg centre for Creative
Computing, Tilburg University,
Warandelaan 2, PO Box 90153,
5000 LE Tilburg, The Netherlands
h.j.vdnherik@uvt.nl

Prof. Dr. S.P. Hoogendoorn

Transport & Planning, Faculty of
Civil Engineering and Geosciences,
Delft University of Technology
s.p.hoogendoorn@tudelft.nl

Wouter N.H. Jansweijer

Human Computer Studies Group,
Informatics Institute, Universiteit
van Amsterdam, Science Park 107,
1098 XG Amsterdam,
The Netherlands
w.n.h.jansweijer@uva.nl

Michiel Kamermans

Thales Research & Technology
Netherlands, Delftechpark 24,
2628 XH, Delft, The Netherlands
Michiel.Kamermans@icis.decis.nl

Marten Kampman

University of Utrecht,
Information & Computing Sciences
martenkampman@gmail.com

Bert Kappen

Radboud University Nijmegen,
Donders Institute for Brain,
Cognition and Behaviour, Geert
Grooteplein 21, 6525 EZ Nijmegen,
The Netherlands
b.kappen@science.ru.nl

Masja Kempen

D-CIS Lab / Thales Research &
Technology Netherlands, P.O. Box
90, 2600 AB, Delft, The Netherlands
masja.kempen@icis.decis.nl

Leon Kester

TNO Defence, Security and Safety,
Oude Waalsdorperweg 63, 2597 AK
The Hague, The Netherlands
leon.kester@tno.nl

Rick van der Kleij

TNO Defense, Security and Safety,
PO Box 21, 3769 ZG, Soesterberg,
The Netherlands
rick.vanderkleij@tno.nl

Dr. J.W.C. Van Lint

Transport & Planning, Faculty of
Civil Engineering and Geosciences,
Delft University of Technology
j.w.c.vanlint@tudelft.nl

Willem J. Muhren

Department of Information
Systems and Management, Tilburg
University, PO Box 90153, 5000 LE
Tilburg, The Netherlands
w.j.muhren@uvt.nl

Martijn Neef

TNO Defense, Security and Safety,
Oude Waalsdorperweg 63, 2597 AK
The Hague, The Netherlands
martijn.neef@tno.nl

Ralph Niels

Donders Institute for Brain,
Cognition & Behavior, Radboud
University Nijmegen,
The Netherlands
r.niels@donders.ru.nl

Anton Nijholt

Human Media Interaction Group,
University of Twente, PO Box 217,
7500 AE Enschede, The Netherlands
anijholt@cs.utwente.nl

Frans A. Oliehoek

Intelligent System Laboratory
Amsterdam, Science Park 107,
1098 XG Amsterdam,
The Netherlands
F.A.Oliehoek@uva.nl

Q. Ou

Transport & Planning, Faculty of
Civil Engineering and Geosciences,
Delft University of Technology
Q.Ou@tudelft.nl

Patrick de Oude

Intelligent Autonomous Systems
Group, Informatics Institute, Faculty
of Science, University of
Amsterdam, Science Park 107,
1098 XG Amsterdam,
The Netherlands
P.deOude@uva.nl

Gregor Pavlin

Thales Research & Technology
Netherlands, Delftechpark 24,
2628 XH, Delft, The Netherlands
Gregor.Pavlin@icis.decis.nl

Mannes Poel

Human Media Interaction Group,
University of Twente, PO Box 217,
7500 AE Enschede, The Netherlands
mpoel@cs.utwente.nl

Marc Ponsen

Department of Knowledge
Engineering, Maastricht University,
Mindebroedersberg 6a, 6211 LK,
Maastricht, The Netherlands
m.ponsen@maastrichtuniversity.nl

Ronald Poppe

Human Media Interaction Group,
University of Twente, Enschede,
The Netherlands
poppe@ewi.utwente.nl

Eric Postma

Tilburg centre for Creative
Computing, Tilburg University,
Warandelaan 2, P.O. Box 90153,
5000 LE Tilburg, The Netherlands
e.o.postma@uvt.nl

Jan Ramon

DTAI, Katholieke Universiteit
Leuven, Celestijnenlaan 200A,
B-3001 Heverlee, Belgium
jan.ramon@cs.kuleuven.be

