

Foundations of Real-World Intelligence

Copyright notice: Excerpted from *Foundations of Real-World Intelligence* edited by Yoshinori Uesaka, Pentti Kanerva, and Hideki Asoh published by CSLI Publications. ©2001 by CSLI Publications. All rights reserved. This text may be used and shared in accordance with the fair-use provisions of U.S. copyright law, and it may be archived and redistributed in electronic form, provided that this entire notice, including copyright information, is carried and provided that CSLI Publications is notified and no fee is charged for access. Archiving, redistribution, or republication of this text on other terms, in any medium, requires the consent of CSLI Publications.

Foundations of Real-World Intelligence



edited by
Yoshinori Uesaka
Pentti Kanerva
Hideki Asoh

Copyright © 2001
CSLI Publications
Center for the Study of Language and Information
Leland Stanford Junior University
Printed in the United States
05 04 03 02 01 1 2 3 4 5

Library of Congress Cataloging-in-Publication Data

Foundations of real-world intelligence / edited by Yoshinori Uesaka,
Pentti Kanerva, Hideki Asoh.

p. cm. -- (CSLI lecture notes ; no. 125)

Includes bibliographical references and index.

ISBN 1-57586-339-1 (acid-free paper) -- ISBN 1-57586-338-3 (pbk. :
acid-free paper)

1. Artificial intelligence. 2. Neural networks (Computer science)
3. Evolutionary programming (Computer science) I. Uesaka, Yoshinori,
1936- . II. Kanerva, Pentti, 1937- . III. Asoh, Hideki, 1958- . IV. Series.
Q335 .F685 2001
006.3--dc21

2001043252

∞ The acid-free paper used in this book meets the minimum
requirements of the American National Standard for Information Sciences—
Permanence of Paper for Printed Library Materials, ANSI Z39.48-1984.

Please visit our web site at
<http://cslipublications.stanford.edu/>
for comments on this and other titles, as well as for changes
and corrections by the author and publisher.

Contents

Preface	x
General Introduction	1
<i>RWI Research Center, Electrotechnical Laboratory</i>	
1 Real-World Intelligence and the Real-World Computing Program	1
NOBUYUKI OTSU	
1.1 Outline of the RWC Program	3
1.2 Real-World Intelligence	4
1.3 Concluding Remarks	8
2 Theoretical and Algorithmic Foundations of Real-World Intelligence	9
HIDEKI ASOH	
2.1 Objective	10
2.2 Approach	10
2.3 Research Issues	11
2.4 Organization of R&D and This Book	13
2.5 Concluding Remark	17
References	17
I Inference and Learning with Graphical Models	19
<i>RWI Research Center, Electrotechnical Laboratory</i>	
3 An Overview of Theoretical Foundation Research in RWI Research Center	21
HIDEKI ASOH, KAZUHISA NIKI, KOITI HASIDA, SHOTARO AKAHO, MASARU TANAKA, YOICHI MOTOMURA, TATSUYA NIWA, AND KENJI FUKUMIZU	
3.1 Models and Algorithms	21
3.2 Frameworks of Learning	27

4	BAYONET: Bayesian Network on Neural Network	28
	YOICHI MOTOMURA	
4.1	Bayesian Networks Based on Neural Networks	29
4.2	Implementation	32
4.3	Application	36
4.4	Conclusion	37
5	Multivariate Information Analysis	38
	KAZUHISA NIKI, JUNPEI HATOU, TOSHIAKI KAWAMATA, AND IKUO TAHARA	
5.1	Expression of Multivariate Analysis	38
5.2	Simulation	43
5.3	Structure Analyses of fMRI Data	49
5.4	Extended Functional Connectivity Analysis	51
5.5	Conclusion	54
6	Dialogue-based Map Learning in an Office Robot	55
	HIDEKI ASOH, YOICHI MOTOMURA, TOSHIHIRO MATSUI, SATORU HAYAMIZU, AND ISAO HARA	
6.1	Dialogue-based Map Acquisition	55
6.2	System and Experiment	60
6.3	Discussion	64
6.4	Related Work	65
6.5	Conclusion and Future Work	66
7	Conclusion	68
	References	68
II	Approximate Reasoning: Real-World Applications of Graphical Models	73
	<i>RWC Theoretical Foundation SNN Laboratory</i>	
	BERT KAPPEN, STAN GIELEN, WIM WIEGERINCK, ALI TAYLAN CEMGIL, TOM HESKES, MARCEL NIJMAN, AND MARTIJN LEISINK	
8	Mean Field Approximations	74
8.1	Mean Field Approximation with Structure	75
8.2	Boltzmann Machine Learning Using Mean Field Theory and Linear Response Correction	79
8.3	Second-order Approximations for Probability Models	85
8.4	Discussion	93
9	Medical Diagnosis	94
9.1	Probabilistic Modeling in the Medical Domain	95
9.2	Promedas, a Demonstration DSS	97
9.3	Discussion	99

10	Automatic Music Transcription	100
10.1	Dynamical Systems and the Kalman Filter	102
10.2	Tempogram Representation	106
10.3	Model Training	109
10.4	Evaluation	111
10.5	Discussion and Conclusions	116
	References	119
III	Evolutionary Computation and Beyond	123
	<i>RWCP Theoretical Foundation GMD Laboratory</i>	
	HEINZ MÜHLENBEIN AND THILO MAHNIG	
11	Analysis of the Simple Genetic Algorithm	125
11.1	Definitions	125
11.2	Proportionate Selection	126
11.3	Recombination	126
11.4	Selection and Recombination	128
11.5	Schema Analysis Demystified	130
12	The Univariate Marginal Distribution Algorithm (UMDA)	133
12.1	Definition of UMDA	134
12.2	Computing the Average Fitness	137
13	The Science of Breeding	139
13.1	Single Trait Theory	140
13.2	Tournament Selection	144
13.3	Analytical Results for Linear Functions	145
13.4	Numerical Results for UMDA	147
13.5	Royal Road Function	147
13.6	Multimodal Functions Suited for UMDA Optimization	150
13.7	Deceptive Functions	151
13.8	Numerical Investigations of the Science of Breeding	152
14	Graphical Models and Optimization	154
14.1	Boltzmann Selection and Convergence	155
14.2	Factorization of the Distribution and the FDA	157
14.3	A New Annealing Schedule for the Boltzmann Distribution	159
14.4	Finite Populations	163
14.5	Population Size, Mutation, and Bayesian Prior	165
14.6	Constraint Optimization Problems	170
15	Computing a Bayesian Network from Data	171
15.1	LFDA—Learning a Bayesian Factorization	172
15.2	Optimization, Dependencies, and Search Distributions	175

