




Effect of network topology on neuronal encoding based on spatiotemporal patterns of spikes

Petra E. Vertes & Thomas Duke

To cite this article: Petra E. Vertes & Thomas Duke (2010) Effect of network topology on neuronal encoding based on spatiotemporal patterns of spikes, HFSP Journal, 4:3-4, 153-163, DOI: [10.2976/1.3386761](https://doi.org/10.2976/1.3386761)

To link to this article: <https://doi.org/10.2976/1.3386761>




View supplementary material 



Published online: 07 Sep 2010.



Submit your article to this journal 



Article views: 510



View related articles 

Effect of network topology on neuronal encoding based on spatiotemporal patterns of spikes

Petra E. Vertes¹ and Thomas Duke²

¹Cavendish Laboratory, University of Cambridge, Cambridge CB3 0HE, UK

²Department of Physics and Astronomy and London Centre for Nanotechnology, UCL, London WC1E 6BT, UK

(Received 20 November 2009; accepted 22 March 2010; published online 7 May 2010)

Despite significant progress in our understanding of the brain at both microscopic and macroscopic scales, the mechanisms by which low-level neuronal behavior gives rise to high-level mental processes such as memory still remain unknown. In this paper, we assess the plausibility and quantify the performance of *polychronization*, a newly proposed mechanism of neuronal encoding, which has been suggested to underlie a wide range of cognitive phenomena. We then investigate the effect of network topology on the reliability with which input stimuli can be distinguished based on their encoding in the form of so-called polychronous groups or spatiotemporal patterns of spikes. We find that small-world networks perform an order of magnitude better than random ones, enabling reliable discrimination between inputs even when prompted by increasingly incomplete recall cues. Furthermore, we show that small-world architectures operate at significantly reduced energetic costs and that their memory capacity scales favorably with network size. Finally, we find that small-world topologies introduce biologically realistic constraints on the optimal input stimuli, favoring especially the topographic inputs known to exist in many cortical areas. Our results suggest that mammalian cortical networks, by virtue of being both small-world and topographically organized, seem particularly well-suited to information processing through polychronization. This article addresses the fundamental question of encoding in neuroscience. In particular, evidence is presented in support of an emerging model of neuronal encoding in the neocortex based on spatiotemporal patterns of spikes.

[DOI: 10.2976/1.3386761]

CORRESPONDENCE

Petra E. Vertes: pv226@cam.ac.uk

Neurons in the cerebral cortex form a network in which information flows in the form of short pulses called spikes. Whether neurons represent information purely through their firing rates (rate code) or whether the precise timing of their spikes is also exploited (temporal code) has been the focus of a long-standing debate. The past 20 years, however, have seen an accumulation of evidence for the ability of neurons to fire precisely-timed spikes and many theoretical studies have explored the use of temporal codes in propagating and storing information (Abeles, 1991; Buzsaki *et al.*, 1994; Bienenstock, 1995; Hopfield, 1995; Miller, 1996; Diesmann *et al.*, 1999). In particular, much attention has been turned

to the role of synchrony and related phenomena such as the propagation of information along hypothetical pools of synchronously firing neurons or *synfire chains* (Abeles, 1991; Bienenstock, 1995; Ikegaya *et al.*, 2004). However, most studies so far have neglected the increasingly well-established (Swadlow, 1985, 1994) existence of differing time-delays associated with each axon. Incorporating axonal delays inevitably destroys the synchrony that synfire chains rely on but such delays have been shown, *in silico*, to give rise to an even richer phenomenon called “polychronization,” which has been proposed as a mechanism for neuronal encoding (Izhikevich, 2006).

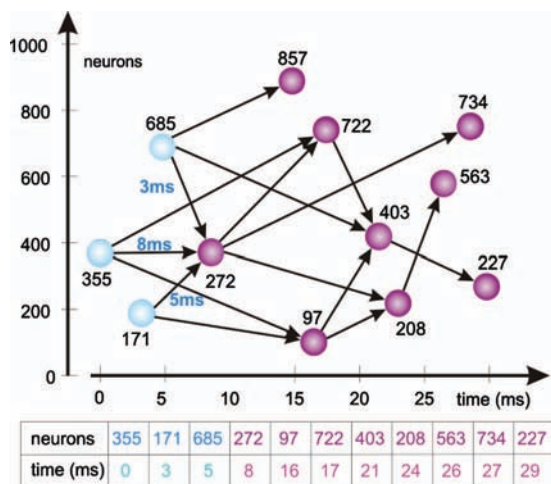


Figure 1. A polychronous group. Input-neurons receiving external stimulation are colored light blue. All other neurons fire due to coincident input (converging arrows) from presynaptic neurons with appropriate axonal delays. The table summarizes the associated firing pattern, listing the neurons due to fire along with their relative firing times.

Neurons typically require the convergence of several inputs in order to be driven to fire. In the presence of axonal delays, the spikes from two neurons firing with a precise time difference may in fact arrive simultaneously on a common postsynaptic neuron. As shown in Fig. 1, a set of input neurons’ distinct but well-timed spiking can therefore provoke a whole avalanche of activity provided that, at each level, appropriate conduction delays compensate for the difference in time between presynaptic activations. Groups of neurons with well-matched axonal delays that can sustain such a cascade of activity are called *polychronous groups* because they give rise to time-locked, but not synchronous, firing of neurons (Izhikevich, 2006).

In a network of spiking neurons with random input, the appearance of polychronous groups manifests itself by their spontaneous and recurrent “activation,” meaning that the characteristic spatiotemporal patterns of spikes associated with the groups can be observed in the overall spiking activity of the network. Although spatiotemporal firing patterns have been observed to repeatedly occur with ms precision both *in vivo* (Ikegaya et al., 2004; Abeles et al., 1993; Lindsey et al., 1997; Prut et al., 1998; Villa et al., 1999; Chang et al., 2000) and *in vitro* (Ikegaya et al., 2004; Mao et al., 2001; Volman et al., 2005), whether or not these indeed correspond to the activation of polychronous groups still remains to be rigorously established. Furthermore, the existence of these repeating spike patterns is currently the only experimental evidence in support of the theory of polychronization. From a theoretical point of view, however, the potentially high information content of such patterns immediately suggests that polychronization could be a key phenomenon for information processing in the brain.

