

Evidence for Response Consistency Supports Polychronous Neural Groups as an Underlying Mechanism for Representation and Memory

Mira Guise, Alistair Knott, and Lubica Benuskova

Dept of Computer Science, University of Otago, Dunedin, New Zealand
{mguise,alikh,lubica}@cs.otago.ac.nz

Abstract. Izhikevich [6] has proposed that certain strongly connected groups of neurons known as polychronous neural groups (or PNGs) might provide the neural basis for representation and memory. Polychronous groups exist in large numbers within the connection graph of a spiking neural network, providing a large repertoire of structures that can potentially match an external stimulus [6,8]. In this paper we examine some of the requirements of a representational system and test the idea of PNGs as the underlying mechanism against one of these requirements, the requirement for consistency in the neural response to stimuli. The results provide preliminary evidence for consistency of PNG activation in response to known stimuli, although these results are limited by problems with the current methods for detecting PNG activation.

Keywords: spiking network, polychronous neural group, activation, representation, memory.

1 Introduction

It is widely assumed that synaptic plasticity provides the neural basis for long-term memory in the brain [1,2,9] although the precise nature of the underlying representation is still unclear [3]. Izhikevich [6] has proposed that certain strongly connected groups of neurons known as polychronous neural groups (or PNGs) might provide this representational mechanism. An understanding of this underlying mechanism is particularly relevant to the developing field of neuromorphic computing, but is also of interest to researchers in machine learning, or even information retrieval [5]. In this report we examine some of the requirements of a representational system and test the idea of PNGs as a mechanism of representation against one of these requirements, the requirement for consistency in the neural response to stimuli.

Polychronous groups arise from an interaction between the precise firing times of spatio-temporal input patterns and the variability of axonal transmission delays between neurons. Figure 1 shows a schematic example of such an interaction. The input stimulus is composed of a sequence of firing events, each representing the firing of a specific neuron at a precise point in time. The stimulus in this example forms an ascending spatio-temporal pattern as shown in Fig. 1 (unfilled circles).

In this model network there exists a polychronous group whose intra-group axonal delays are congruent with the input stimulus (gray-filled circles). As shown in panels A, B and C, the spatio-temporal arrangement of three of the firing events that make up the stimulus (filled black circles in Fig. 1) interacts with the axonal delays, producing convergent input to group neurons. This firing event triplet acts as a trigger for PNG activation, producing a wave of neural firing that propagates throughout the polychronous group (only the first step is shown). Without this convergent input the neurons in the group would fail to reach the firing threshold and the input stimulus would not propagate.

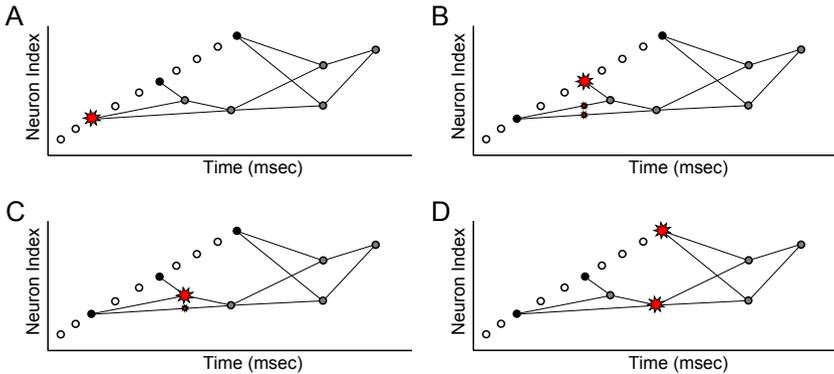


Fig. 1. Polychronous group activation following stimulation with an ascending firing pattern (unfilled circles). The precise timing of three of the firing events in the pattern (denoted by filled black circles) matches the axonal delays between group neurons, producing activation of the polychronous group. Other firing event combinations might produce additional group activations (not shown). Firing events resulting from PNG activation are shown with gray-filled circles. The first of the three firing events is fired in panel A, producing spikes that take time to propagate to PNG neurons. In panel B, the convergence of the propagating spike and the second of the triplet firing events is sufficient for a PNG neuron to reach the firing threshold (panel C). Further group firing events are supported by the axonal delays between group neurons (panel D).

