



**This Is a Publication of
The American Association for Artificial Intelligence**

This electronic document has been retrieved from the
American Association for Artificial Intelligence
445 Burgess Drive
Menlo Park, California 94025
(415) 328-3123
(415) 321-4457
info@aaai.org
<http://www.aaai.org>

*(For membership information,
consult our web page)*

The material herein is copyrighted material. It may not be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from AAAI.

Integration of Knowledge and Neural Heuristics

LiMin Fu

- This article discusses the First International Symposium on Integrating Knowledge and Neural Heuristics, held on 9 to 10 May 1994 in Pensacola, Florida. The highlights of the event are summarized, organized according to the five areas of concentration at the conference: (1) integration methodologies; (2) language, psychology, and cognitive science; (3) fuzzy logic; (4) learning; and (5) applications.

The last few years have seen a rapid growth of interest in combining knowledge-based techniques and computational neural networks as a new paradigm for producing AI. This trend has begun to pick up its momentum since the late 1980s, and both approaches have enjoyed many successful applications to real-world problems. This hybrid idea is largely a consequence of an increasingly strong belief that knowledge and neural models can complement each other beneficially.

The growing community in this area convened at the First International Symposium on Integrating Knowledge and Neural Heuristics (ISIKNH) on 9 to 10 May 1994 in Pensacola Beach, Florida, for the first time on the international level. We received applications from more than 60 research groups from more than 10 countries, which reflected a rather complete time-sliced picture of ongoing research efforts in this area. The high enthusiasm of the participants in the relaxed environment under the beach breeze fostered productive brainstorming during the two-day gathering. This report summarizes the highlights of the event.

Research Topics

The conference presentations were organized into five areas: (1) integration methodologies; (2) language, psychology, and cognitive science; (3) fuzzy logic; (4) learning; and (5) applications.

Integration Methodologies

In his keynote speech, B. Chandrasekaran (The Ohio State University) argued that the debate about the right approach to AI could be clarified by removing many confusing notions with regard to what made something a representation. In one case, content is more important than form, whereas in the other, the reverse is true. He pointed out that proper understanding of Alan Newell's knowledge level versus symbol level distinction could illuminate many phenomena related to representation. As to the issue of integration, the first question is always, "Integrate what and what?" Many integration alternatives could exist, and not all of them make sense in a given context. After all, one cannot be so naive as to overlook the potential that a hybrid inherits the weaknesses, rather than the strengths, of its parents.

The integration or synergism of knowledge-based components and neural networks in a system can be explored from their functional and structural relationships in the system. Five integration architectures can be identified:

First is *completely overlapped*: In this architecture, the system is both a knowledge-based system and a

neural network. It has a dual nature. The system optimizes its performance by combining the strengths of the two forms. Depending on the need, it can be presented to the user as a traditional expert system or a neural network.

Second is *partially overlapped*: The system is a hybrid of a knowledge-based system and a neural network, exhibiting features of both. The two components share some but not all of their own internal variables or data structures. They often communicate through computer internal memory rather than external data files. For example, a neural network augmented with explanation capability is a partially overlapped system.

Third is *parallel*: A knowledge-based system and a neural network work in parallel to solve a common problem. Both can be stand-alone systems. The two components do not share their own internal variables or data structures. They communicate through their input-output devices, such as data files.

Fourth is *sequential*: A knowledge-based system and a neural network operate in sequence to solve a particular problem. Again, both can be stand-alone systems, and they do not share internal variables. The output of one component is passed on to the other for further processing.

Fifth is *embedded*: In this integration, either a knowledge-based component is embedded within a neural network, or vice versa. By saying that *X* is embedded in *Y*, we mean that *X* (a guest) becomes an element of *Y* (a host). Internal information exchange is expected. However, this architecture differs from the partially overlapped architecture in that the system's external features are determined by the host component only. It is arguable that many neural networks already use knowledge in specifying their input-output and structures. Besides, it is worthwhile to embed a neural network within an expert system.

In another perspective, we can categorize integration paradigms according to the nature of coupling: *Fully coupled* corresponds to the completely overlapped architecture. *Tightly cou-*

consists of a premise (antecedent) and a consequent. In the network configuration, the premise is assigned a hidden unit called a conjunction unit, each condition corresponds to an assigned attribute or concept node, and the consequent corresponds to an assigned concept node. Each condition node is connected to the conjunction unit (the premise), which is, in turn, connected to the consequent node. Under such construction, the rule strength corresponds to the weight associated with the connection from the *conjunction unit* to the consequent node. In addition to knowledge-based connections, we might add some hypothetical connections to increase the learning capacity of the network. The neural network so built is referred to as a *rule-based connectionist network*.

