The Importance of Chaos Theory in the Development of Artificial Neural Systems

by Dave Gross

Introduction

Neural networks are a relatively new development in computer science, having survived a brush with the exclusive-or problem while the field was still in its teens in the 1960s and recovered for a renaissance in the 1980s. Chaos is a new mathematical theory, dating back to perhaps the 1960s at the earliest and blooming only in the 1980s. The intersection of chaos with neurobiology dates back perhaps ten years. The use of chaos theory in the development and study of artifical neural systems (a.k.a. neural networks) is newer still.

This paper will briefly introduce the reader to the general concepts of artificial neural networks and of chaos theory, will discuss the research of Dr. Walter J. Freeman and others in the area of chaos and neurobiology, and will discuss the research on chaos and artificial neural systems. Finally, some conclusions will be drawn concerning the importance of chaos theory in the development of artificial neural systems.

This paper is written for the reader with a background in computer science. The discussions of neurobiology and of mathematics are therefore overly simplified, while the discussion of computer science and of artificial neural systems demands some degree of prior knowledge about these disciplines.

I would like to especially thank Dr. Walter J. Freeman (wfreeman@garnet.berkeley,edu) for sending me reprints of some of his papers on chaos in neurobiology, and Ice (ssingh@watserv1.waterloo.edu) for a list of references and abstracts relating to chaos in neurobiology and in artificial neural systems.

Artificial Neural Systems

Artificial neural systems are attempts to model some of the characteristics of the brain in order to capture and explore those qualities of the brain's reasoning power in which the architecture of the brain is assumed to play a major part. This has led to models which use connected local processing elements (neurons) accepting weighted inputs from other such elements and using these weighted inputs to give a single output which is in turn fed to other such processing elements, back to itself, or is given as an output from the system.

Much of the emphasis of neural network research has been in trying to more accurately simulate brain activity both on the microscopic (neuron) level and the macroscopic (overall brain activity) level. This has led to developments in areas such as Hebbian learning and unsupervised learning, which may have seemed counterintuitive to the pure computer scientists, but which had direct biological analogues.

Many of these biologically-oriented or simulation-oriented developments in neural networks have proven to have very practical results from a computer science point of view. Chaos theory has a good chance of being one of these developments.

To give some idea of how nondeterministic behavior might be produced by an artificial neural system, imagine a net with two layers and both feed-forward and feed-back output. One example input neuron in this system feeds its output back to itself with a high weight, as well as feeding its output to the neurons in the output layer, each of which has a low weight on the connection to this sample neuron (or, alternately, a higher threshold). Imagine that an initial input to that system causes that example neuron to fire an output which is not quite high enough to trigger the firing of any of the output layer neurons, but is high enough, when fed back to itself, to re-fire itself.

This neuron, after having been given this initial stimulus, fires itself cyclicly at a low level continuously. Now imagine that this system is given the same input a second time. This time, the example neuron not only gets the input stimulus, but also gets the stimulus that has been feeding back from its own cyclic firing. If this added input increases the output of the neuron significantly, it may trigger a firing of a neuron or neurons in the output layer -- producing a response to a given input that did not occur the first time that input was presented.

There are several variations on this mind-game that can be played. You can imagine, for instance, that instead of one neuron cycling the feedback to itself, that two or more neurons are playing "frisbee" with the feedback. In that case, the output for a given input will not only depend on whether that input has been seen before, but on which neuron is holding the "frisbee" at the time the input is presented to the network.

It's enough to make your own biological neural system spin.

Neurobiologists have found that such low-level activity is always present in the brain, but for a long time assumed that it was just irrelevant electric "noise." Now some believe that this activity, far from being random and irrelevant, is chaotic and essential to healthy brain activity.

In one study, for instance, researchers compared the pattern recognition capabilities of biological and artificial neural systems and commented that while "[p]attern recognition systems based on the perceptron... operate by relaxation to one of a collection of equlibrium states, constituting the minimization of an energy function" on the other hand "[b]iological pattern recognition systems do not go to equilibrium and do not minimize an energy function. Instead, they maintain continuing oscillatory activity, sometimes nearly periodic but most commonly chaotic." (Yao, Freeman, Burke & Yang 1991)

We can imagine "sometimes nearly periodic" activity with the frisbee analogy used earlier, but what is meant by chaotic activity?

Chaos: What is it?

Most computer scientists discover chaos in one way -- through colorful graphic displays of Mandelbrot sets on their terminals. Most of these computer scientists are content to watch the filigree unfold on their CRTs during lunch hour without delving too deeply into the mathematics behind it.

