


Festschrift

LNAI 4850

Max Lungarella Fumiya Iida
Josh Bongard Rolf Pfeifer (Eds.)

50 Years of Artificial Intelligence

Essays Dedicated to the 50th Anniversary
of Artificial Intelligence



Springer

TP18

739

Max Lungarella Fumiya Iida
Josh Bongard Rolf Pfeifer (Eds.)

50 Years of Artificial Intelligence

Essays Dedicated to the 50th Anniversary
of Artificial Intelligence



Springer



E2008000727

Series Editors

Jaime G. Carbonell, Carnegie Mellon University, Pittsburgh, PA, USA
Jörg Siekmann, University of Saarland, Saarbrücken, Germany

Volume Editors

Max Lungarella
Rolf Pfeifer
University of Zurich
Artificial Intelligence Laboratory
Andreasstrasse 15, 8050 Zurich, Switzerland
E-mail: {lunga,pfeifer}@ifi.uzh.ch

Fumiya Iida
Massachusetts Institute of Technology
Robot Locomotion Group Computer Science
and Artificial Intelligence Laboratory
32 Vassar Street, Cambridge, MA 02139, USA
E-mail: iida@csail.mit.edu

Josh Bongard
University of Vermont
Department of Computer Science
329 Votey Hall, Burlington, VT 05405, USA
E-mail: j.bongard@uvm.edu

The illustration appearing on the cover of this book is the work of Daniel Rozenberg (DADARA).

Library of Congress Control Number: 2007941079

CR Subject Classification (1998): I.2, H.3-5, H.2.8, F.2.2, I.6

LNCS Sublibrary: SL 7 – Artificial Intelligence

ISSN 0302-9743
ISBN-10 3-540-77295-2 Springer Berlin Heidelberg New York
ISBN-13 978-3-540-77295-8 Springer Berlin Heidelberg New York

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.

Springer is a part of Springer Science+Business Media

springer.com

© Springer-Verlag Berlin Heidelberg 2007
Printed in Germany

Typesetting: Camera-ready by author, data conversion by Scientific Publishing Services, Chennai, India
Printed on acid-free paper SPIN: 12204854 06/3180 5 4 3 2 1 0

Lecture Notes in Artificial Intelligence

4850

Edited by J. G. Carbonell and J. Siekmann

Subseries of Lecture Notes in Computer Science

Preface

Half a century ago, at the now famous 1956 Dartmouth Conference, the “fathers” of Artificial Intelligence (AI) – among them John McCarthy, Marvin Minsky, Allen Newell, Claude Shannon, Herbert Simon, Oliver Selfridge, and Ray Solomonoff – convened under the premise “that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it.” Fifty years have passed, and AI has turned into an important field whose influence on our daily lives can hardly be overestimated. Many specialized AI systems exist that are at work in our cars, in our laptop computers, and in our personal and commercial technologies. There is no doubt that the impact of AI on our lives in the future will become even more general and ubiquitous.

In this book we provide a representative collection of papers written by the leading researchers in the field of Artificial Intelligence. All of the authors of papers in this volume attended the 50th Anniversary Summit of AI (<http://www.ai50.org>), held at the Centro Stefano Franscini, Monte Verità, Ascona, Switzerland, July 9–14, 2006. The objective of the summit was fourfold: (1) to celebrate the 50th anniversary of AI as a discipline; (2) to look back and assess the field of AI (what has been done, and where we are); (3) to bring together people with different backgrounds (to enhance interaction between groups and foster future collaborations); and (4) to attract young and talented researchers to generate additional momentum in this exciting field. The AI summit combined discussions from a historical standpoint; scientific exchange on the state of the art; speculations about the future; business, political and educational perspectives; contributions by researchers from different but related areas; presentations of the latest research by top scientists in the field; as well as many informal discussions among the participants and visitors. In this volume, we have tried to maintain the breadth of topics presented and discussed at the summit by including chapters focusing on subjects ranging from the history and prospects of AI, to speech recognition and processing, linguistics, bionics, and consciousness.

We would like to thank all the participants of the summit for helping to make it a successful event, the authors for their contributions to this volume, and the reviewers. We would also like to express our gratitude to the Centro Stefano Franscini, Neuronics AG, Swisscom Innovations, Matek, and Migros Kulturprozent for their support.

September 2007

Max Lungarella
Fumiya Iida
Josh C. Bongard
Rolf Pfeifer

Lecture Notes in Artificial Intelligence (LNAI)