This propagation of neural firing across the group is called *group activation*. When activated, the neurons in the polychronous group are said to *polychronize* in a causal chain of firing events that is both precisely timed and reproducible [6,8]. However, not all PNGs are capable of activation. *Structural* PNGs are defined purely topologically, as groups of neurons with connection latencies commensurate with a given input stimulus [10]. For polychronization to occur the synaptic connections converging on each group neuron must be sufficiently strong to allow the post-synaptic neuron to reach the firing threshold. Polychronous groups with compatible synaptic weights can activate when presented with a triggering stimulus at which point they are known as *activated* PNGs.

Izhikevich [6] observed that the number of structural PNGs in a network is typically many times larger than the number of neurons. Given this large

repertoire of structural PNGs, how might we use it to build a representational system? Several attributes present themselves as necessary for a robust system and we will refer to these with the terms *selectivity*, *consistency*, *stability* and *capacity*. A *selective* system produces PNG activations in response to a stimulus that are sufficiently specific to allow the unique identification of the stimulus. A *consistent* system is able to dependably produce PNG activations on every presentation of the stimulus. A *stable* system is able to maintain long-term representations in the form of structural PNGs that are capable of activation, and a system with good *capacity* allows a biologically plausible number of these structural representations.

Activated groups produce distinct spatio-temporal signatures within the flood of firing events generated by the network and the analysis of the firing response to stimuli should therefore allow the detection of PNG activation. Izhikevich has used such a technique to study the selectivity of the neural response to repeated stimulation [6]. In this experiment he tracked the evolution of polychronous groups in response to one of two input patterns and found that different groups were activated for each pattern, suggesting that the underlying structural groups might provide a unique long-term representation of each pattern.

Although this experiment provided some initial evidence in support of selective PNG activation, it did not address any of the other attributes necessary for a representational system based on PNGs. In addition, the method used in Izhikevich [6] for measuring PNG activation is not described, providing some hurdles to the reproduction of these results. The experiments described below employ a template matching technique for detecting PNG activation (methods outlined below and in more detail in a separate technical report [4]).

In the remainder of this report we will focus on the requirement for a *consistent* representational system, using the pattern-specific activation of polychronous groups to measure the dependability of the neural response to known stimuli. Polychronous groups exist in a competitive medium in which the group affiliation of individual neurons is constantly fought over [8] and this dynamic environment therefore calls into question the reliability of a representational system based on PNG activation. Although PNG activation is often described as “stereotypical” and “reproducible”, a specific PNG will not necessarily activate on every presentation of a triggering stimulus [6]. However, other PNG activations may result from the same stimulus and a stimulus-specific neural response consisting of some subset of the set of all stimulus-specific PNG activations may therefore occur with some consistency. In the following experiments we will assess the empirical probability of this stimulus-specific neural response given the presentation of a known stimulus. Does every stimulus presentation produce a relevant group activation, or only some presentations?

2 Methods

Twenty independent networks were created for these experiments, each composed of 1000 Izhikevich neurons (800 excitatory and 200 inhibitory) with parameters as described in [6]. The networks were matured for two hours by exposure

to 1 Hz random input generated by a Poisson process. Following maturation, the networks were trained on one of two input patterns or were left untrained. The current experiments reproduce the few known details of the repeated stimulation experiment described in [6], namely a twenty minute training period, and the use of an *ascending* or *descending* input pattern as the stimulus (see Fig. 2).

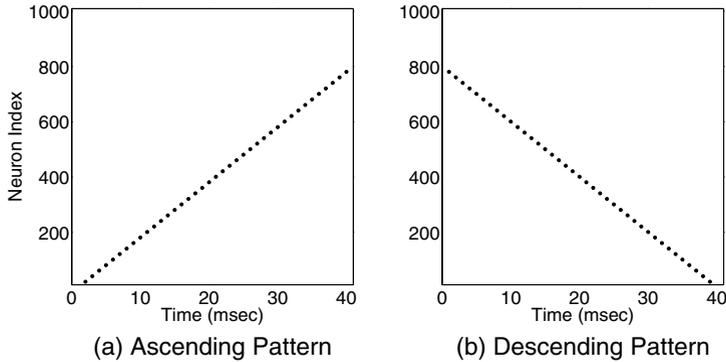


Fig. 2. The ascending and descending patterns: each spatio-temporal input pattern is composed of 40 firing events. Both patterns share the same neurons, differing only in the temporal order of their firing events.

