Emergent Complex Patterns in Autonomous Distributed Systems: Mechanisms for Attention Recovery and Relation to Models of Clinical Epilepsy*

Elan L. Ohayon[1,4,a], Hon C. Kwan[3], W. McIntyre Burnham[1,2], Piotr Suffczynski[4,5], Stiliyan Kalitzin[4]

Invited Paper

Abstract - **Dynamical systems based on distributed elements can exhibit complex autonomous behavior. Simultaneous existence of separate stable dynamic states (attractors) and the transitions between them can model certain forms of epileptic discharge. Multi-stable systems have also been proposed for storage and retrieval of activation patterns. Here we consider systems with alternative types of collective behavior. In these systems emergent intermittency allows for autonomous switching between turbulent (chaotic) and laminar phases. We demonstrate that the distributions of the duration of various phases have distinctive statistical properties, different from those in multi-stable systems that are driven by stochastic processes. These properties are proposed to identify and classify mechanisms that may underlie paroxysmal activity as revealed in electrophysiological recordings of epileptiform activity. Unlike spontaneous stochastically-driven ictal transitions in multi-stable systems, certain features of intermittency-based transitions can, in principle, be forecasted and perhaps even ameliorated. We show that intermittency in a recurrent network does not require plastic connections. At the same time, we argue that an autonomous system with modifiable connections might require intermittent transition mechanisms in order to sustain proper connectivity and function. Networks showing intermittency avoid lockups and at the same time respond robustly and commensurably to dynamical input perturbation. They may thus provide a candidate mechanism for pattern recognition and attention recovery in biological and artificial systems.**

Keywords: Intermittency, recurrent neural network, computer modeling, seizures, autonomous neurodynamics.

[1]University of Toronto Epilepsy Research Program and Institute of Medical Science, [2]Department of Pharmacology, Medical Sciences Building, Room 4303, University of Toronto, 1 King's College Circle, Toronto, Canada (email: ohayon@chass.utoronto.ca; mac.burnham@utoronto.ca)

[3]Department of Physiology, Medical Sciences Building, Room 3232, University of Toronto, 1 King's College Circle, Toronto, Ontario, Canada, M5S 1A8 (email: h.kwan@utoronto.ca)

[4]Stichting Epilepsie Instellingen Nederland, Achterweg 5, 2103 SW Heemstede, The Netherlands (email: psuffa@sein.nl; skalitzin@sein.nl)

[5]Laboratory of Medical Physics, Institute of Experimental Physics, Warsaw University, Hoza 69, 00-681 Warsaw, Poland.

[a]Corresponding author: Elan Liss Ohayon[1].

I. INTRODUCTION

Under most conditions, neural-based organisms can attend and respond to a range of environmental inputs while avoiding permanent lockup of their dynamics. This processing and autonomous behavior is done by highly recurrent neural systems that display continuously changing dynamics. How networks independently change between dynamical states (attractors) has been a long-standing question in neural network theory [25]. It is often stipulated that either these networks are (i) driven between various disjointed attractors via perturbations or (ii) a parametric change to the system is introduced causing a change in the phase space topology such that there is an expansion of an existing basin or the introduction of a new attractor via a bifurcation. The transitions under both scenarios are highly dependent on factors extrinsic to the current dynamics. There is no way for the behavior to start, persist or end autonomously without help from external parameter tuning or stochastic fluctuations. This quandary has its counterpart in our understanding of epileptic phenomena. Epilepsy is the prototypical dynamical ailment characterized by abrupt and significant shifts in dynamical behavior. These changes in the brain are generally considered to come about either through (a) a deformation or bifurcations of the system, (b) through shifts between preexisting attractors, or (c) a combination of the two above scenarios. Transitions in and out of seizures are thus thought to occur by endogenous or external perturbation that may be as random as environmental noise [12], [13], [24]. As such, the issues that arise in epilepsy are remarkably similar to the healthy behavioral case. For example, if we consider attention as a dynamical state in which the system prepares for recognition or action, we can posit that the shifts in attentional states might be triggered by the environment or even an alternate brain structure that modulates alertness. However, here again we have to conceive of a perturbation (contentful or noise) that will trigger these transitions and subsequent reversal so as to bring about a return to the initial condition. But what if the mechanisms that trigger and end seizures or attention are intrinsic to the very dynamical properties of the system? Here we explore the possibility of an intermittency mechanism in a recurrent neural network that does not assume a change in the system topology nor an external input to trigger changes in activity.