16	System Dynamics Approach to Optimization	176
16.1	The Replicator Equation	176
16.2	Boltzmann Selection and the Replicator Equation	178
16.3	Some System Dynamics Equations for Optimization	179
16.4	Optimization of Binary Functions	181
17	Three Royal Roads to Optimization	184
18	Conclusion and Outlook	185
	References	186
IV	Distributed and Active Learning	189
	<i>RWC Theoretical Foundation NEC Laboratory</i>	
19	Distributed Cooperative Bayesian Learning	190
	KENJI YAMANISHI	
19.1	Introduction	190
19.2	Plain Model	192
20	Learning Specialist Decision Lists	209
	ATSUYOSHI NAKAMURA	
20.1	Preliminaries	211
20.2	Algorithm S-Loss-Update	213
20.3	Algorithm SWML	218
20.4	Algorithm S-Fixed-Share-Update	222
21	The Lob-Pass Problem	226
	JUN'ICHI TAKEUCHI, NAOKI ABE, AND SHUN'ICHI AMARI	
21.1	Preliminaries	231
21.2	Upper Bounds on the Expected Regret	236
21.3	Concluding Remarks	248
	References	248
V	Computing with Large Random Patterns	251
	<i>RWC Theoretical Foundation SICS Laboratory</i>	
	<i>Swedish Institute of Computer Science</i>	
22	Analogy as a Basis of Computation	254
	PENTTI KANERVA	
22.1	Computer as a Brain and Brain as a Computer	256
22.2	Artificial Neural Nets as Biologically Motivated Models of Computing	257
22.3	Description vs. Explanation	258
22.4	The Brain as a Computer for Modeling the World, and Our Model of the Brain's Computing	259
22.5	Pattern Space of Representations	260
22.6	Simple Analogical Retrieval	265

22.7	Learning from Examples	266
22.8	Toward a New Model of Computing	269
23	The Sparchunk Code: A Method to Build Higher-level Structures in a Sparsely Encoded SDM	272
	GUNNAR SJÖDIN	
23.1	Encoding Higher-Level Concepts	272
23.2	The SDM Model	273
23.3	Nonscaling for a Constant Error Probability ϵ	274
23.4	Sparse Coding	276
23.5	The Sparchunk Code	276
23.6	Clean-up of the Sparchunk Code	278
23.7	Summary	279
23.8	Appendix	279
24	Some Results on Activation and Scaling of Sparse Distributed Memory	283
	JAN KRISTOFERSON	
24.1	Different Activation Probabilities for Writing and Reading?	283
24.2	Scaling Up the Memory	286
25	A Fast Activation Mechanism for the Kanerva SDM Memory	289
	ROLAND KARLSSON	
25.1	The Jaeckel Selected-Coordinate Design	290
25.2	The New Selected-Coordinate Design	291
25.3	Results	292
25.4	Conclusions	293
26	From Words to Understanding	294
	JUSSI KARLGREN AND MAGNUS SAHLGREN	
26.1	The Meaning of ‘Meaning’	294
26.2	A Case in Point: Information Access	295
26.3	Words as Content Indicators	297
26.4	Latent Semantic Analysis	299
26.5	Random Indexing	300
26.6	What Is Text, from the Perspective of Linguistics?	301
26.7	The TOEFL-Test	302
26.8	Experimental Set-Up	303
26.9	Results and Analysis	303
26.10	Some Cognitive Implications	305
26.11	Implications for Information Access	306
26.12	Meaning in Text	307
	References	308

Preface

As is well known, computers have become increasingly powerful and far exceed human ability to solve well-defined problems such as numerical computation, document processing, and logical inference in ideal worlds where all information is known and there is an algorithm for the solution that may be stated precisely in a programming language. Nevertheless, computers are still far inferior to humans in many other intellectual tasks such as pattern recognition, dealing with ill-posed problems, and learning. If computers are to develop to the point where they can perform humanlike, flexible processing in the real world—thus opening up a new horizon in information-processing technology—we must pursue the fundamentals of humanlike information-processing, come to terms with how humans process information at an intuitive level, and embody this theoretical knowledge in various developing hardware technologies. The development of advanced information systems that exhibit flexible, humanlike intelligence capable of dealing with real-world problems is one of the most important challenges shared by such diverse fields as pattern information processing, knowledge information processing, intelligent robots, and friendly human-machine interfaces.

With the aim of developing such information-processing systems, the Ministry of International Trade and Industry (MITI) of Japan launched a large-scale national program in Real World Computing (RWC) in 1992, with a ten-year budget of over 500 million dollars. The program is outlined in the general introduction to this book. Phase I of the RWC program (1992–1996) explored three general areas: novel functions for applications, theoretical foundations, and parallel and distributed computational bases. The research in Phase II (1997–2001) has been more integrated and focused, allotting resources to two projects: real-world

intelligence (RWI) and parallel and distributed computing (PDC). The aim of the RWI program area is to broaden the application horizons of conventional information-processing technology to include flexible, real-world intelligence. To pursue this objective, we established six core areas of research and development: autonomous learning systems, multi-modal interaction systems, self-organizing information base systems, theoretical and algorithmic foundations, real-world adaptive devices, and intellectual resources.

