LINKING CHAOS IN THE MODEL TO CHAOS IN THE BRAIN

KAOLIN FIRE

ABSTRACT. A plausible correlation is made between a period-doubling artificial neuron pair and experimentally found chaos in the squishy human brain. Preliminary statistical evaluations are done.

1. INTRODUCTION

Chaos is everywhere; it's in the orbit of the planets, it's in our weather, it's used in computer graphics. Since the second century A.D., the brain has been generally accepted as the organ of the mind, and since the 1980's, chaos has been on the brain. Chaos is in our brains. Chaos has been found in how we process external senses, and may be key to memory. It helps explain the figure/ground problem, and has been implicated in at least one theory of the evolution of vocabulary [1] as well as synaesthesia [2]. But most models of the brain do not include chaos. Those that do, don't seem convincingly biological.

On one end of the spectrum, we have a simple computer model for a neuron which has been analytically proven [3] to have regions of period doubling and all that chaotic goodness.

On the other end of the spectrum, we have the squishy human brain. The squishy human brain (or at least the squishy cat brain, a close relative for our purposes) has been shown to have elements of chaotic behavior in neural spike trains in the visual system. [4] Statistical voodoo (see section 1.4) can be used to test the dimension of this information, and show that it is indeed fractal.

Unfortunately, the chaotic goodnes of the simple computer model for a neuron has not been mapped in any way that I can find to the chaotic goodness of the squishy human brain.

So I propose to give an overview of all related factoids while attempting to tie, in a "there exists" sort of manner, one to the other. Preferably with a couple of applets and pictures.

Ideally I will be able to create spike trains out of this simple computer model for a neuron which resolve under the statistical voodoo.

Whether that will mean anything is anyone's guess.

1.1. **Chaos basics.**

So what is this chaos? Chaos is not disorder. Chaos is actually quite ordered. Chaos requires three things [5]:

- elements of regularity (periodic points that are dense)
- elements of oneness (dense orbits)
- elements of unpredictibility (sensitive dependance on initial conditions)

What does that give you? Infinity, of a sort, non-integer dimensions of "arbitrary" complexity.

To understand these things three, we will need to understand what an iterated function is. An *iterated function* is a function that maps from one space back into the same space.

Date: May 22, 2003.

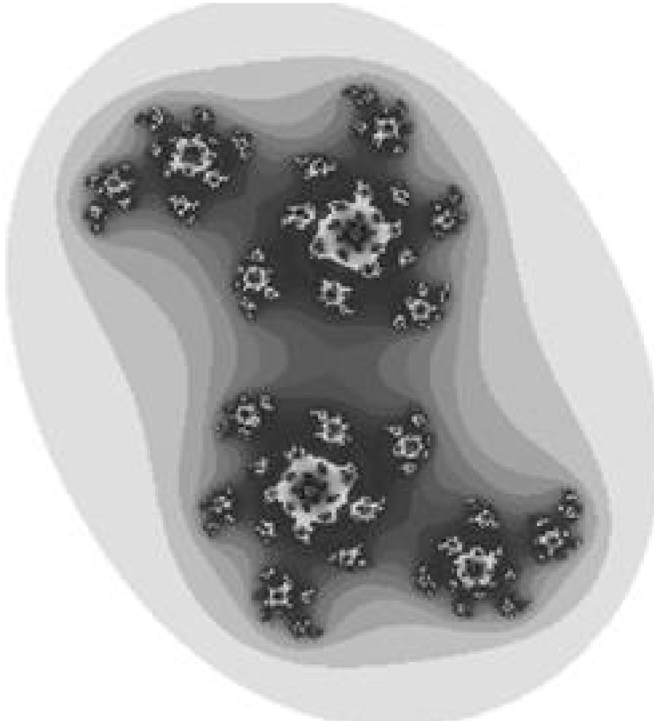

FIGURE 1. a gratuitous shot of *The Julia Set*

This is the "engine" of chaos. We will define the function f to map a point $x \in X \to X$. Iterating the function, we get $f^1, f^2, \ldots$. We define an *orbit* to be the sequence of values followed from some seed x. If for some $m, n, m \neq n$, $f^m(x) = f^n(x)$, then we say the orbit is periodic. A *fixed point* is a point which maps back to itself directly.