II. METHODS

A. The Recurrent Neural Network Model

We used fully interconnected recurrent neural network models including self-feedback (Figure 1). The networks had 5 units with 25 connections $(w_0..w_j)$. The connection strength between units was represented by a weight that ranged from -3 to +3. At any given time the total input (E_i) for a unit was the weighted sum of the activations of the corresponding input units (S_j) such that:

$$E_i = \sum_j w_{ij} S_j \qquad (1)$$

The sum input (E_i) was then fed through a non-linear activation function which yielded the unit's new activation which is also the output for that unit (S_i):

$$S(E_i) = e^{(-E_i)^2} \qquad (2)$$

All network modeling and analysis software were developed in our laboratories using the Labview (National Instruments, Austin, Texas, USA) and Matlab (The MathWorks, Natick, Massachusetts, USA) programming environments. For a review of network modeling see [3], [4])
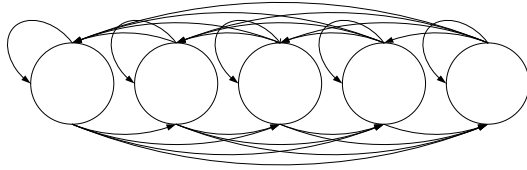

Fig. 1. Illustration of recurrent neural network architecture

B. Analysis: Categorization of Network Dynamics via Close Returns, Derivatives and Lyapunov Exponents

The activity of networks was categorized using a variation of the close return algorithm [7], [14], [16], in which we compared the last state of all units to the history of all their previous states. In *fixed-point* networks all activation settled to a constant value. *Periodic limit cycle networks* returned perfectly to the same value, where the number of iterations between repeats was the period of the network. Networks in which the activities did not repeat perfectly but nevertheless simultaneously returned within 1% of the total range of activation (i.e., $e = 0.01$) were said to have exhibited a *close return*. All other networks were classified as *other* or *turbulent*. In the intermittent case we further categorized various phases by an additional two methods:

(i) We examined the mean in the absolute change in activity for all units, where change was defined as:

$$dS_j = \left(0.5\left|\left(S_{j(t)} - S_{j(t-1)}\right)\right| + 0.5\left|\left(S_{j(t)} - S_{j(t+1)}\right)\right|\right) * dt \quad (3)$$

(ii) We assessed the Lyapunov Exponents of the activity history [5], [8] - [11], [18]. The novelty of our analytic approach is that we computed the spectrum of Lyapunov exponents associated with each individual point of the system's phase trajectory. Most other researchers assume that the attractor of the system is homogeneous, with the same geometry in all of its points, and therefore an average value of the maximal Lyapunov exponent represents the system's stability. In the case of intermittent dynamics such an assumption would be erroneous. As our system has two (in general there can be more) distinct phases, there is no reason to expect self-similarity in the system's attractor. Our algorithm for computing local Lyapunov exponents in the phase space of the system can be summarized by the formulas:

$$\{L_i(t)\} = \log(eig(F_{ij}(t) + F_{ji}(t))/2);$$
$$F_{ij} = U_{ik}(t) H_{kj}^{-1}(t);$$
$$U_{ij}(t) = \left\langle (S_i(t+\boldsymbol{t}) - S_i(t'+\boldsymbol{t}))(S_j(t) - S_j(t'))\right\rangle_{t'}; \quad (4)$$
$$H_{ij}(t) = \left\langle (S_i(t) - S_i(t'))(S_j(t) - S_j(t'))\right\rangle_{t'};$$
$$t': |S(t) - S(t')| < \boldsymbol{e}$$

Where *eig* means the set of eigenvalues of the corresponding matrix, ‖ is the Euclidian norm of a vector and $\boldsymbol{e}$ and $\boldsymbol{t}$ are constants selected as 0.1 and 1 for this application.

III. RESULTS

A. Tracking down biologically relevant dynamics in a population of random networks

We began by assessing the activation distribution in a population of 1000 networks with random weights (set at -3 to 3) and random initial conditions (0 to 1). Each network was run 1000 iterations and assessed using the close return categorization procedure described in the methods. We began by eliminated networks with fixed point or periodic limit cycle dynamics (46.6 ± 2.08%; [SE%=SQRT((p*(100-p)/n))]). We then explored the remaining networks for dynamics that might be akin to biological systems (a detailed description of the outcome of the categorization of random networks was presented in [17] and is being prepared for publication). In this paper we consider in detail a subcategory of networks with particularly salient dynamics, specifically, those that displayed *intermittency*. Table I in the Appendix provides the weight matrix for the construction of such a network.
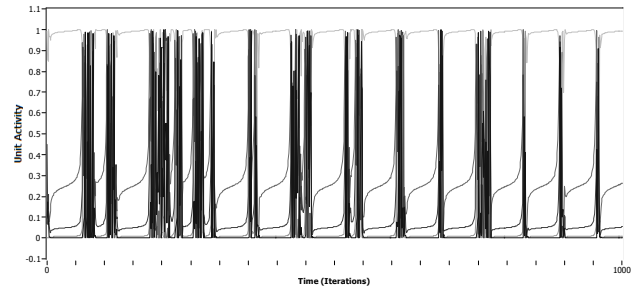

Fig. 2. Superimposed activity traces for 5 units in a recurrent neural network showing intermittent activity.