Indeed, since each neuron can participate in many different groups depending on the context in which it fires, the total number of available groups is far higher than the number of neurons. It is this crucial finding that led Izhikevich to propose polychronous groups as an encoding mechanism, with groups representing memories formed by repeated experience of given stimuli and with the activation of groups corresponding to the recognition of these stimuli (Izhikevich, 2006).

Although polychronous groups are unlikely to develop by chance, they have been shown, in simulations, to emerge based on the self-organization of neurons through spike-time dependent plasticity (STDP), a learning process that has been observed in many systems both *in vivo* and *in vitro* (Caporale and Dan, 2008). The groups emerge because STDP naturally tends to strengthen links whose delays are appropriate for polychronization and to weaken links that might disrupt a cascade (Izhikevich et al., 2004).

Furthermore, it was found (Izhikevich, 2006) that training such a network by repeated presentation of two distinct input stimuli results in the emergence of two sets of groups, each triggered by and corresponding to one of the input signals. Subsequent presentation of the first input will preferentially elicit the activation of groups from the corresponding set and vice-versa. This early work was carried out on networks with random so-called *Erdős-Rényi* (ER) architecture and did not include a quantitative analysis of the phenomenon described above.

In this paper, we have set out to answer the following five basic questions to determine the plausibility of polychronization as a mechanism for neuronal encoding in cortical networks: (1) How reliably is this system able to discriminate between the two stimuli it has been trained with based on the activation of the polychronous groups associated with the recall cue? (2) Do these “correct” activations occur within a tight window of time following the presentation of the cue, as would be expected from a useful representation of the input signal? (3) How does the network’s capacity for discriminating between the two inputs degrade as the recall cues are corrupted? (4) Can the use of different network topologies significantly boost this performance? (5) How does the memory capacity of such a system scale with network size?

RESULTS

Discrimination task and network topology

Our model consists of a network of $N=1,000$ Izhikevich neurons (Izhikevich, 2003), each projecting to $k=100$ neighbors and incorporates axonal delays and STDP as described in the Materials and Methods section. For the discrimination task, the networks were trained with two alternating input stimuli (A and B) and were subsequently analyzed in order to identify the sets of polychronous groups formed during training in response to each stimulus. The method for identifying these groups is described in more detail in the Materials and Methods section. It is

important to distinguish the *existence* of a group from its *activation*. As shown in Fig. 1, the existence of a group can be deduced from the connections and time-delays of the network on a purely anatomical basis, in the absence of any firing. Indeed, the firing pattern corresponding to a group is simply the pattern of spikes, which one would observe if the first few neurons were to be stimulated at specified time intervals in an otherwise silent network. Once the network has developed these polychronous groups, however, appropriate external stimulation may result in the activation of the groups, meaning that their characteristic time-pattern of spikes is actually realized in the overall spiking activity of the network.

During the recall phase, the networks were periodically presented with a cue (stimulus A) in addition to a constant level of background noise. Input A consisted of consecutive stimulation of neurons 1–48 at times 1 ms, 2 ms, up to 48 ms while stimulus B involved similar consecutive stimulation of neurons 447–400. The network's response to the recall cue was recorded and scanned for activations of both “correct” and “incorrect” groups corresponding to signals A and B, respectively.

We found that levels of both A and B group activations were low and roughly equal during periods when no signal was presented to the network and that while B activations were uniformly spread over the recording period, a large majority of A activations started within the first 50 ms after presentation of the cue. We therefore quantify each network's performance with the ratio R_c between the total number of correct (A) and incorrect (B) group activations recorded within the short (150 ms) time-windows following the 60 presentations of the recall cue during the recording time. A large value of R_c thus requires the activations both to be of the correct kind and to occur at the correct times only.

To investigate the effects of network topology, we studied the benchmark ER random networks as well as so-called *scale-free*, *Watts–Strogatz* (WS), and *modular* networks. These four well-known classes of network topology are described in detail in following sections. In contrast to ER networks that can be defined based on their connection density alone, all other networks require one or more further parameters for their characterization. All data shown were obtained at parameter settings that were optimized as described in the Materials and Methods section. We note that, with these optimal parameters, both the WS and the modular networks used here displayed so-called small-world characteristics, featuring both highly local connectivity and short average distances between any pair of neurons. This is of particular interest in light of recent studies that point to mammalian cortical networks sharing these small-world characteristics (Humphries *et al.*, 2006; Hellwig, 2000; Song *et al.*, 2005). Further results for WS networks with different parameter settings shown in [Supplementary Material Fig. S1](#) confirm that

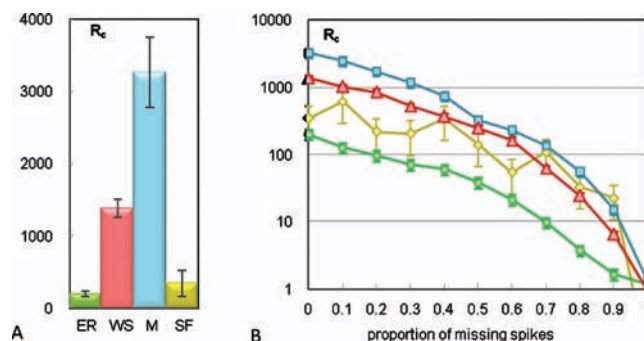


Figure 2. Comparison of recall performance R_c across different architectures. (A) From left to right, performance of ER, WS, modular, and scale-free networks. (B) Degradation of R_c (on a logarithmic scale) with increasing proportion of missing spikes in the recall cue for four different architectures: ER (green circles), WS (red triangles), modular (blue squares), and scale-free (yellow diamonds).

performance is clearly optimized in the region corresponding to a small-world regime.

Figure 2(A) compares performance in the discrimination task across the four different network topologies. Since the ratio R_c was large ($R_c \gg 1$) for all networks, it seems that the activation of appropriate polychronous groups can indeed serve to discriminate between two stimuli that the network has been trained with. However, a key requirement for reliable encoding is that relevant memories are retrieved, even when prompted with incomplete cues. As shown in Fig. 2(B), all networks' performance dropped away as increasingly incomplete cues were presented. In addition, the large error bars associated with scale-free networks indicate that their performance was very unreliable. Crucially however, both small-world networks consistently performed an order of magnitude better than random ER ones. Note that the performance of ER networks with intact recall cues was easily matched by small-world networks with highly degraded cues (up to 60% missing spikes). Our results suggest that the small-world architecture of brain networks is particularly well-suited and perhaps necessary for information processing through polychronization.