This book is about the theoretical and algorithmic foundations of real-world intelligence. The objective of the research in this area has been to establish a theoretical foundation for two key technologies, information integration and learning, drawing ideas from neural networks, genetic algorithms, Bayesian networks, EM algorithms, hidden Markov models, and the like. Conventional artificial intelligence in a narrow sense is a top-down approach; it is too artificial, based on deductive formal logic and symbol manipulation. In contrast, neural networks—that is, connectionist AI—artificial life, genetic algorithms, complex systems, and emergent computation represent a bottom-up approach. They are based on pattern computing or subsymbolic dynamic computation, shifting the focus toward induction, learning, and self-learning. However, what is essential to developing systems with real-world intelligence is understanding the computational aspect of intelligence; that is, identifying and confirming the unifying principles of computation behind both these methods, to deepen them theoretically, and to implement them as a more powerful computational methodology.

We have classified the various themes investigated under theoretical and algorithmic foundation into the following six categories: models for representing knowledge, algorithms for inference and integration, algorithms for learning, system architectures for information integration and learning, frameworks of interaction for learning, and test-bed applications for empirical evaluation. The themes were researched by six research laboratories: NEC Laboratory, GMD Laboratory (German National Research Center for Information Technology), SNN Laboratory (Stichting Neurale Netwerken in the Netherlands), SICS Laboratory (Swedish Institute of Computer Science), Toshiba Laboratory, and ETL (Electrotechnical Laboratory) RWI Research Center.

To promote cooperative R&D among these laboratories, numerous workshops and meetings were held. As the research was drawing to a close, we faced the question of how to present the results because, in general, the fruits of theoretical research mature slowly and are not immediately demonstrated in practical systems or hardware. We finally decided to collect the results into a book, as well as try to exhibit them

in more conventional ways such as prototypes. The introductory survey and the chapters in this book are the results of the research of five of the laboratories, each of which covers multiple research issues from representation of knowledge to test-bed applications.

We hope that this book will be of value to all who have at least a modest background in modern information-processing technology and a genuine interest in the fascinating possibilities that have motivated the research.

Finally, we want to express our thanks to all contributors to the book and to everybody who devoted themselves to the RWC program. In particular, we sincerely thank Professor Hidehiko Tanaka of the University of Tokyo, chairman of the promotion committee of the RWC program; Dr. Junichi Shimada, managing director of the RWC program and director of the Tsukuba Research Center of the Real World Computing Partnership (RWCP); Dr. Yoshikuni Okada, general manager of the research planning department of RWCP; and Dr. Nobuyuki Otsu, chief of the RWI Research Center of ETL. We also thank Dikran Karagueuzian and his staff at CSLI Publications for making it possible to produce and publish the book, and Judith Feldmann for diligent editing of the manuscript.

July 2001
The Editors

General Introduction[†]

RWI Research Center, Electrotechnical Laboratory

NOBUYUKI OTSU AND HIDEKI ASOH

This chapter provides a general introduction to the ideas discussed in this book. First, we introduce the concept of *real-world intelligence*. Then, we overview the Real-World Computing (RWC) program, a ten-year large-scale R&D program promoted by the Ministry of International Trade and Industry (MITI)[†] of the Japanese government since 1992. The program consists of two R&D domains, real-world intelligence (RWI) domain and the parallel distributed computing (PDC) domain. We briefly overview R&D activities in the RWI domain. Finally, we describe in greater detail R&D efforts in the theoretical and algorithmic foundation area of the RWI domain, which is the theme of this book, as introductions to the following chapters.

1 Real-World Intelligence and the Real-World Computing Program

NOBUYUKI OTSU

Supported by the remarkable development of computer and communication technologies, information technology is producing an innovative change in today's society, not only in industrial activities but also in the qualitative improvement of our way of life. The amount of information we handle will increase explosively because of the increasing needs of

[†]In a major reorganization, Electrotechnical Laboratory (ETL) became a part of the new National Institute of Advanced Industrial Science and Technology (AIST) in April 2001, while the Ministry of International Trade and Industry (MITI) was changed to Ministry of Economy, Trade, and Industry (METI) in January 2001.

multimedia information processing and the expansion of new application domains. This means not only an increase in quantity but also an increase in the quality and variety of information. Such social and technological needs are starting to require a new paradigm of information technology, not simply as a linear extension of the conventional one, but as an essentially new underlying framework. In other words, it is necessary to make computers more user-friendly and easy to use by providing them with humanlike flexible and intelligent capabilities in order to assist and collaborate with humans in the diverse information environment of the real world.

We call such intelligent capabilities *real-world intelligence* to stress the contrast with the conventional artificial intelligence technologies, which are mainly based on the explicit description of knowledge and logical inference in well-defined environments such as games, problem solving, and theorem proving.

Today, computers have come to possess enormous computing power that far surpasses human abilities to solve well-defined problems such as numerical computation, document processing, and logical inference in preassumed ideal information worlds where there are algorithms for the solution that can be stated clearly in programming languages. Nevertheless, computers are still inferior to humans in many areas such as pattern recognition, problem solving under incomplete information, and learning. The framework of the information processing done by modern computers is still not as flexible as that of the information processing done by humans in the real world, where many problems are illdefined and hard to describe in algorithms. It might be said that current information technology is still immature in the so-called intuitive or inductive aspect in contrast to the logical and deductive aspect of reasoning.

Therefore, to cope with such real-world problems and to open a new horizon in information processing technology, it is essential to pursue the more flexible ways that humans process information, by investigating the intuitive or subsymbolic level of human information processing and embodying these methods as new information processing technologies on the basis of the developing hardware technologies. The development of information systems that have humanlike, flexible intelligence and that can cope with real-world problems is one of the most important demands common to various fields such as pattern information processing, knowledge information processing, intelligent robotics, and friendly human-machine interface, all of which aim at advancing our methods of intelligent information processing.

