Soft Computing for Soft Technologies: Artificial Neural Networks and Fuzzy Set Theory for Human Services

Jan Steyaert

ABSTRACT. This article outlines the basic concepts of neural networks, gives information on the presently available hardware and software and discusses the possible implications of this technology for the human services and service provision.

KEYWORDS. Neural networks, fuzzy set theory, human service applications

INTRODUCTION

A few months ago, several electronic mail messages were exchanged, both on the CUSSnet and the SOCWORK network. All were labeled "new science" and dealt with the question of whether neural networks and fuzzy set theory, so-called soft computing, had anything to offer the human services. Many messages requested information on the theory and key concepts behind these new technologies. This article outlines the basics of neural network and fuzzy set theory and describes some of the concepts behind them.

Jan Steyaert is Research Assistant, Belgian National Fund for Scientific Research. He is also Secretary of ENITH (European Network for Information Technology in Human Services) and member of the HUSITA-board. Mr. Steyaert is studying the use of client information systems in service provision and social policy. Address correspondence to: University of Antwerp, Department of Sociology, Universiteitsplein 1, B-2610 Antwerp, Belgium. (steyaert@psws.uia.ac.be)

Computers in Human Services, Vol. 10(4) 1994

An introduction to the available software and hardware will also be given, as well as guidelines on where to look for further information and copies of shareware products. Above all, I will try to assess the applicability of neural networks and fuzzy set theory to human service problems. This article is not a report on the results of an empirical research project. I did not try to test the applicability of neural networks on real practice data or in human service agencies. This article is merely an attempt to shed light on some of the most promising new technologies.

BACKGROUND: KEY CONCEPTS AND GENERAL APPLICATIONS

Neural Networks

Neural networks have existed for many years as a special form of machine learning. However, only recently have they evoked much interest. This is partly because of theoretical breakthroughs and partly because of spectacular examples of neural network applications. Currently, neural networks are being used in different areas such as pattern recognition in analyzing visual images, predicting financial data (Bowen & Bowen, 1990), real estate appraisal, prediction of conflict outcomes (Schrodt, 1991), prediction of length of stay in intensive care unit (Tu & Guerriere, 1993) or prediction of length of stay in psychiatric care (Davis, Lowell & Davis, 1993).

A well-known application of neural network technology is risk assessment when granting someone a mortgage loan (Dayhoff, 1990, p. 220). Financial organisations want to lower the risk of people being unable to pay back loaned money. Different methods can be used to assess this risk. One method is statistical analysis of previous cases. This method is problematic because many factors influence the outcome of the risk assessment process. Also, most statistical analyses are conducted with linear regression, and reality seldom is linear. Another method to assess risk is to elicit the experience and knowledge of experts in the field. This method is also problematic because experts are expensive and often have difficulty in explaining why they do certain things in certain circumstances. A third method to analyze risk is to develop a neural network. Neural network software is trained by feeding it previous cases with as many relevant variables as possible (the input layer of the network) and the result of whether or not those people were able to pay back their loan (the output layer). The neural network then learns from these examples and builds a model that it can apply to new cases.

Another application of neural network technology can be seen in the character recognition capabilities of the personal digital assistants (PDAs) entering the marketplace. Some PDAs only work with handwritten commands on the screen (no keyboard). Neural network technology converts the handwriting to typed characters. As one works with a PDA, it learns to read one's handwriting better and better.

When talking about neural networks, most of the time we are talking about artificial neural networks (ANNs). Natural neural networks can be found in the brains of all living species. These brains consist of neurons, synapses and the like. Neurons are linked to each other by synapses, and signals are passed on from one neuron to the other via the synapses. One of the main features of brains is that knowledge is not stored in one or a few neurons, but distributed across a network consisting of a great many neurons. Information gets distributed different ways because the strengths of the synapses differ.

An artificial neural network is a computer simulation of the natural brain. Artificial neural networks consist of one or more layers of nodes (neurons) linked with connections (synapses). These connections have different strengths or weights and pass data on in different ways. Nodes are organized in layers. Each network has at least one input layer and one output layer. In addition, a neural network can have a number of hidden layers. The number of nodes on each layer differs. The input layer of a neural network is a representation of the question to the network; the output layer will contain the network's answer.