Leon J.M. Rothkrantz

Man Machine Interaction, Delft
University of Technology, Mekelweg
4 2628 CD, The Netherlands and
Netherlands Defence Academy,
Faculty of Technical Sciences, Den
Helder, The Netherlands
l.j.m.rothkrantz@tudelft.nl

Harm van Seijen

TNO Defense, Security and Safety,
Oude Waalsdorperweg 63, 2597 AK
The Hague, The Netherlands
harm.vanseijen@tno.nl

Maarten van Someren

Human-Computer Studies,
University of Amsterdam, Science
Park 107, 1098 XG Amsterdam,
The Netherlands
M.W.vanSomeren@uva.nl

Karl Tuyls

Department of Knowledge
Engineering, Maastricht University,
Mindebroedersberg 6a, 6211 LK,
Maastricht, The Netherlands
k.tuyls@maastrichtuniversity.nl

Bob van der Vecht

TNO Defence, Security and Safety,
Oude Waalsdorperweg 63, 2597 AK
The Hague, The Netherlands
bob.vandervecht@tno.nl

Arnoud Visser

Intelligent System Laboratory
Amsterdam, Science Park 107,
1098 XG Amsterdam,
The Netherlands
A.Visser@uva.nl

Louis Vuurpijl

Donders Institute for Brain,
Cognition & Behavior, Radboud
University Nijmegen,
The Netherlands
l.vuurpijl@donders.ru.nl

Bartel Van de Walle

Department of Information Systems
and Management, Tilburg
University, PO Box 90153, 5000 LE
Tilburg, The Netherlands
bartel@uvt.nl

Shimon Whiteson

Informatics Institute, University
of Amsterdam, Science Park 107,
1098 XG Amsterdam,
The Netherlands
s.a.whiteson@uva.nl

Wim Wiegerinck

SNN Adaptive Intelligence, Geert
Grooteplein 21, 6525 EZ Nijmegen,
The Netherlands
w.wiegerinck@science.ru.nl

Bob J. Wielinga

Human Computer Studies Group,
Informatics Institute, Universiteit
van Amsterdam, Science Park 107,
1098 XG Amsterdam,
The Netherlands
B.J.Wielinga@uva.nl

Niek Wijngaards

D-CIS Lab / Thales Research &
Technology Netherlands,
P.O. Box 90, 2600 AB,
Delft, The Netherlands
niek.wijngaards@icis.decis.nl

Don Willems

Donders Institute for Brain,
Cognition & Behavior,
Radboud University Nijmegen
d.willems@donders.ru.nl

Zhenke Yang

Man Machine Interaction, Delft
University of Technology, Mekelweg
4 2628 CD, The Netherlands
z.yang@tudelft.nl

Job Zwiers

Human Media Interaction Group,
University of Twente, PO Box 217,
7500 AE Enschede, The Netherlands
zwiers@cs.utwente.nl

Studies in Computational Intelligence, Volume 281

Editor-in-Chief

Prof. Janusz Kacprzyk
Systems Research Institute
Polish Academy of Sciences
ul. Newelska 6
01-447 Warsaw
Poland
E-mail: kacprzyk@ibspan.waw.pl

Further volumes of this series can be found on our homepage: springer.com

Vol. 260. Edward Szczerbicki and Ngoc Thanh Nguyen (Eds.)
Smart Information and Knowledge Management, 2009
ISBN 978-3-642-04583-7

Vol. 261. Nadia Nedjah, Leandro dos Santos Coelho, and
Luiza de Macedo de Mourelle (Eds.)
Multi-Objective Swarm Intelligent Systems, 2009
ISBN 978-3-642-05164-7

Vol. 262. Jacek Koronacki, Zbigniew W. Ras,
Slawomir T. Wierzbichon, and Janusz Kacprzyk (Eds.)
Advances in Machine Learning I, 2009
ISBN 978-3-642-05176-0

Vol. 263. Jacek Koronacki, Zbigniew W. Ras,
Slawomir T. Wierzbichon, and Janusz Kacprzyk (Eds.)
Advances in Machine Learning II, 2009
ISBN 978-3-642-05178-4

