Reflections on Connectionist Modeling

Robert M. French
LEAD-CNRS UMR 5022, University of Burgundy-Franche Comté
robert.french@u-bourgogne.fr

In the 1980's I was a graduate student in the relatively new field of cognitive science, working in the lab of Douglas Hofstadter at the University of Michigan. The central theme of our research was *emergence*. The key question for us was how high-level, meaningful patterns emerged from a substrate made up of neurons devoid of all meaning. Aside from my work on emergent computational models of analogy making, which was the subject of my thesis, two classes of emergent models were of particular interest to me: genetic algorithms and connectionist models. As a first-year graduate student I took a course on connectionist modeling in which the professor, Stephen Kaplan, had us read basically every article ever written that had anything to do with networks of computational neurons -- something that was still possible before the Connectionist Revolution began in 1986. This revolution -- perhaps more a renaissance than a revolution -- was launched largely by the publication of Rumelhart and McClelland's (1986) seminal work, *Parallel Distributed Processing* (*PDP*). Sometime in 1985 a ring-binder preprint of *PDP* showed up in our lab. It contained, among many other things, the now-famous "feedforward-backpropagation" (FFBP) algorithm that overcame the problems with Frank Rosenblatt's "perceptrons" (very simple neural networks) that Minsky and Papert (1969) had brought to light in their book, *Perceptrons*. Minsky and Papert's book brought the entire embryonic field of neural network research to a screeching halt and it did not revive for another decade and a half. In any event, I coded up the FFBP algorithm in LISP and, even though it ran as slow as molasses, it did, indeed, work. I was incredibly excited by the whole thing, and have been a connectionist modeler ever since!

It has been fun to be able to play a part, however small, in the connectionist movement of the 1980's and early 1990's. The researchers who later became pillars of connectionism were, for the most part, not long into their careers and were all eminently approachable. Debates raged with the "opposing camp" of researchers from traditional artificial intelligence. There was an explosion of new connectionist algorithms to do everything from learning to read aloud to building a Turing Machine out of connectionist parts, from de-blurring images to predicting sunspots and exchange rates, and on and on.

But all was not perfectly rosy with the main actor in the Connectionist Revolution -- namely, feedforward-backpropagation networks. Backpropagation did not exist in real brains, said some. Backpropagation networks were *too* powerful, said others. They couldn't do one-shot learning. They were generally terribly slow and could only be trained by seeing the same items thousands, sometimes hundreds of thousands, of times. They had no attentional mechanisms. Sometimes they never converged at all. And so on. One of the most serious problems with these networks was that, once they had learned some set of patterns, learning a second set very often completely erased all learning of the first set (for a review see French, 1999). In attempting to find a solution to this problem, I met Steve Lewandowsky, co-author of the present book on computational modeling, who was attempting to do the same thing. And that, as Bogart says in the movie, was the beginning of a beautiful friendship.

Perhaps the question most often asked of computational modelers by empirical researchers is, "What's the point?" My reply is always this: Good computational models inform empirical research, and good empirical research informs computational modeling. And one of the most striking examples of this interaction comes from connectionism, long before it was called "connectionism".

In 1956 only a vanishingly small number of vacuum tube-filled electronic "calculators" existed (they weren't even called "computers" yet!) and one of them had been built by IBM. Four IBM researchers, among them, John Holland, co-director of my thesis, decided that the newly minted IBM 704 Electronic Calculator could be used to model the emergence of Hebbian cell assemblies from "an unorganized net of neurons" (Rochester, Holland, Haibt, & Duda, 1956). The neural network they programmed into the machine consisted of a mere 69 neurons. But, most importantly, *it didn't work.* Or at least it didn't work the way Donald Hebb said it should have, if his synaptic strengthening law[1], put forward in his seminal work, *The Organization of Behavior* (Hebb, 1949), was right. So, John told me they went to Montreal to discuss the situation with Hebb himself. The problem was that the neurons in their model, instead of gradually organizing themselves into "cell assemblies," each representing similar sets of inputs, would *all* became active regardless of what the input to the network was. They decided that, in addition to excitatory connections between neurons, the model -- and Hebb's theory -- also required *inhibitory connections* between some of the neurons. Upon returning to the IBM research center where they worked, they added inhibitory connections to their model and -- *mirabile dictu!* -- their modified model worked exactly as it should have. Cell assemblies formed and divided precisely as predicted by Hebb.