B. Intermittency dynamics in a recurrent neural network

Figure 2 shows an example of the activity trace from a network with intermittent dynamics. Intermittent network activity was characterized by the presence of calm *laminar* periods of low activity interspersed with *turbulent* phases in which activity vacillated considerably in all units.

C. Statistical properties of intermittency dynamics in a recurrent neural network

This dramatic network-level switch between stable activity and states of hyper-excited oscillations was highly evocative of recurrent paroxysmal activity seen in certain forms of epilepsy. It was also reminiscent of Intermittency Type I, a category of behavior exhibited by some dynamical systems [18], [21], [23]. One feature of Intermittency Type I is its signature distribution of durations for turbulent and laminar phases. To examine the distributions we divided the activity into the two phases using both (i) the mean of unit activity changes and (ii) by assessing the Lyapunov exponent for all points in the recording. The first method was a direct measure of the local change in activity levels. The Lyapunov exponent measured the stability of the system and tendency of the system trajectories to diverge or converge at each point in time.

Figure 3 is a short segment showing the transition of activity traces from the laminar phase to turbulence. Also shown is the number of Lyapunov exponents greater than zero at each iteration. We see that in the beginning of the turbulent phase the number of positive exponents rises from 1 (the laminar flow) to 2 and higher values. We note that Lyapunov exponents are suitable indicators for early stages of the transition into the turbulent phase.

Figure 4a is the distribution of time lengths of the turbulent phases in a network run for 320,000 iterations that included 4730 phase changes. The turbulent distribution had a sparsely populated long tail, with the longest turbulent event recorded being 132 time steps. Laminar events (Figure 4b, 4c) had two peaks, one at the short durations and the other at long intervals. The distribution had a steep cutoff such that a laminar duration never exceeded 62 iterations. There was broad agreement between the two methods of dividing the activity phases. The division also agreed well with the visual characteristics. We also examined the relation between the duration of laminar and turbulent epochs. Figure 5 (next page) is a scatter diagram looking at the duration of each turbulent event (y) plotted against the preceding laminar duration (x).
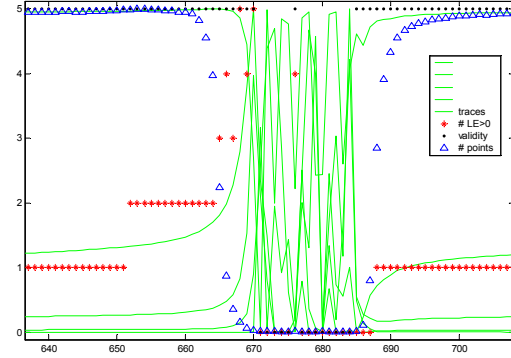

Fig. 3. Transition to turbulent phase and the number of positive Lyapunov exponents (stars). Activity traces (S_i*5) are solid lines. The number of phase-space points (normalized to 5*N/max(N)) used for computation at each time are charted in triangles. Black dots occupy values of either 5 for valid calculation or 0 in cases that the H matrix (see Methods) is degenerate.

D. Sensitivity to external stimulus

Depending on the direction and amplitude of an external stimulus the input could either prolong the length of the laminar phase or hasten the return to turbulence. Given that the activity was constantly progressing along a flow when in the laminar phase, the system was both responsive to *continuous* and *periodic* stimuli. For example, a periodic input of 0.01 presented to unit #2 every 10th iteration reduced laminar durations. Increasing the frequency of the external stimulus increasingly *shortened* the laminar phase. Nonetheless, certain return to the laminar phase was maintained even with constant input of this amplitude. A similar input to unit 0 however had the reverse effect. An input every 10th iteration *prolonged* the laminar event. As frequency of stimulation was increased the stimulus eventually eliminated the turbulent period altogether.

