

Research Article

Open Access

Hendrik Reimann, Jonas Lins, and Gregor Schöner

The dynamics of neural activation variables

Abstract: This paper presents a comprehensive and detailed analysis of the elementary building blocks of neurally inspired architectures for cognitive robotics. It provides a brief outline of the fundamental principles by which biological nervous systems link to the environment in terms of perception, cognition, and behavior. We describe a class of dynamic neural activation variable based on these principles. We show that these dynamic neurons have the appropriate stability properties. Adding even simple connections between a small number of nodes is sufficient to constitute systems that make important decisions. Going through these mechanisms in detail, this paper should facilitate the design of neurally inspired architectures for behavior generation in robotic agents.

Keywords: Cognitive Robotics; Neural Dynamics; Bifurcations; Neural Networks

DOI 10.1515/pjbr-2015-0003

Received January 17, 2014; accepted January 17, 2015

1 Introduction

Cognitive robotics attempts to endow robotic agents with the ability to think and reason about the world and behave in it. The focus lies on developing cognitive architectures that are closely coupled to sensory and motor surfaces. Sensory stimuli from the environment are processed and transformed into appropriate actions without resorting to unnecessary abstractions like symbolic representations or artificial discretization [1]. Forming such symbolic representations of things in the world, and especially keeping them updated in real time, is computationally costly and often practically infeasible. Remarkably, humans and other animals manage to process sensory information and generate appropriate actions at speed and accuracy levels far surpassing that of any robotic agent, probably largely relying on non-symbolic representations

[2]. Cognitive robotics hopes to reproduce this proficiency by using biologically inspired mechanisms and principles.

One big source of such inspiration are the neural networks between the sensors and the effectors in humans and animals – our brains. Over the decades, neuroscientists have made tremendous advances in understanding how brains function to generate the amazing scope and flexibility of behavior exhibited by humans and animals. Although we are far from a complete understanding of human cognition, some overarching principles of neural information processing have emerged: the neural basis of cognition is distributed, sub-symbolic, continuous and noisy [3]. Higher-level functions like discrete symbols and stable representations of stimuli that are no longer directly available to the sensory systems, on the other hand, have to be embedded in and emerge from structures that adhere to these basic principles.

Dynamic neural activation variables can provide the elementary building blocks of systems that possess these desired characteristics. Cognitive architectures based on dynamic neural activation variables have been successfully employed in robotics for a wide variety of applications, ranging from relatively simple tasks like memorizing and generating sequences of actions [4] or dynamical organization of different elementary behaviors [5] to complex architectures performing high-level functions like object recognition [6] and scene representation [7].

In this paper, we give a detailed and comprehensive analysis of dynamic neural activation variables. We start with a brief review of how real neural systems encode and process information (Section 2) and how artificial networks of static activation variables have been used in artificial intelligence research (Section 3). Then we endow these activation variables with dynamics (Section 4) and make a detailed analysis of how the system behaves in time (Section 5) and how it reacts to varying external stimuli (Section 6). We finish with a detailed analysis of the simplest case of a dynamic neural network, two interconnected dynamic activation variables (Section 7).

Hendrik Reimann, Jonas Lins, Gregor Schöner: Institute for Neural Computation, Ruhr-University Bochum, Bochum, Germany,
E-mail: hendrik.reimann@ini.rub.de; jonas.lins@ini.rub.de;
gregor.schoener@ini.rub.de

2 Principles of neural coding

Neural processing is based on the influence of interconnected neurons on each other. Communication between two connected neurons occurs in the form of action potentials, short bursts of electricity propagated along the nerve fibers, usually called *spikes*. As the shape of the voltage fluctuation is largely identical for each spike, the specifics of any low level processes can safely be disregarded here [8].

The stereotypical nature of spikes means that neural processing must be based on the properties of spike trains rather than single spikes. Which aspect of the spike trains is relevant has been discussed controversially. Candidates based on precise timing, such as the exact length of interspike intervals, are rejected by many researchers, who argue that the necessary temporal fidelity is ruled out by the magnitude of random variability in spike timing (e.g., [9, 10]; for review see [11]). There is indeed little support for a constitutive role of spike timing in neural processing (apart from special cases; e.g., [12]; for review, see [13]). We therefore adopt the common view that the relevant functional aspect of single neuron activity is spike *rate*, that is, the number of spikes per unit of time. This is often referred to as *rate code* and is based on the assumption that there is a consistent relationship between the evolution of a neuron's spike rate over time and that of some property of the sensory environment, motor action, or other mental processes. In other words, spike rate represents certain aspects of these events or conditions.

Depending on the brain region, spike rate may in fact be related to all kinds of sensory, motor, or cognitive parameters. One common case is physical space. The classical concept of *receptive fields* [14] captures the fact that neurons in sensory brain areas only respond to stimuli within specific spatial regions of the sensory surfaces. When recording the spiking activity of a neuron in visual cortex, for instance, it is common to find an increase in spike rate whenever a dot of light is present within a certain region of the visual field. Depending on the exact area of the visual cortex, this region may span less than a degree of visual angle or include half the visual field. If the association between the presence of the stimulus in this region and the neuron's spike rate is consistent over trials, we may say that the spike rate of this neuron represents whether there is a visual stimulus present in the cell's receptive field.

This statement can be further refined through additional measurements and experimental conditions. For instance, we may map out a neuron's spike rate for each stim-

ulus position, plotting the so called *receptive field profile* [15]. Such experiments have shown that receptive fields have inner structure. For example, a common scheme in early visual cortex is that a neuron responds with maximum spike rate for a specific stimulus position and with lower and lower spike rate as the stimulus is moved away from that position. In this case, visualizing the receptive field profile by plotting spike rate against retinal space will result in a bell curve similar to a Gaussian. Given this curve, we may say that the neuron's spike rate does not simply represent the presence versus absence of a stimulus in its receptive field, but rather that different spike rates represent different distances of the stimulus from the peak of the receptive field profile.

Neurons with similar behavior are found throughout the sensory cortices, such as in the somatosensory cortex, where neurons have restricted receptive fields on the skin surface [14]. On the other hand, neural responses in motor areas are dependent in a similar manner on the spatial properties of movements. In the motor cortex, for example, the spike rate of many neurons is maximal for particular directions of hand movement and decreases with rising distance of the current movement angle from that 'preferred' direction [16, 17]. Another example is provided by the superior colliculus, a subcortical region, where neurons preferably respond to specific spatial targets of saccadic eye movements [18].

The fact that cells associated with motor action show similar spatial response properties as sensory cells hints that receptive field profiles are merely a special case of a more universal response scheme. The more general concept is known as *neural tuning* and applies to non-spatial feature dimensions just as it applies to physical space in the case of receptive fields. The non-spatial analogues to spatial receptive profiles are called *tuning curves*, which also tend to take bell shapes. In visual cortex, for example, many neurons respond strongly to lines of a particular orientation within their receptive fields, while they spike less for other orientation angles [19, 20]. Other examples are color [21] and shape [22] tuning in visual cortex, tuning to the direction of visual motion [23], pitch tuning in auditory cortex [24], or tuning to the orientation of tactile objects in somatosensory cortex [25]. These examples are prototypical for conditions throughout the brain, in that the basic principle of restricted sensitivity and preferred values remains unaltered regardless of the domain or brain area. Note, however, that tuning curves may also take more intricate shapes, such as differences of Gaussians (e.g., in the case of complex cells in visual cortex; [26]) or even more complex ones, particularly in 'higher' cortical areas. An additional complication is that neurons are typically sen-

sitive to and tuned along multiple dimensions at the same time, such as the two dimensions of retinal space, color, and orientation.