The technique used by Izhikevich [6] for detecting PNG activations in the firing data was not described and therefore needed to be redeveloped for the current experiments. It was clear that this technique needed to discriminate pattern-specific PNG activations from unrelated PNG activations, and from other spiking events generated by the network. The original method was assumed to make use of the Izhikevich search algorithms [7] (see also [10]) to find structural PNGs in the network, suggesting the use of a template matching technique for the detection of PNG activation. The default behavior of these algorithms is to initiate a PNG search based on all combinations of three neurons in the network. However, we created a small modification that limits the search to triplet combinations that occur only in the training patterns. Any polychronous groups found by this modified algorithm are referred to as pattern-specific PNGs, as the activation of these groups is initiated by a firing event triplet that occurs in the input pattern.

The assumed template matching technique involves isolating PNGs from a trained network and using them as templates to probe for group activation. The technique is reproduced as follows: first, a network is trained with a specified input pattern and pattern-specific structural PNGs are isolated from the network at regular intervals; the isolated PNGs are then used as spatio-temporal templates to match the firing data. For convenience, the experiment is split into two phases: in an initial *training phase*, the network is repeatedly stimulated with the ascending or descending pattern at 5 or 25 Hz for twenty minutes; in the following *test phase* of the experiment, the network is stimulated with the

same ascending or descending pattern at 1 Hz, and pattern-specific templates isolated during the training phase are used to probe for group activation.

At one minute intervals throughout the training phase a search is initiated for structural PNGs that can act as pattern-specific templates. The search involves testing all triplet combinations (i.e. combinations of three firing events) from the input pattern for their ability to discover structural PNGs in the network [4]. However, not all PNGs will be found as the algorithm is limited for performance reasons to testing combinations of just three firing events.

The test phase involves scanning the stream of firing events generated by the stimulated network for template matches. For each temporal offset in the network firing data, each of the templates is matched in sequence and successful matches are saved to a file. A matching threshold of 50% means that at least half the firing events in each template must match the the firing time and the neuron fired in the network event stream (although the firing time is allowed a jitter of ± 2 milliseconds). Each successful template match provides evidence of PNG activation in response to the stimulus.

The use of a 1 Hz stimulation frequency in the test phase creates a well-defined temporal frame for each stimulus and its response. Stimulus onset occurs at $t = 0$ in each one second *response frame*, and the remainder of the frame has sufficient temporal length to include all of the firing events in the resulting neural response. A 1 Hz random background pattern is also presented throughout each test period. For a more detailed description of the methods see the accompanying technical report [4].

3 Results

Together the training and testing phases of the experiment produce a large set of data that supports multiple analyses. Training phase data provides a view of the evolution of structural PNGs in response to the stimulus, while test phase data provides a snapshot of the process of PNG activation. Figure 3 uses a combination of both datasets to show a selection of three matching templates following low-intensity (1 Hz) test stimulation of a network. These matching templates are sampled from a larger pool of pattern-specific templates that match PNG activations triggered by some triplet combination from the ascending input pattern. The first few firing events in each of the templates in Fig. 3 are therefore upward-sloping, reflecting the isolation of the template from a network trained on the ascending pattern. Each group consists of multiple convergent connections that support the propagation of neural firing across the members of the group before terminating at an inhibitory neuron (gray-filled circles).

Temporal alignment of just these first few firing events for all matching templates (and with all other firing events removed) produces sloping firing patterns that can be seen in Fig. 4. Recall that the first few firing events in each template (the initial triplet) correspond to the stimulus trigger that leads to PNG activation. The gray-scale intensity in this figure encodes the number of times the corresponding firing event acted as a trigger for the initiation of PNG activation, where activation was measured by the number of matching templates

