

Aperiodic Dynamics and the Self-Organization of Cognitive Maps in Autonomous Agents

Derek Harter

Department of Computer Science and Information Systems
Texas A&M University - Commerce
Commerce, TX 75429
dharter@tamu-commerce.edu

Robert Kozma

Department of Mathematics and Computer Science
University of Memphis
Memphis, TN 38152
rkozma@memphis.edu

Abstract

When we look at the dynamics produced by biological neuronal populations, we are immediately struck by the fact that aperiodic, chaotic like dynamics, appear to be the normal operating state of such systems. Recent work has shown that such aperiodic dynamics, at least in perceptual systems, may not only be the result of random perturbations experienced by the system from external stimulation, but that the brain itself generates aperiodic dynamics in order to deal more flexibly and reliably with noisy environmental stimulation. Complex systems concepts are helping us to understand the properties of nonlinear systems that are fundamental for the emergence of complex spatio-temporal patterns in natural and biological systems. Advances in neuroscience and computational neurodynamics are applying these concepts of self-organization to understanding the spatio-temporal patterns observed in biological brains. In this paper we introduce a neural population model that is capable of replicating the generation of these types of aperiodic dynamics observed in biological brains. We use the model to self-organize cognitive maps in an autonomous agent through the agents interaction with its environment. We show how such high-dimensional spatio-temporal dynamics may be shaped by environmental input and learning to form chaotic attractors that come to represent 'meanings' for the agent. We discuss how the internal generation of such aperiodic dynamics may aid

in the formation and recognition of such noisy environmental stimuli in biological organisms in general and in our simulated agents specifically.

Keywords

Self-Organization; Nonlinear Dynamics; Cognitive Maps

February 22, 2005

~5100 words; 25 refs; 7 figures;

1 Introduction

Some researchers in dynamical cognition and neurodynamics have speculated on the possibilities that aperiodic, chaotic-like dynamics may play in the role of adaptive behavior [1, 2, 3, 4, 5]. Chaotic dynamics have been observed in the formation of perceptual states of the olfactory sense in rabbits [1]. Mathematical theories of the nonconvergent neurodynamics of perception and decision making have been proposed based on the principles of the olfactory neurodynamics [6, 7]. Other researchers have analyzed activity patterns of primate and human cortex and reported on the dynamics of large-scale neural organization [8, 9].

Skarda and Freeman [1] have speculated that chaos may play a fundamental role in the formation of perceptual 'meanings', that is in learning patterns that have significance to the organism, whether harmful or beneficial. Chaos provides the right blend of stability and flexibility needed by the system, with swift and robust transitions from one cognitive state to the other using first order phase transitions. According to Skarda and Freeman, the normal background activity of neural systems is a chaotic state. In the perceptual systems, input from the sensors as well as internal feedback perturbs the neuronal ensembles from the chaotic background. The result is that the system transitions into a new attractor that represents the 'meaning' of the sensory input to the organism, given the context of the state of the organism and its environment.

The normal chaotic background state is not like noise. Noise cannot be easily stopped and started, whereas chaos can essentially switch immediately from one attractor to another. Also, aperiodic spatio-temporal dynamics may play an important role in the rapid learning abilities of biological organisms. The chaotic background state essentially offers many choices of attractors that can be rapidly shaped by small synaptic changes to form relevant attractor basins. This type of dynamics may be a key property in the flexible production of behavior in biological organisms

[2].

The study of nonlinear dynamics has expanded in all areas of science in the past decades for many reasons. Nonlinear dynamics provide new conceptual and theoretical tools that allow us to understand and examine complex phenomena that we have never been able to model before. Nonlinear dynamics show up everywhere, in physical systems such as electrical circuits, lasers, optical and chemical systems. Such dynamics are especially prevalent in the biological world, from fractal growth patterns in biological development to the self-organizing characteristics of population models, and the importance in regulating healthy biological rhythms such as the beating of the heart.

Nonlinear systems in critical states have many interesting properties. Phenomenon such as stochastic and chaotic resonance are known which enable such systems to detect the presence of signals much better in noisy environments than linear systems are capable of doing [5]. Their greatest interest lies, however, in their fundamental relationship to self-organization and emergence of complex patterns and behaviors in complex environments. Aperiodic dynamics are both an indication of and a mechanism for the emergence of such self-organizing properties.

Insights in nonlinear systems theory are beginning to be applied to understanding the dynamics of the brains, and how such processes produce cognition [2, 10, 11]. Aperiodic dynamics are believed to play a fundamental role in the mechanisms for the self-organization of meaning in mammalian perceptual systems [1, 12]. Neurological evidence has shown that perceptual meanings (of recognized smells) are created through the formation and dissolution of chaotic attractors in the olfactory bulb. We will discuss this example of the self-organization of a perceptual pattern of meaning. We use this type of organization in aperiodic systems to model the formation of cognitive maps in the hippocampus of biological organisms.