Differences between network topologies

In order to shed some light on the origin of these differences between the four networks topologies, we analyzed the number, the size distribution, and the structure of the groups for each network type.

In Table 1, we show the average number of groups identified for stimuli A and B in each of the four network topologies. These numbers vary greatly between topologies, contributing to the difference in performance between them but, as can be seen by comparing the first two columns in the table, they are roughly equal for different stimuli within a given architecture. The next two columns show the cumulative number of active A and B groups found within windows of 150 ms after each of the 60 presentations of stimulus A.

Table I. Number of groups found for stimuli A and B. The first two columns describe the average number of groups identified for stimulus A and stimulus B, respectively. The next two columns contain the average of the total number of groups found to be active in the 150 ms time-windows immediately following all 60 presentations of the recall cue.

Network	Groups			
	A	B	Active A	Active B
Random	38±5	37±5	650±90	3.1±0.4
WS	1,220±19	1,220±18	14,400±400	10.3±0.9
Modular	610±12	580±13	7,800±300	2.6±0.3
Scale-free	38±4	38±4	350±40	1.0±0.5

Furthermore, we analyzed the structure of the groups for each network type. In the first column of Table II, we report the average fraction of nodes in a group that are also part of the pool of 48 neurons receiving the stimulus, while in the second column we show the fraction of nodes in each group that are part of the “neighborhood” of 100 neurons surrounding the pool of neurons where the input arrives. In the case of modular networks, this neighborhood corresponds to the whole module in which the input is contained. As discussed later, the values in these first two columns and their relevance to the difference in performance apparent in Fig. 2 and Table I can be understood based on the differences in the density of connections in the subgraphs formed by the pools of 48 neurons receiving input. These values are reported in the last two columns of Table II. Finally, the size distribution of groups for each network topology is shown in Fig. 3.

Pattern localization

The precise identity of the neurons that are stimulated as part of the input signals does not affect the performance of ER or

Table II. Structure of the groups and input-subgraphs associated with each topology. The “stim” values in the first column correspond to the average fraction of neurons in a group—excluding the three initial input-neurons—that are also part of the pool of 48 neurons receiving the stimulation. Similarly, the “neb” values correspond to the fraction of neurons in a group that are also part of the neighborhood of 100 neurons surrounding this pool of neurons. The subgraph density, denoted by “sub. dens.” corresponds to the density of links within the subgraph formed by the pool of 48 neurons receiving the input. The average subgraph densities before and after STDP are shown in the third and fourth columns, respectively, and are extremely accurate estimates so the errors are omitted.

Network	Stim (%)	Neb (%)	Sub. dens. before (%)	Sub. dens. after (%)
Random	1.3±0.3	7.4±0.6	10	3
WS	26±4	75±4	68	16
Modular	45±7	61±7	82	16
Scale-free	0.8±0.2	6±0.2	6.5	2

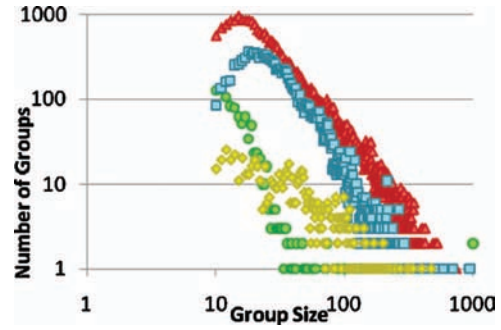


Figure 3. Distribution of group sizes. Distribution of the number of neurons constituting each group: ER (green circles), WS (red triangles), modular (blue squares), and scale-free (yellow diamonds). The roughly linear variation on this log-log plot potentially indicates a scale-free nature for these distributions. All values correspond to the numbers of groups of a given size, summed over ten different network realizations.

scale-free networks in which all neurons are equivalent when averaged over different realizations of the network architecture. However, other architectures are sensitive to this “localization” of the input. The standard stimuli used in the previous section and described in the Materials and Methods section were expected to be optimal for all architectures. In the context of WS and modular architectures, their defining characteristics were that the relevant neurons were spatially contiguous and that, for the modular network, the two inputs were contained in separate modules. Removing the constraint on the spatial adjacency of input-neurons had significant consequences for both modular and WS architectures.

Figure 4 shows the performance of both these networks with localized inputs as well as with spatially disparate stimuli involving the consecutive activation of a randomly chosen set of excitatory neurons. In this discrimination task, the recall cues were 60% incomplete (similar data for a range of values of cue corruption can be found in [Supplementary Material Fig. S2](#)). Although small-world networks significantly outperformed ER ones even with randomized input stimuli, it is worth noting that localization of inputs consid-




Figure 4. Effect of signal localization on performance with 60% incomplete recall cues. From left to right: ER networks, WS networks with both delocalized and localized inputs, and modular networks with delocalized and localized inputs.

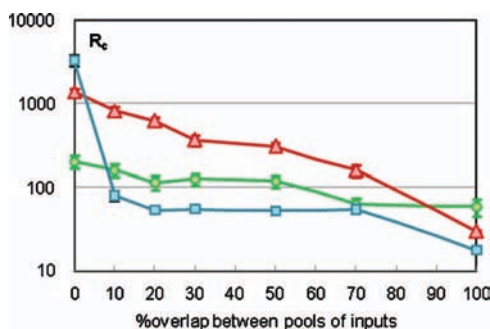


Figure 5. Degradation of recall performance with increasing overlap between pools of input neurons: ER (green circles), WS (red triangles), and modular (blue squares)

erably boosted performance and that, as discussed in following sections, many real cortical areas may in fact benefit from such a boost.

Memory capacity

The results described so far involved only two input signals, which were also maximally different in the sense that the pools of stimulated neurons were distinct for different inputs. In a qualitative investigation of the memory capacity of our system, we have found that network topology, once again, plays a crucial role. Figure 5 shows the degradation of performance for the small-world and ER networks with a growing proportion of overlap between the pools of neurons involved in the two training stimuli. Note that even for 100% overlap in the pools of neurons involved in stimuli A and B, the two stimuli themselves will still be different in that the neurons will receive input in a different order. As discussed in the next section, the performance of modular networks was dramatically decreased at only 10% overlap. However, WS networks again fared significantly better than ER ones, with a 70% overlap in the two pools of input-neurons yielding similar performance to that of ER networks with two completely distinct inputs. This suggests that a WS network of size N should be able to store significantly more memories of a given length than an ER (or highly modular) network.

