Implications of Polychronous Neuronal Groups for the Nature of Mental Representations

William Benjamin St. Clair (wst.clair@ucmerced.edu)
David C. Noelle (dnoelle@ucmerced.edu)
Cognitive and Information Sciences
University of California, Merced
5200 North Lake Road; Merced, CA 95343 USA

Abstract

How are concepts represented in the human mind? Vector space representations based on the instantaneous firing rates of neurons have been used with great success. However, there is growing evidence, both empirical and computational, that relevant information is encoded in spatiotemporal patterns of spikes called *polychronous neuronal groups* (PNGs). In this paper, we consider the philosophical implications of PNG representations with regard to their temporal extension, grounding, compositionality, and similarity. We suggest that the temporally extended nature of PNGs implies that conceptual-level dynamics may only be coherent at coarse time scales. We introduce the notion of PNG trigger sets as a way to ground the meaning of PNG representations, and we discuss potential approaches to compositionality. Finally, we identify the open problem of how to define an appropriate similarity metric for PNG-based mental representations.

Keywords: Philosophy of Cognitive Science; Neuroscience; Representation; Dynamical Systems.

Introduction

How are concepts represented in the human mind? One highly productive approach to this question has involved the application of a continuous dynamical systems perspective to the problem (Spivey, 2008). From this perspective, the currently active concepts of a cognitive system (or subsystem) are jointly encoded as a point in a high dimensional vector space (Churchland, 1989). Nearby points in this space, according to some distance metric, are seen as representing similar conceptual states, allowing regions and manifolds within this space to capture more general concepts or categories (Gärdenfors, 2000). The evolution of mental states over time becomes a trajectory in this vector space (Yoshimi, 2012), driven by mechanistic cognitive processes (Churchland, 1996).

This *vector space* approach to conceptual representation has many strengths. It supports accounts of the biological basis of cognition by viewing the dimensions of the vector space as the activity of neural units, such as their firing rates, offering a framework for mapping from the physical state space of the brain to the conceptual state space of the mind (Spivey, 2008). The approach provides a straightforward way to discuss the cognitive state of a system at any instant in time, as well as how those states change over time. Issues surrounding the grounding of representations are well addressed from the perspective of *conceptual role semantics* (Greenberg & Harman, 2006), with the "meaning" of a representation being a function of the inputs that activate it and the other representations that it produces through causal relationships, eventually making causal contact with the world through sensory and motor processes. When using the vector space approach, these cognitive processes are well described in neurocomputational terms, with nearby points in a vector space tending to produce similar results when presented to models of downstream neural circuits. While the vector space approach has been criticized as lacking support for compositional and structured conceptual representations (Fodor & Pylyshyn, 1988), and there continues to be extensive work on addressing this critique (Gayler & Levy, 2011), highly promising approaches to compositionality have been proposed, making use of vector space operations of superposition, convolution, and sparse coding (Smolensky, 1990; Plate, 2003; OReilly, Bhattacharyya, Howard, & Ketz, 2011). In general, the vector space approach to conceptual representation has been very productive.

Past challenges to the vector space approach have come from above: from more abstract and symbolic characterizations of cognitive content and cognitive processing. More recently, a challenge has arisen from below: from insights into the neural coding of information (Rieke, Warland, Steveninck, & Bialek, 1999). There is increasing empirical evidence that, in at least some neural systems, relevant information is encoded in the spatiotemporal pattern of discrete action potentials, or spikes, produced by neurons in a given nucleus (Rolston, Wagenaar, & Potter, 2007; Madhavan, Chao, & Potter, 2007; Pasquale et al., 2008). While information, in some cases, may be carried by synchronous or coherent firing of neurons (Fries et al., 2005), computational considerations have suggested that content may frequently be encoded in complex asynchronous patterns of spikes (Izhikevich, 2006). These complex spike patterns have been called *polychronous neuronal groups* (PNGs). The PNG approach to representation differs substantially from the vector space approach. A PNG is a temporally extended pattern of discrete spiking events over a collection of neurons, and it is not clear how such a pattern could be mapped to a point in a continuous vector space so as to preserve relevant aspects of similarity between representations. In the vector space approach, temporally extended trajectories capture dynamic changes in cognitive content, while the PNG account encodes individual conceptual states in such trajectories. A PNG need not be oscillatory, so it does not make sense to extract features like frequency or phase to map a PNG into a vector space. In some important ways, the PNG approach is fundamentally

different than the vector space approach.