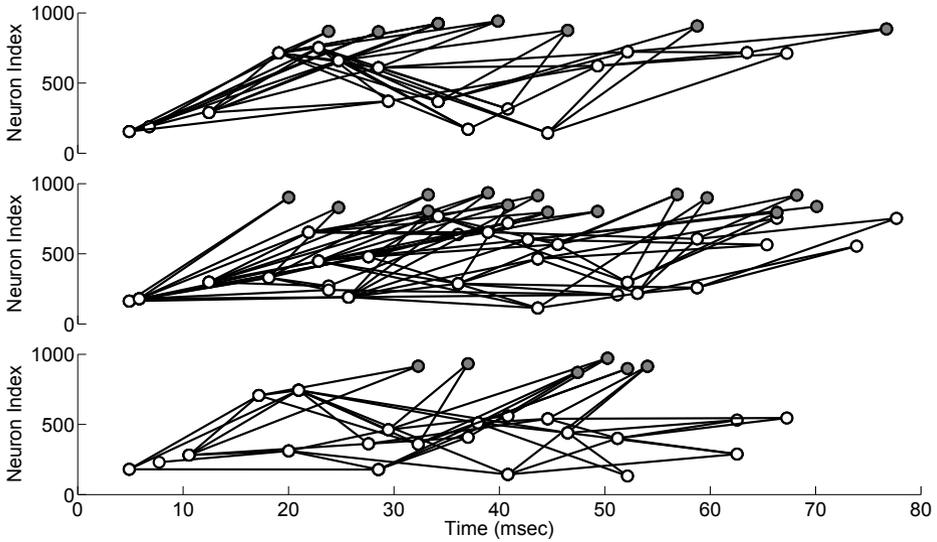


Fig. 3. A selection of three templates that match the firing data following stimulation with the ascending input pattern. The x- and y-axes for each template represent time in milliseconds and neuron index respectively (the y-axis is ordered so that inhibitory neurons are at the top of the graph). Nodes depict firing events generated by excitatory or inhibitory neurons and are drawn using either open circles (excitatory neurons) or gray-filled circles (inhibitory neurons). Lines between nodes represent causal connections between firing events.

accumulated across twenty independent networks. The figure therefore provides a picture of which of the input pattern firing events succeeded or failed at initiating PNG activation. Many of the forty firing events that make up each input pattern failed to initiate a responding group over the ten minutes (six hundred response frames) of the testing phase. Significantly, the majority of these failures are clustered in the later stages of the input pattern, suggesting that group response is concentrated on the early part of each stimulus presentation.

Nevertheless, the PNG activation response as a whole exhibits a high degree of consistency. Figure 5 shows the activation response of 40 networks (20 trained on the ascending pattern and twenty untrained networks) in the first 100 seconds of the 10 minute test run (only the first 100 of 600 response frames are shown in Fig. 5). The stimulus is presented at the start of each frame and any templates that match the firing events in the remainder of the frame are taken as evidence of PNG activation. Each row in Fig. 5 represents a single network and is divided into one hundred segments representing each of the one hundred response frames. The presence of a filled circle in each segment indicates the detection of a PNG activation response in the corresponding response frame. If there was no response, or the method was unable to detect the response, the segment is left empty.

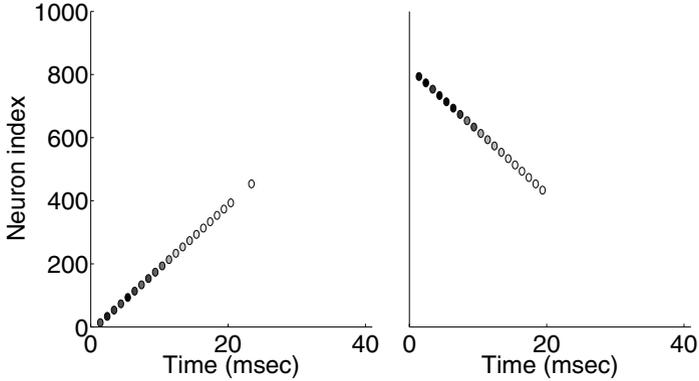


Fig. 4. The initial triplets from all templates that match the ascending input pattern (left) or the descending input pattern (right). The first three firing events from each matching template were extracted and aligned in order to show the coverage of the input pattern firing events. Firing events are represented by filled circles; the intensity of the fill color for each firing event represents the number of templates that matched PNG activations triggered by the event. This number, accumulated across twenty independent networks, is greatest in the early stages of each input pattern (darker fill color) and decreases in later stages of the input pattern (lighter fill color). The missing firing events in the later stages correspond to input pattern firing events that failed to initiate a group response during the test period.

The first 25 frames in this experiment used the ascending pattern, the next 25 used the descending pattern, the third group of 25 frames repeated the use of the ascending pattern, and in the final 25 frames no input pattern was provided (the null pattern). Using a combined pool of all templates to measure the PNG activation response, the twenty trained networks at the top of Fig. 5 show a consistent response to the ascending pattern but little or no response to the descending pattern or the null pattern. In contrast, the twenty untrained networks at the bottom of Fig. 5 show only sporadic activation and no apparent correlation with the type of input pattern. Comparing the activation response of the trained networks with the response of the untrained networks, we see a high degree of consistency in the response to the ascending pattern only where the network has been previously trained on the ascending pattern.

