First BMI Press hardcover edition, 2012, the 2nd print

Printed and bound in the United States of America.

Library of Congress Control Number: 2012952697

ISBN 978-0-9858757-2-5

Natural and Artificial Intelligence
Introduction to Computational Brain-Mind

Juyang Weng

BMI Press
Okemos, Michigan

Contents

Acknowledgments

Our work for autonomous intelligence development has gone through several systems: Cresceptron (1991 - 1995), SHSOLIF (1993 - 2000), SAIL (1996 - 2010), Dav (1999 - present), MILN (2004-2008), WWN (WWN-1 through WWN-7, 2007-present). Many colleagues and friends have provided insights, encouragements and suggestions during this period. I would like to thank my colleague Prof. Anil K. Jain in the Department of Computer Science and Engineering at Michigan State University (MSU) who encouraged and supported my early-stage ideas. I would like to thank my colleague Prof. George Stockman in the same department whose critical comments are so useful for a less-guarded mind. My early understanding about animal learning benefited from Prof. Mark E. Rilling in the Department of Psychology at MSU who kindly accepted me as one of his students in his animal learning course in the fall of 1997. My knowledge about neuroscience was greatly enriched by attending MIT 9.011 Principles of Neuroscience, a course in the Department of Brain and Cognitive Sciences, in the fall of 1999, while I was on sabbatical at MIT, hosted by Prof. Sandy Pentland. This six-credit-hour per week intensive course for graduate students was co-taught by Profs. Ann Graybiel, Guosong Liu, Richard Wurtman, Peter Chiller, Chris Brown, Emilio Bizzi, Mriganka Sur, Earl Miller, and Suzanne Corkin, from whom I have learned a lot and with whom I had a great time discussing issues, inside and outside the classroom. I had many extensive discussions with Ida Stockman, a well known, extremely nice developmental psychologist and expert in special education, who taught me many of her perspectives about development. I especially like to mention late Esther Thelen, a well know developmental psychologist, who suggested the term "autonomous" before "mental development" which now has become a well known acronym as AMD. During Fall 2006 and Spring 2007, Prof. Mriganka Sur at MIT hosted me as a visiting professor at the Department of Brain and Cognitive Sciences. This was an invaluable period where I learned about many latest results in neuroscience from faculty, post-doctoral associates and graduate students and got some rich experience. During this full academic year, I was also fortunate to have the opportunity to learn from the courses MIT 7.012 Introductory Biology taught by Eric Lander and Robert Weinberg, MIT 9.036 the Visual System taught by Peter H. Schiller, and MIT 9.285J Neural Coding and Perception of Sound taught by Bertrand A.R. Delgutte. During the summer of 2012, I took and passed BMI 811 Biology for Brain-Mind Research, taught by Dr. Jon Stoltzfus, and BMI 821 Neuroscience for Brain-Mind Research, taught by Prof. Lynwood G. Clemens. The Brain-Mind Institute gave me an opportunity to formally take these two courses, as the oldest student in each course.

I would also like to extend thanks to my colleagues in related fields at several institutions, for collaborations and idea exchanges. They are: Prof. Narendra Ahuja with UIUC, Prof. Ronald Arkin with Georgia Tech., Prof. Dean Aslam with MSU, Prof. Dana Ballard with Rochester U., Prof. Neil E. Berthier with U. Mass, Prof. Rodney A. Brooks with MIT, Prof.. Mingmin Chi with Fudan U., Prof. Chris Brown with Rochester U., Prof. Susan Carey with New York U. and later Harvard U., Prof. Rama Chellappa with U. Maryland, Prof. Rachel Clifton with U. Mass, Prof. Gedeon Deak with UCSD, Prof. Gerald M. Edelman with The Scripps Research Institute, Prof. John R. Deller with MSU, Prof. Jeff Elman with UCSD, Prof. Jianfeng Feng with U. of Warwick, Prof. Stan Franklin with U. of Memphis, Prof. Stephen Grossberg with Boston U., Prof. John Henderson with MSU, Prof. Patrick Dickson with MSU, Prof. Michael Gabriel with UIUC, Prof. Tony Jebara with Columbia U., Prof. Susan A. Gelman with U. Michigan, Prof. Roderic Grupen with U. Mass, Prof. Thomas S. Huang with UIUC, Prof. Xuanjing Huang with Fudan U., Prof. Takeo Kanade with CMU, Prof. Stephen Levinson with UIUC, Prof. Taosheng Liu with MSU, Prof. Hong Lu with Fudan U., Prof. Sridhar Mahadevan with U. Mass, Prof. Maja Mataric with USC, Prof. James McClelland with Stanford U., Dr. Jun Miao with the Institute of Computing Technology of the Chinese Academy of Sciences, Dr. Javier R. Movellan with UCSD, Prof. Ranjan Mukherjee with MSU, Prof. Michael M. Merzenich with UCSF, Prof. Alex Pentland with MIT, Prof. Jiaguo Qi with MSU, Prof. Brian Scassellati with Yale U., Prof. Tomaso Poggio with MIT, Prof. Fathi Salam with MSU, Prof. Klaus J. Schulten with UIUC, Prof. Terrence J. Sejnowski with Salk Institute, Prof. Tiande Shou with Fudan U., Prof. Olaf Sporns with Indiana U., Prof. David A. Stewart with MSU, Prof. Ida Stockman with MSU, Prof. Mriganka Sur with MIT, Prof. Lal Tammala with MSU, Prof. Esther Thelen with Indiana U., Prof. Eric Torng with MSU, Prof. Xiaofeng Wu with Fudan U., Prof. Ning Xi with MSU, Prof. Xiangyang Xue with Fudan U., Prof. Hongbo Yu with Fudan University, Prof. Jun Zhang with U. Michigan, Prof. Liming Zhang with Fudan U., Prof. Yang Wang with MSU, Prof. Wenqiang Zhang with Fudan U., Dr. Paul J. Werbos with NSF, Drs. Cha Zhang and Dr. Zhengyou Zhang with Microsoft Research, and many others.