Finally, there is an additional scheme of neural responsiveness which, confusingly, is often referred to simply as ‘rate coding’ (disregarding that the concept of neural tuning as well is based on spike rate). It is characterized by a simple monotonic mapping between the underlying dimension and spike rate, and is more commonly found in the peripheral nervous system than in the brain. For example, the spike rate of pressure sensors in the skin increases with the degree of pressure and decreases as pressure is reduced [27]. A similar relationship holds between the spike rate of spinal motor neurons and muscle force [28]. In the brain, this ‘pure’ rate coding is often mixed with tuning properties in the same neurons (e.g., [29, 30]). For example, visual neurons tuned to orientation respond with increased spike rate to increases in stimulus contrast, while their orientation tuning properties remain unaffected by this [30]. Keeping all other parameters constant while mapping out the spike rate for each contrast value thus yields a monotonically increasing curve. The fundamental difference of this response scheme to neural tuning is that each value of spike rate represents a different value of the parameter dimension, so that in principle a single neuron is capable of signaling any given value along that dimension. In neural tuning, in contrast, each neuron only signals the distance of the current value from its preferred value.

This difference between rate coding neurons and tuned neurons has important consequences for the nature of representations supported by each scheme. Disregarding the problem of noise in neural firing, a single rate coding neuron is sufficient to represent every possible value along the underlying dimension. In the case of tuned neurons, representing a specific value along the underlying dimension rather than only a distance requires multiple neurons, each of them tuned to a different preferred value [31]. In this arrangement, a different subset of cells responds for each parameter value, determined by the overlap of the current value with the cells’ tuning curves. Specific parameter values are then represented by which neurons are active or, in other words, by the location of activity within the array of neurons, rather than directly by spike rate. For this reason, the scheme is known as *space coding*.

The notion of space coding is supported by the fact that there is typically an abundance of differentially tuned neurons for each sensory or motor dimension. This is especially obvious in regions where neurons are spatially ordered to conform to the structure of the dimensions to which they are tuned, which is known as *topographical*

organization. In primary visual cortex, for instance, neurons at neighboring horizontal positions of the cortical sheet have systematically displaced receptive fields in retinal space [32, 33]. This map of retinal space is interlaced with orderly maps of stimulus orientation for each retinal location [32, 34, 35]. Topography is not strictly necessary to enable space coding, though, as has been shown in the motor cortex [17, 29]. This is because a neuron’s position in the cortex does not necessarily correlate with its *connectivity*, which is the decisive factor for determining its role in the larger context of the network. Although topography may enable more economic connections, a neuron may also receive synapses from and send synapses to spatially remote neurons.

Note that the paramount relevance of connectivity is a general principle of neural organization. What a neuron ‘does’ – which stimuli it responds to, which effectors it affects, and the exact nature of such relationships – depends solely on its connectivity, say, from which set of retinal receptors it receives input, or which muscle fibers it is physically connected to. Even complex response patterns, such as neural tuning, are brought about by specific patterns of connectivity (e.g., orientation tuning; see [26]). From the perspective of an external observer, one might say that connectivity comes closest to being the one aspect of neural organization that assigns ‘meaning’ to the activity of neurons, as it determines the co-variation of neural activity with other variables in the world.

3 Neural activation patterns

The first step towards analyzing the connectivity structure in neural networks is descriptive. One lesson from the body of neurological research is that a wide range of sensory stimulus patterns are related to spatial patterns across whole regions of neurons rather than activation patterns of single neurons. Relatively simple features like the *redness* currently present in some region of the retina can be directly linked to the firing rate of a single neuron. For more complex features like *color*, the only relationships that have been successfully established are those with activation patterns over whole populations of neurons.

How can the structure of the relationship between sensory stimulus patterns and spatial patterns of neural activation be analyzed? How is the transformation from firing of photoreceptors on the retina to localized activation patterns in topographic color maps in the visual cortex realized? One attempt to answer these questions has led to the so-called *connectionist models* of neural activation.




Figure 1: Sigmoid function from Equation 2 rising smoothly from 0 to 1 as a given threshold u_0 is approached. The slope at the threshold is defined by the parameter β , larger β -values correspond to steeper growth.

These models attempt to explain the functional features of neural networks mathematically. Many neurophysiological details, on the other hand, are treated as implementation quirks and neglected. The following paragraphs provide a short overview of this approach and what can be accomplished with connectionist models.

The connectionist approach starts out with a functional description of a single node of activation, often simply referred to as a *neuron*. The state of this elementary unit of computation is expressed by a single parameter u , called the *activation variable*. This activation variable of a connectionist neuron can be loosely interpreted as the firing rate of a biological neuron. This is dangerous, though, as such detailed identifications inevitably break down at some point. For instance, the activation variable of the connectionist neuron can be negative, while the firing rate cannot. One should keep in mind that while some of the main functions of neural substrates in the brain are captured remarkably well by connectionist models, direct identifications of mathematical terms with biological implementation details are problematic and bound to break down at some point.

How can the connectivity patterns in networks of such nodes represent complex features and perform high-level tasks such as recognizing objects from a visual input pattern on the retina? Consider neural populations sensitive to the orientation of edges within their receptive field, similar to what has been found in the visual cortex [26]. As a simple example, take two neurons with preferred directions of 0 and 90°, encoding the presence of horizontal and vertical bars in their common receptive field. Now add another neuron that receives synaptic input from both these orientation detectors given by

$$v = \sigma(w_1 u_1 + w_2 u_2). \quad (1)$$

Here the weight parameters w_i represent the strength of the connection from u_i into v . The σ is a gating function that is 0 for small values, 1 for large values, and rises




Figure 2: Simple example of a connectionist network with two input neurons u_i , $i = 1, 2$, and one output neuron v . The strength of the connections are determined by the weight parameters w_i .

monotonously in between. This functional shape is called *sigmoid*, as it resembles a form of the greek letter sigma (compare Figure 1). The exact mathematical function used for this transformation is somewhat arbitrary as long as the described characteristics are met; a frequently used form is

$$\sigma(u) = \sigma_{\beta, u_0}(u) = \frac{1}{1 + e^{-\beta(u-u_0)}}. \quad (2)$$

What kind of features can be encoded by the neuron v , receiving synaptic input from u_1 and u_2 ? With the proper parameter settings, v will be active when both u_1 and u_2 provide input of a certain strength, i.e. when *both* horizontal *and* vertical bars are present in the receptive field. If the size of the receptive field is small enough, then any such bars will cross and form a corner – the high-level neuron v is a *corner detector*. A diagram of this network is shown in Figure 2.

This network is a simple example of a *perceptron* and was first introduced by Rosenblatt [36] as a model of information storage in neural systems. Its function is that of a classifier, as it basically makes one decision: is a certain pattern present in the input, or not? In our example, the system decides whether it perceives a corner or not.

The perceptron sparked increased interest and research activity in neural networks. Researchers asked what kind of classifications could be done by perceptrons, and how. While the simplistic example shown above can easily be parameterized by hand to perform a desired function, this quickly becomes infeasible for more complex cases. Indeed, perceptrons appealed to many people because it is possible for them to autonomously learn a given classification problem [37].

As for the *what* question, a theoretical analysis showed that regardless of the number of neurons, perceptrons are limited to classification problems that are *linearly separable*, i.e. a hyperplane can be found in the input space such that all members of one class lie on one side




Figure 3: A connectionist network with three input neurons, two output neurons and one hidden layer containing five neurons.

and those of the other class lie on the other side [38]. This limitation can be overcome by endowing the neural network with a richer connectivity structure, like adding layers between the input and output neurons, usually called *hidden* layers. Figure 3 shows a schematic of such a network, often referred to as *multi-layer perceptron*.