- Vol. 4874: J. Neves, M.F. Santos, J.M. Machado (Eds.), *Progress in Artificial Intelligence*. XVIII, 704 pages. 2007.
- Vol. 4850: M. Lungarella, F. Iida, J. Bongard, R. Pfeifer (Eds.), *50 Years of Artificial Intelligence*. X, 399 pages. 2007.
- Vol. 4845: N. Zhong, J. Liu, Y. Yao, J. Wu, S. Lu, K. Li (Eds.), *Web Intelligence Meets Brain Informatics*. XI, 516 pages. 2007.
- Vol. 4830: M.A. Orgun, J. Thornton (Eds.), *AI 2007: Advances in Artificial Intelligence*. XIX, 841 pages. 2007.
- Vol. 4828: M. Randall, H.A. Abbass, J. Wiles (Eds.), *Progress in Artificial Life*. XII, 402 pages. 2007.
- Vol. 4827: A. Gelbukh, Á.F. Kuri Morales (Eds.), *MICAI 2007: Advances in Artificial Intelligence*. XXIV, 1234 pages. 2007.
- Vol. 4826: P. Perner, O. Salvetti (Eds.), *Advances in Mass Data Analysis of Signals and Images in Medicine, Biotechnology and Chemistry*. X, 183 pages. 2007.
- Vol. 4819: T. Washio, Z.-H. Zhou, J.Z. Huang, X. Hu, J. Li, C. Xie, J. He, D. Zou, K.-C. Li, M.M. Freire (Eds.), *Emerging Technologies in Knowledge Discovery and Data Mining*. XIV, 675 pages. 2007.
- Vol. 4798: Z. Zhang, J.H. Siekmann (Eds.), *Knowledge Science and Engineering and Management*. XVI, 669 pages. 2007.
- Vol. 4795: F. Schilder, G. Katz, J. Pustejovsky (Eds.), *Annotating, Extracting and Reasoning about Time and Events*. VII, 141 pages. 2007.
- Vol. 4790: N. Dershowitz, A. Voronkov (Eds.), *Logic for Programming, Artificial Intelligence, and Reasoning*. XIII, 562 pages. 2007.
- Vol. 4788: D. Borrajo, L. Castillo, J.M. Corchado (Eds.), *Current Topics in Artificial Intelligence*. XI, 280 pages. 2007.
- Vol. 4775: A. Esposito, M. Faundez-Zanuy, E. Keller, M. Marinaro (Eds.), *Verbal and Nonverbal Communication Behaviours*. XII, 325 pages. 2007.
- Vol. 4772: H. Prade, V.S. Subrahmanian (Eds.), *Scalable Uncertainty Management*. X, 277 pages. 2007.
- Vol. 4766: N. Maudet, S. Parsons, I. Rahwan (Eds.), *Argumentation in Multi-Agent Systems*. XII, 211 pages. 2007.
- Vol. 4755: V. Corruble, M. Takeda, E. Suzuki (Eds.), *Discovery Science*. XI, 298 pages. 2007.
- Vol. 4754: M. Hutter, R.A. Servedio, E. Takimoto (Eds.), *Algorithmic Learning Theory*. XI, 403 pages. 2007.
- Vol. 4737: B. Berendt, A. Hotho, D. Mladenic, G. Semeraro (Eds.), *From Web to Social Web: Discovering and Deploying User and Content Profiles*. XI, 161 pages. 2007.
- Vol. 4733: R. Basili, M.T. Pazzienza (Eds.), *AI*IA 2007: Artificial Intelligence and Human-Oriented Computing*. XVII, 858 pages. 2007.
- Vol. 4724: K. Mellouli (Ed.), *Symbolic and Quantitative Approaches to Reasoning with Uncertainty*. XV, 914 pages. 2007.
- Vol. 4722: C. Pelachaud, J.-C. Martin, E. André, G. Chollet, K. Karpouzis, D. Pelé (Eds.), *Intelligent Virtual Agents*. XV, 425 pages. 2007.
- Vol. 4720: B. Konev, F. Wolter (Eds.), *Frontiers of Combining Systems*. X, 283 pages. 2007.
- Vol. 4702: J.N. Kok, J. Koronacki, R. Lopez de Mantaras, S. Matwin, D. Mladenič, A. Skowron (Eds.), *Knowledge Discovery in Databases: PKDD 2007*. XXIV, 640 pages. 2007.
- Vol. 4701: J.N. Kok, J. Koronacki, R. Lopez de Mantaras, S. Matwin, D. Mladenič, A. Skowron (Eds.), *Machine Learning: ECML 2007*. XXII, 809 pages. 2007.
- Vol. 4696: H.-D. Burkhard, G. Lindemann, R. Verbrugge, L.Z. Varga (Eds.), *Multi-Agent Systems and Applications V*. XIII, 350 pages. 2007.
- Vol. 4694: B. Apolloni, R.J. Howlett, L. Jain (Eds.), *Knowledge-Based Intelligent Information and Engineering Systems, Part III*. XXIX, 1126 pages. 2007.
- Vol. 4693: B. Apolloni, R.J. Howlett, L. Jain (Eds.), *Knowledge-Based Intelligent Information and Engineering Systems, Part II*. XXXII, 1380 pages. 2007.
- Vol. 4692: B. Apolloni, R.J. Howlett, L. Jain (Eds.), *Knowledge-Based Intelligent Information and Engineering Systems, Part I*. LV, 882 pages. 2007.
- Vol. 4687: P. Petta, J.P. Müller, M. Klusch, M. Georgeff (Eds.), *Multiagent System Technologies*. X, 207 pages. 2007.
- Vol. 4682: D.-S. Huang, L. Heutte, M. Loog (Eds.), *Advanced Intelligent Computing Theories and Applications*. XXVII, 1373 pages. 2007.
- Vol. 4676: M. Klusch, K.V. Hindriks, M.P. Papazoglou, L. Sterling (Eds.), *Cooperative Information Agents XI*. XI, 361 pages. 2007.
- Vol. 4667: J. Hertzberg, M. Beetz, R. Englert (Eds.), *KI 2007: Advances in Artificial Intelligence*. IX, 516 pages. 2007.
- Vol. 4660: S. Džeroski, L. Todorovski (Eds.), *Computational Discovery of Scientific Knowledge*. X, 327 pages. 2007.