It may be tempting to view the vector space account as supervening on the PNG account, with PNGs implementing vector space states at some lower level of analysis. In this brief article, we argue that such a view either is untenable, with no mapping from the complexities of PNG representations to points in a vector space being possible, or, at least, is in dire need of an explanation of how such a reduction might be accomplished. Specifically, we raise four problematic issues that arise when shifting from a vector space approach to a PNG approach: (1) *Temporal Extension* — If conceptual representations involve temporally extended PNGs, to what degree can a conceptual state be said to be active at a particular instant or actively maintained over an interval? How does the use of the PNG framework change the characterization of the evolution of conceptual states over time? (2) *Grounding* — If conceptual role semantics is to be used to understand the "meaning" of a PNG representation, what are the biologically realistic causal mechanisms that link PNGs in an inferential cascade? (3) *Compositionality* — Does the PNG approach fare better or worse than the vector space approach in accommodating compositional or structured representations? (4) *Similarity* — Is there a distance measure for PNG representations that could be used to capture conceptual similarity while reflecting the way in which downstream neural circuits would naturally generalize across disparate PNGs?

Each of these four issues is elaborated in the following sections, and preliminary insights into how these issues could be addressed are provided. The goal is to highlight how the PNG approach challenges the prevailing vector space account of mental representation, while offering some clues concerning how this challenge might be met.

We begin by offering a review of PNG representations, providing a foundation for exploring each of the four issues that we find problematic. We then conclude with a brief discussion of open questions in this domain.

Polychronous Neuronal Groups

Polychronous Neuronal Groups (PNGs) have been proposed as a possible unit of representation in the human brain (Izhikevich, 2006). A PNG is a reproducible, time-locked, spatiotemporal spike-timing pattern over a collection of neurons. They are reproducible in the sense that the sequence of spike times tends to replay when the input conditions experienced by the neural network are repeated. They are time-locked in the sense that, once the PNG begins, the times between the spikes within the spiking pattern are the same whenever the PNG is triggered. They are spatiotemporal in the sense that they are defined in terms of a specific set of neurons that participate in the pattern (spatial) as well as the specific times at which spikes appear in the pattern (temporal). PNGs spontaneously emerge in spiking neural networks that incorporate variance in the amount of time it takes for an action potential to reach its receiving neurons (conductance delays), and they are reinforced by mechanisms of spike

timing dependent plasticity (STDP) (Izhikevich, 2006).

To understand the information-bearing properties of PNGs, it is important to understand how they are generated and propagated. An individual neuron remains at its resting potential until it receives, or "observes", a sufficient number of spikes in a short enough period of time, at which point this coincident input causes the neuron to generate an action potential of its own. This action potential is then, in turn, observed by the neurons to which this neuron projects. However, since it takes time for action potentials to propagate down axonal connections, there is a delay between when a spike is generated and when it is received. For example, in the cat brain, this delay can be as short as 0.1 ms, or as long as 44 ms (Swadlow, 1992). Since a cortical neuron may project to anywhere between 1,000 and 10,000 other neurons, a single action potential will be received at many different times. Thus, spikes that are synchronized on generation will not necessarily be synchronized on their receipt.

Typically, a single input spike is insufficient to drive the receiving neuron to fire an action potential, and the membrane potential of such a neuron is constantly decaying toward its resting potential. Within just a few milliseconds after receiving a single spike, the membrane potential of a neuron will return to its equilibrium state, removing the electrical effects of the spike (Cessac, Paugam-Moisy, Viéville, et al., 2010). This highlights the need for synchrony in the arrival of spikes to initiate firing, but it is important to remember that spikes that are synchronized at the time of receipt will not necessarily be synchronized at the time of their initiation, due to variance in conductance delays.
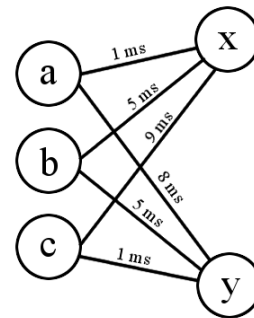

Figure 1: A small neural network with time delays.