2 K-Sets: A Neurodynamical Population Model of Brain Dynamics

2.1 Aperiodic Dynamics in Olfactory Systems

In their influential paper, Skarda and Freeman argued that chaos, as an emergent property of intrinsically unstable neural masses, is very important to brain dynamics [1]. In experiments carried out on the olfactory system of trained rabbits, Freeman was able to demonstrate the presence of chaotic dynamics in EEG recordings and mathematical models. In these experiments, Freeman and associates conditioned rabbits to recognize smells, and to respond with particular behaviors for particular smells (e.g. to lick or chew). They performed EEG recordings of the activity in the olfactory bulb, before and after training for the smells.

The EEG recordings revealed that in fact chaotic dynamics, as shown by the observed strange attractors, represented the normal state when the animal was attentive, in the absence of a stimulus. These patterns underwent a dramatic (nonlinear) transition when a familiar stimulus was presented and the animal displayed recognition of a previously stored memory, through a behavioral response. The pattern of activity changed, very rapidly, in response to the stimulus in both space and time. The new dynamical pattern was much more regular and ordered, very much like a limit cycle, though still chaotic of a low dimensional order. The spatial pattern of this activity represented a well defined structure that was unique for each type of odor that was perceptually significant to the animal (e.g. conditioned to recognize). Figure 1 shows an example of such a recorded pattern after recognition of a stimuli of the EEG signals and the associated contour map. In this figure after recognition, all of the EEG waves are firing in phase, with a common frequency which Freeman called the carrier wave. The pattern of recognition is encoded in the heights (amplitude modulations) of the individual areas. The amplitude patterns, though regular, are not exact limit cycles and exhibit low dimensional chaos. In other words, different learned stimuli were




Figure 1: EEG carrier wave patterns (left) and contour map (right) of olfactory cortex activity in response to a recognized smell stimulus (from Freeman, 1991, p. 80)

stored as a spatio-temporal pattern of neural activity, and the strange attractor characteristic of the attention state (before recognition) was replaced by a new, more ordered attractor related to the recognition process. Each (strange) attractor was thus shown to be linked to the behavior the system settles into when it is under the influence of a particular familiar input odorant.

Figure 2 shows the effects on the spatial attractor pattern due to learning. Every time a new odor was learned by the animal, all of the existing attractor patterns changed. In this figure the contour pattern of activity for sawdust is shown (before learning the banana odor) on the left, for the newly learned banana odor (center), and then again for sawdust (right, after learning new banana odor). These maps are snapshots of the EEG activation of an 8x8 grid of electrodes implanted in the olfactory bulb. Notice that the spatial pattern for sawdust no longer resembles its previous pattern.

How, then, do we know that the right most pattern does in fact capture some meaning to the animal about the sawdust odor? We only know this through behavioral observation. The animal was trained to produce some behavior when it recognized a smell (such as licking in response to sawdust odor), and rewarded when it displayed the appropriate behavior for the stimuli (by being given food, for example). As external observers of both the behavior of the organism and the internal dynamics of the olfactory region of the animals brain, we only know that the EEG activation pattern captured in the olfactory bulb has some association with the sawdust odor because the




Figure 2: Change in contour maps of olfactory bulb activity with the introduction of a new smell stimulus (from Freeman, 1991, p. 81). Left, snapshot of activation after recognition of sawdust order. Center, activation of same populations after learning banana order. Right, activation of population for same sawdust order after learning new banana order.

animal reliably displays the externally conditioned behavior whenever the EEG pattern is formed.

Another interesting question is, given that the sawdust pattern changes so drastically upon formation of a new attractor in the olfactory bulb, how is it that the behavioral response remains associated with the now deformed attractor dynamics for sawdust? The answer to this question is certainly not completely known yet. The existing attractors are modified when a new attractor is formed, but they are not destroyed. So the dynamics we observe as AM patterns in the EEG signals can look very different, but this only represents a squeezing or stretching of the attractor. Somehow parts of the attractor remain invariant such that other portions of the brain (e.g. those responsible for the generation of behavior in response to the stimuli) recognize and react to them.

Whenever an odor becomes meaningful in some way, changes in the synaptic connections between neurons in different parts of the olfactory cortex take place. Just as in the Hopfield model and other neural networks, these changes are able to create another attractor, and all existing attractors are modified as a result of this learning. However, in real brains, the attractors of perceptual meaning are not simple point attractors, but are specific strange attractors.

Freeman suggests “an act of perception consists of an explosive leap of the dynamic system

from the basin of one (high dimensional, in the attentive state) chaotic attractor to another (low dimensional state of recognition)” [12]. These results suggest that the brain maintains many chaotic attractors, one for each odorant an animal or human being can discriminate. Freeman and Skarda speculate on many reasons why these chaotic dynamics may be advantageous for perceptual categorization. For one, chaotic activity continually produces novel activity patterns which can provide a source of flexibility in the individual. But since chaos is a ordered state, such flexibility is under control. As Kelso remarks [13], such fluctuations continuously probe the system, allowing it to feel its stability and providing opportunities to discover new patterns. Another advantage of chaos is that it allows for very rapid switching between attractors, which random activity is not able to do. Freeman also proposed that such patterns are crucial to the development of nerve cell assemblies.

