

Verifiable Autonomous Systems

How can we provide guarantees of behaviours for autonomous systems such as driverless cars? This tutorial text for professionals, researchers, and graduate students explains how autonomous systems, from intelligent robots to driverless cars, can be programmed in ways that make them amenable to formal verification. The authors review specific definitions, applications, and the unique future potential of autonomous systems, along with their impact on safer decisions and ethical behaviour. The topics discussed in this book include the use of rational cognitive agent programming from the Beliefs–Desires–Intentions paradigm to control autonomous systems, and the role of model-checking in verifying properties of this decision-making component. Several case studies concerning both the verification of autonomous systems and extensions to the framework beyond the model-checking of agent decision-makers are included, along with complete tutorials for the use of the freely available verifiable cognitive agent toolkit Gwendolen, written in Java.

DR LOUISE A. DENNIS is leader of the Autonomy and Verification research group at The University of Manchester and conference coordinator for the ACM Special Interest Group for Artificial Intelligence (AI). She studied mathematics and philosophy at the University of Oxford and received her PhD from the University of Edinburgh in using AI techniques to prove mathematical theorems; her interest in the overlap between mathematics, philosophy, and AI has continued ever since. Her current research encompasses the programming of autonomous systems, the development of agent programming languages, reasoning about systems and programs via formal mathematical techniques, and the ethical implications of AI. Beyond the university setting, Dr Dennis is active in public engagement and spends a lot of time taking Lego Robots into schools to introduce robotics programming to children.

DR MICHAEL FISHER is a professor of Computer Science at The University of Manchester. He holds a Royal Academy of Engineering Chair in Emerging Technologies and is a fellow of both the British Computer Society and the Institution of Engineering and Technology. He was previously a professor of logic and computation in the Department of Computing & Mathematics at Manchester Metropolitan University and a professor of computer science at University of Liverpool. Dr Fisher's research concerns autonomous systems, particularly software engineering, formal verification, safety, responsibility, and trustworthiness. He has been involved in over 200 journal and conference papers and authored the book *An Introduction to Practical Formal Methods Using Temporal Logic* (Wiley) in 2011.

Cambridge University Press & Assessment
978-1-108-48499-2 — Verifiable Autonomous Systems
Louise A. Dennis , Michael Fisher
Frontmatter
[More Information](#)

Verifiable Autonomous Systems
Using Rational Agents to Provide Assurance about
Decisions Made by Machines

LOUISE A. DENNIS
University of Manchester

MICHAEL FISHER
University of Manchester



CAMBRIDGE
UNIVERSITY PRESS

Cambridge University Press & Assessment
978-1-108-48499-2 — Verifiable Autonomous Systems
Louise A. Dennis, Michael Fisher
Frontmatter
[More Information](#)

CAMBRIDGE UNIVERSITY PRESS

Shaftesbury Road, Cambridge CB2 8EA, United Kingdom
One Liberty Plaza, 20th Floor, New York, NY 10006, USA
477 Williamstown Road, Port Melbourne, VIC 3207, Australia
314–321, 3rd Floor, Plot 3, Splendor Forum, Jasola District Centre, New Delhi – 110025, India
103 Penang Road, #05-06/07, Visioncrest Commercial, Singapore 238467

Cambridge University Press is part of Cambridge University Press & Assessment, a department of the University of Cambridge.

We share the University's mission to contribute to society through the pursuit of education, learning and research at the highest international levels of excellence.

www.cambridge.org

Information on this title: www.cambridge.org/9781108484992

DOI: 10.1017/9781108755023

© Louise A. Dennis and Michael Fisher 2023

This publication is in copyright. Subject to statutory exception and to the provisions of relevant collective licensing agreements, no reproduction of any part may take place without the written permission of Cambridge University Press & Assessment.

First published 2023

A catalogue record for this publication is available from the British Library.

Library of Congress Cataloging-in-Publication Data

Names: Dennis, Louise, author. | Fisher, Michael, 1962- author.

Title: Verifiable autonomous systems : using rational agents to provide assurance about decisions made by machines / Louise A. Dennis,

University of Manchester, Michael Fisher, University of Manchester.

Description: Cambridge, United Kingdom ; New York, NY, USA : Cambridge University Press, 2023. | Includes bibliographical references and index.

Identifiers: LCCN 2023006704 | ISBN 9781108484992 (hardback) | ISBN 9781108755023 (ebook)

Subjects: LCSH: Robust control. | Automatic control. | Machine learning. | Computer programs—Verification.

Classification: LCC TJ217.2 .D46 2023 | DDC 629.8—dc23/eng/20230302

LC record available at <https://lcn.loc.gov/2023006704>

ISBN 978-1-108-48499-2 Hardback

Cambridge University Press & Assessment has no responsibility for the persistence or accuracy of URLs for external or third-party internet websites referred to in this publication and does not guarantee that any content on such websites is, or will remain, accurate or appropriate.

Cambridge University Press & Assessment
978-1-108-48499-2 — Verifiable Autonomous Systems
Louise A. Dennis , Michael Fisher
Frontmatter
[More Information](#)

This book is dedicated to Gwendolen Sellers.
The original, you might say.

