Implications of polychronous neuronal groups for the continuity of mind

William Benjamin St. Clair & David C. Noelle

Cognitive Processing
International Quarterly of Cognitive Science

ISSN 1612-4782
Volume 16
Number 4

Cogn Process (2015) 16:319-323
DOI 10.1007/s10339-015-0645-5

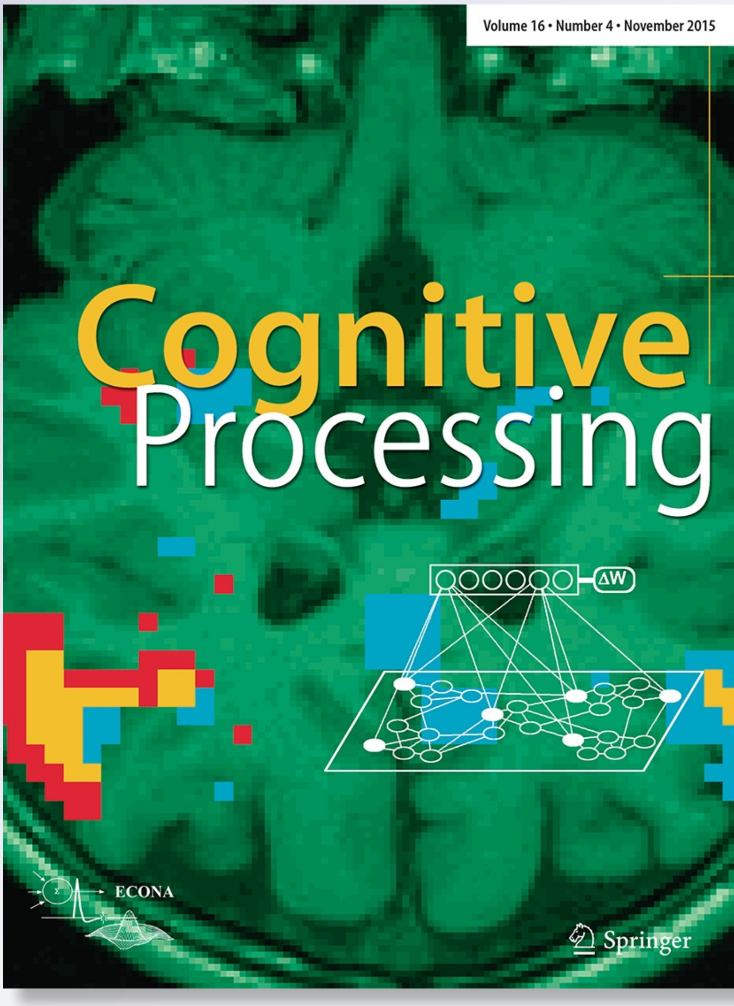

Springer

Cogn Process (2015) 16:319–323
DOI 10.1007/s10339-015-0645-5

SHORT REPORT

Implications of polychronous neuronal groups for the continuity of mind

William Benjamin St. Clair · David C. Noelle

Received: 6 October 2014 / Accepted: 16 January 2015 / Published online: 29 January 2015

Abstract Is conceptual space continuous? The answer to this question depends on how concepts are represented in the brain. Vector space representations, which ground conceptual states in the instantaneous firing rates of neurons, have successfully captured cognitive dynamics in a broad range of domains. There is a growing body of evidence, however, that conceptual information is encoded in spatiotemporal patterns of neural spikes, sometimes called polychronous neuronal groups (PNGs). The use of PNGs to represent conceptual states, rather than employing a continuous vector space, introduces new challenges, including issues of temporally extended representations, meaning through symbol grounding, compositionality, and representational similarity. In this article, we explore how PNGs support discontinuous transitions between concepts. While the continuous dynamics of vector space approaches require such transitions to activate intermediate and blended concepts, PNGs offer the means to change the activation of concepts discretely, introducing a form of conceptual dynamics unavailable to vector space models.

Keywords Neural representation · Dynamical systems · Philosophy of mind

This article is part of the Special Issue on 'Complexity in brain and cognition' and has been edited by Cees van Leeuwen.