- Vol. 4659: V. Mafik, V. Vyatkin, A.W. Colombo (Eds.), *Holonic and Multi-Agent Systems for Manufacturing*. VIII, 456 pages. 2007.
- Vol. 4651: F. Azevedo, P. Barahona, F. Fages, F. Rossi (Eds.), *Recent Advances in Constraints*. VIII, 185 pages. 2007.
- Vol. 4648: F. Almeida e Costa, L.M. Rocha, E. Costa, I. Harvey, A. Coutinho (Eds.), *Advances in Artificial Life*. XVIII, 1215 pages. 2007.
- Vol. 4635: B. Kokinov, D.C. Richardson, T.R. Roth-Berghofer, L. Vieu (Eds.), *Modeling and Using Context*. XIV, 574 pages. 2007.
- Vol. 4632: R. Alhajj, H. Gao, X. Li, J. Li, O.R. Zaiane (Eds.), *Advanced Data Mining and Applications*. XV, 634 pages. 2007.
- Vol. 4629: V. Matoušek, P. Mautner (Eds.), *Text, Speech and Dialogue*. XVII, 663 pages. 2007.
- Vol. 4626: R.O. Weber, M.M. Richter (Eds.), *Case-Based Reasoning Research and Development*. XIII, 534 pages. 2007.
- Vol. 4617: V. Torra, Y. Narukawa, Y. Yoshida (Eds.), *Modeling Decisions for Artificial Intelligence*. XII, 502 pages. 2007.
- Vol. 4612: I. Miguel, W. Ruml (Eds.), *Abstraction, Reformulation, and Approximation*. XI, 418 pages. 2007.
- Vol. 4604: U. Priss, S. Polovina, R. Hill (Eds.), *Conceptual Structures: Knowledge Architectures for Smart Applications*. XII, 514 pages. 2007.
- Vol. 4603: F. Pfenning (Ed.), *Automated Deduction – CADE-21*. XII, 522 pages. 2007.
- Vol. 4597: P. Perner (Ed.), *Advances in Data Mining*. XI, 353 pages. 2007.
- Vol. 4594: R. Bellazzi, A. Abu-Hanna, J. Hunter (Eds.), *Artificial Intelligence in Medicine*. XVI, 509 pages. 2007.
- Vol. 4585: M. Kryszkiewicz, J.F. Peters, H. Rybinski, A. Skowron (Eds.), *Rough Sets and Intelligent Systems Paradigms*. XIX, 836 pages. 2007.
- Vol. 4578: F. Masulli, S. Mitra, G. Pasi (Eds.), *Applications of Fuzzy Sets Theory*. XVIII, 693 pages. 2007.
- Vol. 4573: M. Kauers, M. Kerber, R. Miner, W. Windsteiger (Eds.), *Towards Mechanized Mathematical Assistants*. XIII, 407 pages. 2007.
- Vol. 4571: P. Perner (Ed.), *Machine Learning and Data Mining in Pattern Recognition*. XIV, 913 pages. 2007.
- Vol. 4570: H.G. Okuno, M. Ali (Eds.), *New Trends in Applied Artificial Intelligence*. XXI, 1194 pages. 2007.
- Vol. 4565: D.D. Schmorow, L.M. Reeves (Eds.), *Foundations of Augmented Cognition*. XIX, 450 pages. 2007.
- Vol. 4562: D. Harris (Ed.), *Engineering Psychology and Cognitive Ergonomics*. XXIII, 879 pages. 2007.
- Vol. 4548: N. Olivetti (Ed.), *Automated Reasoning with Analytic Tableaux and Related Methods*. X, 245 pages. 2007.
- Vol. 4539: N.H. Bshouty, C. Gentile (Eds.), *Learning Theory*. XII, 634 pages. 2007.
- Vol. 4529: P. Melin, O. Castillo, L.T. Aguilar, J. Kacprzyk, W. Pedrycz (Eds.), *Foundations of Fuzzy Logic and Soft Computing*. XIX, 830 pages. 2007.
- Vol. 4520: M.V. Butz, O. Sigaud, G. Pezzulo, G. Baldassarre (Eds.), *Anticipatory Behavior in Adaptive Learning Systems*. X, 379 pages. 2007.
- Vol. 4511: C. Conati, K. McCoy, G. Paliouras (Eds.), *User Modeling*. XVI, 487 pages. 2007.
- Vol. 4509: Z. Kobti, D. Wu (Eds.), *Advances in Artificial Intelligence*. XII, 552 pages. 2007.
- Vol. 4496: N.T. Nguyen, A. Grzech, R.J. Howlett, L.C. Jain (Eds.), *Agent and Multi-Agent Systems: Technologies and Applications*. XXI, 1046 pages. 2007.
- Vol. 4483: C. Baral, G. Brewka, J. Schlipf (Eds.), *Logic Programming and Nonmonotonic Reasoning*. IX, 327 pages. 2007.
- Vol. 4482: A. An, J. Stefanowski, S. Ramanna, C.J. Butz, W. Pedrycz, G. Wang (Eds.), *Rough Sets, Fuzzy Sets, Data Mining and Granular Computing*. XIV, 585 pages. 2007.
- Vol. 4481: J. Yao, P. Lingras, W.-Z. Wu, M. Szczuka, N.J. Cercone, D. Ślęzak (Eds.), *Rough Sets and Knowledge Technology*. XIV, 576 pages. 2007.
- Vol. 4476: V. Gorodetsky, C. Zhang, V.A. Skormin, L. Cao (Eds.), *Autonomous Intelligent Systems: Multi-Agents and Data Mining*. XIII, 323 pages. 2007.
- Vol. 4460: S. Aguzzoli, A. Ciabattini, B. Gerla, C. Manna, V. Marra (Eds.), *Algebraic and Proof-theoretic Aspects of Non-classical Logics*. VIII, 309 pages. 2007.
- Vol. 4457: G.M.P. O'Hare, A. Ricci, M.J. O'Grady, O. Dikenelli (Eds.), *Engineering Societies in the Agents World VII*. XI, 401 pages. 2007.
- Vol. 4456: Y. Wang, Y.-m. Cheung, H. Liu (Eds.), *Computational Intelligence and Security*. XXIII, 1118 pages. 2007.
- Vol. 4455: S. Muggleton, R. Otero, A. Tamadon-Nezhad (Eds.), *Inductive Logic Programming*. XII, 456 pages. 2007.
- Vol. 4452: M. Fasli, O. Shehory (Eds.), *Agent-Mediated Electronic Commerce*. VIII, 249 pages. 2007.
- Vol. 4451: T.S. Huang, A. Nijholt, M. Pantic, A. Pentland (Eds.), *Artificial Intelligence for Human Computing*. XVI, 359 pages. 2007.
- Vol. 4442: L. Antunes, K. Takadama (Eds.), *Multi-Agent-Based Simulation VII*. X, 189 pages. 2007.
- Vol. 4441: C. Müller (Ed.), *Speaker Classification II*. X, 309 pages. 2007.
- Vol. 4438: L. Maicher, A. Sigel, L.M. Garshol (Eds.), *Leveraging the Semantics of Topic Maps*. X, 257 pages. 2007.
- Vol. 4434: G. Lakemeyer, E. Sklar, D.G. Sorrenti, T. Takahashi (Eds.), *RoboCup 2006: Robot Soccer World Cup X*. XIII, 566 pages. 2007.
- Vol. 4429: R. Lu, J.H. Siekmann, C. Ullrich (Eds.), *Cognitive Systems*. X, 161 pages. 2007.
- Vol. 4428: S. Edelkamp, A. Lomuscio (Eds.), *Model Checking and Artificial Intelligence*. IX, 185 pages. 2007.

