

Natural Computing Series

Peter R. Lewis
Marco Platzner
Bernhard Rinner
Jim Tørresen
Xin Yao *Editors*




Self-aware Computing Systems

An Engineering Approach

 Springer

Natural Computing Series



Series Editors: G. Rozenberg

Th. Bäck A.E. Eiben J.N. Kok H.P. Spaink

Leiden Center for Natural Computing

Advisory Board: S. Amari G. Brassard K.A. De Jong C.C.A.M. Gielen
T. Head L. Kari L. Landweber T. Martinetz Z. Michalewicz M.C. Mozer
E. Oja G. Päun J. Reif H. Rubin A. Salomaa M. Schoenauer
H.-P. Schwefel C. Torras D. Whitley E. Winfree J.M. Zurada

More information about this series at <http://www.springer.com/series/4190>

Peter R. Lewis • Marco Platzner • Bernhard Rinner
Jim Tørresen • Xin Yao
Editors

Self-aware Computing Systems

An Engineering Approach

 Springer

Editors

Peter R. Lewis
School of Engineering & Applied Science
Aston University
Birmingham, United Kingdom

Marco Platzner
Department of Computer Science
Paderborn University
Paderborn, North Rhine-Westphalia
Germany

Bernhard Rinner
Institute of Networked and Embedded Systems
Alpen-Adria-Universität Klagenfurt
Klagenfurt am Wörthersee
Austria

Jim Tørresen
Department of Informatics
University of Oslo
Oslo, Norway

Xin Yao
School of Computer Science
University of Birmingham
Birmingham, United Kingdom

ISSN 1619-7127

Natural Computing Series

ISBN 978-3-319-39674-3

ISBN 978-3-319-39675-0 (eBook)

DOI 10.1007/978-3-319-39675-0

Library of Congress Control Number: 2016942574

© Springer International Publishing Switzerland 2016

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

The use of general descriptive names, registered names, trademarks, service marks, 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.

The publisher, the authors and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, express or implied, with respect to the material contained herein or for any errors or omissions that may have been made.

Printed on acid-free paper

This Springer imprint is published by Springer Nature
The registered company is Springer International Publishing AG Switzerland

Foreword

This book considers the design of new computation systems that are in some ways more responsive to the environment and their own state than current system designs and aim to be more reliable through the creation of self-aware and self-expressive systems. One of the driving forces of this work is the realisation of the growth in system complexity and the difficulty of using current “standard” methods and designs to continue to create working systems. This is certainly relevant, as the interest in the design and understanding of complex computing systems in technical applications has been growing significantly in various research initiatives like autonomic, organic, pervasive or ubiquitous computing and in the multi-agent system community. There have been many novel applications demonstrating a wide range of self-* properties, as well as studies looking also at emerging global behaviour due to self-organised local interaction. The authors of this book present the results of a large European cooperative project focusing specifically on self-awareness, which may be seen as one of the essential backgrounds for developing and supporting the other self-* properties, which is addressed here by the term “self-expression”.

Ever since researchers have realised that machines could be programmed to have increasingly adaptive behaviors, there has been much research on how to introduce adaptive behaviour and more biological like capabilities into systems – more types of reasoning, more types of awareness, and more types of intelligent processing. Particularly important in adaptation is that the system has the knowledge and the capabilities that allow it to do these adaptations in novel situations and at runtime. There are many examples of large-scale programmes to foster the understanding of the necessary attributes and architectures of systems capable of these adaptations. Hence there were programs on adapting routers and networks in real time (e.g., DARPA’s Active Networks), platforms and other plug and play architectures with robust real time services (e.g., DARPA’s META program; Europe’s AUTOSAR (AUTomotive Open System ARchitecture), programs that worked to understand emergent behaviour and make use of it (Europe’s Organic Computing), systems with computational reflection used for resource management (e.g., reflective architectures), and, of course, an enormous amount of work on multi-agent systems and autonomous computing.

In this landmark EU project, these slowly developing themes, drawn from a wide diversity of fields, have been brought together and further developed with both thoughtful discussions on foundations and new research and developments in the engineering of several application areas.

One particularly important aspect of this book is the way in which it builds up our repertoire of engineering methods for self-awareness by purposely drawing its concepts for self-awareness from a diversity of fields and its examples from a diversity of applications. Most importantly, these applications span across different levels of computational systems, from agents and applications (interactive music systems in Chapter 14) to middleware services (Chapter 11) to adaptive networks (Chapter 10) and even hardware (Chapters 8, 9 and 12.)