Elements of regularity, or periodic points that are dense, bespeaks of the ability to find a seed that is periodic arbitrarily (in the mind of the beholder) close to any other seed you'd like. That is, within any interval around an aperiodic seed, a periodic seed can be found. For instance, presuming the space of $\mathbb{R}$, we could pick our (nonperiodic) seed x to be 1.349156, and having periodic points that are dense would imply we could find a periodic seed y within a happy ε of x.

Elements of oneness, or "dense orbits", implies that there exists at least one orbit that gets arbitrarily close to *every* point in the set. This is also referred to as a "transitive orbit". While this may be hard to believe, the quadratic map, where chaos is easy to find, is easily mapped into symbolic space, where a transitive orbit tends to be trivial to construct.

The final point, sensitivity on initial conditions, is the one that gave rise to the popularization of "the butterfly effect". That is, any two points arbitrarily close to each other will eventually be split by an arbitrary distance. The slightest difference in the seed, even just in the *precision* used to represent the seed (as Edward Lorenz found in his differential models of weather patterns [6]).

A chaotic system is inherently dynamic. It is simple to conceive it as a system constantly "in motion". Chaotic systems can be both robust and sensitive. The slightest *nudge* can cause a periodic or nearly periodic orbit to essentially jump states. Or the state can resist stern perturbation. It depends on the system, and the state.

This becomes important later on.

For now, suffice it to say that this makes for some very pretty pictures, most famously, the Mandelbrot and Julia sets (as seen in Figure 1). Also interesting is the study of *piecewise*

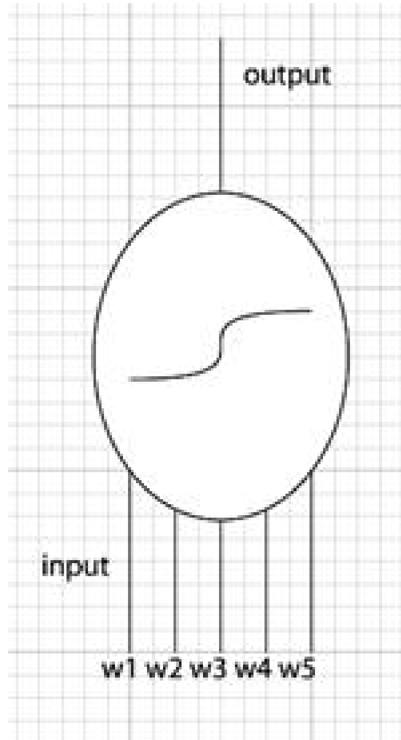

FIGURE 2. A clunky looking perceptron

isometries [7], systems mapping geometric constructions by rotations and translations back onto themselves; the simple trick is to represent the figures as vectors of their points.

1.2. **Artificial Neural Net basics.** Artificial neural nets are inherently discrete time beasties, a directed, typically acyclic, graph of perceptrons. Any individual perceptron has a number of inputs and outputs, and an internal state, all of which tend to be represented by elements of the $\mathbb{R}$ (see Figure 2). The outputs tend to be functions of the summation of weighted inputs. A typical "firing", or "activation" function is the *sigmoid*:

$$\frac{1}{1+e^{-z}}$$

where z is the sum of weighted inputs. This function determines the next state of a given perceptron, which is generally synonymous with its output.

1.2.1. *Learning in Neural Networks.* Imagine if you will, this directed, acyclic graph. It has inputs, and it has outputs. The question becomes, "What does it *do*?" And that all depends on what you want it to do.

It is difficult to create a neural net and match it to any complex task entirely by hand. This is where the simplest form of artificial intelligence comes into play. The human decides what inputs exist, what the output is supposed to signify, and generally how many neurons in how many 'layers' the neurons will be divided into. Layers other than input and output are referred to as *hidden layers*; they are 'beyond our ken' (or at least, care). We only care about inputs and outputs, supposedly. The machine is a black box.

Random values are chosen for all weights, and the system is run. The output is tested against the desired output, and (typically) the weights are updated based upon that difference. These differences are then *backpropogated* through the network, and the system is run again. The system is run on a large set of patterns so that it will 'know' all of them, and it is hoped that then the system has 'learned' some generalized pattern and will be able to respond appropriately to inputs that were not part of its original test set.

