

Lecture Notes in Artificial Intelligence 7716

Subseries of Lecture Notes in Computer Science

LNAI Series Editors

Randy Goebel

University of Alberta, Edmonton, Canada

Yuzuru Tanaka

Hokkaido University, Sapporo, Japan

Wolfgang Wahlster

DFKI and Saarland University, Saarbrücken, Germany

LNAI Founding Series Editor

Joerg Siekmann

DFKI and Saarland University, Saarbrücken, Germany

Joscha Bach
Ben Goertzel
Matthew Iklé (Eds.)

Artificial General Intelligence

5th International Conference, AGI 2012
Oxford, UK, December 8-11, 2012
Proceedings

Series Editors

Randy Goebel, University of Alberta, Edmonton, Canada
Jörg Siekmann, University of Saarland, Saarbrücken, Germany
Wolfgang Wahlster, DFKI and University of Saarland, Saarbrücken, Germany

Volume Editors

Joscha Bach
Humboldt Universität Berlin
Raumerstr. 11
10437 Berlin, Germany
E-mail: joscha.bach@gmail.com

Ben Goertzel
Aidyia Ltd., Unit 612, 6/F
Lu Plaza, 2 Wing Yip Street
Kwun Tong, Hong Kong
E-mail: ben@goertzel.org

Matthew Iklé
Adams State University, Suite 3060
Alamosa, CO 81101, USA
E-mail: moikle@adams.edu

ISSN 0302-9743 e-ISSN 1611-3349
ISBN 978-3-642-35505-9 e-ISBN 978-3-642-35506-6
DOI 10.1007/978-3-642-35506-6
Springer Heidelberg Dordrecht London New York

Library of Congress Control Number: 2012953483

CR Subject Classification (1998): I.2, F.4.1, I.5, F.1-2, D.2

LNCS Sublibrary: SL 7 – Artificial Intelligence

© Springer-Verlag Berlin Heidelberg 2012

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, re-use of illustrations, recitation, broadcasting, reproduction on microfilms or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer. Violations are liable to prosecution under the German Copyright Law.

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

Typesetting: Camera-ready by author, data conversion by Scientific Publishing Services, Chennai, India

Printed on acid-free paper

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

Preface

This volume collects the research papers contributed to AGI-12, the 5th Conference on Artificial General Intelligence.

The AGI conference series is the premier international forum for cutting-edge research and thinking regarding the original goal of the AI field — the creation of thinking machines. AGI—artificial general intelligence—refers to generally intelligent capabilities at the human level and ultimately beyond. Like its predecessors, AGI-12 brought together researchers in AGI and related disciplines, to present and discuss the current state of approaches, architectures, algorithms, and ideas relevant to the advancement of AGI. In honor of the Alan Turing centenary year 2012, this was the first AGI conference to be held in the UK. The conference took place at the University of Oxford, St. Anne’s College, December 8–11, 2012. It was hosted by the “Future of Humanity Institute” at Oxford University through its “Programme on the Impacts of Future Technology”.

A total of 80 contributed articles were submitted to AGI-12, of which 34 (42.5%) were accepted. Contributions covered a wide array of AGI research and development aspects, with the key proviso that each paper should somehow contribute specifically to the development of AGI. As in previous years, we had a host of papers covering practical proto-AGI software systems and architectures, as well as papers on the mathematical theory of AGI, and connections between AGI and neuroscience and/or cognitive science.

The AGI-12 conference program also included four invited keynote lectures. David Hanson, an American robotics designer and researcher, responsible for the creation of a series of realistic humanoid robots and founder of Hanson Robotics, delivered a talk on “Humanoid Robotics and AGI.” Angelo Cangelosi, Professor of Artificial Intelligence and Cognition and Director of the Centre for Robotics and Neural Systems at Plymouth University, UK, discussed “Cognitive Robotics.” Margaret Boden, a researcher in the fields of artificial intelligence, psychology, philosophy, cognitive and computer science, and Research Professor of Cognitive Science at the Department of Informatics at the University of Sussex, talked about “Creativity and AGI.” Finally, Nick Bostrom, Professor of Philosophy at Oxford, Director of the Future of Humanity Institute, and Director of the Programme on the Impacts of Future Technology, discussed the future evolution of advanced AGIs and the dynamics of AGI goal systems.

