

Natural Computing Series

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

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 Nebehay, 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.

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

self-awareness is an enabling property for effective self-adaptation. When self-adaptation behaviour is based on self-awareness, it is a form of self-expression.

self-organising system A system which changes its organisation (e.g., its structure, architecture, topology), without external control.

Chapter 1

Self-aware Computing: Introduction and Motivation

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

1.1 Self-aware Computing: A New Paradigm

Designing and operating computing and communication systems are becoming increasingly challenging tasks, due to a multitude of reasons. First, compute nodes are evolving towards parallel and heterogeneous architectures to realise performance gains while minimising their power consumption. Progress in micro(nano)-electronics allows us to integrate more and more functionality on a single compute node, but at the same time requires us to deal with increasing numbers of faulty and unreliable components. Second, distributed systems are growing in the numbers and heterogeneity of nodes and must be able to cope with an increasing level of dynamics. The network topology and the collective resources of a distributed system can vary strongly during runtime since nodes may leave and enter the network dynamically. The position, functionality and available resources of each node may also change dynamically. Third, future challenging application domains have quite divergent requirements with respect to functionality and flexibility, performance, resource usage and costs, reliability and safety, and security. Fuelled by technological progress, applications with exciting levels of user interaction will be possible, and these dynamic socio-technical systems bring with them numerous additional runtime trade-offs to consider. Fourth, the size and complexity of decentralised com-

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

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

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

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

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

puting systems have grown at an increasingly fast rate, posing new challenges in terms of scalability and complexity. Based on our experience, we believe that current design and operation principles and methods will neither be able to scale with future systems, nor efficiently handle the variety and changing nature of requirements and optimisation goals. Novel design and operation principles and methods, such as those incorporating self-awareness (SA) and self-expression (SE), are needed.

Self-aware computing describes a new paradigm for systems and applications that proactively gather information; maintain knowledge about their own internal states and environments; and then use this knowledge to reason about behaviours. This paradigm is well suited for advanced intelligent decision making in dynamic and uncertain environments, which can in turn support effective and explainable autonomy and self-adaptation. Self-expression describes behaviours that are based on the knowledge acquired through system self-awareness, such as self-adaptation and self-explanation. A self-expressive system can adapt to its environment including its users, and thus limits the need for users to adapt to fixed system behaviours [390].

Self-awareness is not a new concept and has been studied for a long time in the fields of psychology and cognitive science [275]. However, a clear understanding and interpretation of self-awareness in computer science and engineering is lacking. There has been no universally agreed and accepted definition of self-aware computing in spite of frequent use of the word “self-aware” in different contexts. Although there are systems that are declared to be self-aware in one sense or another, little has been said about engineering methodologies that can help to build such systems. It has not previously been clear what properties a self-aware system could and should have, and what capabilities such a system might have.

This book attempts to bridge the gaps in the literature related to self-aware computing, in the spirit of recent work to translate concepts of self-awareness from psychology to computing [236]. It focuses on key ideas from self-awareness theory, leading to working definitions and pragmatic principles of *computational self-awareness* that can be used in engineering self-aware systems and understanding their behaviours. While it surveys a range of views concerning self-awareness, it does not attempt to engage in more philosophical debates on the ability of machines to achieve so-called “true” self-awareness, or what that might mean. As an example of the pragmatic approach taken in this book, building on notions from psychology, different levels of self-awareness as they apply to computing systems are described, examples are then used to illustrate what capabilities a system could have with which level(s) of self-awareness, how these might be implemented, and what benefits and costs are associated with such functionality. An engineering methodology is then introduced to facilitate the design of self-aware systems with different required capabilities.

This book builds on and extends earlier work in the related fields including autonomic computing [105] and organic computing [277] and architectures like MAPE-K and Kramer and Magee’s three-layered architecture [225]. It takes an engineering approach to the design of self-aware computing systems, that includes considering different levels of self-awareness and the introduction of a reference architecture for designing self-aware and self-expressive computing systems. As a more con-

in distributed applications and briefly describes a dedicated middleware implementation.

Part IV demonstrates how self-awareness and self-expression are useful in the three widely different application domains of hardware acceleration of financial computation, object tracking in multi-camera networks, and active music systems, respectively. Chapter 12 demonstrates how complex financial models can be speeded up using reconfigurable hardware combined with optimisation algorithms. Object tracking in multi-camera networks is the topic of Chapter 13, where autonomous monitoring of each camera in a network is combined with learning mechanisms to adapt its behaviour to changing conditions. Finally, Chapter 14 illustrates how persons without musical skills can influence music in interactive music systems using nature and socially-inspired methods.