Christian Omlin and Lee Giles (both of NEC Research Institute) described a technique that integrated temporal symbolic knowledge into recurrent neural networks. Ron Sun (University of Alabama) presented a technique for implementing schemes and logics in connectionist models. I. Hatzilygeroudis (University of Patras, Greece) showed how to integrate symbolic rules with neurocomputing. Ricardo Machado (IBM Rio Scientific Center, Brazil) and Armando da Rocha (University of Patras, Greece) described a hybrid system with incremental learning capability. David Opitz and Jude Shavlik (both of University of Wisconsin at Madison) reported a genetic algorithm for refining knowledge-based neural networks. Jeffrey Mahoney and Raymond Mooney (both of University of Texas at Austin) described a hybrid system that could adapt weights, learn new rules, and modify network architecture. K. D. Nguyen and R. C. Lacher (both of Florida State University) presented a system that could learn the semantic interpretation of the nodes in an expert network by back propagation.

In addition, T. S. Dillon, S. Sestito, M. Witten, and M. Suing (all of La Trobe University, Australia) described a method for extracting rules from an unsupervised learning network by weight thresholding. Selwyn Pira-

muthu and Michael Shaw (both of University of Florida) reported that decision trees could be used as the feature selector for a feed-forward neural network.

Applications

The number of hybrid intelligence applications has snowballed in recent years. The momentum grows as more successful applications are developed. Whether or not a new theory is created seems to be of less concern in this discipline.

Financial-market prediction is a challenging application area where neural networks appear to be promising. However, one can be cynical about neural network performance. Consider a neural network for predicting the rise or fall of a stock price. Suppose the neural network is supplied with the past-month price data, which showed a continuously rising trend. After training on the data, it is not a surprise that the network will predict a price rise in the future—abysmal ignorance about the market world.

Lawrence Bookman (Sun Microsystems Lab Inc.) described a technique for automatically constructing a connectionist network knowledge base for unstructured domains. This technique made use of online corpora and applied information theory to extract statistical relationships between words in the text.

Kazuhiro Kohara and Tsutomu Ishikawa (both of NTT Network Information Systems Laboratories, Japan) presented a system that used prior knowledge and neural heuristics to predict stock prices. The input parameters of the system included market indexes such as dollar-to-yen exchange rate, interest rate, crude oil price, and New York stock prices and event knowledge extracted from daily headlines of newspapers. His results showed that neural networks outperformed multiple regression analysis, recurrent networks outperformed feed-forward networks, networks with prior knowledge did better than ones without such knowledge, and networks input with longer history data did better than ones with shorter history data. The best result was

about 70-percent correct with respect to the direction of change in prices.

In engineering applications, Vivek Goel and Jianhua Chen (both of Louisiana State University) reported an expert network, based on my knowledge-based conceptual neural networks, for material selection in machining domains. The network could automatically update its knowledge and reason about whether a connection was semantically correct.

Other application examples include the following: Sylvian Ray (University of Illinois at Urbana-Champaign) described a technique for multichannel signal analysis; Thomas English (Texas Tech University) showed how to define learning tasks by assessing situational awareness in air battlefields; Steven Walczak (University of Tampa) presented a hybrid system for resource allocation in academic admissions; Walter Johnson, Khaled Kafel, and John Forde (Suffolk University) used a constraint-based feed-forward neural network for game playing; Hsu-Huang Hsu, LiMin Fu, and Jose Principe (all of University of Florida) reported an expert network for sleep staging that required a significantly lower amount of data for training; and Jurgen Rahmel and A. von Wangenheim (both of University of Kaiserslautern, Germany) reported a case-based diagnostic system coupled with Kohonen networks.

Panel Discussion

The panel, made up of Chris Lacher (chair), James Anderson, Ronald Yager, Steve Gallant, Ron Sun, and Lawrence Bookman, discussed the future directions of AI. Anderson pointed out that the merger of simple neural ideas and classical AI was interesting and would be a future direction. He added that the neural network approach should concentrate more on flexibility and programmability, and the classical AI approach should address more of the representational and dynamic aspects. Yager assumed a similar attitude. He expressed that AI provided a nice paradigm that the neural network approach could accommodate well, as illustrated by the integrated