E. Robustness against noise

Noise increased the turbulence within the laminar periods, thereby interfering with the visual and automated differentiation of phases from turbulent epochs. The network's intermittent activity structure however was preserved in the face of continuous injection of additive noise to all units, at amplitudes exceeding 10% of the activation range (uniform white noise with ±0.1 amplitude). Although the introduction of noise shortened the duration of the laminar phase, no matter how much noise was injected the system returned to its original behavior once the noise was removed.
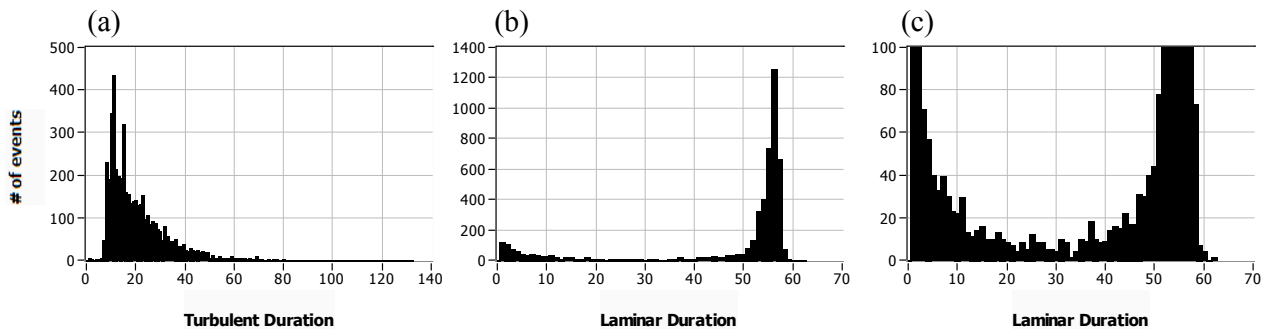

Fig 4. Event duration histograms (a) Distribution of turbulent durations (b) Distribution of laminar inter-turbulent durations (c) Zoom of laminar events histogram highlighting the U-shaped features of Intermittency Type I distributions.
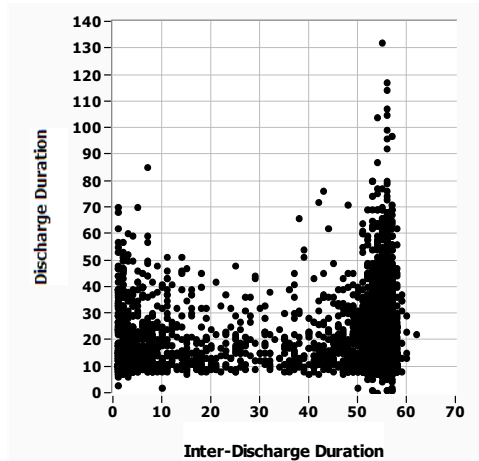
Fig. 5. Scatter diagram of successive event durations. Y-axis is duration of turbulent epoch. X-axis indicates duration of preceding laminar event.

IV. DISCUSSION

1. Types of Epilepsy and a Potentially New Category:

Following Lopes da Silva et al [12], [13], [24], the introduction outlined three general routes to epilepsy. Here we demonstrate the possibility of a fourth, autonomous, route. The recurrent neural network model presented here exhibited transitions between dynamical phases resembling the changes between ictal and interictal patterns seen in biological systems. By using either the mean of the changes of these activations or their Lyapunov exponents we were able to differentiate the observed dynamics into two well-defined phases that matched the observed turbulent and laminar patterns.

2. Intermittency as a Mechanism for Autonomous Transitions in Epilepsy and Behavior

Critical transitions in dynamical states are observed in pathological cases (epilepsy) but are also an important feature of healthy cognitive processes including attention. These findings show that even a simple recurrent neural network can autonomously generate the transitions between dynamics required for this type of intermittent behavior. The ability to autonomously switch between phases without intervention is important for understanding how dynamical transitions take place in biological systems. The possibility that such intermittency mechanisms could be instantiated in the brain suggests: (i) that there is a viable alternative to the conception of deformation theories and separate attractor scenarios and (ii) analysis techniques that attempt to characterize behavior over extended intervals without regard to the multiplicity of phases could fail to properly capture the underlying dynamics.

3. Autonomy from the Environment

The transitional features of this system are important in accounting for two basic and competing requirements of a dynamical system that must interact with the environment while remaining independent: (i) the system has to be responsive to input, and (ii) it has to avoid having its dynamics

locked following exposure to the environmental input. A dramatic example of how neural ensembles can be locked by environmental perturbations is the phenomenon of photosensitive epilepsy in which epileptic discharges are induced by a visual stimulus. In these cases a trigger from the environment alters the brain's dynamics and once triggered, the discharge can persist long after the stimulus is removed. In contrast, in the normal perceptual case, a stimulus will draw the system's attention but the healthy organism is able to react in a time delimited manner so that it can attend to other events and avoid being locked into the single attentional state. The collective behavior of the recurrent neural network units allowed for just this sort of autonomous transition between dynamical phases and consequent decoupling from the environment.