The curiosities of the Mandelbrot set or other graphs which display chaotic behavior{1} illustrate some of the interesting features of chaos theory. The boundaries of the commonly-pictured figures are irregular and intricate, and any attempt to magnify them only creates depictions just as magnificently irregular and intricate as the original. In fact, any two connected points on this boundary have an infinite length of boundary between them -- that's some measure of how convoluted this boundary is!

That such complicated patterns can result from seemingly simple mathematics is one feature of chaos theory. Chaos is statistically indistinguishable from randomness, and yet it is deterministic and not random at all. While it is deterministic in the sense that a chaotic system (on a computer, for instance) will produce the same results if given the same inputs, it is unpredictable in the sense that you can not predict in what way the system's behavior will change for any change in the input to that system.

One description, given by researchers who found chaotic activity in the brain, is that "[c]haos is controlled noise with precisely defined properties" (Skarda & Freeman 1987). A more complex definition is that "[i]n a dynamic system, chaos is a steady state solution of the system, but it is not an equilibrium solution, or a periodic solution, or a quasiperiodic solution" (Yao & Freeman 1990). The gist of these definitions is that chaos lies somewhere between periodic, predictable behavior and totally random behavior. It is random-appearing, and yet has a large degree of underlying order.

Chaos in the Brain

The existance of chaos in the brain has only been a major topic of discussion among researchers for less than ten years. In that time, chaotic behavior has been discovered both on the microscopic (neural) level and the macroscopic level in the brain.

One group of researchers, commenting on the discovery of chaos at the neural level, theorized that perhaps chaotic behavior could be responsible for schzophrenia, insomnia, epilepsy, and other disorders (Guevara, Glass, Mackey & Shrier 1983). Here, we will be most interested in the discovery of chaos on the macroscopic level in the brain{2}.

As a sharp contrast to earlier beliefs that chaos represented a possible source of harmful disorder in the brain, later researchers held that chaos was essential to proper brain functioning.

Dr. Walter Freeman of U.C. Berkeley's Department of Physiology-Anatomy has led the way in researching the role of chaos on the macroscopic level in the brain. Freeman's discovery of chaotic behavior in the electroencephalogram (EEG) tracings of olfactory bulbs in rabbits has led to a wealth of research on the role of chaos in the brain and in artificial neural systems.

Freeman noted that for some well-known but complex stimuli, recognition is almost instantaneous. A person recognizes a familiar face, or the scent of a barbeque, or the taste of chocolate almost as soon as that stimulus is presented to her.

"How does such recognition," Freeman asks, "happen so accurately and quickly, even when the stimuli are complex and the context in which they arise varies" (Freeman 1991). The answer he proposes is chaos.

Freeman found that there is constant activity in the olfactory cortex and that this activity is chaotic (Skarda & Freeman 1987). He believes that it is likely that the rest of the brain behaves in a similar fashion, and has proposed some possible reasons for this: "Chaos constitutes the basic form of collective neural activity for all perceptual processes and functions as a controlled source of noise, as a means to ensure continual access to previously learned sensory patterns, and as the means for learning new sensory patterns" (ibid), furthermore chaos "provides the system with a ready state so that it is unnecessary for the system to `wake up' from or return to a `dormant' equlibrium state every time that an input is given" (Yao & Freeman 1990).

A chaotic system in general, and the chaos exhibited in the brain, often alternates in a seemingly random way between various areas (or groups of behaviors) of the phase-space. These areas, known as chaotic attractors, are often called "wings" because an early model used in the discovery of chaos theory (the Lorenz attractor{3}) had two such areas that when graphically represented resembled butterfly wings.

The way the brain uses chaos to ensure continual access to previously learned patterns is to develop these wings for different learned inputs. According to researchers, the background chaotic activity enables the system to jump rapidly into one of these wings when presented with the appropriate input. "The transition back and forth between the wings or between the central part and one wing stands for phase transition{4} in the sense of physics and for pattern recognition in the sense of neural networks" (Yao & Freeman 1990).

If the input does not send the system into one of these wings, it is considered a novel input (e.g. an unfamiliar scent) and "instead of producing one of its previously learned activity patterns, the system falls into a high-level chaotic state rather than into the basin for the background odor. This `chaotic well' enables the system to avoid all of its previously learned activity patterns and to produce a new one" (Skarda & Freeman 1987).

Some researchers believe that this sort of chaotic background behavior is in fact necessary for the brain to engage in continual learning -- categorizing a novel input into a novel category rather than trying to fit it into an existant category.

"Without such a mechanism the system cannot avoid reproducing previously learned activity patterns and can only converge to behavior it has already learned" (ibid).

Chaos in Neural Networks

Once Freeman decided that chaos "may be the chief property that makes the brain different from an artificial-intelligence machine" (Freeman 1991), it was up to the artificial neural system researchers to narrow the gap.

