


Intelligent Hybrid Systems



Edited by

Suran Goonatilake · Sukhdev Khebbal

Intelligent Hybrid Systems

Edited by

Suran Goonatilake and Sukhdev Khebbal

University College London, UK



JOHN WILEY & SONS

Chichester · New York · Brisbane · Toronto · Singapore

Copyright © 1995 by John Wiley & Sons Ltd.
Baffins Lane, Chichester
West Sussex PO19 1UD, England

National Chichester (01243) 779777
International (+44) 1243 779777

Chapter 2 © Honeywell Inc.
All rights reserved.

No part of this publication may be reproduced by any means,
or transmitted, or translated into a machine language
without the written permission of the publisher.

Other Wiley Editorial Offices

John Wiley & Sons, Inc., 605 Third Avenue,
New York, NY 10158-0012, USA

Jacaranda Wiley Ltd, 33 Park Road, Milton,
Queensland 4064, Australia

John Wiley & Sons (Canada) Ltd, 22 Worcester Road,
Rexdale, Ontario M9W 1L1, Canada

John Wiley & Sons (SEA) Pte Ltd, 37 Jalan Pemimpin #05-04,
Block B, Union Industrial Building, Singapore 2057

British Library Cataloguing in Publication Data

A catalogue record for this book is available from the British Library

ISBN 0 471 94242 1

Produced from camera-ready copy supplied by the editors
Printed and bound in Great Britain by
Redwood Books, Trowbridge, Wiltshire.

Foreword

“If the oak tree flowers before the ash, prepare for a splash: if the ash tree flowers before the oak, prepare for a soak”

(Source: English folklore)

This rule is an example drawn from hundreds of apparently simple and often contradictory rules for predicting future observations of the notoriously fickle English weather. On closer examination the rule is less simple than it first appears. The variables in the rule are ‘fuzzy’; there is no definition of either what ‘rain’ is or what is meant by flower. Even the observation frequency is left to the user’s discretion. People, of course, regularly use sets of rules to approximate to and infer from observations, the principles which govern the evolution of complex systems like the weather. Although this example is taken from folklore weather forecasting, it could just as easily have been drawn from the financial markets, where combinations of simple rules based on the analysis of time-series are routinely used to support trading decisions, or from medicine, where classifications of symptoms must be made in the light of incomplete readings and often conflicting information. It is in solving these types of difficult real-world problems that a new generation of intelligent computing methods, Intelligent Hybrid Systems, are now bearing fruit.

The grander claims for rule-based AI (Artificial Intelligence) in the 80’s seemed to fail in part because specialist practitioners either failed to understand or failed to flag the limitations of their single approach. The contributors to this volume have supported a different view: there is no one definitive approach which will explain cognition or solve complex problems, and there is no ‘philosopher’s stone’; instead, there are a number of tools and models that can be applied under different circumstances. The proof of this philosophy is to be found in the following chapters that detail successes ranging from aircraft design optimisation to stock portfolio selection.

Increasingly, complex industrial and business problems are driving the research and development of new intelligent computing methods. At Reuters we use an intelligent system for checking data-errors in the flood of data we collect from Stock Exchanges around the world. Some of our financial market customers have recently

experimented with neural network and genetic algorithm forecasting tools which draw data from our financial data-feeds. In other companies, credit evaluation and credit card fraud detection are now routinely subject to the analysis of intelligent techniques to help operators make better decisions. With the convergence of media and technology one expects further uses of hybrid systems to analyse, distil and help make decisions on the so-called information superhighway.

This timely book reviews the current state of intelligent hybrid systems. It offers both a classification scheme to analyse and describe hybrid models and presents clear arguments for their need. It is above all a book about how to move from theory to practical implementations. I believe it will inspire academics and technologists alike to take a dip in the waters of this emerging field.

John Taysom

Vice President
Reuters New Media Inc.
Palo Alto

Acknowledgements

This book owes much to many people. Firstly, we would like to take this opportunity to thank all of our contributors for their time and effort in preparing their papers. A very special thanks goes to Bob Ijaz for his enthusiasm and tireless effort in formatting the book. Many thanks to the staff at John Wiley for making the book possible, especially to Gaynor Redvers-Mutton for her patience and guidance throughout the project. We are also grateful to Charles, Sinclair and Co. for the use of their printing facilities.