In order to investigate precisely how the memory capacity of WS networks scales with network size, we have set a threshold in performance above which we considered a network successful in recalling an input pattern it had been trained with. This enabled us to evaluate the minimum network size required for successfully storing m distinct memories or input patterns. As shown in Fig. 6, the required network size increased in a sublinear fashion as the number of input signals to be “remembered” was increased. In other words, modest increases in network size yielded large increases in memory capacity. See the Materials and Methods section and the SI text for details on how Fig. 6 was obtained and for a demonstration that varying the (somewhat arbitrary) choice of the threshold value of R_c has no impact on our conclusions.

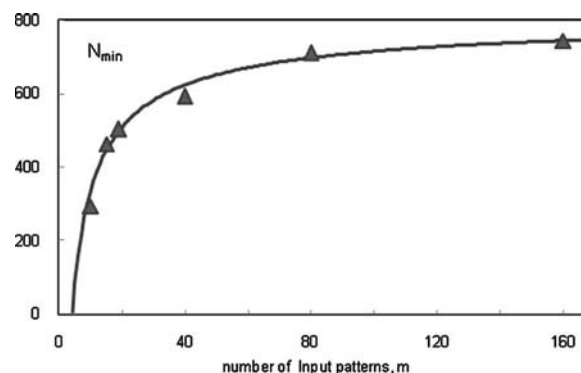


Figure 6. Memory capacity. Evolution of the minimum network size N_{\min} required to successfully store and recall an increasing number m of memories (for WS networks, with threshold $R_c > 100$).

DISCUSSION

Polychronization as a neuronal encoding mechanism

The only requirements for polychronization to occur in a network are the existence of time-delays and STDP. Since both of these are well-established properties of mammalian cortical networks, it was noted that polychronous groups are expected to emerge in the cortex, unless an unknown mechanism exists to suppress them (Izhikevich, 2006). These groups therefore seem a natural mechanism for generating the precise spiking patterns observed *in vivo* (Ikegaya *et al.*, 2004; Abeles *et al.*, 1993; Lindsey *et al.*, 1997; Prut *et al.*, 1998; Villa *et al.*, 1999; Chang *et al.*, 2000) and *in vitro* (Ikegaya *et al.*, 2004; Mao *et al.*, 2001; Volman *et al.*, 2005) and polychronization thus seems a strong candidate for a neuronal assembly-based encoding. It is important to note however that outside of the existence of such repeating spike patterns, there is currently no experimental evidence validating either the existence of polychronous groups or their potential role in cortical information processing.

From a theoretical point of view, however, polychronous groups display some key characteristics that make this proposition qualitatively plausible. As described in Izhikevich *et al.* (2004), there exists constant competition and association between different groups just as different stimuli or experiences compete for our attention and different ideas or experiences often get associated to one another in memory. Further characteristics of polychronous groups that tie in with our experience of perception, memory formation, and recall are that the system can be made insensitive to input scaling (Izhikevich, 2006) and that not only does the network require training before groups emerge but a prolonged period without the activation of a given group will result in its dissolution. Yet, in spite of this seeming fragility, polychronous groups are extremely robust in the face of biologically realistic synaptic turnover (Izhikevich *et al.*, 2004), which may, for instance, arise as a result of structural plasticity (Butz *et al.*, 2009).

Here, we have shown that the activation of appropriate polychronous groups can indeed serve to discriminate between different stimuli that the network has been trained with. However, in random ER networks, the reliability of this discrimination quickly drops away as increasingly incomplete cues are presented, thus calling into question the plausibility of polychronization as a neuronal encoding mechanism. Our main result is that the network's topology has a dramatic influence on performance with small-world networks yielding an order of magnitude improvement, including in the case of heavily degraded input stimuli.

Note that real cortical networks are likely to display a combination of properties, which may include a highly skewed or possibly scale-free degree distribution and modularity as well as small-world characteristics (Reijneveld *et al.*, 2007; van den Heuvel *et al.*, 2009). The precise anatomy of cortical neural networks, however, has not yet been sufficiently firmly established for us to test polychronization in that context. Instead, we have focused on studying the performance of polychronous encoding in four simple, commonly used network types in order to gauge the effect of different topological features on polychronization. Our aim was to identify network characteristics that could make polychronization sufficiently reliable for it to be a plausible encoding mechanism. In addition, we hoped that the network properties, if any, which would successfully support polychronization would also be biologically plausible in the context of cortical networks. As discussed in the sections below, both these hypotheses were borne out by our study.

The four architectures we chose to investigate reflect four broad categories of networks that are commonly used in network-science literature and can be characterized by widely used network measures such as clustering coefficient, path length, small-world coefficient, and modularity. The values of these network measures for each topology, both before and after training, are reported in [Supplementary Material Table S1](#) and may give us some insight as to the topologies and encoding strategies that could be used in real cortical networks, as discussed below. Once a full characterization of cortical networks is possible at a cellular level, our methods may also be applied to real cortical network topologies, providing a detailed test of the hypothesis that polychronous firing patterns encode information in the cortex.

Scale-free networks

Scale-free networks are characterized by a power law decay in degree distribution. These networks have the desirable property of being very robust to random removal of nodes or links. In addition, many large, real-world networks are known to be scale-free. In our system, scale-free topologies yield, in some cases, dramatic improvements in performance on a simple discrimination task. However, in some 25% of cases, the network showed no or very little activation in response to a stimulus it had been trained with, resulting in

very large error bars (Fig. 2). Preliminary investigation indicates that this is largely due to the fact that, depending on the network realization, there may sometimes be insufficient synapses in the area relevant to the stimulus to support the emergence of corresponding groups. The optimization of parameters has allowed us to reduce the occurrence of such failures; however, we have found that performance quickly drops away as reliability increases. Further optimization of the parameters may reveal a narrow range where the performance is reliable but still improved when compared to ER networks, although it is unclear how such a delicate optimization of parameters may be achieved in nature.

Watts–Strogatz networks

A Watts–Strogatz type network (Watts and Strogatz, 1998) such as the one we implemented is characterized by a parameter β , which allows the smooth tuning of network topology from the fully ordered to the random state. Intermediate values give rise to a range of so-called small-world topologies, which preserve some important attributes from both ends of the spectrum. Similarly to random ER networks, they display a short path length between any pair of nodes. However, they also exhibit a high clustering coefficient, as the ordered $\beta=0$ networks do. This coefficient measures the number of shared nearest neighbors for each pair of neighboring nodes and thus reflects the increased proportion of local rather than long-range links.