The related projects would not be possible without contributions from my research advisees who worked on various components. They are: Gil Abramovich, Hamid Alavi, Leif E. Alton, Micky L. Badgero, Laura A. Blackwood, Greg J. Bloy, David A. Bordoley, Kevin A. Brown, Yi Chen, David M. Cherba, Paul Cornwell, Matthew P. Ebrom, Shaoyun Chen, Yuntao Cui, Colin H. Evans, Raja Ganjikunta, Laura Grabowski, Kamen Guentchev, Xiao Huang, Jianda Han, Mao Hu, Sally Howden, Wey-Shiuan Hwang, Zhengping Ji, Ameet V. Joshi, Jeffrey A. Kohler, Gongjun Li, Yong-Beom Lee, Quo L. Liu, Matthew D. Luciw, Tianyu Luwang, Jason L. Massey, Kenneth Marable, Kajal Miyan, Rossen Nedeltchev, Lisa Ossian, Stephen Paslaski, Rebecca Smith, Mojtaba Solgi, Xiaoying Song, Jason Sperber, Benjamin Syzek, Daniel L. Swets, Keng Yew Tham, Wei Tong, Nikita N. Wagle, Jingliang Wang, Yuekai Wang, Jamal A. Wills, Changjiang Yang, Shuqing Zeng, Nan Zhang, Yilu Zhang, Zejia Zheng, and Yang Zhou. They were with MSU or with Fudan U. where I served as a visiting professor.

I would also like to extend my special thanks to the students in the MSU CSE902 Special Topics course in Mental Development and Online Learning, MSU CSE941 Special Topics in Autonomous Mental Development, MSU CSE841 Artificial Intelligence, and MIT 9.915/6.887 Computational Cognitive and Neural Development that I was fortunate to teach using earlier versions of the volume as the textbook. They all provided valuable written and verbal comments and suggestions, which significantly helped me in improving the writing and reducing typos.

I would also like to acknowledge NSF and DARPA for their support of my new line of work. I would especially like to thank Drs. Mike Lesk, Howard Moraff, and Larry Reeker at NSF and Drs. Mark Swinson and Doug Gage at DARPA for their encouragement and support. I would also like to thank Jing Xiao at NSF and Mark Swinson at DARPA for their support of the "Workshop on Development and Learning" (WDL), held at MSU, April 5-7, 2000, which has grown into a regularly scheduled conference series called the International Conference on Development and Learning (ICDL) which later led to the *AMD Newsletter* and the *IEEE Transactions on Autonomous Mental Development* with much support from the interested IEEE societies and the growing research community. The NSF/DARPA support also enabled us to create the web based *Mental Development Repository* at www.mentaldev.org. I am very grateful to the financial supports from Microsoft Research, General Motors, and Toyota Research Institute North America. Such supports were indeed visionary and I wish that scientific rules could be put in place for our peer reviews so that our peer reviewers can be visionary too.

The great participation and support of the research community and professional societies to the *Brain-Mind Institute* and the *Brain-Mind Magazine* are also greatly appreciated, as we hope that such educational and research activities can bridge the wide communication gaps between disciplines and between the research community and general public.

I extend thanks to Shelly Villarreal, Kimberly Thompson and Stephanie Beal who read the manuscript for typos and grammar. However, all of the remaining typos are mine and not theirs. Please report errors to me at weng@cse.msu.edu.

Last, but not the least, I'm grateful to my wife, Min Guo, for her supports, encouragements, and assuming many of my chores to give me time for this endeavor, as well as to my two lovely boys, Colin and Rodney, some of whose early learning behaviors were invaluable to me while I was fascinated by development.

Juyang Weng

October, 2012
East Lansing, USA

Preface

The mind is what the brain does. This volume tries to map a mind model to the corresponding brain so as to not only deepen our understanding of both the brain and the mind, but also unveil computational underpinnings. That is why the words "Brain-Mind" are hyphenated in the title.