¥495.00

Table of Contents

Historical and Philosophical Issues

AI in the 21 st Century – With Historical Reflections	1
<i>Max Lungarella, Fumiya Iida, Josh C. Bongard, and Rolf Pfeifer</i>	
The Physical Symbol System Hypothesis: Status and Prospects	9
<i>Nils J. Nilsson</i>	
Fifty Years of AI: From Symbols to Embodiment - and Back	18
<i>Luc Steels</i>	
2006: Celebrating 75 Years of AI - History and Outlook: The Next 25 Years	29
<i>Jürgen Schmidhuber</i>	
Evolutionary Humanoid Robotics: Past, Present and Future	42
<i>Malachy Eaton</i>	
Philosophical Foundations of AI	53
<i>David Vernon and Dermot Furlong</i>	
On the Role of AI in the Ongoing Paradigm Shift within the Cognitive Sciences	63
<i>Tom Froese</i>	

Information Theory and Quantification

On the Information Theoretic Implications of Embodiment – Principles and Methods	76
<i>Rolf Pfeifer, Max Lungarella, Olaf Sporns, and Yasuo Kuniyoshi</i>	
Development Via Information Self-structuring of Sensorimotor Experience and Interaction	87
<i>Christopher L. Nehaniv, Naeem Assif Mirza, and Lars Olsson</i>	
How Information and Embodiment Shape Intelligent Information Processing	99
<i>Daniel Polani, Olaf Sporns, and Max Lungarella</i>	
Preliminary Considerations for a Quantitative Theory of Networked Embodied Intelligence	112
<i>Fabio P. Bonsignorio</i>	