We would also like to thank our colleague Konrad Feldman for proof reading all the chapters and for making constructive technical comments. We also like to acknowledge the Department of Computer Science at University College London and our colleagues in the Intelligent Systems Group, for providing a stimulating environment for our explorations into hybrid systems.

Finally, we would like to thank Jim Hendler, Lawrence Davis, Tariq Samad, Henry Tirri and Andreas Scherer for encouraging the very initial idea of a book on intelligent hybrid systems.

Contributors

Venkat Ajjanagadde,

Wilhelm-Schickard-Institut,
University of Tuebingen,
Sand 13 W-7400 Tuebingen,
Germany.

Frederic Gruau,

Centre d'etude nucleaire de Grenoble
DRFMC SP2M PSC,
BP85X 38041
France.

Charles Karr,

U.S. Bureau of Mines,
Tuscaloosa Research Center,
P.O. Box L,
University of Alabama Campus,
Tuscaloosa, AL 35486-9777,
U.S.A.

Casimer C. Klimasauskas,

Neuralware,
Penn Center West IV-227,
Pittsburgh,
PA 15276,
U.S.A.

David Montana,

Bolt Beranek and Newman
70 Fawcett Street
Cambridge,
MA 02138,
U.S.A.

Wendy Foslien,

Honeywell SSDC,
3660 Technology Drive,
Minneapolis, MN 55418,
U.S.A.

Vasant Honavar,

Assistant Professor of Computer Science
and Neuroscience,
Computer Science Dept.
226 Atanasoff Hall,
Iowa State University,
Ames, IA 50011,
U.S.A.

Randy Kerber,

Lockheed AI Center,
3251 Hanover Street,
Palo Alto,
CA 94394
U.S.A.

Brian Livezey,

Lockheed AI Center,
3251 Hanover Street,
Palo Alto,
CA 94394
U.S.A.

Khair Minh Pham,

inferOne,
Immeuble Les Maradas,
1 Bd de l'Oise,
95030 Cergy Pontoise Cedex,
France.

David J. Powell,

Corporate Research and Development,
General Electric Company,
Schenectady, NY,
U.S.A.

Andreas Scherer,

Praktische Informatik I,
FernUniversitaet,
Postfach 940,
W - 5800 Hagen,
Germany.

Danny Shamhong,

Department of Computer Science
University College London
Gower Street
London WC1E 6BT,
U.K.

Evangelos Simoudis,

Lockheed AI Center,
3251 Hanover Street,
Palo Alto,
CA 94394
U.S.A.

Henry Tirri

Department of Computer Science
P.O. Box 26
SF-00014 University of Helsinki
FINLAND.

Leonard Uhr

Computer Sciences Dept.,
University of Wisconsin,
Madison,
Wisconsin,
U.S.A.

Tariq Samad,

Honeywell SSDC,
3660 Technology Drive,
Minneapolis, MN 55418,
U.S.A.

Gunter Schlageter,

Praktische Informatik I,
FernUniversitaet,
Postfach 940,
W - 5800 Hagen,
Germany.

Lokendra Shastri

International Computer Science Institute
1947 Center Street,
Suite 600 Berkeley,
CA 94704-1105,
U.S.A.

Michael M. Skolnick

Assistant Professor
Department of Computer Science
Rensselaer Polytechnic Institute
Troy, NY., 12180
U.S.A.

Siu Shing Tong,

Corporate Research and Development,
General Electric Company,
Schenectady, NY.
U.S.A.

About The Editors

Dr. Suran Goonatilake is at the Department of Computer Science at University College London. He is a co-founder of SearchSpace Ltd., a London-based company that specialises in applying intelligent systems in finance and business. He holds a PhD in Computer Science from University College London and a bachelor's degree in Computing and Artificial Intelligence from the University of Sussex. His PhD research was on the use of genetic-fuzzy hybrid systems for supporting financial trading. His current research interests include the use of genetic algorithms for decision rule induction and their application in insurance risk assessment, direct marketing and portfolio management. Suran is also the co-editor of the book "Intelligent Systems for Finance and Business".