Vol. 264. Olivier Sigaud and Jan Peters (Eds.)
*From Motor Learning to Interaction
Learning in Robots*, 2009
ISBN 978-3-642-05180-7

Vol. 265. Zbigniew W. Ras and Li-Shiang Tsay (Eds.)
Advances in Intelligent Information Systems, 2009
ISBN 978-3-642-05182-1

Vol. 266. Akitoshi Hanazawa, Tsutomu Miki,
and Keiichi Horio (Eds.)
Brain-Inspired Information Technology, 2009
ISBN 978-3-642-04024-5

Vol. 267. Ivan Zelinka, Sergej Celikovský, Hendrik Richter,
and Guanrong Chen (Eds.)
Evolutionary Algorithms and Chaotic Systems, 2009
ISBN 978-3-642-10706-1

Vol. 268. Johann M.Ph. Schumann and Yan Liu (Eds.)
Applications of Neural Networks in High Assurance Systems,
2009
ISBN 978-3-642-10689-7

Vol. 269. Francisco Fernández de de Vega and
Erick Cantú-Paz (Eds.)
Parallel and Distributed Computational Intelligence, 2009
ISBN 978-3-642-10674-3

Vol. 270. Zong Woo Geem
Recent Advances In Harmony Search Algorithm, 2009
ISBN 978-3-642-04316-1

Vol. 271. Janusz Kacprzyk, Frederick E. Petry, and
Adnan Yazici (Eds.)
*Uncertainty Approaches for Spatial Data Modeling and
Processing*, 2009
ISBN 978-3-642-10662-0

Vol. 272. Carlos A. Coello Coello, Clarisse Dhaenens, and
Laetitia Jourdan (Eds.)
Advances in Multi-Objective Nature Inspired Computing,
2009
ISBN 978-3-642-11217-1

Vol. 273. Fatos Xhafa, Santi Caballé, Ajith Abraham,
Thanasis Daradoumis, and Angel Alejandro Juan Perez
(Eds.)
*Computational Intelligence for Technology Enhanced
Learning*, 2010
ISBN 978-3-642-11223-2

Vol. 274. Zbigniew W. Ras and Alicja Wierzchowska (Eds.)
Advances in Music Information Retrieval, 2010
ISBN 978-3-642-11673-5

Vol. 275. Dilip Kumar Pratihar and Lakhmi C. Jain (Eds.)
Intelligent Autonomous Systems, 2010
ISBN 978-3-642-11675-9

Vol. 276. Jacek Mańdziuk
*Knowledge-Free and Learning-Based Methods in Intelligent
Game Playing*, 2010
ISBN 978-3-642-11677-3

Vol. 277. Filippo Spagnolo and Benedetto Di Paola (Eds.)
*European and Chinese Cognitive Styles and their Impact on
Teaching Mathematics*, 2010
ISBN 978-3-642-11679-7

Vol. 278. Radomir S. Stankovic and Jaakko Astola
From Boolean Logic to Switching Circuits and Automata, 2010
ISBN 978-3-642-11681-0

Vol. 279. Manolis Wallace, Ioannis E. Anagnostopoulos,
Phivos Mylonas, and Maria Bielikova (Eds.)
Semantics in Adaptive and Personalized Services, 2010
ISBN 978-3-642-11683-4

Vol. 280. Chang Wen Chen, Zhu Li, and Shiguo Lian (Eds.)
*Intelligent Multimedia Communication: Techniques and
Applications*, 2010
ISBN 978-3-642-11685-8

Vol. 281. Robert Babuška and Frans C.A. Groen (Eds.)
Interactive Collaborative Information Systems, 2010
ISBN 978-3-642-11687-2