1.1 Outline of the RWC Program

The Real World Computing (RWC) program started in 1992 as the successor to the Fifth Generation Computer project. This is the large-scale ten-year Japanese national project launched by MITI with a budget of over \$500 million for ten years. Whereas the Fifth Generation project pursued the logical (symbol-based) aspect of information processing (or intelligence) of humans, the RWC program is rather pursuing the intuitive (pattern-based) aspect of information processing and also aiming at unifying both aspects in a bottom-up manner within a new framework and foundation for next-generation information processing.

The primary objective of the RWC program is to lay the theoretical foundation necessary to pursue the technological realization of human-like flexible, intelligent information processing, capable of directly and flexibly processing various kinds of information in the real world. The aim is to build a new paradigm of information processing for the highly information-based society of the twenty-first century (MITI, 1992; Otsu, 1993).

In July 1992 the *RWC partnership* (RWCP) was founded, and fifteen Japanese companies including almost all major electronics firms have since joined, with more than thirty contract research themes making it through the reviewing process. In October 1992, the RWCP founded its own central laboratory, Tsukuba Research Center (TRC), near ETL (Electrotechnical Laboratory) in Tsukuba City, expecting close cooperation with ETL and inviting about twenty researchers from laboratories of each company. The RWC program also opened the door to foreign countries, and four research institutes—GMD: German National Research Center for Information Technology, SNN: Stichting Neurale Netwerken (the Netherlands), SICS: Swedish Institute of Computer Science, and Kent Ridge Digital Laboratories (Singapore)—participate as contractors, and some others as subcontractors (details are available at <http://www.rwcp.or.jp/home-E.html>). ETL, which belongs to MITI and played an important role in the early conceptualization of the program, continues to support and lead the program, having sent some researchers to the main positions in TRC and also carrying out its own leading-edge research on RWC with a group of about fifty researchers.

In Phase I (1992–1996), exploratory research was done in three layered fields:

1. novel functions for application;
2. theoretical foundation; and
3. parallel and distributed computational bases.

Via the interim review in 1996, Phase II (1997–2001) has focused on more

integrative and intensive R&D, allotting resources to two domains:

- RWI (real-world intelligence) and
- PDC (parallel and distributed computing).

The RWI domain, which is promoted by RWI Research Center in ETL/MITI, aims to add real-world intelligence (learning and information integration capabilities) to the conventional information-processing technology in order to expand its application horizon. Main R&D topics include theoretical bases, application systems implementing novel functions (multi-modal computer-human interface systems, autonomous learning, mobile systems, etc.), and supporting hardware (adaptive devices such as evolvable hardware and smart pixels) and software (intellectual resources such as databases and library; Otsu, 1998). In the PDC domain, which is promoted by TRC/RWCP, aims to provide parallel and distributed computational bases (seamless computing systems) for improving general use of the currently expanding global computer network infrastructure. Main R&D topics include computer architecture (such as PC/WS cluster and optical interconnection), software environments (such as SCORE, MPC++), and some PDC applications such as large-scale simulation, data retrieval, and parallel protein information analysis (PAPIA).

1.2 Real-World Intelligence

1.2.1 Objectives

As has already been stated, the goal of the research and development of real-world intelligence in the RWC program is to develop fundamental technologies for modeling and understanding humanlike flexible intelligence, including abilities of information integration and learning, and combining these with the conventional information technologies, thereby broadening the potential and application domain of information processing.

Unlike the computer, humans and other living creatures develop quite flexible intelligence that functions adaptively in the real world. The brain acquires such real-world intelligence through interactions with the real world via pattern information. Symbolic intelligence is actually formed on this real-world basis.

The essential functions of real-world intelligence are

- information integration, and
- learning/self-organization.

Information integration involves processing various kinds of information in the real world that contain ambiguity and uncertainty, such as im-

ages and sounds, in an integrated and parallel manner, so that such information can be used in recognition, understanding, overall judgment, and decision-making about behavior. *Learning/self-organization* involves systems' adapting or evolving their own functions by autonomously collecting information through interactions with the real world.

Information integration and learning/self-organization have been studied separately so far in theoretical contexts, but the special accomplishment of real-world intelligence lies in exercising these abilities in a consistent and integrated manner. In addition, information integration and learning/self-organization should be integrated with recognition and reasoning functions. Their integration in a theoretical framework requires stochastic and statistical methods to deal with uncertain information in the real world. Using these information integration and learning/self-organization functions as a common framework, we need a system that closely ties the research on theoretical foundations and research and development with actual applications of the theory.

1.2.2 Approach and Methodology

What kind of approach or methodology should we take in trying to realize real-world intelligence? To put it briefly and directly, we need to use the stochastic (probabilistic and statistical) and parallel/dynamic computing approach, in order to cope with large-scale but partial, uncertain, and incomplete information in the real world. Conventional methods of information processing are based on the computer's logic, i.e. procedural and serial processing (a linear method), whereby input information is converted into output information as an algorithm. This is a direct method with a rigid framework and is very efficient at handling well-defined problems; however, it lacks flexibility and often breaks down when dealing with real-world situations. In contrast, multilayer neural networks (Rumelhart et al., 1986) or multivariate data analysis are based on a feedback loop for optimization with evaluation by including degrees of freedom as parameters in the processing. With this, the method learns adaptively the optimal processing from examples. It is, in a sense, a forward adaptation method. On the other hand, model fitting (Geman and Geman, 1984; Hopfield and Tank, 1985) or regularization theory (Poggio and Koch, 1985) may be considered backward adaptation methods. These assume a parameterized model as an ideal output. With this, the method evaluates the discrepancy of the model, updating it based on input information, achieving optimal though implicit processing.