We have observed that the system's performance is spectacularly enhanced in small-world networks, such that a 60% degradation of the recall cue still yields results comparable to that of ER networks with intact cues, and discernable discriminatory ability remains even at 80–90% degradation. This can be intuitively understood by noticing that a high clustering coefficient favors the formation of polychronous groups by increasing the likelihood of interaction between the neighbors of any neuron that has fired. At the same time, the short path length and the existence of many equivalent short paths ensures the rapid and wide spreading of activity through the network and thus the use of its full representational capacity.

Our results are in line with recent studies on artificial neural networks of binary neurons performing associative memory tasks (Morelli *et al.*, 2004; Bohland and Minai, 2001). The exceptional performance demonstrated here in biologically plausible networks performing polychronization is especially interesting in view of several recent studies that have discovered hallmarks of small-world topology in neural networks. These include the discovery of small-world network properties in the medial reticular formation of the vertebrate brain (Humphries *et al.*, 2006), as well as the preferential short-range linking (Hellwig, 2000) and high clustering (Song *et al.*, 2005) of synaptic connections in the rat visual cortex. These small-world properties of neural net-

works thus seem well-suited for supporting polychronization and are instrumental in establishing it as a plausible mechanism for neuronal encoding.

Modular networks

The modular networks we investigated are composed of ten loosely interconnected, dense modules, which are in turn subdivided into ten even denser submodules. Given the fixed overall connectivity of 10%, the parameters determining the density of connections between and inside modules were optimized for the case where the two inputs are segregated in different modules. In fact, such optimization leads to pathological parameters (see Materials and Methods section) with zero connectivity between modules. We have instead used close parameter values maintaining some, albeit sparse, between-module connections.

These networks display distinctly small-world characteristics (see [Supplementary Material Table S1](#)) and perform twice as well as the WS ones, as shown in Fig. 2. However, the extremely high density of connections inside the modules causes interference between input stimuli located inside the same module, thus reducing the area's representational capacity. This results in very poor performance when the two stimuli are made to overlap by as little as 10% (see Fig. 5) since the network storing and recalling the stimuli is effectively reduced to the small, highly connected module containing both inputs. As the intermodule connectivity is increased at the expense of intramodule connections, the performance for segregated stimuli drops away while the degradation associated with overlapping and other suboptimal stimuli is tempered. We therefore suggest that the optimal cortical architecture for the purposes of polychronization would be modular on a large scale with different modules processing different types of input. This conforms to our current picture of the large-scale hierarchical organization of the cortex as well as the existence of different cortical areas for different sensory modalities. At a lower-level, however, further hierarchical modularity would only be advantageous if each submodule were dedicated to detecting a single stimulus type. Where this is not the case, we suggest a simpler, more homogenous small-world topology within each large-scale module.

Small-world networks are energy efficient

Enhanced performance can also be achieved for ER networks, provided the degree k of each node is increased ([Supplementary Material Fig. S5](#)). In order to equal the performance of WS networks, random networks require each neuron to make $k=300$ synapses instead of the usual connectivity of $k=100$ used elsewhere in this paper. Increasing the number of synapses is, however, energetically costly in a real neural network and small-world networks are therefore preferable, offering good performance at much lower cost. Furthermore, small-world networks involve more local connections,

which are also energetically favorable and easier to develop since they are shorter. When considering the energy cost corresponding to an encoding mechanism, one must also factor in the cost of producing and recovering from spikes since it has been shown that spiking accounts for a very substantial portion of the brain's metabolism ([Attwell and Laughlin, 2001](#)). The average spiking rate of neurons during memory recall is roughly equal for all small-world network architectures (~ 5.5 Hz) but is almost double that (~ 9 Hz) for ER networks, making them highly unfavorable from an energetic point of view. Our findings are in line with previous studies on the attractive properties of small-world networks for information transfer at low wiring cost [for review see [Bassett and Bullmore \(2006\)](#) and references therein].

Small-world networks make use of topography

Interestingly, the small-world architectures continue to substantially outperform random ER networks in the discrimination task even when the input stimuli used involve randomly selected neurons not localized to an area with above-average interconnection (Fig. 4 and [Supplementary Material Fig. S2](#)). This effect is especially marked for highly corrupted recall cues. However, localized input signals yield a significant further increase in performance, implying that such a system would greatly benefit from the topographical arrangement known to exist in many sensory and motor cortical areas. Indeed, in many such areas, neurons are ordered according to similarity in their receptive fields such that, for example, the touch of a finger on our arm would result in localized stimulation of predominantly contiguous sets of neurons in the primary somatosensory area. In contrast to small-world networks, random ER architectures would derive no benefit from such an organization.

Mechanism underlying variation in performance

After having examined each network topology in detail, let us make some general remarks about the topological characteristics which determine a network's performance. As shown in Table II, the density of connections in the subgraph corresponding to the pool of 48 input-neurons is one of the key factors in increased network performance. Indeed, the higher the subgraph density, the more likely it is that neurons stimulated by a particular signal will share enough neighbors with well-matched time-delays to support a cascade of firings. High subgraph density therefore leads to more groups, which translate into more group activations in response to the correct cue. This is consistent with our observation that increasing the overall connectivity of a graph also boosts performance. At fixed levels of overall connectivity however, higher subgraph densities are achieved by more strongly clustered network architectures such as the modular or WS networks. Again, this is consistent with our finding that delocalization of inputs in such networks leads to a drop in performance as the subgraphs become less dense. Note, how-

ever, that overly high levels of ordering on the other hand can lead to pathological behavior. Indeed, for WS networks with very low rewiring parameters the performance ratio collapses because the system becomes oversensitive and a large numbers of both A and B activations occur at all times, including when no stimulus is present. Interestingly, the same oversensitivity is not observed for the small, highly ordered modules that lead to optimal performance in the case of modular networks.

Beyond the effects of local connectivity, it is important to note the role of global topology as illustrated by the case of modular networks. Despite having the highest subgraph density and clustering coefficient, modular networks yield less groups and therefore fewer correct activations than WS networks. As previously described, the network is effectively subdivided into dense modules, each of which has sharply decreased representational capacity. This is consistent with our finding that memory capacity increases superlinearly with network size, at least for WS networks. Note that, as shown in Table II, a very large fraction of each group in the modular network involves the same neurons as those contained in the original stimulus. This implies that during training, various potential groups are in intense competition with each other—each trying to achieve optimal synaptic weights to favor the neuronal firing order they require. In addition, during recall, the refractory periods that follow each neuron's spiking when a group is active are likely to interfere with the activation of other groups.