2.2 K-Set Model of Aperiodic Dynamics

The K-set hierarchy, developed by Freeman and associates [14, 2, 1, 12], is both a model of neural population dynamics and a description of the architectures used by biological brains for various functional purposes. The original purpose of the K-set was to model the dynamics observed in the olfactory perceptual system. The lowest level of the hierarchy, the K0 set, provides a basic unit that models the dynamics of a local population of tens of thousands of neurons. The dynamics of the K0 set are described by a second order ordinary differential equation feeding into an asymmetric sigmoid function:

$$ab \frac{d^2 x(t)}{dt^2} + (a + b) \frac{dx(t)}{dt} + x(t) = f(t) \quad (1)$$

This equation was determined by measuring the electrical responses of isolated neural populations to stimulation and other conditions. The a and b parameters are time constants that were determined through such physiological experiments. $x(t)$ is the pulse density of the modeled neu-

ral population, in other words the average number of neurons that are pulsing in the population at any given point in time. $f(t)$ is a nonlinear asymmetric sigmoid function describing the influence of incoming activation, and is given in equation 2.

$$f(t) = k[1 - \exp(-\frac{e^{x(t)} - 1}{k})] \quad (2)$$

A K0 unit models the dynamics of an isolated neural population. From the basic K0 unit can be built up architectures that capture the observed dynamics of increasingly larger functional brain areas. The KI models excitatory-inhibitory feedback populations. KII models interacting excitatory-inhibitory populations and correspond to organized brain regions such as the olfactory bulb (OB) or the prepyriform cortex (PC). KIII combine 3 or more KII populations to model functional brain areas such as perceptual cortex or hippocampus, and are capable of aperiodic dynamics of the type observed in these regions to, for example, derive meaning from perceptual senses. In the simulations presented in this paper, we use a discretized version of the K-model (described in [15, 16]) developed for use in large-scale autonomous agent simulations.

In the original K model, the purpose of the KIII set was to model the chaotic dynamics observed in rat and rabbit olfactory systems. [17, 18, 19] KII are capable of oscillatory behavior, as described above. When three or more oscillating systems (KII) of different frequencies are connected through positive and negative feedback, the incommensurate frequencies can result in aperiodic dynamics. The dynamics of the KIII are produced in just this manner, by connecting three or more KII units of differing frequencies together. The KIII set was not only capable of producing time series similar to those observed in the olfactory systems under varying conditions of stimulation and arousal, but also of replicating power spectrum distributions characteristics of biological and natural systems in critical states. [20, 21]

The power spectrum is a measure of the power of a particular signal (or time series as for example that obtained from an EEG recording of a biological brain) at varying frequencies. The typical power spectrum of a rat EEG (see Figure 3, top) shows a central peak in the 20-80 Hz range, and a $1/f^\alpha$ form of the slope. The measured slope of the power spectrum varies around $\alpha = -2.0$. $1/f^\alpha$ type power spectra are abundant in nature and are characteristic of critical states, between order and randomness, at which chaotic processes operate. Power spectra of biological brains have been observed to vary from $\alpha = -1.0$ to $\alpha = -3.0$. The atypical part of the experimental EEG spectra is the central peak, indicating stronger oscillatory behavior in the γ frequencies. This central peak in the 20-80 Hz range is known as the γ frequency band, and is associated with cognitive processes in biological brains. The K-models are capable of replicating the power spectra of biological EEG signals, as shown in Figure 3, bottom. [15, 12]

Power spectrum of in vivo EEG recordings have the typical $1/f^\alpha$ slope power spectrum, which looks similar to systems in critical states (except for the peak in the gamma range). Systems in critical states are interesting because they are indicative of self-organization. They are not tuned externally in order to determine a critical point, natural systems have no one tuning the parameters, they organize themselves. They are near critical states in the sense that the collective behavior (of neurons in this case) cannot be extrapolated from behavior of the individual elements. Some driving events have little effect on the system but some driving events cause a collapse. It appears that the neural system self-organizes in such a way that "collapse" in the critical states are precipitated by significant cues. In the case of perception, these are perceptual stimuli, like the critical presence of a small but significant number of some molecules indicative of a particular odorant.

Recently, a new class of chaotic behavior, called chaotic itinerancy, has been introduced




Figure 3: The power spectrum of a rat Olfactory Bulb EEG is simulated with the KA-III model. The calculated “ $1/f$ ” slope of the EEG and model is approximately -2.0 . Rat OB data from (Kay 04), KA power spectrum from (Harter 2004)

[22, 7, 10], which is related to dynamical behavior of K-sets. Chaotic itinerancy is observed in high-dimensional dynamical systems with trajectories evolving through successions of “attractor ruins”, with each attractor being destroyed as soon as it is reached, and the system continuously remains unstable, as in a search mode. Results by the KIII model indicate that the complex, intermittent spatio-temporal oscillations in KIII are possible manifestations of Tsuda’s attractor ruins and chaotic itinerancy in a biologically plausible neural network model [5, 23].