As we consider real-world intelligence, we must take such new techniques into account. Key issues are a mechanism for evaluation and a feedback loop for optimization, and both should be installed in the sys-

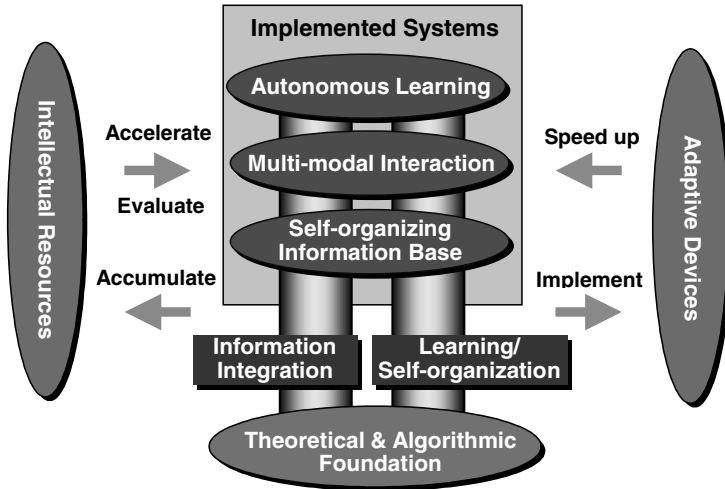


FIGURE 1.1. Scheme of R&D in the RWI domain.

tem if the system is to be adaptive and autonomous. It is also necessary to consider pattern recognition in thinking how the brain deals with pattern information in the real world. The basis for learning and inference, pattern recognition is very important for intelligence. It is located at the “front end” of intelligence, where it meets the real world, bridging patterns and symbols. From a logical perspective, pattern recognition has an inductive phase of learning and also a deductive phase of decision making. In addition, pattern recognition focuses essentially on parallel processing or overall judgment.

Behind the flexible and intelligent processing of information, in areas such as pattern recognition, neural computation, regularization, and stochastic inference, lies a Bayesian framework (Otsu, 1982, 1989; Geman and Geman, 1984). Multivariate data analysis represents such things most simply as a linear model; what allows for nonlinear functionality to some extent is a neural network. Its extremity is the Bayesian inference (Otsu, 1982).

1.2.3 Research and Development

For pursuing the above objectives, we have established six core areas for research and development (figure 1.1). Three of them are system oriented, covering basic information-processing abilities (such as recognition and understanding, reasoning, and control) and typical application domains, and the other three areas support the system-oriented areas.

To promote and encourage the cooperative R&D in those areas, workshops are being organized and meetings are often held within each area and across areas. Participants are researchers from contracted companies, university professors as subcontractors, and researchers from the RWI Research Center in ETL as chairs.

Autonomous Learning System. Autonomous learning system is a physical agent (e.g. a mobile robot) that moves around autonomously in real environments, acquires and learns information related to its environment (including people) by sensing and interacting with the environment, and can be developed into a system that offers services according to needs. This system is, in a sense, considered as a test bed for integrating various novel functions that are being developed. R&D topics include sensor- and vision-based navigation, map learning, vision for object recognition, understanding of speech and sounds in noisy environments, knowledge acquisition and representation, multi-agent architecture, action planning, and so on. A prototype system, Jijo2, is being developed in ETL.

Multi-modal Interaction System. A multi-modal interaction system is a personified agent-type computer-human interface that enables users to communicate with computers or information systems in a natural way, by integrating multi-modal information (speech, images, gestures, etc.). R&D topics are understanding and synthesis of images (in 2-D and 3-D), speech sounds, facial expressions, hand sign language, and inter-modal learning for concept formation by complementary use of imagery information and language (e.g. the way infants and children learn concepts of objects in the real world with assistance from mothers). It is hoped that this kind of system will become a next-generation interface system that will allow everybody to communicate easily with computers and information systems everywhere.

Self-organizing Information Base System. A self-organizing information base system is one that can retrieve, sort, summarize, and present diverse and vast amounts of information in the real world or in information networks in a self-organizing manner for supporting users' intellectual activities. This technology is urgently required in the current information society, especially given the flood of information typically encountered on the Internet. R&D topics are clustering of large-scale data such as contents of newspapers or TV, 3-D browsing and user interface, and mutual retrieval between different media such as texts, images, and speech sounds. Work is also being done to develop a common format in which to represent and treat such multimedia information in a unified manner.

Theoretical and Algorithmic Foundations. Theoretical and algorithmic foundation is the target area of this volume. The area concerns the theoretical foundations of information integration, learning and self-organization, and optimization techniques that support RWI systems. The theoretical bases of the technology advances are of great importance in bringing about a true breakthrough and prevent the end result of these technologies from becoming a simple collection of heuristics. Extensive research is being carried out on neural networks, Bayesian networks, genetic algorithms, EM algorithm, HMM, ICA, and so on. A truly new technology will be created through the close interaction and collaboration of empirical work with application-oriented R&D.

Real-World Adaptive Devices. Systems in the real world need to operate adaptively in real time. It is often difficult to achieve this through only the application of software; hardware support is essential. This does not mean that we need supercomputers, but we may need certain specialized hardware devices (chips) that contribute to the realization of real-time and adaptive processing. New types of hardware are being developed, such as reconfigurable hardware (RHW: next generation FPGA), evolvable hardware (RHW+GA), and smart pixel (optical vision chip). Real-world application systems implemented with the adaptive devices are also being developed in ETL, for example, evolver (an autonomous mobile robot with gene), EMG-controlled prosthetic hands, and so on.

Intellectual Resources. Support from the software side is also important. That is, we need to provide and maintain various databases, benchmarks, and a library of software, in order to accelerate and evaluate the R&D of real-world intelligence technology. This area is called the intellectual resources for R&D. For instance, real-world databases such as images, speeches, and texts (corpus) are important for the design and evaluation of the real-world intelligence systems and techniques being developed. It is also important to archive acquired common techniques and programs in the software library; doing so will accelerate R&D and will also leave the results of R&D in a stable form for coming generations. Databases already developed are available in CD-ROM, and some of the software library will be open to the public on the Web.

