

promoting access to White Rose research papers



Universities of Leeds, Sheffield and York
<http://eprints.whiterose.ac.uk/>

This is an author produced version of a paper published in **Proceedings of the Eighth International Conference on the Simulation of Adaptive Behavior.**

White Rose Research Online URL for this paper:

<http://eprints.whiterose.ac.uk/7960/>

Published paper

Buckley, C., Bullock, S. and Cohen, N. (2004) *Toward a dynamical systems analysis of neuromodulation*. In: Schaal, S., Ijspeert, A.J., Billard, A., Vijayakumar, S., Hallam, J. and Meyer, J-A., (eds.) *From Animals to Animats 8: Proceedings of the Eighth International Conference on the Simulation of Adaptive Behavior. SAB'04, 13 - 17 July 2004, Los Angeles, CA, USA.* MIT Press , Cambridge, Mass., pp. 334-343.

Toward a Dynamical Systems Analysis of Neuromodulation

Chris Buckley, Seth Bullock and Netta Cohen

Biosystems Group, School of Computing

University of Leeds, UK

christb|seth|netta@comp.leeds.ac.uk

Abstract

This work presents some first steps toward a more thorough understanding of the control systems employed in evolutionary robotics. In order to choose an appropriate architecture or to construct an effective novel control system we need insights into what makes control systems successful, robust, evolvable, etc. Here we present analysis intended to shed light on this type of question as it applies to a novel class of artificial neural networks that include a neuromodulatory mechanism: GasNets.

We begin by instantiating a particular GasNet subcircuit responsible for tuneable pattern generation and thought to underpin the attractive property of “temporal adaptivity”. Rather than work within the GasNet formalism, we develop an extension of the well-known FitzHugh-Nagumo equations. The continuous nature of our model allows us to conduct a thorough dynamical systems analysis and to draw parallels between this subcircuit and beating/bursting phenomena reported in the neuroscience literature.

We then proceed to explore the effects of different types of parameter modulation on the system dynamics. We conclude that while there are key differences between the gain modulation used in the GasNet and alternative schemes (including threshold modulation of more traditional synaptic input), both approaches are able to produce tuneable pattern generation. While it appears, at least in this study, that the GasNet’s gain modulation may not be crucial to pattern generation, we go on to suggest some possible advantages it could confer.

1. Introduction

Evolutionary roboticists employ a range of control architectures as varied as the research questions they address. Particular control systems are adopted for many different reasons. Some researchers wish to model a particular

biological phenomenon and place emphasis on incorporating this particular aspect (Alexander & Sporns, 2002). Other researchers, (Beer, 2000; Tuci, Quinn, & Harvey, 2002) choose controllers for their perspicuity, attempting to incorporate as few *a priori* assumptions as possible. From an engineering perspective, however, we wish to employ highly evolvable control architectures—ones that facilitate fast and effective evolution of high-quality solutions to the kinds of task we are typically faced with. Unfortunately, as of yet there are few guiding principles to help researchers choose an appropriate class of control system.

Recent successes of one novel type of artificial neural network control architecture, the GasNet (Smith, Husbands, & O’Shea, 2001), have led researchers to consider the evolvability of this architecture, its causes and generality. Unlike most artificial neural networks, GasNets include processes inspired by biological *neuromodulation* in addition to those modelled after the more traditionally understood process of *neurotransmission*. Although GasNets have only been tested on a small range of tasks to date, their success suggests that neuromodulation may increase the evolvability of this class of control system across a range of real-world tasks (Smith, Husbands, Philippides, & O’Shea, 2002).

In order to assess this hypothesis, it is important to try to identify which aspect or set of aspects of the GasNet architecture might be responsible for this increased evolvability. In doing so, we hope to be able to reveal design principles underlying successful control systems. Here we report the initial steps in a dynamical systems study of the GasNet system. Our goal is to gain a deeper insight into the possible root causes of the increase in evolvability exhibited by the GasNet. We will start by briefly outlining the idea of neuromodulation and suggesting a working definition to be used in the rest of this study.

2. Neuromodulation in Neuroscience

The traditional bias of modern neuroscience is to assume that fast synaptic transmission constitutes the primary means of neuronal interaction. Neuronal communication

is taken to be a composite of three dominating ideas (Katz, 1999):

1. Fast—pulses or on-off responses act on the 10 millisecond time scale.
2. Point-to-point—a neuron’s neighbourhood is completely specified by the incoming synaptic connections and the outgoing neuronal branches of its dendritic tree.
3. Inhibitory/excitatory—synaptic connections either increase or decrease the activation of a target neuron.

The transmission of synaptic pulses is confined to the synaptic cleft and is mediated by neurotransmitters. Hence the amalgamation of these notions is commonly known as *neurotransmission*, an idea with historical roots in a 1954 paper by Hodgkin and Huxley (HH) on the squid giant axon (Hodgkin & Huxley, 1954). This seminal paper is still regarded as the progenitor of all modern neuronal models. Hodgkin and Huxley were particularly interested in the electrical signalling or firing of nerve cells, primarily because of its ubiquity throughout the nervous system, but also because the idea of electrical circuits were very much a dominating paradigm of the era. Consequently, but perhaps not intentionally, studies of the synapse has dominated subsequent research. Although it has often been recognised that synaptic pulses do not tell the whole story and that many more subtle processes take place, these processes have been largely ignored. Recently, this shortcoming is being comprehensively challenged. Many neuroscientists have voiced dissent against the exclusiveness of neurotransmission (Harris-Warrick, Nagy, & Nusbaum, 1992; Abbot, LeMasson, Siegel, & Marder, 2003; Turrigiano, 1999). The “patch clamp” is an experimental technique which has generated an avalanche of very detailed data on neuronal activation (Poggio & Glaser, 1993, Ch 1). Researchers discovered that many phenomena do not fit easily into the picture provided by neurotransmission. For example, Harris-Warrick remarks that “it is no longer possible to discuss sensory processing or motor coordination without considering the role that non-traditional forms of neuronal activity and communication play” (Harris-Warrick et al., 1992).