AGI-12 was held together with the first conference on AGI Safety and Impacts. While AGI-12 focused, like all the AGI conferences, on the technical business of designing and building AGI systems, AGI-Impacts pursued related issues regarding the potential future of AGI: What will be the impacts of AGI on the world? Which directions of research should be most enthusiastically explored,

and which should be de-emphasized or avoided? What can we say now about the future impact of AGI, and how should we act in consequence? AGI-Impacts papers are not included here, but abstracts may be found on the AGI-Impacts website.

In both the contributed articles and the invited keynotes, and the collaboration between AGI-12 and AGI-Impacts, we see the cross-disciplinarity and diversity that make the AGI field so fascinating at this stage of its development. As our technology and understanding progress, significant advances in the creation of AGI systems become viable, which is reflected in a large diversity of approaches by an increasing number of researchers in AI and related fields. The AGI conference series embraces this diversity and the creative inventiveness that it fosters.

Producing a conference of such high quality was made possible only through the support of a large community of volunteers. We thank the local Organizing Committee members for all of their advice and help in preparations and arrangements. We thank all the Program Committee members for their dedicated service to the review process. We especially thank all of our contributors, participants, and keynote speakers. The presentations, demos, and tutorials ultimately provide the material for generating thoughtful, interesting, and stimulating discussions toward the ultimate goal of achieving AGI.

Finally, we honor the support of all our sponsors: Oxford University, Future of Humanity Institute; Kurzweil AI; Rick Schwall; and Novamente LLC.

October 2012

Joscha Bach
Ben Goertzel
Matt Iklé

Organization

Program Committee

Sam Adams	IBM Research
Itamar Arel	University of Tennessee
Joscha Bach	Humboldt-University of Berlin
Sarah Bull	National ICT Australia Ltd
Cristiano Castelfranchi	Institute of Cognitive Sciences and Technologies
Antonio Chella	Università di Palermo
Jonathan Connell	IBM Research
David A. Dalrymple	MIT Media Lab
Stan Franklin	University of Memphis
Deon Garrett	IIIM, University of Reykjavik
Nil Geisweiller	Novamente LLC
Ben Goertzel	Novamente LLC
J. Storrs Hall	Institute for Molecular Manufacturing/Foresight Institute
Louie Helm	Singularity Institute
Eva Hudlicka	Psychometrix Associates
Marcus Hutter	Australian National University
Matt Iklé	Adams State University
Benjamin Johnston	University of Technology, Sidney
Cliff Joslyn	Portland State University
Randal Koene	Boston University
Kai-Uwe Kühnberger	University of Osnabrück
Christian Lebiere	Carnegie Mellon University
Shane Legg	IDSIA
Moshe Looks	Google Inc.
Stacy Marsella	University of Southern California
Günter Palm	University of Ulm
Stephen Reed	texai.org
Paul Rosenbloom	University of Southern California
Ute Schmid	University of Bamberg
Juergen Schmidhuber	IDSIA
Zhongzhi Shi	Chinese Academy of Sciences
Bas Steunebrink	IDSIA
Rich Sutton	University of Alberta

VIII Organization

Kristinn Thorisson

Julian Togelius

Mario Verdicchio

Pei Wang

CADIA, Reykjavik University

IT University of Copenhagen

Università degli Studi di Bergamo

Temple University

Additional Reviewers

Hadi Afshar

Alex Altair

Tarek R. Besold

Tor Lattimore

Javier Snaider

Table of Contents

Creativity, Cognitive Mechanisms, and Logic	1
<i>Ahmed M.H. Abdel-Fattah, Tarek Besold, and Kai-Uwe Kühnberger</i>	
MicroPsi 2: The Next Generation of the MicroPsi Framework	11
<i>Joscha Bach</i>	
An Extensible Language Interface for Robot Manipulation	21
<i>Jonathan Connell, Etienne Marcheret, Sharath Pankanti, Michiharu Kudoh, and Risa Nishiyama</i>	
Noisy Reasoners: Errors of Judgement in Humans and AIs	31
<i>Fintan Costello</i>	
2012: The Connectome, WBE and AGI	41
<i>Diana Deca</i>	
Logical Prior Probability	50
<i>Abram Demski</i>	
A Representation Theorem for Decisions about Causal Models	60
<i>Daniel Dewey</i>	
Modular Value Iteration through Regional Decomposition	69
<i>Linus Gisslen, Mark Ring, Matthew Luciw, and Jürgen Schmidhuber</i>	
Perception Processing for General Intelligence: Bridging the Symbolic/Subsymbolic Gap	79
<i>Ben Goertzel</i>	
On Attention Mechanisms for AGI Architectures: A Design Proposal ...	89
<i>Helgi Páll Helgason, Eric Nivel, and Kristinn R. Thórisson</i>	
A Framework for Representing Action Meaning in Artificial Systems via Force Dimensions	99
<i>Paul Hemeren and Peter Gärdenfors</i>	
Avoiding Unintended AI Behaviors	107
<i>Bill Hibbard</i>	
Decision Support for Safe AI Design	117
<i>Bill Hibbard</i>	