Consider the network portrayed in Figure 1. If neurons *a*, *b*, and *c* spike at the same time, Time 0, those spikes will be received by neuron *x* at Times 1, 5, and 9, respectively, and those same spikes will be received by neuron *y* at Times 8, 5, and 1, respectively. In this case, neither *x* nor *y* receive the coincident spikes needed to fire. The difference in arrival times are caused by differences in axonal propagation times. If, instead, neurons *a*, *b*, and *c* spike at Times 8, 4, and 0, respectively, neuron *x* will receive all three of these spikes at Time 9, potentially allowing the cell to fire. In contrast, neuron *y* would receive the three spikes at Times 16, 9, and

1, respectively, providing it with no coincident spikes to drive an action potential. Alternatively, if neurons *a*, *b*, and *c* fire in the *reverse* order, neuron *y* will may spike, while neuron *x* will remain silent. Thus, the effects of spikes from neurons *a*, *b*, and *c* on the firing of neurons *x* and *y* is critically dependent on the timing of the spikes.

In larger, more connected, networks, like those found in mammalian brains, a particular stimulus will cause a chain reaction of spikes over time. This group of neurons firing with precise timing is what forms a corresponding PNG. Importantly, PNG patterns can be strengthened with repetition, with the strengthening well explained by spike timing dependent plasticity (STDP) (Izhikevich, 2006). Synapses that exhibit STDP are *strengthened* whenever the post-synaptic neuron fires just *after* it receives evidence of a pre-synaptic spike. Conversely, whenever the post-synaptic neuron fires just *before* it receives evidence of a pre-synaptic spike, then the synapse is *weakened* (Dan, Poo, et al., 2004). Thus, as a PNG unfolds, STDP strengthens the synapses participating in the PNG's generation and weakens the synapses that were active but did not facilitate the firing of neurons participating in the PNG. Thus, every time a particular PNG unfolds, and hence becomes strengthened via the mechanism of STDP, it becomes easier for that PNG to be reproduced.

To restate, a PNG is a reproducible, time-locked, spatiotemporal spike-timing pattern. A PNG is reproducible in the sense that, when the neurons participating in a PNG are stimulated in a similar way, the PNG will unfold in a similar way. Furthermore, each reproduction causes a PNG to become more stable through STDP, making the PNG increasingly robust to timing noise (i.e., some input spikes may be omitted or added without substantially effecting the generation of the PNG). A PNG is time-locked due to the fact that the conductance delays between the participating neurons are fixed by the anatomy of the network. A PNG is spatiotemporal in the sense that it necessarily occurs at many times (polychronous) and involves many neurons. Once stabilized via STDP, subtle variations in spike timing due to noise do not lead to unpredictably different PNGs, but generate a member of a family of related PNGs (Izhikevich, 2006). Also, it is important to note that many PNGs may be simultaneously active in a common neural network without interacting, due to the low probability that two arbitrary PNGs will overlap substantially in their precise spike times. In addition to these properties, PNGs also minimize redundancy through the weakening of synapses via STDP, and they are more energy efficient than vector space representations that depend on neural firing rate (Levy & Baxter, 1996). It is also interesting to note that the "small world" connectivity structure of the mammalian brain gives rise to stable PNGs much more readily than networks of neurons that are connected uniformly at random (Sporns & Zwi, 2004; Vertes & Duke, 2010).

Since their introduction, PNGs have been utilized extensively in computational neuroscience models of cognitive information processing. The intricate dynamics of PNGs have been used in combination with models of NMDA receptors and neurotransmitter reuptake to produce a promising account of working memory function (Szatmáry & Izhikevich, 2011). PNGs have been incorporated into a formal account of the dopamine system in order to produce a candidate model of neural reinforcement learning that addresses the problem of temporally distal reward (Izhikevich, 2007). In addition to their use in computational neuroscience models, empirical evidence for PNGs has been reported, with reproducible, time-locked, spatiotemporal patterns of spikes being observed in cortical slices (Rolston et al., 2007). It is clear that PNGs show great promise as a form of representation in the brain.