Dr. Sukhdev Khebbal is a co-founder of SearchSpace Ltd. and a project leader in the Department of Computer Science, University College London. His current research interests include the theoretical aspects of intelligent hybrid systems, object-oriented programming and intelligent agents. His PhD research was on the design of object-oriented environments for integrating neural networks and expert systems. He is the principal designer of the HANSA architecture (Heterogeneous Application geNerator Standard Architecture) which is the core framework of the largest European Community funded project in hybrid systems. HANSA allows software developers to rapidly generate decision support systems using industry standard tools and intelligent systems. To complement this research background, Sukhdev has ten years of practical experience in the application of computing techniques to business problems.

The authors can be contacted at:

S.Goonatilake@cs.ucl.ac.uk

S.Khebbal@cs.ucl.ac.uk

Contents

Foreword	vii
Acknowledgements	ix
Contributors	xi
1 Intelligent Hybrid Systems : Issues, Classifications and Future Directions	1
SURAN GOONATILAKE AND SUKHDEV KHEBBAL	
<hr/>	
PART ONE FUNCTION-REPLACING HYBRIDS	21
<hr/>	
2 Fuzzy Controller Synthesis with Neural Network Process Models	23
WENDY FOSLIEN AND TARIQ SAMAD	
3 Replacing the Pattern Matcher of an Expert System with a Neural Network	47
HENRY TIRRI	
4 Genetic Algorithms and Fuzzy Logic for Adaptive Process Control	63
CHARLES KARR	
5 Neural Network Weight Selection Using Genetic Algorithms	85
DAVID J. MONTANA	
<hr/>	
PART TWO INTERCOMMUNICATING HYBRIDS	105
<hr/>	
6 A Unified Approach For Engineering Design	107
DAVID J. POWELL, MICHAEL M. SKOLNICK AND SHI SHING TONG	
7 A Hybrid System for Data Mining	121
RANDY KERBER, BRIAN LIVEZEY AND EVANGELOS SIMOUDIS	
8 Using Fuzzy Pre-processing with Neural Networks for Chemical Process Diagnostic Problems	143
CASIMER C. KLIMASAUSKAS	

9	A Multi-Agent Approach for the Integration of Neural Networks and Expert Systems	153
	ANDREAS SCHERER AND GUNTER SCHLAGETER	
<hr/>		
PART THREE POLYMORPHIC HYBRIDS		175
<hr/>		
10	Integrating Symbol Processing Systems and Connectionist Networks	177
	VASANT HONAVAR AND LEONARD UHR	
11	Reasoning with Rules and Variables in Neural Networks	209
	VENKAT AJJANAGADDE AND LOKENDRA SHASTRI	
12	The NeurOagent: A Neural Multi-agent Approach for Modelling, Distributed Processing and Learning	221
	KHAI MINH PHAM	
13	Genetic Programming of Neural Networks: Theory and Practice	245
	FRÉDÉRIC GRUAU	
<hr/>		
PART FOUR DEVELOPING HYBRID SYSTEMS		273
<hr/>		
14	Tools and Environments for Hybrid Systems	275
	SUKHDEV KHEBBAL AND DANNY SHAMHONG	
	Glossary	309
	Intelligent Systems Resource Guide	313
	Index	321

1

Intelligent Hybrid Systems: Issues, Classifications and Future Directions

Suran Goonatilake and Sukhdev Khebbal

1.1 Introduction

Humans are hybrid information processing machines. Our actions are governed by a combination of genetic information and information acquired through learning. Information in our genes hold successful survival methods that have been tried and tested over millions of years of evolution. Human learning consists of a variety of complex processes that use information acquired from interactions with the environment. It is the combination of these different types of information processing methods that has enabled humans to succeed in complex, rapidly changing environments.

