ELSEVIER

Artificial Intelligence 119 (2000) 295–299

Artificial Intelligence

www.elsevier.com/locate/artint

Book Review

Gossiping Nets ☆

Rolf P. Würtz [1]

Ruhr-Universität Bochum, Institut für Neuroinformatik, D-44780 Bochum, Germany

The question of how the human mind works and what the basis of its admirable capabilities is probably goes back to the first human who thought about how the very act of thinking comes about. In the meantime, psychologists have acquired plenty of material on human behavior, learning, and perception, while physiologists have uncovered an immense collection of properties of nerve cells and the structures they form in the brain.

Starting with a paper by McCulloch and Pitts [1] these biological findings have been abstracted to a series of simplified models, which can be defined in mathematical terms and studied in their own right. According to this model, the brain is "nothing but" a set of simple elements interconnected to a huge graph with some laws governing the time course of the distribution of activity across the elements. Some of the elements are connected to sensors like eyes or ears and are considered input, others are attached to muscles and are considered output. For simplicity, the elements are called neurons, the graphs they form "neural networks" or "neural nets".

Since 1946, this basic idea has been followed, refined, tested and put to use. It turned out that the "nothing but" is a huge euphemism, both because the model is much too simple to cover the biological properties of a brain and even small neural networks can be extremely complicated systems. As for the question of how the mind works, many more open questions have been raised than definitive answers given. Formalized systems have been studied in considerable detail and some of them have become standard tools in statistics, control theory, and computer science.

Researchers working on this idea found themselves in the situation that no academic discipline was really responsible for their research and some even rejected it quite fiercely. Consequently, the now established research discipline of Neurocomputing continues to have strong influences from electrical engineering, mathematics, computer science, psychology, neurobiology, and physiology. Clearly the development of such a field cannot proceed without tension.

☆ Book Review of J.A. Anderson and E. Rosenfeld (Eds.), *Talking Nets*: *An Oral History of Neural Networks*. MIT Press, Cambridge, MA, 1998. 448 pp. $39.95 (cloth), $22.95 (paper). ISBN 0-262-01167-0. http://mitpress.mit.edu/book-home.tcl?isbn=0262011670.

[1] Email: Rolf.Wuertz@neuroinformatik.ruhr-uni-bochum.de.

0004-3702/00/$ – see front matter
PII: S0004-3702(00)00015-1

Anderson and Rosenfeld present a book of interviews with 17 persons, who have shaped that development in one way or the other. The interviews are very personal, highly subjective, and no attempt has been made by the editors to justify choosing one person or leaving out another. The complete list of interviewees is: Jerome Lettvin, Walter Freeman, Bernard Widrow, Leon Cooper, Jack Cowan, Carver Mead, Teuvo Kohonen, Stephen Grossberg, Gail Carpenter and Stephen Grossberg, Michael Arbib, James Anderson, David Rumelhart, Robert Hecht-Nielsen, Terry Sejnowski, Paul Werbos, Geoffrey Hinton, and Bart Kosko. Each interview tells a story of personal intellectual growth, of triumphs and setbacks, of struggle and achievement.

"Talking Nets" is a collection that may one day prove useful for writing the history of Neurocomputing. This purpose is served well by the index and glossary, both well done and very valuable. What would be desirable is a list of original papers to accompany the interviews, with citations at the appropriate points. Ten years ago, Anderson and Rosenfeld had already contributed to the material collection by compiling two thick volumes of original papers [2,3], but "Talking Nets" would benefit from a much longer reference list.

Their new book is different. Despite what might be guessed from the subtitle and from the choice of publisher "Talking Nets" is neither a history book nor a scientific textbook. Objectivity or scientific standards for many of the claims made in the book should not be expected. Such standards are much better served in, e.g., Arbib's Handbook [4], which covers most subjects treated in "Talking Nets" and adds plenty of references to fill in the details. The book under review here concentrates on gossip, highly elaborate gossip by highly distinguished scholars, but gossip nonetheless.

It is a book about people, and some of them are quite enigmatic. Parts are very personal, up to the point of expressing personal disregard for fellow interviewees. Depending on taste and perspective this can be perceived as anything between funny and disgusting. The same holds for some passages of self-promotion. It is safe to conclude from the book that the development of Neurocomputing has been dominated by male scientists. At least, Gail Carpenter is the only woman interviewed, but her having to share her chapter with her husband Stephen Grossberg, who also has his own separate interview, is certainly poor style.

Talking about style, the single-sided page layout of the book is plain ugly, and accidents such as depicting Bart Kosko in mirror image simply should not happen.

So much said, I can recommend reading the book for sheer fun. Take a moment to value the multiple meanings of the book's title, where "nets" may serve as subject or object, plus the reference to NETtalk, a simple yet very successful text-to-speech system, discussed in detail by Terry Sejnowski.

The single interviews make excellent bedtime entertainment, and a lot of interesting and important facts can be learned along the way. Whoever has played with the intriguing thought that the human brain might be mathematically tractable will enjoy listening to some of the gurus who have believed in and elaborated this hypothesis throughout most of their lives. It is clearly a book about people and their ideas, not the other way around. If you enjoyed "Cantor's Dilemma" by Carl Djerassi, "Talking Nets" is probably a good choice of book to take on a long flight.

The book contains some of the hilarious anecdotes about the group around McCulloch, told many times before but to my knowledge not available in print. Especially Jerome

Lettvin, Michael Arbib, and Jack Cowan give their views on why Norbert Wiener loathed Warren McCulloch, the rise and fall of the tragic genius Walter Pitts, and the games he played with security personnel.