The KIII sets are capable of organizing perceptual categories in the fashion observed in biological perceptual systems. The KIII used as such a pattern classifier is very robust and compares well with more standard methods of pattern classification. [5]

3 Hippocampal Simulation

The formation of new aperiodic attractors appears to be the mechanism by which mesoscopic level population dynamics form representations of external stimuli [12, 2]. Such ‘representations’ are, however, very different in character from traditional notions of static symbols in symbolic systems. For one thing, the recognition process is very much a dynamical process, by which noisy and incomplete perceptions, along with internal states and expectations of the organism, combine into a mechanism that allows for a rapid collapse or shifting of the dynamics when enough evidence has accumulated. Such a recognition can be very sudden and occur over a large area of population. And, when internal expectations bias the system to look for very faint or small perceptual cues, the result can be a very sensitive mechanism that can detect very faint stimuli among very noisy environmental conditions.

The normal baseline background state in perceptual systems appears as a high-dimensional chaotic attractor. When sensory information is received and impinges on the receptors the dynamics of the neuronal populations may change dramatically. When the sensory information is

remembered as being something experienced in the past, the neuronal dynamics fall into a new chaotic attractor that 'represents' the recognition of the stimuli. In biological EEG recordings, attractors after recognized stimuli are more regular and show stronger periodic components, though still chaotic, indicative of a lower-dimensional chaotic state. Failure of the perceptual system to recognize the stimuli results in the attractor remaining in the high-dimensional basal background state. If this unrecognized stimuli coincides with a pain or pleasure signal, a new attractor will be formed to capture the 'meaning' of the new stimuli.

The same basic mechanisms of attractor formation in perception are also believed to be used by brains in other areas to form longer-term memory and behavior producing structures [24]. We use the basic KA-III architecture, described previously, to simulate the formation of cognitive maps in the hippocampus of an autonomous agent. One function of the hippocampus appears to be the formation of long-term representations of the environment, which can be used to navigate to goal locations, remember where food sources and shelter are located, etc. These environmental representations are usually referred to as cognitive maps. In this experiment, we show how aperiodic attractors can be shaped and associated with locations in the agents environment, in a way that we speculate is similar to the processes used by the real hippocampus to form cognitive maps. These attractors have interesting spatio-temporal properties. For example, attractors for locations that are close to one another will be more similar to each other than locations further apart. In this experiment we demonstrate the formation of aperiodic attractors in a simulated organism, and how they can come to be associated with environmental locations. This experiment does not address how such attractors can come to be associated and connected up to meaningful behaviors. Future work is planned to begin to link such attractors to the problems of behavior generation. We will next show how these types of representations can be formed in an autonomous agent, and how they

might be used to perform goal-directed navigation and route planning tasks.

3.1 Experimental Architecture

In this experiment, we used the Webots virtual environment to simulate a Khepera robot moving in and learning an environment [25]. Figure 4 (bottom left) shows the morphology of the Khepera agent. The Khepera robot is a simple agent that contains 8 infra-red and 8 light sensors. The sensors are positioned around the periphery of the body, with most sensors concentrated in the front of the robot. The agent has two independently controlled wheels that allow it to move forward, backward, and turn left or right in place. The environment for this experiment is also shown in Figure 4. In the environment we place 8 light sources, which will be used as salient environmental locations. The light sources can be thought of as locations of food, or some other type of positive environmental feature. The light sources are detectable to the agent at a distance, and the detection range is indicated in Figure 4. In addition to the 8 salient environmental locations, there are 4 landmarks. The landmarks are always detectable to the agent, and it knows the distance and direction to each of the 4 landmarks as part of its sensory information. These landmarks provide a means of localization for the agent, so that it can judge its relative position in the environment.

The architecture of the simulated hippocampus is shown in Figure 5. The portions of the architecture that form the cognitive map of the environment are simulated by a KA-III. These are the CA1, CA2 and CA3 areas, and are based on biological evidence of the structure of the hippocampus. Though the model captures some aspects of the architecture of the hippocampus, many details are omitted in the model. The CA3 and CA1 areas each contain an 8x8 array of KA-II units (for a total of 64 units in each CA region). The CA2 layer has a single KA-II unit. Within the CA3 and CA1 area, the KA-II are connected to one another via lateral connections. Each of the four nearest neighbors are connected to one another (the edges wrap around to technically form a torus out of




Figure 4: Agent morphology (bottom left) and environmental setup for hippocampal simulations. The environment contains landmarks, used as allocentric reference points by the agent, and salient environmental locations, such as food sources. The agent is only able to detect the presence of a food source when it is within a particular range of it.

the area). The e_1 units are connected to the four closest neighboring e_1 's, and similarly the i_1 units are also connected to the four closest neighbors.