A Quantitative Investigation into Distribution of Memory and Learning
in Multi Agent Systems with Implicit Communications 124
*Roozbeh Daneshvar, Abdolhossein Sadeghi Marascht,
Hossein Aminaiee, and Caro Lucas*

Morphology and Dynamics

AI in Locomotion: Challenges and Perspectives of Underactuated
Robots 134
Fumiya Iida, Rolf Pfeifer, and André Seyfarth

On the Task Distribution Between Control and Mechanical Systems: A
Case Study with an Amoeboid Modular Robot 144
Akio Ishiguro and Masahiro Shimizu

Bacteria Integrated Swimming Microrobots 154
Bahareh Behkam and Metin Sitti

Adaptive Multi-modal Sensors 164
Kyle I. Harrington and Hava T. Siegelmann

Neurorobotics

What Can AI Get from Neuroscience? 174
Steve M. Potter

Dynamical Systems in the Sensorimotor Loop: On the Interrelation
Between Internal and External Mechanisms of Evolved Robot
Behavior 186
*Martin Hülse, Steffen Wischmann, Poramate Manoonpong,
Arndt von Twickel, and Frank Pasemann*

Adaptive Behavior Control with Self-regulating Neurons 196
Keyan Zahedi and Frank Pasemann

Brain Area V6A: A Cognitive Model for an Embodied Artificial
Intelligence 206
*Fattori Patrizia, Breveglieri Rossella, Marzocchi Nicoletta,
Maniadakis Michail, and Galletti Claudio*

The Man-Machine Interaction: The Influence of Artificial Intelligence
on Rehabilitation Robotics 221
*Alejandro Hernández Arieta, Ryu Kato, Wenwei Yu, and
Hiroshi Yokoi*

**Machine Intelligence, Cognition, and Natural
Language Processing**

Tests of Machine Intelligence 232
Shane Legg and Marcus Hutter

A Hierarchical Concept Oriented Representation for Spatial Cognition in Mobile Robots	243
<i>Shrihari Vasudevan, Stefan Gächter, Ahad Harati, and Roland Siegwart</i>	

Anticipation and Future-Oriented Capabilities in Natural and Artificial Cognition	257
<i>Giovanni Pezzulo</i>	

Computer-Supported Human-Human Multilingual Communication	271
<i>Alex Waibel, Keni Bernardin, and Matthias Wölfel</i>	

Human-Like Intelligence: Motivation, Emotions, and Consciousness

A Paradigm Shift in Artificial Intelligence: Why Social Intelligence Matters in the Design and Development of Robots with Human-Like Intelligence	288
<i>Kerstin Dautenhahn</i>	

Intrinsically Motivated Machines	303
<i>Frédéric Kaplan and Pierre-Yves Oudeyer</i>	

Curious and Creative Machines	315
<i>Hod Lipson</i>	

Applying Data Fusion in a Rational Decision Making with Emotional Regulation	320
<i>Benjamin Fonooni, Behzad Moshiri, and Caro Lucas</i>	

How to Build Consciousness into a Robot: The Sensorimotor Approach	332
<i>J. Kevin O'Regan</i>	

Robot Platforms

A Human-Like Robot Torso ZAR5 with Fluidic Muscles: Toward a Common Platform for Embodied AI	347
<i>Ivo Boblan, Rudolf Bannasch, Andreas Schulz, and Hartmut Schwenk</i>	

The iCub Cognitive Humanoid Robot: An Open-System Research Platform for Enactive Cognition	358
<i>Giulio Sandini, Giorgio Metta, and David Vernon</i>	

Intelligent Mobile Manipulators in Industrial Applications: Experience and Challenges	370
<i>Hansruedi Fröh, Philipp Keller, and Tino Perucchi</i>	

Art and AI

The Dynamic Darwinian Diorama: A Landlocked Archipelago Enhances
Epistemology 386
 Adrienne Wortzel

Author Index 399

AI in the 21st Century – With Historical Reflections

Max Lungarella¹, Fumiya Iida², Josh C. Bongard³, and Rolf Pfeifer¹

¹ Dept. of Informatics, University of Zurich, Switzerland

² Computer Science and Artificial Intelligence Lab, MIT, USA

³ Dept. of Computer Science, University of Vermont, USA

lunga@ifi.uzh.ch, iida@csail.mit.edu, josh.bongard@uvm.edu,
pfeifer@ifi.uzh.ch

Abstract. The discipline of Artificial Intelligence (AI) was born in the summer of 1956 at Dartmouth College in Hanover, New Hampshire. Half of a century has passed, and AI has turned into an important field whose influence on our daily lives can hardly be overestimated. The original view of intelligence as a computer program – a set of algorithms to process symbols – has led to many useful applications now found in internet search engines, voice recognition software, cars, home appliances, and consumer electronics, but it has not yet contributed significantly to our understanding of natural forms of intelligence. Since the 1980s, AI has expanded into a broader study of the interaction between the body, brain, and environment, and how intelligence emerges from such interaction. This advent of embodiment has provided an entirely new way of thinking that goes well beyond artificial intelligence proper, to include the study of intelligent action in agents other than organisms or robots. For example, it supplies powerful metaphors for viewing corporations, groups of agents, and networked embedded devices as intelligent and adaptive systems acting in highly uncertain and unpredictable environments. In addition to giving us a novel outlook on information technology in general, this broader view of AI also offers unexpected perspectives into how to think about ourselves and the world around us. In this chapter, we briefly review the turbulent history of AI research, point to some of its current trends, and to challenges that the AI of the 21st century will have to face.