Finally note that, as shown in Fig. 3, group sizes (defined as the number of neurons constituting each group) may possibly follow power law distributions. It may be interesting, in later work, to develop a formal analysis of the number and size distribution of groups expected to arise from varying subgraph densities. Furthermore, the origin and potential significance of these power laws are also interesting areas for further exploration. Potential parallels with the scale-free distributed neuronal avalanches that have been found experimentally by local field potential measurements in cultured and acute cortical slices may be especially revealing as these have been widely interpreted to have significant implications for information processing (Beggs and Plenz, 2003, 2004).

Computation with polychronous groups

It was noted early on (Izhikevich, 2006) that polychronization as a common coding strategy across all modalities could ensure easy integration of cross-modal input in association areas. Here, we have shown that for small-world networks the number of neurons required for reliable recall increases in a highly sublinear fashion with increasing number of input patterns (Fig. 6), such that the network can form and activate appropriate polychronous groups even in the case of each neuron serving, in parallel, as an input neuron. This supports the plausibility of hierarchically organized cortical regions with higher areas forming high-level groups that represent




Figure 7. Schematic representation of high-level computation with polychronous groups. The coactivation of groups corresponding to the black and white piano keys (in visual cortical areas) and those corresponding to the musical notes (in auditory areas) triggers the activation of higher-level groups in association areas corresponding to the concepts of “piano” and that of a particular song. The cascading activation of the groups corresponding to this song trigger the recall of the subsequent notes in the song. This prediction, represented by the notes depicted in light gray, can then be compared with the actual sensory input flowing up the auditory pathway.

and are triggered by synchronous or time-locked input from several groups active in lower regions (see Fig. 7). Furthermore, we have shown that in small-world networks stimuli are easily recognized even based on highly degraded recall cues. This not only implies that polychronization is robust to noise but it further suggests that the system is able to recognize the beginning of stimuli it has been trained with and predict the rest of the stimulus or even other stimuli that are expected to follow, based on the cascading nature of group activations. We speculate that feeding back these predictions from higher-level areas to lower regions where inputs are streaming in could form the basis of cognitive functions, which require constant comparison between expected and actual sensory input.

In summary, we have shown that polychronization possesses several properties fundamental to viable encoding mechanisms. We have also shown that the network's topology has a dramatic influence on its performance in discrimination tasks and that small-world networks dramatically reduce the energy expended, in both spiking and synapse formation, for a given performance level. Furthermore, small-world architectures set realistic requirements on the nature of input stimuli, performing best with topographic inputs. These topological constraints on both networks and stimuli are in line with recent evidence on the architecture of real cortical networks, suggesting that these are exceptionally well-suited for supporting polychronization as a mechanism for encoding and retrieving stimuli. Further work is required to analyze the system's performance with increasingly realistic cortical networks and inputs, including the hierarchical scenario of Fig. 7. As a first step, two-

and three-dimensional small-world networks as well as “spatially embedded” networks where conduction delays are roughly proportional to spatial separation would be worth investigating.

MATERIALS AND METHODS

Neuron and network model

Our networks consist of $N=1,000$ neurons, 80% of which are excitatory and 20% of which are inhibitory, preserving the ratio found in the mammalian neocortex (Braitenberg and Schuz, 1991). Inhibitory interneurons are only allowed to project to excitatory neurons and, unless otherwise stated, the number of projections per neuron is $k=100$. While inhibitory synapses have fixed and equal weights, the excitatory synapses are plastic and evolve following the STDP rule as described in Song *et al.* (2000) and used in Izhikevich (2006). Weights are manually clipped at 10 mV. Delays are modeled to be constant over time and are randomly chosen integer values between 1 ms and 20 ms for each excitatory synapse while inhibitory synapses all have 1 ms delays.

To model each neuron, we have adopted Izhikevich's equations [see SI or Izhikevich (2003)], which, under appropriate choice of parameters, have been shown to accurately simulate 20 fundamental features of spiking activity associated with various types of biological neurons (Izhikevich, 2004). While all the assumptions above are biologically plausible, we have also had to make some simplifications for the sake of computational tractability. In particular, as in Izhikevich (2006), we consider the coincidence of two spikes on a neuron as a sufficient condition for inducing it to fire although in reality many more simultaneous inputs are required.

Network architectures and parameters

For all network architectures, care was taken to avoid self-links and repetition in links. For ER networks, the $k=100$ projections of each neuron were chosen at random.

The modular networks consist of ten modules of 100 neurons each, defined by an intramodule pairwise connection probability p . These are further subdivided into ten submodules with interior connection density $q > p$. The remaining connections available in order to reach a degree $k=100$ for each node are assigned as links between different modules. Each submodule contains eight excitatory neurons and two inhibitory ones to preserve the overall ratio across the entire structure. For the purpose of optimization, q was increased in units of 0.1 in the range $[0.2, 0.6]$ and then in units of 0.05 up to 1.0 while p was held at 0.2. The same sweep was carried out for p with q held near its best value at 0.9. The parameters used for the figures in this paper were not the optimal but pathological parameters $p=q=1.0$ but the close values of $p=0.8$ and $q=0.9$.

The WS networks are based on the one-dimensional Watts–Strogatz model (Watts and Strogatz, 1998), where nodes are aligned along the perimeter of a circle and each

node is linked to its k nearest neighbors ($k/2$ in each direction). In our model, the excitatory and inhibitory neurons are aligned along separate, concentric circles. Each excitatory neuron is linked to 80 excitatory neurons as previously described but also to 20 inhibitory neurons adjacent to each other on the inner circle. The next three excitatory neurons along the outer circle are also linked to the same 20 inhibitory ones but every fourth excitatory neuron shifts by one unit its set of target inhibitory neurons, so that the pattern of linking is circularly symmetric. Correspondingly, each inhibitory neuron is linked to 100 adjacent excitatory ones and the next inhibitory neuron along the circle has its pool of excitatory targets shifted by four units. This regular pattern of connections is then relaxed, as defined by Watts–Strogatz, through the random rewiring of a proportion β of the existing links. By varying the parameter β in units of 0.1 between 0 and 1.0, the region $0.1 < \beta < 0.5$ was found to be optimal. All data presented in the main text used $\beta=0.3$. Further results for different parameter settings are shown in [Supplementary Material Fig. S1](#).