This is a textbook example of how modeling, theory and experimentation go hand-in-hand. In this case, empirical work by Hebb with Wilder Penfield and Karl Lashley led to his theory of how cell assemblies formed in the brain. His theory, however, did not include inhibitory connections between neurons, because at the time their existence had not been established experimentally. Rochester et al. (1956) then tried to build a computational model based on Hebb's theory and they found that, in order to produce the neuronal clustering that Hebb had predicted, inhibitory connections between neurons, in addition to excitatory connections, were necessary. This caused Hebb to modify his theory of how cell assemblies formed by including inhibitory connections (Hebb, 1959). It also stimulated neuroscientists to look for, and ultimately, find, inhibitory connections between real neurons in the brain. Examples of this kind of interaction between models and experimental research abound, although they are perhaps not as clear cut as this one.

After a slump in interest in connectionist models in the 2000's in favor of Bayesian learning techniques, an interest in connectionist models was reignited by the discovery of efficient algorithms to do "deep learning" (Hinton, Osindero, & Teh, 2006). The idea was to build a stack simple connectionist networks, each one feeding its output into the input of the next one. As information went from the input layer through each successive layer to the output, the features and structure that characterized the input were gradually discovered by the layers. Deep-learning connectionist networks have produced extraordinarily good classification/recognition of hand-written digits, faces, speech, animals, and so on. Many current speech-recognition applications, for example, rely on these networks. Once again, these networks are bio-inspired from real brains where information is processed by many layers before it reaches output effectors.

In short, I would argue that the recent advent of deep-learning neural networks is every bit as exciting as the mid-1980's when connectionist models returned to the scene after a decade and a half in the wilderness. And the sometimes acrimonious debates of the 1980's between connectionists and symbolic AI researchers are, fortunately, a thing of the past. Most importantly, computing power has gone from a mere 12,000 operations/second of the IBM 704 to the Chinese Sunway TaihuLight that in 2016 was clocked at 93 petaflops, i.e., 93,000,000,000,000,000 operations/second. Digital Research just built the world's largest neural network with 160 billion weights. And this is just the beginning....

References:

French, R. M. (1999). Catastrophic forgetting in connectionist networks. *Trends in Cognitive Sciences*, *3*(4), 128–135.

Hebb, D. O. (1949). *The Organization Of Behaviour*. New York: John Wiley and Sons.

[1] Hebb's Law: When an axon of cell A is near enough to excite a cell B and repeatedly or persistently takes part in firing it, some growth process or metabolic change takes place in one or both cells such that A's efficiency, as one of the cells firing B, is increased" (Hebb, 1949, p. 62). Or more succinctly, "Neurons that fire together, wire together."

Hebb, D. O. (1959). A neuropsychological theory. In S. Koch (Ed.), *Psychology: A Study Of A Science. Volume1: Sensory, Perceptual, And Physiological Foundations* (pp. 622-643). New York, NY: McGraw-Hill.

Hinton, G. E., Osindero, S. & Teh, Y.-W. A fast learning algorithm for deep belief nets. Neural Comp. 18, 1527–1554 (2006).

Minsky, M. and Papert, S. (1969). *Perceptrons*. Cambridge, MA: MIT Press.

Rochester, N., Holland, J. H., Haibt, L. H., & 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, IT-2,* 80-93.

Rumelhart, D.E., J.L. McClelland and the PDP Research Group (1986). *Parallel Distributed Processing: Explorations in the Microstructure of Cognition.* Cambridge, MA: MIT Press.