Starting from insights into “self-awareness” achieved by other disciplines like psychology and philosophy, the notions of “computational self-awareness” and “self-expression” are systematically developed. The majority of the book focuses on computational systems that require some form of anticipation where the new algorithms and methods are needed to provide the appropriate anticipatory behaviour. In practice, these methods can include different forms of self-awareness (such as awareness of goals, of the current state and readiness of system resources, of one’s planning process and of the ordering of events), such that the system is not simply reacting to events and changes, but can anticipate them. The ideas and mechanisms outlined are applied to a number of interesting applications: Computational finance applications using heterogeneous computing clusters are investigated and include self-adaptive algorithms that are supported by hardware; low-latency adaptive network processing; run-time reconfigurable hardware acceleration; heterogeneous computing and hardware/software co-processing for algorithmic trading and reconfigurable hardware acceleration of self-optimisation of reconfigurable hardware designs. Self-awareness in distributed smart camera networks is considered for both single cameras at a node level and multiple camera systems within a network. Interesting bio-inspired methods are aimed at the network level, artificial pheromones are employed to construct a local neighbourhood graph, allowing adaptation in the network as topologies change. A hypermusic demonstrator is considered as a third application. This considers various methods and techniques to enable adaptability (self-expression) in musical output. Three methods within this application are the focus of this work, each providing different input information and overall levels of information: *SoloJam* provides rather overarching rhythmic shaping; *Funky Sole Music* provides what might be considered more specific, lower level, inputs such as walking tempo, movement types and foot activity; *PheroMusic* considers more links between musical soundscapes.

Thus, the book provides a comprehensive introduction to self-aware computation providing a broad range of new theoretical background and foundation before moving on to consider details of architectures and techniques to help design self-aware computational systems, from nodes to networks. Many of the problems that have been addressed in this book will continue to be timely for many years to come and could well provide the focus of research strands within many research fields. Particular challenges remain with respect to performance, safety and security properties

of such systems. Although self-awareness is supposed to improve the performance of computational systems in complex environments, there is still a lack of formal frameworks for rigorously arguing about the behaviour of such systems.

The authors are all well known in this research area and the editors, Lewis, Platzner, Rinner, Tørresen and Yao have done an excellent job in pulling together what is an excellent book.

Los Angeles
Karlsruhe
York

Kirstie Bellman
Hartmut Schmeck
Andy Tyrrell

March 2016

Preface

Self-aware computing is an emerging field of research. It considers systems and applications able to proactively gather and maintain knowledge about aspects of themselves, learning and reasoning on an ongoing basis, and finally expressing themselves in dynamic ways, in order to meet their goals under changing conditions. The aspects they might be aware of include their own internal state, capabilities, goals, environment, behaviour and interactions. The presence of gathered knowledge permits advanced intelligent decision making leading to self-expression: that is, effective, autonomous and adaptive behaviour, based on self-awareness. Self-awareness and self-expression capabilities are key to designing and operating future computing systems that will inherently and autonomously deal with high levels of dynamics and uncertainty, heterogeneity, scalability, resource constraints and decentralisation. Concepts of self-awareness have been established in psychology, philosophy and cognitive science but are relatively new to computing. In computing systems, our concepts of self-awareness and self-expression integrate and enhance a number of recent approaches dealing with systems with so-called self-* properties, e.g., self-adaptation, self-organisation and self-healing.

This book is the first ever to focus on the emerging field of self-aware computing from an engineering perspective. It first comprehensively introduces fundamentals for self-awareness and self-expression in computing systems, proposing the new notion of *computational self-awareness*. It then focuses on architectures and techniques for designing self-aware computing systems at the node and network levels. Finally, the effectiveness of these techniques is demonstrated on a variety of case studies. While a number of books on related topics such as self-adaption and self-organisation, and even self-awareness concepts in computing, have already been published, this book is unique as it provides a holistic view of self-aware computing including its relationship with self-expression, and the process of engineering such systems, i.e., a thorough understanding of how to model and build self-aware computing systems based on design patterns and techniques.

This book targets graduate students and professionals in the fields of computer science, computer engineering, and electrical engineering, but also practitioners and scientists from other fields interested in engineering systems with advanced proper-

ties relying on their ability to reason about themselves in a complex environment. The authors and editors of this book are active researchers in various aspects related to self-aware computing systems. They have a strong track record in successfully collaborating on this topic, for example, through the European FET project “Engineering Proprioception in Computing Systems (EPiCS)”. The extensive joint experience of the contributors makes this edited book consistent and well integrated. Therefore, we specifically recommend this book as reading material for the graduate level or for self-study on self-aware computing systems.

The book reports some of the latest results in self-aware and self-expressive computing, and we hope it serves as a launchpad for further research discussions and new ideas in the future.

Birmingham
Paderborn
Klagenfurt
Oslo
Birmingham

Peter R. Lewis
Marco Platzner
Bernhard Rinner
Jim Tørresen
Xin Yao

March 2016

Acknowledgements

The research leading to many results in this book was conducted during the EPiCS project (Engineering Proprioception in Computing Systems) and received funding from the European Union Seventh Framework Programme under grant agreement no. 257906.

The contributors would like to acknowledge additional support for research performed in individual chapters of this book.