As this review of PNGs shows, the information carried by a PNG in a neural network is critically dependent on the timing of individual action potentials. This contrasts with vector space accounts of mental representation which map vector space dimensions onto the instantaneous firing rates of neurons. The PNG approach highlights the way in which individual spike times can carry information, with spiking rates lacking sufficient spatiotemporal detail to discriminate between different representational states. It is this shift that gives rise to a number of potential problems with viewing PNGs as the foundation of mental representation.

Temporally Extended Representations

The Challenge

In vector space representation schemes that make contact with biology, each dimension corresponds to the instantaneous activity of a neural element, such as the instantaneous firing rate of a neuron. This provides us with a natural way to capture the mental state of the agent at any given point in time. The active maintenance of a conceptual state involves a relative lack of change in these firing rates, and the evolution of mental states over time are captured in trajectories through the vector space.

In contrast, the PNG approach inherently involves temporally extended representations. A PNG is a spatiotemporal pattern of spikes. If mental representations are to be associated with PNGs, to what degree can any concept be seen as active at any given instant? Can a representation be actively maintained over time if the physical substrate of the representation is changing over time? How does the use of the PNG framework change the characterization of the evolution of conceptual states over time?

Addressing the Challenge

While some philosophical work may be needed to fully appreciate the nature of temporally extended mental representations, we do not see this challenge as insurmountable.

The activation level of a particular PNG at an arbitrary point in time, in the midst of a sequence of spikes, does not have a clear definition. We can identify, however, the degree to which recently produced spikes match portions of a PNG, as well as the propensity for the neural network to continue

with the production of further spikes in the PNG. Thus, the notion of the activation of a concept is only coherent at time scales that match the time scale of the PNG. For example, if the PNG π refers to a 30 ms long spike sequence, it may be asked if, over the last 30 ms, π appeared. For that same π, it may also be asked at what times over the last 60 ms π was present. In this way, activation of a mental state represented by a PNG only makes sense at relatively coarse time scales in comparison to the time scales used for common vector space representations. The PNG approach does not admit to a coherent sense of a truly instantaneous mental state.

The mental state encoded by a PNG may be actively maintained for a period of time longer than the duration of the spike sequence that makes up the PNG. For example, PNG models of working memory have involved the repeated activation of a PNG, with the spike pattern being sequentially reinitiated, allowing it to persist for arbitrarily long periods, as needed (Szatmáry & Izhikevich, 2011). Once again, this notion of active maintenance is limited to a time scale corresponding to the temporal length of the PNG, but the PNG approach does not rule out the possibility of persisting in a mental state for a longer period of time.

Grounding Of Representations

The Challenge

The *symbol grounding problem* highlights the need for representational schemes to provide some account of the *meaning* of mental representations (Harnad, 1990). Understanding in an ungrounded representational system is analogous to the content of a dictionary, defining words in an ultimately circular fashion, in terms of other words. It has been argued that this problem can be overcome by grounding internal representations in reflections of the world, mediated by *iconic representations* associated with direct sensations (Harnad, 1990). Similarly, the meaning of an internal representation can be seen as arising from the role it plays in a cognitive inferential process, as described by the theory of *conceptual role semantics*, with causal and inferential chains eventually connecting to the world through sensation and action (Greenberg & Harman, 2006).

In the vector space approach, representations may be grounded in the causal processes in which they participate, both in terms of inputs that give rise to a representation (eventually leading back to iconic representations) and the effects of that representation on downstream neural circuits. These causal relationships can be characterized in terms of functional mappings between vector spaces. For example, if the transduction of sensory information from the world directly results in a pattern of neural firing rates, this pattern corresponds to a point in a sensory vector space, and neural circuits can be seen as mapping this point to corresponding points in the vector spaces for other neural populations, encoding the corresponding conceptual content. In this way, the mappings implemented by neurocomputational mechanisms ground internal representations.