These non-traditional processes have been collectively grouped under the term *neuromodulation*. Although the word has been used for over 20 years, the ubiquity of such processes has only just been acknowledged. Throughout this paper we will employ a working definition of neuromodulation provided by Katz that casts neuromodulation as the antithesis or complement of neurotransmission:

“Any communication between neurons, caused by the release of a chemical that is either not fast,

or not point-to-point or not simply excitation or inhibition” (Katz, 1999, p.3).

Neuromodulators act over a range of spatial and temporal scales. Neurotrophines and hormones, for instance, are large macromolecular chemicals that are transported via the circulatory system. They can change many properties of the neural tissue that they come into contact with, affecting the rates of synaptic growth and intrinsic properties of neurons (Turrigiano, 1999). Although they act over large volumes of tissue, they are relatively slow and diffuse. In comparison amphipathic neurotransmitters are small molecules that can pass through lipid tissue. Consequently they act over small volumes but are relatively fast (Changeux, 1993). They have been implicated in short term changes to synaptic strengths and the process of learning. Crudely, whereas neurotransmission has been conceived of as analogous to the operation of an electrical circuit, one can visualize neuromodulation as waves of gases and liquids diffusing from neurons or perhaps neuronal modules. They affect volumes of neural tissue and change the functionality of the neurons and synapses within it. In contrast to the dominating paradigm of electrical circuitry, a colourful term sometimes used to convey this alternative idea is “the liquid brain” (Changeux, 1993; Husbands, Philippides, Smith, & O’Shea, 2001).

3. Neuromodulation in Robotics, Gas-Nets

Generally, dynamic artificial neural networks are represented as the simplest extension of the feedforward McCulloch-Pitts perceptron (Arbib, 1998, p.4-11). They were first popularised by Hopfield in the 1970’s (Haykin, 1999; Hopfield, 1982), and are described by equations of the following form:

$$u_i(t+1) = F \left(\sum_{ij} \omega_{ij} u_j(t) + \theta_i \right) + I_i \quad (1)$$

The neural activation, u_i , of a network node, i , represents the mean firing rate of either a single neuron or an average over an ensemble of neurons. Each such node possesses a threshold, θ_i , and receives synaptic input from any neighbour, j , weighted by a synaptic link, w_{ij} , as well as external input, I_i . The function F is known as the transfer function and is generally of sigmoidal form, i.e. approximately linear in its mid-range with saturating limits. Note, the only dynamic variables are the neural activations, every other parameter is fixed before the network operates.

In the GasNet formalism this network of nodes is embedded in a 2-dimensional continuous space (Husbands et al., 2001). The neural activation of a node is now described by an equation of the following form:

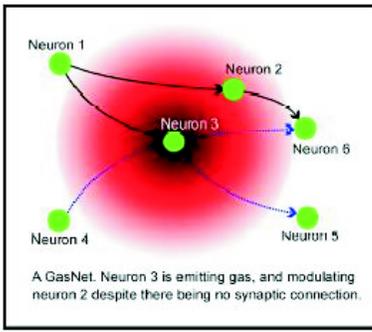


Figure 1: Gas diffuses across a control network embedded in a 2-d plane, from (Husbands et al., 2001)

$$u_i(t+1) = F(\kappa_i[\sum_{ij} \omega_{ij}u_j(t)] + \theta_i) + I_i \quad (2)$$

This differs from equation (1) only by the presence of a prefactor, κ_i , which is proportional to the concentration of “gas” in the vicinity of the node. Furthermore, each node may possess the ability to emit slowly diffusing gas under certain conditions—e.g., when either gas concentration at the node’s location, or the node’s neural activation, exceeds some fixed, node-specific threshold. Thus, as gas ebbs and flows across the plane in which the node’s are located, it modulates the sum of the inputs in a multiplicative fashion, deforming the network’s weight space (see Figure 1).

This gas mechanism exhibits all three qualities of neuromodulation as defined above. It is not point-to-point but diffuses in nature, it acts on a slower timescale than the electrical signals, and finally it neither simply excites nor inhibits neurons, but rather *modulates* the gain, or the steepness, of their transfer functions. Given this trivium of novel functionality, plus the embedding of the network in a continuous space, it is unclear exactly which element or combination of elements in the design of a GasNet is responsible for the increased evolvability that they exhibit.

Smith et al. (2002) made the first attempt at analysis of the GasNet architecture. In particular, they identified a subcircuit of the GasNet control system that had been evolved to solve a robot control problem, and showed that it contributed to the “temporal adaptivity” of the evolved solution.¹ This subcircuit was capable of tunable pattern generation, in that it could produce electrical oscillations that were slowly modulated by gas such

¹The term “temporal adaptivity” refers to the ability of an agent to respond to changes on a wide range of timescales. Biological agents are faced with an environment that exhibits regularities at scales ranging from seconds and minutes to diurnal and annual rhythms. They must adapt their internal structures to mirror and exploit these regularities. The ability to do this has also been suggested as a vital aspect of artificial control systems.

that rhythmic behaviours were readily achieved. However the discrete nature of the GasNet system hindered the completion of a more comprehensive dynamical systems analysis.