The grandfather of learning in artificial neural networks is Donald Hebb, whose 1949 paper [8] laid the foundations. The Hebb learning rule has a maximum *capacity*, the number of patterns it can remember, of about 0.14 the number of neurons included. Hebb learning suffers from *catastrophic forgetting*[1]; if the hebbian neural network reaches capacity, it becomes unable to retrieve any of the previously stored patterns.

'Forgetful', or 'Palimpsest', learning rules have been invented to avoid this catastrophic memory failure; they fail more gracefully, replacing older memories with new patterns. Unfortunately, nonchaotic palimpsests have been shown to have a capacity of only up to 0.05n.

Much more can be read in the canonical usenet neural network FAQ. [9]

1.2.2. *Chaos in Neural Networks.* Your typical neural network, the non-cyclic, or feed-forward, can not exhibit chaos. No matter how many perceptrons you toss into the system, there is no recursion; there is no self-reference. Recurrent Neural Networks (RNN's) introduce closed paths in the graph topology. Thus, we have a self-referencing, iterative system, and the *possibility* of chaos.

Now we jump to a minimal, two perceptron RNN, fully connected. N_1 is an input to N_2 and vice versa, and they are each "connected" to themselves. There are very solid areas of chaos in this system. [3]. Forgetful learning rules that ride the edge of chaos have been shown to have a capacity as high as 0.25n! [10] This ties in, somehow, to [11], though I'm not sure how much. And here's a nifty article on "The speed of information transfer". [12]

1.3. **Biological Neuron basics.** A neuron is a cell. It has all the things that animal cells have, only it's a little more specialized. It has inputs called *dendrites* and typically has one output *axon*, which may connect to many dendrites over *synapses*. A synapse is a "gap" between the axon and the *dendritic spine*. Spine formation is constantly in flux, which may be key in short- and long-term memory.

The adult human brain, weighing in at roughly 1.3kg, has on the order of 10^{11} neurons [13]. The cerebellum, at 150 grams, has roughly 10^{10} neurons. These together have on the order of 10^{14} dendritic spines, or an average of ten thousand synapses apiece.

By Hebbian updating, were these artificial neurons, the entire brain would be able to hold 1.4×10^{10} patterns, and after that extreme senility approaching death would set in. Linear 'forgetful' learning rules would see only 5.0×10^{9} patterns. With a dash of chaos, though, we're approaching 2.5×10^{10} patterns. Again, the question remains "How does all this relate? Does it really at all?"

The average adult knows only between 10,000 and 20,000 words, and typically uses no more than 3000 different words in daily conversation. On the other hand, words often represent fairly complex ideas; language is a complicated beast, and sometimes it's amazing that we've managed to make any progress with it. [1], [14]. On top of that, it has to be recognized how many patterns are built up of hundreds or thousands (or more) of other patterns.

Let us take the visual system, for instance. Any picture you can imagine has blobs of color, lines, shading... There are areas of the brain that code for, essentially, pixels. These feed into areas that code for edges at any point in your visual field, which then feed into

[1]The more patterns it learns, the less "distance" there is between any two patterns. At the point of catastrophic forgetfulness, the distance has dissolved to noise. All patterns have essentially been lost.

detectors for lines at any point in your visual field of any orientation (360 degrees divided into a resolution on the order of a single degree!). These then feed into motion detectors that only respond to motion of certain orientations in certain directions. Line detectors also feed into cells that only respond to, say, right angles of a certain orientation. And it goes on from there.

Which doesn't touch on the auditory system, tactile system, or the motor system, late alone actual 'memories'.

Before moving on to (or through) the chaotic definition of "elements of regularity", it is interesting to note that the brain appears to like, in that it requires less effort, elements of geometrical regularity. This has been investigated thoroughly in the visual system [15] and is likely present in each of the senses. Areas of the brain tend to be found to have evolved into structures requiring a minimum amount of energy to represent and navigate "the real world". For instance, studies on cats show that neurons are far less active when viewing a "real world" scene than white noise. [16]

1.3.1. *Chaos in Biological Networks.* With the numbers waved about vaguely, we move on and ask what chaos has to do with the brain? Does chaos have anything to do with a functioning human brain?

Though there is some dissent [17], early research [18][19] is still bearing fruit.