On Measuring Social Intelligence: Experiments on Competition and Cooperation	126
<i>Javier Insa-Cabrera, José-Luis Benacloch-Ayuso, and José Hernández-Orallo</i>	
Toward Tractable AGI: Challenges for System Identification in Neural Circuitry	136
<i>Randal A. Koene</i>	
CHREST Models of Implicit Learning and Board Game Interpretation	148
<i>Peter Lane and Fernand Gobet</i>	
Syntax-Semantic Mapping for General Intelligence: Language Comprehension as Hypergraph Homomorphism, Language Generation as Constraint Satisfaction	158
<i>Ruiting Lian, Ben Goertzel, Shujing Ke, Jade O'Neill, Keyvan Sadeghi, Simon Shiu, Dingjie Wang, Oliver Watkins, and Gino Yu</i>	
An Intelligent Theory of Cost for Partial Metric Spaces	168
<i>Steve Matthews and Michael Bukatin</i>	
Fractal Analogies for General Intelligence	177
<i>Keith McGreggor and Ashok Goel</i>	
Pattern Mining for General Intelligence: The FISHGRAM Algorithm for Frequent and Interesting Subhypergraph Mining	189
<i>Jade O'Neill, Ben Goertzel, Shujing Ke, Ruiting Lian, Keyvan Sadeghi, Simon Shiu, Dingjie Wang, and Gino Yu</i>	
Pursuing Artificial General Intelligence by Leveraging the Knowledge Capabilities of ACT-R	199
<i>Alessandro Oltramari and Christian Lebiere</i>	
Space-Time Embedded Intelligence	209
<i>Laurent Orseau and Mark Ring</i>	
Memory Issues of Intelligent Agents	219
<i>Laurent Orseau and Mark Ring</i>	
What Is It Like to Be a Brain Simulation?	232
<i>Eray Özkural</i>	
Extending Universal Intelligence Models with Formal Notion of Representation	242
<i>Alexey Potapov and Sergey Rodionov</i>	

Differences between Kolmogorov Complexity and Solomonoff Probability: Consequences for AGI	252
<i>Alexey Potapov, Andrew Svitenkov, and Yurii Vinogradov</i>	
Deconstructing Reinforcement Learning in Sigma	262
<i>Paul S. Rosenbloom</i>	
Extending Mental Imagery in Sigma	272
<i>Paul S. Rosenbloom</i>	
Binary Space Partitioning as Intrinsic Reward	282
<i>Wojciech Skaba</i>	
Perceptual Time, Perceptual Reality, and General Intelligence	292
<i>Leslie S. Smith</i>	
Transparent Neural Networks: Integrating Concept Formation and Reasoning	302
<i>Claes Strannegård, Olle Häggström, Johan Wessberg, and Christian Balkenius</i>	
Optimistic AIXI	312
<i>Peter Sunehag and Marcus Hutter</i>	
On the Functional Contributions of Emotion Mechanisms to (Artificial) Cognition and Intelligence	322
<i>Serge Thill and Robert Lowe</i>	
Stupidity and the Ouroboros Model	332
<i>Knud Thomsen</i>	
On Ensemble Techniques for AIXI Approximation	341
<i>Joel Veness, Peter Sunehag, and Marcus Hutter</i>	
Motivation Management in AGI Systems	352
<i>Pei Wang</i>	
Hippocampal Formation Mechanism Will Inspire Frame Generation for Building an Artificial General Intelligence	362
<i>Hiroshi Yamakawa</i>	
Fuzzy-Probabilistic Logic for Common Sense	372
<i>King-Yin Yan</i>	
Author Index	381