Here we will attempt to extend and generalise this analysis to some degree by developing and analysing a simple, idealised, continuous neuromodulatory network. Our model is based on the well-known FitzHugh-Nagumo equation (Murray, 1989, p.161-166), but inherits some important properties from the original GasNet formalism. We will explore the extent to which behaviours exhibited by GasNet sub-systems are characteristic of neuromodulatory systems more generally. In doing so, we can take advantage of both the continuous nature of the FitzHugh-Nagumo equation, which readily admits of dynamical systems analysis, and the body of existing analysis of the FitzHugh-Nagumo system. Moreover, the ubiquity of the FitzHugh-Nagumo model in neuroscience ensures that we have some chance of relating our findings on neuromodulation both to the adaptive behaviour and neuroscience communities.

4. Neuromodulation in the FitzHugh-Nagumo Equations

4.1 The FitzHugh-Nagumo Equations

While the (HH) equations were extraordinarily successful at reproducing key experimental data, a deep understanding of the underlying dynamics was hindered by their inherent complexity. In the early 1960’s, the FitzHugh-Nagumo model (FHN), a simple, analytically tractable, yet non-trivial reduction of HH, was developed (FitzHugh, 1961; Nagumo & Yoshizawa, 1962). Electrical pulses are derived from the differential permeability of the neural tissue to chemical ions. The dominant ionic species in nerve membranes are potassium and sodium, but in general there are many ionic species acting over many timescales. Since the timescales of the ionic species in the HH model are not of the same order, simplification could be achieved by assuming that sufficiently fast variables would settle to their steady state values almost instantaneously. This allowed certain variables to be eliminated from the Hodgkin-Huxley equation. The FHN model is given in equations (3) and (4).

$$F(v, w) \equiv \frac{dv}{dt} = f(v) - w + I_\alpha, \quad f(v) = v(a-v)(v-1) \quad (3)$$

$$G(v, w) \equiv \frac{dw}{dt} = \frac{bv - \gamma w}{\tau} \quad (4)$$

Here, v is the membrane potential, while w plays the role of the ionic currents. The remaining terms I_α , b , γ and τ are all positive constants and $0 < a < 1$. This model has been extraordinarily successful and displays

many of the key phenomena discovered in the original Hodgkin-Huxley model. For example FHN readily exhibits excitable and oscillatory behaviour (Murray, 1989, p. 164). Consequently FHN seems a natural choice for investigating some of the issues raised by GasNet research.

Although the FHN equation involves only two free variables (alongside time) we cannot solve it directly. Instead progress can be made by investigating the stability of equilibrium states of the model. Equilibrium behaviour is defined as the state, or sequence of states, that the system settles to after some finite transient. The equilibrium points correspond to points in the phase space where all the derivatives with respect to time of the system are equal to zero. For FHN, the equilibrium points can be found by setting the derivatives of (3) and (4) to zero, giving equations (5) and (6).

$$w = f(v) + I \quad (5)$$

$$w = \frac{b}{\gamma}v \quad (6)$$

If we plot these equations we obtain the nullclines of the system (see Figure 2). For the whole system to be in equilibrium, derivatives of both v and w must equal zero, this point is represented by the intersection of the nullclines.

These equilibria can be stable (if perturbed from this value the system quickly returns) or unstable (if perturbed from this point the system may not return). With the aid of linear stability analysis it is possible to derive a condition for the local stability of any given equilibrium point. For details of such analysis the reader is referred to (Strogatz, 1994; Murray, 1989). For a very good introduction to dynamical systems theory for robot control systems see (Beer, 1995). In the case of FHN, stability is determined by a single quantity known as the *trace*. For this system the condition for stability is given below. Where v^* is the value of v at equilibrium.

$$Tr(A) = \frac{\partial f(v^*)}{\partial v} - \frac{\gamma}{\tau} < 0. \quad (7)$$

The FHN equations exhibit three classes of behaviour which are determined by the number and stability of the equilibrium points. Figures 2a and 2b have a single intersection which is locally and globally stable. All initial conditions of the equations will relax to this point. Such configurations represent excitable systems—perturbations generate short-lived spiking followed by a return to quiescence. Alternatively if the nullcline cross at their centres they can exhibit fixed point or cyclic behaviour, Figure 2c and Figure 2d respectively.

The empty circles drawn on all plots represent the bifurcation points of the system. They delimit the region, for a particular set of parameters, within which

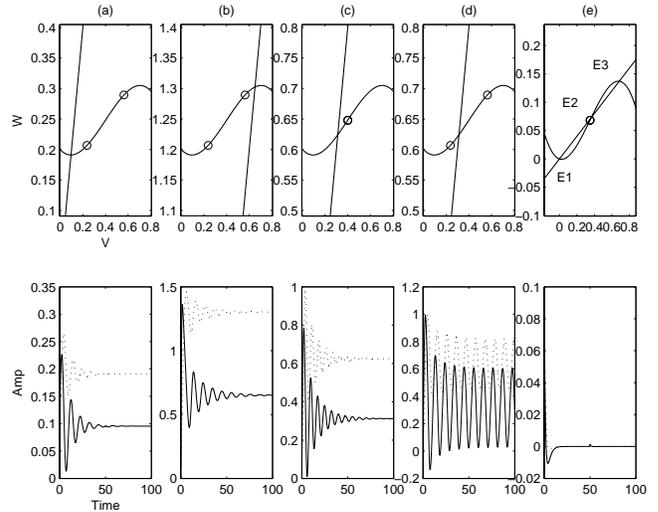


Figure 2: Modes of FHN behaviour: (a) and (b) represent non-oscillatory fixed-point behaviour but are excitable under perturbation. If the nullclines cross at their centres then the system can exhibit a fixed point (c) or cyclic behaviour (d). (e) shows a configuration with multiple equilibrium points. In general E1 and E3 are stable, while E2 is unstable. The circles drawn on all plots represent the bifurcation points of the system. They indicate the region within which the value of v^* is such that the first term of equation (7) makes the *trace* positive and hence unstable (which gives oscillations in this system). Outside this region the *trace* is negative and stable.

