

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

A Conceptual Ladder from Spikes to Behavior: Toward the Neural Basis of Dynamic Choices at Multiple Scales

Permalink

<https://escholarship.org/uc/item/3st5161c>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 40(0)

Authors

Shea, Timothy M

Noelle, David C

Publication Date

2018

A Conceptual Ladder from Spikes to Behavior: Toward the Neural Basis of Dynamic Choices at Multiple Scales

Timothy M. Shea (tshea2@ucmerced.edu)

David C. Noelle (dnoelle@ucmerced.edu)

Cognitive and Information Sciences, University of California, Merced
5200 Lake Road, Merced, CA 95343 USA

Abstract

Reducing cognitive phenomena to neural activity is seen by many as lacking in scientific utility. The conceptual chasm between electrochemical activity and the act of making a choice is too broad to span in a single step. Instead, we adopt a multi-scale approach to cognitive neuroscience by constructing a conceptual ladder that incrementally climbs from neuronal spikes to cognitive processes with each step offering theoretic reductions. Here we propose a sequence of intermediate neurocomputational processes that are promising for understanding an array of cognitive phenomena. We illustrate this approach in the context of the dynamics of choice. These dynamics emerge from serial evaluation mediated by systems in frontal cortex and the basal ganglia. The effect is to promote neural oscillations that provide a substrate for communication through coherence. Both empirical and simulation studies are described to support this view of emergent behavior.

Keywords: spiking neural networks; synchrony; coherence

Introduction

In this paper, we construct a conceptual ladder upon which we climb from the foundations of neuroscience -- spiking neurons -- to the dynamical patterns of brain activity and behavior that arise during decision making. In spite of broad acceptance of the general idea that, in some way, high level cognitive processes are grounded in the electrical and chemical processes of neurons, there are relatively few instances in which cognition has been clearly mapped onto cellular actions. In part, this may be due to the tremendous difference in scale: the human hand is a thousand times larger than a neuron. Pushing a button takes hundreds of times longer than sending an electrical impulse down an axon. Nevertheless, we describe here an account which ascends these spatial and temporal scales, one rung at a time, by drawing on emerging perspectives at various levels of computational and empirical neuroscience, as well as modern behavioral research. We review some literature related to each of these levels of analysis and discuss our own computational simulations of relevant phenomena.

While this conceptual ladder may be a generally useful framing for a broad array of cognitive processes, we focus on a category of behavioral experimental paradigms which target value-based decision making. A simple example is a two-armed bandit task, in which a participant selects one of two actions on each trial, guided by the predicted reward for each action. These kinds of choices can be made reflexively

and can be learned by non-human primates and rodents. They arguably offer a clear and simple model of the broader problem of action selection. There is also extensive research on the contributions of brain waves, brain regions, cell types, neurotransmitters, and various other biological factors to these types decision tasks. This large body of research provides multi-scale perspectives on the neural basis of value-based decision making and provides constraints on any proposed account of choice.

Our conceptual ladder starts with spiking neural networks. We explain how synchronization in spiking neural networks offers a general mechanism to understand the emergence of oscillations. We discuss how oscillations modulate regional processing in the brain. We argue that the interaction of oscillations provides a useful active substrate for neural communication. Lastly, we describe how these properties can give rise to characteristic patterns of brain activity and behavior observed in decision making.

Spiking Neural Networks Synchronize

One of the first lessons learned by every scientist who attempts to build a spiking neural network model is that spiking networks synchronize (Brunel, 2000). Synchronous spiking is the generation of a large fraction of spikes during narrow, semi-regular temporal bands, separated by relative silence. This fact often goes unmentioned in the literature, perhaps because it is so basic and pervasive to the approach, but it is virtually assured that every successful spiking neural network model has undergone a period of tuning to avoid excessive synchrony (c.f. Hahn, et al., 2014). Excessive synchrony is undesirable, as it reduces the representational capacity of a spiking network and impairs spike timing dependent plasticity (STDP). It also leads to visibly artificial spike raster plots that fail to capture observed biological spiking patterns.

A common cause of synchronous spiking is recurrent synaptic connectivity associated with a "reservoir" (Maass, Natschläger, & Markram, 2002) or cortical layers (Tiesinga & Sejnowski, 2009). In these cases, feedback loops from recurrent connectivity drive excitatory spiking interspersed with inhibitory silencing. Without careful tuning, this feedback-induced synchrony can overpower activity related to meaningful stimuli in the network. Nevertheless, recurrent connections in biological neuronal networks are ubiquitous.