The activation of a PNG and the downstream effects of the initiation of a PNG are mechanistically and computationally quite different than standard neurocomputational mechanisms that can be easily cast as functions between vector spaces. Given this difference in the causal structure of PNGs, how can PNG representations be grounded?

Addressing the Challenge

We assert that the PNG approach to mental representation requires only a slightly different understanding of the nature of the relevant causal relationships. Rather than being characterized as functional mappings between vector spaces, we posit that the causal relationships between PNGs are best described in terms of *trigger sets*. Let us first consider the definition of a σ-triggered polychronous neuronal group.

Definition A σ-triggered polychronous group refers to the set of neurons that can be activated by a chain reaction whenever trigger neurons $N_k (1 \leq k \leq \sigma)$ fire according to the timing pattern $t_k (1 \leq k \leq \sigma)$, where σ is the size of the stimulus required to trigger the PNG (Martinez & Paugam-Moisy, 2009).

Here, we recognize that a given PNG can have more than one stimulus trigger. This distinction motivates the definition of a *trigger set*.

Definition For a given PNG π, its *trigger set*, τ_π, is the set of spike-time patterns that trigger the existence of π. Each *spike-time pattern* in the trigger set of π will give rise to π when presented in the absence of interfering spikes.

Note that each element of a trigger set may be a PNG or a portion of a PNG. Thus, the presence of a PNG spike-time pattern may trigger, or help trigger, other PNGs. The set of PNGs that have the potential of being triggered by a given PNG can form the core of a formal characterization of the causally grounded *meaning*, μ_π, of that PNG.

Definition If a PNG, π_0, is seen as a set of spikes (with each spike indexed by the identity of the spiking neuron and the time of the spike), and $\mathbb{P}$ is the set of all possible PNGs in the neural network, then the *meaning* of π_0, called μ_{π_0}, is defined as the set of PNGs whose trigger set, τ_π, contains an element with a nonempty intersection with π_0, or

$$\mu_{\pi_0} = \{\pi : \pi \in \mathbb{P}, \ \exists \pi_t \in \tau_\pi, \ \pi_t \cap \pi_0 \neq \emptyset\}.$$

By this definition, μ_{π_0} includes any PNG for which π_0 contributes some spikes that may contribute to the triggering of the PNG. Thus, μ_{π_0} includes PNGs that may only be triggered by π_0 in the *context* of other spike-time patterns. In this way, the meaning of a PNG is *context sensitive*. While constrained by the network's topology, the size of μ_{π_0} may be very large.

These definitions describe the causal relationships between PNGs, providing the basis for a *conceptual role semantics* approach to symbol grounding in the PNG framework. PNGs

triggered by sensory neurons can be considered iconic representations, as can PNGs that trigger motor responses. PNGs triggered by other neurons in a network form internal representations that are ultimately grounded in these iconic representations through the causal connections of their trigger sets.

Compositionality Of Representations

The Challenge

A representation that exhibits compositionality, simply put, is one where "the meaning of a compound expression is a function of the meanings of its parts" (Janssen, 1996). In the vector space approach, compositional representations have historically been difficult to capture (Fodor & Pylyshyn, 1988), but progress has been made (Van Gelder, 1990; Gayler & Levy, 2011). The most common solutions involve either representational components being maintained in subspaces of a parent vector space, or components being superimposed or convolved to form compound representations like tensor product codes, holographic reduced representations, or sparse codes (Smolensky, 1990; Plate, 2003; OReilly et al., 2011).

Do PNG representations suffer from the same problems of compositionality as vector space representations? Are current approaches to compositionality in the vector space framework also appropriate for PNG representations? How might compositionality be captured in PNG representations?

Addressing the Challenge

The PNG approach offers two ways to capture compositional representations that are impossible, or at least of limited utility, in the vector space approach. These two methods include sequential concatenation of PNG component representations and the superposition of PNG components.