- Chapters 6 and 7 were also supported by EPSRC Grants (Nos. EP/I010297/1, EP/K001523/1 and EP/J017515/1).
- Chapter 8 was also supported by the German Research Foundation (DFG) within the Collaborative Research Centre “On-The-Fly Computing” (SFB 901) and the International Graduate School on Dynamic Intelligent Systems of Paderborn University.
- Chapter 9 was also supported in part by HiPEAC NoE, by the European Union Seventh Framework Programme under grant agreement numbers 287804 and 318521, by the UK EPSRC, by the Maxeler University Programme, and by Xilinx.
- Chapter 12 was also supported in part by the China Scholarship Council, by the European Union Seventh Framework Programme under grant agreement numbers 287804 and 318521, by the UK EPSRC, by the Maxeler University Programme, and by Xilinx.
- Chapter 13 was also supported by the research initiative Mobile Vision with funding from the Austrian Institute of Technology and the Austrian Federal Ministry of Science, Research and Economy HRSMV programme BGBl. II no. 292/2012.
- Chapter 14 was also supported by the Research Council of Norway under grant agreement number 240862/F20.
- Peter Lewis would like to thank the participants of the Dagstuhl Seminar “Model-Driven Algorithms and Architectures for Self-aware Computing Systems”, Seminar Number 15041, for many insightful discussions on notions of self-aware computing.

Contents

1	Self-aware Computing: Introduction and Motivation	1
	Peter R. Lewis, Marco Platzner, Bernhard Rinner, Jim Tørresen, and Xin Yao	
1.1	Self-aware Computing: A New Paradigm	1
1.2	Organisation of This Book	4
Part I Concepts and Fundamentals		
2	Self-awareness and Self-expression: Inspiration from Psychology . . .	9
	Peter R. Lewis, Arjun Chandra, and Kyrre Glette	
2.1	Introduction to Self-awareness	9
2.2	Key Concepts for Self-aware Systems	11
2.2.1	Public and Private Self-awareness	12
2.2.2	Levels of Self-awareness	13
2.2.3	Self-awareness in Collective Systems	15
2.2.4	Self-expression	15
2.3	Computational Self-awareness	16
2.3.1	Private and Public Computational Self-awareness	16
2.3.2	Levels of Computational Self-awareness	17
2.3.3	Collective and Emergent Computational Self-aware Systems	19
2.4	Summary	21
3	Relationships to Other Concepts	23
	Kyrre Glette, Peter R. Lewis, and Arjun Chandra	
3.1	Introduction	23
3.2	Self-awareness in Artificial Intelligence	25
3.3	Self-awareness in Collective Systems	25
3.4	Formal Models for Self-awareness	26
3.5	Self-awareness in Engineering	27
3.6	Self-awareness in Pervasive Computing	29

- 3.7 Self-awareness in Robotics 29
- 3.8 Self-awareness in Autonomic Systems 30
- 3.9 Self-awareness in Organic Computing 32
- 3.10 Self-expression in Computing 33
- 3.11 Summary 34

4 Reference Architecture for Self-aware and Self-expressive

- Computing Systems 37**
- Arjun Chandra, Peter R. Lewis, Kyrre Glette, and Stephan C. Stilkerich
- 4.1 Introduction 37
- 4.2 Architectures for Designing Self-adaptive Systems 38
- 4.3 Generic Reference Architecture for Designing Self-aware and
Self-expressive Computing Systems 42
 - 4.3.1 Reference Architecture for Agents 43
 - 4.3.2 Architecting Collectives 47
- 4.4 Reference Architecture in Practice 49

Part II Patterns and Techniques

**5 Design Patterns and Primitives: Introduction of Components and
Patterns for SACS 53**

- Tao Chen, Funmilade Faniyi, and Rami Bahsoon
- 5.1 Introduction and Motivation 53
- 5.2 Patterns for Self-aware Architecture Style 54
 - 5.2.1 Basic Notations 54
 - 5.2.2 The Self-aware Patterns 56
- 5.3 Architectural Primitives and Attributes for Self-aware Systems . . . 70
 - 5.3.1 Taxonomy of Primitives 71
 - 5.3.2 List of Architectural Primitives and Attributes 71
- 5.4 Discussion 73
 - 5.4.1 Phase 1: Collect Requirements and Constraints 73
 - 5.4.2 Phase 2: Propose Candidate Architecture 74
 - 5.4.3 Phase 3: Select the Best Pattern(s) 74
 - 5.4.4 Phase 4: Fit the Selected Pattern(s) 75
 - 5.4.5 Step 5: Determine the Important Primitives and the
Possible Alternatives for Non-functional Requirements . . . 75
 - 5.4.6 Step 6: Create Scenarios 76
 - 5.4.7 Step 7: Score the Alternative of Primitives Against
Each Non-functional Attribute Using Analytical or
Simulation Models 76
 - 5.4.8 Step 8: Find the Best Alternatives for the Final
Architecture View 78
- 5.5 Conclusion 78

6 Knowledge Representation and Modelling: Structures and Trade-Offs 79
 Leandro L. Minku, Lukas Esterle, Georg Nebehay, and Renzhi Chen