For the scale-free networks, which cannot, by definition, respect $k=100$, the degree k of each node is set to be above a minimum value k_{\min} and is drawn from the distribution $P(k) = Ak^{-\alpha}$ where A was chosen such that for a given α and k_{\min} , the probability of k exceeding about 650 vanishes. The neighbors of each neuron are then chosen randomly. For optimization, the value of α was chosen from among 3, 2.2, 2, 1.7, and 1.5 and $k_{\min}=50$ as well as $k_{\min}=20$ was tested for each of these. For the optimal $\alpha=2$, we also tested $k_{\min}=10, 15, 17, 18, 19, 20$, and 50. All data presented used $\alpha=2$ and $k_{\min}=19$.

Network measures

All network measures reported in [Supplementary Material Table S1](#) were calculated using the MATLAB BGL package (http://www.stanford.edu/dgleich/programs/matlab_bgl/).

As described in the relevant documentation, this package relies on methods published in Fagiolo (2007) for the calculation of clustering coefficients and a generalized algorithm for calculating modularity in directed networks first published in Leicht and Newman (2008).

Training and inputs

We begin with every excitatory and inhibitory synaptic value set to $s=6$ and $s=-5$ respectively. We provide random input by injecting 20 mV to a randomly selected neuron every ms for 4 h (model time). This is followed by a training period during which, in addition to this background noise, the network is presented with two distinct coherent stimuli of excitatory neurons alternating every second. The training period lasts 2 h (model time). For the basic discrimination task, stimulus A consisted of consecutive stimulation of neurons 1–48 with 50 mV at times 1 ms, 2 ms up to 48 ms. Stimulus B involved similar consecutive stimulation of neurons 447–400 during the first 48 ms of the next second.

Definition and identification of polychronous groups

In general, polychronous groups can be defined only with respect to a specific number of “anchor” neurons whose initial pattern of activation results in the cascading activity of the group. The anchor neurons can be understood as the first “layer” of neurons receiving external input directly. To each configuration of anchor neurons, there could potentially belong a group, so restricting the set of anchors considered confines us to studying only a sample of the existing polychronous groups. As in the original work on polychronization (Izhikevich, 2006) we have focused on groups originating from anchors composed of three neurons only, as larger numbers of anchor neurons render the task of identifying polychronous groups computationally unfeasible. Further, in the context of groups emerging in response to coherent stimuli, the groups that are said to correspond to a particular input stimulus are found by choosing as anchors every possible triplet of neurons present in the input and preserving the timing difference in their firing (Izhikevich, 2006). To identify all possible groups corresponding to a given stimulus, we consider every eligible combination of triplets of neurons and record potential cascades resulting from the appropriate initial input to the anchor neurons in an otherwise silent network.

Detection of polychronous groups

After the training period, we present a single recall cue every second for a period of two min (in addition to the background noise). We record the first minute of activity and search for activated groups corresponding to both stimuli. We say that a group is activated if 50% of its excitatory neurons fire with the appropriate timing (± 1 ms). Note that raising this threshold to 70% still yields the same picture in terms of the differences in performance between the four network topologies. About 95% of spikes in the network’s activity are “noise” and do not correspond to the activation of any group (Izhikevich, 2006), however, active groups can be detected by comparing the templates of all previously identified polychronous groups or firing patterns to the recorded activity of the network. The number of active groups for each stimulus is counted cumulatively over the 60 signal presentations.

Corruption of the cue, overlapping and multiple inputs

Corruption of the recall cue is implemented by omitting a percentage p of spikes normally present in the input signal used as a cue and replacing these spikes with a 1 ms pause. The omitted spikes are randomly chosen in each run and all results presented are averaged over a hundred runs. In the section on overlapping pools of neurons, the set of neurons corresponding to signal B (normally neurons 447–400) is chosen instead so as to have a given overlap with the pool of neurons (1–48) used for signal A. A 100% overlap implies that signal B involves the same neurons (1–48) but firing patterns consist of reverse-order consecutive stimulations of

neurons 48–1. For multiple input patterns, smaller pools of only 20 contiguous excitatory neurons are used for each input. The sets of neurons are equally spaced over the entire range of 800 excitatory neurons. Each input signal involves stimulation of one of the 20 neurons every ms, in random order. Results were gathered for a range of network sizes ($N=400, 600, 800, 1,000, 1,200$) and various numbers of input patterns ($m=16, 20, 40, 80, 160, 320$). Here, the performance ratio R_c was taken to be the ratio between the number of group activations corresponding to the signal being recalled and the number of erroneous activations corresponding to another, randomly chosen, stimulus that the network has been trained with. Plotting the performance versus overlap size for each combination of (N, m) yields a complex plot (Supplementary Material Fig. S3) with six clear trendlines corresponding to the six values of m . Setting a threshold in performance ($R_c > 100$) and taking its intersection with the trendlines yields the minimum values of network size N_{\min} corresponding to successful recall for each value of m (Fig. 6; see also Supplementary Material Fig. S4).

ACKNOWLEDGMENTS

We thank Eugene Izhikevich, Danielle Bassett, and Naaman Tammuz for helpful advice and discussion. The research was designed by T.D. and P.E.V. and was carried out by P.E.V. who was funded by an EPSRC studentship. The key disciplines relevant to this work are Network Science, Complex Systems, and Neuroscience.