1 Introduction

For a long time, humans have been romanced by the idea of understanding how thinking works, or how to construct intelligent machines and replicate the intelligent behavior displayed by many natural systems. Traditional Jewish mysticism, for instance, includes tales of the Golem, a thinking automaton made from the sticky clay of the bank of the river Moldau. In the 17th century, philosopher Gottfried Wilhelm von Leibniz outlined plans for a thinking machine by conceiving an artificial universal language composed of symbols, which could stand for objects or concepts, and logical rules for their manipulation. A little more than half a century ago (when Norbert Wiener was devising Cybernetics [1] and Gray Walter was building robotic tortoises [2], the English mathematician Alan Turing proposed a much-discussed imitation game used as a yardstick for assessing if a machine is intelligent or not, which since then has been known as the Turing Test for artificial intelligence [3, 4].

The advent of the general-purpose computer in the fifties of last century did substantially alter the dreams and ambitions of people and made AI an even more tangible possibility. So, in 1956, the fathers of “modern” AI, Marvin Minsky, John McCarthy, Allen Newell, Nathaniel Rochester, Claude Shannon, and Herbert Simon came together for a summer school at Dartmouth College (Hanover, New Hampshire) under the premise “that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it” [5]. That date can be considered the birth of AI because it is only thereafter that numerous research groups around the world began to engage in the construction of artificial systems with the professed goal of emulating, equaling, or even surpassing human mental and physical abilities (a different view on the history of AI can be found in [6]). The many attempts to synthesize intelligence or higher cognition (by formalizing knowledge and crystallizing cognitive principles obtained from the study of human beings) have resulted in many specialized AI systems that nowadays are at work in specific problem domains, such as knowledge discovery and data mining (for example, for the identification of customer profiles), fuzzy logic, probabilistic reasoning, and artificial neural networks (as used in intelligent control for robots, home appliances and consumer electronics), evolutionary optimization (for the design of antennae, circuits, and turbines), as well as statistical and machine learning (often employed in genomics and proteomics).

Although, as these examples show, AI has proven successful in many cases, it is clear that the original goals set out by the first generation of AI visionaries have not yet been reached. The holy grail of AI, the creation of general (human-like) machine intelligence (think of HAL from Stanley Kubrick’s *2001* or *Bladerunner*’s Roy Batty), has remained elusive. Natural intelligence is far from being understood and artificial forms of intelligence are still so much more primitive than natural ones. Seemingly simple tasks like object manipulation and recognition – tasks easily solved by a 3-year-old – have not yet been realized artificially. A look at the current research landscape reveals how little we know about how biological brains achieve their remarkable functionalities, how these functionalities develop in the child, or how they have arisen in the course of evolution. Also, we do not understand the cultural and social processes that have helped shape human intelligence. Because basic theories of natural intelligence are lacking and – despite impressive advances – the required technologies for building sophisticated artificial systems are yet not available, it is not surprising that the capabilities of current robots fall far short of even very simple animals. The inevitable conclusion then is that something important must have been missed or still needs to be discovered [7]. There is still no overall theory, no framework, explaining what thinking is and how to build intelligent machines.

Although we do not have any clear cut answer as to what aspects of natural intelligence are key for creating artificial intelligence, we will point at some avenues that might be worth pursuing in realizing the dreams of our ancestors.

2 Trends

The field of artificial intelligence has changed dramatically over the last 50 years. Back in the 1950s, the brain was conceptualized as some kind of powerful computer,

and intelligence was thought to be the result of a symbol-crunching computer program located somewhere in our brain (this perspective is also known as the “brain as a computer metaphor” or the “information processing view”; [8, 9]). Since then, the by now “classical”, computational, or cognitivistic approach has grown into a large discipline with many facets, and has achieved many successes in applied computer science and engineering. Think, for instance, of the victory of IBM’s chess-playing supercomputer (Deep Blue) on reigning world chess champion, Garry Kasparov, on May 1997. While this early view of intelligence is probably adequate in formal or computational domains like chess playing or logical reasoning, in the past 20 years it has become increasingly evident that conceptualizing intelligence as a mere computational process cannot explain natural, adaptive forms of intelligence. The latter kind of intelligence requires a complete physical organism interacting with the real world: in other words, intelligence requires a body [10-15].

