HYBRID NEURAL NETWORK AND EXPERT SYSTEMS

Larry R. Medsker

Kluwer Academic Publishers

HYBRID NEURAL NETWORK AND EXPERT SYSTEMS

by

Larry R. Medsker

Department of Computer Science and Information Systems
The American University

Kluwer Academic Publishers

Boston/Dordrecht/London

Distributors for North America:
Kluwer Academic Publishers
101 Philip Drive
Assinippi Park
Norwell, Massachusetts 02061 USA

Distributors for all other countries:
Kluwer Academic Publishers Group
Distribution Centre
Post Office Box 322
3300 AH Dordrecht, THE NETHERLANDS

Library of Congress Cataloging-in-Publication Data

Medsker, Larry.
Hybrid neural network and expert systems / Larry R. Medsker.
p. cm.
Includes bibliographical references and index.
ISBN 0-7923-9423-2 (alk. paper)
1. Neural networks (Computer science) 2. Expert systems (Computer science) I. Title.
QA76.87.M43 1994
006.3--dc20 93-38572
CIP

Printed on acid-free paper.

Printed in the United States of America

HYBRID NEURAL NETWORK AND EXPERT SYSTEMS

Dedicated
to

Masud and Dalila

Harold

David

Alan, Fran, Mike, Silvana, Xiaojing

Sean, Monty, Geoff, Mark, Kirk, Kevin, Nabin, Mary and Bob

Preface

Three years ago, when I started presenting tutorials on the integration of neural networks and expert systems, I could uncover only enough work in this area to fill one page of references. Today we see rapidly growing interest and an order of magnitude more projects on hybrid systems that combine neural networks and expert systems. Working systems have been developed for demonstrating feasibility and some are actually in use in practical situations.

Several developments have stimulated these activities. Attention has grown in the research community, including a recognition among some AI researchers that combined approaches will be necessary if the remaining tough problems in AI are to be solved. The rapid developments in R&D for neural networks have produced many applications and development tools. Both the expert system and the neural network technologies are at stages in which useful hybrid systems are conveniently possible. The current opportunity is to develop more efficient and effective design and development techniques to enable widespread production of useful and practical systems. Work is also needed to clarify the range of appropriate problems. Progress in hybrid systems will advance more rapidly as we share ideas and experiences, allowing practical models and guidelines to emerge.

This book is a step toward summarizing the state of hybrid systems and disseminating information about working systems that illustrate the issues and opportunities in this field. While other intelligent technologies should be considered eventually, the current focus is on the two that have a peculiarly complementary nature: the logical, reasoning aspect of expert systems and the visual, pattern-oriented nature of neural networks. Together they represent the range of human intelligence that is difficult or impossible to simulate with either technology alone.

The first part of the book summarizes the concepts and principles of neural networks and expert systems that are relevant to the integration of the two. The status of research in hybrid systems is summarized, and initial models for integration are presented.

Next, five case studies are presented in detail so that the nature of the problems and the design and implementation processes are clear. The applications cover product design using the analysis of technical data, control systems using the analysis of detector and sensor data, and the monitor and control of chemical processes. An important thread is the need to process data, often in large quantities, for efficient use in decision making.

The final part of the book covers guidelines and methodologies for hybrid system development and surveys the current choices for development tools and environments. While systems are currently being developed and used, the next few years should see dramatic improvements in these areas, and the final chapter presents ideas about what we can expect.

This book contains the results of my literature survey and analysis of hybrid systems research and applications development. Some of the material also derives from my experience developing hybrid systems, and for that I gratefully acknowledge the contributions of my students and colleagues. In particular, the collaborations with Masud Cader, Dalila Benachenhou, and Harold Szu were essential. Students in my NSF-sponsored Research Experiences for Undergraduates (REU) programs and American University graduate students Fran Labate, Silvana Rubino, and Michael Bramante also made important contributions. Thanks go to Ron Sun at the University of Alabama for information on the AAAI-92 Workshop.