The computation or information processing capability of a single neuron is limited to transforming its inner activation level to an output value. While this may seem almost simplistic when seen in isolation, the nonlinearity of this transformation is an essential ingredient for the processing power of the network as a whole. It enables a perceptron with the appropriate connectivity structure to make decisions based on features that are present in an input pattern, but not salient. Artificial neural networks have been successfully applied to a wide range of different problems [37].

One defining feature of connectionist networks like the one shown in Figure 3 is that the neurons are organized in layers. When an input pattern is chosen for the lowest layer, the activation values subsequently be computed for the higher layers until the output layer at the top is reached. The flow of information in this layout is unidirectional from bottom to top, giving rise to the name *feed-forward network*. The output pattern is determined by a combination of the input pattern and the connectivity structure within the network. While this is a purely deterministic mapping, it is highly nonlinear due to the gating sigmoid in the activation function of each single neuron. The specifics of the mapping from input to output pattern performed by these networks are encoded in the connection weights w_{ij} , hence the term *connectionist*.

In general, feed-forward networks are very good for tasks that require a direct mapping from an input to a desired output pattern, like pattern recognition or function approximation. In the nervous system, though, the as-

sumption of uni-directional flow of activation is only occasionally satisfied. Most neural systems make decisions that are not uniquely determined by the input pattern, but factor in the internal state of the system itself. One obvious example for this is working memory. If an object is presented and then hidden in one of two identical containers in plain sight, we have no problem finding it, because we saw where it was hidden and remember it. Though seemingly automatic, this ability is learned in infancy and can be perturbed in very young babies [39]. This information is present neither in the visual input pattern, because the two containers look the same, nor in the connectivity structure, as synapses cannot adapt this fast, so it must be represented in the current pattern of activation of some sub-network.

In order for models to deal with these phenomena, we need to give up the feed-forward limitation and introduce recurrence and feedback into the network. The activation of a given neuron can in principle depend upon the output of any other neuron in the network, not only ones that are located in layers above it. In other words, these networks contain loops.

Neurologically, this is the normal case. It was estimated by Braitenberg that over 95% of all connections between cortical neurons are not part of a feed-forward stream [40]. In the highly interconnected central nervous system, it takes only between three and five synapses to get from any neuron to any other neuron. The typical neuron in the cortex receives about 10 000 synaptic inputs and projects to a similar number of other neurons. It is not surprising that feed-forward networks can model only a small fraction of the functionality of such a system.

4 A neuron in time

Systems whose state depends not only upon some other variable but also upon their own past are abundant in nature – indeed the case of a physical system not depending on time makes sense only as a rare limit case where the system moves so fast that dependence upon past states are very hard to measure: the amount of light in a windowless room directly depends upon the state of the light switch, but only if the electrostatics of the power connections and the time it takes for the filament in the light bulb to heat up and start glowing are neglected.

How can the dependence of a system upon the past be represented in a functional description? How can the activation level of a neuron be affected by the state of other neurons and its own history? One way to do this is to in-




Figure 4: Example trajectories for neural activation variables with two different categories of change rate. The system with a constant rate of change from Equation 5 grows in a straight line (red). The system with only random fluctuations from Equation 6 grows in the pattern of a random walk (green and blue for two different noise parameter settings).

roduce a functional dependence of the *change* of a neural activation variable upon its current state. The change of the activation is expressed by its temporal derivative

$$\frac{du(t)}{dt} = \dot{u}(t). \quad (3)$$

The dependence of the change of neural activation on a synaptic input is then simply written as

$$\dot{v}(t) = w\sigma(u(t)), \quad (4)$$

where as before u, v are the activation variables of two neurons and the weight parameter w gives the synaptic strength.

How does the activation level of such a dynamic neuron develop in time? Assume the simple case of a constant stimulus. For the neural activation variable u that changes according to

$$\dot{u}(t) = C, \quad (5)$$

with some constant C , the time course of the activation variable $u(t)$ will grow unchecked, as shown in Figure 4 (red line).

As another example, assume the case of a neuron that does not receive input with a well-defined structure over time, but is subject to random activations

$$\dot{u}(t) = \xi, \quad (6)$$

where ξ is gaussian white noise. Two example time courses of this system are shown in Figure 4 (green and blue). These trajectories also tend away from the initial state. It can be shown that for longer periods of time, the expected distance from the initial state grows larger.

It is obvious how this is problematic: for the not unusual case of a constant synaptic input and given enough time, the activation of the receiving neuron would grow infinitely large. Even in the absence of constant input, the activation of neurons that are only subject to random fluctuations would also tend towards infinity. From a dynamical systems perspective, this kind of behavior is called *unstable*. Obviously this is at odds with the actual dynamics of neural activation variables: in the absence of constant input, a neuron is usually at a resting level. Driven out of this state by transient random perturbations, the neural activation relaxes back towards the resting level very soon.

This resetting force that drives the neural activation towards the resting level can be modeled as a factor that depends upon the current activation relative to the resting level h

$$\dot{u}(t) = -(u(t) - h) = -u(t) + h, \quad (7)$$

where h is usually negative. The rate of change $\dot{u}(t)$ is proportional to the difference $u(t) - h$ between the activation and the resting level, but points in the opposite direction: when the activation is larger than the resting level, the rate of change is negative, and vice versa, so the rate of change will always drive the system back towards h . The dynamics of this system are illustrated Figure 5.

5 Stability and relaxation time

Points to which the activation variable will always return after small perturbations are called *stable fixed points*, or *attractors* of a system. For a noisy information processing system like the brain, stability is a very important feature: to be used reliably in mental operations, an internal state must be reasonably stable – otherwise the whole system would fluctuate and lead to erratic behavior.

Things can be more stable or less stable. Take a person standing upright upon a support surface of varying size: the mechanical system of the body in space is in an attractor that is maintained by coordinated muscle forces counteracting the pull of gravity. Now assume that the support surface is subjected to small perturbations, e.g. abrupt movements with fixed amplitude in a random direction. If the support surface is large, the subject will have no trouble retaining balance after a medium-sized perturbation. If the support surface is very small, though, the same perturbation might be sufficient to make the subject fall, and if he does remain standing, he will have much more trouble recovering from the movement.

Let us take a closer look at the recovery process after a perturbation. The three example trajectories of the lin-

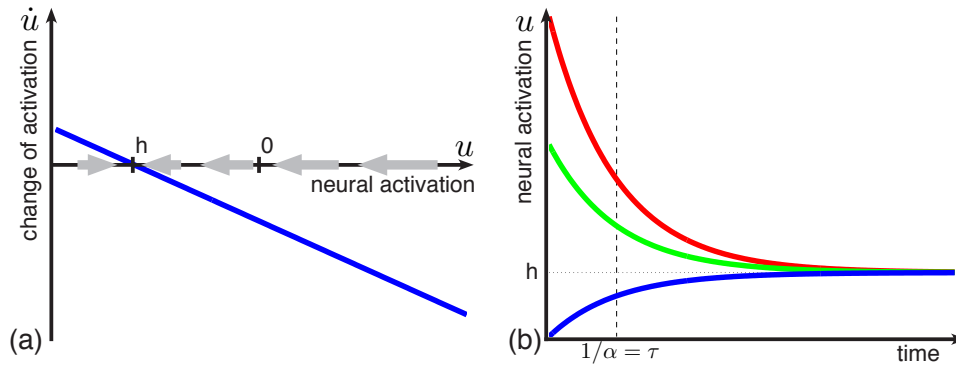


Figure 5: Dynamics of a single neural activation variable (Equation 7). Panel (a) shows the phase plot (blue), the grey arrows sketch the vector field over u . Note how the length of the arrows corresponds to the value of the plotted function at that point; both vanish at the resting level h . Panel (b) shows three example trajectories relaxing to the resting level h from different initial values. The time interval τ after which all trajectories have reduced the initial distance to the attractor h by 36.8% is marked by a dashed vertical line.