Contents

Part I: Reinforcement Learning

Approximate Dynamic Programming and Reinforcement Learning

Lucian Buşoniu, Bart De Schutter, Robert Babuška

1	Introduction	3
2	Markov Decision Processes: Exact Dynamic Programming and Reinforcement Learning	6
2.1	Markov Decision Processes and Their Solution	6
2.2	Exact Value Iteration	9
2.3	Exact Policy Iteration	10
3	The Need for Approximation in Dynamic Programming and Reinforcement Learning	12
4	Approximate Value Iteration	13
4.1	Approximate Model-Based Value Iteration	13
4.2	Approximate Model-Free Value Iteration	15
4.3	Convergence and the Role of Nonexpansive Approximators	15
5	Approximate Policy Iteration	20
5.1	Approximate Policy Evaluation	20
5.2	Policy Improvement: Approximate Policy Iteration	23
5.3	Theoretical Guarantees	24
5.4	Actor-Critic Algorithms	28
6	Finding Value Function Approximators Automatically	29
6.1	Resolution Refinement	30
6.2	Basis Function Optimization	31
6.3	Other Methods for Basis Function Construction	32

7	Approximate Policy Search	33
8	Comparison of Approximate Value Iteration, Policy Iteration, and Policy Search	36
9	Summary and Outlook	37
	References	39
Learning with Whom to Communicate Using Relational Reinforcement Learning		45
<i>Marc Ponsen, Tom Croonenborghs, Karl Tuyls, Jan Ramon, Kurt Driessens, Jaap van den Herik, Eric Postma</i>		
1	Introduction	46
2	Reinforcement Learning	47
3	Relational Reinforcement Learning	50
4	Multi-agent Relational Reinforcement Learning	51
5	Empirical Evaluation	53
	5.1 Learning Task	54
	5.2 Experimental Results	55
6	Conclusions	60
	References	61
Switching between Representations in Reinforcement Learning		65
<i>Harm van Seijen, Shimon Whiteson, Leon Kester</i>		
1	Introduction	65
2	Background on Factored MDP	67
3	Feature Selection in Factored MDPs	68
	3.1 Feature Types	68
	3.2 Candidate Representation	69
4	Representation Selection for a Contextual Bandit	71
	4.1 A Contextual Bandit Example	71
	4.2 Constructing the Switch Representation	72
	4.3 Evaluation of a Representation	74
	4.4 Improving Performance by Off-Policy Updating ...	74
5	Representation Selection for an MDP	77
	5.1 Off-Policy Updating of the Unselected Representations	77
	5.2 Off-Policy Updating of the Unselected Switch Actions	78
6	Experimental Results and Discussion	78
	6.1 Contextual Bandit Problem	78
	6.2 MDP Task	81
7	Conclusions	82
8	Future Work	83
	References	83

Part II: Collaborative Decision Making

A Decision-Theoretic Approach to Collaboration: Principal Description Methods and Efficient Heuristic Approximations

87

Frans A. Oliehoek, Arnoud Visser

1	Introduction	87
1.1	Forms of Uncertainty	89
1.2	Decision-Theoretic Approach to MASS	90
1.3	Overview	92
2	The Objective Approach: Dec-POMDPs.....	93
2.1	Decentralized POMDPs.....	93
2.2	Histories and Policies	96
2.3	Solving Dec-POMDPs	99
2.4	Special Cases and Generalization.....	103
3	The Subjective Approach.....	105
3.1	Interactive POMDPs	106
3.2	Solving I-POMDPs.....	108
3.3	The Complexity of Solving I-POMDPs.....	109
4	Application of Decision-Theoretic Models and the Need to Scale up	109
4.1	An Example: RoboCup Rescue as a Dec-POMDP	109
4.2	Aggregation and Hierarchical Decompositions	113
4.3	Modeling and Exploiting Independence between Agents	114
4.4	Compression of the Considered Policy Space.....	115
5	Efficient Heuristic Approaches for Teams of Agents	116
5.1	Allocating Pre-specified Roles Sensibly.....	116
5.2	Frontier Selection in Exploration	117
6	Conclusions	119
	References	120

Efficient Methods for Near-Optimal Sequential Decision Making under Uncertainty.....

125

Christos Dimitrakakis

1	Introduction	125
2	Decision Making under Uncertainty.....	127
2.1	Utility, Randomness and Uncertainty	127
2.2	Uncertain Outcomes.....	128
2.3	Bayesian Inference	129
2.4	Distribution-Free Bounds	131
3	Sequential Decision Making under Uncertainty	133
3.1	Stopping Problems	133
3.2	Bandit Problems	135