The book is a result of extensive teamwork through the EU-funded research project “Engineering Proprioception in Computing Systems (EPiCS)”. Eight research groups in five different countries collaborated to develop concepts and foundations for self-awareness and self-expression in computing systems and tested and demonstrated their usefulness in highly different domains. While the book includes some of the research results, it also, and more importantly, serves the purpose of stimulating further research into the field of self-aware and self-expressive computing systems.

Part I

Concepts and Fundamentals

Part I motivates the concepts of self-awareness and self-expression for engineering computing systems by looking into other disciplines and related concepts. It introduces a reference architecture for describing and engineering computational self-awareness and self-expression in computing systems. The architecture provides a common language which paves a way for identifying architectural patterns influencing the engineering of computational self-awareness and self-expression capabilities across a range of applications. Chapter 2 translates concepts from psychology to the domain of computing, introducing key ideas in self-aware computing. Chapter 3 relates our concepts of computational self-awareness and self-expression to other efforts in computer science and engineering under the self-awareness label. Depending on the fields, the term self-awareness may have different meanings. Chapter 4 concludes the first part by presenting our reference architecture for describing self-aware and self-expressive computing systems.

Chapter 2

Self-awareness and Self-expression: Inspiration from Psychology

Peter R. Lewis, Arjun Chandra, and Kyrre Glette

Abstract Self-awareness concepts from psychology are inspiring new approaches for engineering computing systems which operate in complex dynamic environments. There has been a broad and long-standing interest in self-awareness for computing, but only recently has a systematic understanding of self-awareness and how it can be used and evaluated been developed. In this chapter, we take inspiration from human self-awareness to develop new notions of computational self-awareness and self-expression. We translate concepts from psychology to the domain of computing, introducing key ideas in self-aware computing. In doing so, this chapter therefore paves the way for subsequent work in this book.

2.1 Introduction to Self-awareness

The *Oxford English Dictionary* defines *awareness* as “knowledge or perception of a situation or fact.” Informally, we might typically consider that humans build up knowledge, or become aware of things, by perceiving the world around them. We observe interactions, listen to other people, watch television, read books, and, particularly in early life, learn through play. When considering awareness in humans, it is common to consider that all the knowledge we possess, all of our awareness, is acquired through perception. This idea was first postulated by Hume [187], who argued that all human knowledge is induced from experience. What then does it mean for a human to be *self-aware*? For Hume, the “self” is not a defined physical entity, but instead describes the bundle of experiences or perceptions unique to an

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

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

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

individual. A Humean form of self-awareness might then be considered to consist of an individual's knowledge of its experiences. Kant [210] criticised Hume's view, extending the scope of the self significantly, arguing that there is some entity which is the subject of these experiences, and is common through space and time. This Kantian self synthesises information from experiences with concepts held in the mind and with the imagination. Kant further argued that as an individual performs actions within the world, since its actions are based on its synthesised knowledge, they represent its self, giving rise to the self also as an object. This object in turn is something which can be perceived and experienced.

Though there is a long history of analysis of the nature of the self in philosophy, more recently, psychology has made a more pragmatic attempt to develop an understanding of the varieties of knowledge individuals possess concerning themselves. The notion of self-awareness first appears in the literature around the turn of the twentieth century [25, 382], perhaps most importantly with James [197] making the distinction between two forms of self based on the differences between the Humean and Kantian views described above. First, the *implicit* self, often referred to as the self-as-subject, or the "I", is the self which is the subject of experiences. These experiences are unique to the individual, and they are from the individual's own point of view, determined by factors such as their sensing apparatus, their situation within the world, and other factors associated with their own state. Second, the *explicit* self, or self-as-object, can be discerned. Here the self is an object of knowledge. It is a thing which can be recognised, modelled and reasoned about, including in relation to other objects in the world. An individual's awareness of its explicit self is often considered the more advanced form of self-awareness in this distinction, building on implicit self-awareness. Indeed, implicit self-awareness emerges much earlier in the lives of human infants than its explicit counterpart does [231].

One commonly considered form of self-awareness is that as measured by the so-called *mirror test* [140]. A subject being evaluated is presented with a mirror, to which it is then allowed to get accustomed. The subject is then distracted and, without its knowledge, a visible change is made to its appearance. This is usually done by marking its face, e.g., putting a spot on its cheek or forehead. The subject is then presented with the mirror again. Any behaviour directed towards this marker by the subject implies self-recognition, which is seen as being enabled by a mental representation of oneself (also known as a secondary representation). As Asendorpf et al. [18] put it:

"[secondary representation] is not a perception of oneself but rather a constructed mental model of oneself that can be manipulated in fantasy. Therefore, the ability to recognise oneself in a mirror that requires linking a mirror image (a primary representation) with one's self marks the capacity for secondary representation."