There are various projecting connections between the CA areas. For simplicity, the areas in this experiment have been fully connected to one another, and we indicate this using bi-directional arrows in Figure 5. When a layer projects to another layer, each of the e_1 (or sometimes i_1) units projects to multiple units in the other layer. We use a fan-out parameter to determine the level of connectivity. For example, if we have a fan-out of 10 between areas CA1 to CA3, each unit in CA1 is connected to 10 other units in CA3. The projecting connections that fan-out to other areas are chosen at random. So, in the previous example, the 10 projecting fan-out connections would have been chosen at random in the target CA3 layer. The fan-out parameter for projecting connections is chosen so that the areas are fairly sparsely connected, and reflect to some degree the observed projecting connectivity in the biological hippocampus.

Orientation beacons are fed into the hippocampal simulation through the DG region (Figure 5, left). The DG layer contains an 8x8 matrix of KA-0 units. Orientation signals from the 4 landmarks are fed into the DG units. Each of the 4 landmarks has 8 units associated with the direction to the landmark, and 8 units associated with the distance. Directions are broken into 8 cardinal units, North, NorthEast, East, SouthEast, South, SouthWest, West and NorthWest. Units are sensitive to the direction of a particular landmark, though we use a graded response with a normal distribution, instead of a simple winner-takes-all configuration. Similarly there are 8 cardinal distance values VeryClose, Close, MediumClose, Medium, MediumFar, Far, VeryFar, Distant. Again a graded response with normal distribution is applied to the units. The DG area feeds into the CA3 area.

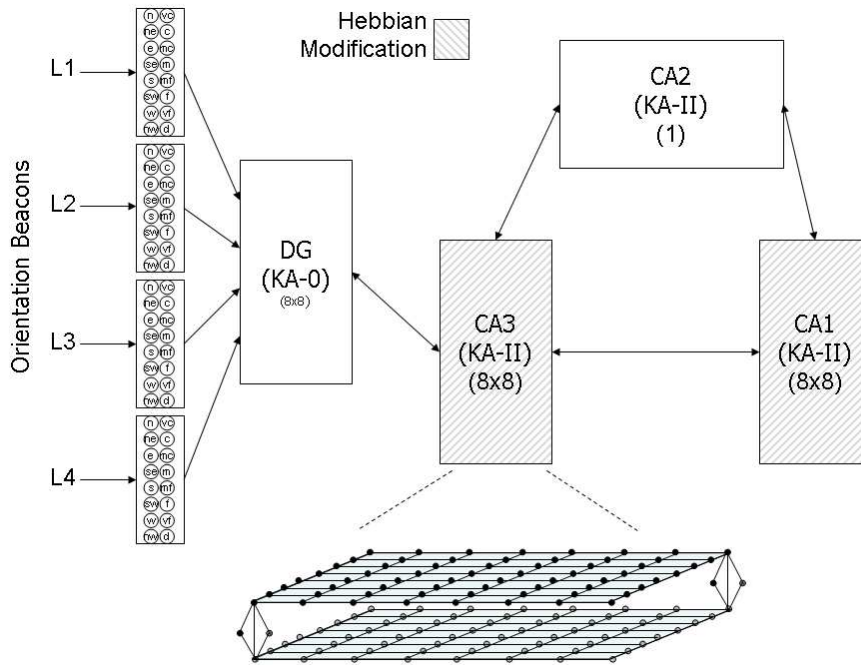


Figure 5: Architecture of KA-III hippocampal simulations. The CA1, CA2 and CA3 areas form a KA-III. CA1 and CA3 are 8x8 matrices of KA-II units. CA2 contains a single KA-II. An example of the CA3 layer is shown below. Each e_1 unit (top sheet) is connected to its four nearest neighbors as is each i_1 unit (bottom sheet). The units at the edges have connections that wrap around, so technically the sheets form tori. The top sheet of excitatory units form KA-II with the bottom sheet of inhibitory units by connecting with the unit immediately above/below along with two other units. We show an example of the KA-II formations in the lower-left and upper-right (only two examples of the actual 64 KA-II formed are shown in the figure). Hebbian modification occurs only among lateral connections in CA1 and CA3. Input from the landmarks feeds into an 8x8 matrix of KA-0 units (DG). Projecting connections between the CA1, CA2 and CA3 are formed by fan-out connectivity, where each unit fans-out and connects with a number of units in the projecting layer. See text for full description.

3.2 Method

Learning is performed in these experiments using Hebbian modification. The lateral connections between units in the CA1 and CA3 areas are the only ones subjected to Hebbian modification. All other weights, including KA-II internal weights and weights projecting between areas, are held constant. All weights are initialized to small random values within some range. The modification of the lateral connections in the CA1 and CA3 areas causes the attractor dynamics to be shaped in response to and by the stimulation from the environmental landmarks. Hebbian modification, however, is not turned on constantly. Hebbian modification only occurs when the agent is within the detection range of an environmental location. Proximity to a salient location causes a type of positive reinforcement signal. Therefore, attractors are only shaped and affected when the robot is in environmentally salient regions.