4. Heterogeneous Dynamics and their Topology

The manner in which these transitions take place is perhaps best understood by considering the heterogeneous topology of the trajectory. Figure 6 is a reconstruction of the system. It shows the trajectory of the network embedded in a 3-dimensional space over 10,000 iterations by plotting the activity of 3 units, each on an axis. The activity plot clearly shows the structure of the laminar flow as a narrow curved tube passing through a 5-dimensional hypercube of which 3 dimensions are shown. When the activity is in the vicinity of the entrance to the flow it is attracted inward and proceeds through the curved structure until it exits at a lower point along the flow. The dashed arrow indicates the direction of the flow. Upon exiting, the turbulent activity commences. The turbulent activity is similar to billiard ball type chaos bouncing in a 5-dimensional hypercube [1] until it once again finds the entrance to the tube. The activity thus does not have the homogenous properties associated with most dynamical systems. The move from the flow to a chaotic state does not require an outside perturbation or deformation of the system.
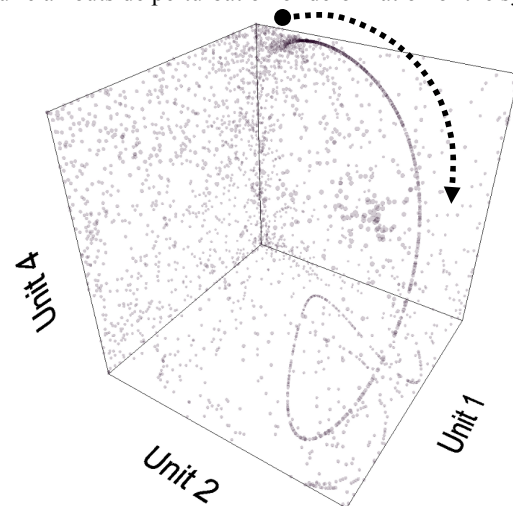

Fig. 6. Topology of intermittent neural network activities. The state of units 1,2 and 4 indicate position in 3-dimensional activation space.

5. Distribution and Predictability of Durations in Intermittent Network Activity

The maximal length of the laminar phase is clearly illustrated in the Figure 4 distributions histogram. Knowing this limit and the shape of the distribution has predictive value. As there is a buildup-like phenomena taking place, the duration to the next turbulent event is statistically predictable and bounded. With each time step from the past event there is a known probability that precisely follows the laminar distribution. In the interval immediately following an exit from discharge the probability of entry back into a turbulent event is relatively high but then drops as the system makes it to the mid-section of the distribution. As time progresses the discharge becomes inevitable with an interval between such events never exceeding 62 iterations.

These predictable properties follow directly from the topology. The laminar flow's finite length reflects the fact that there is a maximal duration for the laminar phase. Unlike multi-stable systems and systems undergoing deformations, prediction in this case is not entirely dependent on an external, possibly random or unknowable factor. The left segment in the distribution illustrated in Figure 4 marks the proclivity to reenter a turbulent phase immediately upon exit, the right-hand rise marks the maximal inter-paroxysmal intermission period. In other words, possible evidence for an intermittency mechanism of this sort could be obtained by searching epileptiform recordings for inter-paroxysmal duration distribution similar to the one described.

Figure 5 indicates that the longest turbulent periods (>90 iterations) might only follow the longest laminar (>30 periods). However, due to the rarity of long turbulent events (only 10 such occurrences out of 4730 events) and the bias toward long laminars it is difficult to establish the significance of this correlation even given the large sample size (320,000 iterations). The relation could be a peculiarity of this particular data set and would certainly be difficult to recognize in a biological counterpart, where it might take years to reliably accumulate a comparable amount of data.

As of late, there has been much debate as to whether seizures can be predicted. This model suggests that it may very much depend on the mechanisms underlying the progression. If certain seizures turn out to follow a Type I intermittency distribution then there is hope for distinct kinds of prediction such as time to upcoming seizure. Although the relationship between the laminar epoch and the length of the following turbulent phase were weak in this particular case, it is also conceivable that a situation exists with a strong coupling between two phases. Specifically, the topology of the laminar could include a well-defined short trajectory with a return exit close to the entrance to the flow. If such an intermittent topology were found it could shed light on two phenomena: (i) in the realm of epilepsy research, the idea that short ictal events actually have a protective effect by averting long turbulent events or full-scale seizure; (ii) in attention studies, to help explain the relation between the delay following stimulus presentation and the attentional readiness of a system. If the time past since the last reaction (here conceived of as a turbulent event) is short, then the system will only repeat that state for a short time. However, if there has been a lengthy delay (long laminar duration) since the last presentation the system's reaction it is more likely to be extended. These behaviors could also be changed with the modification of the topology via learning. The specifics would depend on the topology, and that such a structure could exist, remains to be shown.