One of the longest running debates in computational neuroscience involves the role of spike timing in neuronal

information processing. Although it is widely agreed that information is conveyed by neural spikes in aggregate, there is disagreement over whether that information is encoded in the precise timing of spikes or only in the rate at which spikes are generated. In some specific cases, highly temporally precise sensory information that can be decoded from spike times (e.g. in binaural spatial cues, Chase & Young, 2006) beyond that which can be extracted from firing rates. We suggest that emerging theories that emphasize the role of spike synchrony in robust neuronal communication offer a potential route to resolving these tensions.

Spike Synchrony Links Spikes to Oscillations

The excitability of neurons in a brain region is correlated with the phase of ongoing neural oscillations in that region (Fries, 2005). Local field potentials (LFP) are generated by fluctuations in membrane potential across many thousands of neurons in a region (Gold et al., 2006). These fluctuations are the result of various processes that occur at different timescales. While spike generation is a distinct feature of individual neurons, LFPs deemphasize high-frequency fluctuations, at the timescale of spikes, so that only very dense concentrations of spikes affect the LFP amplitude. For this reason, synchronous spiking across a neural population will give rise to an ongoing LFP oscillation, while irregular spiking activity will not. Alternative oscillatory mechanisms are certainly plausible (e.g. endogenously generated oscillatory currents, Brunel, Hakim, & Richardson, 2003), but synchronous spiking offers a parsimonious account for the emergence of oscillations while simultaneously explaining why LFP oscillation phase is correlated with neuronal excitability. The patterns of connectivity which give rise to synchrony result in periods of inhibition (low excitability) and disinhibition (high excitability) and these periods will necessarily be phase-locked to resulting oscillations. Thus, spike synchrony provides a functional link from the activity of individual neurons to neural oscillations.

An important consideration in the emergence of oscillations from synchronous spiking is the role of neuronal and network parameters in shaping the oscillations which emerge. Properties such as neuron leak conductance and excitatory-inhibitory ratio of the network directly affect the gain in recurrent synaptic pathways. Both the mean and variance in spike transmission delays alter the length of feedback loops. These and other structural properties determine the frequency, phase, and amplitude of intrinsic neural oscillations (**Figure 1**).




Figure 1: This spike raster plot demonstrates synchrony in a simulated spiking neural network as a result of inhibitory recurrence. Each row represents a neuron and each dot represents a spike at a given time. The blue line indicates the average firing rate of the population in time.

Neural Oscillations Provide an Active Substrate for Communication

As described previously, the phase of neural oscillations is correlated with neuronal excitability. These findings have led to an emerging consensus concerning the functional role of neural oscillations. This consensus holds that the effectiveness of communication between brain regions is modulated by the coherence of oscillations in the regions. This hypothesis is called communication-through-coherence (CTC) (Fries, 2005).

The CTC hypothesis makes clear predictions concerning neuronal communication in the context of a single dominant frequency of oscillation (**Figure 2**). However, biological neural oscillations typically exhibit power in many distinct frequency bands. In this case, it is unclear which frequency might have a dominant effect on neuronal excitability. One potential solution to this problem involves a mechanism through which multiple oscillations are entrained, such that the phase relations between regions are consistent across multiple frequency bands. Indeed, just such entrainment has been observed. Cross-frequency-coupling (CFC) describes the phenomena where ongoing neural oscillations in one frequency are coupled to oscillations in a different frequency. CFC is particularly strong between alpha-band (8-12 Hz) phase and gamma-band (>30 Hz) amplitude (phase-amplitude coupling). This means that the strength of the high frequency component of neural oscillations in a region is modulated by the phase of the low frequency component. The degree of phase-amplitude coupling in a region has been linked to the strength of intra-regional communication versus inter-regional communication. Other forms of CFC include phase-phase coupling and amplitude-amplitude coupling (Canolty & Knight, 2010).




Figure 2: Communication-through-coherence between two regions with the same dominant frequency. In this scenario, the effectiveness of communication (likelihood of a transmitted spike to induce a spike at the receiving neuron) is predominantly affected by the relative phase and spike transmission delays. The sending population (**blue**) exhibits a strong influence on neurons in the receiving population (**orange**) with a compatible phase and has minimal impact on neurons in the inhibitory period of an oscillation. Receiving neurons are shown with staggered phase to illustrate the effect of CTC.