1.3 Concluding Remarks

Conventional artificial intelligence in a narrow sense is a top-down approach. Based on deductive formal logic and symbol manipulation, it is too artificial to model human intelligence by itself. In contrast, the recent paradigm shift in neural networks (connectionist AI) and in artificial life (such as genetic algorithms, complex systems, and emergent

computation) is a bottom-up approach, based on pattern recognition or subsymbolic dynamic computation, shifting the focus toward the inductive aspect of intelligence—toward learning and self-organization. However, what is essential to the project of modeling real-world intelligence is the computational (or cognitive) aspect of intelligence; we need to clarify and confirm the unified principle of computation behind these methods, to deepen them theoretically, and to implement them as more powerful computational paradigms. There already exists a framework of Bayesian inference, decision-making (Otsu, 1989), and other types of stochastic computation. The project also tends toward furthering this existing research.

A consequent and natural question is the following: how will AI research in the twenty-first century progress? An initial answer is that it will further investigate distributed and cooperative intelligence based on embedded systems, which corresponds to the further development of the global information network and of ubiquitous, wearable intelligent information appliances. The RWC program, in particular the R&D in the RWI domain, is actually providing technological bases for these future developments.

2 Theoretical and Algorithmic Foundations of Real-World Intelligence

HIDEKI ASOH

What are the theoretical and algorithmic foundations for realizing intelligent systems?

About ten to twenty years ago, the main answer was *symbolic logic* and *production rules*. Logic programming (declarative programming, constraint programming) and production systems were considered a very powerful tool for representing knowledges about the world and making inferences. Many kinds of logics, inference algorithms, and expert systems were investigated. The 5th Generation Computer Project developed special hardware that could handle very fast logical inference.

Logical statements are very powerful as representations of knowledge and rules that can be explicitly written down in language, and of the inferences we make with them. However, there are many types of implicit knowledge that are difficult to represent as logical statements or linguistic rules. For one thing, such knowledge often includes noise and contradictions. We human beings have huge amount of implicit knowledge that is indispensable to our surviving in the real world. For example, we can recognize a visual image of face, but we cannot explain why we can or just how we recognize it. Pattern recognition research has mainly

been tackling such implicit knowledge. One of its major theoretical foundations is statistics.

Real-world intelligence systems should handle complex, uncertain, dynamic, multi-modal information in the real world. Both explicit and implicit types of knowledge are important here. Hence we need to develop a novel integrated framework that can represent our knowledge and inferential abilities using both types of knowledge. In addition, given the enormous complexity and variety of the real-world environment, it is impossible to preprogram all knowledge. Hence a learning capability is crucial to modeling real-world intelligence. Learning, like the process of evolution, is a kind of meta-programming strategy; instead of writing target programs themselves, we implement learning programs that themselves generate and modify target programs according to the interaction between system and environment.

2.1 Objective

The objective of the research activities in the area of the theoretical and algorithmic foundations of RWI is to establish the theoretical foundations of the two key technologies, *information integration* and *learning/self-organization*. Our aim is to provide novel schemes and algorithms for acquiring knowledge, representing knowledge, making inferences with that knowledge, and using that knowledge to interact with the environment and other users. Not only proposing new schemes and methods but also analyzing the performance of the methods theoretically and empirically, and efficiently implementing the methods and providing them to the researchers creating prototype systems of real-world intelligence are important objectives.

2.2 Approach

Approaches to the theoretical foundations of intelligent systems can be divided into two groups: rule-based and normative, or model-based. The former approach, based on symbolic logic, is used mainly in conventional AI in areas such as knowledge engineering, natural language processing, game playing, and so on. The latter is used in pattern recognition and is based on statistics and probability theory.

The two approaches are coming closer together recently, as the necessity of handling real-world information becomes consensus among the researchers of intelligent systems. Real-world information is multi-modal, noisy, uncertain, dynamic, and complex. To handle such information, both approaches need to be extended. For example, in some rule-based expert systems such as MYCIN, measures of certainty of information called certainty factor have been introduced. More recently, inductive

logic inference is being investigated to try to make the systems learn logical statements from experiences.

Simple statistics based on Gaussian distributions or mixture of Gaussian distributions, which were major tools in statistical pattern recognition, turned to be not powerful enough to treat complex information such as speech signals or motion images. For example, in speech recognition systems, it seems that more complex probability models such as the hidden Markov model or neural networks play an important role. Probability distributions with combinatorial structure, called graphical models, have been investigated intensively in recent years.

The main approach we take here is the latter one: a normative approach based on statistics and probability theory. One reason to choose it is its theoretical soundness. At the same time, we remain open to novel alternatives. Explorative studies seeking new schemes of representing information, making inferences, and learning are also being done.

2.3 Research Issues

In order to review various research themes investigated in the RWI project, we will organize them in the following six categories and then briefly describe major research issues.

- Models for representing knowledge
- Algorithms for inference and integration
- Algorithms for learning/self-organization
- System architectures for information integration and learning
- Frameworks of interaction for learning
- Test-bed applications for empirical evaluation

2.3.1 Model

A model is a general scheme or language for representing various constraints between variables. It is also used to describe the generation process of observed data. Here we focus mainly on probabilistic relationships between variables and exploit as models families of probability distributions with graphical structure.

There is no universally superior model. Each model has its own characteristics. Hence our mission is to clarify the characteristics of models and make predictions of their performance in various specific but somehow general situations.

In addition to probabilistic models, we also investigate novel information representation schemes that are appropriate for information integration.

2.3.2 Algorithms for Inference and Integration

Since computations of probability using a structured probability distribution model is computationally intensive, efficient approximative algorithms are necessary to realize the real-time response of systems. Anytime algorithms, which can output approximative results at any requested moment, are preferable. The active control of the probability computing process is also important for the efficiency of the computation.

Integrating information from multiple sources can be considered a special case of inference. However, the vast variety of the information sources makes the problem more difficult.

2.3.3 Algorithms for Learning/Self-Organization

To learn is to acquire proper constraints to behave properly in environment. If a criterion for evaluating the appropriateness of constraints is given, the learning problem reduces to an optimization problem. A well-known criterion in statistics is maximum likelihood. However, the likelihood function is not appropriate as a criterion for selecting structure of models because it will always select the most complex model.