6.1 Introduction 80

6.2 Adaptivity 80

 6.2.1 Definition and Examples 81

 6.2.2 Implications 83

6.3 Robustness 89

 6.3.1 Definitions and Examples 89

 6.3.2 Implications 92

6.4 Multi-objectivity 96

 6.4.1 Definition and Examples 96

 6.4.2 Implications 99

6.5 Decentralisation 104

 6.5.1 Definitions and Examples 105

 6.5.2 Implications 108

6.6 Summary 111

7 Common Techniques for Self-awareness and Self-expression 113
 Shuo Wang, Georg Nebehay, Lukas Esterle, Kristian Nymoen, and Leandro L. Minku

7.1 Introduction 114

7.2 Online Learning 114

 7.2.1 Example Application 115

 7.2.2 Benefits and Challenges at Levels of Self-awareness 119

 7.2.3 Other Related Techniques 120

7.3 Nature-Inspired Learning 125

 7.3.1 Example Application 125

 7.3.2 Benefits and Challenges at Levels of Self-awareness 130

 7.3.3 Other Related Techniques 132

7.4 Socially-Inspired Learning in Collective Systems 133

 7.4.1 Example Application 134

 7.4.2 Benefits and Challenges at Levels of Self-awareness 139

 7.4.3 Other Related Techniques 141

Part III Nodes and Networks

8 Self-aware Compute Nodes 145
 Andreas Agne, Markus Happe, Achim Lösch, Christian Plessl, and Marco Platzner

8.1 Heterogeneous Multi-cores 146

8.2 Related Work on Self-aware Compute Nodes 147

8.3 Reference Architecture for Self-aware Compute Nodes 150

8.4 ReconOS 151

 8.4.1 Architecture and Programming 152

 8.4.2 Partial Reconfiguration 153

8.4.3	Sensors and Actuators	154
8.4.4	Availability of ReconOS	156
8.5	Case Study for a Self-aware Heterogeneous Multi-core	156
8.5.1	Self-expression Under Performance Constraints	158
8.5.2	Self-expression Under Conflicting Constraints	162
8.5.3	Comparison of Self-expression Strategies	163
8.6	Discussion and Conclusion	165
9	Self-adaptive Hardware Acceleration on a Heterogeneous Cluster . .	167
	Xinyu Niu, Tim Todman, and Wayne Luk	
9.1	Overview of Heterogeneous Computing	168
9.1.1	Heterogeneous Clusters: Performance	168
9.1.2	Heterogeneous Clusters: Verification	172
9.2	Architectures of Heterogeneous Clusters	173
9.2.1	Overview of Existing Heterogeneous Clusters	173
9.2.2	Software Layers in Heterogeneous Clusters	174
9.3	Self-aware and Self-adaptive Applications for Heterogeneous Clusters	177
9.3.1	Self-awareness in Heterogeneous Clusters	177
9.3.2	Runtime Scenarios	178
9.3.3	Monitoring	179
9.3.4	Adaptive Strategies in Heterogeneous Clusters	182
9.3.5	Computational Capacity	182
9.3.6	Workload Distribution	183
9.3.7	Communication Scheduling	183
9.4	Evaluation Results	184
9.4.1	Benchmark Applications	184
9.4.2	Self-adaptive Temperature Control	186
9.4.3	Self-adaptivity for Resource Availability Variations	186
9.5	Verification of Heterogeneous Clusters	188
9.5.1	Verification of Hardware-Software Codesign	189
9.5.2	Runtime Verification by In-Circuit Statistical Assertions	191
9.5.3	Results	192
9.6	Summary	192
10	Flexible Protocol Stacks	193
	Markus Happe and Ariane Trammell-Keller	
10.1	Introduction	194
10.2	Concepts and Methodologies	194
10.2.1	Self-aware/expressive Network Node Architecture	197
10.2.2	Protocol Stack Negotiation and Adaptations	198
10.2.3	Dynamic Hardware/Software Mapping	199
10.3	EmbedNet Execution Environment	200
10.4	Case Studies	202
10.4.1	Sensor Network	203
10.4.2	Smart Camera Network	204

- 10.5 Comparison to Related Research Projects 211
- 10.6 Conclusion..... 213
- 11 Middleware Support for Self-aware Computing Systems 215**
 Jennifer Simonjan, Bernhard Dieber, and Bernhard Rinner
- 11.1 Introduction to Middleware Systems 216
 - 11.1.1 Middleware Basics 216
 - 11.1.2 Application Example of a Distributed Self-aware
 Computing System 217
- 11.2 Middleware Requirements 219
- 11.3 Middleware Paradigms 221
 - 11.3.1 Host-Centric Middleware 222
 - 11.3.2 Content-Centric Middleware 225
 - 11.3.3 Requirements Conformity of Middleware Paradigms 227
- 11.4 Publish/Subscribe 228
 - 11.4.1 Publish/Subscribe Flavours 230
 - 11.4.2 Decoupling..... 231
 - 11.4.3 Publish/Subscribe for SACS 231
- 11.5 Ella: A Publish/Subscribe-Based Hybrid Middleware 232
 - 11.5.1 Architecture 232
 - 11.5.2 SACS-Specific Features in Ella 235
 - 11.5.3 Ella in Practice..... 236
- 11.6 Conclusion..... 237