the value of v^* is such that the first term of in equation (7) makes the *trace* positive and hence the system is unstable (which gives oscillations in this configuration). Outside this region the *trace* is negative and the system is stable. Note the size of this region will be altered by the right hand term of equation (7). We will come back to this effect later in this work. Figure 2e presents a multistable configuration of the nullclines. It exhibits three equilibrium points with two stable (E1 and E3) and one unstable (E2). It can be shown that cyclic behaviour is not possible in this configuration.

Our work will focus on the case depicted in Figure 2d, the only case in this system that supports oscillation and is thus capable of rhythmic pattern generation. This configuration of nullclines has been implicated in the production of effective robot control. For example (Mathayomchan & Beer, 2002) evolved recurrent neural networks to perform a simple ball catching task. It was discovered that seeding the initial evolutionary population with “centre crossing” networks increased the speed and reliability of the evolutionary process.

Furthermore the ability to alternate between oscillatory and non-oscillatory behaviour has been identified in a host of neural systems. A parameter change

that causes a qualitative change in a system’s dynamics (e.g., a transition from non-oscillatory to oscillatory behaviour) is known as a bifurcation. The bifurcation in this particular case is known as the Hopf bifurcation and has been the focus of a great deal of investigation, both in its own right (Guckenheimer, Gueron, & Harris-Warrick, 1993) and as a model for biological systems (Rinzel & Ermentrout, 1989). In general it is thought that systems existing at or near bifurcations in their dynamics have greater potential for exhibiting interesting behaviour—they are more computationally rich (Poggio & Glaser, 1993, Ch 3). Bifurcation points have also been directly implicated in modern research on neuromodulation. For example, researchers argue that it might be advantageous for systems to dwell near bifurcation boundaries, thus making them sensitive to neuromodulatory input (Fellous & Linster, 1998).

4.2 A Dynamical Pattern Generator

A GasNet “discrete dynamical pattern generator” was outlined by (Smith et al., 2002). It is a simple system of coupled neurons that utilizes a gas dynamic to modulate electrical oscillations such that tuneable rhythmic patterns of activity are generated. This key behaviour is obtained by slowly switching the system from an unstable to a stable dynamic. In order to explore a similar mechanism using the FHN equation, we added to it a modulatory mechanism analogous to the gas employed in the GasNet.

Our simple system comprises one node governed by the FitzHugh-Nagumo equation that emits a modulator, M_1 , when its electrical activation rises above a fixed threshold, T_v . A second node is modelled in a much more simple fashion. Should the concentration of M_1 at this node rise above a fixed threshold, T_C , it emits a second modulator, M_2 . The first node is sensitive to M_2 , in that the concentration of this modulator affects the parameters of the node’s FHN equation. Note that no electrical activity is modelled for the second node, which is merely a source of modulator that can be switched on and off (see figure 3). Our equations for modulator growth and decay at both nodes are given in (8)-(11).

$$\frac{dC_1}{dt} = H_1(v)G_1 + (H_1(v) - 1)D_1 \quad (8)$$

$$H_1(v) = \begin{cases} 1, & \text{if } v > T_v \\ 0, & \text{otherwise} \end{cases} \quad (9)$$

$$\frac{dC_2}{dt} = H_2(C_1)G_2 + (H_2(C_1) - 1)D_2, \quad (10)$$

$$H_2(C_1) = \begin{cases} 1, & \text{if } C_1 > T_C \\ 0, & \text{otherwise} \end{cases} \quad (11)$$

The concentration of each modulator is represented by C_i , with their specific growth and decay rates denoted

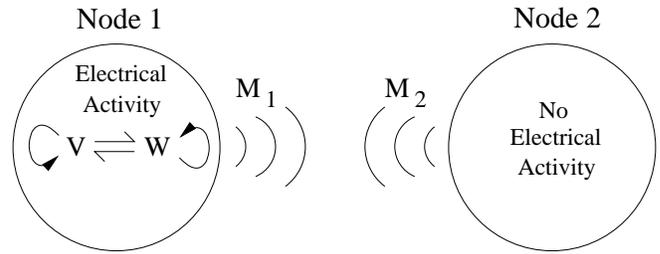


Figure 3: A dynamical pattern generator circuit: The electrical behaviour of node 1 is described by the FHN equation. It releases a modulator, M_1 , when its electrical potential, v , rises above its electrical threshold, T_v . Node 2 is triggered to release a second modulator, M_2 when the concentration of M_1 rises above its modulator threshold T_C . Reciprocally, the concentration of this second modulator affects the parameters of node 1’s FHN equation.

G_i and D_i , respectively. Each node’s heavyside function returns unity when it is emitting modulator, zero otherwise.