Indeed. There is chaos here, chaos there; chaos everywhere. And out of chaos, order (or apparent disorder, depending on how you live your life).

This section needs more fleshing out, but to poke the highlights: the 1950's Hodgkin-Huxley model of everyone's favorite neuron, that of the giant squid, displays a parameter range where chaotic dynamics appear, though it is not certain whether they appear under 'normal operating conditions'. It *is* certain that the dynamics of many 'normally operating' neural systems are fractal.

A biological neuron does not simply "sum and fire" based on its inputs; firing tends to follow what appears to be a poisson distribution. Where does this perceived randomness come from? Sensitivity to initial conditions.

1.4. **Chaos detector (Statistical Voodoo).** This is laid out ad nauseum in [20], and a bit more simply in the introduction of [4]. Essentially, one takes the "rate estimate of neural firing" over a set of counting intervals T. Fractal processes exhibit 'slow' power-law convergence: The standard deviation of the firing rate decreates more slowly than $1/\sqrt{T}$ as the averaging time increases. Nonfractal signals, on the other hand, decrease precisely as $1/\sqrt{T}$. Randomly reording the neural spike intervals loses long-term correlation, the fractal nature of the train, and thus falls at precisely at $1/\sqrt{T}$, while the original (natural, fractal) signal falls more slowly.

2. Putting together Brain and Computer

It is shown that the chaos of the artificial neural net can be usefully interpreted in two equally valid ways. The first treats their values at any given clock cycle as the "next" time period between two spikes, equalized against a max and min from sample experimental data. The second uses a threshold function to discretize whether or not the artificial neurons can be said to have "fired". Time between firings is then inherently measured in "clock cycles". Statistical Voodoo is used to show that regardless of the interpretation, the fractal dimension of the information remains constant.

3. Proofs or Experimental Methods or Simulations... Stuff.

Here I put the chaotic 2-perceptron network through its paces. First, I'll state what I expected to see:

The first data set is from a run of the chaotic 2-perceptron network for a 'simulated' 5000 seconds, counting each 'cycle' as one hundredth of a second; the hope was for between 20 and 120 spikes per second. [21] A spike is considered 'generated' when the perceptron's state/output is greater than .5 (on the sigmoid scale of 0 to 1).

The second data set is from a run of the chaotic 2-perceptron network for a 'simulated accumulated' 5000 seconds, where at each clock cycle, a value of 0 is taken as a rate of '20 spikes per second', or a delay of .05 seconds until the next spike; a value of 1 is taken as a rate of '120 spikes per second', or a delay of .0083 seconds until the next spike, with the range between interpolated linearly.

Data set one aleph is the first data set randomly reordered; data set two aleph is the second data set randomly reordered.

All four data sets were then put through the power law paces outlined in section 1.4 (Statistical Voodoo).

For completeness sake, another four data sets would have been created from a nonchaotic range of the perceptron pair period-doubling analog, put through their paces paces as well; however, both methods outlined failed to provide a significant standard deviation. The deviance shrinks from 'small' to 'zero' too quickly to catch an accurate reading. It is likely that this is due to many rounding errors in various calculations; they require further study. It is quite premature to question the validity of the chaos (as proven analytically to exist), and also premature to question the validity of the statistical voodoo.

4. Discussion

What have we proved? Apparently nothing, though this interdisciplinary review hopefully brings many ideas together, and can be a launching point for further research and ideas. More reading needs to be done regarding the *Information Theoretic Limit* [22] and how that relates to chaos [23], and other such information theory things regarding neural spike trains [24], [25], [26], [27], [28], [29].

Adding chaos to an active model of the brain could help in reconstruction of images based on electrical probes. [30] And more interesting models could be made and tested in all sorts of ways.

⋮

fizzle

References

[1] Tao Yang. Dynamics of vocabulary evolution. *International Journal of Computational Cognition*, 1(1):1–19, March 2003.

[2] V. S. Ramachandrann and E. M. Hubbard. Synaesthesia: A window into perception, thought, and language. *Journal of Consciousness Studies*, 9(3):3–34, 2001.

[3] Xin Wang. Period-doublings to chaos in a simple neural network: An analytical proof. *Complex Systems*, 5(4):425–44, 1991.