Special thanks go to the case study contributors for their pioneering work in hybrid systems applications and for their work in sharing the results in this book. I am particulary grateful to Ray Foss at DuPont; Alper Caglayan, Jim Mazzu, and Paul Gonsalves at Charles River Analytics, Inc.; Laveen Kanal and Srinivasan Raghavan at LNK Corp.; and Jim Hendler and Anne Wilson at the University of Maryland.

Finally, I appreciate the assistance of Alper Caglayan at Charles River Analytics, Inc. and Steve Ward and Marge Sherald at Ward Systems, Inc. in the use of NueX and NeuroShell, respectively. I also appreciate the use of some figures and text from my chapters in *Hybrid Architectures for Intelligent Systems* (CRC Press), *Expert Systems and Applied Artificial Intelligence* (Macmillan Publishing Co.), and *Design and Development of Expert Systems and Neural Networks* (Macmillan Publishing Co.).

Table of Contents

Part I Fundamentals of Hybrid Systems

Chapter 1

Overview of Neural and Symbolic Systems

Expert system and neural network technologies have developed to the point that the advantages of each can be combined into more powerful systems. In some cases, neural computing systems are replacing expert systems and other AI solutions. In other applications, neural networks provide features not possible with conventional AI systems, and they could provide aspects of intelligent behavior that have thus far eluded the AI symbolic/logical approach.

Because of recent advances in neural network technology, hybrid intelligent systems can now address new problems. As these systems grow in number and importance, developers need frameworks for understanding the combination of neural networks and expert systems and will need models and guidelines for effective implementation. This chapter provides the fundamental concepts about neural networks and expert systems and gives examples to illustrate the potential of each of the technologies. On this basis, subsequent chapters address the combination of neural network and expert systems to produce hybrid systems.

1.1 Expert System Strengths and Limitations

Expert systems perform reasoning using previously-established rules for a well-defined and narrow domain. They combine knowledge bases of rules and domain-specific facts with information from clients or users about specific instances of problems in the knowledge domains of the expert systems.

An important advantage of expert systems is the ease with which knowledge bases can be modified. This is a result of the architecture (see Figure 1.1), which separates the knowledge base from the inference engine. As a result, changing the knowledge base does not require programming but can be done via word processing or an editor. This feature makes knowledge engineering accessible to a wider variety of analysts, end users, and experts.

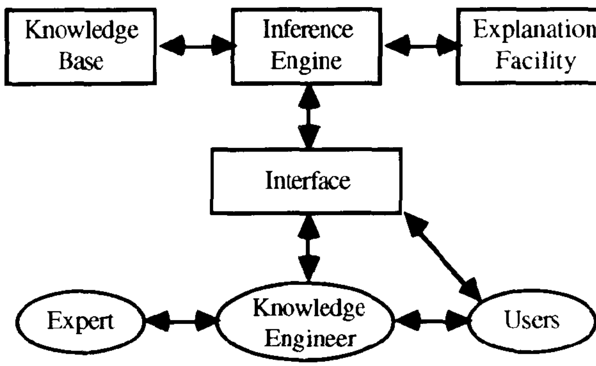

Figure 1.1. The architecture of expert systems

Ideally, reasoning can be explained, and most expert system development tools allow the creation of explanation systems to help the user understand questions being asked or conclusions. Thus, the system can function more like human experts who explain the reasoning processes behind their recommendations.

Expert systems are especially good for closed-system applications for which inputs are literal and precise, leading to logical outputs. For stable applications with well-defined rules, expert systems can be easily developed to provide good performance. Such systems can take advantage of the wealth of techniques developed in expert systems research to perform different types of reasoning using very efficient, systematic techniques. Thus, people are able to inspect and understand these systems because they have familiar structures similar to the logical frameworks humans use.