The concentration of M_2 affects the behaviour of node 1 in the following manner. Increasing concentration of M_2 decreases the γ parameter of node 1 in the manner described by equation (12), where C_{2max} represents a ceiling concentration value for M_2 , and $[\gamma_{min}, \gamma_{max}]$ describes a legal range of values for γ . This γ modulation affects the node’s nullclines in a fashion analogous to the gain modulation described in section 3. That is, γ could be considered to play a similar role to that played by κ in the GasNet equation (2).

$$\gamma = \gamma_{max} - \frac{C_2}{C_{2max}}(\gamma_{max} - \gamma_{min}) \quad (12)$$

The system is initialised in an oscillatory configuration. The initial nullclines of the system are depicted by the solid lines and crosses (the bifurcation points) in Figure 4d. Figures 4a and 4b represent the v and w time traces of the system with the fixed electrical threshold T_v superimposed over v . Figure 4c shows the build-up and decay of modulators M_1 and M_2 , while figure 4d displays how the nullclines change as a result of modulation. The w nullcline oscillates between the solid and dashed lines, altering the system’s equilibrium *position*. This change in the equilibrium position alters the first term of equation (7) from a negative to positive value and hence from an unstable oscillating solution (inside the crosses) to a stable fixed point (outside the crosses). Furthermore, because γ also affects the *trace* through the second term in equation (7), the modulation also alters the *range* of the oscillatory region. As modulation moves the w nullcline from the solid to the dashed line, the portion of the v nullcline associated with oscillatory behaviour grows from the section spanned by the two crosses to that spanned by the two circles.

The general effect of this modulation is to produce a beating/bursting system, with fast oscillation of the v and w variables within a low-frequency, modulated packet. Initially, as the system oscillates, M_1 builds, stimulating the emission of M_2 . As the concentration of M_2 rises, it decreases γ such that node 1’s equilibrium position lies outside the oscillatory region. The delay between the build up of M_1 and M_2 produces the low frequency packet. There is also a smoothing effect on M_2 , since while M_1 displays small amplitude, high-frequency oscillations as it builds, these are not present in the dynamic of M_2 .

This behaviour strongly resembles that of the discrete dynamical pattern generator identified within an evolved GasNet solution (Smith et al., 2002). The amplitude and frequency of the slow packet can easily be tuned by altering the speed and maximum concentration of the modulators. This shows that important aspects of the GasNet functionality can be instantiated and understood in a wider class of continuous systems which offer an advantage in terms of tractable analysis. Hence, we have demonstrated that dynamical pattern generation of the kind described for the GasNet is not particular to the specifics of that formulation of neuromodulation.

Furthermore the continuous nature of this system allows us to draw parallels between its behaviour and bursting behaviours referred to in the neuroscience literature (Rinzel & Ermentrout, 1989). The extent to which the behaviour of our simple model can be used to explore bursting remains an open question. However, in light of this observation, it may be interesting to use bursting ideas from neuroscience as a tool for understanding and designing robot control systems.

5. Not Inhibitory or Excitatory

We will now take a closer look at the movements of the nullclines during modulation of several kinds. In particular we wish to understand whether multiplicative gain modulation is necessary to effect dynamical pattern generation.

One of the key tenets of neuromodulation is that it “is not simply excitatory or inhibitory” but modulatory—i.e. it alters behavioural parameters, rather than merely activation levels. Given the tractability of our model, we are in a position to explore this assertion more closely. First of all let us assess what we mean by “not simply excitatory or inhibitory”. In the GasNet formulation this idea is quite clear—the gain parameter acts as a *multiplicative* factor in the update function in contrast to the merely additive character of synaptic inputs, see equation (2).

The idea of multiplicative connections is not new. Indeed, Pollack (1987) argued that greater computational power would not come about without their use. Typically the output of a particular node is calculated as

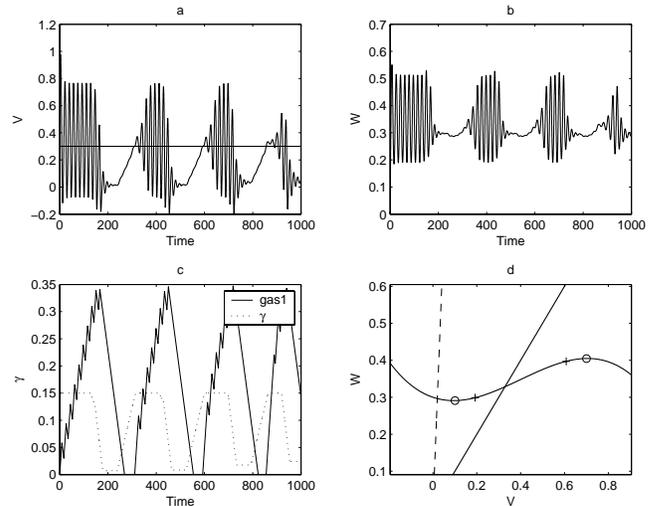


Figure 4: The output of the dynamical pattern generator: The figure shows the output of the system under γ modulation. Panels a and b show v and w , respectively. Figure c depicts the growth and decay of modulators M_1 and M_2 . Figure d displays the movement of the nullclines and bifurcation points—the solid lines and crosses depict their positions in the absence of M_2 , while the dashed line and circles denote their locations in the presence of a maximum concentration of M_2 .

a function of the sum of the synaptic inputs. Pollack thought that connections that multiplied the sum of the synaptic inputs were equally as important. Furthermore he believed that full Turing-complete computability could not be realized without such connections. This was later shown not to be the case (Siegelmann & Sontag, 1995). Nevertheless, Pollack goes on to state that “multiplicative connections remain a critical and under appreciated component of neurally inspired computing” (Pollack, 1999). Although the addition of multiplicative connections may not confer new functionality on simple additive networks, they may alter what behaviour is easily obtained, and hence network *evolvability*. The next section will look closely at the different types of parameter modulation possible in our augmented FHN model.

5.1 Modulation in the FHN Equation

Figure 5 shows the change in the nullclines under various parameter modulations of equations (3) and (4). The circles mark the oscillatory region on the v nullcline.