Freeman himself was working on a computer simulation of the olfactory cortex by 1988, in part to allow for closer and more sustained monitoring of activity than was possible with EEGs on biological models (Eisenberg, Freeman & Burke 1989). That model, based on what was then known about the olfactory bulb and using only eight artificial neurodes, replicated many of the features Freeman found in the biological counterpart.

Other researchers created a simple artificial neurode model in which the individual neurons display chaotic behavior, modeling the behavior of biological neurons (Aihara, Takabe & Toyoda 1990; see also Ikeguchi, et al. 1990). At this point, however, the utility of single neurodes with chaotic dynamics is unknown, and macroscopic chaotic behavior can be modeled with more traditional artificial neurode models.

Some of the earliest research into macroscopic chaotic behavior in artificial neural systems discussed how chaos might crop up as an unintentional by-product of a system with feed-forward and feed-back neurode outputs (Hopfield nets, for instance). It was found that many such systems, if they have both excitatory and inhibitory connections between neurodes, can display chaotic behavior (Choi & Huberman 1983). Fukai & Shiino (1990) found similar results by assigning specific neurodes the task of either excitation or inhibition, rather than making the neurodes neutral and having the weighted connections either inhibitory or excitatory{5}.

Attempts to take advantage of chaos in artificial neural systems to reproduce benefits like those that Freeman and others have speculated are produced by chaos in the brain have met with some success. One researcher found that by adding chaos to a Hopfield-type net{6} it could be made to only recognize certain classes of inputs and not form patterns for others, thus engaging in selective learning (Sandler 1990).

The best indication that chaos can be practically utilized in artificial neural systems is in the performance of one that has already been developed. This chaotic system, designed to optically recognize four different types of industrial parts and determine whether or not they appear to be defective, was compared to non-chaotic artificial neural system implementations of the same problem{7} and was found to have significantly superior performance in positively identifying both acceptable and unacceptable parts (Yao, Freeman, Burke & Yang 1991).

Conclusion

Artificial neural systems were designed to capture some of the useful brain functions by modeling the features of the brain. Research into the function of the brain has led researchers to conclude that continuing background chaotic activity and chaotic dynamics in information processing are essential elements of biological neural systems.

The questions, then, are whether chaos theory is necessary for artificial neural systems which seek to duplicate the brain's abilities, and to what extent chaos can be exploited to improve the performance of artificial neural systems.

To the first question, there is as yet no answer. Dr. Freeman believes that chaos is essential for brain activity, and "is a quality that makes the difference in survival between a creature with a brain in the real world and a robot that cannot function outside a controlled environment" (Bower 1988).

Researchers like Freeman believe that systems that settle to equilibrium states or low-level oscillations rather than wells of chaotic activity are doomed to failure. They make the analogy to biological neural systems, in which these non-chaotic behaviors are indicative of coma, seizure, or death.

Others are not convinced. They see chaos as an understandable by-product of complicated systems like the brain or artificial neural systems, but one which in itself does not necessarily add to the efficacy of the system. Others, such as adaptive resonance theory creators Gail Carpenter and Stephen Grossberg, believe that the benefits that are supposedly offered by chaotic systems can be achieved in other ways, at least in artificial neural systems (ibid).

The evidence seems to show, however, both that chaotic activity in the brain provides specific advantages to the biological creature, and that chaotic activity in artificial neural systems has the potential to provide specific advantages to that system.

Some of the components of a successful artificial neural system displaying usefully chaotic behavior are: Inter-field as well as intra-field connections, and both inhibitory and excitatory weights. Other components which may prove useful are: Neurodes which are either wholly excitatory or wholly inhibitory, the ability to switch weights from positive to negative based on the state of the system, and neurodes which themselves display chaotic behavior.

Some of the beneficial behaviors we could expect from such systems are: Selective memorization, faster pattern recognition, recognition of new patterns as such and the development of new categories for these new patterns, and the ability to better distinguish patterns from background noise.

Many of these features have already been demonstrated (Yao, Freeman, Burke & Yang 1991; Sandler 1990), but only in very specific applications. The widespread use of chaos in artificial neural systems may be some time in coming, yet it seems unlikely that chaos theory will not play a part in the future development of these systems.