REFERENCES

- Abeles, M (1991). *Corticonics*, Cambridge University Press, Cambridge.
- Abeles, M, Vaadia, E, Bergman, H, Prut, Y, Haalman, I, and Slovin, H (1993). “Dynamics of neuronal interactions in the frontal cortex of behaving monkeys.” *Concepts Neurosci.* **4**, 131–158.
- Attwell, D, and Laughlin, SB (2001). “An energy budget for signalling in the grey matter of the brain.” *J. Cereb. Blood Flow Metab.* **21**, 1133–1145.
- Bassett, D, and Bullmore, E (2006). “Small-world brain networks.” *Neuroscientist* **12**, 512–523.
- Beggs, JM, and Plenz, D (2003). “Neuronal avalanches in neocortical circuits.” *J. Neurosci.* **23**, 11167–11177.
- Beggs, JM, and Plenz, D (2004). “Neuronal avalanches are diverse and precise activity patterns that are stable for many hours in cortical slice cultures.” *J. Neurosci.* **24**, 5216–5229.
- Bienenstock, E (1995). “A model of neocortex.” *Networks* **6**, 179–224.
- Bohland, JW, and Minai, AA (2001). “Efficient associative memory using small-world architecture.” *Neurocomputing* **38–40**, 489–496.
- Braitenberg, V, and Schuz, A (1991). *Anatomy of the Cortex: Statistics and Geometry*, Springer, Berlin.
- Butz, M, Wörgötter, F, and van Ooyen, A (2009). “Activity-dependent structural plasticity.” *Brain Res. Rev.* **60**, 287–305.
- Buzsaki, G, Llinas, R, Singer, W, Berthoz, A, and Christen, Y (1994). *Temporal Coding in the Brain*, Springer, New York.
- Caporale, N, and Dan, Y (2008). “Spike timing-dependent plasticity: a hebbian learning rule.” *Annu. Rev. Neurosci.* **31**, 25–46.
- Chang, E, Morris, K, Shannon, R, and Lindsey, B (2000). “Repeated sequences of interspike intervals in baroresponsive respiratory related neuronal assemblies of the cat brain stem.” *J. Neurophysiol.* **84**, 1136–1148.

- Diesmann, M, Gewaltig, MO, and Aertsen, A (1999). "Stable propagation of synchronous spiking in cortical neural networks." *Nature (London)* **402**, 529–533.
- Fagiolo, G (2007). "Clustering in complex directed networks." *Phys. Rev. E* **76**, 026107.
- Hellwig, B (2000). "A quantitative analysis of the local connectivity between pyramidal neurons in layers 2/3 of the rat visual cortex." *Biol. Cybern.* **82**, 111–121.
- Hopfield, J (1995). "Pattern recognition computation using action potential timing for stimulus representation." *Nature (London)* **376**, 33–36.
- Humphries, M, Gurney, K, and Prescott, T (2006). "The brainstem reticular formation is a small world, not scale-free, network." *Proc. R. Soc. London, Ser. B* **273**, 503–511.
- Ikegaya, Y, Aaron, G, Cossart, R, Aronov, D, Lampl, I, Ferster, D, and Yuste, R (2004). "Synfire chains and cortical songs: temporal modules of cortical activity." *Science* **304**, 559–564.
- Izhikevich, E (2003). "Simple model of spiking neurons." *IEEE Trans. Neural Netw.* **14**, 1569–1572.
- Izhikevich, E (2004). "Which model to use for cortical spiking neurons?" *IEEE Trans. Neural Netw.* **15**, 1063–1070.
- Izhikevich, E (2006). "Polychronization: computation with spikes." *Neural Comput.* **18**, 245–282.
- Izhikevich, E, Gally, J, and Edelman, G (2004). "Spike-timing dynamics of neuronal groups." *Cereb. Cortex* **14**, 933–944.
- Leicht, EA, and Newman, MEJ (2008). "Community structure in directed networks." *Phys. Rev. Lett.* **100**, 118703.
- Lindsey, B, Morris, K, Shannon, R, and Gerstein, G (1997). "Repeated patterns of distributed synchrony in neuronal assemblies." *J. Neurophysiol.* **78**, 1714–1719.
- Mao, BQ, Hamzei-Sichani, F, Aronov, D, Froemke, RC, and Yuste, R (2001). "Dynamics of spontaneous activity in neocortical slices." *Neuron* **32**, 883–898.
- Miller, R (1996). "Neural assemblies and laminar interactions in the cerebral cortex." *Biol. Cybern.* **75**, 253–261.
- Morelli, LG, Abramson, G, and Kuperman, MN (2004). "Finding and evaluating community structure in networks." *Eur. Phys. J. B* **38**, 495–500.
- Prut, Y, Vaadia, E, Bergman, H, Haalman, I, Slovin, H, and Abeles, M (1998). "Spatio-temporal structure of cortical activity, properties and behavioral relevance." *J. Neurophysiol.* **79**, 2857–2874.
- Reijneveld, J, Ponten, S, Berendse, H, and Stam, CJ (2007). "The application of graph theoretical analysis to complex networks in the brain." *Clin. Neurophysiol.* **118**, 2317–2331.
- Song, S, Miller, KD, and Abbott, LF (2000). "Competitive hebbian learning through spike-timing-dependent synaptic plasticity." *Nat. Neurosci.* **3**, 919–926.
- Song, S, Sjöström, PS, Reigl, M, Nelson, S, and Chklovskii, DB (2005). "Highly nonrandom features of synaptic connectivity in local cortical circuits." *PLoS Biol.* **3**, 507–519.
- See supplementary material at <http://dx.doi.org/10.2976/1.3386761>.
Supplementary information touches on the following topics: (1) details of the neuron model used; (2) the performance of WS networks as a function of the rewiring parameter; (3) the effect of stimulus delocalization on WS and modular networks over the whole range of degradation of the recall cue; (4) details on how Fig. 6 was obtained; (5) the performance of ER random networks with increasing connectivity; and (6) the effect of STDP on network topology.
- Swadlow, HA (1985). "Physiological properties of individual cerebral axons studied *in vivo* for as long as one year." *J. Neurophysiol.* **54**, 1346–1362.
- Swadlow, HA (1994). "Efferent neurons and suspected interneurons in motor cortex of the awake rabbit: axonal properties, sensory receptive fields, and subthreshold synaptic inputs." *J. Neurophysiology* **71**, 437–453.
- van den Heuvel, M, Stam, CJ, Kahn, RS, and Hulshoff, HE (2009). "Efficiency of functional brain networks and intellectual performance." *J. Neurosci.* **29**, 7619–7624.
- Villa, A, Tetko, I, Hyland, B, and Najem, A (1999). "Spatiotemporal activity patterns of rat cortical neurons predict responses in a conditioned task." *Proc. Natl. Acad. Sci. U.S.A.* **96**, 1106–1111.
- Volman, V, Baruchi, I, and Ben-Jacob, E (2005). "Manifestation of function-follow-form in cultured neuronal networks." *Phys. Biol.* **2**, 98–110.
- Watts, D, and Strogatz, S (1998). "Collective dynamics of 'small-world' networks." *Nature (London)* **393**, 440–442.