Artificial General Intelligence

Third Conference on Artificial General Intelligence
AGI 2010
Lugano, Switzerland
March 5-8, 2010

Editors
Eric Baum, Marcus Hutter,
Emanuel Kitzelmann

ATLANTIS PRESS

Eric Baum, Marcus Hutter, Emanuel Kitzelmann (Editors)

Artificial General Intelligence

Proceedings of the Third Conference on Artificial General Intelligence, AGI 2010, Lugano, Switzerland, March 5-8, 2010

Amsterdam-Beijing-Paris

ISBN: 978-90-78677-36-9

Advances in Intelligent Systems Research

volume 10

In Memoriam
Ray Solomonoff (1926-2009)

The Great Ray Solomonoff, pioneer of Machine Learning, founder of Algorithmic Probability Theory, father of the Universal Probability Distribution, creator of the Universal Theory of Inductive Inference, passed away on Monday 30 November 2009, from complications in the wake of a broken aneurism in his head. He is survived by his loving wife, Grace.

Ray Solomonoff was the first to describe the fundamental concept of Algorithmic Information or Kolmogorov Complexity, and the first to prove the celebrated Invariance Theorem. In the new millennium his work became the foundation of the first mathematical theory of Optimal Universal Artificial Intelligence.

Ray intended to deliver an invited lecture at AGI 2010, the Conference on Artificial General Intelligence (March 5-8, 2010) in Lugano (where he already spent time in 2001 as a visiting professor at the Swiss AI Lab IDSIA). The AGI conference series would not even exist without his essential theoretical contributions. With great sadness AGI 2010 was held «In Memoriam Ray Solomonoff».

Ray will live on in the many minds shaped by his revolutionary ideas.

Eulogy by Jürgen Schmidhuber

This book is part of the series ***Advances in Intelligent Systems Research*** (ISSN: 1951-6851) published by Atlantis Press.

Aims and scope of the series
During the past decade computer science research in understanding and reproducing human intelligence has expanded from the more traditional approaches like psychology, logics and artificial intelligence into multiple other areas, including neuroscience research. Moreover, new results in biology, chemistry, (surface) physics and gene technology, but also in network technology are greatly affecting current research in computer science, including the development of intelligent systems. At the same time, computer science's new results are increasingly being applied in these fields allowing for important cross-fertilisations. This series aims at publishing proceedings from all disciplines dealing with and affecting the issue of understanding and reproducing intelligence in artificial systems. Also, the series is open for publications concerning the application of intelligence in networked or any other environment and the extraction of meaningful data from large data sets.

Research fields covered by the series include: * Fuzzy sets * Machine learning * Autonomous agents * Evolutionary systems * Robotics and autonomous systems * Semantic web, incl. web services, ontologies and grid computing * Biological systems * Artificial Intelligence, incl. knowledge representation, logics * Neural networks * Constraint satisfaction * Computational biology * Information sciences * Computer vision, pattern recognition * Computational neuroscience * Datamining, knowledge discovery and modelling for e.g. life sciences.

© ATLANTIS PRESS, 2010
http://www.atlantis-press.com

ISBN: 978-90-78677-36-9

This book is published by Atlantis Press, scientific publishing, Paris, France.

All rights reserved for this book. No part of this book may be reproduced, translated, stored or transmitted in any form or by any means, including electronic, mechanical, photocopying, recording or otherwise, without prior permission from the publisher.

Atlantis Press adheres to the creative commons policy, which means that authors retain the copyright of their article.

Artificial General Intelligence

Volume Editors

Marcus Hutter
CSL@RSISE and SML@NICTA
Australian National University
Room B259, Building 115
Corner of North and Daley Road
Canberra ACT 0200, Australia
Email: marcus.hutter@gmx.net
WWW: http://www.hutter1.net

Eric B. Baum
Azure Sky Research Inc.
386 Riverside Drive
Princeton NJ 08540 USA
Email: ebaum@fastmail.fm
WWW: http://whatisthought.com/

Emanuel Kitzelmann
Cognitive Systems Group, WIAI
University of Bamberg
Feldkirchenstrasse 21
D-96052 Bamberg, Germany
Email: emanuel.kitzelmann@uni-bamberg.de
WWW: http://www.uni-bamberg.de/kogsys/members/kitzelmann/