The amazing thing about artificial neural networks is that they are not programmed but actually learn to perform certain tasks. This learning takes place in the training stage when the neural network learns from actual cases. At the start of the training process, the network is initialized with random weights for all connections between the nodes. The next important task is to "feed" the network with many case examples of the task you want the network to learn. Once a neural network has been set up and trained, it can be used in the application stage, to apply its "experience" on new problem situations.

Artificial neural networks learn to perform tasks by learning associations between certain input patterns and output results from cases. The learning process basically comes down to an adjustment of the weights between the nodes in the network. Different learning techniques can be used. These learning rules are the core ingredient of any neural network. The most known is called backpropagation. Every time the network learns about a difference between what it predicted and what actually happened, the weights of the connections are adjusted. "The network learns by making corrections to the connections, based on the error at the output. Corrections signals propagate back through the network during training, hence the name" (Lawrence, 1993, p. 121). As a rule, one can say that the more cases presented to the network, the better the results. However, the cases need to be presented more than once to the network. Training is an iterative process, in which cases are fed into the network several times, until the results of the network are within a certain error margin. Another rule is that the more learning cycles, the better the network. However, this is not a linear relation, and in some cases adding more learning cycles can lessen the network effectiveness. One should therefore always analyze the learning process by testing it at regular intervals, for example, every 50 cycles. Normally, this is done by setting a percentage of the training cases aside for testing and using only part of the available data for training. If the network is trained on, for example, 90% of your data, you can use the other 10% to test it.

After the network has been trained and tested and the weight between its nodes has been adjusted to the available examples, you can use the network to perform the task it has learned. In the above example, if you present a new case, the network will assess the risk of the person not paying back the loan. Neural networks appear to be quite good at these tasks. They also have some major advantage over the development of expert systems, in the sense that you no longer need to go through the difficult task of eliciting rules from experts. By using a neural network approach, you can elicit knowledge from client files and databases. The problem then comes down to the quality of your files and databases.

Fuzzy Logic

The Third International Fuzzy Systems and Intelligent Control Conference, held in March 1994 in Kentucky, had as its theme "fuzzy logic, neural networks and soft computing." The University of California at Berkeley has an initiative known as Soft Computing, headed by L. Zadeh, the inventor of fuzzy logic. Zadeh described "the distinguishing characteristic of soft computing is that its primary aim is to exploit the tolerance for imprecision and uncertainty to achieve tractability, robustness and low cost. . . . The major components of soft computing are fuzzy logic, neural network theory and probabilistic reasoning techniques" (Zadeh, 1992).

Fuzzy logic is a special type of logic that introduced the concept of something not only being completely true or false, but true or false in degrees. Fuzzy logic is often an important aspect of neural networks, as fuzzy logic is used in constructing the algorithms. Fuzzy sets is a concept that was introduced in mathematics several years ago, introducing the notion that language does not work with precise, clear-cut definitions. The words we use in everyday language or in our daily practice have vague definitions.

Most of us understand, if I talk about a client and say that I gave him some advice. However, someone else would have said in exactly the same situation that he/she gave the client some psycho-social help. The differences are vague. Where does giving advice stop and giving psycho-social help start? The example may sound trivial, but has strong implications if, for example, different funding sources need to be used for giving advice and psycho-social help, as may become the case in The Netherlands. Clear definitions of the concepts then have important financial consequences. In order to define the concepts, one can sit down and discuss these matters and come up with elaborate definitions explaining the precise differences between activities of treatment, such as giving information, giving advice, and giving psycho-social help. But in the end, no matter how many distinctions we make and how many words we use to distinguish between the possible activities of treatments in social work, there will still be vague boundaries between words.

The concept of having to work with some degree of vagueness in our language, rather than having the ability to make our meanings exactly correct, differs from the aims of classification systems such as DSM-III-R or ICD. The idea of these classifications is precisely to define terms beyond doubt. The concept of fuzzy sets also differs from the binary nature of conventional computers, which works with yes-no classifications. Fuzzy set theory prefers to work with "maybes." Fuzzy set theory therefore defines concepts or terms as continuums. Can a certain situation be classified as child neglect? In fuzzy set theory (and many clinical practice situations), this is not a yes-no question, but a continuum.