This type of *hybrid* information processing is now being replicated in a new generation of *adaptive machines*. The applications range from aircraft control systems that diagnose and repair themselves to systems that can successfully trade in foreign exchange markets. At the heart of these adaptive machines are *intelligent* computing systems, some of which are inspired by the mechanics of nature.

Neural networks, for example, are inspired by the functionality of nerve cells in the brain. Like humans, neural networks can *learn* to recognise patterns by repeated exposure to many different examples. They are good at recognising complex patterns, whether they be hand-written characters, profitable loans or good financial trading decisions. Genetic algorithms are also naturally inspired and based on the biological principle of “survival of the fittest”. The main idea behind a genetic algorithm is the

evolution of a problem's solution over many generations, with each generation having a better solution than its predecessor.

While these intelligent techniques have produced encouraging results in particular tasks, certain complex problems cannot be solved by a single intelligent technique alone. Each intelligent technique has particular computational properties (e.g. ability to learn, explanation of decisions) that make them suited for particular problems and not for others. For example, while neural networks are good at recognising patterns, they are not good at explaining how they reach their decisions. Fuzzy logic systems, which can reason with imprecise information, also have particular strengths and limitations. They are good at explaining their decisions but they cannot automatically acquire the rules they use to make those decisions. These limitations have been a central driving force behind the creation of *intelligent hybrid systems* where two or more techniques are combined in a manner that overcomes the limitations of individual techniques.

For example, there is now considerable interest in applying fuzzy logic for tasks such as loan evaluation. Typically, an expert in loan evaluation has to specify all the rules needed for the fuzzy logic system to make a decision on whether to grant a loan or not depending on the application details. However, this is a time consuming and error-prone process that should ideally be automated. Neural networks with their learning capabilities can be used to automatically learn these fuzzy decision rules, thus creating a hybrid system which overcomes the limitations of fuzzy systems.

Hybrid systems are also important when considering the varied nature of application domains. Many complex domains have many different component problems, each of which may require different types of processing. For example, the Kobe Steel Plant in Japan controls blast furnaces for making iron and steel by using a hybrid of different intelligent techniques to solve sub-tasks of the problem [Otsu92]. Neural networks are used to predict heat levels from sensor data gathered from the furnace. These predictions are then used by an expert system to infer the correct control adjustments so as to maintain the furnace at its optimal operating level. This complex control task was previously performed by a highly trained operator who had many years of experience.

Intelligent hybrid systems represent not only the combination of different intelligent techniques but also the integration of intelligent techniques with conventional computing systems such as spreadsheets and databases. For intelligent systems to add value to organisational decisions they must be able to extract and use information from a wide variety of sources. In addition, the decisions or results produced by the intelligent systems should be disseminated to existing applications or other systems for further processing. For these reasons it is vital that there are methods and protocols for integrating intelligent systems with other conventional computing systems. One such method is object-oriented programming [Wien88] which is a software engineering methodology that can provide the "glue" to join together different processing techniques. This forms a natural model for intelligent hybrid systems because individual techniques can be defined as objects which interact by sending a common set of messages.

This book presents a collection of work in this emerging area of intelligent hybrid systems, written by leading-edge researchers from around the world. It spans the combination of many different techniques including neural networks, genetic algorithms, expert systems, fuzzy logic systems and rule induction. Application examples are drawn from domains including industrial control, financial modelling, and cognitive modelling. The book's aim is to inform researchers and application developers, in both industry and academia, about the potential of intelligent hybrid systems for solving real-world problems.

The next section examines the strengths and weaknesses of different intelligent techniques which determine the need for hybrid systems. Section 1.3 introduces a classification scheme for hybrid systems and section 1.4 details a typical hybrid systems development cycle. Section 1.5 concludes by discussing potential areas of future development.

1.2 Properties of Intelligent Systems

In this section we compare and contrast different intelligent systems on four key information processing capabilities. These are knowledge acquisition, brittleness, higher and lower level reasoning and explanation.

