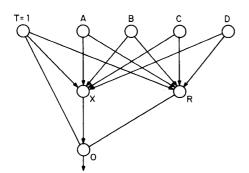Fig. 4. A network for the "rule-plus-exception" problem.

were estimated during course of learning were:

| | First Layer Sensitivities | | | | |
	T	A	B	C	D
X	0.15	0.09	0.05	0.02	0.02
R	0.16	0.42	0.27	0.06	0.03

| Second Layer Sensitivities | | |
T	X	R
0 0.39	0.22	0.72

If one hidden unit has to be removed, we have to compare S_{OX} to S_{OR}, to realize that unit X, the "exception" is the one to be eliminated. Further, we look at the sensitivities of the connection from the inputs to the remaining unit R. Among them S_{RC} and S_{RD} are the smallest, so w_{RC} and w_{RD} can be pruned. This leaves a small size network that can compute the "rule" (AB) only, and it is the best two input approximation to the problem at hand.

SUMMARY AND COMMENTS

We have devised a simple procedure that takes advantage of pieces of data (namely, the gradient and increments to the weights) that are available during the normal course of neural net training. Learning remains intact, and the sensitivity of the (unchanged) error function to the elimination of each synapse is concurrently evaluated by a "shadow process" that demands only a negligible computational overhead. A decision on the connections that should be pruned is made only after the completion of the training phase.

The advantages of this method over previous approaches have been described and also demonstrated by examples. However, one should not overlook the fact that this procedure, like previous approaches, is only a heuristic that provides some reasonable estimates of the true sensitivities (as defined in (7) above). As we explained in the derivation of (12), we are restricted to the learning path (the dashed line in Fig. 2), otherwise we have to specially train and retrain for each synapse that is a candidate for elimination. This would result in a prohibitively long training which our simple procedure avoids.

REFERENCES

[1] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, "Learning internal representations by error propagation," in *Parallel Distributed Processing*, vol. 1, D. E. Rumelhart and J. L. McCleland, Eds. Cambridge, MA: M.I.T. Press, 1986, pp. 318–362.

[2] D. G. Luenberger, *Linear and Nonlinear Programming*. Reading, MA: Addison-Wesley, 1984.

[3] J. Sietsma and R. J. F. Dow, "Neural net pruning—Why and how?," in *Proc. IEEE Int. Conf. Neural Networks*, vol. 1 (San Diego, CA), 1988, pp. 325–332.

[4] S. J. Hanson and L. Y. Pratt, "Comparing biases for minimal network construction with back-propagation," in *Advances in Neural Information Processing I*, D. S. Touretzky, Ed. Morgan Kaufmann, 1989, pp. 177–185.

[5] M. C. Mozer and P. Smolensky, "Skeletonization: A technique for trimming the fat from a network via relevance assessment," in *Advances in Neural Information Processing I*, D. S. Touretzky, Ed. Morgan Kaufmann, 1989, pp. 107–115.

Neural Networks for Control Systems

PANOS J. ANTSAKLIS, SENIOR MEMBER, IEEE

Abstract—This letter describes 11 papers from the April 1990 Special Issue of the IEEE CONTROL SYSTEMS MAGAZINE on Neural Networks in Control Systems.

INTRODUCTION

Ever-increasing technological demands of our modern society require innovative approaches to highly demanding control problems. Artificial neural networks with their massive parallelism and their learning capabilities offer the promise of better solutions, at least to some problems. By now, the control community has heard of neural networks and wonders if these neural networks can be used to provide better control solutions to old problems or perhaps solutions to control problems which have withstood its best efforts.