W. B. St. Clair · D. C. Noelle (✉)
University of California Merced, Merced, CA, USA
e-mail: dnoelle@ucmerced.edu

W. B. St. Clair
e-mail: wst.clair@ucmerced.edu

Introduction

One highly productive approach to mental representation has involved the application of a continuous dynamical systems perspective (Spivey 2008). From this perspective, currently active concepts in a cognitive system (or subsystem) are jointly encoded as a point in a high dimensional conceptual vector space (CVS) (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 (Shepard 1987). The evolution of mental states over time becomes a trajectory in this vector space (Yoshimi 2012), driven by mechanistic cognitive processes (Churchland 1989). In general, the CVS approach has been very productive.

Past challenges to the CVS approach have come from above: from more abstract and symbolic characterizations of cognitive processing. More recently, a challenge has arisen from below: from insights into the neural coding of information. 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 et al. 2007; Madhavan et al. 2007; Pasquale et al. 2008). While information may be carried by synchronous or coherent firing of neurons, as in synfire chains (Bienenstock 1995), 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 CVS approach (St. Clair and Noelle 2013). A PNG is a temporally extended pattern of discrete spiking events over a collection of neurons. 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. 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 continuous vector space. In some important ways, the PNG approach is fundamentally different than the CVS approach. In this work, we will focus on how the PNG approach differs from the CVS approach with respect to the continuity of conceptual transitions. We conclude that if it turns out that parts of the nervous system depend on PNGs to encode content, then those systems may exhibit discrete transitions from one concept to the next, without activating intermediate and blended representations that are implied by CVS approaches.

Polychronous neuronal groups

Polychronous neuronal groups (PNGs) have been proposed as a possible unit of representation in the 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 neuronal network are repeated. They are time-locked in the sense that, once the PNG begins, the times between the spikes within the 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 occur (temporal). PNGs spontaneously emerge in spiking neuronal networks that incorporate variance in the amount of time it takes for an action potential to reach its receiving neurons (conduction delays). There is a growing literature on cognitive processing using PNGs (Izhikevich 2007; Chorley and Seth 2011; Szatmáry and Izhikevich 2011), which builds on earlier work with synfire chains (Bienenstock 1995; Hayon et al. 2005; Trengove et al. 2013).

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 a sufficient number of spikes in a short enough period of time, at which point the neuron generates its own spike. This spike is then, in turn, received by the neurons to which this neuron projects. However, since it takes time for spikes to propagate down axons, there is a delay between when a spike is generated and when it is received. Since a cortical neuron may project to thousands of other neurons, a single spike 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, and the membrane potential of such a neuron is constantly and quickly decaying toward its resting potential. 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 axonal propagation delays.

Consider the network portrayed in Fig. 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 is 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* 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* are critically dependent on the timing of the spikes.

A PNG is a reproducible, time-locked, spatiotemporal pattern of spikes. PNGs become more stable through mechanisms of synaptic plasticity, such as spike-timing dependent plasticity (STDP) (Dan and Poo 2004; Izhikevich 2006). This increases the likelihood that the PNG will be triggered in similar situations in the future, and it makes the PNG increasingly robust to noise. Once stabilized via STDP, subtle variations in spike timing due to noise do not lead to unpredictablly 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 neuronal network without interacting, due to the low probability that two

Fig. 1 A small neuronal network with time delays

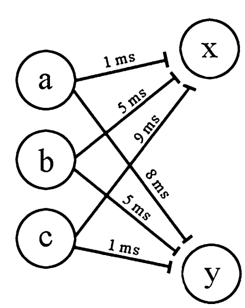

arbitrary PNGs will overlap substantially in both the set of neurons involved and their precise spike times.

PNG triggering and conceptual transitions

The conditions under which a given PNG appears may be characterized by its trigger set, formalized as follows:

Definition A *PNG*, π, is a spatiotemporal spike pattern in a neuronal network captured by a set of neuron-time pairs. The trigger set of π, 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.

The definition of a PNG's trigger set highlights the fact that the activation of a PNG can be caused in multiple ways. Note that each element of a trigger set is, itself, a PNG, since a PNG is defined simply in terms of its neuron-time pairs. This fact allows us to formally characterize the potential causal relationships between PNGs.