The new approach to understanding intelligence has led to a paradigm shift which emphasizes the physical and information-theoretical implications of embodied adaptive behavior, as opposed to the disembodied view of the computational framework [16-19]. The focus in this new paradigm is on systems acting in the real, physical and social world [20]. The implications of this change in perspective are far-reaching and can hardly be overestimated. With the fundamental paradigm shift from a computational to an embodied perspective, the kinds of research areas, theoretical and engineering issues, and the disciplines involved in AI have also changed substantially. The research effort in the field, for instance, has shifted towards understanding the lower level mechanisms and processes underlying intelligent behavior, as well as realizing higher forms of intelligence by first starting with simpler ones [21]. Cognition and action are viewed as the result of emergence and development rather than something that can be directly built (i.e. programmed) into a robot (such research draws heavily on insights from child development; e.g. [17, 22-24]. Automated design methods using ideas borrowed from biological evolution and ontogeny have also provided novel insights into the general nature of intelligence [25].

Physical agents in the real world, whether biological or artificial, are highly complex, and their investigation requires the cooperation of researchers from many different fields. In terms of research disciplines contributing to AI, we observe that in the classical approach computer science, psychology, philosophy, and linguistics played major roles, whereas in the embodied approach, computer science and philosophy [26] are now complemented by robotics, biomechanics [17, 27], material science and biology [28] and neuroscience [29]. The shift from high-level considerations (as raised in psychology and linguistics) to lower level sensory-motor processes – with the neurosciences now offering explanations at both the sensory-motor and cognitive levels of intelligence – is evident [30, 31]. There have also been changes concerning the notions used for describing the research area: a significant number of researchers dealing with embodiment no longer refer to themselves as working in artificial intelligence but rather in robotics, engineering of adaptive systems, artificial life, adaptive locomotion, bio-inspired systems, or neuroinformatics. But more than that, not only have researchers in artificial intelligence migrated into neighboring fields, but researchers in these neighboring fields have also started to contribute to artificial intelligence [29, 32]. This brings us to the issues of cross-disciplinarity and cross-fertilization.

3 Cross-Disciplinarity

Although most researchers in AI today are still working within the boundaries of the field as it was defined 50 years ago, the trend seems to be an expansion of AI into a broader study of the interaction between body, brain, and environment, and how intelligence emerges from such interaction [18, 33]. It has become evident that many fields (linguistics, cognitive sciences, robotics, material science, neuroscience, and morphogenesis) are highly relevant in order to advance the state of the art. The community is interested in locomotion, manipulation, and, in general, how an agent can act successfully in a dynamic and complex world. Concepts such as embodiment, autonomy, situatedness, interactive emergence, development, semiotic dynamics, and social interaction have come to the forefront and have spawned a spectrum of new and interdisciplinary approaches to the study of natural and artificial intelligence. There are important implications of the embodied view of intelligence for science, engineering, economics, and society at large. Here, we only mention but a few.

One of the important insights is that intelligence is not a “box” sitting inside the brain, but is distributed throughout the organism and requires the organism to interact with and explore its local environment. With respect to engineering, it has been demonstrated that systems which are not centrally controlled or do not feature a strict hierarchical structure tend to be more robust and adaptive against external perturbations [34]. It is important to realize that the underlying concepts might be directly mapped onto robot design (e.g. for de-mining operations, waste cleanup in hazardous environments, environmental exploration, service in hospitals and homes, and so on) and onto the design of embedded systems (i.e. systems equipped with sensors and motors which interact autonomously and continuously with their surrounding environment). The application of concepts from embodied AI in embedded systems could have profound economical and technological implications, because such systems are extremely widespread: they are used, for instance, in fuel injection devices, water purification plants, air conditioning systems, remote monitoring and control systems, as well as in many systems designed for human computer interaction. Finally, companies, societies and political bodies with an organizational structure based on local interactions and self-organization have been shown to react more robustly to unpredictable market forces than hierarchically organized ones (a good example of this is the financial difficulties currently experienced by many major airlines, compared to the proliferation and success of smaller, regional airlines).

4 Challenges and Outlook

Niels Bohr once famously quipped: “It’s hard to predict, especially the future”. Similarly, it is difficult to make predictions, especially about the future of artificial intelligence. However, after 50 years of explorations into artificially intelligent systems (and many false predictions), we have learned enough about the nature of intelligence to venture plausible guesses as to where AI could go as a field and which challenges it might face in the coming decades. In this section, we list some of these challenges and point to chapters in this volume that discuss them.

It seems to be generally accepted nowadays that the behavior of any system physically embedded within a particular environment is not the mere outcome of an

internal control structure (e.g. the central nervous system) but is also affected by the system's morphology and the its material properties . Increasing evidence suggests that a direct link exists between embodiment and information: coupled sensory-motor activity and body morphology induce statistical regularities in sensory input and within the control architecture and therefore enhance internal information processing [35]. Mechanisms need to be created for the embodied and embedded systems to collect relevant learning material on their own [36] and for learning to take place in an "ecological context" (i.e. with respect to the environment). These and related issues are addressed in the chapters by Pfeifer et al. [18] and Polani et al. [19]. Potential avenues of how to exploit the interaction between control structure and physical properties (morphology and material characteristics) are described in the chapters by Iida et al. [16] and Behkam and Sitti [18].