CONTROL TECHNOLOGY DEMANDS

The use of neural networks in control systems can be seen as a natural step in the evolution of control methodology to meet new challenges. Looking back, the evolution in the control area has been fueled by three major needs: the need to deal with increasingly complex systems, the need to accomplish increasingly demanding design requirements, and the need to attain these requirements with less precise advanced knowledge of the plant and its environment; that is, the need to control under increased uncertainty. Today the need to better control increasingly complex dynamical systems under significant uncertainty has led to a reevaluation of the conventional control methods, and it has made apparent the need for new methods. It has also led to a more general concept of control, one which includes higher level decision making, planning, and learning, which are capabilities necessary when higher degrees of system autonomy are desirable. These ideas are elaborated upon in [1]. In view of this, it is not surprising that the control community is seriously and actively searching for ideas to deal effectively with the increasingly challenging control problems of our modern society. Need is the mother of invention and this has been true in control since the times of Ktesibios and his water clock with its feedback mechanism in the third century B.C.[2], the earliest feedback device on record. So the use of neural networks in control is rather a natural step in its evolution. Neural networks appear to offer new, promising directions toward better understanding and perhaps even solving some of the most difficult control problems. History, of course, has made clear that neural

networks will be accepted and used if they solve problems which have been previously impossible or very difficult to solve. They will be rejected and will be just a novelty fading fast, if they do not prove useful. The challenge is to find the best way to utilize fully this powerful new tool in control; the jury is still out as their best uses have not yet been decided. It is hoped that this special issue will raise interest in neural networks in the control community and elsewhere, and will provide challenges and food for thought.

Special Issue

The special issue contains 11 papers.[1,2] Early versions of most of these papers were presented in conferences on control, robotics, or neural networks in 1989. In selecting these papers, the emphasis was placed on presenting as varied and current a picture as possible of the use of neural networks in control. Additional papers were commissioned specifically for this special issue to make the exposition more complete and self-contained. Applications were emphasized, but rigor was also praised. Complete proofs, however, of the results were not included. Nevertheless, the authors take full responsibility for their claims! Please remember that this is a window with a view towards control applications of neural networks. It was opened originally to include papers from the 1989 American Control Conference, and then it was widened to give a more comprehensive picture. However, it is still a window. This is not a survey issue. It is a special issue designed to raise interest, to be thought provoking, and to generate new ideas.

There are several topics covered by the papers in this special issue. The first paper, by *A. N. Michel and J. A. Farrell*, introduces mathematical models of neural networks and discusses algorithms to assign the weights in associative memories. The next paper, by *D. Nguyen and B. Widrow*, introduces applications by using neural networks to model and control a highly nonlinear system, a trailer truck backing up to a loading dock. Modeling of chemical processes is addressed in the third paper by *N.V. Bhat, P. Minderman, T. McAvoy, and N. Wang*. Such processes are typically very complex and neural networks do offer a very attractive alternative, as these models are perhaps better learned than fully detailed. System identification in the time and frequency domains is the topic of the next paper by *R. Chu, R. Shoureshi, and M. Tenorio*. In order to effectively use neural networks in control problems, the neural controllers must be compared to conventional ones; this is the direction taken in the fifth paper by *L. G. Kraft and D. P. Campagna*, where a neural controller and certain conventional adaptive controllers are applied to the same simple system and the results are compared. The sixth paper, by *F.-C. Chen*, discusses a method to introduce neural networks to enhance self-tuning controllers to deal with large classes of nonlinear systems; the backprogagation learning algorithm is used. In the seventh paper, by *S. R. Naidu, E. Zafiriou, and T. J. Mc Avoy*, neural networks and backpropagation are used for sensor failure detection in chemical process control systems. Additional information about learning algorithms in neural networks is given in the next paper by *S. C. Huang and Y. F. Huang*; backpropagation is discussed and certain extensions are introduced. The next two papers are experimental applications of neural networks to control complex systems in real time. The pitch attitude of an underwater telerobot is regulated in the ninth paper by *R. M. Sanner and D. L. Akin*, and the experimental results are presented. Mobile robots with many sensors learn to interact in the next paper by *S. Nagata, M. Sekiguchi, and K. Asakawa*; the robots demonstrate their abilities by playing a form of the cops-and-robbers game. The interaction of rule-based systems and neural networks is studied by *D. A. Handelman, S. H. Lane, and J. J. Gelfand* in the last paper and a con-

troller integrating the two is developed; it is used to teach a two-link robot manipulator a tennis-like swing. A more detailed description of the papers follows.