The models treated in this project have a graphical structure and so structure selection is indispensable for them to provide a good representation of knowledge. Hence we need to determine criteria for model selection and learning.

When the model is complex, the search for good representation tends to become computationally intractable. Thus we need heuristics that work in typical settings. Two of current major trends in machine learning research are distributed learning and active, explorative learning such as boosting. Genetic algorithms can be considered a kind of active distributed learning mechanism.

In many learning procedures, computing statistics from data is the most computationally costly part. Approximation techniques for statistical inference introduced from statistical mechanics are effective here, as well.

2.3.4 System Architecture

For the extensible systematic realization of a multi-modal information integration system, a flexible system architecture is necessary. Multi-agent software architecture is a candidate for implementing complex information processing.

2.3.5 Frameworks of Interaction for Learning

Although distributed learning or active learning is a powerful learning mechanism, acquiring the proper model structure is still a very difficult

problem. To overcome this difficulty, we need to develop novel frameworks for learning.

2.3.6 Test-Bed Applications

To evaluate algorithms, we need test-bed applications. Several application problems are selected from various domains such as genetic information processing, decision support systems, natural language processing, and robotics. “Test-bed” does not mean just “toy problems” or “small-scale problems.” Some of research issues are directly related to systems oriented research in other RWI research areas, and some of them treat large-scale real-world data.

2.4 Organization of R&D and This Book

As will be shown in the following chapters, various research issues in the theoretical and algorithmic foundations area have been investigated by five distributed laboratories in RWC partnership (RWCP) and one laboratory in RWI research center at ETL. Each laboratory covers multiple research issues from representation of knowledge to test-bed application. In what follows, we briefly summarize their goals and major results as an introduction to the following chapters.

2.4.1 Inference and Learning with Graphical Models (ETL Lab)

The ETL Lab aims to establish the foundations of learning and integrated information processing for real-world intelligence that can interact closely with the real world, including humans, and execute flexible information processing. Various models, such as Bayesian networks, probabilistic constraint programs (an extension of constraint program that can treat statements of the first order predicate logic with probability), and mixture of Gaussians, are explored by the lab.

A combination of Bayesian networks and neural networks developed by this lab is able to extend the applicability of the Bayesian networks to real-world continuous-valued data. The algorithm is also being implemented in Java language with an elegant graphical user interface and flexible data-base access capability. Another remarkable result is a new compilation mechanism for probabilistic constraint programs.

Another approach to graphical modeling, multivariate information analysis, is also being pursued. Several methods have been proposed to determine the relationship between random variables, including analyzing information-theoretic measures such as mutual entropy between variables. The methods are applied to the brain image data from functional MRIs, which describe the activity of the brain during cognitive tasks to reveal the relations between the activities of several parts of a brain.

Investigating novel learning schemes that interact more closely with the environment is another major research issue in this lab. Two new schemes have been proposed and are being implemented. The first is dialogue-based learning, or socially embedded learning, which uses communication between learning systems and human users intensively. The second is intermodal learning, which exploits the multimodality of the sensory-motor flow of real-world intelligent systems. There it is expected that integrating various parts of the structure that are acquired from different modalities of the sensory-motor flow will result in a complete structure of knowledge.

The performance and characteristics of the methods are analyzed theoretically and empirically. The developed methods are implemented efficiently and provided to the researchers to build prototypes of real-world intelligence.

In chapter I of this book, following an overview of the R&D in the ETL Lab, three major results, i.e. the combination of the Bayesian networks and neural networks, multivariate information analysis, and dialogue-based learning, are described in detail.

2.4.2 Probabilistic Knowledge Representation and Active Decision (SNN Lab)

Chapter II discusses the work of the SNN Lab. The aims of the research in this laboratory are to develop novel theory, techniques, and implementations for learning and reasoning in a complex dynamic multisensory environment. The approach to reasoning and learning is based on the axioms of probability theory and Bayesian statistics. Boltzmann Machines and Bayesian networks are investigated as main target models. The lab pursues efficient approximative algorithms for such graphical models, incorporating elegant methods from statistical mechanics such as mean field approximation. Efficient learning algorithms for Boltzmann Machines and Bayesian networks are developed and evaluated.

These novel methods are demonstrated in two real-world applications: medical diagnosis and music transcription. Large-scale probabilistic networks are constructed and applied to the diagnosis of anemia under the collaboration with University Hospital Utrecht.

2.4.3 Reflective Teams: Active Learning and Information Integration in Open Environments (GMD Lab)

The goal of the GMD Lab research on reflective teams (R-Teams) is to build complex artificial reality systems that provide new insights into real-world processes. This requires:

- building the theoretical foundations for real-world intelligent systems by developing new heuristics and algorithms based on soft computing embedded in R-Teams architecture;
- demonstrating the benefits of R-Teams architectures for the integration of real-world technologies; and
- implementing real-world applications like distributed problem solving in a network of computational agents, mobile radio networks, and city traffic simulations in R-Teams architectures.

The lab has proposed a new optimization algorithm. As a first step toward designing and evaluating decentralized heuristics for agents interconnected by a network, GMD Lab. has chosen a benchmark application—the optimization of decomposed discrete functions. They use two different agent types—probabilistic and intelligent—for generic decentralized solution methods. They have developed the probabilistic methods UMADA and FDA (factorized distribution algorithm), extensions of genetic algorithms. The lab has analyzed the dynamic behaviors of algorithms UMADA and FDA.

In a next step, the lab implements the R-Teams optimization method (*RTeam-Opt*) based on game theory. The agents are active and reflective. Each agent is responsible for a component of the fitness function. The global fitness function is unknown to the individual agent. The agents use local optimization, cooperation with neighbors, and reflection. RTeam-Opt more efficiently optimizes most of the fitness functions solved by FDA. This method is applied to the optimal antenna placement problem and its effectiveness is demonstrated.