Another challenge will be to devise a systematic theory of intelligence [14]. As argued by Nehaniv et al. [17] and Polani et al. [19], a promising approach may have to place special emphasis on information theory as a descriptive and predictive framework linking morphology, perception, action, and neural control. It will also be important to understand how raw (uninterpreted) information from unknown sensors can be used by a developing embodied agent with no prior knowledge of its motor capabilities to bootstrap cognition through a process of autonomous self-structuring in response to a history of self-motivated interaction with a rich environment [18, 19, 37]. This is closely connected to the issue of self-motivation and curiosity which are discussed in the chapters by Kaplan and Oudeyer [22] and Lipson [25]. Research on general capacities such as creativity, curiosity, motivations, action selection, and prediction (i.e. the ability to foresee consequence of actions) will be necessary because, ideally, no tasks should be pre-specified to a robot, but only an internal "abstract" reward function, some core knowledge, or a set of basic motivational (or emotional) "drives." Robots need to be endowed with the capacity to explore which are the activities that are maximally fitted to their current capabilities. Robots controlled by some core motives will hopefully be able to autonomously function in their environments not only to fulfill predefined tasks but also to search for situations where learning happens efficiently.

In the same vein, it will be necessary to address the issue of how robots (and embodied agents in general) can give meaning to symbols and construct meaningful language systems (semiotic systems). A fascinating new insight – explored under the label of "semiotic dynamics" – is that such semiotic systems and the associated information structure may not be static entities, but are continuously invented and negotiated by groups of people or agents which use them for communication and information organization [9]. This quite naturally leads to investigation of the collective dynamics of groups of socially intelligent robots and systems [20].

A further challenge will be how to integrate artificial intelligence and other fields. For instance, Potter [29] argues that neuroscience and artificial intelligence could profit from each other and invites researchers in both fields to venture across the divide (e.g. tools developed in the context of neuroinformatics could be used to design and manage neural databases). Similarly, Steels [9] predicts that many of the insights gained in fields such as nonlinear dynamical systems theory or evolutionary biology may have a direct bearing on the foundations of artificial intelligence as well as on how communication systems and other forms of interaction self-organize. A closer

(symbiotic) interaction between artificial and natural systems will also benefit rehabilitation engineering as shown by recent work on human-machine interaction [38, 39]. As envisioned by Fattori et al. [31], robotics and neurophysiology could meet in a research field where bioelectrical signals obtained by single or multi-electrode recordings can be used to drive a robotic device [40]. As a last challenge, one could envision situations where AI could be used in the context of computer-assisted human-human interaction technologies [41].

5 Epilogue

The reconsideration of brain and body as a fundamental unit, physically and informationally, as well as the emergence of a new quantitative framework that links the natural and artificial domains, has begun to produce new insights into the nature of intelligent systems. While much additional work is surely needed to arrive at or even approach a general theory of intelligence, the beginnings of a new synthesis are on the horizon. Perhaps, finally, we will come closer to understanding and building human-like intelligence. The superior adaptability of embodied, distributed systems has been acknowledged for a long time, but there is now theoretical evidence from artificial intelligence research and corroboration from computer simulations supporting this point. In summary, the ideas emerging from the modern, embodied view of artificial intelligence provide novel ways of approaching technological, social, and economic problems in the rapidly changing world of the 21st century.

References

1. Wiener, N.: *Cybernetics or Control and Communication in the Animal and the Machine*. MIT Press, Cambridge, MA (1948)
2. Walter, G.W.: An imitation of life. *Scientific American* 182(5), 42–45 (1950)
3. Turing, A.M.: Computing machinery and intelligence. *Mind* 59(236), 433–460 (1950)
4. Legg, S., Hutter, M.: *Tests of machine intelligence*. LNCS, vol. 4850. Springer, Heidelberg (2007)
5. McCorduck, P.: *Machines Who Think*, 2nd edn. A.K. Peters Ltd., Natick, MA (2003)
6. Schmidhuber, J.: 2006: Celebrating 75 years of AI – history and outlook: the next 25 years. LNCS, vol. 4850. Springer, Heidelberg (2007)
7. Brooks, R.A.: The relationship between matter and life. *Nature* 409, 409–411 (2001)
8. Nilsson, N.: The physical symbol system hypothesis: status and prospects. LNCS, vol. 4850. Springer, Heidelberg (2007)
9. Steels, L.: Fifty years of AI: from symbols to embodiment – and back. LNCS, vol. 4850. Springer, Heidelberg (2007)
10. Brooks, R.A.: New approaches to robotics. *Science* 253, 1227–1232 (1991)
11. Chiel, H., Beer, R.: The brain has a body: adaptive behavior emerges from interactions of nervous system, body, and environment. *Trends in Neurosciences* 20, 553–557 (1997)
12. Clark, A.: *Being There – Putting Brain, Body, and World Together Again*. MIT Press, Cambridge, MA (1997)
13. Pfeifer, R., Scheier, C.: *Understanding Intelligence*. MIT Press, Cambridge, MA (1999)
14. Pfeifer, R., Bongard, J.C.: *How the Body Shapes the Way we Think – A New View of Intelligence*. MIT Press, Cambridge, MA (2007)