

SpringerBriefs in Computer Science

For further volumes:
<http://www.springer.com/series/10028>

Gerardo I. Simari • Simon D. Parsons

Markov Decision Processes and the Belief-Desire-Intention Model

Bridging the Gap for Autonomous Agents

 Springer

Gerardo I. Simari
Department of Computer Science
University of Oxford
United Kingdom
gerardo.simari@cs.ox.ac.uk

Simon D. Parsons
Department of Computer
and Information Science
Brooklyn College
City University of New York
New York
USA
parsons@sci.brooklyn.cuny.edu

ISSN 2191-5768 e-ISSN 2191-5776
ISBN 978-1-4614-1471-1 e-ISBN 978-1-4614-1472-8
DOI 10.1007/978-1-4614-1472-8
Springer New York Dordrecht Heidelberg London

Library of Congress Control Number: 2011937480

© The Author(s) 2011

All rights reserved. This work may not be translated or copied in whole or in part without the written permission of the publisher (Springer Science+Business Media, LLC, 233 Spring Street, New York, NY 10013, USA), except for brief excerpts in connection with reviews or scholarly analysis. Use in connection with any form of information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed is forbidden.

The use in this publication of trade names, trademarks, service marks, and similar terms, even if they are not identified as such, is not to be taken as an expression of opinion as to whether or not they are subject to proprietary rights.

Printed on acid-free paper

Springer is part of Springer Science+Business Media (www.springer.com)

Acknowledgements

This work was possible due to generous support from *Comisión de Investigaciones Científicas de la Provincia de Buenos Aires* (Argentina) granted to Gerardo Simari during his time at *Laboratorio de Investigación y Desarrollo en Inteligencia Artificial* (LIDIA) at *Universidad Nacional del Sur* (Bahía Blanca, Argentina).

The authors are grateful to Martijn Schut for sharing code with them, and to Agent Oriented Software for a license for JACK that allowed them to use that code. Finally, the authors would also like to thank the anonymous reviewers who provided valuable feedback on earlier versions of this work; their comments and suggestions helped improve this manuscript.

Contents

| | | |
|----------|--|----|
| 1 | Introduction | 1 |
| 2 | Preliminary Concepts | 3 |
| 2.1 | The BDI Model | 3 |
| 2.2 | Markov Decision Processes | 6 |
| 3 | An Empirical Comparison of Models | 11 |
| 3.1 | The TILEWORLD Domain | 11 |
| 3.2 | Obtaining Basic Approximate Solutions | 12 |
| 3.2.1 | Reducing State Information | 13 |
| 3.2.2 | Using Optimal Sub-solutions | 14 |
| 3.3 | Experimental Results | 16 |
| 3.3.1 | Results for Dynamism | 18 |
| 3.3.2 | Results for Determinism | 19 |
| 3.3.3 | Results for Accessibility | 20 |
| 3.3.4 | A Larger Environment | 21 |
| 3.4 | Conclusions: Limitations and Significance of the Empirical Evaluation | 23 |
| 4 | A Theoretical Comparison of Models | 27 |
| 4.1 | Preliminary Mappings | 28 |
| 4.2 | Intentions, <i>i</i> -plans, and Policies | 28 |
| 4.3 | Potential Equivalence of Models | 33 |
| 4.4 | From Policies to <i>i</i> -plans | 34 |
| 4.4.1 | Obtaining ‘more universal’ <i>i</i> -plans | 39 |
| 4.4.2 | State Spaces and Heuristics | 40 |
| 4.5 | From <i>i</i> -plans to Policies | 42 |
| 4.5.1 | Theoretical Mapping | 42 |
| 4.5.2 | General Mapping | 45 |
| 4.5.3 | Mapping Deliberation and Means-ends Reasoning into Policies | 46 |

- 5 Related Work** 49
 - 5.1 The BDI model 49
 - 5.2 The MDP model 51
 - 5.3 BDI and MDP: Bridging the Gap 52

- 6 Conclusions, Limitations, and Future Directions** 55
 - 6.1 Limitations 55
 - 6.2 Future Work 56

- References** 57

- Index** 63