ear system (7) shown in Figure 5 can be interpreted as relaxation to the attractor at h after perturbations of varying size. Note that the shape of the curve is similar for all three trajectories – it has the same shape, only scaled. This scale solely depends on the size of the initial perturbation. For a general linear system of the form

$$\dot{u} = \alpha(-u + h), \tag{8}$$

the relationship between time and distance to the attractor can be described by the exponential function

$$u(t) = e^{-\alpha t} (u_0 - h) + h, \tag{9}$$

where u_0 is the initial state at $t = 0$ [41]. The time it takes for the system to reduce the perturbation by a fixed ratio r , say half of the initial distance from the attractor ($r = 0.5$), is the same in all three cases. It can be shown that this time only depends upon the parameter α . As a convention, this ratio is taken to be $r = e^{-1} \sim 0.368$, and it can be shown that the time needed for a system of the form (8) to reduce a perturbation to $r = 36.8\%$ of the initial value is exactly $\frac{1}{\alpha}$. This is called the *relaxation time* of a system.

These notions about relaxation time allow us to take a more formal look at the stability of a state like upright stance on a narrow support platform or the resting level of a neural activation variable. Intuitively we called a state more stable if the system returns to it faster after a perturbation. Defined this way, a measure of stability is given by the reciprocal of the relaxation time. For the linear system in Equation 8, the stability of the attractor at h is $(\frac{1}{\alpha})^{-1} = \alpha$.

It is important to note that for the system in Equation 8, the parameter α corresponds to the slope of the phase plot in Figure 5a. A higher α -value means that the same perturbation will elicit a larger change back towards the attractor. The state changes faster, so relaxation time is

shorter and the system is more stable. The slope of the phase plot is a good way to quickly assess the stability of an attractor.

This formal notion of stability even allows us to generalize the notion of an attractor by considering attractors with *negative* stability. Such a point corresponds to a zero crossing of the phase plot with positive slope. If the variable is exactly at that point, the rate of change is zero and the system will stay there indefinitely. Once it deviates slightly to one side, say becomes slightly larger, the rate of change becomes positive. This means that the activation will grow further, resulting in still further deviation, and so on – the state will grow with increasing speed. Such points are called *repellers*, because the activation tends to move away from them.

For the relaxation time of neural activation variables, we use the convention of always moving the parameter that determines it to the *left* side of the equation

$$\tau \dot{u}(t) = -(u + h). \tag{10}$$

Comparing Equation 10 with Equation 8, we see that in this form the relaxation time of the neural activation variable u is given by τ , as indicated in Figure 5b.

6 Varying inputs

Usually the input that a neuron receives is not constant but has a distinct time structure. Any external stimulus can be written as an additional term $s(t)$ adding to the rate of change of the neural activation

$$\tau \dot{u}(t) = -u(t) + h + s(t). \tag{11}$$

As a simple first case, consider a stimulus that is constant at first, then abruptly changes to a different value – for instance a photoreceptor when the room light is switched on. Figure 6b shows the trajectory of the activation variable u in these conditions. After the stimulus appears, u leaves the resting level h and relaxes to a new stable state $h + s_1$. A bit later, the stimulus is slightly reduced, and again the activation variable relaxes to the new attractor $h + s_2$. The phase plot in Figure 6a provides more insight into these dynamics. Before the stimulus is presented, the system is the same one as before, corresponding to the dotted line. Turning on the stimulus adds a constant term to the rate of change \dot{u} everywhere, essentially shifting the phase plot upwards. This results in the zero-crossing moving to the right along the u -axis from h to $h + s_1$. This new fixed point has the same stability characteristics as the old fixed point at h . Reducing the stimulus results in a downward shift of the phase plot, the zero-crossing moves to the left to the new value $h + s_2$.

Note from the point of view of other neurons receiving activation from this one, only the first change is relevant. The blue line in Figure 6b shows the gated output $\sigma(u(t))$ of the neuron. It switches from 0 to 1 after the stimulus is presented initially. Although the stimulus is reduced later, it is still sufficiently strong for the attractor to remain in the positive region of the activation variable, so the neuron remains in the *on*-state.

This system concludes what can be called the dynamic version of the feed-forward networks. While the systems presented so far have temporal structures that are distinct from the stimulus pattern, this difference is limited to a certain sluggishness in tracking the stimuli: the output pattern takes some time to emerge, but it is still uniquely determined by the input pattern. To loosen this bond and allow internal states to play a role, we need to introduce recurrent connections and feedback effects.

The simplest possible recurrent network is a single neuron whose activity feeds back to itself, called *self-excitation*. Figure 7 shows a diagram of this system with added input. The dynamics of the activation variable u of this single neuron are given by

$$\tau \dot{u}(t) = -u(t) + h + s(t) + w\sigma(u(t)), \quad (12)$$

where the last term represents the self-excitatory input with weight w . Figure 8a shows a phase plot of this system for a medium-strength stimulus. The phase plot shows clearly that this system is not linear, due to the nonlinear gating function $\sigma(u)$. In most regions of the u -line, though, where the sigmoid is almost constant, the phase plot is approximately linear. In these cases where the self-excitatory term is either negligible because the neuron is inactive,

or resembles a constant input because the neuron is active, the dynamics of this system are very similar to the dynamics of a neuron with constant input that we analyzed above. For values significantly below zero, the dynamics are similar to the case without input, while for values significantly above zero it resembles the case with input – the dynamics of both these cases are sketched again as dotted lines in Figure 8a. Around zero, where the sigmoid is neither clearly *on* or *off*, the dynamics are different from these linear cases.

Several example trajectories are shown in Figure 8b. We can see that for initial values significantly below zero, the activation variable relaxes to $h + s$, while for the values significantly above zero, it relaxes towards $h + s + w$. These two values are both attractors of the self-exciting neuron with the given stimulus s . Which attractor the system relaxes to is determined by the initial state: if the neuron is active, it will stay active, if it is inactive, it will also remain so. A system like this with two attractors is called *bistable*. This bistable dynamic neuron is a first example of a system that can account for hysteresis effects: the stable state that the neuron relaxes to depends upon the history of the activation state.

How does this single neuron with self-excitation react to varying input? There are two other distinct situations: a very strong input, and no input at all. Exemplary phase plots for these cases are shown in Figure 9a (solid blue lines). The phase plots reveal that in these cases of no stimulus and strong stimulus, the stable states of the single neuron with self-excitation are similar to those of a single neuron without self-excitation. In both cases, there is exactly one stable state, which is significantly above threshold in the presence of stimulus and significantly below threshold with no stimulus. The exact position of the stable state in the presence of input might be different, but remember that the activation gets transformed by a sigmoid, so from a functional perspective this difference is irrelevant.