Preface

Artificial General Intelligence (AGI) research focuses on the original and ultimate goal of AI - to create broad human-like and transhuman intelligence, by exploring all available paths, including theoretical and experimental computer science, cognitive science, neuroscience, and innovative interdisciplinary methodologies. Due to the difficulty of this task, for the last few decades the majority of AI researchers have focused on what has been called *narrow AI* - the production of AI systems displaying intelligence regarding specific, highly constrained tasks. In recent years, however, more and more researchers have recognized the necessity - and feasibility - of returning to the original goals of the field. Increasingly, there is a call for a transition back to confronting the more difficult issues of *human level intelligence* and more broadly *artificial general intelligence*.

The Conference on Artificial General Intelligence is the only major conference series devoted wholly and specifically to the creation of AI systems possessing general intelligence at the human level and ultimately beyond. Its third installation, AGI-10, in Lugano, Switzerland, March 5-8, 2010, attracted 66 paper submissions. Of these submissions, 29 (i.e., 44%) were accepted as full papers, additional 12 were accepted as short position papers, which was a more selective choice than for AGI-09 in Arlington, Virginia. Both full and short papers are included in this collection. The papers presented at the conference and collected in these proceedings address a wide range of AGI-related topics such as Universal Search, Cognitive and AGI Architectures, Adaptive Agents, special aspects of reasoning, the formalization of AGI, the evaluation of AGI systems, machine learning for AGI, and implications of AGI. The contributions range from proven theoretical results to system descriptions, implementations, and experiments to general ideas and visions.

The conference program also included a keynote address by Richard Sutton and an invited lecture by Randal Koene. Richard Sutton is a Professor at the University of Alberta. The co-author of the textbook «Reinforcement Learning: an Introduction», has made numerous contributions to the fields of AI and learning. His talk was on *Reducing Knowledge to Prediction*. The idea is to formalize and reduce knowledge about the world to predictive statements of a particular form that is particularly well suited for learning and reasoning. He presented new learning algorithms in this framework that his research group has developed. Randal Koene is the Director of the Department of Neuroengineering at Tecnalia. His talk was on *Whole Brain Emulation: Issues of scope and resolution, and the need for new methods of in-vivo recording*.

Finally, the conference program included a number of workshops on topics such as formalizing AGI and the future of AI, pre-conference tutorials on various AGI-related topics, and an AGI-systems demonstration.

Producing such a highly profiled program would not have been possible without the support of the community. We thank the (local) organizing committee members for their advise and their help in all matters of actually preparing and running the event. We thank the program committee members for a smooth review process and for the high quality of the reviews - despite the tight review phase and the fact that due to the high number of submissions the review load per PC member was considerably higher than originally expected. And we thank all participants for submitting and presenting interesting and stimulating work, which is the key ingredient needed for a successful conference. We also gratefully acknowledge the support of a number of sponsors:

- Association for the Advancement of Artificial Intelligence (AAAI)
- KurzweilAI.net (Kurzweil Best Paper Prize and Kurzweil Best Idea Prize)
- The University of Lugano, Faculty of Informatics (scientifically endorses AGI-10)

March 2010

Marcus Hutter (Conference Chair)

Eric Baum, Emanuel Kitzelmann (Program Committee Chairs)

Conference Organization

Chairs

Marcus Hutter (Conference Chair)	Australian National University, Australia
Jürgen Schmidhuber (Local Chair)	IDSIA, Switzerland
Eric Baum (Program Chair)	Azure Sky Research Inc., USA
Emanuel Kitzelmann (Program Chair)	University of Bamberg, Germany