[4] Malvin Teich, Conor Heneghan, Steven Lowen, Tsuyoshi Ozaki, and Ehud Kaplan. Fractal character of the neural spike train in the visual system of the cat. *Journal of the Optical Society of America*, 14(3):529–46, 1997.

[5] T. Y. Li and J. A. Yorke. Period three implies chaos. *American Mathematical Monthly*, 82(1):985–92, 1975.

[6] E. N. Lorenz. Deterministic nonperiodic flow. *Journal of the Atmospheric Sciences*, 20:130–41, 1963.
[7] Arek Goetz. Piecewise isometries – an emerging area of dynamical systems. http://math.sfsu.edu/goetz/Research/graz/graz.pdf, June 2001.
[8] D. O. Hebb. *The organization of behavior*. Wiley, New York, 1949.
[9] W.S. Sarle editor. Neural network faq, part 1 of 7: Introduction, 1997. periodic posting to the Usenet newsgroup comp.ai.neural-nets.
[10] Amos Storkey. Palimpsest memories: a new high-capacity forgetful learning rule for hopfield networks, 1998.
[11] Clemens Wagner and Jörg W. Stucki. Construction of an associative memory using unstable periodic orbits of a chaotic attractor. *Journal of Theoretical Biology*, 215:375–84.
[12] C. Wagner and R. Stoop. The small world of fractal coupling. in preparation.
[13] Ph.D. Eric H. Chudler. *Fun Facts About the Brain*. http://faculty.washington.edu/chudler/ffacts.html, 2003.
[14] Christian Balkenius, Peter Grdenfors, and Lars Hall. The origin of symbols in the brain.
[15] Manfredo Massironi, Pala Rocchi, and Cesare Cornoldi. Does regularity affect the construction and memory of a mental image in teh same way it affects a visual trace? *Psicológica*, 22:115–42, 2001.
[16] Y. Dan, J. J. Attick, and R. C. Reid. Efficient coding of natural scenes in the lateral geniculate nucleus: experimental test of a computational theory. *Journal of Neuroscience*, 16(10):3351–62, 1996.
[17] Milan Palus. Nonlinear dynamics in the eeg analysis: Disappointments and perspectives.
[18] Walter J. Freeman. Simulation of chaotic eeg patterns with a dynamic model of the olfactory system. *Biological Cybernetics*, 56:139–150, 1987.
[19] Walter J. Freeman. Strange attractors that govern mammalian brain dynamics shown by trajectories of electroencephalographic (eeg) potential. *IEEE Transactions on Circuits and Systems*, 35(7), July 1988.
[20] S. Thurner, S. B. Lowen, M. Feurstein, C. Heneghan, H. G. Feichtinger, and M. C. Teich. Analysis, synthesis, and estimation of fractal-rate stochastic point processes. *Fractals*, 5:565–95, 1997.
[21] Fred Rieke, David Warland, Rob de Ruyter van Steveninch, and William Bialek. *Spikes: Exploring the Neural Code*. Computational Neurosciences series. MIT Press, 1997. review by Cosma Shalizi.
[22] http://www.ieeeits.org/review/01csi.pdf.
[23] http://frt.fy.chalmers.se/cs/cas/courses/infotheory/pdfs/itcs_09.pdf.
[24] Rob R. de Ruyter van Steveninck, Geoffrey D. Lewen, Steven P. Strong, Roland Koberle, and William Bialek. Reproducibility and variability in neural spike trains. *Science*, 275:1805–8, March 1997.
[25] Emery N. Brown, Riccardo Barbieri, Valérie Ventura, Robert E. Kass, and Loren M. Frank. The time-rescaling theorem and its application to neural spike train data analysis. *Neural Computation*, 14:325–46, 2001.
[26] Peter E. Latham. Yet another method for calculating information from neural data, 2000.
[27] Peter E. Latham. Associative memory in realistic neuronal networks, 2001.
[28] Peter E. Latham. Computation and memory in recurrent networks, 2002.
[29] Jean-Pierre Nadal. Information theoretic approach to neural coding and parameter estimation: a perspective.
[30] G. Stanley, F. Li, and Y. Dan. Reconstruction of natural scenes from ensemble responses in the lateral geniculate nucleus. *Journal of Neuroscience*, 19(18):8036–42, 1999.