Figure 5a shows the effect of modulating γ . The *trace* is affected in two ways: because γ affects the gradient of the w nullcline, it changes the equilibrium position, v^* , and hence the first term of equation (7). Furthermore, it also changes the second term of equation (7), which scales the oscillatory region (note the difference between the locations of crosses and circles).

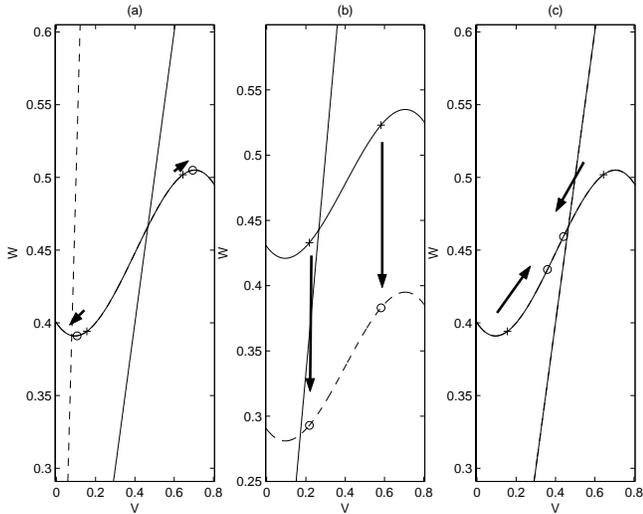


Figure 5: A transition from oscillatory to non-oscillatory behaviour caused by three types of modulation: (a) γ modulation, (b) I_γ modulation, and (c) τ modulation. Each panel displays the movement of nullclines and bifurcation points—solid lines and crosses depict their positions in the absence of modulation, while dashed lines and circles denote their locations in the presence of modulation. See text for details.

In contrast the size of the oscillatory region does not change under I_γ modulation, which merely translates one nullcline (see figure 5b). The only change to the *trace* is due to the first term of equation (7). This type of modulation is closely analogous to synaptic input in equation (1). Perturbing I_γ is analogous to raising or lowering a GasNet node’s electrical threshold, i.e. increasing or decreasing its level of activation. Hence, this type of change is not traditionally associated with neuromodulation since, at root, it is merely “excitatory or inhibitory”.

Perhaps the most interesting type of modulation considered here is shown in Figure 5c, for the parameter τ . Perturbing this parameter produces no change in the nullclines and hence no change in the position of the equilibrium point, v^* . Nevertheless, this kind of modulation affects the second term of the *trace* equation, and hence alters the range of the portion of the v nullcline associated with oscillatory behaviour. The effect of manipulating τ is to scale the w axis while preserving the nullcline configuration.

In general there are two ways the nullclines can change, first the equilibrium point can be *translated*, which corresponds to a change in the first term of equation (7). Second, altering certain parameters can change the *size* of the oscillatory region, which corresponds to a change in the second term of equation (7). While γ (gain) modulation achieves a mixture of both effects, I_γ (threshold) modulation produces pure translation, and τ (time constant) modulation achieves pure scaling of the

oscillatory region, leaving the equilibrium position of the system unchanged.

5.2 Comparing Different Modulation Types

If the neuromodulatory mechanisms in GasNets are responsible for enhanced evolvability (perhaps through facilitating the construction and tuneability of dynamic pattern generators of the kind explored above) it is important to understand how the different aspects of neuromodulation contribute. Is the fact that GasNets employ gain modulation critical? Our definition of neuromodulation (not fast, not point-to-point, not inhibition/excitation) suggests that it might be since alternative styles of modulation (e.g., threshold modulation) are merely equivalent to excitation/inhibition. Here we explore the ability of different types of modulation to effect dynamic pattern generation in an attempt to discover whether the particular kinds of “multiplicative” modulation feted in the literature are in some sense behaviourally superior. Might simple threshold modulation do the job (despite not counting as neuromodulation)?

Figure 6 shows our FHN model system under I_γ modulation (i.e. equation (12) is modified such that the γ, γ_{min} and γ_{max} terms are replaced by equivalent I_γ terms). It successfully produces beating behaviour analogous to that seen under γ modulation. Variation in M_2 causes a vertical translation of the v nullcline such that the equilibrium point, v^* , lies at times inside, and at other times outside, the region associated with oscillatory behaviour. The size of the oscillatory region remains unchanged.

Figure 7 shows the system under τ modulation. Again, it successfully exhibits beating behaviour. However, this is not achieved by translation of the nullclines, but rather by scaling the region of the v nullcline associated with oscillatory behaviour. In the absence of M_2 , oscillatory behaviour is associated with the portion of the v nullcline spanned by the two crosses. As the concentration of M_2 increases, these points move together, reducing the size of the oscillatory region, until they collide at a point indicated by the open circle. At or above this level of M_2 concentration, no oscillatory behaviour is possible.

In both cases (as well as the case of γ modulation described earlier), slowly modulating a particular system parameter achieves beating by allowing the system to alternate between non-oscillatory and oscillatory modes of behaviour. The manner in which this alternation is achieved is all that varies. Thus, even though we have discovered some fundamental differences between the different forms of modulation, each remains able to support a dynamical pattern generator circuit. In particular, we have shown that threshold modulation (that is merely inhibitory/excitatory) is sufficient in this regard, despite not satisfying our definition of neuromodulation.