Program Committee

Igor Aleksander	Imperial College London, UK
Lee Altenberg	University of Hawaii at Manoa, USA
Itamar Arel	University of Tennessee, Knoxville, USA
Sebastian Bader	Rostock University, Germany
Eric Baum	Azure Sky Research Inc., USA
Anselm Blumer	Tufts University, USA
Hugo de Garis	Xiamen University, China
Wlodek Duch	Nicolaus Copernicus University, Poland
Artur Garcez	City University London, UK
J. Storrs	Hall Institute for Molecular Manufacturing, USA
Benjamin Johnston	University of Technology, Sydney, Australia
Bert Kappen	Radboud University, The Netherlands
Emanuel Kitzelmann	University of Bamberg, Germany
Kai-Uwe Kühnberger	University of Osnabrück, Germany
Christian Lebiere	Carnegie Mellon University, USA
Shane Legg	University College London, UK
Moshe Looks	Google Research, USA
András Lörincz	Eötvös Loránd University, Hungary
Hassan Mahmud	Australian National University, Australia
Eric Nivel	Reykjavik University, Iceland
Jan Poland	ABB Research, Zurich, Switzerland
Brandon Rohrer	Sandia National Laboratory, USA
Sebastian Rudolph	University of Karlsruhe, Germany
Robert Schapire	Princeton University, USA
Lokendra Shastri	Infosys Technologies Ltd, India
Ray Solomono	Oxbridge Research, Cambridge, USA
Rich Sutton	University of Alberta, Canada
Kristinn Thorisson	Reykjavik University, Iceland

Lyle Ungar	University of Pennsylvania, USA
Les Valiant	Harvard University, USA
Marco Wiering	University of Groningen, The Netherlands
Mary-Anne Williams	University of Technology, Sydney, Australia
David Wolpert	NASA, USA

Organizing Committee

Tsvi Achler	University of Illinois at Urbana Champaign, USA
Eric Baum	Azure Sky Research Inc., USA
Sarah Bull	NICTA, Australia
Ben Goertzel	Novamente LLC, USA
Marcus Hutter	Australian National University, Australia
Emanuel Kitzelmann	University of Bamberg, Germany
David Orban	Singularity University, USA
Stephen Reed	Texai.org, USA

Local Organizing Committee

Carlo Lepori	CFO of IDSIA
Mauro Pezze	Dean of the Faculty of Informatics of USI
Jürgen Schmidhuber	IDSIA
Albino Zgraggen	CFO of the University of Lugano/USI

External Reviewers

Subhadip Bandyopadhyay

Arijit Laha

Srinivas Narasimhamurthy

Bintu Vasudevan

Table of Contents

Full Articles.

Efficient Constraint-Satisfaction in Domains with Time 1
Perrin Bignoli, Nicholas Cassimatis, Arthi Murugesan

The CHREST Architecture of Cognition: The Role of Perception in General Intelligence . 7
Fernand Gobet, Peter Lane

A General Intelligence Oriented Architecture for Embodied Natural Language Processing 13
Ben Goertzel, Cassio Pennachin, Samir Araujo, Ruiting Lian, Fabricio Silva, Murilo Queiroz, Welter Silva, Mike Ross, Linas Vepstas, Andre Senna

Toward a Formal Characterization of Real-World General Intelligence 19
Ben Goertzel

On Evaluating Agent Performance in a Fixed Period of Time 25
Jose Hernandez-Orallo

Artificial General Segmentation 31
Daniel Hewlett, Paul Cohen

Grounding Possible Worlds Semantics in Experiential Semantics 37
Matthew Ikle, Ben Goertzel

The Toy Box Problem (and a Preliminary Solution) 43
Benjamin Johnston

Playing General Structure Rewriting Games 49
Lukasz Kaiser, Lukasz Stafiniak

Towards Automated Code Generation for Autonomous Mobile Robots 55
Dermot Kerr, Ulrich Nehmzow, Stephen A. Billings

Searching for Minimal Neural Networks in Fourier Space 61
Jan Koutnik, Faustino Gomez, Juergen Schmidhuber

Remarks on the Meaning of Analogical Relations 67
Ulf Krumnack, Helmar Gust, Angela Schwering, Kai-Uwe Kuehnberger

Quantitative Spatial Reasoning for General Intelligence 73
Unmesh Kurup, Nicholas Cassimatis

Cognitive Architecture Requirements for Achieving AGI 79
John Laird, Robert Wray

Sketch of an AGI Architecture with Illustration 85
Andras Lorincz, Zoltan Bardosi, Daniel Takacs

GQ(λ): A General Gradient Algorithm for Temporal-Difference Prediction Learning with Eligibility Traces 91
Hamid Maei, Richard Sutton

A Generic Adaptive Agent Architecture Integrating Cognitive and Affective States and their Interaction 97
Zulfiqar A. Memon, Jan Treur