1.2.1 Knowledge Acquisition

Knowledge acquisition is a crucial stage in the development of intelligent systems. As a process, it involves eliciting, interpreting and representing the knowledge from a given domain [Kidd87]. Knowledge acquisition for expert systems (from domain experts) is time consuming, expensive and potentially unreliable [Feig79]. Experts find it difficult to articulate particular types of “intuitive” knowledge [Part86] and are sometimes unwilling to participate in lengthy knowledge elicitation exercises. Other problems include the possible existence of gaps in an expert's knowledge and the correctness of an expert's knowledge [Jack86].

Furthermore, expert systems do not have mechanisms to deal with any changes in their decision making environment — they cannot adapt and learn from changes in their operating environment. Thus the *maintenance of knowledge* in expert systems is time consuming and expensive.

Due to these problems, techniques such as neural networks and genetic algorithms, which can learn from domain data, have certain advantages. In expert systems, the decision boundaries — the bounds used to make particular decisions — are specified by a domain expert, while in neural networks and genetic algorithms these decision boundaries are *learned* [Weis91]. Changes in the operating environment cause the decision boundaries to be shifted or changed. Systems that learn can detect and adapt to these changes.

Rule induction systems are another example of systems that have been applied to automate the knowledge acquisition process. They have been used to learn rules and decision trees from “raw” domain data and have been applied successfully in industrial and commercial domains [Quin86, Mich80].

1.2.2 Brittleness

Although there are notable successes of expert systems, many of these systems operate in very narrow domains under limited operational conditions. Holland [Holl86] refers to this phenomena in expert systems as *brittleness*:

The systems are *brittle* in the sense that they respond appropriately only in narrow domains and require substantial human intervention to compensate for even slight shifts in domain.

An operational view of the brittleness problem can be seen as the inability of an intelligent system to cope with inexact, incomplete or inconsistent knowledge. Causes of this brittleness problem are twofold — inadequate representation structures and reasoning mechanisms. In expert systems, knowledge is represented as discrete symbols and reasoning consists of *logical* operations on these constructs.

In contrast, reasoning in neural networks involves the numeric aggregation of representations over the whole network [Beal90]. As Smolensky [Smol87] points out,

Each connection represents a soft constraint; the knowledge contained in the system is the set of all such constraints. If two units have an inhibitory connection, then the network has the knowledge that when one is active the other ought not be. Any of the soft constraints can be overridden by others, they have no implications singly; they only have implications collectively.

This distributed representation and reasoning allows these systems to deal with incomplete and inconsistent data and also allow the systems to *gracefully degrade* [Rume86]. That is, even if some parts of a neural network are made non-operational, the rest of the neural network will function and attempt to give an answer.

This type of inherent fault tolerance contrasts strongly with expert systems which usually fail to function even if a single processing part is non-operational. This manifests itself in many ways, particularly as problems in the maintenance of large, complex knowledge bases [Mcde83].

Rule induction systems such as ID3 are also susceptible to the problems of brittleness. They have difficulties in inducing relationships in data which has contradictory and incomplete examples [Fors86]. However, recent enhancements of the ID3 algorithm have overcome some of the limitations with respect to its handling of noisy data [Quin93].

Fuzzy logic deals with the problem of brittleness by adopting novel knowledge representation and reasoning methods. Fuzzy sets, the form in which knowledge is represented, diffuse the boundaries between concepts. There are no sharp divisions where one concept ends and the next begins. This fuzzy data representation, in conjunction with fuzzy reasoning mechanisms, allows the processing of data which are non-exact or *partially correct* [Zade88].

The key argument for genetic algorithms being able to cope with brittleness is put forward by Holland [Holl86]. Holland argues that it is the maintenance of a *population of solutions* which makes genetic algorithms and classifier systems non-brittle. Each rule in the classifier system population contains a relationship describing

the system being modelled. The system's flexibility arises from the rules representing a wide range of competing, conflicting hypotheses. The selection of the appropriate rule to fire is dependent on its past performance — a statistical aggregation of its correct performance. Similar to neural networks, it is this *statistical reasoning* property, based on past performance that gives genetic algorithms their ability to cope with brittleness.

1.2.3 Higher and Lower Level Reasoning

For a complete theory of cognition, there need to be explanations of how we can do “low-level”, parallel, pattern recognition tasks as well as how we perform sequential, “high-level” cognitive tasks.