A PNG representation might be seen as being composed of subsequences of spikes, giving rise to representations at multiple time scales. In essence, a PNG may be seen as containing many smaller PNGs within it, or, inversely, it may be a part of a sequence of other PNGs. Consider the meaning of a particular PNG, μ_{π_0}. If a PNG $\pi_1 \in \mu_{\pi_0}$ has an element of its trigger set, τ_{π_1}, contained completely within π_0, then π_0 will reliably trigger π_1, in the absense of interfering spikes. Note that $\mu_{\pi_0} \neq \mu_{\pi_1}$, so these two PNGs do not have the same meaning. These two PNGs may be combined by simple concatenation, producing a new PNG, $\pi_2 = \pi_0 \cup \pi_1$. Importantly, the meaning of this compound representation, μ_{π_2}, is a simple function of the meanings of its parts: $\mu_{\pi_2} = \mu_{\pi_0} \cup \mu_{\pi_1}$. (Note that this is the case even if π_0 does *not* reliably trigger π_1.)

An alternative approach to compositionality involves directly superimposing PNGs over the same time interval. Since spikes may be sparse over time, the probability of superposition producing interference between PNGs is relatively small. In this way, the composition of PNG representations may simply involve the simultaneous activation of the component PNGs.

Similarity Of Representations

The Challenge

Distances in a vector space have been fruitfully used to capture dissimilarity between representations, providing a useful mechanism for generalization. Common distance metrics, like inner-product distance (related to angular distance), are tightly related to the kinds of functional mappings between vector spaces that are are easily implemented by neural circuits. How might similarity between representations be captured in the PNG approach?

Addressing the Challenge

We see this as an important open question. There are many existing metrics for evaluating the similarity of spike trains (Victor & Purpura, 1996; Naud, Gerhard, Mensi, & Gerstner, 2011). Developing a good similarity metric has proven difficult, however, as spike sequences are inherently non-Euclidian (Aronov & Victor, 2004). The metrics that have been presented focus on comparing spike trains generated by a single neuron, recorded over many trials. Current approaches make the assumption that the significance of a spike can be treated as independent from other spikes produced in the same neural network. Because of these and other issues, existing similarity metrics for comparing groups of spike trains do not predict well whether two PNGs will have a similar effect on downstream neurons.

It is our suspicion that similarity metrics based on trigger sets may overcome some of the obstacles described here. Determining a meaningful similarity metric for PNG representations is a focus of our future work.

Conclusions and Future Work

Ongoing work in computational neuroscience is uncovering the powerful capabilities of polychronous neuronal groups, and empirical studies are starting to find evidence for this kind of encoding in biological neural circuits. If PNGs emerge as a dominant means of representation in the brain, the vector space account of conceptual states will need to be reconsidered. While many of the strengths of the vector space account appear to transfer, with some modifications, to the PNG framework, there remain challenges for viewing complex spike patterns as conceptual representations. Perhaps the most substantial challenge faced by the PNG approach to conceptual representation involves the nature of representational similarity in this framework. This is the primary focus of our ongoing research program.

Much philosophical work has been done in order to clarify how the vector space account of mental representation provides a bridge between brain processes and cognitive processes. If further neuroscientific investigations suggest that polychronous neuronal groups carry conceptual content in at least some brain systems, similar work will need to be pursued for this representational framework. By highlighting several potentially problematic issues with PNG representations, and sketching promising solutions for some of these

issues, we hope to have helped launch this philosophical effort.

Acknowledgments

References

Aronov, D., & Victor, J. (2004). Non-euclidean properties of spike train metric spaces. *Physical Review E*, *69*(6), 061905.

Cessac, B., Paugam-Moisy, H., Viéville, T., et al. (2010). Overview of facts and issues about neural coding by spikes. *Journal of physiology, Paris*, *104*(1), 5.

Churchland, P. (1989). *The neurocomputational perspective: Philosophical essays on the mind, brain, and science.* Cambridge, USA: MIT Press.

Churchland, P. (1996). *The engine of reason, the seat of the soul: A philosophical journey into the brain.* Cambridge, MA: MIT Press.

Dan, Y., Poo, M., et al. (2004). Spike timing-dependent plasticity of neural circuits. *Neuron*, *44*(1), 23–30.

Fodor, J., & Pylyshyn, Z. (1988). Connectionism and cognitive architecture: A critical analysis. *Cognition*, *28*(1), 3–71.