SOFTWARE, HARDWARE, AND INFORMATION RESOURCES

This section provides a limited overview of the available neural network and fuzzy set products, both commercial and shareware. A selection of the shareware products is available from the CUSSN disk copy service.[1]

Software

Some shareware products allow one to learn more about neural networks and the way they work. A very useful program in this sense is the Visible Neural Network. It is a free version of a commercial product that comes without the manual and source code. It offers a very clear demonstration of how neural networks behave. The program uses a neural network to teach the computer to recognize colours. The input of the network is a combination of three base colours (red, yellow, blue). During the training phase, different combinations of these colours together with the desired output are presented to the network. They are visually presented on the screen along with the output patterns of the network. As more and more cases are processed with each training cycle, the degree of error in recognizing the colours diminishes. Through the visual presentation of the output pattern, the notion of fuzzy set is easy to grasp. The network does not recognize a certain colour as red, or green or whatever, but as being a lot of red and only a little of blue and yellow, etc. The visible neural network also offers the possibility of numerical and statistical analyses of the network. You can do this at various stages during training, and see how the network actually learns to recognize colours and gains experience.

Another useful product is AutoNet. The demonstration version is freely available. The restriction of the demo as compared to the production version is that data files cannot be modified. AutoNet comes with example data files on several subjects, such as forecast of air travel in coming months, stock exchange predictions, and treatment planning of rabies. AutoNet is very simple to use. One of the reasons is that it offers very few options. As its name suggests, AutoNet automatically constructs the network and determines all parameters without the user's involvement. It does so by using training examples of solutions to problems in order to learn the most appropriate network structure, whereas most networks only use these examples to adjust the weights of the different nodes of the network. A drawback of AutoNet is that you cannot change the network architecture or parameters in any way. There is no possibility to change the default values that the program generates.

One commercial product that is very useful in learning about neural networks is Brainmaker. For approximately $199 US (version 3–either DOS or Windows), Brainmaker comes with an excellent manual and an even better introductory book on neural network theory, design and applications (Lawrence, 1993). Brainmaker actually consists of two programs. The first one, Netmaker, allows you to manipulate your data in a spreadsheet format. Netmaker can read your data from different file formats, including Lotus or dBase files. Netmaker allows you to display and change your data and do some manipulation with your rows/columns. Basically, its aim is to generate and save a file that can be read by Brainmaker to construct, train, and apply the neural network. When offering data to Brainmaker, the program teaches the network in a number of cycles. A very useful option during the training is the test/save option, that allows you to test the network every n cycles and save the network at that stage. Later on, you can determine which network is most accurate in use. Once training is finished, you can use the network by presenting new cases. The network will then predict the outcome.

A limitation to Brainmaker is its restriction to backpropagated neural networks. No other learning algorithms can be used. However, backpropagation is by far the most effective algorithm, and if you are not very familiar with neural network theory, you will probably not want to play around with other algorithms.

Other well-known commercial products include Neuroshell, Neuralworks, and Explorenet 3000. Special mention should be made of Braincel, a shell that adds neural network intelligence to your spreadsheet. It allows you to handle your data within your spreadsheet and process it through the neural network, either in training or actual use, without having to leave the spreadsheet program. Unfortunately, you do pay a price in reduced speed.

Most of the programs mentioned have a base version selling at a few hundred dollars and a professional version with a considerably higher price. Base versions are definitely not primitive versions; on the contrary, they take you a long way.

Hardware

Most of the real applications of neural networks require a lot of computer resources. Although the available software takes you a long way, especially if you want to do a simulation of a neural network, actual use requires special hardware such as an accelerator board that can be plugged into standard PCs or an Intel i80170NX chip. This chip is also called ETANN or Electrically Trainable Analog Neural Network. Networks developed with BrainMaker can be implemented directly onto this chip. Other available hardware includes the ANZA neural network coprocessors or the DENDROS neural network chip. Some specifically designed workstations are also available. Unfortunately, most of this neural hardware comes with a price of several thousand dollars.