6. Sensitivity to stimulus frequency and intensity

The proportional effects of the stimulus seen in this study are simply related to their ability to interfere with the flow along the laminar. The stronger, longer or more frequent the intervention, the better chance it has of pushing the system out of the laminar phase. A stimulus that hastens the trajectory through the flow or is orthogonal to the flow will precipitate and aggravate the turbulent epoch. Conversely, a stimulus that biases the activity in the direction of the entrance to the flow will delay and might even stop the turbulence.

7. Ameliorating Seizures

The topology of the system and the attendant ability of the model to respond differentially to a stimulus depending on its intensity, frequency and duration helps elucidate how stimulation might delay, stop or even hasten a seizure in an intermittency-based system. If we can reconstruct the properties of the laminar phase it might become possible to perturb the system so as to increase its likelihood of remaining out of the paroxysmal state.

8. Modulation of attention

From an attention perspective these properties demonstrate that intermittency allows for a paradigm in which stimulus intensity or frequency can proportionally affect response. These are important features for attention systems in which the responses should be graded in relation to the salience, duration and amplitude of the relevant environmental or intrinsic triggers (whether they signal danger or opportunity). Although the possibility still exists that at extremes stimuli will saturate the activity and the system might remain locked in one of the two phases indefinitely, the experiment in which we applied noise to the network demonstrated that the system can still recuperate and return to the initial intermittent behavior once the extreme stimulus is removed.

9. Attention and context

The observation that these networks can respond commensurably to stimulus intensity, duration and frequency is also an important feature for systems in which processing occurs in parallel. These properties allow for neural elements from a distal location in the brain to have a graded influence on the dynamics of local processing without the local processing being locked down by the incoming dynamics. Such distributed modulation is critical both in allowing for context-sensitive attentional processing and avoiding lock-ups due to synchronization.

10. Endogenous Response Variability

The fact that the system displays a sensitivity to initial conditions might explain some of the difficulties of predicting seizures. However, in terms of behavioral reactions a certain amount of unpredictability might actually provide an organism with an advantageous response. The endogenous variability guarantees that the system will respond in a proportionally beneficial manner without the liabilities of excessive repetition and over-predictability.

11. Attention Recovery, the Homunculus Problem and Autonomy

Cognitive models often conceive of the tracking and release of attention as being moderated by an external trigger. Limiting the factors that can change attention to perturbation events is neither conceptually helpful nor practically achievable. A central difficulty with the neuroscience literature on attention, in particular, is that it tends to avoid addressing the question of how a local ensemble's dynamics might shift independently of input from a stimulus or another brain structure (for a recent review of the attention literature see [22]). Models usually focus on an already partially preprocessed signal or the addition of an external cue. That is, the response starts and stops in direct relation to the activity of another brain location or the stimulus. The problem is thus only pushed back leading to an infinite regress. This attribution of the attention mechanism to an as-yet-to-be-identified location is a traditional homunculus fallacy (see Dennett's Cartesian theatre [2]).

The analogous error in dynamical research is to pre-assign to the incoming sensory information the dynamics that will drive the system, such that the workings of the target neural ensemble are, at best, a simple modulation of this signal. This approach leaves the dynamical processing and decoupling questions forever unanswered. How are different dynamical regimes initiated autonomously in preparation for the signal? How does the location avoid lock-up once the signal is present? The intermittency model can begin to account for attention recovery by building the dynamical transitions into the collective behavior rather than appealing to external salvation.

12. Recurrent Network Attention Does Not Require Plasticity thereby Ensuring Rapidity of Response

An extremely important, possibly counter-intuitive element to the intermittency model is that it does not require plasticity. All transitions in the activity of our model occurred in the absence of changes to the network structure. This may be counter-intuitive given that neuroscientists mostly conceive of a change in behavior as being underwritten by an alteration in either the intrinsic cell properties or the structural features of the network; what else could explain the fact that an identical stimulus could affect the systems differently in two successive exposures? As we have shown, the system is indeed in a different state, but this change in state is not due to an alteration in its units' intrinsic properties or the network's connectivity properties. Change in activity alone alters the system's location in activity space, and thus is sufficient to generate a differentiated response. Moreover, assuming that most plastic events are slower than activity, the intermittency model offers an exceedingly quick activity-based method for responding to a stimulus such that the network need not wait for alterations to its connectivity or intrinsic cell properties to modulate its output. This opens up the possibility for laminar flow geometry to act as a preprogrammed recognition or motor response that can be triggered by an outside event while still ensuring a finite response.