Depending on different stimulus strength, we have determined three cases in which the dynamics of the neural activation variable is significantly different: two monostable cases for no (or very weak) stimulus and strong stimulus, and the bistable case for medium-strength stimuli. Let us now consider a stimulus that is not constant, but starts very weak and only slowly gains strength over time. In the phase plot, this corresponds to the whole curve slowly shifting upward, as sketched by the dashed blue lines in Figure 9a. At first, nothing much happens, the system is monostable. As the stimulus gets stronger, the apex in the positive region crosses the u -axis (orange dot), creating two new zero crossings that did not exist before: a

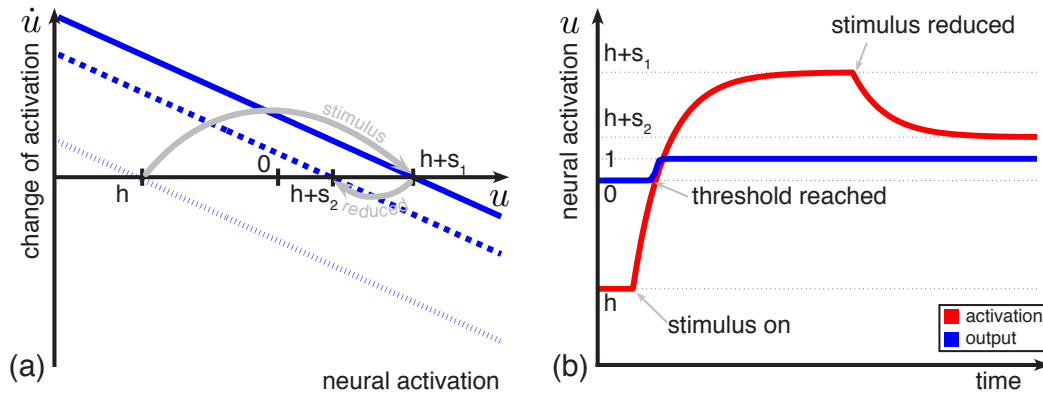


Figure 6: Dynamics of a single neural activation variable with stimulus (Equation 11). Panel (a) shows the phase plot for two different stimulus strengths s_1 (solid) and s_2 (dashed). The grey arrows indicate how the attractor is shifted from the dynamics without stimulus (dotted line) to new values determined by the stimulus strength. The example trajectory plotted in panel (b) shows how the activation variable u (red) tracks the stimulus by relaxing to the new attractor after each shift. The blue line shows the profile of the gated output $\sigma(u)$.




Figure 7: Connectivity diagram for a single dynamic neuron with self-excitation strength w and external stimulus s (Equation 12).

second attractor appears, the system is now bistable. As the stimulus gains more strength, the apex in the negative region also crosses the u -axis (magenta dot), and two zero-crossings disappear, the system is monostable again. These different cases that depend upon values of the stimulus parameter are called *regimes* of the system. The single neuron with self-excitation has two monostable regimes and one bistable regime.

What happens to the *state* of the neural activation variable as the stimulus strength increases and the system moves through these regimes? When the stimulus is absent, the neuron is inactive, as that is the only stable state. As the system enters the bistable regime, the second attractor appearing in the positive region, the activation variable still remains at the first attractor in the negative range: this is still a stable state and small perturbations due to noise are resisted by the system. Only after the stimulus gains sufficient strength for the first attractor to disappear will the activation variable leave the negative region: the rate of change \dot{u} over this range is consistently positive in that region, so u will grow until it relaxes to the attractor. The trajectory shown in Figure 9b (solid line) illustrates this pro-

cess. First the activation variable increases smoothly along with the stimulus. Then the attractor in the negative region disappears and the state rapidly relaxes to the attractor in the positive region. When the stimulus pattern is reversed in time, i.e. the stimulus *decreases*, the activation variable follows a markedly different pattern: it stays *on* over the bistable regime and only turns *off* when the attractor in the positive region disappears, as indicated by the dashed line in Figure 9b.

The appearance of new stable states as system parameters change, or the disappearance of old ones, is called a *bifurcation* of the system. The appearance of new attractors is usually a quiet process, because the system is always in different attractors and will only relax to the newly created one when it is driven out of the current state by a strong perturbation or when that state becomes unstable. In contrast to the appearance of an attractor, the loss of stability of a previously attractive state (as happened to the *off* state in the example above) has observable signatures.

The process described in the example above, where a stimulus gains strength until the *off*-state of the neural activation variable becomes unstable and the system relaxes to the stable *on*-state, is called the *detection instability*. The stimulus has been there for a while, slowly and continuously gaining in strength. The loss of stability of the *off*-state and subsequent relaxation to the *on*-state marks a quasi-discrete event in time that corresponds to the detection of the stimulus by the neural activation variable. It is important to note that this detection event is brought about by the nonlinear dynamics of the activation variable, transforming the continuously changing stimulus parameter into a discrete categorization of *there*

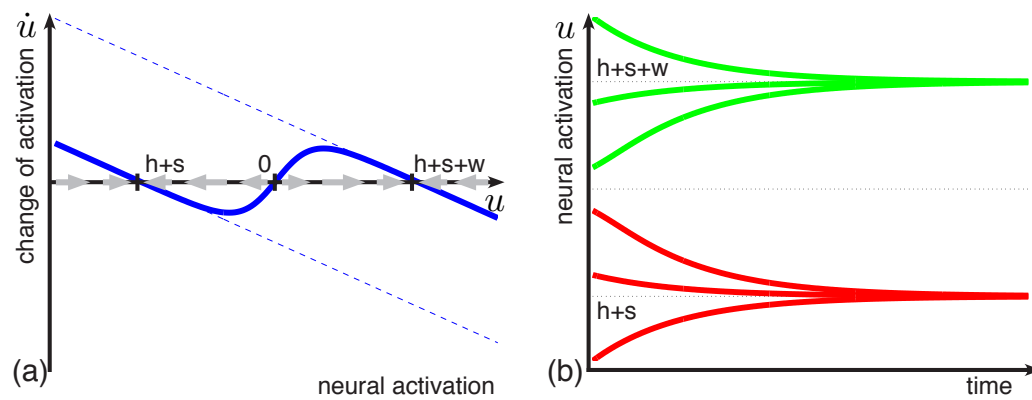


Figure 8: Dynamics of a single neural activation variable with self-excitation (Equation 12) under a medium stimulus. Panel (a) shows the phase plot (solid) and as a reference the phase plots for the system without self-excitation, with and without stimulus (dashed). The grey arrows sketching the vector field show how the direction of the rate of change switches between negative and positive at the repeller 0. Panel (b) illustrates how several example trajectories starting above (green) and below zero (red) relax to different attractors.

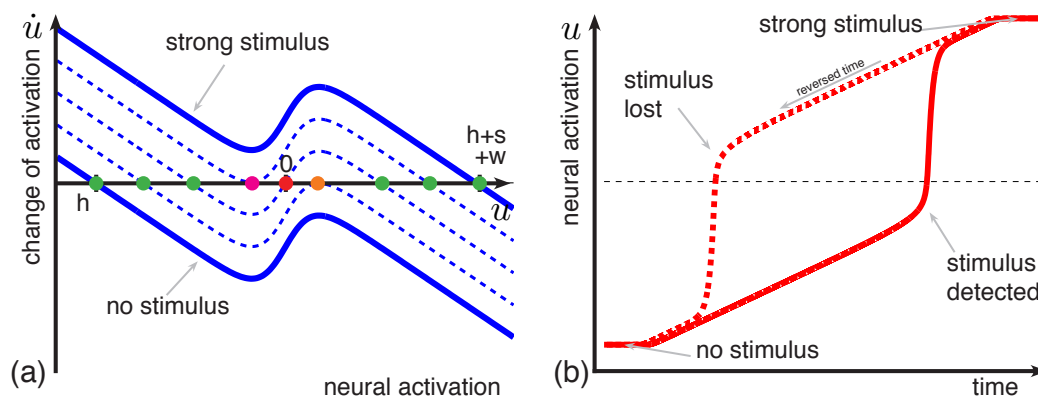


Figure 9: Dynamics of a single self-exciting neuron with systematically increasing stimulus strengths. The phase plot (a) is given for the limit cases of no stimulus and a strong stimulus (solid lines) and three intermediate values (dashed). The attractors and repellers are indicated by the green and red dots for each case, the magenta and orange dots are a pair of attracting/repelling fixed points that are on the verge of appearing (or disappearing) in a bifurcation. Panel (b) shows two example trajectories for increasing (solid) and decreasing (dashed) stimulus, the latter plotted reversed in time.

and *not there*. For this reason, we also say that the system makes a *detection decision*.