A Cognitive Architecture for Knowledge Exploitation 103
Gee Wah Ng, Yuan Sin Tan, Loo Nin Teow, Khin Hua Ng, Kheng Hwee Tan, Rui Zhong Chan

An Artificial Intelligence Model that Combines Spatial and Temporal Perception 109
Jianglong Nan, Fintan Costello

A Conversion Between Utility and Information 115
Pedro Ortega, Daniel Braun

A Bayesian Rule for Adaptive Control based on Causal Interventions 121
Pedro Ortega, Daniel Braun

Discovering and Characterizing Hidden Variables 127
Soumi Ray, Tim Oates

What we Might Look for in an AGI Benchmark 133
Brandon Rohrer

Towards Practical Universal Search 139
Tom Schaul, Juergen Schmidhuber

Artificial Scientists & Artists Based on the Formal Theory of Creativity 145
Juergen Schmidhuber

Algorithmic Probability, Heuristic Programming and AGI 151
Ray Solomonoff

Frontier Search 158
Yi Sun, Tobias Glasmachers, Tom Schaul, Juergen Schmidhuber

The Evaluation of AGI Systems 164
Pei Wang

Designing a Safe Motivational System for Intelligent Machines 170
Mark Waser

Position Statements.

Software Design of an AGI System Based on Perception Loop 176
Antonio Chella, Massimo Cossentino, Valeria Seidita

A Theoretical Framework to Formalize AGI-Hard Problems 178
Pedro Demasi, Jayme Szwarcfiter, Adriano Cruz

Uncertain Spatiotemporal Logic for General Intelligence 180
Nil Geisweiller, Ben Goertzel

A (hopefully) Unbiased Universal Environment Class for Measuring Intelligence of Biological and Artificial Systems 182
Jose Hernandez-Orallo

Neuroethological Approach to Understanding Intelligence 184
DaeEun Kim

Compression Progress, Pseudorandomness, and Hyperbolic Discounting 186
Moshe Looks

Relational Local Iterative Compression 188
Laurent Orseau

Stochastic Grammar Based Incremental Machine Learning Using Scheme 190
Eray Ozkural, Cevdet Aykanat

Compression-driven Progress in Science 192
Leo Pape

Concept Formation in the Ouroboros Model 194
Knud Thomsen

On Super-Turing Computing Power and Hierarchies of Artificial General Intelligence Systems 196
Jiri Wiedermann

A Minimum Relative Entropy Principle for AGI 198
Antoine van de Ven, Ben Schouten

Author Index **200**

Author Index

Araujo, Samir 13
Aykanat, Cevdet 190

Bardosi, Zoltan 85
Bignoli, Perrin 1
Billings, Stephen A. 55
Braun, Daniel 115, 121

Cassimatis, Nicholas 1, 73
Chan, Rui Zhong 103
Chella, Antonio 176
Cohen, Paul 31
Cossentino, Massimo 176
Costello, Fintan 109
Cruz, Adriano 178

Demasi, Pedro 178

Geisweiller, Nil 180
Glasmachers, Tobias 158
Gobet, Fernand 7
Goertzel, Ben 13, 19, 37, 180
Gomez, Faustino 61
Gust, Helmar 67

Hernandez-Orallo, Jose 25, 182
Hewlett, Daniel 31

Ikle, Matthew 37

Johnston, Benjamin 43

Kaiser, Lukasz 49
Kerr, Dermot 55
Kim, DaeEun 184
Koutnik, Jan 61
Krumnack, Ulf 67
Kuehnberger, Kai-Uwe 67
Kurup, Unmesh 73

Laird, John 79
Lane, Peter 7
Lian, Ruiting 13
Looks, Moshe 186
Lorincz, Andras 85

Maei, Hamid 91
Memon, Zulfiqar A. 97
Murugesan, Arthi 1

Nan, Jianglong 109
Nehmzow, Ulrich 55
Ng, Gee Wah 103
Ng, Khin Hua 103

Oates, Tim 127
Orseau, Laurent 188
Ortega, Pedro 115, 121
Ozkural, Eray 190

Pape, Leo 192
Pennachin, Cassio 13

Queiroz, Murilo 13

Ray, Soumi 127
Rohrer, Brandon 133
Ross, Mike 13

Schaul, Tom 139, 158
Schmidhuber, Juergen 61, 139, 145, 158
Schouten, Ben 198
Schwering, Angela 67
Seidita, Valeria 176
Senna, Andre 13
Silva, Fabricio 13
Silva, Welter 13
Solomonoff, Ray 151
Stafiniak, Lukasz 49
Sun, Yi 158
Sutton, Richard 91
Szwarcfiter, Jayme 178