Definition Given a PNG, π_0, and that $\mathbb{P}$ is the set of all possible PNGs in the neuronal network, then the *meaning* of π_0, μ_{π_0}, is a pair of sets: its *causal component*, $\mu^C_{\pi_0}$, defined as the set of PNGs whose trigger set, τ_π, contains an element with a nonempty intersection with π_0, or

$$\mu^C_{\pi_0} = \{\pi : \pi \in \mathbb{P},\ \exists\ \pi_t \in \tau_\pi,\ \pi_t \cap \pi_0 \neq \emptyset\};$$

and its *existential component*, $\mu^E_{\pi_0}$, defined as the set of PNGs that contain a nonempty intersection with an element of the trigger set of π_0, τ_{π_0}, or

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

By this definition, the forward-looking causal component of μ_{π_0}, $\mu^C_{\pi_0}$, includes any PNG for which π_0 contains some spikes that may contribute to the triggering of that PNG. This means that $\mu^C_{\pi_0}$ includes PNGs that may only be triggered by π_0 in the context of other spike-time patterns occurring in the network. The backward-looking existential component of $\mu^E_{\pi_0}$ includes any PNG that contains spikes that may contribute to the triggering of π_0. This means that $\mu^E_{\pi_0}$ includes PNGs that may only trigger π_0 in the context of other spike-time patterns occurring in the network. Thus, context sensitivity is intrinsic to the *meaning* of a PNG.

These definitions form the foundation of a neural account of conceptual representation. One leading philosophical account of the grounding of conceptual meaning is Conceptual Role Semantics (Block 1998). In brief, this theory characterizes the meaning of a representation in terms of the causal relationships surrounding it. Thus, the meaning of a representation depends on the representations that can cause it, as well as the representations that it can cause, eventually making contact with sensory-motor processes interacting with the world. The formal definition of the meaning of a PNG, μ_π, makes explicit and precise the causal relationships between PNGs, allowing them to act as grounded representations according to Conceptual Role Semantics (St. Clair and Noelle 2013).

There are additional features of PNGs that make them attractive for theories of neural representation. For example, a special case of PNGs, called synfire chains, has been studied extensively as a unit of representation. Generalizing from this research, spatiotemporal spiking patterns have been argued to have the essential properties of conceptual representation, including stability, reproducibility, learnability, storage capacity, and even compositionality (Bienenstock 1995; Abeles et al. 2004; Hayon et al. 2005). Indeed, there are natural mechanisms for composing PNGs, potentially producing complex compositional representations, simply by superimposing them, with multiple PNGs being simultaneously active. Using our definition of PNG meaning, the meaning of the union of two simultaneous PNGs will be systematically related to the meanings of the two individual PNGs, providing an approach to compositional semantics.

Our definition of PNG meaning is not intended to suggest that every PNG in a neuronal network corresponds to a psychological concept. Indeed, since any reliably produced set of spikes fits our simple definition of a PNG, even very short spike sequences are formally PNGs. Our definition of PNG meaning, however, does provide a perspective on which specific PNGs might correspond to useful concepts. Specifically, a PNG should be considered as a useful unit of representation to the degree that it arises in many contexts (i.e., the cardinality of its μ^E_π is large) and it participates in the triggering of a wide variety of other PNGs (i.e., the cardinality of its μ^C_π is large). From this perspective, useful conceptual units appear relatively independently. Thus, every concept is captured by a PNG, but not every PNG encodes a concept.

To visualize the dynamics of PNGs, consider Fig. 2. On the left, 4 neurons fire in a precise time pattern. This pattern is sufficient to trigger PNG 1. As PNG 1 unfolds, it includes a 4 spike sequence that is sufficient to trigger PNG 2. This second PNG can, itself, trigger further PNGs. If we now consider PNGs as encoding conceptual information, then PNG triggering dynamics can be seen as underlying conceptual transitions. Note that the triggers for PNG 2 do not need to be fully contained within PNG 1. We can imagine a different PNG 2 that is triggered with the aid of co-occurring contextual spikes in addition to those present in PNG 1. While this case is not depicted in the cartoon of Fig. 2, in any complex neuronal network, the triggering of a PNG by multiple preceding PNGs will be the most