Another result is multiagent-based distributed software architecture that has learning capabilities. The basic software units of the architecture are agents that are activated periodically and thus perform cycles as if they were making *ticks* like a clock. They can access tag boards for posting or reading messages, called tags. Moreover, they control when a tag board is to be *flipped*, i.e. turned over from the current display to an updated display. This collection of organizing principles and generic functional units is called the *flip-tick architecture* (FTA) for R-Teams applications. The architecture is cycle-based. It supports one of the most common operation principles of real-world intelligent systems: to efficiently perform cycles in large data spaces as well as in complex control loops. *R-Teams architecture* denotes the structure and design of application systems that are based on teams of reflective FTA agents. Such agents are able to assess their situation and the current context of op-

eration and to draw conclusions for their future behavior, e.g. through adaptation or learning. This multiagent microsimulation architecture is applied to the problem of simulating city traffic in Bonn.

Based on the various activities of the GMD Lab described above, the theoretical foundation of evolutionary algorithms for optimization is fully described in chapter III.

2.4.4 Distributed and Active Learning (NEC Lab)

Chapter IV presents the work of the NEC Lab. Rigorous theoretical investigations on three topics, distributed cooperative Bayesian learning, learning specialist decision lists, and the lob-pass problem, are described. The goals of the research in this laboratory are to investigate and develop efficient and robust methods of distributed and active learning. New committee-based query learning methods and distributed cooperative Bayesian learning strategies are recognized as important progresses in this area. The performance of the methods is evaluated both theoretically and empirically, and rigorous theoretical analyses of the efficiency of the algorithms are given.

A new criterion, extended stochastic complexity, that extends the model selection criterion of MDL (minimum description length) from a Bayesian decision theoretic point of view has been applied to distributed cooperative learning strategies.

As a test-bed application, the development of the efficient word-clustering method and applications to syntactic disambiguation is a remarkable result. The method is applied to large scale tagged text corpus to show applicability to the real-world data and the world record performance are demonstrated. Probability distributions suited for natural language processing are explored. Another test-bed application is genetic information processing.

2.4.5 Stochastic Pattern Computing (SICS Lab)

Chapter V presents the research of the SICS Lab on computing with large random patterns and its application to natural language understanding. Goals of this laboratory are

- to understand natural (i.e. the brain's) mechanisms of intelligence, and
- to build computers that employ such mechanisms, based on holistic representation (combining syntax and semantics in a uniform representation) and on stochastic computing (combining the strengths of numeric and symbolic computing in a system that computes with very large pseudorandom patterns and achieves reliability by statistical laws of large numbers).

Sparse distributed representation is proposed and investigated. One major development is the *spatter code* and the *sparchunk code*. The spatter code is a method of forming analyzable holistic representations of higher-level concepts using dense binary vectors, i.e. vectors with roughly the same number of 1s as 0s. The sparchunk code solves the same problem with sparse codes. One advantage with this code is that the SDM memory can be used for breaking a concept down into its constituents. Both these methods can be used hierarchically. An algorithm for associating information represented by sparse vectors is proposed and applied to natural language understanding.

2.4.6 Symbol-Pattern Integration (Toshiba Lab)

The Toshiba Lab studies fundamental philosophical issues on the integration of symbols and patterns. Specifically, they aim at modeling the logical reasoning of patterns. Here a pattern is represented by a function and approximated by a neural network. Then the neural network is interpreted as a multilinear function. To make reasoning out of patterns, the lab investigates the logical structure embedded in the space of multilinear functions, collaborating with JAIST and Chiba University. Some fundamental relationships between multilinear function space and BCK logics have been revealed, and a complete many-valued logic with product-conjunction is selected as a suitable logic for making inferences with pattern information. The result is then applied to various real-world data such as financial time series.

2.5 Concluding Remark

This section provides an overview of the R&D in the theoretical and algorithmic foundations of the RWC program. Many important research topics have been covered by six distributed laboratories and five of them are described in more detail in the following chapters.

We hope that the readers will enjoy the following chapters and that the results of this project will become fundamental tools to open the door to the new world of information technology.

References

- Geman, S. and Geman, D. (1984). Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence* PAMI-6:721–741.
- Hopfield, J. and Tank, D. (1985). Neural computation of decisions in optimization problems. *Biological Cybernetics* 52:141–152.
- MITI (1992). *The Master Plan of RWC Program*. Edited by Industrial Electronics Division, Machinery and Information Industries Bureau, Ministry of International Trade and Industry.

- Otsu, N. (1982). Optimal linear and nonlinear solutions for least-square discriminant feature extraction. *Proceedings of 6th International Conference on Pattern Recognition*, pp. 557–560.
- Otsu, N. (1989). Toward soft logic for the foundation of flexible information processing. *Bulletin of ETL* 53(10):75–95.
- Otsu, N. (1993). Toward flexible intelligence: MITI's new program of real-world computing. *Proceedings of 13th International Joint Conference on Artificial Intelligence*, pp. 786–791.
- Otsu, N. (1998). Project for fundamental information technology of the next generation: Real-world intelligence technology. RWC-98 Symposium (in RWC News 4-E, available at <http://www.rwcp.or.jp/rwc-news/04e/ohtsu.html>).
- Poggio, T. and Koch, K. (1985). Ill-posed problems in early vision: From computational theory to analogue networks. *Proceedings of Royal Society of London* B226:303–323.
- Rumelhart, D., Hinton, G., and Williams, R. (1986). Learning representations by back-propagating errors. *Nature* 323(9):533–536.
- RWCP (1997). *Proceedings of 1997 Real World Computing Symposium.*, Tokyo: Real World Computing Partnership.
- RWCP (1998). *Proceedings of 1998 Real World Computing Symposium.* Tokyo: Real World Computing Partnership.
- RWCP (2000). *Proceedings of 2000 Real World Computing Symposium.* Tokyo: Real World Computing Partnership.