Cambridge University Press & Assessment
978-1-108-48499-2 — Verifiable Autonomous Systems
Louise A. Dennis , Michael Fisher
Frontmatter
[More Information](#)

Contents

	<i>Acknowledgements</i>	<i>page xi</i>
1	Introduction	1
	1.1 What is an Autonomous System?	2
	1.2 Why Autonomy?	4
	1.3 Why use <i>Formal Verification</i> ?	8
	 Part I Foundations	
2	Autonomous Systems Architectures	13
	2.1 Architectures for Autonomous Systems	13
	2.2 Agent Architectures	16
	2.3 Modularity in Modern Robotic Software	20
	2.4 Practical Systems	21
3	Agent Decision-Maker	24
	3.1 Agents and Cognitive Agents	24
	3.2 Agent Programming	26
	3.3 GWENDOLEN Programming Language	33
	3.4 Agent Environments	40
4	Formal Agent Verification	42
	4.1 What is Verification?	43
	4.2 From Testing to Formal Verification	43
	4.3 Varieties of Formal Verification	44
	4.4 Understanding Program Model-Checking	45
	4.5 Program Model-Checking with Java Pathfinder	50
	4.6 Logical Agent Requirements	53
	4.7 Discussion	57

5	Verifying Autonomous Systems	58
5.1	Modular Architectures for Autonomous Systems	61
5.2	Overview	62
5.3	Verifying Autonomous Choices	63
5.4	Onward	67
6	Agent-Based Autonomous System Verification	68
6.1	From Operational Semantics to Model-Checking	69
6.2	The Property Specification Language	69
6.3	Where Does the Automaton Representing a BDI Agent Program Branch?	71
6.4	The Problem with Environments	72
6.5	Example: Cars on a Motorway	73
6.6	Moving on to Applications	79
 Part II Applications		
7	Multi-Agent Auctions	83
7.1	What is the System?	83
7.2	What <i>Properties</i> Do We Want to Establish?	88
7.3	GWENDOLEN Code	88
7.4	Environments	99
7.5	Example Verification	99
7.6	Issues	103
8	Autonomous Satellite Control	104
8.1	What is the System?	104
8.2	What <i>Properties</i> Do We Want to Establish?	106
8.3	GWENDOLEN Code	110
8.4	Environments	117
8.5	Example Verification	123
8.6	Issues	126
9	Certification of Unmanned Air Systems	127
9.1	What is the System?	127
9.2	What <i>Properties</i> Do We Want to Establish?	131
9.3	GWENDOLEN Code	137
9.4	Environments	145
9.5	Example Verification	146
9.6	Issues	148

Contents

ix

10	Ethical Decision-Making	150
	10.1 What is the System?	150
	10.2 What <i>Properties</i> Do We Want to Establish?	158
	10.3 ETHAN/GWENDOLEN Code	162
	10.4 Environments	167
	10.5 Example Verification	168
	10.6 Issues	169
	10.7 Cognitive Agents as Ethical Governors	170
Part III Extensions		
11	Compositional Verification: Widening Our View beyond the Agent	177
	11.1 Example Systems	177
	11.2 Why use a Compositional Approach to Verification?	180
	11.3 Other Formal Tools	182
	11.4 Some Verification Results	182
	11.5 How Do We Combine These Results?	194
	11.6 Discussion	202
12	Runtime Verification: Recognising Abstraction Violations	203
	12.1 Example System	204
	12.2 Formal Machinery	206
	12.3 Integration of Monitor in the MCAPL Framework	214
	12.4 Verification Results	215
	12.5 Discussion	216
13	Utilising External Model-Checkers	218
	13.1 Example Systems	219
	13.2 Other Model-Checkers	224
	13.3 Translation Approach	226
	13.4 Verification Results	231
	13.5 Discussion	237
Part IV Concluding Remarks		
14	Verifiable Autonomous Systems	241
15	The Future	244
<i>Appendix A</i>	Gwendolen Documentation	246
<i>Appendix B</i>	AIL Toolkit Documentation	297

x

Contents

<i>Appendix C</i>	AJPF Documentation	330
	<i>References</i>	362
	<i>Index</i>	373

Acknowledgements

Thanks are due to the many people who worked with us on the material in this book: Davide Ancona, Martin Mose Bentzen, Rafael Bordini, Paul Bremner, Neil Cameron, Rafael C. Cardoso, Marie Farrell, Angelo Ferrando, Mike Jump, Maryam Kamali, Nick Lincoln, Felix Lindner, Alexei Lisitsa, Matt Luckuck, Viviana Mascardi, Owen McAree, Marija Slavkovik, Sandor Veres, Matt Webster, and Alan F. Winfield. Thanks also to Chris Anderson for proof reading the early chapters.

Cambridge University Press & Assessment
978-1-108-48499-2 — Verifiable Autonomous Systems
Louise A. Dennis , Michael Fisher
Frontmatter
[More Information](#)