Part IV Applications and Case Studies

- 12 Self-aware Hardware Acceleration of Financial Applications on a
 Heterogeneous Cluster 241**
 Maciej Kurek, Tobias Becker, Ce Guo, Stewart Denholm,
 Andreea-Ingrid Funie, Mark Salmon, Tim Todman, and Wayne Luk
- 12.1 Introduction 242
 - 12.1.1 Overview of Techniques and Tools 242
- 12.2 Rule-Based Algorithmic Trading 243
- 12.3 Model-Based Algorithmic Trading 245
- 12.4 Market Data Feed Arbitration 245
- 12.5 In Detail: ARDEGO — Machine Learning-Based Optimisation
 of Reconfigurable Systems 246
 - 12.5.1 Background 247
 - 12.5.2 ARDEGO Approach 251
 - 12.5.3 Acceleration of ARDEGO 255
 - 12.5.4 Evaluation 256
- 12.6 Conclusion..... 260

13 Self-aware Object Tracking in Multi-Camera Networks 261
 Lukas Esterle, Jennifer Simonjan, Georg Nebhay, Roman Pflugfelder,
 Gustavo Fernández Domínguez, and Bernhard Rinner

13.1 Smart Camera Networks 262

13.2 Object Tracking 263

13.3 Multi-camera Tracking Coordination 264

13.4 Self-aware and Self-expressive Building Blocks 265

13.4.1 Object Tracking 266

13.4.2 Object Handover 267

13.4.3 Topology Learning 269

13.4.4 Strategy Selection 269

13.4.5 Resource Monitoring 270

13.4.6 Constraints and Objectives 270

13.5 Camera Network Case Study 271

13.5.1 Camera Network Setup 271

13.5.2 Tracking Results 272

13.5.3 Topology Learning 274

13.5.4 Communication and Utility Trade-off 274

13.6 Conclusion and Outlook 275

14 Self-awareness in Active Music Systems 279
 Kristian Nymoen, Arjun Chandra, and Jim Tørresen

14.1 Introduction 279

14.2 Decentralised Circulation of Musical Control 281

14.2.1 SoloJam Algorithmic Details 282

14.2.2 SoloJam Implementation 284

14.3 Adaptive Mapping in Active Music Systems 286

14.3.1 Gesture Recognition in Active Music Systems 288

14.3.2 Pheromone-Inspired Gait Recognition 288

14.3.3 Music Synthesis in Funky Sole Music 290

14.3.4 Adaptive Mapping 290

14.4 Pheromone Trails in a Musical Space 292

14.4.1 Flexible Musical Scenes 292

14.4.2 Pheromone Mechanism 294

14.5 Conclusion 296

15 Conclusions and Outlook 297
 Peter R. Lewis, Marco Platzner, Bernhard Rinner, Jim Tørresen, and
 Xin Yao

15.1 Computational Self-awareness 298

15.2 Challenges and Research Questions 299

References 301

Index 323

List of Contributors

Andreas Agne

Paderborn University, Germany. e-mail: agne@upb.de

Rami Bahsoon

University of Birmingham, UK. e-mail: r.bahsoon@cs.bham.ac.uk

Tobias Becker

Imperial College London, UK. e-mail: tobias.becker@imperial.ac.uk

Arjun Chandra

Studix, Norway. e-mail: arjun@studix.com

Renzhi Chen

University of Birmingham, UK. e-mail: rx332@cs.bham.ac.uk

Tao Chen

University of Birmingham, UK. e-mail: t.chen@cs.bham.ac.uk

Stewart Denholm

Imperial College London, UK. e-mail: stewart.denholm10@imperial.ac.uk

Bernhard Dieber

Alpen-Adria-Universität Klagenfurt, Austria. e-mail: bernhard.dieber@county.at

Lukas Esterle

Alpen-Adria-Universität Klagenfurt, Austria. e-mail: lukas.esterle@aau.at

Gustavo Fernández Domínguez

Austrian Institute of Technology, Austria. e-mail: gustavo.fernandez@ait.ac.at

Funmilade Faniyi

University of Birmingham, UK. e-mail: f.faniyi@gmail.com

Andreea-Ingrid Funie

Imperial College London, UK. e-mail: andreea.funie09@imperial.ac.uk

Kyrre Glette

University of Oslo, Norway. e-mail: kyrrehg@ifi.uio.no

Ce Guo

Imperial College London, UK. e-mail: ce.guo10@imperial.ac.uk

Markus Happe

ETH Zurich, Switzerland. e-mail: markus.happe@alumni.ethz.ch

Maciej Kurek

Imperial College London, UK. e-mail: mk306@imperial.ac.uk

Peter R. Lewis

Aston University, UK. e-mail: p.lewis@aston.ac.uk

Achim Loesch

Paderborn University, Germany, e-mail: achim.loesch@upb.de

Wayne Luk

Imperial College London, UK. e-mail: w.luk@imperial.ac.uk

Leandro L. Minku

University of Leicester, UK. e-mail: leandro.minku@leicester.ac.uk

Georg Nebehay