More often-heard remarks not regularly seen in print become (sort of) citeable through this book, such as Karl Pribram's purported definition of the "Four Fs", which build the basis of intention and are controlled by the limbic system, namely Feeding, Fighting, Fleeing, and Reproduction.

Then there are the numerous stories about the unofficial mechanisms of science and their effects. Michael Arbib reports how he had to be kind enough to submit a different Ph.D. thesis, after the project he was assigned turned out to be too simple. Paul Werbos could not get his version of backpropagation across to the committee, because it was "too trivial and simple an idea to earn a Ph.D. at Harvard".

If planning and submitting a Ph.D. thesis can already be bad, wait for the stories about journal articles. Stephen Grossberg tells the story about submitting 10 articles to one journal and having them all rejected. Paul Werbos once received a review stating "We can't publish this because it is a challenge to good people like Rumelhart". Because the set of good ideas is probably of measure zero, surrounded by a huge mass of nonsense, scientific practice has several tough filters. The darker sides of peer review are laid out in several interviews quite explicitly. If a novel and brilliant idea is presented in an incomprehensible way it has a good chance of being rejected. Once it is boiled down to an understandable essence, chances are that this essence is picked up by others with better access to journals. This process may, but need not, include malice. It is sometimes very hard to track the origin of an idea in one's own head.

In the history of Neurocomputing the probably necessary conservatism inherent in peer review has been consequently circumvented by strong and at times ruthless attempts by individuals at setting up their own private communication channels, creating their own journals and organizing their own conferences. Refer to Robert Hecht-Nielsen and Bart Kosko to learn how to force IEEE into hosting a conference they do not want.

And, of course, the insults and injuries from the Ph.D. committees and the journal editors continue when it comes to employment and tenure. What makes the continued stories of rejection and unjust assignment of ideas bearable to the reader is the fact that the interviewees do not confine themselves to whining over their own rejections. Many go to considerable lengths to make the case for colleagues who have been treated unfairly. For anyone who considers questions such as "Who invented backpropagation?" or "Who invented Hopfield nets?" to be important, here is an abundance of data and views about them.

Another difficult matter is the relation of brain modelers and biologists. Arbib, Rumelhart, and Werbos complain quite explicitly about biologists' reluctance to accept ideas from the Neurocomputing field. This difficulty is probably more than another quarrel between academic disciplines; it has to do with the poorly defined relation between experiment and theory in the field, which are known to interact brilliantly in physics.

And, of course, there is the story of Minski and Papert's infamous "Perceptrons" book told from the victims' perspective. In that book they rigorously proved the limitations of the simplest of neural networks, the two-layer perceptron. Today it is well-known that by adding a third "hidden" layer of sufficient size, perceptrons can approximate every

function, which makes the above limitations quite irrelevant. However, the book created an atmosphere, which made it very difficult for neural network research to be published, let alone funded.

While it is established that "Perceptrons" had political influence far exceeding the relevance of its theorems, some of the battle between symbolic AI and subsymbolic techniques such as Neural Networks still persists. This is a pity because both aim at shedding light on the workings of the human mind, and that goal is so far in the distance that it appears unwise to rule out serious attempts on the grounds of taste. Mixing personal taste with the authority gained from mathematical proofs can be quite a potent poison, as the example shows. But the ways influence and funding are gained and redistributed are hardly ever entirely palatable.

Funding is another important theme covered from various aspects. Unfortunately, the philanthropic idealists who dream about improving the *conditio humana* by thoroughly understanding the human mind and getting better political decisions from that understanding (as Paul Werbos describes his motivation as a student) are usually not the ones who receive multimillion dollar grants. After several pages describing his sufferings, he consequently accepts a job at the NSF securing him influence on directing such grants.

Probably not very surprising for those familiar with the history of Artificial Intelligence is the high level of military involvement in funding and shaping the field of Neurocomputing. Illustrative examples include the definition of bionics as "the life sciences in the service of the death sciences" and the application of statistical methods to the prediction of military conflicts.

As for the very modern policy of preferably funding those technological developments which are guaranteed to lead to a consumer product within five years, the little story told by Leon Cooper should be read by every politician in charge of making funding decisions. On the other hand, the pressure to present results that really work has probably had a stimulating effect on the field and reduced the danger of drifting off into a purely abstract endeavor. The great fun and high motivation drawn from the ability to build systems doing something useful can be shared in many of the interviews.

Although some of the applications of neural networks are considered huge progress by their creators and commentators alike, it is clear throughout the book that the major drive behind the development of the field is the question of how human brains and minds work. Among the various goals of Neurocomputing, this is the one where the ratio between progress and significance of the problem is lowest. None of the interviewed appears particularly handicapped by an excess of humility. Nevertheless, they are not optimistic about rapid completion of the project of understanding the workings of the brain. Lettvin assumes this will never be the case, Hecht-Nielsen estimates a requirement of 300 years. This, together with the enthusiasm they all nevertheless show, amounts to maybe the best reason to read this book. There are many deep mysteries out there, and Neurocomputing offers a new and original path to pursue answers to a central human question.

References

[1] W. McCulloch, W. Pitts, A logical calculus of the ideas immanent in nervous activity, Bull. Math. Biophysics 5 (1943) 115–133. Reprinted in [2].

[2] J.A. Anderson, E. Rosenfeld (Eds.), Neurocomputing. Foundations of Research, MIT Press, Cambridge, MA, 1988.
[3] J.A. Anderson, A. Pellionisz, E. Rosenfeld (Eds.), Neurocomputing II. Directions for Research, MIT Press, Cambridge, MA, 1990.
[4] M.A. Arbib (Ed.), The Handbook of Brain Theory and Neural Networks, MIT Press, Cambridge, MA, 1995.