The PNG activation response to each stimulus presentation is assumed to occur in the early portion of each response frame, shortly after stimulus presentation at $t = 0$. Over this period, one or more PNG activations triggered by the stimulus have the opportunity to match the pool of PNG templates. Some insight into the temporal evolution of PNG activation is provided by computing the proportion of matches that occur at each temporal offset within the frame (the *template match ratio*) to produce an empirical measure of the likelihood of PNG activation at each offset. Firstly, each one second response frame is sliced into 1000 consecutive sub-frames and the number of template matches at each

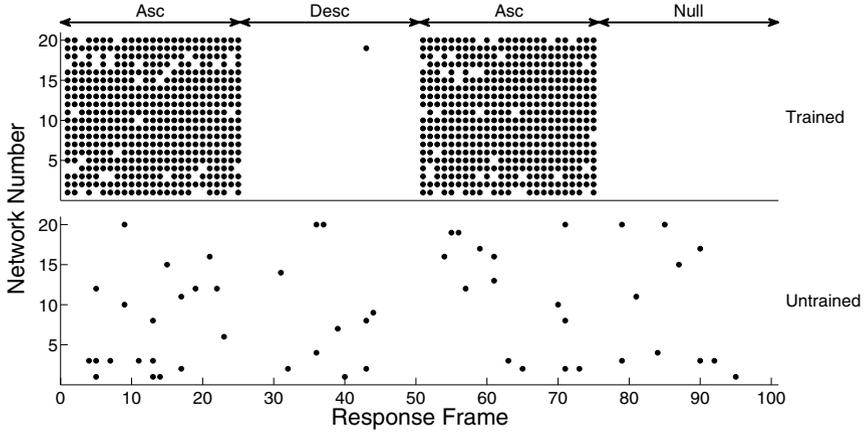


Fig. 5. The PNG activation response of twenty trained networks and twenty untrained networks over one hundred response frames. Trained networks were trained on the ascending pattern. A filled circle represents a positive response to the stimulus while an empty space denotes a lack of response. The stimulus for the first and third quarter of the one hundred frames was the ascending pattern and the stimulus for the second quarter was the descending pattern. No stimulus was provided in the fourth quarter (null pattern). The top figure shows the measured response for a network trained on the ascending pattern at 5 Hz and the bottom figure shows the result using an untrained network. The trained networks in the top figure were derived from the untrained networks in the corresponding row of the bottom figure.

one millisecond sub-frame is counted. The template match ratio for each offset is then computed by aggregating the number of matches for each offset across all response frames. Using this procedure we expect to see an isolated peak in the number of matches at a short delay following the stimulus at time $t = 0$, reflecting the transient activation of a responding PNG. However, due to limitations in the template matching method the delay can only be calculated to within half the length of each template (i.e. ± 15 milliseconds), depending on where on each template the match occurs.

Figure 6 shows the template match ratios for each network distributed over the first twenty sub-frames of each response frame. As predicted there is an isolated peak that consistently occurs in the first ten milliseconds following the stimulus. Within this small temporal window the likelihood of a template match typically reaches 50% or more, indicating that PNG activation is in full swing. As PNG activation comes to an end, the likelihood of a template match decreases to zero and remains at zero for the remainder of the response frame.

Although these positive results support the consistency of PNG activation, it is worth noting that the majority of templates are ineffective in matching the firing data. Here, we define an effective template as one that is able to match the firing data at least once during a ten minute period of stimulation with the corresponding input pattern. On average, just 32% of ascending templates and