Austrian Institute of Technology, Austria. e-mail: gnebehay@gmail.com

Xinyu Niu

Imperial College London, UK. e-mail: niu.xinyu10@imperial.ac.uk

Kristian Nymoen

University of Oslo, Norway. e-mail: kristian.nymoen@imv.uio.no

Roman Pflugfelder

Austrian Institute of Technology, Austria. e-mail: roman.pflugfelder@ait.ac.at

Marco Platzner

Paderborn University, Germany. e-mail: platzner@upb.de

Christian Plessl

Paderborn University, Germany. e-mail: christian.plessl@uni-paderborn.de

Bernhard Rinner

Alpen-Adria-Universität Klagenfurt, Austria. e-mail: bernhard.rinner@aau.at

Mark Salmon

University of Cambridge, UK. e-mail: mhs39@cam.ac.uk

Jennifer Simonjan

Alpen-Adria-Universität Klagenfurt, Austria. e-mail: jennifer.simonjan@aau.at

Stephan C. Stilkerich

Airbus Group Innovation, Germany. e-mail: stephan.stilkerich@airbus.com

Tim Todman

Imperial College London, UK. e-mail: timothy.todman@imperial.ac.uk

Jim Tørresen

University of Oslo, Norway. e-mail: jimtoer@ifi.uio.no

Ariane Trammell-Keller

ETH Zurich, Switzerland. e-mail: ariane.trammell@alumni.ethz.ch

Shuo Wang

University of Birmingham, UK. e-mail: s.wang@cs.bham.ac.uk

Xin Yao

University of Birmingham, UK. e-mail: x.yao@cs.bham.ac.uk

Acronyms

ACO	Ant Colony Optimisation
AES	Advanced Encryption Standard
ALA	Ant Learning Algorithm
API	Application Programming Interface
BSD	Berkeley Software Distribution
CDC	Concept Drift Committee
CDT	Correct Detected Track
CMT	Consensus-Based Matching and Tracking
CPU	Central Processing Unit
CUDA	Compute Unified Device Architecture
CV	Computer Vision
DDD	Diversity for Dealing with Drifts
DDM	Drift Detection Method
DPS	Dynamic Protocol Stack
DWM	Dynamic Weight Majority
EA	Evolutionary Algorithm
EDDM	Early Drift Detection Method
EGO	Efficient Global Optimisation
FAT	False Alarm Track
FB	Functional Block
FF	Flip-Flop
FPGA	Field-Programmable Gate Array
FPS	Frames per Second
FMC	FPGA Mezzanine Card
FOV	Field of View
FSR	Force-Sensitive Resistor
GA	Genetic Algorithm
GPSP	General Purpose Sensor Platform
GP	Gaussian Process
GPU	Graphics Processing Unit
H2S	Hardware-to-Software

HLS	High Level Synthesis
HMM	Hidden Markov Models
HPC	High Performance Computing
ICAP	Internal Configuration Access Port
IDP	Information Dispatch Point
ILP	Integer Linear Programming
IP	Internet Protocol
IPC	Inter-process Communication
LUT	Look-up Table
MAC	Media Access Protocol
MLO	Machine Learning Optimiser
MOEA/D	Multi-objective Evolutionary Algorithm Based on Decomposition
MOP	Multi-objective Optimisation Problem
MPI	Message Passing Interface
MTBF	Mean Time Between Failures
NoC	Network-on-Chip
OSC	Open Sound Control
OT	Object Tracking
PE	Processing Element
RAP	Redundancy Allocation Problem
RAM	Random Access Memory
RTM	Reverse Time Migration
S2H	Software-to-Hardware
SA	Self-aware
SACS	Self-aware Computing Systems
SDRAM	Synchronous Dynamic Random Access Memory
SE	Self-expression
SIMD	Single Instruction, Multiple Data
SMT	Satisfiability Modulo Theories
SVM	Support Vector Machine
SoC	System-on-Chip
SOP	Single-Objective Optimisation Problem
SSE	Streaming SIMD Extensions
STEPD	Statistical Test of Equal Proportions
TCP	Transmission Control Protocol
TDF	Track Detection Failure
TPOT-RL	Team-Partitioned Opaque-Transition Reinforcement Learning
Todi	Two Online Classifiers for Learning and Detecting Concept Drift
TLD	Tracking-Learning-Detection
UDP	User Datagram Protocol
VHDL	Very High Speed Integrated Circuit Hardware Description Language

Glossary

This glossary lists important terms used in this book, in particular in Part I “Concepts and Fundamentals”, with accompanying descriptions or definitions. The glossary is organised into four sections: concepts of self-awareness and self-expression, engineering self-aware systems, related approaches, and general terms. The terms in each of the sections are listed alphabetically.