Although expert systems and related methods have provided plausible models to describe “high-level” cognitive tasks such as language generation and comprehension [Bode88], and goal directed reasoning [Bode88], their ability to explain “low-level” pattern recognition type tasks has been limited. In contrast, neural networks offer the exact complementary ability. They are very good at modelling pattern recognition tasks such as visual processing and motor control but are not well equipped to model sequential, high-level cognitive tasks. Because of these limitations, some researchers have criticised neural networks as inadequate models of cognition [Fodo88].

An analogue to cognitive modelling can be found in complex real world industrial and commercial problems. That is, most complex real world tasks are easily decomposed into logical, sequential operations and parallel, pattern recognition operations [Hand89]. For example in industrial robotics, neural networks are good candidates for performing low-level signal processing tasks such as the processing of sensory data while expert systems are good candidates for high-level path planning tasks [Thor90].

Similar to neural networks, fuzzy logic systems are also good at low-level signal processing tasks such as industrial control. They have recently been applied to hand-written character recognition, blast furnace control, robotic-arm control and camera auto-focus control [Mcne93].

Other tasks which have combined signal processing and high-level reasoning subproblems can be found in financial data analysis, medical diagnosis, and autonomous vehicle navigation.

1.2.4 Explanation

The ability to provide users with explanations of the reasoning process is an important feature of intelligent systems [Clan83]. Explanation facilities are required both for user acceptance of the solutions generated by an intelligent system, and for the purpose of understanding whether the reasoning procedure is sound [Davi77]. Good examples of this requirement can be found in medical diagnosis, loan granting, and legal reasoning. There have been fairly successful solutions to the explanation problem by expert systems, symbolic machine learning and case-based reasoning systems. In expert systems and rule induction systems, explanations are typically provided by tracing the *chain of inference* during the reasoning process [Sout91].

In a fuzzy logic system the final decision is generated by *aggregating* the decisions of all the different rules contained in the fuzzy rule-base. In these systems a chain of inference cannot be easily obtained, but the rules are in a simple to understand “IF-THEN” format which users can easily inspect.

Genetic algorithms, especially in the form of *classifier systems*, can build reasoning models in the form of rules [Holl87]. As in the case of expert systems, it is possible to trace a chain of inference and provide some degree of explanation of the reasoning process.

In contrast, in neural networks it is difficult to provide adequate explanation facilities. This is due to neural networks not having explicit, declarative knowledge representation structures but instead having knowledge encoded as *weights* distributed over the whole network [Died89]. It is therefore more difficult to find a *chain of inference* which can be used for producing explanations. This can be a particular problem in certain business and medical applications. For example, the provision of details about the mortgage lending decision process to potential borrowers, is now a legal requirement in certain countries [Hugh92]. This suggests that neural networks cannot be used. Similarly, in medical diagnosis physicians want to be absolutely sure about the detailed reasoning procedures used by an automated diagnostic system before a particular treatment is prescribed.

There is now a small but growing number of researchers who are attempting to provide neural networks with explanation facilities by extracting rules from their internal weight structures [Gall88]. Additionally there has also been progress in using “feature significance estimation” methods that identify the most important variables in a given neural network model [Intr92, Jabr92].

Table 1.1, shows a property assessment table that summarises the above computational properties with respect to individual intelligent techniques.

Technologies	Properties				
	Automated Knowledge Acquisition	Coping with Brittleness	High-level Reasoning	Low-level Reasoning	Explanation
<i>Expert Systems</i>	✓	✓	✓✓✓✓✓	✓	✓✓✓✓✓
<i>Rule Induction</i>	✓✓✓✓	✓✓	✓✓✓	✓✓	✓✓✓
<i>Fuzzy Systems</i>	✓	✓✓✓✓✓	✓✓✓	✓✓✓✓✓	✓✓✓✓
<i>Neural Networks</i>	✓✓✓✓✓	✓✓✓✓✓	✓	✓✓✓✓✓	✓
<i>Genetic Algorithms</i>	✓✓✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓

Table 1.1 — Property assessment of different intelligent techniques.