Chaos, symbols, and connectionism

John A. Barnden

Extracted from *Behavioral and Brain Sciences* *10* (2):174-175 (1987).

(It's a commentary on the target article by Christine A. Skarda & Walter J. Freeman, "How brains make chaos in order to make sense of the world", in the same issue of the journal, pp.161–195.)

As I have no doubt that the study of chaotic behavior in neural networks is an interesting and fruitful line of research, I shall confine my comments largely to some philosophical claims Skarda & Freeman (S&F) make and to the relationship of their model to connectionism.

Governing metaphors should be kept constantly under review as a matter of principle. But I am not convinced that S&F's findings and arguments constitute a serious threat to the "digital computer metaphor," or, more precisely, to views of the brain as a symbol-manipulation device. There is, first, the danger of extrapolating from findings and theories about low-level sensory mechanisms to high-level cognition. S&F are sensitive to the danger and freely admit that their model is extremely limited psychologically (see end of Sect. 3.6), but they would nevertheless like us to allow the extrapolation. More important, however, S&F do not give us any new reason to be worried about viewing higher-level cognitive processes as based on symbol manipulation. As far as I know, the "symbol manipulationists" have in any case always presumed that low levels of perception are at least largely based on specialized mechanisms that are probably not to be regarded profitably as manipulating symbols in any conventional sense. [The fact that AI (artificial intelligence) researchers and others simulate such mechanisms on digital computers is of course only weakly relevant here.] To show, therefore, that the olfactory bulb is best described as operating in a way foreign to symbol manipulation is not to push the backs of the symbol manipulationists any nearer to the wall.

The question of what the rest of the brain does with the output of the olfactory bulb (and directly connected brain centers) is significant. S&F themselves come near to suggesting that the bulb produces symbols when they say the inhalation of a learned odor pushes the bulb into a qualitatively distinctive, stereotypic state of activity. What is to stop us regarding these patterns as symbols? In what way is the idea that the rest of the brain uses these patterns in a symbol-manipulation style rendered implausible? I am not arguing that the rest of the brain does so use them, but only wondering what light is thrown on the issue by the S&F model.

S&F might reply that the activity pattern resulting from a specific learned odor varies somewhat in response to environmental context and preexisting internal state, and therefore cannot be regarded as a symbol. This point has some force, but there appears to be nothing to stop me from retreating slightly and saying that the rest of the brain proceeds to extract a symbol - corresponding to some invariant part or aspect of the pattern. (That such a part or aspect exists is surely at the basis of S&F's model.) Also, no retreat at all might be necessary if we allowed symbols to embody a certain amount of "fuzz." This would depart from the conventional view of symbols in artificial intelligence and cognitive science, but it is not clear that the unfuzziness of symbols is crucial to those fields, even if most researchers in those fields think it is. The most crucial aspect of the symbol-manipulation view seems to me to be the ability to form complex structures out of basic symbols, to analyze such structures, to compare symbols, and to associate symbols with symbols and other entities. None of these abilities requires unfuzziness of symbols in principle [Nelson Goodman's (1968) views on notational systems notwithstanding].

S&F would do well to be more careful about nomenclature when making their philosophical claims. The metaphor that their attack is directed at is surely the symbol-manipulation metaphor, not a metaphor of the brain as a digital computer as such, since it is clear that the brain is not like a computer at a low level of description. Now, at the electronic level of description a computer does not operate by symbol manipulation any more than a neural net does at the neurophysiological level of description, so that any terminology that confuses levels is likely to be misleading. When, for instance, S&F say in Sect. 4.1 that the process of odor recognition and discrimination can be conceived in terms of dynamic interactions at the level of the neural mass without appeal to symbols, we might well respond that the behavior of a program running on a computer can be conceived in terms of dynamic interactions at the level of the electronic mass without appeal to symbols. That this response would not (I take it) get at the heart of what the authors are saying would be their own fault, to put it abruptly.

While we are on the subject of levels, I dispute the implication in Sect. 4.2 that "self-organized" is correlated with "bot□tom-up" or that "rule-driven" is antithetical to "self-organized." A rule-driven system can be self-organized at the level of rules (since rules can modify themselves and other rules), and a top□down decomposition of a system can involve elements of self-organization at any level. Actually, it is not clear that either connectionists or S&F are adopting a bottom-up approach. To be sure, they are suggesting particular low-level mechanisms to explain particular high-level behaviors, but that does not make the approaches bottom-up. It is more that they are adopting a top-down approach different from those taken by certain other researchers. I am a little puzzled at the claimed divergence from connectionism with respect to types of feedback (Sect. 4.3). There seems to be nothing in the spirit of connectionism that disallows "locally dense feedback." Also, inhibitory feedback, which is claimed by S&F to have been given scant attention in connectionism, has played a very significant role in connectionist thinking for some time. One need only look, for instance, at the model of McClelland and Rumelhart (1981) and at the importance given to lateral inhibition in the Kohhonen (1984) book cited by S&F. On the other hand, I do agree that connectionists would do well to look more closely at the transmission delays and temporal dispersion effects on connections (whether or not they are feedback connections). I think it is really best to regard S&F's reliance on chaos and certain feedback effects as constituting a (most intriguing) extension of present-day connectionism rather than as diverging from it.

Finally, I would be interested to know what happens in the olfactory bulb and in S&F's model when several individually learned odors are presented simultaneously. Can a spurious output result, by virtue of the combined odors pushing the system into an activity state corresponding to another learned odor? Does this, if it happens, have any correlation to observed behavior? What new light is thrown on whether the output patterns can be usefully viewed as taking part in symbol manipulation?