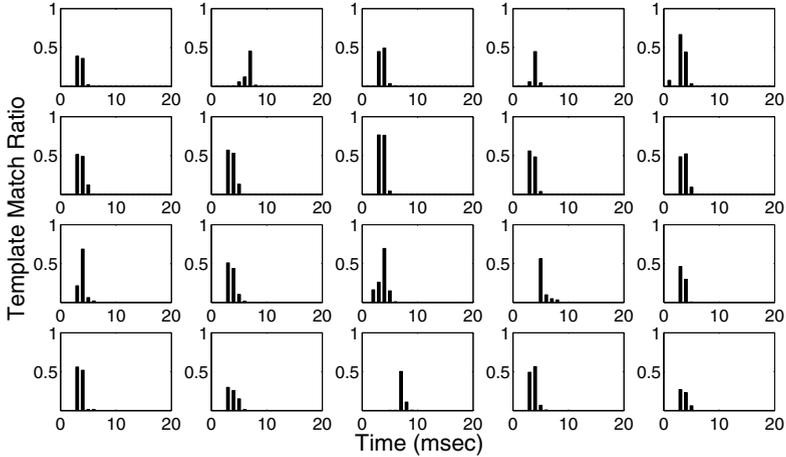


Fig. 6. Template match ratios distributed over each one second frame for each of twenty independent networks. The template match ratio was computed for each one millisecond slot in the response frame, accumulated over multiple frames. The response for each network is confined to the first ten milliseconds following the stimulus and therefore only the first twenty milliseconds of the frame are shown.

43% of descending templates were effective at finding a match (averaged across twenty independent networks). The template matching performance between networks is also very variable with some networks averaging as few as three matches in each response frame. In some frames, the evidence for PNG activation is based on a single template match suggesting that the template matching method is near to its sensitivity limit for some networks.

4 Discussion

The template matching method attempts to match spatio-temporal templates derived from the structural PNGs found in a trained network with the sequence of firing events that are produced when the network is stimulated with the same pattern. To ensure that template matches were pattern-specific, the selected templates were restricted to structural PNGs that were triggered by triplet combinations of the input pattern firing events. We can imagine that structural groups exist in the network that require larger, more complex, triggering patterns although it seems likely that the probability of finding groups with larger triggers decreases with the size of the triggering pattern. Templates that match the firing data such as those shown in Fig. 3 provide an impression of the corresponding PNG activations that occur in the milliseconds following each stimulus. However, looking at a selection of matching templates creates only a partial picture of the complex pattern of neural firing in response to spatio-temporal stimuli. Visualization of all of the PNG activations that are initiated by combinations of firing

events from the input pattern produces a complex graph in which individual PNGs interact and merge (results not shown).

Izhikevich [8] has proposed that competitive interactions occur between the polychronous groups in the network, with neurons that are shared by multiple PNGs synchronising their firing times with different polychronising pathways at different times. However, cooperative interactions are also possible in which firing events generated by separate PNG activations together produce the required spatio-temporal initiators for additional PNG activations. The emerging picture is one in which the activation response to complex stimuli is a composition of individual PNG activations that interact and merge in a complex manner.

Interestingly, all of the templates that found a match in the neural response were initiated by firing event triplets from just the early portion of the input pattern. This effect was found across all networks and for both the ascending and descending input patterns. A possible explanation is that competition during PNG formation for use of shared neurons creates an interference effect between early PNG activations and those that come later, with the earlier activating groups forming first and therefore dominating the available neural resources.

This explanation has implications for the maximum number of simultaneous activations that a network of a given size is able to support, and might in turn impact the maximum number of representations that can simultaneously be “held in mind” in a representational system based on polychronous groups. However, note that this explanation does not contradict the extraordinary potential capacity of a PNG representational system [6] because any potential limitation in the number of *simultaneous* activations supported by a representational system does not necessarily affect the network capacity i.e. the total number of representations that can be stored within the network.

Despite any interference caused by interactions between simultaneous activations, the template matching method provides good support for the consistency of a PNG-based representational system. Using a combined pool of all templates, one or more template matches are detected in almost every response frame, suggesting a consistent PNG activation response following each stimulus presentation. The best single template for each network is also able to show quite a high degree of consistency, although most individual templates match only rarely.

Computing the template match ratio for each one millisecond time-slot in the response frame shows that all matches are confined to a narrow temporal window following each stimulus presentation (see Fig. 6). This strong interaction between the time of the stimulus and the time of template matching supports the view that template matching reflects the causal relationship between stimulus onset and subsequent PNG activation. The template match ratio can also be computed at frame level (i.e. the proportion of matching frames), producing a value that reflects the empirical likelihood of PNG activation given the stimulus. With the combined templates, this likelihood value approaches certainty for many networks, although there is considerable inter-network variation in performance.

Together these results indicate a high degree of consistency in the PNG activation response following a stimulus. However, despite this consistent response there are occasional response frames where no neural response is detected, despite the presence of a known stimulus. The lack of a detectable response does not mean that PNG activation did not occur and may instead be due to limitations in the template matching method. Examination of the precise timing of the firing events in consecutive response frames shows considerable jitter in the spike times of PNG neurons between frames (results not shown). Competition for neural resources between activating groups may increase this jitter to the point where the corresponding template fails to match.