Another important clue in deciphering the functional role of oscillations comes from the spatial scales over which neural communication occurs. Over long neurophysiological distances, spike transmission delays and variability both increase. These increases correspond to a wider temporal window in which postsynaptic potentials might arrive. To accommodate this, the frequency of neural oscillation must be lower, producing increased durations of the periods of neuronal excitability. In contrast, oscillations at a high frequency (e.g. gamma-band), with more brief periods of neuronal excitability, provide a substrate for communication over short neurophysiological distances. These predictions are consistent with the behavioral correlates of neural oscillations (**Table 1**).

Table 1: Behavioral correlates of neural oscillations (Canolty & Knight, 2010).

	Frequency	Correlates
Gamma	> 30 Hz	Local network processing
Beta	20 - 25 Hz	Motor planning, rhythm
Alpha	8 - 12 Hz	Visual attention
Theta	4 - 8 Hz	Conflict detection
Delta	< 4 Hz	Multisensory integration

Together, CTC and CFC suggest a framework in which communication across neuronal networks is modulated by a broad range of interacting oscillations at various spatial scales. These interacting oscillations provide an *active substrate* for communication: a constantly shifting network

of channels vying for the available bandwidth. Rather than a central coordinator allocating influence to streams of neural processing, the streams dynamically entrain and diverge from one another according to learned task demands and new inputs. As the relative power of oscillation components fluctuate, and as they shift in and out of relative phase, the effectiveness of particular channels of communication increase and decrease. Because the oscillations emerge from spiking activity, neural spikes drive these shifts in phase and power. The detailed relationship between spikes and waves is highly complex, however, making oscillation coherence a useful intermediate level of analysis.

This framework of activity-driven modulation of neural communication offers a plausible mechanism for variation in functional connectivity between larger neural regions (Akam & Kullman, 2012; Salinas & Sejnowski, 2001). A vast literature based on fMRI, EEG, and other imaging methodologies has demonstrated that functional changes in coupling between regions can occur rapidly in response to task and context cues (Bullmore & Sporns, 2009). While functional connectivity between regions can reflect common inputs, it is often taken to imply selective communication. These observations have been made without knowledge of the mechanism by which neural activity might be selectively processed or ignored. Given that the coherence of oscillations among regions is rapidly altered by interacting neuronal activity, CTC and CFC could underlie task- and context-dependent changes in functional connectivity throughout the brain.

Dynamics of Choice Arise from Properties of the Communication Substrate

Perhaps counterintuitively, coherent oscillations supporting communication between regions do not necessarily improve information processing between those regions and produce associated facility in behavior. It is natural to expect greater amplitudes in task-linked oscillations to be associated with better communication and, thus, better task performance. Almost paradoxically, though, high amplitude oscillations are regularly associated with a lack of strong information processing. For example, alpha amplitude increases in visual areas when the eyes are closed and there is (apparently) no visual attention. Also, conflict-related high amplitude theta waves in medial frontal and subcortical regions is linked to slow responses and poor performance (Zavala, Zaghoul, & Brown, 2015).

In a computational model of the role of thalamocortical loops in signal selection (Shea, Rodny, Warlaumont, & Noelle, 2017), we found that this inverse relationship can emerge from a trade-off between cognitive stability and sensory signal integration. During periods of low conflict, such as when a sensory stimulus is clearly distinguished from its background, the strong stimulus input suppressed background noise, leading to a stable selection of the signal. During periods of high conflict, such as when the stimulus was not clearly distinct from background noise, competition caused increased oscillations, promoting greater

communication. However, if the increased oscillations led to conflict being quickly resolved, the oscillations then rapidly faded. If conflict persisted, the oscillations remained strong as the network failed to pull the signal from the noise. Thus, in our model, the strength of oscillations is inversely related to the neuronal network's performance.




Figure 3: Architecture of a spiking neural network model of thalamocortical signal selection. The network is based on recurrent synaptic circuits which connect the thalamus, cortex, and basal ganglia. Arrow-heads denote excitatory pathways, circles denote inhibitory pathways. STIM = stimulus, THL = thalamus, CTX = cortex, DA1 = Striatum (Da1-Receptors), DA2 = Striatum (Da2-Receptors), STN = Subthalamic Nucleus, GPI = Globus Pallidus Internal, GPE = Globus Pallidus External (Shea, Rodny, Warlaumont, & Noelle, 2017).