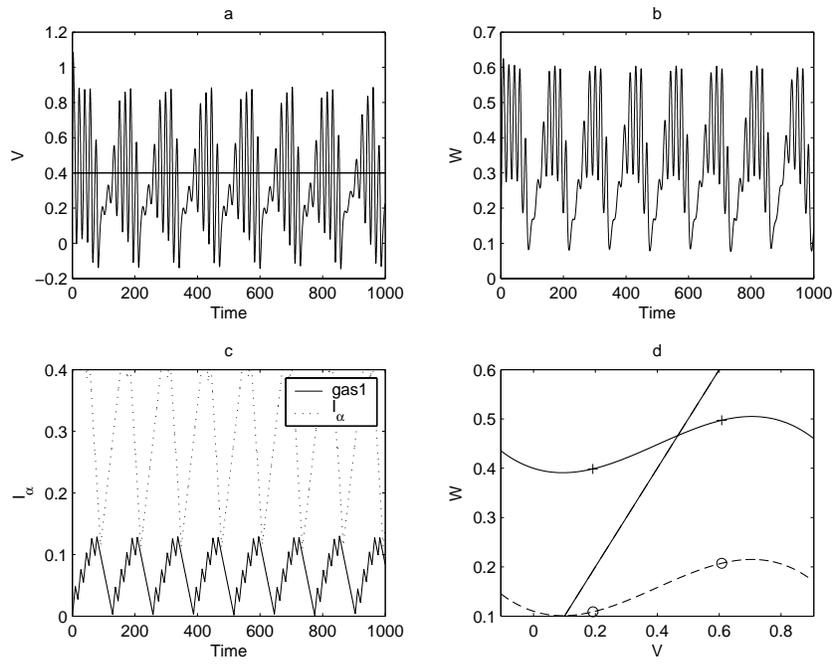


Figure 6: Output of the dynamical pattern generator under I_γ modulation.

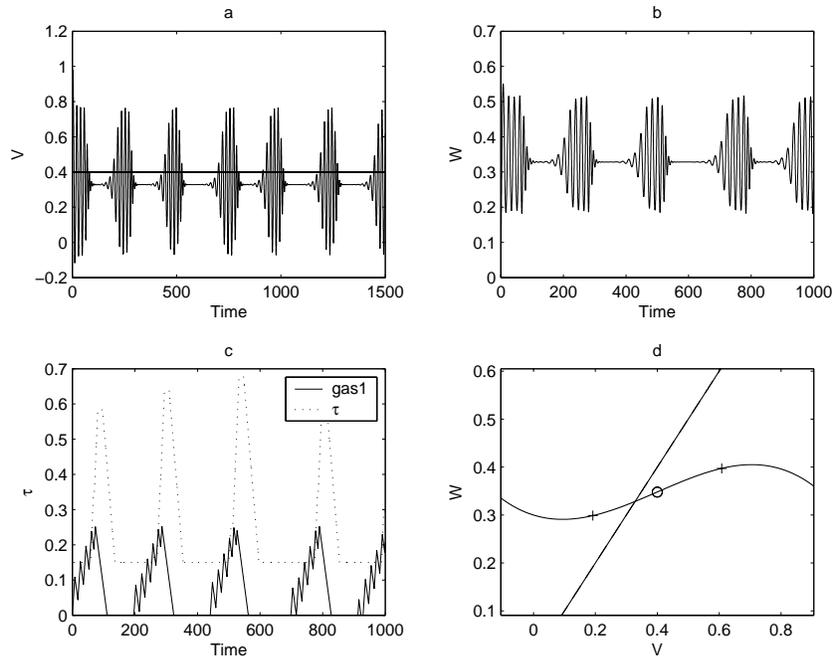


Figure 7: Output of the dynamical pattern generator under τ modulation.

However, while we have seen that it is possible for pattern generation to be achieved via a number of modulatory mechanisms, our current work suggests that the differences between gain modulation and threshold modulation are significant in some respects. Specifically, imagine an arbitrarily configured system of the kind described here that happens to be quiescent. Our analysis shows that gain modulation will be capable of exciting oscillatory behaviour in *any* such system. By contrast, mere threshold modulation will in general be unable to do so. Future papers will seek to characterise this and other differences between the behavioural and evolutionary potential of different modulatory mechanisms.

6. Conclusion

We have constructed a simple model of neuromodulation based on the FHN equation. Using this model we were able to obtain very similar behaviour to that produced by dynamical pattern generators observed in evolved GasNets. This lends weight to the idea that important GasNet behaviours are not specific to the particular idiosyncrasies of that formalism, but are more generic features of neuromodulation. We also noted that the observed behaviours were reminiscent of beating/bursting system studies in the neuroscience literature (Rinzel & Ermentrout, 1989).

We gave a working definition of neuromodulation suggested by (Katz, 1999). The GasNet contains all three characteristics of this definition. In particular gain modulation in the GasNet is concordant with the idea that neuromodulation is neither excitatory nor inhibitory, but rather modulatory. We investigated what this meant for our model system, observing how the system’s nullclines changed under different kinds of parameter modulation. We concluded that different parameter modulations exhibited key differences in their mechanics. While I_γ modulation, which we argued was analogous to modulation of electrical threshold or synaptic input, could only *translate* the nullclines, γ modulation, which is analogous to gain modulation, could also affect the *size* of the oscillatory region (our third mechanism, τ modulation, was only able to affect this aspect).

Given these differences we explored their effect on the ability of the system to reproduce the behaviour of a dynamical pattern generator subcircuit. We discovered that while each type of modulation was able to effect the dynamic pattern generation behaviour, each achieved this in a different manner. Initially this leads us to believe that the slow dynamic of the modulation (its growth and decay) is perhaps more important than its additive or multiplicative nature. However, our most recent work suggests that multiplicative interactions may nevertheless be important for system evolvability.