Explicit self-awareness requires a subject to possess the capacity to construct such a secondary conceptual representation of itself. What then does the mirror test tell us about self-awareness? Humans, primates and some other animals have "passed" the mirror test [140, 18], however, Haikonen [155] showed that very little sophistication in computing machinery can enable a computational system with visual sensors to also pass. Haikonen therefore goes on to suggest that the ability or

2.2.3 *Self-awareness in Collective Systems*

So far, we have just considered self-awareness in the context of a single individual. However, Mitchell [272] notes that self-awareness can also be observed in collective systems, where there is no central point at which such self-knowledge is located. Examples of these collective systems include those comprised of individuals that might normally be considered either organisms in their own right (e.g., ants in a colony) or constituent cells of a larger organism (e.g., neurons in the brain). In these cases, it appears from an external perspective that such biological collective systems are self-aware at the level of the collective, even though this property may not be present at the level of the individual component. This awareness, Mitchell describes as being concerned with

*“information about the global state of the system, which feeds back to adaptively control the actions of the system’s low-level components. This **information about the global state is distributed** and statistical in nature, and thus is difficult for observers to tease out. However, the system’s components are able, collectively, to use this information in such a way that **the entire system appears** to have a coherent and useful sense of its own state.” [272]*

We have added the emphasis here, in order to highlight that a system which behaves as if it were self-aware is not necessarily required to possess a single “mind-like” component¹. Indeed, in many cases, the entire system appears self-aware, despite only local knowledge being present at constituent parts of the collective. Self-awareness might be considered the product of emergence.

This is a key observation which can contribute to the design of self-aware systems: one need not require that such a system possess a global omniscient controller. Indeed, many natural systems appear to have been favoured by evolution which do not have such a central point of control, and rely upon relevant knowledge being available at appropriate locations within the system. It is highly likely that this idea can improve the robustness and adaptability of such systems; these are desirable properties for natural and artificial systems alike.

2.2.4 *Self-expression*

As we have seen, self-awareness is concerned with knowledge synthesised and held by an individual about itself and its experiences. This knowledge may be centrally held, or else distributed in nature. However, in studying the self-awareness properties of natural and computational systems, we have found it advantageous to explicitly and separately consider the related process of an individual determining its behaviour as a result of this knowledge. This process we call self-expression. In

¹ Indeed, while the brain has long been known to be a collective system composed of neurons, consciousness in the human mind is itself thought by some [98] to also be a distributed phenomenon, with nothing like what we might call global knowledge.

- **Private self-awareness:** This is a system's ability to obtain knowledge based on phenomena that are internal to itself. A system needs internal sensors to achieve this.
- **Public self-awareness:** This is 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.

Some prior work in self-aware computing has considered only private self-awareness, i.e., a system's awareness of its own internal state. However, the importance of the availability of external sources of information to self-awareness processes should be emphasised: self-awareness is not only concerned with sources of information internal to the individual. We argue that a full consideration of computational self-awareness also includes public aspects, where a system's knowledge of itself in relation to its social and physical environment can be synthesised with private self-awareness, in order to produce integrated conceptual models. In this book, we hope to demonstrate that the distinction, inclusion and synthesis of both public and private self-awareness raise many important questions for engineers of self-aware computing systems.

2.3.2 Levels of Computational Self-awareness

We now describe our computational framing of the levels of self-awareness, which were first presented by Faniyi et al. [127] and elaborated upon by Lewis et al. [236]. It is possible to relate the levels of self-awareness to the concepts of private and public self-awareness, and hence the sources of the relevant knowledge (i.e., based on internal or external sensors). The relevance of each level to these concepts is also described. The relationship provides an indication of the architecture that will be required in order to realise each of the levels. In each case, Neisser's level of self-awareness is given on the left, and our corresponding level of computational self-awareness is on the right.

1. Ecological self → Stimulus-aware

A system is stimulus-aware if it is able to obtain knowledge of stimuli acting upon it, enabling the ability to respond to events. It does not have knowledge of past/future stimuli. Since stimuli may originate both internally and externally, stimulus-awareness can be **private**, **public** or **both**.

2. Interpersonal self → Interaction-aware

A system is interaction-aware if it can obtain knowledge that stimuli and its own actions constitute interactions with other systems and the environment. It is able to obtain knowledge via feedback loops that its actions can provoke or cause