The model contained seven populations of leaky integrate-and-fire spiking neurons and eleven synaptic pathways comprising a recurrent thalamo-cortico-basal ganglia circuit (Figure 3). Shaped input currents were injected at the model thalamus to simulate a stimulus, and these signal inputs were overlaid with random background noise. Thalamic activity propagated forward through the cortex and basal ganglia, where an off-center on-surround inhibitory dynamic gave rise to neuronal competition. For weak stimuli, the resulting competition was unstable, and the excitatory-inhibitory recurrence of the network caused neurons in every region to synchronize. For stronger stimuli, suppression of background noise kept recurrent feedback under control, allowing a stable representation of the signal to emerge in the cortex (Figure 4).

The task performed by our model is simple: a signal, represented as increased thalamic activity in a region of an abstract neural map, is obscured by noise. The network attempts to amplify the signal by suppressing the noise. This reflects a single choice task: a Go task as opposed to a Go/No-Go task. Although this task seems simpler than many animal behaviors, very complex decisions can be mapped onto such a signal selection mechanism through a combination of affordance competition and serial evaluation.

Many neural network models of two-alternative forced choice tasks represent relevant features of the alternatives simultaneously in two different populations of neurons, with




Figure 4: During each trial, the model is presented with a noisy bump-attractor stimulus for 2s. In trials with high stimulus signal-to-noise ratio (SnR) (4 - 10 s above), the network selects the stimulus and inhibits background noise, yielding a high evoked SnR. In trials with low stimulus SnR (0 - 4 s above), competition between signal and noise leads to the emergence of oscillations. The local field potential (LFP) is modeled as the leak, excitatory, and inhibitory currents of all neurons in a population inversely scaled by the distance to a fixed point in the population.

choices also represented separately and in parallel. The neurons representing the alternatives then compete through recurrent projections, producing some form of a “winner take all” dynamic (Usher & McClelland, 2001). This approach is problematic when considering how it might scale up to situations in which there are many more than two options. It is unlikely that the brain is equipped with parallel neural pathways for each possible choice in a large and novel set, with inhibition appropriately configured to cause the alternatives to compete when a choice is presented.

A different account of choice involves serial evaluation. Alternatives are represented and entertained one at a time, with relative value information represented, and there is a process of shifting from one alternative to the next, with potentially many repeated considerations of a given choice. Rather than making a parallel forced-choice, there is only an “accept-reject” decision for the currently represented choice, with acceptance resulting in the selection of that alternative and rejection entailing continued serial consideration. This account scales to multiple alternatives and is consistent with the proposal that decision making capabilities of this kind initially evolved in the context of foraging, where potential choices (e.g. bushes that might contain good berries) are considered serially, resulting in repeated “accept-reject” decisions (Hayden, 2018).

There is increasing evidence that serial evaluation is used in the brain. For example, Rich & Wallis (2016) recorded ensemble spiking activity and local field potentials from the orbitofrontal cortex, finding neurons representing the value of options in a 2-alternative forced choice task, alternatively encoding the value of one choice and then the other.

If input signals in our computational model are seen as encoding value, perhaps with noise capturing a background value standard arising from previously considered options, then successful signal selection can be an “accept” choice, while continued strong oscillations could drive a “reject”, prompting input of another potential choice.

Our computational model of choice is incomplete. Any form of attentional switching -- triggered by oscillations or other factors -- is exogenous to the network. The model illustrates how recurrent communication between cortical and midbrain decision structures can support both signal selection, and can promote oscillations necessary for neuronal communication. Thus, prolonged “deliberation” in such a circuit will give rise to more effective spike propagation across a broader region of neural tissue.

Discussion

The problem of grounding cognitive processes in the activity of the brain is at the core of cognitive science, yet the mechanisms of brain function which are well understood operate at a finer scale than what is relevant to behavior. In investigating the neural mechanisms of serial decision-making, we have identified important abstractions, at different scales, with which we can climb from spikes to choices. We have proposed that synchrony is a fundamental

property of neuronal networks and argued that synchrony is the primary mechanism of neural oscillation. There are counter-arguments to this proposal, yet there is no alternative which enjoys as much empirical support or as robustly explains the bidirectional interactions of spikes and brain waves. We also reviewed evidence that neural oscillations which emerge from spike synchrony provide a flexible, dynamic mechanism for selective communication across spatial and temporal scales: an *active substrate*. We then explored how this communication across recurrently connected brain regions may underlie some of the cognitive processes of choice.