7. Future Work

The analysis and results presented here are only a first step toward a complete answer to the question: “what is the source of GasNet evolvability”? One obvious next step is to consider the role of spatial embedding in GasNet evolvability. To what extent is the low-dimensionality of the space within which GasNet nodes are embedded implicated in their ability to solve robot control tasks. Second, we wish to extend our simple model to consider small assemblies of multiple neurons, and the role of neuromodulation in their self-organisation and self-regulation.

We believe that in order to answer these questions, we need to confront some of the deeper issues concerning neuromodulation, not only with respect to behaviour-based robotics but also in the wider context of the neuroscience community. In particular, given growing evidence that multiple parallel processes at different spatial and temporal scales co-operate (operate together) across neural systems, we would like to know, as Eve Marder puts it, “what factors stabilise network operation in the face of so many competing influences?” (Poggio & Glaser, 1993, chapter 2)

References

- Abbot, L. F., LeMasson, G., Siegel, M., & Marder, E. (2003). Activity-dependent modification of intrinsic neuronal properties., 151–166.
- Alexander, W. H., & Sporns, O. (2002). An embodied model of learning, plasticity, and reward. *Adaptive Behavior*, 10(3/4), 143–159.
- Arbib, M. A. (Ed.). (1998). *The Handbook of Brain Theory and Neural Networks*. MIT Press, Cambridge, MA.
- Beer, R. D. (2000). Dynamical approaches to cognitive science. *Trends in Cognitive Sciences*, 4(3), 91–99.
- Beer, R. (1995). On the dynamics of small continuous-time recurrent neural networks. *Adaptive Behavior*, 3(4), 471–511.
- Changeux, J. (1993). Chemical signaling in the brain. *Scientific American*, 8, 58.
- Fellous, J.-M., & Linster, C. (1998). Computational models of neuromodulation. *Neural Computation*, 10(4), 771–805.
- FitzHugh, R. (1961). Impulses and physiological states in theoretical models of the nerve membrane. *Journal of Biophysics*, 1, 445–466.

- Guckenheimer, J., Gueron, S., & Harris-Warrick, R. (1993). Mapping the dynamics of a bursting neuron. *Philosophical Transactions of the Royal Society of London, Series B*, 341.
- Harris-Warrick, R. M., Nagy, F., & Nusbaum, M. P. (1992). Neuromodulation of the stomatogastric networks by identified neurons and transmitters. In Harris-Warrick, R. M., Marder, E., Selverston, A. I., & Moulins, M. (Eds.), *Dynamic Biological Networks: The Stomatogastric Nervous System*, pp. 251–291. MIT Press, Cambridge, MA.
- Haykin, S. (1999). *Neural Networks: A Comprehensive Foundation*. Prentice Hall, New Jersey.
- Hodgkin, A., & Huxley, A. F. (1954). A quantitative description of membrane current and its application to excitation and conduction in nerve. *Journal of Physiology (London)*, 117, 500–544.
- Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, 79, 2554–2558.
- Husbands, P., Philippides, A., Smith, T. M. C., & O’Shea, M. (2001). The shifting network: Volume signalling in real and robot nervous systems. In Kelemen, J., & Petr, S. (Eds.), *Fifth European Conference on Artificial Life: ECAL2001*, pp. 23–36. Springer, Heidelberg.
- Katz, P. S. (Ed.). (1999). *Beyond Neurotransmission: Neuromodulation and its Importance for Information Processing*. Oxford University Press, Oxford.
- Mathayomchan, B., & Beer, R. (2002). Center-crossing recurrent neural networks for the evolution of rhythmic behavior. *Neural Computation*, 14, 2043–2051.
- Murray, J. D. (Ed.). (1989). *Mathematical Biology*. Springer, Heidelberg.
- Nagumo, J. S., Arimoto, S., & Yoshizawa, S. (1962). An active pulse transmission line simulating nerve axon. *Proceedings of the Institute of Radio Engineers*, 50, 2061–2071.
- Poggio, T. A., & Glaser, D. A. (Eds.). (1993). *Exploring Brain Functions: Models in Neuroscience*. John Wiley and Sons, New York.
- Pollack, J. B. (1987). On connectionist models of natural language processing. PhD Dissertation in Computer Science, University of Illinois.
- Pollack, J. B. (1999). Connectionism: Past, present, and future. *Artificial Intelligence Review*, 3, 3–20.
- Rinzel, J., & Ermentrout, B. (1989). Analysis of neural excitability and oscillations. In Koch, C., & Sergev, I. (Eds.), *Methods in Neuronal Modeling*, pp. 251–291. MIT Press, Cambridge, MA.
- Siegelmann, H. T., & Sontag, E. D. (1995). On the computational power of neural nets. *Journal of Computer and System Sciences*, 50(1), 132–150.
- Smith, T., Husbands, P., & O’Shea, M. (2001). Not measuring evolvability: Initial exploration of an evolutionary robotics search space. In *Congress on Evolutionary Computation: CEC2001*, pp. 9–16. IEEE Press.
- Smith, T., Husbands, P., Philippides, A., & O’Shea, M. (2002). Neuronal plasticity and temporal adaptivity: GasNet robot control networks. *Adaptive Behavior*, 10(3/4), 161–184.
- Strogatz, S. H. (1994). *Nonlinear Dynamics and Chaos*. Addison-Wesley, Reading, MA.
- Tuci, E., Quinn, M., & Harvey, I. (2002). Evolving fixed-weight networks for learning robots. In *Congress on Evolutionary Computation: CEC2002*, pp. 1970–1975. IEEE Press.
- Turrigiano, G. G. (1999). Homeostatic plasticity in neuronal networks: The more things change, the more they stay the same. *Trends in Neuroscience*, 22, 221–227.