We use a simple sensory-motor mechanism to cause wandering and exploration behavior in the agent. This exploration behavior is not affected by the cognitive map learning, but we will talk later about how the cognitive maps built by the agent using aperiodic attractors may be used to perform goal-directed navigation. The wander behavior is implemented using KA-0 units and is described in [15]

The agent is allowed to roam in the environment at random for 60 minutes. While moving in the environment, the agent learns and builds a cognitive map. At the end of this learning period, we examine the attractors that have been formed to discover their properties.

We use two types of learning in the simulation, Hebbian modification as we have indicated and habituation. Hebbian modification only occurs when the robot is within a certain range of a light source. When the robot is not within proximity to a light source, no reinforcement signal is produced. During these times habituation of the stimulus occurs. This has the effect of lessening

the response of the simulated hippocampus to unimportant regions in the environment. [5]

The expected effects of this simulation with the two types of learning, Hebbian modification and habituation, is to form 8 distinct attractors (one for each location). The aperiodic attractors will display complex spatio-temporal characteristics. However, they should be identifiable such that when the agent is within proximity to a particular location, the observed attractor that is generated in the CA1 layer is unique. We discuss the results of this simulation next.

3.3 Results

We first examine the spatio-temporal patterns produced in the CA3 region after learning. Figure 6 shows a representation of the spatio-temporal patterns formed in the CA1 region. We measured the activity in the following manner. Around each of the 8 locations in the environment we chose 4 test points at random within the proximity range of the locations (see Figure 6 top left corner). We placed the agent at each of these 4 test points at each of the 8 locations for a half of a second (500 time steps of activity). The distance and direction information is fed in as input to the system and we capture the activity of the e_1 units in the CA3 layer. Figure 6, left bottom, shows an example of a half second of activity captured in the e_1 units in the CA1 layer in response to test point a at location E8.

In order to visualize the spatio-temporal dynamics of our 32 tests, we simplify our representation of the captured time series. Each of the 32 tests consists of 64 time series with 500 values each. We represent the amplitude of each of the 64 units for the half a second of activity by taking the standard deviation of the time series. This collapses the $8 \times 8 \times 500$ array into an $8 \times 8 \times 1$ array of amplitude values. Each measured standard deviation gives us an idea of how active or quiescent each unit was in response to the input from the environmental beacons. We can then plot our 8×8 array of amplitudes as a contour plot, which we have done in Figure 6, right.

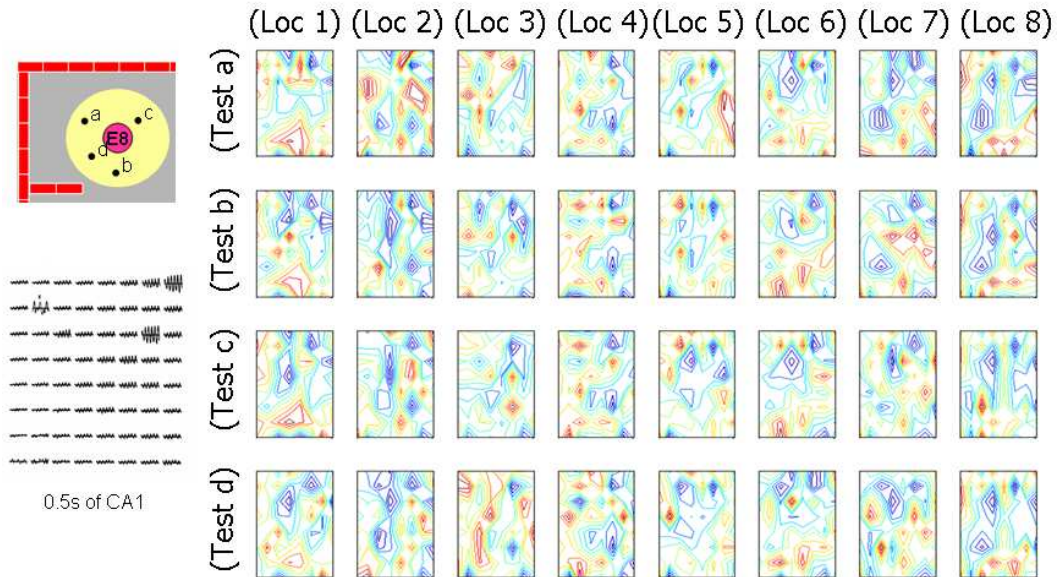


Figure 6: Amplitude Modulation (AM) pattern formation (right) in the CA1 area for the simulation in response to four test points in proximity to each of the eight environmental locations. (Left top) shows an example of 4 test points chosen at random in proximity to location E8. (Left bottom) is an example of a half second of activity of the e_1 units in the CA1 layer in response to test point a at location E8. Amplitudes for each of these 64 time series were calculated simply by using the standard deviation. These 64 measures of amplitude were the values used to produce the contour maps that are shown. See text for full description.




Figure 7: Result of cluster analysis on AM vectors. This figure demonstrates that the spatio-temporal dynamics of the CA1 layer cluster nicely into 8 attractors. The cluster analysis was performed using the 64 vector of the amplitudes for each of the 32 test cases (4 test points at each of the 8 locations). The attractors also form higher-level basins that capture, to some extent, the proximity of the locations experienced in the environment. See text for full description.