We have seen that for medium-strength stimuli, the state of the single self-exciting neuron depends upon the activation history, active neurons remain active, and vice versa. This implies that the *direction* in which the stimulus parameter moves through the bistable regime determines the state of the activation variable. Indeed, looking at Figure 9b, we can compare the time profiles of the activation variable for two different cases: one where the stimulus was absent and then slowly turned on, and the opposite case where it was present and then slowly turned off. In the monostable regimes, the trajectories are approximately equal, because the stable state is uniquely determined by the stimulus strength. In the bistable regime,

though, this resemblance is broken by the dependency of the state upon the history of the activation variable. For the increasing stimulus, the activation variable is in the *off*-state in the bistable regime, while for the decreasing stimulus, it is in the *on*-state. As the stimulus decreases, the *on*-state loses stability as the attractor in the positive region disappears, the state relaxes to the *off*-state in the now monostable system.

Analyzing the dynamics of the neural activation system with self-excitation given in Equation 11, we have systematically varied the input parameter $s(t)$, while assuming the self-excitation parameter w is fixed at a medium value. What happens if that value is much larger or smaller? For small self-excitation strengths, the system is barely different from the case without self-

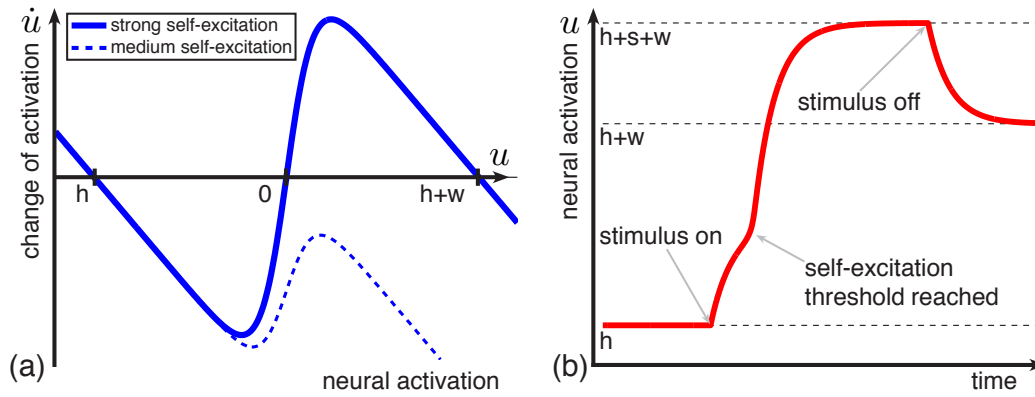


Figure 10: Dynamics of a single neuron with strong self-excitation (Equation 12). Panel (a) shows the phase plot (solid) without external stimulus, for comparison the phase plot of the same system with medium self-excitation is repeated (dashed). Panel (b) shows an example trajectory for a case in which the stimulus is turned on and off again.




Figure 11: Connectivity diagram for two dynamic neurons with mutual inhibition receiving the same stimulus (Equation 13).

excitation, which we analyzed in detail above. For strong self-excitation, the system behaves differently. Figure 10a shows the phase plot for the system with strong self-excitation with and without a stimulus. When a stimulus is present, the system has one attractor in the positive region of the activation variable, where the neuron is active. If no stimulus is present, though, the picture is changed: the self-excitation alone is sufficiently strong to create a second zero crossing of the phase plot in the positive region of u , making the system bistable with one *on*-state and one *off*-state.

Even in the absence of input, this system with strong self-excitation has a stable state in which the neuron is active. The implications of this are shown by the example trajectory given in Figure 10b. As before, the neuron is inactive at first, then the stimulus is turned on for a short while, activating it. Note the sharp increase in the slope of the trajectory as it crosses the zero-line: this is where self-excitation kicks in, adding a large term to the rate of change of u , so it rises faster than before. When the stimulus is removed, the attractor in the positive region does not disappear – the strong self-excitation is sufficient to keep the neural activation positive even without the presence of

a stimulus. This effect can be seen as a basic form of memory, the neural activation generated by a stimulus that was there in the past but not any longer. This is another, even stronger example of a dynamic neural system that is not purely input driven but capable of sustaining stable representations of stimuli that are no longer present.

7 Connected dynamic neurons

In the last section we introduced self-excitation as the simplest form of feedback. By adding this connection, single activation variables can already model systems that are capable of making rudimentary decisions driven by a combination of internal and external states. In this section, we make another step towards a connected system of dynamic neural activation variables by taking the simple case of two connected neurons and analyzing the patterns generated by this system in detail.

A connection between two neurons is always directed, the activation of one neuron is used as a stimulus for the second neuron. As before, the stimulus strength is the activation of the neuron transformed by a sigmoid and multiplied by a parameter representing the connection strength. This system is represented by

$$\begin{aligned} \tau_1 \dot{u}_1(t) &= -u_1(t) + h + s(t) + w_{12} \sigma(u_2(t)) + \xi_1, \\ \tau_2 \dot{u}_2(t) &= -u_2(t) + h + s(t) + w_{21} \sigma(u_1(t)) + \xi_2, \end{aligned} \quad (13)$$

where we use two activation variables u_1 and u_2 receiving the same stimulus s . The second term from the right in each equation corresponds to the connection from the other neuron, the indices of the weight parameter w_{12} indicates that u_1 is receiving a stimulus from u_2 with this

strength, and respectively u_2 is receiving input from u_1 with connection strength w_{21} – the index of the receiving neuron comes first, the index of the stimulating neuron second. In the example we discuss, these parameters will be negative, meaning the neurons inhibit each other. We refer to these as *interaction terms*. The rightmost terms ξ_i represent Gaussian white noise.

To understand the dynamics of this system of two neural activation variables with mutual inhibition, we start by looking at one variable alone. For u_1 , excitation coming from u_2 is indistinguishable from excitation by an external stimulus – what matters is the total strength of the input. If we treat this ‘total stimulus strength’ as one parameter, this is the same system as the one given by Equation 11, illustrated in Figure 6. By the detailed analysis above, this system has a single attractor at $h + s_1$, where $s_1 = s(t) + w_{12}\sigma(u_2(t))$ is the total stimulus strength.

The important thing to note here is that the location of the attractor for u_1 depends upon the state of u_2 . Because the connection between the neurons is gated by the sigmoid, there are essentially two distinct cases to consider: the other neuron is *on*, or it is *off*. In the latter case, the interaction term will be approximately zero. If u_2 is *on*, though, the inhibition it generates will cancel out the excitation from the external stimulus, pushing the fixed point of u_1 down below zero. So if u_2 is *on*, the only stable state for u_1 is *off*. As the connections are symmetric, the same holds for u_2 . Because of the strong mutual inhibition between them, u_1 and u_2 cannot be *on* at the same time.