Concepts of Self-awareness and Self-expression

self-awareness Self-awareness is a broad concept which describes the property of a system (typically a human) which has knowledge of “itself”, based on its own senses (perceptual) and internal models (conceptual). This knowledge may take different forms (cf. levels of self-awareness), and be based on perceptions of both internal and external phenomena (cf. public vs. private self-awareness). It can be a property of single systems (e.g., agents) and collective systems.

collective self-awareness Collective self-awareness refers to the self-awareness property of a collective system, i.e., as opposed to a single agent. Levels of, and public/private self-awareness apply also at this abstraction. This means that a self-aware system is not required to have a central “knowledge” component (though it may have, if desired).

computational self-awareness Computational self-awareness is a notion we have developed to refer to a computational interpretation of self-awareness. Since much of the literature on self-awareness does not readily make sense to engineers or applies directly to technical systems, aspects of computational self-awareness are designed to describe self-awareness properties of computational systems, inspired by self-awareness in humans.

emergent self-awareness This is a special case of collective self-awareness, when the collective self-awareness properties are present, but it is not obvious how this comes about by simply examining the behaviour of individual nodes within a collective.

level(s) of self-awareness A very common theme in self-awareness theory is the distinction between several levels of self-awareness, to describe different aspects or capabilities which comprise a system's complex self-awareness. There are many examples of "sets of levels" to be found in the literature. In developing our notion of computational self-awareness, we have based a set of levels of computational self-awareness on the set of levels (for humans) proposed by Ulric Neisser. Note that our levels are not hierarchical, do not build on each other, nor are they in any particular order, save that the ecological self/stimulus awareness is the most basic, and the conceptual self/meta-self-awareness is typically the most complex.

meta-self-awareness Meta-self-awareness is one of the levels of computational self-awareness we propose, indeed the highest one in our framework. It refers to the capability of a system to be aware of its own self-awareness. This can be very useful, since it means a system has knowledge, obtained at run time, about its own self-awareness processes, including, for example, how effective its learning is at present, or how much resource is being spent to maintain its knowledge. Meta-self-awareness is closely related to, and permits, meta-reasoning. It is a concept inspired directly from human psychology.

private self-awareness Private self-awareness refers to a system's ability to obtain knowledge based on phenomena that are internal to itself. A system needs internal sensors to achieve this. Again, this is a notion which exists in human self-awareness theory, and also features in computational self-awareness.

public self-awareness Public self-awareness refers to a system's ability to obtain knowledge based on phenomena external to itself. Such knowledge depends on how the system itself senses/observes/measures aspects of the environment it is situated in, and includes knowledge of its situation and context, as well as (potential) impact and role within its environment. This is a notion which exists in human self-awareness theory, and also features in computational self-awareness.

scope of self-awareness The scope of self-awareness refers to the domain of phenomena able to be sensed and modelled by the self in question. For a system which is only privately self-aware, the scope may be the same as the span (i.e., it has no perception of its environment). For a system which has some private and some public self-awareness, the scope would be larger than the span, and include external social or physical aspects of the environment. The term scope can be useful to avoid having to use the word "level" to mean multiple things simultaneously in a passage of text.

self-aware system We do not formally define this in the book, however we generally consider a self-aware system to be one which (at least) obtains and maintains knowledge relating to itself (including its perspective of its environment), without external control.

self-awareness capability When a particular level of self-awareness is present in a system, we refer to this as the system having that particular self-awareness capability. For example, a node may have a time-awareness capability, indicating that

it implements the time-awareness level. Levels may be realised in different ways simultaneously in the same system, meaning that, for example, a system may have several time-awareness capabilities.

self-explanation Another form of self-expression when based on self-awareness, self-explanation is the ability of a system to explain/justify its behaviour to an entity on the outside (such as a user or another system).

self-expression Self-expression is, in the general sense, behaviour based on self-awareness. It may include a wide range of different actions, enacted through a system's actuators, including self-adaptation, self-explanation, or just normal system behaviour. Self-expression can also be considered as a property of a collective, since a collective's behaviour can also be based on collective self-awareness. Examples of this might include the adaptive behaviour of a flock of birds in response to an external (to the flock) stimulus.

self-expression capability As with self-awareness capabilities, self-expression capabilities refer to the presence of an implementation of self-expression in a system. For example, a system which adapts its parameters in response to its goal-awareness, would have a self-expression capability. Again, multiple self-expression capabilities may be present simultaneously.

self-knowledge Self-knowledge is a general term for knowledge (usually held in a learnt model) concerning the system itself, which typically is produced as part of a self-awareness process. Note that this can include objective self-knowledge (i.e., about the system as an object in the world, how it interacts with others, how its internal state changes, etc.) and also subjective self-knowledge (i.e., about its experiences, sensor data, changing context, etc.).

self-optimisation Self-optimisation is a form of self-expression; self-optimisation is the ability of a system to optimise itself by improving metrics such as performance or power consumption.

span of self-awareness We use this term to refer to the domain of the subject of the self-awareness, i.e., it is the answer to the question: who is the self here? For example, if a single agent is self-aware, then the span is the agent. If we are considering the collective self-awareness of a network of smart sensors, then the span would be the network. The term span can be useful to avoid having to use the word "level" to mean multiple things simultaneously in a passage of text.