The lack of tolerance of the template matching method to temporal jitter is just one of the flaws of this method for detecting PNG activations. Although this technique is able to respond selectively to substantially different stimuli (e.g. discriminating between the ascending and descending patterns, or the ascending and null patterns in Fig. 5), the low matching threshold used in these experiments potentially allows templates to match unrelated spatio-temporal patterns. The template matching method may therefore have difficulty in resolving stimuli that are too closely related.

Another problem with the template matching method is that it treats matching as a local process when it is likely to be a global one. The neural response to a complex stimulus is a unique *set of PNG activations*; it is therefore the set as a whole and not individual activations that provide a unique signature of the stimulus. Given a set-oriented view of the neural response, if a single template happens to match a single PNG activation, does this provide good evidence of the presence of the stimulus? For example, two stimuli with partial overlap in their spatio-temporal firing patterns could both match the same template and may therefore not be individually resolvable. In recognition of a set-oriented view of the neural response, the template matching method makes use of a pool of templates that are able to detect multiple PNG activations. However, this method does not take into account the number of unique matches in each response frame and is therefore unable to counter the problem of overlapping stimuli.

Each of the templates generated in the training phase contribute to the time it takes to scan the firing data in the testing phase. It is therefore a problem that the majority of templates are ineffective, with less than half of the templates ever able to generate a match. Although the single best template for each network matches the neural response very consistently, the majority of templates that match at all do so only rarely. In addition, the number of matches in each response frame is sometimes very low suggesting that this method is close to the threshold for maximum sensitivity for some networks.

It is likely that Izhikevich [6] used a similar technique to show selectivity in the activation response, despite the flaws of the template matching method. The issues with this method, while limiting the scope and accuracy of the current results, do not invalidate our overall finding. Here we provide preliminary evidence for the consistency of PNG activation in response to stimuli, suggesting that polychronous groups may be able to meet at least one of the necessary

criteria for a representational system. The neural response to complex stimuli appears to involve multiple interacting PNG activations suggesting that an alternative method for measuring the neural response must treat any single PNG activation as only partial evidence in favor of a particular stimulus. Work is in progress on such an alternative technique that will address these limitations.

References

1. Abraham, W.C.: Metaplasticity: tuning synapses and networks for plasticity. *Nature Reviews Neuroscience* 9(5), 387–387 (2008)
2. Caporale, N., Dan, Y.: Spike timing-dependent plasticity: a hebbian learning rule. *Annu. Rev. Neurosci.* 31, 25–46 (2008)
3. Caroni, P., Donato, F., Muller, D.: Structural plasticity upon learning: regulation and functions. *Nature Reviews Neuroscience* 13(7), 478–490 (2012)
4. Guise, M., Knott, A., Benuskova, L.: Consistency of polychronous neural group activation supports a role as an underlying mechanism for representation and memory: detailed methods and results. Tech. rep., Dept of Computer Science, University of Otago, Dunedin (2013)
5. Hoffmann, H., Howard, M.D., Daily, M.J.: Fast pattern matching with time-delay neural networks. In: *The 2011 International Joint Conference on Neural Networks (IJCNN)*, pp. 2424–2429. IEEE (2011)
6. Izhikevich, E.M.: Polychronization: computation with spikes. *Neural Computation* 18(2), 245–282 (2006)
7. Izhikevich, E.M.: Reference software implementation for the Izhikevich model: minimal spiking network that can polychronize (2006), <http://www.izhikevich.org/publications/spnet.htm>
8. Izhikevich, E.M., Gally, J.A., Edelman, G.M.: Spike-timing dynamics of neuronal groups. *Cerebral Cortex* 14(8), 933–944 (2004)
9. Martin, S., Grimwood, P., Morris, R.: Synaptic plasticity and memory: an evaluation of the hypothesis. *Annual Review of Neuroscience* 23(1), 649–711 (2000)
10. Martinez, R., Paugam-Moisy, H.: Algorithms for structural and dynamical polychronous groups detection. In: Alippi, C., Polycarpou, M., Panayiotou, C., Ellinas, G. (eds.) *ICANN 2009, Part II. LNCS*, vol. 5769, pp. 75–84. Springer, Heidelberg (2009)