This volume strives to unify natural intelligence with artificial intelligence. It approaches intelligence through not only what intelligence *is* but also how intelligence *arises*.

In terms of natural intelligence, instead of modeling what a brain *does*, this volume explains how the brain *develops* — how the brain circuits and mind functions emerge from interactions between the DN and the external environments.

In terms of artificial intelligence, instead of modeling an intelligent agent as a probabilistic version of a *static*, task-specific, and handcrafted finite automaton (FA), this volume explains how a genome-like program (Developmental Program, DP) autonomously grows and adapts a brain-like network (Developmental Network, DN).

In contrast with each FA, the DP is task-nonspecific, general purpose, and capable of growing a wide variety of intelligent agents (DNs). Each DN is an attentive, dynamically changing, and probabilistically optimal brain, natural or artificial, that is drastically smaller (e.g., a large constant) than its static FA-equivalent in terms of the number of states (i.e., exponential in the number of concepts to model).

In contrast with traditional artificial intelligence (AI) approaches, humans are not in the loop of handcrafting an FA. Instead, they are part of the external environments, interacting with the DNs as they interact with human children. Different training paths lead to different careers of those DN brains, natural or artificial.

Fundamentally different from traditional artificial neural networks, each DN can abstract, at least in principle, as well as any symbolic AI system.

Like all other books on this highly challenging subject, this volume has not resolved all the details of the brain-mind problem. It is only an introduction to a new departure. Much of the material in this volume seems to be ahead of time, so I beg the reader to tolerate the facts, approaches, methods, analyses, and views expressed in this volume.

Constrained by my limited exposure to the related disciplines, examples of fundamental discipline questions discussed or implied in this volume include:

Biology: How could autonomous and individual cells interact to give rise to animal behaviors, and what cellular roles could the genome likely play?

Neuroscience: From an overarching perspective, how could a brain self-wire, perform top-down attention, and develop its functions?

Psychology: How does an integrated brain architecture accomplish multiple psychological learning models and develop behaviors?

Computer Science: How does a brain-like network compute, adapt, reason, and generalize, and how is the automaton theory related to the brain-like network?

Electrical Engineering: How does a brain-like network perform general-purpose, nonlinear, feedback sensing-and-control, beyond traditional nonlinear control?

Mathematics: How does a brain-like network perform general-purpose, nonlinear optimization, and how does a brain realize emergent functionals?

Physics: How do meanings arise from physics, and how does a brain-like network treat space and time in a unified way, reminiscent of relativity?

Social sciences: How do computational principles of human brains provide insight into possible solutions to a variety of social and political problems?

A basic reason for such a concise volume to be able to discuss the above wide variety of discipline problems is that nature is governed by basic laws that are fewer and more abstract than the number of observable concrete natural phenomena, analogous to how newton's laws of motion explain rich phenomena of natural motions.

As both a research monograph and a textbook, the problems at the end of each chapter are meant for senior undergraduate students, graduate students, researchers and other interested readers to practice. The problems marked with asterisk "*" are relatively more challenging. The required mathematical background has been reviewed in the Appendix. For those readers who like to see more concrete DN examples, I have cited the related publications. A less analytical reader should be able to get the basic ideas if he pays attention to rich meanings of the mathematical formulations.

Juyang Weng

October, 2012
East Lansing, USA

Chapter 1

Agents and Tasks

An intelligent system is for performing tasks. However, there are many different tasks in all walks of life, from playing chess, to driving a car, to performing scientific research. The subjects of artificial intelligence (AI) include almost all other disciplines. Suppose that one wants to build an intelligent system for a task T in discipline D; then we may expect that he must first learn knowledge in discipline D so that he knows how to construct an intelligent machine for task T. If we treat AI in this task-specific way, we would probably never be able to finish this course, since we would have to learn almost all disciplines. Therefore, we need to approach AI in a systematic way. The first basic concept of this systematic way is agents. We discuss agents and their environments. From these basic agent-related concepts, we will study various properties of a task, which fall into five categories: external environment, input, internal environment, output and goal.

1.1 Agents

An *agent* is anything that senses and acts, as shown in Fig. 1.1. Thus, any task executor is an agent, regardless if it is natural or artificial. A cat or a human is a natural agent. A computer program or a robot is an artificial agent. A mixed agent is also possible, such as a human whom is controlled by a computer program or a robot remotely controlled by a human. agent

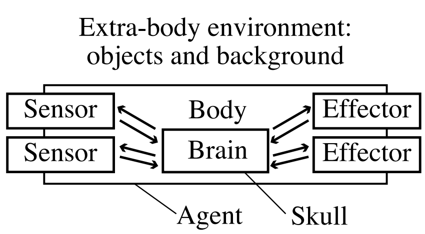

Figure 1.1: The abstract model of an agent