Fries, P., et al. (2005). A mechanism for cognitive dynamics: neuronal communication through neuronal coherence. *Trends in cognitive sciences*, *9*(10), 474–480.

Gärdenfors, P. (2000). *Conceptual spaces-the geometry of thought.* Cambridge, MA: Bradford Books, MIT Press.

Gayler, R., & Levy, S. (2011). Compositional connectionism in cognitive science ii: the localist/distributed dimension. *Connection Science*, *23*(2), 85–89.

Greenberg, M., & Harman, G. (2006). Conceptual role semantics. In *Oxford handbook of philosophy of language.* Oxford: Oxford University Press.

Harnad, S. (1990). The symbol grounding problem. *Physica D: Nonlinear Phenomena*, *42*(1), 335–346.

Izhikevich, E. (2006). Polychronization: Computation with spikes. *Neural computation*, *18*(2), 245–282.

Izhikevich, E. (2007). Solving the distal reward problem through linkage of stdp and dopamine signaling. *Cerebral Cortex*, *17*(10), 2443–2452.

Janssen, T. (1996). Compositionality. In *Handbook of logic and linguistics.* North Holland, Amsterdam: University of Amsterdam.

Levy, W., & Baxter, R. (1996). Energy efficient neural codes. *Neural Computation*, *8*(3), 531–543.

Madhavan, R., Chao, Z., & Potter, S. (2007). Plasticity of recurring spatiotemporal activity patterns in cortical networks. *Physical biology*, *4*(3), 181.

Martinez, R., & Paugam-Moisy, H. (2009). Algorithms for structural and dynamical polychronous groups detection. *Artificial Neural Networks–ICANN 2009*, 75–84.

Naud, R., Gerhard, F., Mensi, S., & Gerstner, W. (2011). Improved similarity measures for small sets of spike trains. *Neural computation*, *23*(12), 3016–3069.

OReilly, R., Bhattacharyya, R., Howard, M., & Ketz, N. (2011). Complementary learning systems. *Cognitive science*.

Pasquale, V., Massobrio, P., Bologna, L., Chiappalone, M., Martinoia, S., et al. (2008). Self-organization and neuronal avalanches in networks of dissociated cortical neurons. *Neuroscience*, *153*(4), 1354–1369.

Plate, T. (2003). *Plate, holographic reduced representation: Distributed representation for cognitive structures.* CSLI Publications, Stanford, CA.

Rieke, F., Warland, D., Steveninck, R. van, & Bialek, W. (1999). *Spikes: exploring the neural code*. Cambridge, MA: MIT press.

Rolston, J., Wagenaar, D., & Potter, S. (2007). Precisely timed spatiotemporal patterns of neural activity in dissociated cortical cultures. *Neuroscience*, *148*(1), 294.

Smolensky, P. (1990). Tensor product variable binding and the representation of symbolic structures in connectionist systems. *Artificial intelligence*, *46*(1), 159–216.

Spivey, M. (2008). *The continuity of mind* (Vol. 40). USA: Oxford University Press.

Sporns, O., & Zwi, J. (2004). The small world of the cerebral cortex. *Neuroinformatics*, *2*(2), 145–162.

Swadlow, H. (1992). Monitoring the excitability of neocortical efferent neurons to direct activation by extracellular current pulses. *Journal of neurophysiology*, *68*(2), 605–619.

Szatmáry, B., & Izhikevich, E. (2011). Spike-timing theory of working. *BMC Neuroscience*, *12*(Suppl 1), O3.

Van Gelder, T. (1990). Compositionality: A connectionist variation on a classical theme. *Cognitive Science*, *14*(3), 355–384.

Vertes, P., & Duke, T. (2010). Effect of network topology on neuronal encoding based on spatiotemporal patterns of spikes. *HFSP journal*, *4*(3-4), 153–163.

Victor, J., & Purpura, K. (1996). Nature and precision of temporal coding in visual cortex: a metric-space analysis. *Journal of Neurophysiology*, *76*(2), 1310–1326.

Yoshimi, J. (2012). Active internalism and open dynamical systems. *Philosophical Psychology*, *25*(1), 1–24.