Each contour plot gives us an idea of the spatio-temporal dynamics produced in the CA1 region in response to the particular test point at the particular locations. The AM pattern contour plots, therefore, give us an idea of which units are more highly stimulated (higher amplitudes in their activity) in response to which locations. Examining the AM patterns, it can be seen that the spatio-temporal dynamics at each of the 8 locations are more similar to one other than to patterns at the other locations. This may not be immediately obvious but, for example, if you examine the 4 test points at location 1 you will notice a peak of activity at the mid-bottom part of the layer units. Other similarities in the dynamics among points close to a location may also be found when examining the contour maps.

We can show analytically that the spatio-temporal dynamics shaped in response to the environmental locations are more similar to each other than to those in other locations. In Figure 7 we performed a standard cluster analysis of the 32 test patterns. We treated each 8x8 matrix as a 64 dimensional vector of amplitudes, and used euclidean distance to measure the closeness of the patterns to one another. As shown in the figure, each of the 4 test points at a location are closer to each other than to any point in another location. This figure shows that, in fact, the patterns produced within a region are consistently more similar to one another, than those produced in another environmental region. The attractors also form higher-level basins that capture, to some extent, the proximity of the locations to each other experienced in the environment, for example locations 8, 1 and 3 and also 5, 6 and 7.

4 Discussion/Conclusion

The KA-III hippocampal simulation described here forms distinct AM patterns for the 8 salient environmental regions. These patterns are aperiodic spatio-temporal activity in the CA regions. The characteristic activity peaks in the AM patterns are consistent with data known about so called

“place cells” that form in the hippocampus. Place cells are single cells that have been observed to fire selectively for environmental location. Data on place cells have only been recorded for single-cell observation. However, in our population model, we develop characteristic peaks of amplitude selectively for location. We would expect to see similar amplitude patterns in biological hippocampus as developed in our model, and cells that participate in the populations that are highly active for particular locations would indeed look similar to observed place cells [26].

The experiments discussed in this paper use immediate reward signals to shape aperiodic attractors into a kind of cognitive map. It is known that cognitive maps are not, however, really formed in this manner. Cognitive maps can be formed without immediate reward being delivered to the animal. There is a type of value judgment associated with such formation, however. Being in an unfamiliar location heightens awareness of the organism, a type of low-level fear or anxiety response. These altered states may trigger processes as have been described in this paper that guide learning. The response diminishes as the animal becomes familiar with the environment. So a more realistic mechanism would need to use recognition and familiarity of the current location to guide and heighten the fear/anxiety response, which would in turn affect learning of the location.

Learning in the hippocampus of the environmental layout is so effective that it appears to take only one or two experiences to become memorized. This fast or one-shot learning is also typical of another function of the hippocampus, the formation of so-called episodic memories. In memorizing and recalling episodes, our brains are capable of remembering an episode with only a single exposure to the experience. Of course, more emotionally charged experiences tend to form much stronger episodic memories. Future research with these models will aim to take these types of mechanisms more into account.

The next step in this research is to begin to understand how such AM patterns might be used

in the service of goal-directed navigation. It is known that if you measure the onset time of place cells in a biological brain, this time gradually shifts back in phase as the animal moves through the environment [27]. This phase shift of the onset of the place cells may be evidence of the formation of navigation planning in the biological brain. One possible interpretation is that when the animal forms an intention to travel to a goal location, a sequence of AM patterns cycle through the hippocampus. This sequence can be interpreted as sequences of locations the animal intends to visit, from the current one to the next one, etc. in order to reach the goal. As the animal moves through the environment, its idea of the current location changes, and thus this whole sequence shifts back in phase in real-time to represent the next few intended steps the animal is planning to take. For this type of mechanism to be organized, the AM patterns must not simply form in an isolated way, but connections between adjacent locations must be incorporated into the mechanism. If the agent learns which AM patterns are co-located to which others, it may be possible to set up such a mechanism to produce a goal-directed planning for navigating in the environment. These transitional associations may be produced by a hierarchical organization, where different areas form dynamics that associate between the spatio-temporal dynamics of location representations. We are currently examining how such an associational dynamics might be learned and used to perform navigation. The transition between the aperiodic attractors may resemble the chaotic itinerancy phenomenon [10].

The self-organization of spatio-temporal patterns in nonlinear systems are essential to cognitive mechanisms in biological brains. We need to better understand how such mechanisms operate in order to build better models of cognition and smarter autonomous agents. This paper has demonstrated one such self-organizational mechanism for the creation of AM patterns in a cognitive map of an agents environment.

Acknowledgment

This work was supported by NASA Intelligent Systems Research Grant NCC-2-1244. We thank the anonymous reviews whose comments helped to greatly improve this paper.