Takacs, Daniel 85
Tan, Kheng Hwee 103
Tan, Yuan Sin 103
Teow, Loo Nin 103
Thomsen, Knud 194
Treur, Jan 97

van de Ven, Antoine 198
Vepstas, Linas 13
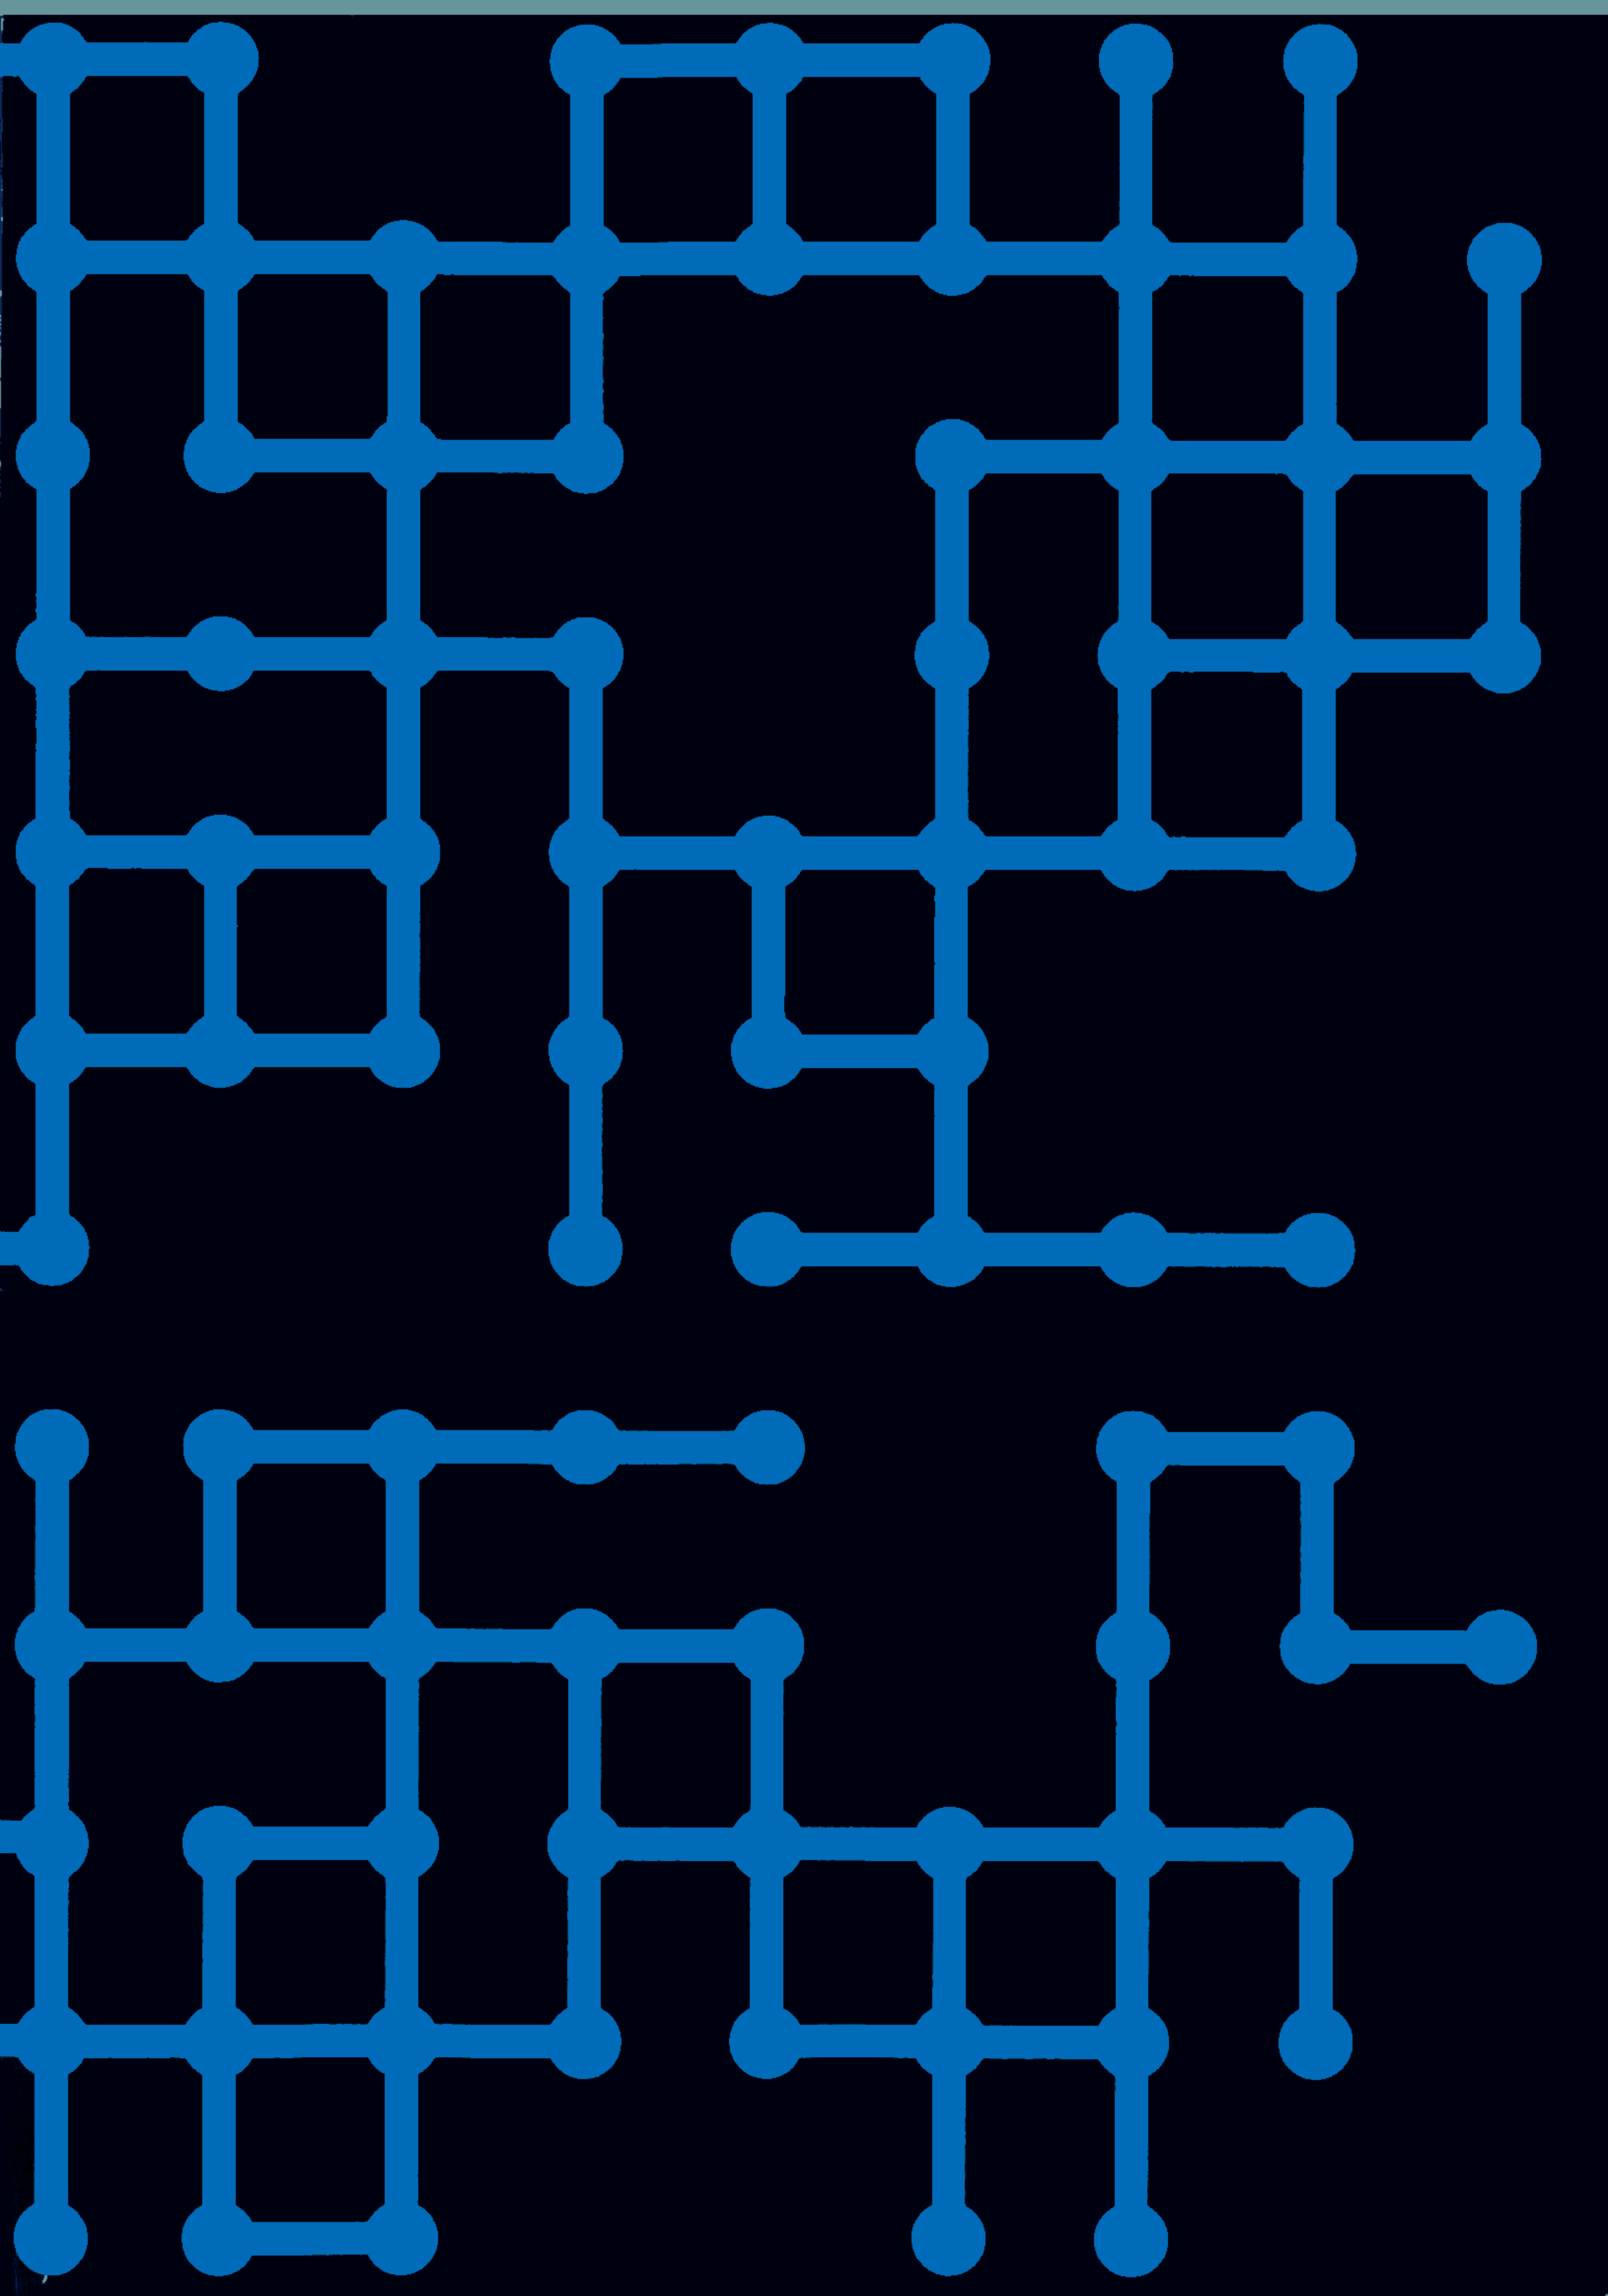
Wang, Pei 164
Waser, Mark 170
Wiedermann, Jiri 196
Wray, Robert 79