The dynamics of this system are illustrated in Figure 12 by plotting vector fields of the direction and magnitude of the rate of change at possible states (u_1, u_2) of the system. Panel (a) shows the case where no external stimulus is presented. The system has a single attractor in which both neurons are *off* (green dot). If a sufficiently strong external stimulus is given, this picture changes to the case described in the previous paragraph: due to the mutual inhibition, there are two stable states, in each of which one neuron is *on* and the other one is *off*.

What kind of trajectories arise from this system? Consider the usual case where both activation variables are off when a stimulus appears. Figure 13b shows trajectories from this setup for both neurons. After the stimulus is switched on, both activation variables rise at first. In this period, both systems get excitation from the stimulus, but no inhibition from the other neuron yet, corresponding to the dashed line in Figure 13a. As the activation variables approach the threshold of the sigmoid, though, the picture changes: *both* neurons are switching to the *on* state and start to inhibit each other. From the phase plot in Figure 12 we know that this state cannot be stable, though.

How is this symmetric stand-off between the two activation variables resolved? The simple solution is that whoever falters first will lose. A slight difference in activation, say u_2 slightly lower than u_1 , translates into a difference in inhibition of the other neuron. Note that in this transitory region the sigmoid is approximately linear and neurons are neither *on* nor *off*. Less inhibited than u_2 , the activation of u_1 will decrease slower than u_2 , which increases the difference in activation. After a short time, u_2 drops below threshold, reducing the inhibition of u_1 to almost zero. Now with only excitation from the positive external stimulus, u_1 relaxes to its stable *on*-state. This is exactly what happened in the example trajectory shown in Figure 13b. With inhibitive connections of equal strength, which neuron turns out to be active and which passive is essentially down to random chance from noise.

This system of two neurons with mutual inhibition is another example where the state is not purely input-driven. Presenting a stimulus moves the system into a bistable regime, random fluctuations determine which attractor the state variable relaxes to. Once it has relaxed, though, the system dynamics stabilize the state variable at that value against further perturbations. We say that the system makes a *stable selection decision*.

With different parameter settings, this same system of two interconnected neurons can generate oscillatory solutions [42]. This is a simple model of a central pattern generator (CPG), which have been widely used to generate locomotor patterns for legged robots [43].

8 Conclusions

We have presented a detailed analysis of dynamic neural activation variables. For a single neural node, we discussed the stability properties and analyzed the dependency upon the external stimulus parameter. In the case of self-excitation, we identified the bifurcation between the mono-stable and the bistable regimes and discussed how these bifurcations correspond to detection decisions. The fundamental characteristics and properties of networks of neural nodes were demonstrated using the simple example of two interconnected neurons with mutual inhibition and showed how increasing stimulus leads this system through a bifurcation from one to two stable states, the choice between which corresponds to a selection decision. These features provide the elementary building blocks to design cognitive architectures for robotic systems that connect the available sensory systems to the effectors with which the agent acts upon the world.

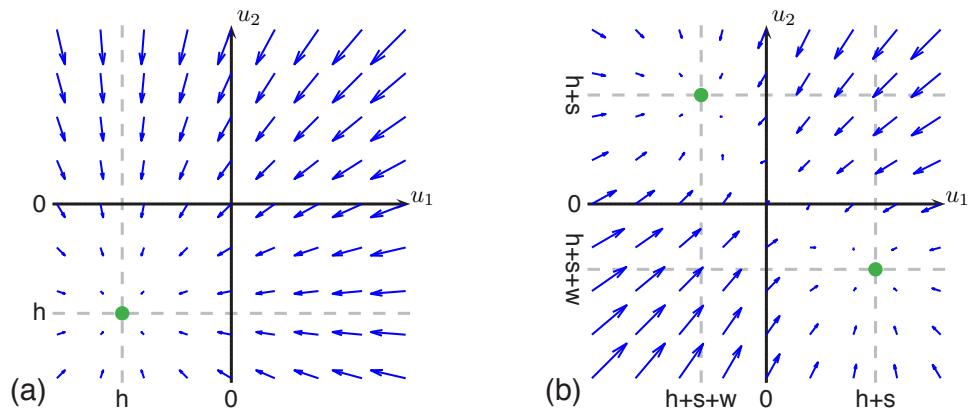


Figure 12: Phase plot of two activation variables in a two-dimensional state space. The vectors indicate the direction and magnitude of the rate of change at the state where they start. Panel (a) shows the case without external stimulus with one fixed point attractor (green dot), while in panel (b) the external stimulus is active, resulting in two fixed point attractors.

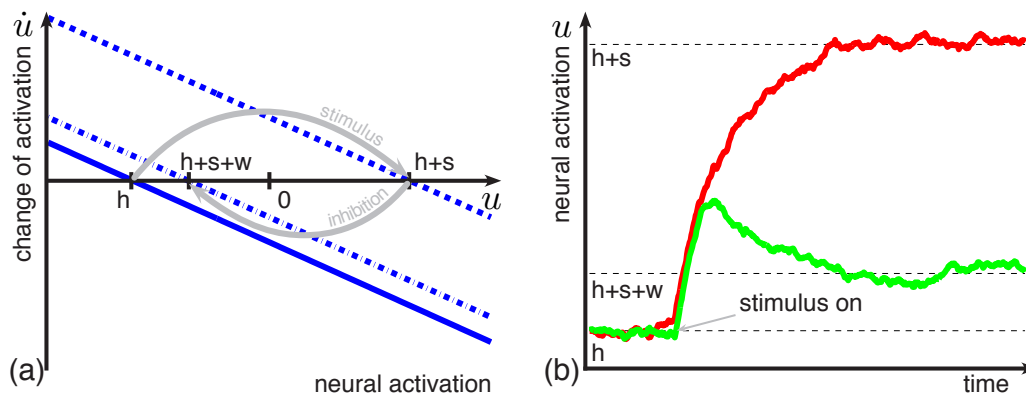


Figure 13: Panel (a) shows example phase plots of one out of two neurons with mutual inhibition (Equation 13) for three different parameter settings. With no stimulus and no inhibition from the other neuron (solid line) the attractor is at the resting level. Presenting a stimulus (dashed line) shifts the attractor to the positive region. Inhibition from the other neuron (dash-dotted line) shifts it back to the negative region again, as indicated by the grey arrows. Example trajectories for both neurons are shown in panel (b).

References

- [1] A. Clark and R. Grush, "Towards a Cognitive Robotics," *Adaptive Behavior*, vol. 7, pp. 5–16, Jan. 1999.
- [2] L. W. Barsalou, "Grounded cognition," *Annual Review of Psychology*, vol. 59, no. 1, pp. 617–645, 2008.
- [3] G. Schöner, "Dynamical Systems Approaches to Cognition," in *The Cambridge Handbook of Computational Psychology* (R. Sun, ed.), Cambridge University Press, 2008.
- [4] Y. Sandamirskaya and G. Schöner, "Serial order in an acting system: a multidimensional dynamic neural fields implementation," in *International Conference on Development and Learning*, 2010.
- [5] M. Richter, "A robotic architecture for action selection and behavioral organization inspired by human cognition," *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2012.
- [6] C. Faubel and G. Schöner, "Learning to recognize objects on the fly: a neurally based dynamic field approach," *Neural networks: the Official Journal of the International Neural Network Society*, vol. 21, pp. 562–76, May 2008.
- [7] S. K. U. Zibner, C. Faubel, I. Iossifidis, and G. Schöner, "Dynamic neural fields as building blocks of a cortex-inspired architecture for robotic scene representation," *IEEE Transactions on Autonomous Mental Development*, vol. 3, no. 1, pp. 74–91, 2011.
- [8] J. J. Eggermont, "Is there a neural code?," *Neuroscience and Biobehavioral Reviews*, vol. 22, no. 2, pp. 355–370, 1998.
- [9] M. N. Shadlen and W. T. Newsome, "Noise, neural codes and cortical organization," *Current opinion in neurobiology*, vol. 4, no. 4, pp. 569–579, 1994.
- [10] M. N. Shadlen and W. T. Newsome, "The variable discharge of cortical neurons: implications for connectivity, computation, and information coding," *The Journal of Neuroscience*, vol. 18, no. 10, pp. 3870–3896, 1998.
- [11] A. A. Faisal, L. P. J. Selen, and D. M. Wolpert, "Noise in the nervous system," *Nature Reviews Neuroscience*, vol. 9, no. 4,