Description of Papers

The mathematical framework necessary for in-depth studies of several system and control applications of neural networks is set in the first paper by *A. N. Michel and J. A. Farrell* titled *"Associative Memories via Artificial Neural Networks,"* where mathematical models are introduced and methods are described to design associative memories using feedback neural networks. Neural networks with full feedback interconnections are of interest here. Their dynamical behavior, studied via differential equations, exhibits stable states which act as basins of attraction for neighboring states as they develop in time. This time evolution towards these equilibrium points can be seen as the attraction of an imperfect pattern towards the correct one, stored as a stable equilibrium. Several design methods are presented to appropriately assign the weights, so that the resulting network will behave as an associative memory. A neural network so designed can be useful in control as, for example, an advanced look-up dictionary of different control algorithms; when certain operating conditions are present, they are matched to stored conditions and the control action which corresponds to conditions that most closely match the current operating conditions are selected. Other applications of associative memories to control are, of course, possible.

A method to use neural networks to control highly nonlinear systems is presented by *D. Nguyen and B. Widrow* in their paper titled *"Neural Networks for Self-Learning Control Systems."* Feedthrough, multilayered neural networks are used; and learning, via the backpropagation algorithm, is implemented to determine the neural network weights to first model the plant and then design the controller. First, a neural network emulator learns to identify the dynamic characteristics of the system. The controller, another multilayered network, then learns to control the emulator. The self-trained controller is then used to control the actual dynamic system. The learning continues as the emulator and controller improve as they track the physical process. The power of this approach is demonstrated by using the method to steer a trailer truck while backing up to a loading dock.

The main emphasis in the next two papers is on system modeling. The modeling of nonlinear chemical systems using neural networks and learning is addressed by *N. V. Bhat, P. Minderman, T. McAvoy, and N. Wang* in *"Modeling Chemical Process Systems via Neural Computation."* Backpropagation is used for the system to learn the nonlinear neural network model from plant input–output data, and for interpreting biosensor data. Typical chemical processes to be controlled are rather complex and frequently the relationships are perhaps better learned than fully detailed out. Two reactor examples are considered, a steady-state reactor and a dynamic pH stirred tank system; the interpretation of sensor data is illustrated using a fluorescence spectra example.

Two methods for identification of dynamical systems are described in the paper *"Neural Networks for System Identification"* by *R. Chu, R. Shoureshi, and M. Tenorio*. First a technique for assigning weights in a Hopfield network is developed to perform system identification in the time domain; it involves the minimization of least-mean-square of error rates of estimates of state variables. System identification in the frequency domain is also illustrated, and it is shown that transfer functions of dynamical plants can be identified via neural networks.

Conventional adaptive controllers and neural network based controllers are compared in the paper by *L. G. Kraft and D. P. Campagna* titled *"A Comparison Between CMAC Neural Network Control and Two Traditional Adaptive Control Systems."* If neural network controllers are to be used in the control of dynamic systems, they must be evaluated against controllers designed using conventional control theory. A self-tuning regulator and a model reference adaptive controller are compared to a neural cerebellar

model articulation controller. They are all used to control the same simple system and the results are tabulated and discussed at length.

A method to provide adaptive control for nonlinear systems is introduced in *"Back-Propagation Neural Networks for Nonlinear Self-Tuning Adaptive Control"* by *F.-C. Chen*. The author uses a neural network and the backpropagation algorithm to alter and enhance a self-tuning controller so that it can deal with unknown, feedback linearizable, nonlinear systems. Simulations of a nonlinear plant controlled by such a neural controller are included to illustrate the method.

Neural networks and backpropagation are proposed by *S. R. Naidu, E. Zafiriou, and T. J. McAvoy* for sensor failure detection in *"The Use of Neural Networks for Sensor Failure Detection in a Control System."* The ability to reliably detect failures is, of course, essential if a certain degree of autonomy is to be attained. Process control systems are of main interest here. Backpropagation is used for sensor failure detection and the algorithm is compared via simulations to other fault detection algorithms.

Most of the neural network applications seem to incorporate some form of learning. Learning is discussed by *S. C. Huang and Y. F. Huang* in *"Learning Algorithms for Perceptrons Using Back-Propagation with Selective Updates."* The ability to learn is one of the main advantages of neural networks. Learning algorithms are discussed in general with main emphasis on supervised algorithms. The backpropagation algorithm, used in feedforward types of networks, is discussed at length and an extension is presented. These learning algorithms are applied for illustration to a perceptron associative memory.

