To appear in *Biological Psychology* (2001)

Introduction to Neural and Cognitive Modeling

Daniel S. Levine

Reviewed by
Robert M. French (rfrench@ulg.ac.be),
Quantitative Psychology and Cognitive Science,
University of Liège, Liège, Belgium

Overview

To call Daniel Levine's book an "introduction" to neural and cognitive modeling is somewhat of a misnomer. In fact, the book is an admixture of introductory material on neural network models and an overwhelming amount of material about the particular class of neural network models that the author favors — namely, those developed by Stephen Grossberg and his colleagues. The book attempts to lay out a systematic development of neural network modeling and is reasonably, if not completely, successful in the endeavor. Many major types of neural network models are at least briefly discussed, but an outsider reading this book would inevitably come away from it with a highly non standard view of the field.

The first and second chapters are exceptionally well-written, informative and clear. In fact, Levine takes the time to do what so many authors of the post-*Parallel Distributed Processing* generation do not do: he makes a reasonable attempt to trace the roots of the neural network modeling back to McCulloch & Pitts, Hull, and Hebb. The section on Rosenblatt's perceptron (Rosenblatt, 1962), the forerunner of neural networks is one of the best I've seen. Most books on connectionism discuss only the best known of Rosenblatt's perceptrons, those with an input layer, an output layer and a single layer of feedforward-only weights and an error correcting learning rule. Levine, on the other hand, goes into a considerably more detail and explicitly discusses a number of important but less well known aspects of Rosenblatt's critical contribution to early neural network research. This effort is much appreciated.

One regrettable oversight in this section was any discussion of Rosenblatt's remarkable (and unexpected) result, the Perceptron Convergence Theorem, that was perhaps the single most important reason for the interest in the early sixties in the perceptron as a pattern recognition device. The theorem states that if some combination of weights exists that allows the perceptron to recognize a particular set of patterns, then it is *guaranteed* to find a set of weights that will allow it to recognize the set of patterns (i.e., to converge to a solution). The problem, as Minsky and Papert (1969) showed, was with the requirement of the existence of some solution-set of weights. For some relatively simple patterns — in particular, those that were not linearly separable — they showed that either there were no solutions or that, if there were, they would require perceptrons of unreasonable size to recognize the pattern. And, while they put their finger on other weaknesses of the perceptron, the problem of the non-existence of solutions for certain important classes of patterns, rightly or wrongly, brought the entire thriving enterprise of perceptron research to a sudden and virtually complete halt. The difficulty is that, without some discussion of the Perceptron Convergence Theorem when presenting perceptrons, the reader cannot really appreciate the impact of *Perceptrons*, Minsky and Papert's detailed critique of perceptrons. Levine specifically attempts to minimize these results and writes that "Theorems of this sort [highlighting the limitations of the perceptron] were widely interpreted as discrediting perceptron-like devices as learning machines." He goes on to say that even Minsky and Papert admitted that the collective discrediting of perceptrons was an overreaction. (For a more accurate assessment, see Minsky and Papert's preface to the re-edited version *Perceptrons* that appeared in 1988.) The point is that *Perceptrons* had an enormously negative impact on the field of perceptron research that lasted for the better part of a decade and a half, even if certain research, most notably that of the most often cited neural network modeler in Levine's book, Stephen Grossberg, continued to work in the area. One has the distinct feeling that because Levine didn't like what Minsky and Papert's book did to the field — namely, brought it to a standstill — he deliberately downplays their critique. This annoying form of mild revisionism mars an otherwise excellent presentation of perceptrons.

After this introduction, Levine then takes us through a wide variety of neural nework learning techniques and strategies. The treatment is, on the whole, relatively thorough and informative. However, there are a number of problems that bear mentioning. One has the impression in reading this book that it would have been more appropriately entitled: *Grossberg et al.'s neural network models and some other minor models.* Consider, for example, the fact that there are 65 references in the author index to Grossberg, accounting for some 200 pages of text (in a book barely 400 pages long), compared to 28 references for Rumelhart, a scant 17 for McClelland, 12 for Hinton, 7 for Smolensky, etc. This is clearly a very skewed presentation of the field. Not wrong, mind you, but certainly non standard and intentionally so.

Further, I would have liked to have seen more discussion of McCulloch-Pitts neurons, since they serve as the basis for an enormous number of current neural network models. Just as connectionism did not spring whole cloth from the 1986 *Parallel Distributed Processing* volumes of Rumelhart, McClelland and the PDP research group, neither did it leap from Hebb (1949) to Rosenblatt (1962) in a single bound. There were numerous interesting attempts at neural modeling from that period (among them, for example, Rochester, Holland, Haibt, & Duda, 1956, etc. See Anderson and Rosenfeld, 1988) of which we are told nothing at all. When we reach association learning, we again encounter the "Grossberg problem." Now, I readily admit that the error-backpropagation learning algorithm is not the be-all and end-all of neural network modeling. Problems with its behaviorist stimulus-response nature, to say nothing of its neurobiological plausibility, were recognized and criticized early on (see, for example, Kaplan, Weaver & French, 1992). The issue of neurobiological plausibility is, indeed, discussed at some length later in this book. Having said that, is it reasonable to devote a mere 4 pages of the chapter to the error-backpropagation algorithm, which accounts for probably 90% of all neural network models, and four times as many pages to Grossberg's outstar algorithm? Kohonen's important work on unsupervised learning is included, but not enough time is spent developing the utility of these algorithms. Radial basis function models, a particularly widely used model loosely based on this approach, are not even mentioned.