Notes

- {1} For a simple example, if you plot initial values for newton's method of solving for roots of the equation (x{4}-1=0) with a color corresponding to which of the four solutions the method finally converges to for that initial value, you will find wide regions of uniform coloration. Between these regions, however, will be borders which display a fascinating pattern of colors with seemingly little relation to their distance from the associated root. See page 6 of the color illustrations in Gleick (1987).
- {2} Some of the research on chaos at the neuron level is briefly summarized in Aihara, et al. They write, for instance, that "it has been clarified not only experimentally with squid giant axons but also numerically with the Hodgkin-Huxley equations that responses of a resting nerve membrane to periodic stimulation are not always periodic and that the apparently nonperiodic responses can be understood as deterministic chaos." A number of references to papers relating to chaotic neuron behavior are included.
- {3} See page one of the color illustrations in Gleick (1987) for a picture of the Lorenz attractor, or page 50 of the text for illustrations of how phase-space portraits are made.
- {4} An example of a phase transition in physics is that between the liquid and solid states of matter. There is a temperature, for instance, at which a small change in that temperature will result in a dramatic change (from liquid to ice) in the properties of water. Similarly, Freeman (1991) found that "neural collectives in the [olfactory] bulb and cortex ... jump globally and almost instantly from a nonburst to a burst state and then back again... [D]ramatic changes in response to weak input are, it will be recalled, another feature of chaotic systems."
- {5} This was to simulate the "Dale hypothesis" that in the brain each neuron has only an excitatory or an inhibitory nature.
- {6} For purposes of this discussion, consider a Hopfield net to be simply an artificial neural system with both feed-forward and feed-back connections. Sandler also included in his paper the requirement that for some states of the network, the weights of connections between neurodes be able to switch abruptly from positive to negative, and this was necessary for his results. Sandler found that such neuron connections have been known to appear in nature, such as in the chlorine synapses

of some chordless animals, and suggested that researchers try to find similar neurons in mammilian brains. This is an interesting case of neurobiological research into the brain prompting computer science research into brain simulation which in return prompts (one is tempted to say "backpropagates") further lines of inquiry to the neurobiologists.

- {7} Described as a neural network binary autoassociator, a three-layer feedforward network with back-propagation, the olfactory bulb model described in (Eisenberg, et al. 1989), as well as a standard Bayesian statistical method.

References

- Aihara, K., Takabe, T., & Toyoda, M (1990) Chaotic Neural Networks *Physics Letters A*, 144, 333-340
- Babloyantz, A., Salazar, J.M., & Nicolis, C. (1985) Evidence of chaotic dynamics of brain activity during the sleep cycle *Physics Letters*, 111A, 152-156
- Bower, B. Chaotic Connections (1988) *Science News*, 133 58-59
- Choi, M.Y. & Huberman, B.A. (1983) Dynamic behavior of nonlinear networks *Physical Review A*, 28, 1204-1206
- Eisenberg, J., Freeman, W. J., & Burke, B. (1989) Hardware Architecture of a Neural Network Model Simulating Pattern Recognition by the Olfactory Bulb *Neural Networks*, 2 315-325
- Freeman, W. J., Yao, Y., & Burke, B. (1988) Central Pattern Generating and Recognizing in Olfactory Bulb: A Correlation Learning Rule *Neural Networks*, 1, 277-288
- Freeman, W. J. & Yao, Y. (1990) Model of Biological Pattern Recognition with Spatially Chaotic Dynamics *Neural Networks*, 3, 153-170
- Freeman, W. J. (1991) The Physiology of Perception *Scientific American*, 264/2, 78-85
- Fukai, T. & Shiino, M. (1990) Asymmetric Neural Networks Incorporating the Dale Hypothesis and Noise-Driven Chaos *Physical Review Letters*, 64, 1465-1468
- Gleick, J. (1987) *Chaos: Making a New Science* New York: Viking Penguin
- Guevara, M.R., Glass, L., Mackey, M.C., & Shrier, A. (1983) Chaos in Neurobiology *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-13, 790-798
- Ikeguchi, T., Itoh, S., Utsunomiya, T., & Aihara, K. (1990) A dimensional analysis on chaotic neural networks *Electronics and Communications in Japan*, Part 3, V. 73 89-97
- Sandler, Yu. M. (1990) Model of neural networks with selective memorization and chaotic behavior *Physics Letters A*, 144 462-466
- Schoner, G. & Kelso, J. A. S. (1988) Dynamic pattern generation in behavioral and neural systems *Science*, 239, 1513-1519
- Skarda, C. A. & Freeman, W. J. (1987) How brains make chaos in order to make sense of the world *Behavioral and Brain Sciences*, 10, 161-195 with Open Peer Commentary

- Wang, L., Pichler, E. E., & Ross, J. (1990) Oscillations and chaos in neural networks: An exactly solvable model *Proceedings of the National Academy of Sciences*, 87 9467-9471
- Yao, Y., Freeman, W. J., Burke, B., & Yang, Q. (1991) Pattern Recognition by a Distributed Neural Network: An Industrial Application *Neural Networks*, 4, 103-121