- pp. 292–303, 2008.
- [12] C. E. Carr and M. Konishi, “A circuit for detection of interaural time differences in the brain stem of the barn owl,” *The Journal of Neuroscience*, vol. 10, no. 10, pp. 3227–3246, 1990.
- [13] M. N. Shadlen, “Rate versus Temporal Coding Models,” in *Encyclopedia of Cognitive Science*, John Wiley & Sons, Ltd, 2006.
- [14] C. S. Sherrington, *The integrative action of the nervous system*. New Haven: Yale University Press, 1906.
- [15] J. P. Jones and L. A. Palmer, “The two-dimensional spatial structure of simple receptive fields in cat striate cortex,” *Journal of Neurophysiology*, vol. 58, no. 6, pp. 1187–1211, 1987.
- [16] A. P. Georgopoulos, J. F. Kalaska, R. Caminiti, and J. T. Massey, “On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex,” *Journal of Neuroscience*, vol. 2, no. 11, pp. 1527–1537, 1982.
- [17] A. B. Schwartz, R. E. Kettner, and A. P. Georgopoulos, “Primate motor cortex and free arm movements to visual targets in three-dimensional space. I. Relations between single cell discharge and direction of movement,” *The Journal of neuroscience : the official journal of the Society for Neuroscience*, vol. 8, no. 8, pp. 2913–2927, 1988.
- [18] C. Lee, W. H. Rohrer, and D. L. Sparks, “Population coding of saccadic eye movements by neurons in the superior colliculus,” *Nature*, vol. 332, no. 6162, pp. 357–360, 1988.
- [19] D. H. Hubel and T. N. Wiesel, “Receptive fields of single neurones in the cat’s striate cortex.,” *The Journal of Physiology*, vol. 148, pp. 574–591, 1959.
- [20] D. H. Hubel and T. N. Wiesel, “Receptive fields and functional architecture of monkey striate cortex,” *The Journal of Physiology*, vol. 195, no. 1, pp. 215–243, 1968.
- [21] B. R. Conway and D. Y. Tsao, “Color-tuned neurons are spatially clustered according to color preference within alert macaque posterior inferior temporal cortex,” *Proceedings of the National Academy of Sciences*, vol. 106, no. 42, pp. 18034–18039, 2009.
- [22] A. Pasupathy and C. E. Connor, “Shape representation in area V4: position-specific tuning for boundary conformation,” *Journal of Neurophysiology*, vol. 86, no. 5, pp. 2505–2519, 2001.
- [23] K. H. Britten and W. T. Newsome, “Tuning bandwidths for near-threshold stimuli in area MT,” *Journal of neurophysiology*, vol. 80, no. 2, pp. 762–770, 1998.
- [24] D. Bendor and X. Wang, “The neuronal representation of pitch in primate auditory cortex,” *Nature*, vol. 436, no. 7054, pp. 1161–1165, 2005.
- [25] P. J. Fitzgerald, “Receptive field properties of the macaque second somatosensory cortex: Representation of orientation on different finger pads,” *Journal of Neuroscience*, vol. 26, no. 24, pp. 6473–6484, 2006.
- [26] D. H. Hubel and T. N. Wiesel, “Receptive fields, binocular interaction and functional architecture in the cat’s visual cortex,” *The Journal of Physiology*, vol. 160, pp. 106–154, 1962.
- [27] E. D. Adrian and Y. Zotterman, “The impulses produced by sensory nerve endings: Part 3. impulses set up by touch and pressure,” *Journal of Physiology*, vol. 61, no. 4, pp. 465–483, 1926.
- [28] R. S. Person and L. P. Kudina, “Discharge frequency and discharge pattern of human motor units during voluntary contraction of muscle,” *Electroencephalography and Clinical Neurophysiology*, vol. 32, no. 5, pp. 471–483, 1972.
- [29] A. B. Schwartz and D. W. Moran, “Arm trajectory and representation of movement processing in motor cortical activity,” *The European Journal of Neuroscience*, vol. 12, no. 6, pp. 1851–1856, 2000.
- [30] G. Sclar and R. D. Freeman, “Orientation selectivity in the cat’s striate cortex is invariant with stimulus contrast,” *Experimental Brain Research*, vol. 46, no. 3, pp. 457–461, 1982.
- [31] R. Erickson, “Parallel ”Population” Neural Coding in Feature Extraction,” in *The Neurosciences* (F. Schmitt and F. Worden, eds.), pp. 155–169, Cambridge, Massachusetts: MIT Press, 1974.
- [32] D. H. Hubel and T. N. Wiesel, “Ferrier lecture. Functional architecture of macaque monkey visual cortex,” *Proceedings of the Royal Society of London. Series B, Containing papers of a Biological character. Royal Society (Great Britain)*, vol. 198, no. 1130, pp. 1–59, 1977.
- [33] R. B. Tootell, M. S. Silverman, E. Switkes, and R. L. d. Valois, “Deoxyglucose analysis of retinotopic organization in primate striate cortex,” *Science*, vol. 218, no. 4575, pp. 902–904, 1982.
- [34] G. G. Blasdel and G. Salama, “Voltage-sensitive dyes reveal a modular organization in monkey striate cortex,” *Nature*, vol. 321, pp. 579–585, 1986.
- [35] D. H. Hubel and T. N. Wiesel, “Uniformity of monkey striate cortex: a parallel relationship between field size, scatter, and magnification factor,” *The Journal of Comparative Neurology*, vol. 158, no. 3, pp. 295–305, 1974.
- [36] F. Rosenblatt, “The perceptron: a probabilistic model for information storage and organization in the brain.,” *Psychological Review*, vol. 65, no. 6, pp. 386–408, 1958.
- [37] S. Haykin, *Neural Networks and Learning Machines*. Pearson Education, 2011.
- [38] M. Minsky and S. Papert, *Perceptrons: An Introduction to Computational Geometry*. MIT Press, 1988.
- [39] J. Piaget, *The Construction of Reality in the Child*. New York: Basic, 1954.
- [40] V. Braitenberg, “Cortical architechtonics: General and areal,” in *IBRO Monograph Series, Architechtonics of the Cerebral Cortex* (M. A. B. Brazier and H. Petsche, eds.), vol. 3, pp. 443–466, 1978.
- [41] L. Perko, *Differential Equations and Dynamical Systems*. Heidelberg: Springer, 1991.
- [42] S.-i. Amari, “Dynamics of Pattern Formation in Lateral-Inhibition Type Neural Fields,” *Biological Cybernetics*, vol. 27, pp. 77–87, 1977.
- [43] A. J. Ijspeert, “Central pattern generators for locomotion control in animals and robots: a review.,” *Neural networks : the official journal of the International Neural Network Society*, vol. 21, pp. 642–53, May 2008.