Engineering Self-aware Systems

(architectural) pattern We produced eight architectural patterns, which are derived from the reference architecture and describe how various capabilities (such as levels of self-awareness, etc.) can be included or excluded as appropriate to the application need.

methodology for engineering self-aware systems We developed a methodology for engineering self-aware systems, based on the reference architecture and the derived architectural patterns.

primitive A primitive is a particular block in the reference architecture, representing, for example, a level of self-awareness, self-expression and a sensor. They are instantiated for particular applications.

reference architecture We developed a reference architecture which captures the core aspects of computational self-awareness. The aim is to provide a common, principled basis on which researchers and practitioners can structure their work. We have argued that the psychological foundations, while not strictly necessary, can provide a means of channelling a wide range of ideas, which would perhaps otherwise not have occurred to engineers, acting to inspire the design of future computing systems. The architecture can also be used as a template for identifying common ways of implementing self-awareness capabilities. Different implementations of the same capability can thereby be compared and evaluated. Further, we have derived a set of architectural patterns from the reference architecture.

(self-aware) node We use the term self-aware node to refer to various types of system that are self-aware, e.g., an agent, a robot and a camera. Agent is an alternative term, but node can be used when not wanting to be specific about a particular system being an agent. We also claim that self-aware collectives (see next entry) can be viewed as self-aware nodes, at a higher level of abstraction. A node may or may not correspond to a physical system—this is not a requirement, but it may often make sense to make it correspond.

tactic/algorithm/technique A tactic is a particular instantiation of a primitive in the reference architecture, typically referred to as a particular algorithm, technique, etc. These are application specific. Multiple tactics may be suitable for a particular primitive, and some tactics may implement multiple primitives simultaneously.

Related Approaches

autonomic (computing) Autonomic computing is a vision originally pioneered by IBM, of engineered systems which manage themselves. This self-management is stated to include: self-configuration, self-optimisation, self-healing and self-protection. The aim is to reduce the need for human involvement in the management of complex computing systems. Some autonomic computing literature mentions the need for self-awareness as a characteristic to support self-management, though the literature on autonomic computing does not significantly expand on this. (Not to be confused with autonomous.)

autonomous (system) Autonomy is a broad notion with much disagreement surrounding it. However, in general, an autonomous system is one which acts without any external direction. Examples include robots, vehicles and software agents. In many cases, this ability to make decisions is based on a method of decision making pre-programmed into the system, in other cases it is learnt online at run time. The

types of systems we are concerned with in this book are ones which would typically be considered to be autonomous to a greater or lesser extent. (Not to be confused with autonomic.)

metacognition/metareasoning Metareasoning is reasoning about reasoning, and has been the topic of a significant amount of research primarily in the US, where it has been primarily led by DARPA. Metareasoning relies on meta-self-awareness, and again the metareasoning community has discussed self-awareness as being important, but not expanded on the notion significantly.

organic computing This is a vision from a long-running (primarily) German research project to create “life-like” engineered systems, in which self-organising emergent behaviour is controlled (by an observer/controller component), to ensure desirability in the self-organisation. The Organic Computing literature also mentioned self-awareness as beneficial, but again does not expand on this significantly.

General Terms

adaptability In high level terms, this is similar to adaptivity, but describes a system’s potential for adaptation, rather than actual realised adaptivity.

adaptivity In high level terms, this concerns the amount to which a system adapts, e.g., in the presence of a changing environment, or as a result of its learning.

collective We use the term collective to refer to various types of distributed systems, typically without central control. Examples include swarms, systems-of-systems, populations, multi-agent systems, interwoven systems, etc. The term can be used when there is a need to talk generally of these types of systems, without restricting the discussion to a specific one.

learnt model A learnt model is a model which has been induced through a process of (typically online) learning, based on data from sensors and other existing models. Learnt models hold the conceptual knowledge a self-aware system has concerning itself, its interactions, history, expectations, goals, etc.

model We use the term model in a very general way, to refer to a conceptual representation of some knowledge, typically obtained through sensors. A model could simply be a direct representation of some data, or could be abstractions of that data, or further data synthesised from sensory input.

online learning Online learning is the process of learning a model from data on an ongoing basis. Typically, not all data is available in advance (e.g., it arrives in a streaming fashion from sensors), and the concept being learnt may change over time (i.e., concept drift). In online learning, models are often used (e.g., through self-expression in this case) before learning “completes”, if indeed it ever does. Hence most online learning algorithms also need to be anytime algorithms, implying that models are used and improved continuously as time goes by.

self-adaptive system A system which adapts (typically its behaviour) in response to external or internal changes, but without external control. We have argued that