Written Resources

Neural networks and fuzzy set theory are both areas that draw a lot of attention and change rapidly. Traditional publications are, therefore, inadequate for following recent developments and keeping up to date. Those having access to the internet and electronic mail can subscribe to some electronic journals and newsgroups. Such a journal is the *Neuron-Digest*, now in its 11th volume. To subscribe, send an email to neuron_request@cattell.psych.upenn.edu with a request to be added to their list, and you will receive a copy of the journal regularly. Back issues are available through the internet by anonymous ftp. Discussions on neural networks and fuzzy sets are held continuously in comp.ai.neural-nets and comp.ai.fuzzy. Anyone can tap into these lively discussions through the internet facilities. Both newsgroups also have a FAQ-file (Frequently Asked Questions). These files often contain elaborate answers to the questions everyone asks when new to a certain newsgroup.

A large collection of neural network papers and shareware or public domain software is also available from the ftp-site of the Finnish University Network at the address funic.funet.fi. This site tries to assemble all the files available from other sites, so most of the publicly available material is there. However, it is also useful to check me.uta.edu of the University of Texas at Arlington that is a mirror site of the Central Neural System bulletin board. Among other things, it offers copies of the well-known ANNSIM file set (Artificial neural network simulations).

HUMAN SERVICE APPLICATIONS

As many human services have been both praised and decried for their "soft technology" (Hasenfeld, 1982), one can wonder what this soft computing has to offer the soft technology of human services. Neural networks are relatively new. So, one should not expect to have these applications consistently used by clinical practitioners or other human service professionals soon. Although the founding fathers of neural networks worked in the 1940s and 50s, the real interest started only a decade ago and applications were available only a few years ago. Indexing and abstracting services in applied science, such as Medlin for medicine, only started to reference publications on neural networks since 1992. A journal such as *Computers and biomedical research* had its first articles on neural networks in 1993. Within human services, apart from the articles in this issue of *Computers in Human Services*, only a very limited number of publications make reference to neural networks. Benbenishty (1992) mentions neural networks when discussing ways of eliciting knowledge from clinical practitioners and Mattaini and Kirk (1991) make reference to neural networks when discussing their possible use in assessment. Lutz and Flory (1993) react to this by indicating one of the major drawbacks of neural networks: "moreover, like the human mind, neural networks are essentially 'black boxes' that, although capable of using great amounts of data to make a decision, provide little information as to how that decision was made, leaving us in the dark as to the importance of any given factor."

One of the areas in which neural network applications are being introduced at great speed is medical diagnosis. Examples include sexually transmitted diseases, dermatological diseases, breast cancer, and many others. The input layer of these systems consists of either values on a number of variables (sex, age, medical variables) or images from medical equipment, such as electrocardiography or electroencephalograms. The output layer is the medical diagnosis. These networks have been trained on numerous examples of diagnoses made by human experts, and are able to reproduce the acquired experience on new cases. In a similar way, assessment instruments could be developed for use within human services, thus expanding the range of existing assessment approaches as described by Mattaini and Kirk (1991). The article by Brodzinski and Crable in this issue gives a good example of such an application.

Practitioners often face the important decision on which treatment to apply in specific cases. The range of available treatment techniques is vast, and it is no easy task to match these with the client's characteristics and problems. One can rely on theoretical frameworks, empirical clinical practice, or on one's own experience. However, many agencies have a collection of previous cases to draw upon in making decisions. Unfortunately, the necessary knowledge is too often hidden in the agency's client information database and difficult to elicit. Neural networks can help by "learning" from these databases and gaining reproducible experience in treatment planning by matching the client's information (input layer) with the treatment's results (output layer). In using such an approach, the quality of the client information becomes tremendously important.

An example of such a treatment planning system using neural network technology is currently being developed in Amsterdam, The Netherlands. The Municipal Bureau for Drug Treatment is constructing a system that will advise on whether to send a client to a physician, a hospital, a polyclinic, or whether to give him/her a methadone treatment. Other examples include the already mentioned applications on prediction of length of stay in the intensive care unit and psychiatric residential care.