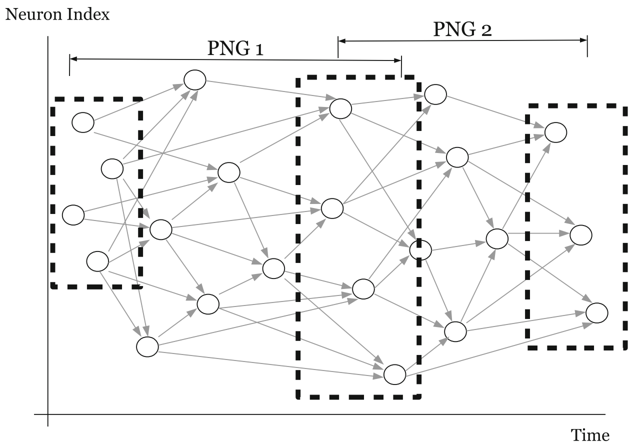

Fig. 2 The *vertical dimension* specifies individual neurons in a group, and the *horizontal dimension* is time. *Circles* correspond to spikes generated by a particular neuron at a specific time. The *arrows* between spikes display causal effects between events. The *dashed* outlines contain groups of spikes whose existence is sufficient to trigger the following PNG

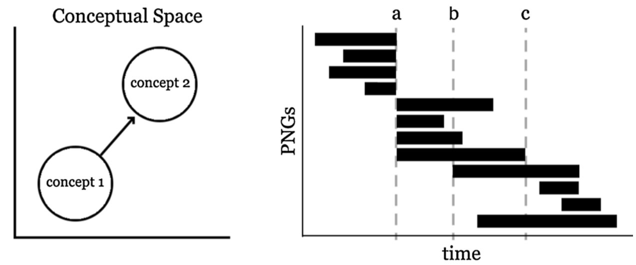

Fig. 3 *Left* A CVS, with concepts encoded as regions, with dimensions either being arbitrary or related to the firing rates of neurons. Conceptual transitions require motion through intermediate regions of conceptual space, which necessarily correspond to intermediate and blended concepts. *Right* Both discontinuous and blended transitions between complex concepts, encoded as collections of PNGs, are shown. At time *a*, there is a discrete transition from an initial concept to a second one. Between times *b* and *c*, some PNGs from one concept are active alongside PNGs from another concept, capturing a blended conceptual transition during this period

common situation. This means that the activation of a PNG can be highly context sensitive. Also, recall that many PNGs can be simultaneously occurring without significant interaction, implying that PNGs can robustly coexist without undermining their heteroclinic stability (Izhikevich 2006). This point about representational capacity is also supported by work on synfire chains (Trengove et al. 2013; Schrader et al. 2011).

Blended transitions

We are specifically concerned with the transitions from one conceptual state to another. Many researchers have suggested that concepts are best seen as regions in a CVS. Similarity is captured by distance metrics in this space. The space is also seen as a state space, through which dynamic trajectories unfold during cognitive processing (Spivey 2008). The dimensions of the CVS are sometimes related to the instantaneous firing rates of groups of neurons.

In the CVS approach, a conceptual transition necessarily requires motion through an intermediate region of conceptual space, as portrayed in Fig. 3. Along this dynamic trajectory, the system temporarily represents concepts that bear a steadily decreasing degree of similarity to the first concept and a steadily increasing degree of similarity to the next concept. In this way, conceptual transitions inherently involve briefly activating intermediate and blended concepts. Extensive evidence has been gathered supporting CVS approaches (Shepard 1987; Rolls and Tovee 1995; Spivey 2008), and a variety of psychological phenomena, including priming effects, have been addressed (Mirman and Magnuson 2009; Cree et al. 1999).

No intermediate representations need arise in the PNG approach. When concepts are taken to be encoded as individual PNGs, such conceptual transitions arise when the existence of one PNG (often in the context of other, simultaneously active, PNGs) *triggers* the existence of another PNG, as caricatured in Fig. 2. The unfolding of the initial PNG triggers the subsequent one, but at no point during this transition is there a spike-time pattern that is some sort of "blend" between the two spatiotemporal spike-time patterns that make up any of these PNGs. The only overlap between the PNGs is the sequence of spikes that trigger the second PNG, and that overlap can be vanishingly small compared to the extent of the PNGs. In this way, PNG representations allow for discrete conceptual transitions without the activation of blended concepts. This is possible because, with PNGs, the mechanisms driving conceptual transitions are decoupled from measures of conceptual similarity.