The other chapters proceed roughly along the same lines as the second chapter. The subject is presented, a number of authors' work is presented and then a large part of the chapter is devoted to explaining how Grossberg and colleagues handled the problem. Now, while I do not have to be convinced of the importance of Grossberg's (and colleagues') work, this book is largely, if not exclusively, a presentation of Grossberg's models.

One of the most interesting aspects of this book is that Levine does attempt to keep the focus, as his title suggests, on neural and cognitive modeling. There is thankfully little (except in Chapter 8, "Recent Technical Advances," which is more of an afterthought than a real chapter) in this book on the practical uses of neural network technology for mine-versus-rock detection, for airport bomb sniffing, for determining bank-loan risks, etc.

Arguably, the most interesting chapter in the book is Chapter 6, "Coding and Catgorization." The author deals nicely with both supervised and unsupervised learning, discusses Anderson's (and colleagues') Brain-State-in-a-Box model and variants, the ubiquitous back-propagation models and, inevitably, about a third of the chapter is devoted to Grossberg et al's ART, ARTMAP, etc., family of models. It is a shame that Kohonen's work on self-organizing maps is not included in this chapter, its rightful place, or the considerable literature on radial-basis function networks. Another particularly glaring omission is the complete absence in this chapter (or elsewhere in this book) of models of exemplar-based categorization, the best known of which is John Kruschke's ALCOVE model. Of course, in a work of this type, one cannot possibly include all possible neural network models. However, one of the central questions in the field of categorization is whether categorization is best modeled by prototype models or exemplar-based models. The best known exemplar-based model is undoubtedly the Generalized Context Model (Nosofsky, 1986) and Kruschke's ALCOVE is a highly successful instantiation of that model. It has made a significant contribution to the categorization debate and, consequently, it is somewhat strange to see a chapter explicitly devoted to neural network models of categorization not include any mention of ALCOVE.

Another network that represented a significant advance in the neural network research enterprise was the simple recurrent network developed by Elman (1990). This model is referred to in passing on only two occasions in the book. However, the Elman network, compared to standard backpropagation networks used in the late 1980's, represented a significant step forward, providing the first plausible way to achieve learning of sequences in which context information was required. This type of network has been applied extensively to various sequential learning tasks and, most significantly, for the last decade it has been one of the neural network models at the center of the debate on implicit learning (see, for example, Cleeremans, 1993; Cleeremans & McClelland, 1991).

Chapter 7 ostensibly concerns "Optimization, Decision, and Knowledge," and might seem, therefore, to be of less interest to the reader interested in neural and cognitive issues. This, however, is not the case. The chapter, fortunately, ventures into areas that are clearly of cognitive interest. A case in point is a particularly clear discussion of a model that treats one of the major problems of high-level reasoning, the binding problem. When we see an object, such as a red ball, the brain must somehow bind the color ("red") to the form of the object ("spherical"), even though these are known to be processed in different parts of the brain (see Sougné, 1998, for a review). Even though binding is generally considered to pose major problems for neural networks, the problem is frequently overlooked in books on neural networks. It was, therefore, good to see that Levine had given Shastri & Ajjanagadde's (1993) solution the space it deserved.

There is also a welcome appendix on neurobiology that is worth mentioning. The author has managed to achieve a good level of presentation for his readership. It is particularly difficult to write this kind of section well. There is simply so much that can be said that it often happens that authors who attempt the task at all get bogged down in a welter of details about ion channels and neural interconnection patterns in different areas of the cortex. But here the chapter provides the naive reader with detail, but not so much to create bewilderment.

In conclusion, the one major criticism of Daniel Levine's *Introduction to Neural and Cognitive Modeling* is its non standard treatment of the subject. But *caveat lector,* the emphasis of this book is very much on neural and cognitive modeling *as done by Grossberg and colleagues.* Consequently, if that is what you are looking for, this is the book for you. Models from the Grossberg camp have traditionally been poorly explained and have had difficulty reaching a wider audience. This book is perfect for the person who wants to know more about Grossberg *et al.*'s models without plowing through the original papers. The explanations are clear and the mathematics is well presented. However, if you are looking for a good overview of current trends in neural network modeling, take this book with a grain of salt.

Anderson, J. & Rosenfeld, E. (1988). *Neurocomputing Foundations of Research* Cambridge, MA: The MIT Press.

Cleeremans, A. & McClelland, J.L. (1991). Learning the structure of event sequences. Journal of Experimental Psychology: General, 120, 235-253.

Cleeremans, A. (1993). *Mechanisms of Implicit Learning: Connectionist Models of Sequence Processing*. Cambridge, MA: MIT Press.

Kaplan, S., Weaver, M., & French, R. (1992). *Active Symbols and Internal Models: Towards a Cognitive Connectionism*. Reprinted in A. Clark & R. Lutz (eds.), *Connectionism in Context*. Springer-Verlag, 91-110.

Kohonen, T. (1984). *Self-organization and Associative Memory*. Berlin: Springer-Verlag.

Kruschke, J. K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review, 99,* 22-44.

Minsky, M. & Papert, S. (1969, re-edited 1988). *Perceptrons.* Cambridge, MA: The MIT Press.

Nosofsky, R. M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General, 115,* 39-57.

Rochester, N., Holland, J. H., Haibt, L. H., and Duda, W. L. (1956). Tests on a cell assembly theory of the action of the brain, using a large digital computer. *IRE Transactions on Information Theory, 2*, 80-93.

Rosenblatt, F. (1962). *Principles of Neurodynamics*. Washington, D.C.: Spartan Books.

Shastri, L. & Ajjanagadde, V. (1993) From simple associations to systematic reasoning. *Behavioral and Brain Sciences 16:3*. 417-494.
Sougné, J. (1998) Connectionism and the problem of multiple instantiation. *Trends in Cognitive Sciences*, 2, 183-189.