

Contents

Self-Modification of Policy and Utility Function in Rational Agents	1
<i>Tom Everitt, Daniel Filan, Mayank Daswani, and Marcus Hutter</i>	
Avoiding Wireheading with Value Reinforcement Learning	12
<i>Tom Everitt and Marcus Hutter</i>	
Death and Suicide in Universal Artificial Intelligence	23
<i>Jarryd Martin, Tom Everitt, and Marcus Hutter</i>	
Ultimate Intelligence Part II: Physical Complexity and Limits of Inductive Inference Systems	33
<i>Eray Özkural</i>	
Open-Ended Intelligence: On the Role of Individuation in AGI.	43
<i>David (Weaver) Weinbaum and Viktoras Veitas</i>	
The AGI Containment Problem.	53
<i>James Babcock, János Kramár, and Roman Yampolskiy</i>	
Imitation Learning as Cause-Effect Reasoning	64
<i>Garrett Katz, Di-Wei Huang, Rodolphe Gentili, and James Reggia</i>	
Some Theorems on Incremental Compression.	74
<i>Arthur Franz</i>	
Rethinking Sigma’s Graphical Architecture: An Extension to Neural Networks	84
<i>Paul S. Rosenbloom, Abram Demski, and Volkan Ustun</i>	
Real-Time GA-Based Probabilistic Programming in Application to Robot Control	95
<i>Alexey Potapov, Sergey Rodionov, and Vita Potapova</i>	
About Understanding.	106
<i>Kristinn R. Thórisson, David Kremelberg, Bas R. Steunebrink, and Eric Nive</i>	
Why Artificial Intelligence Needs a Task Theory: And What It Might Look Like	118
<i>Kristinn R. Thórisson, Jordi Bieger, Thröstur Thorarensen, Jóna S. Sigurðardóttir, and Bas R. Steunebrink</i>	

Growing Recursive Self-Improvers	129
<i>Bas R. Steunebrink, Kristinn R. Thórisson, and Jürgen Schmidhuber</i>	
Different Conceptions of Learning: Function Approximation vs. Self-Organization	140
<i>Pei Wang and Xiang Li</i>	
The Emotional Mechanisms in NARS	150
<i>Pei Wang, Max Talanov, and Patrick Hammer</i>	
The OpenNARS Implementation of the Non-Axiomatic Reasoning System	160
<i>Patrick Hammer, Tony Lofthouse, and Pei Wang</i>	
Integrating Symbolic and Sub-symbolic Reasoning	171
<i>Claes Strannegård and Abdul Rahim Nizamani</i>	
Integrating Axiomatic and Analogical Reasoning	181
<i>Claes Strannegård, Abdul Rahim Nizamani, and Ulf Persson</i>	
Embracing Inference as Action: A Step Towards Human-Level Reasoning . . .	192
<i>John Licato and Maxwell Fowler</i>	
Asymptotic Logical Uncertainty and the Benford Test	202
<i>Scott Garrabrant, Tsvi Benson-Tilsen, Siddharth Bhaskar, Abram Demski, Joanna Garrabrant, George Koleszarik, and Evan Lloyd</i>	
Towards a Computational Framework for Function-Driven Concept Invention	212
<i>Nico Potyka, Danny Gómez-Ramírez, and Kai-Uwe Kühnberger</i>	
System Induction Games and Cognitive Modeling as an AGI Methodology	223
<i>Sean Markan</i>	
Integrating Model-Based Prediction and Facial Expressions in the Perception of Emotion	234
<i>Nutchanon Yongsatianchot and Stacy Marsella</i>	
A Few Notes on Multiple Theories and Conceptual Jump Size	244
<i>Grace Solomonoff</i>	
Generalized Temporal Induction with Temporal Concepts in a Non-axiomatic Reasoning System	254
<i>Tony Lofthouse and Patrick Hammer</i>	

Introspective Agents: Confidence Measures for General Value Functions	258
<i>Craig Sherstan, Adam White, Marlos C. Machado, and Patrick M. Pilarski</i>	
Automatic Sampler Discovery via Probabilistic Programming and Approximate Bayesian Computation	262
<i>Yura Perov and Frank Wood</i>	
How Much Computation and Distributedness is Needed in Sequence Learning Tasks?	274
<i>Mrwan Margem and Ozgur Yilmaz</i>	
Analysis of Algorithms and Partial Algorithms	284
<i>Andrew MacFie</i>	
Estimating Cartesian Compression via Deep Learning	294
<i>András Lőrincz, András Sárkány, Zoltán Á. Milacski, and Zoltán Tőssér</i>	
A Methodology for the Assessment of AI Consciousness	305
<i>Harry H. Porter III</i>	
Toward Human-Level Massively-Parallel Neural Networks with Hodgkin-Huxley Neurons.	314
<i>Lyle N. Long</i>	
Modeling Neuromodulation as a Framework to Integrate Uncertainty in General Cognitive Architectures	324
<i>Frédéric Alexandre and Maxime Carrere</i>	
Controlling Combinatorial Explosion in Inference via Synergy with Nonlinear-Dynamical Attention Allocation	334
<i>Ben Goertzel, Misgana Bayetta Belachew, Matthew Ikle', and Gino Yu</i>	
Probabilistic Growth and Mining of Combinations: A Unifying Meta-Algorithm for Practical General Intelligence	344
<i>Ben Goertzel</i>	
Ideas for a Reinforcement Learning Algorithm that Learns Programs.	354
<i>Susumu Katayama</i>	
Author Index	363

Artificial General Intelligence

9th International Conference, AGI 2016, New York, NY,

USA, July 16-19, 2016, Proceedings

Steunebrink, B.; Wang, P.; Goertzel, B. (Eds.)

2016, XI, 364 p. 55 illus., Softcover

ISBN: 978-3-319-41648-9