While spatiotemporal codes allow completely discrete transitions, they do not rule out transitions that activate intermediate or blended representations. Neuronal networks in the PNG framework may very well produce such blended states, particularly when the full conceptual state of the system is captured by multiple simultaneously active PNGs. As previously noted, since PNGs can be superimposed without much interference, the simultaneous unfolding of multiple PNGs would be a natural code for more complex compositional concepts. A conceptual transition from one complex concept to another might involve the deactivation of PNGs that make up the initial concept during the activation of PNGs that make up the following concept. If the initial concept PNGs all ended just as the following concept PNGs started, the transition would be discontinuous. If, however, the initial concept PNGs dropped out incrementally over time, and the following concept PNGs began at a staggered schedule, there would be an intermediate time at which some of the PNGs for the initial concept would be active alongside the PNGs for the following concept. This is a natural characterization of an intermediate or blended complex concept. An illustration of both discontinuous and blended transitions of these kinds appears in the right panel of Fig. 3. In this way, the PNG framework allows for blended conceptual transitions, but, unlike the continuous vector space approach, it does not require them.

Finally, it is important to note that discontinuous jumps in conceptual space are possible in other spiking neuronal network frameworks, such as synfire chain representations (Bienenstock 1995; Hayon et al. 2005). While we have used PNGs to highlight this difference between certain spike codes and CVS approaches, other frameworks could depart from continuous firing-rate models in similar ways.

Conclusions

In summary, attending to certain features of the spike generation process in the nervous system, such as the discrete nature of spikes and the variance of axonal delays, suggests that an appropriate unit of conceptual representation might be the PNG. If some parts of the nervous system depend on PNGs to encode content, those systems may exhibit discrete transitions from one concept to the next, without activating intermediate or blended representations. In this way, PNGs offer an alternative to accounts of cognition in which the dynamics of mental processing are seen as consistently continuous.

The shift from a CVS account to a PNG approach raises issues in the philosophy of mental representations beyond what has been addressed here. Specifically, we see PNGs as presenting challenges and opportunities for questions of temporal extension, meaning through representational grounding, compositionality, and conceptual similarity (St. Clair and Noelle 2013). While the PNG approach presents new challenges, it also presents new solutions to problems of mental representation, and, thus, shows promise for bridging the gap between neural dynamics and cognitive processing.

References

Abeles M, Hayon G, Lehmann D (2004) Modeling compositionality by dynamic binding of synfire chains. J Comput Neurosci 17(2):179–201

Bienenstock E (1995) A model of neocortex. Netw Comput Neural Syst 6(2):179–224

Block N (1998) Semantics, conceptual role. In: Craig E (ed) Routledge Encyclopedia of Philosophy, vol 8. Routledge, London, pp 652–657

Chorley P, Seth AK (2011) Dopamine-signaled reward predictions generated by competitive excitation and inhibition in a spiking neural network model. Front Comput Neurosci 5:21

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

Cree GS, McRae K, McNorgan C (1999) An attractor model of lexical conceptual processing: simulating semantic priming. Cogn Sci 23(3):371–414

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

Hayon G, Abeles M, Lehmann D (2005) A model for representing the dynamics of a system of synfire chains. J Comput Neurosci 18(1):41–53

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

Izhikevich E (2007) Solving the distal reward problem through linkage of STDP and dopamine signaling. Cereb Cortex 17(10):2443–2452

Madhavan R, Chao Z, Potter S (2007) Plasticity of recurring spatiotemporal activity patterns in cortical networks. Phys Biol 4(3):181

Mirman D, Magnuson JS (2009) Dynamics of activation of semantically similar concepts during spoken word recognition. Mem Cogn 37(7):1026–1039

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

Rolls ET, Tovee MJ (1995) Sparseness of the neuronal representation of stimuli in the primate temporal visual cortex. J Neurophysiol 73(2):713–726

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

Schrader S, Diesmann M, Morrison A (2011) A compositionality machine realized by a hierarchic architecture of synfire chains. Front Comput Neurosci 4:154

Shepard RN (1987) Toward a universal law of generalization for psychological science. Science 237(4820):1317–1323

Spivey M (2008) The continuity of mind, vol 40. Oxford University Press, New York

St. Clair W, Noelle D (2013) Implications of polychronous neuronal groups for the nature of mental representations. In: Proceedings of the 35th annual conference of the cognitive science society, pp 1372–1377

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

Trengove C, van Leeuwen C, Diesmann M (2013) High-capacity embedding of synfire chains in a cortical network model. J Comput Neurosci 34(2):185–209

Yoshimi J (2012) Active internalism and open dynamical systems. Philos Psychol 25(1):1–24