Mental and behavioral phenomena occurring at the level of human experience find their causes not only in brain states, but in the interactions of brains, bodies, and environments. Specific mental states, it might be argued, emerge from these interactions. Much of the modern understanding of neuronal activity has been developed at the scale of individual cells, whereas it is likely that qualitatively different properties emerge from brain-scale collections of neurons. These possibilities invite criticism of a pure reductionist account of cognition. Nevertheless, we propose that our conceptual ladder, providing a multi-scale perspective on neural phenomena, can help bridge the gulf of understanding between brain and behavior. In our model of signal selection, we have demonstrated that such a reductive approach can provide insight.

Finally, we recognize that the precise mechanisms of any decision task will certainly involve processes beyond those discussed here. Nevertheless, much cognitive neuroscience research operates solely at one scale, and we hope that this limited exploration of choice highlights the benefits of work at multiple scales. By tying the mechanisms of neuronal activity to those of oscillatory interactions to those of large-scale functional activity, we can begin to offer neural accounts of behavior which extend beyond a single level of abstraction. In ongoing work, we plan to explore these possibilities.

Acknowledgments

We gratefully acknowledge the advice and support of the Cognitive & Information Sciences Graduate Group at UC Merced.

References

- Akam, T. E., & Kullmann, D. M. (2012). Efficient “communication through coherence” requires oscillations structured to minimize interference between signals. *PLoS Computational Biology*, 8(11), e1002760.
- Brunel, N. (2000). Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *Journal of Computational Neuroscience*, 8(3), 183-208.
- Brunel, N., Hakim, V., & Richardson, M. J. (2003). Firing-rate resonance in a generalized integrate-and-fire neuron with subthreshold resonance. *Physical Review E*, 67(5), 051916.

- Bullmore, E., & Sporns, O. (2009). Complex brain networks: graph theoretical analysis of structural and functional systems. *Nature Reviews Neuroscience*, *10*(3), 186.
- Canolty, R. T., & Knight, R. T. (2010). The functional role of cross-frequency coupling. *Trends in Cognitive Sciences*, *14*(11), 506-515.
- Chase, S. M., & Young, E. D. (2006). Spike-timing codes enhance the representation of multiple simultaneous sound-localization cues in the inferior colliculus. *Journal of Neuroscience*, *26*(15), 3889-3898.
- Fries, P. (2005). A mechanism for cognitive dynamics: neuronal communication through neuronal coherence. *Trends in Cognitive Sciences*, *9*(10), 474-480.
- Gold, C., Henze, D. A., Koch, C., & Buzsaki, G. (2006). On the origin of the extracellular action potential waveform: a modeling study. *Journal of Neurophysiology*, *95*(5), 3113-3128.
- Hahn, G., Bujan, A. F., Frégnac, Y., Aertsen, A., & Kumar, A. (2014). Communication through resonance in spiking neuronal networks. *PLoS Computational Biology*, *10*(8), e1003811.
- Hayden, B. Y. (2018). Economic choice: the foraging perspective. *Current Opinion in Behavioral Sciences*, *24*, 1-6.
- Maass, W., Natschläger, T., & Markram, H. (2002). Real-time computing without stable states: A new framework for neural computation based on perturbations. *Neural Computation*, *14*(11), 2531-2560.
- Rich, E. L., & Wallis, J. D. (2016). Decoding subjective decisions from orbitofrontal cortex. *Nature Neuroscience*, *19*(7), 973.
- Salinas, E., & Sejnowski, T. J. (2001). Correlated neuronal activity and the flow of neural information. *Nature Reviews Neuroscience*, *2*(8), 539.
- Shea, T. S., Rodny, J. J., Warlaumont, A. S., & Noelle, D. C. (2017). Emergent oscillations in the subthalamic nucleus of simulated basal ganglia. *Neuroscience 2017*. Poster.
- Tiesinga, P., & Sejnowski, T. J. (2009). Cortical enlightenment: are attentional gamma oscillations driven by ING or PING?. *Neuron*, *63*(6), 727-732.
- Usher, M., & McClelland, J. L. (2001). The time course of perceptual choice: the leaky, competing accumulator model. *Psychological Review*, *108*(3), 550.
- Zavala, B., Zaghloul, K., & Brown, P. (2015). The subthalamic nucleus, oscillations, and conflict. *Movement Disorders*, *30*(3), 328-338.