13. Plasticity May Require Intermittency

Although an intermittency mechanism may not require plasticity there is good reason to believe that the same dynamical problems that affect activity may also affect synaptic dynamics. As different as they may seem, the same dynamical issues apply both to non-plastic and plastic systems. (The one difference being that for systems of comparable size the plastic system has considerably more degrees of freedom.) Given that the need to avoid lock-ups extends to plasticity, intermittency may offer an autonomous means for transition between learning phases and safeguard against runaway synaptic dynamics. Conversely, there needs to be an autonomous mechanism that can periodically initiate plasticity even after it has settled. In general, all the advantages of intermittency as applied thus far to activity can be extended to plasticity. The extension of intermittency principles to plasticity may prove to be the most interesting application of all, suggesting new ways of looking at computational system learning and the dynamics of long-term representation [6], [20].

14. Clinical Implications: Prediction, Amelioration and Reversal of Susceptibility to Paroxysms

It is important to note that there are other types of intermittency that have very different distribution and features from the Type I intermittency examined thus far. In particular Type III intermittency has been implicated in epilepsy [19]. The topological properties of the laminar flow studied here imply that if intermittency-based mechanisms are responsible for certain epileptiform phenomena then aspects of the dynamical behavior could in principle be predicted. The fact that intermittent systems are responsive to stimuli in a way that is proportional to frequency and intensity without necessarily being locked by the input dynamics supports the idea that we may be able to go beyond prediction and into the realm of amelioration by stimulating paroxysmal ensembles electrically or chemically. Most interesting perhaps is the idea that we may actually be able to reshape the topology of the system. That is, even small changes in the weights, or connectivity, of a network might directly affect the system's response characteristics without requiring continuous stimulation. This possibility, though highly speculative, offers at least a conceptual route to the permanent reversal of the epileptic condition in the biological system.

15. Future Directions, Provisos and Conclusions:

The model in this paper offers a new way to explain some of the dynamical transitions seen in epilepsy. We also explain why these autonomous transitions might be important in elucidating and modeling the dynamics of healthy phenomena such as attention. There are many other biological phenomena which are characterized by largely autonomous dynamical shifts. Sleep is perhaps the most interesting candidate for exploration given its direct relation to attention. The presence of these dynamics now needs to be searched for in epileptiform data. Similarly, the applicability of intermittency to attention requires examination of healthy data and modeling in cybernetic systems. We need to study various forms of plasticity and see what role intermittency might play in long-term learning. It will also be important to establish the effects of a range of additive and multiplicative noise factors on the distribution of laminar and turbulent epochs. The impact on general network theory should also be considered (for a recent review of the subject see [15]).

Caution however must be taken in interpreting the biological correlates of the laminar and turbulent phases and their distribution. Though conceived here as the quiescent period, the laminar phase could represent a synchronous epileptiform state. With respect to attention, either phase could represent the waiting state or the execution of a preset response. Whatever the physical interpretation might be, it is the ability of intermittency to explain autonomous switching between dynamical regimes that recommends it for further investigation as a mechanism for elucidating the dynamical features of epilepsy and attention recovery.

REFERENCES

[1] L. A. Bunimovich, and J. Rehacek, "Nowhere dispersing 3D billiards with non-vanishing Lyapunov exponents," *Commun. Math. Phys.* vol. 189, pp. 729–757, 1997.

[2] D. Dennett, *Consciousness Explained.* Boston: Little, Brown and Company, 1991.

[3] B. Ermentrout, "Phase-plane analysis of neural activity," in *The Handbook of Brain Theory and Neural Networks*, M. A. Arbib, Ed. Cambridge, MA: MIT Press, 1995, pp. 339-343.

[4] J. Hertz, A. Krogh, and R. G. Palmer, *Introduction to the Theory of Neural Computation,* Redwood City, CA: Addison-Wesley, 1991.

[5] L.D. Iasemidis, D. S. Shiau, W. Chaovalitwongse, J.C. Sackellares, P. M. Pardalos, J.C. Principe, P.R. Carney, A. Prasad, B. Veeramani, and K. Tsakalis, "Adaptive epileptic seizure prediction system," IEEE Trans Biomed Eng. vol. 50:5, pp. 616-27, May 2003.

[6] S. Kalitzin., B. W. van Dijk, and H. Spekreijse, "Self-organized dynamics in plastic neural networks: bistability and coherence," *Biological Cybernetics*, vol. 83, pp 139-150, 2000.

[7] H. C. Kwan, E. L. Ohayon, P. A. Hwang, and W. M. Burnham, "Identification of close returns in human epileptiform EEG," *Society for Neuroscience 2002 Abstracts,* Program No. 793.1. Washington, DC: SFN, 2002.