References

- [1] Christine A. Skarda and Walter J. Freeman. How brains make chaos in order to make sense of the world. *Behavioral and Brain Sciences*, 10:161–195, 1987.
- [2] Walter J. Freeman. *How Brains Make Up Their Minds*. Weidenfeld & Nicolson, London, 1999.
- [3] Walter J. Freeman, Robert Kozma, and Paul J. Werbos. Biocomplexity: Adaptive behavior in complex stochastic dynamical systems. *BioSystems*, 59:109–123, 2000.
- [4] Robert Kozma and Walter J. Freeman. Encoding and recall of noisy data as chaotic spatio-temporal memory patterns in the style of the brains. In *Proceedings of the IEEE/INNS/ENNS International Joint Conference on Neural Networks (IJCNN'00)*, pages 5033–5038, Como, Italy, July 2000.
- [5] Robert Kozma and Walter J. Freeman. Chaotic resonance - methods and applications for robust classification of noisy and variable patterns. *International Journal of Bifurcation and Chaos*, 11(6):1607–1629, 2001.
- [6] H. Liljenstrom. Global effects of fluctuations in neural information processing. *International Journal of Neural Systems*, 4:497–505, 1996.
- [7] I. Tsuda and A. Yamaguchi. Singular-continuous nowhere differentiable attractors in neural systems. *Neural Networks*, 11:927–937, 1998.

- [8] S. L. Bressler and J. A. Scott Kelso. Cortical coordination dynamics and cognition. *Trends in Cognitive Sciences*, 5(1):26–36, 2001.
- [9] H. L. Liang, M. Z. Ding, and S. L. Bressler. Detection of cognitive state transitions by stability changes in event-related cortical field potentials. *Neurocomputing*, 38:1423–1428, 2001.
- [10] Ichiro Tsuda. Towards an interpretation of dynamic neural activity in terms of chaotic dynamical systems. *Behavioral and Brain Sciences*, 24(4):793–847, 2001.
- [11] Walter J. Freeman. The wave packet: An action potential for the 21st century. *Journal of Integrative Neuroscience*, 2003. in press.
- [12] Walter J. Freeman. The physiology of perception. *Scientific American*, 264(2):78–85, 1991.
- [13] J. A. Scott Kelso. *Dynamic Patterns: The Self-organization of Brain and Behavior*. The MIT Press, Cambridge, MA, 1995.
- [14] Walter J. Freeman. *Mass Action in the Nervous System*. Academic Press, New York, NY, 1975.
- [15] Derek Harter and Robert Kozma. Chaotic neurodynamics for autonomous agents. *IEEE Transactions on Neural Networks*, 2005. in press.
- [16] Derek Harter and Robert Kozma. Simulating the principles of chaotic neurodynamics. In *Proceedings of the 6th World Multi-Conference on Systemics, Cybernetics and Informatics (SCI 2002)*, volume XIII, pages 598–603, Orlando, FL, July 2002.
- [17] Walter J. Freeman. Simulation of chaotic EEG patterns with a dynamic model of the olfactory system. *Biological Cybernetics*, 56:139–150, 1987.

- [18] K. Shimoide, M. C. Greenspon, and Walter J. Freeman. Modeling of chaotic dynamics in the olfactory system and application to pattern recognition. In F. H. Eeckman, editor, *Neural Systems Analysis and Modeling*, pages 365–372. Kluwer, Boston, 1993.
- [19] Walter J. Freeman and Koji Shimoide. New approaches to nonlinear concepts in neural information processing: Parameter optimization in a large-scale, biologically plausible corticle network. In Zornetzer, editor, *An Introduction to Neural and Electronic Networks*, chapter 7, pages 119–137. Academic Press, 1994.
- [20] Ricard Solé and Brian Goodwin. *Signs of Life: How Complexity Pervades Biology*. Basic Books, New York, NY, 2000.
- [21] P. Bak, C. Tang, and K. Wiesenfeld. Self-organized criticality: An explanation of 1/f noise. *Physical Review Letters*, 59(4):381–384, 1987.
- [22] K. Kaneko and I. Tsuda. Constructive complexity and artificial reality: An introduction. *Physica D*, 75:1–10, 1994.
- [23] Robert Kozma. On the constructive role of noise in stabilizing itinerant trajectories of chaotic dynamical systems. *Chaos*, 1193:1078–1090, 2003.
- [24] Robert Kozma, Walter J. Freeman, and Peter Erdi. The KIV model - nonlinear spatio-temporal dynamics of the primordial vertebrate forebrain. *Neurocomputing*, 52-54:819–826, 2003.
- [25] Olivier Michel. Khepera simulator package version 2.0. Downloaded from WWW at <http://wwwi3s.unice.fr/om/khep-sim.html>, 1996. Freeware mobile robot simulator written at the University of Nice Sophia-Antipolis.

-
- [26] Michael A. Arbib, Péter Érdi, and János Szentágothai. *Neural Organization: Structure, Function and Dynamics*. The MIT Press, Cambridge, MA, 1997.
- [27] M.V. Tsodyks, W.E. Skaggs, T.J. Sejnowski, and B.L. McNaughton. Population dynamics and theta rhythm phase precession of hippocampal place cell firing: A spiking neuron model. *Hippocampus*, 6(3):271–280, 1996.