Another advantage of expert systems is the number and variety of commercial development systems that have become available over the last several years. These tools and associated techniques allow exploratory studies rapid prototyping for use in knowledge engineering. Prototypes are especially useful in gaining the interest and attention of experts so that the knowledge acquisition process is more productive and amicable.

A fundamental limitation of the expert system approach arises from the fact that experts do not always think in terms of rules. Thus, an expert system does not in these cases mimic the actual reasoning process of human experts. The computerized system is an attempt to produce performance that resembles human reasoning in some limited domain. The mechanism, however, may or may not resemble the actual biological or cognitive process. Thus, extensions of any expert system technique may not carry over into behavior that is similar to that of a human. Cognitive science, AI, and expert system research is still needed to produce a fundamental approach that models actual human reasoning.

In the meantime, a variety of approaches have been developed to address deficiencies in current expert system techniques, and many of these approaches show improvement in those limited areas. The major techniques of current interest are surveyed in the next section of this chapter. A comprehensive technique based on a good understanding of human reasoning is still needed.

A specific problem with expert systems is what is referred to as the knowledge acquisition bottleneck. While development tools have become very sophisticated and effective, expert systems still require extensive effort for eliciting knowledge from humans as well as from written material. Knowledge acquisition is still primarily a human-intensive activity requiring the usual system analysis abilities plus additional interviewing and interpersonal skills that are tailored to interacting with human experts. Furthermore, human experts may be too busy or otherwise difficult to deal with, so that a whole project may be threatened or delayed. Research and development efforts are focussing on understanding and automating the knowledge acquisition process to the extent possible. However, this activity remains a serious problem for the rapid and pervasive development of large expert systems.

Another difficulty with expert systems is in the area of large systems development. For real-world applications, the development process becomes difficult to manage. Working with experts and dealing with the complexity of large systems leads to prolonged, expensive development and delivery times. Furthermore, validation and verification of systems becomes difficult, if not impossible, as many lines of reasoning must be checked.

Other limitations involve fundamental uncertainties about the expert system approach. For example, more work is needed on how to represent commonplace knowledge, which humans deal with so well and so often. Also, expert systems do not automatically benefit from experience with their use and thus do not learn for failures or their use with novel examples.

While research and development are proceeding toward the improvement of expert system technology, some researchers question the underlying philosophy, at least in terms of the deficiencies described above. The alternate technologies discussed next address many of these concerns and may, along with solving certain practical problems, give insight into fundamental research questions that need to be answered before the goals for expert systems are fully realized.

1.2 Neural Computing

The state-of-the-art in neural computing is inspired by our current understanding of biological neural networks; however, after all the research in biology and psychology, important questions remain about how the brain and the mind work. Advances in computer technology allow the construction of interesting and useful artificial neural networks that borrow some features from the biological systems. Information processing with neural computers consists of analyzing patterns of activity, with learned information stored as weights between neurode connections (see Figure 1.2). A common characteristic is the ability of the system to classify streams of input data without the explicit knowledge of rules and to use arbitrary patterns of weights to represent the memory of categories. Together, the network of neurons can store information that can be recalled in order to interpret and classify future inputs to the network. Because knowledge is represented as numeric weights, the rules and reasoning process in neural networks are not readily explainable.

Neural networks have the potential to provide some of the human characteristics of problem solving that are difficult to simulate using the logical, analytical techniques of expert system and standard software technologies. For example, neural networks can analyze large quantities of data to establish patterns and characteristics in situations where rules are not known and can in many cases make sense of incomplete or noisy data. These capabilities have thus far proven too difficult for the traditional symbolic/logic approach.

Neural networks rely on training data to "program" the systems. Thus, neural network components can be useful for hybrid systems by using an appropriate training set that allows the system to learn and generalize for operation on future input data. Inputs exactly like training data are recognized and identified, while new data (or incomplete and noisy versions of the training data) can be put into the closest matches to patterns learned by the system.

Neural network components can be useful when rules are not known, either