R. M. Sanner and D. L. Akin in *"Neuromorphic Pitch Attitude Regulation of an Underwater Telerobot,"* present the experimental results of using trained neural networks to regulate the pitch attitude of an underwater telerobot. These experimental results are a follow-up of their previous work involving computer simulations only. The neural network performed as predicted in simulations, however, it was observed that unacceptable delays can be introduced if a single serial microprocessor is used to calculate the control action. Hardware implementations of neural networks are seen as necessary.

The control of mobile robots is the topic addressed by *S. Nagata, M. Sekiguchi, and K. Asakawa* in *"Mobile Robot Control by a Structured Hierarchical Neural Network."* Neural networks are used to process data from many sensors for the real-time control of mobile robots and to provide the necessary learning and adaptation capabilities for responding to the environmental changes in real time. For this, a structured hierarchical neural network and its learning algorithm are used, and the network is divided into two parts connected with each other via short memory units. This approach is applied to several robots which learn to interact and participate in a form of the cops-and-robbers game.

D. A. Handelman, S. H. Lane, and J. J. Gelfand in *"Integrating Neural Networks and Knowledge-Based Systems for Intelligent Robotic Control,"* address the issues involved when integrating these quite distinct systems, which offer very different capabilities. To demonstrate the integration technique and the interaction of the two systems, a two-link robot manipulator is taught how to make a tennis-like swing. The rule-based system first determines how to make a successful swing using rules alone. It then teaches a neural network to perform the task. The rule-based system continues to evaluate the neural network performance and if changes in the operating conditions make it necessary, it retrains the neural network.

If there is a message stressed in this special issue, it is this: Neural networks in control must be studied using mathematical rigor in the tradition of the discipline. Only in this way can the control community harvest the full benefits of these powerful new tools. Only in this way can something lasting and useful for the years to come be created.

REFERENCES

[1] P. J. Antsaklis, K. M. Passino, and S. J. Wang, "Towards intelligent autonomous control systems: Architecture and fundamental issues," *J. Intelligent Robotic Syst.* vol. 1, pp. 315–342, 1989; also, a shorter version appeared in *Proc. Amer. Contr. Conf.* (Atlanta, GA), June 15–17, 1988, pp. 602–607.

[2] O. Mayr, *The Origins of Feedback Control.* Cambridge, MA: M.I.T. Press, 1970.

Standardization of Neural Network Terminology

RUSSELL C. EBERHART

Abstract—It is desirable to move toward commonly accepted terminology in the neural network field. This letter outlines the initial activities of an Ad Hoc Standards Committee established by the IEEE Neural Networks Council to pursue this effort.

INTRODUCTION

A great diversity currently exists in the terminology and notation used in neural networks literature. This can create misunderstanding and confusion for the reader, even a reader relatively experienced in the field. If the reader is new to the field, the communications problems are potentially even more severe.

It is particularly difficult for a person new to neural networks to understand references in articles that refer to nodes, neurodes, processing elements, units, processing units, etc., all of which refer to essentially or exactly the same thing. There is also an unfortunate inconsistency in the diagrammatic representation of networks, and in neural network notation. For example, is the connection weight from node i to node j designated as w_{ij}, or as w_{ji}?

AD HOC STANDARDS COMMITTEE

To address the terminology/notation problem, the IEEE Neural Networks Council has established an Ad Hoc Standards Committee. While in more mature technologies, standards committees typically address issues such as standardization of measurements and procedures, it is felt that neural networks technology is still in such an actively evolving state that an attempt to standardize terminology and notation must take precedence. The Ad Hoc Standards Committee will meet at each of the two International Joint Conferences on Neural Networks (IJCNN's) held each year.

GLOSSARY

The first efforts at terminology and notation standardization will focus on terms and symbols that are frequently used in the neural networks literature. Terminology has originated in a variety of domains such as engineering, mathematics, biology, physics, etc. Terms often seem to be used without a solid understanding of their exact meaning.

A proposed list of terms to be considered by the Ad Hoc Standards Committee is presented below. The list is not meant to be exhaustive, but rather to represent basic, frequently used terms:

Activation function, activation rule, activation state, activation value, adaptive resonance, architecture, associative