A major difference in applying neural networks rather than expert systems in building treatment planning systems is the representation of the knowledge. When building expert systems, knowledge is elicited from clinical experts and is represented in a number of IF-THEN rules and conditions. These rules are fairly understandable and can be discussed. By doing so, clinical expertise can be captured (Nurius & Nicoll, 1992). When using a neural network, the knowledge is elicited from the client information database and represented as weights between nodes in the network. This is hardly understandable and cannot be discussed. Neural networks may be able to reproduce knowledge. However, they are certainly not able to explain why or to give insight into the inference flow of clinical practice. This is a major drawback if one's aim is to use information system development as a tool to research professional knowledge rather than reproduce it.

While few applications exist, the potential can be seen by examining successful business applications. For example, the risk assessment system described earlier suggests that the risk assessment process during a child abuse investigation could be supported by neural network technology.

CONCLUSIONS

Applications of this new science of neural networks and fuzzy logic will not be available in the near future, at least not in the human services. Still, there seems to be some promising developments in soft technology and soft computing coming together. I therefore agree completely with Benbenishty, when he writes about neural networks in clinical decision making: "this is a dynamic and growing field, however, and we should be alert to advances that might support attempts to model expertise in the human services" (1992, p. 608).

NOTES

1. The area of these new technologies is a rapidly evolving one, and its products therefore change at a tremendous speed. By the time this article is published and read, which takes a minimum time span of 6 months to one year, a comprehensive overview would be highly out of date. I therefore wish to give some guidelines towards existing products. Those still in need of a more comprehensive list of products are referred to general journals such as BYTE (e.g., January 1991, pp. 300-301) or more specific journals such as AI-expert (e.g., February 1993, pp. 48-55) or Neural Networks which has a list of suppliers in number 1 of each volume. The CUSSN Shareware, Freeware, and Demo Disk Copy Service can be contacted through D. Schoech, University of Texas at Arlington, School of Social Work, Box 19129, Arlington, TX 76019.

REFERENCES

Benbenishty, R. (1992). An overview of methods to elicit and model expert clinical judgment and decision making, Social Service Review, 66, 598-616.

Bowen, J. & Bowen, W. (1990). Neural nets vs. expert systems: predicting in the financial field, Sixth Conference on Artificial Intelligence Applications, IEEE Computer Society Press, Washington, 72-77.

Brodzinksi, J. & Crable, E. (1994). Using artificial intelligence to model juvenile recidivism patterns, Computers in Human Services, 10, in press.

Davis, G.E., Lowell, W.E. & Davis, G.L. (1993). A neural network that predicts psychiatric length of stay, MD Computing, 9, 87-92.

Dayhoff, J. (1990). Neural network architecture, an introduction, New York: Van Nostrand Reinhold.

Hasenfeld, Y. (1983). Human service organizations, Englewood Cliffs: Prentice Hall.

Kosko, B. (1991). Neural networks and fuzzy systems, Englewood Cliffs: Prentice Hall.

Kosko, B. (1993). Fuzzy thinking, New York: Hyperion.

Kosko, B. & Isaka, S. (1993). Fuzzy logic, Scientific American, 248, 62-67.

Lawrence, J. (1993). Introduction to neural networks: design, theory and applications, Nevada City: California Scientific Software Press.

Lutz, M. & Flory, J. (1993). Instruments and psychometrics: a response to Mattaini and Kirk, Social Work, 38, 229-230.

Mattaini, M. & Kirk, S. (1991). Assessing assessment in social work, Social work, 36, 260-266.

Nurius, P. & Nicoll, A. (1992). Capturing clinical expertise: an analysis of knowledge "miming" through expert system development, Clinical psychological review, 12, 705-717.

Schrodt, P. (1991). Prediction of interstate conflict outcomes using a neural network, Social Science Computer Review, 9, 359-380.

Tu, J. & Guerriere, M. (1993). Use of a neural network as a predictive instrument for length of stay in the intensive care unit following cardiac surgery, Computers and biomedical research, 26, 220-229.

Zadeh, L. & Kacprzyk, J. (1992). Fuzzy logic for the management of uncertainty, New York: Wiley.