[8] Y.C. Lai, M. A. Harrison, M.G. Frei, and I. Osorio, "Inability of Lyapunov exponents to predict epileptic seizures," Phys Rev Lett., vol. 8, pp. 91-96, 2003.

[9] K. Lehnertz, "Non-linear time series analysis of intracranial EEG recordings in patients with epilepsy--an overview," Int J Psychophysiol. vol. 34:1, pp. 45-52, Oct 1999.

[10] K. Lehnertz, F. Mormann, T. Kreuz, R. G. Andrzejak, C. Rieke, P. David, C. E. Elger, "Seizure prediction by nonlinear EEG analysis," IEEE Eng Med Biol Mag. vol. 22:1, pp. 57-63, 2003.

[11] J. Lian, J. Shuai, P. Hahn, and D. M. Durand, "Nonlinear dynamic properties of low calcium-induced epileptiform activity," Brain Res. vol 2:890, pp. 246-54, Feb 2001.

[12] F. H. Lopes da Silva, W. Blanes, S. N. Kalitzin, J. Parra, P. Suffczynski, D. N. Velis, "Dynamical diseases of brain systems: different routes to epilepsy," *IEEE, TBME*, vol. 50:5, pp. 540-549, 2003.

[13] F. H. Lopes da Silva, W. Blanes, S. N. Kalitzin, J. Parra , P. Suffczynski, D. N. Velis "Epilepsies as dynamical diseases of brain systems: basic models of the transition between normal and epileptic activity," Epilepsia, vol. 44 Suppl 12, pp. 72-83, 2003

[14] G. B. Mindlin, and R. Gilmore, "Topological analysis and synthesis of chaotic time series," *Physica D*, vol. 58, pp. 229-242, 1992.

[15] M. E. J. Newman, "The Structure and Function of Complex Networks," SIAM Review vol. 45:2, pp 167–256, 2003.

[16] E. L. Ohayon, R. Liu, W. M. Burnham, and H. C. Kwan "Characterization of human EEG by close return," *Society for Neuroscience 2001 Abstracts,* Program No. 83.9 Washington, DC: SFN, 2001.

[17] E. L. Ohayon, S. Kalitzin, P. Suffczynski, F.Y. Jin, P. W. Tsang, S. Borrett, W. M. Burnham, and H. C. Kwan, "Charting Epilepsy by Searching for Intelligence in Network Space with the Help of Evolving Autonomous Agents," *2003 Ladislav Tauc Conference in Neurobiology. Decoding and Interfacing the Brain: From Neuronal Assemblies to Cyborgs.* Gif-sur-Yvette Cedex, France.

[18] H.O. Peitgen, H. Juergens, and D. Saupe, *Chaos and Fractals: New Frontiers of Science.* New York: Springer-Verlag, 1992.

[19] J.L. Perez Velazquez, H. Khosravani, A. Lozano, B. L. Bardakjian, P. L. Carlen, and R. Wennberg, "Type III Intermittency in Human Partial Epilepsy," *Euro. J. Neurosci.*, vol. 11, pp. 2571-2576, 1999.

[20] A. D. Priesol, D. S. Borrett, and H. C. Kwan, *Dynamics of a Chaotic Neural Network in Response to a Sustained Stimulus.* Technical Report # RBCV-TR-91-38, University of Toronto, Toronto. 1991.

[21] H. G. Schuster, *Deterministic Chaos: An Introduction,* 2nd ed. New York: John Wiley & Sons, 1988.

[22] S. Shipp, "The brain circuitry of attention," *Trends in Cognitive Sciences,* vol 8:5, pp. 223-230, May 2004.

[23] S. H. Strogatz, *Nonlinear dynamics and chaos: With applications to physics, biology, chemistry, and engineering.* Cambridge, MA: Perseus Books, 1994.

[24] P. Suffczynski, S. Kalitzin and F. H. Lopes da Silva, "Dynamics of non-convulsive epileptic phenomena modeled by a bistable neuronal network," *Neuroscience*, vol. 126:2, pp. 467-484, 2004.

[25] X. Wang, and E. K. Blum, "Dynamics and bifurcation of neural networks," in The Handbook of Brain Theory and Neural Networks, M. A. Arbib, Ed. Cambridge, MA: MIT Press, 1995, pp. 339-343.

APPENDIX

TABLE 1
WEIGHT MATRIX FOR AN INTERMITTENT NETWORK

To unit	From Unit				
	0	1	2	3	4
0	1.418	2.750	-1.829	1.878	-2.537
1	-2.225	-1.086	-1.707	2.509	2.212
2	-1.854	-0.038	2.134	-1.737	-1.708
3	1.322	2.718	-2.349	-2.491	2.703
4	-2.404	1.902	0.603	0.962	-0.347