We are IntechOpen,
the world's leading publisher of
Open Access books
Built by scientists, for scientists

6,900
Open access books available

186,000
International authors and editors

200M
Downloads

154
Countries delivered to

Our authors are among the

TOP 1%
most cited scientists

12.2%
Contributors from top 500 universities

BOOK CITATION INDEX
CLARIVATE ANALYTICS
INDEXED

WEB OF SCIENCE™

Selection of our books indexed in the Book Citation Index
in Web of Science™ Core Collection (BKCI)

Interested in publishing with us?
Contact book.department@intechopen.com

Numbers displayed above are based on latest data collected.
For more information visit www.intechopen.com

Attractor Hypothesis of Associative Cortex: Insights from a Biophysically Detailed Network Model

Mikael Lundqvist, Pawel Herman and
Anders Lansner

Additional information is available at the end of the chapter

1. Introduction

Ever since Hebb proposed that cells that fire together wire together, the idea that memories are formed by distributed cell assemblies capable of self-sustained activity [1] has been one of the main hypothesis regarding memory formation and recall. It has laid the foundation for a theory of attractor memory extensively exploited in computational neuroscience. Memory representations, manifested as the selective activations of these cell assemblies, serve as attractors in simulated neural networks and can be retrieved as a result of external stimulation or intrinsic system dynamics.

Despite major efforts in neuroscience to investigate the attractor hypothesis experimentally, which have produced some supporting evidence, no conclusive result to prove or reject it has been provided. This current status can largely be attributed to the limitations in data collection and the distributed nature of Hebbian cell assemblies. For the attractor hypothesis of associative cortex to be validated, simultaneous spiking data from a vast number of cells over a large spatial scale should be recorded. In slices [2] and cell cultures [3], more accessible for such recordings, evidence for cell assemblies capable of self-sustained activity has been provided. In vivo however the task is more challenging since the use of intrusive techniques is limited. In addition, activity related to attractor dynamics can be obscured by spiking contributions reflecting other, parallel processes in behaving animals. In consequence, we must at this point rely on indirect evidence. Simulations in biophysically detailed attractor networks can provide useful insights in this regard and help to address questions relevant to a hypothesis of attractor computations in cortical circuits, for example:

- Is the known cortical connectivity with relatively sparse cell-to-cell connectivity sufficient to support the globally coherent phase-transitions and sustained activity states associated with attractor networks?

- Likewise, is the observed sparse and low rate cortical activity consistent with the activation of recurrently connected cell assemblies?

- What features of the neural activity in vivo could be linked to and interpreted in light of the simulated attractor dynamics?

- In what capacity can additional phenomena, such as oscillations in various frequency bands and cross-frequency coupling effects, be explained by the presence of attractors in the biological system?

In addition, models compliant with known biological data can then make several testable predictions and guide further experimental work. The last few decades, attractor networks have been used extensively as models for cortical memory in various paradigms [4-15]. The major distinguishing feature of the model presented here is that it operates in an oscillatory regime and has a modular structure [16-20]. Throughout this chapter we demonstrate evidence for biological relevance of these features and motivate functional advantages of oscillations in our attractor network.

2. Basic hypothesis

The ad hoc hypothesis adopted here is that layers 2/3 of associative cortex provide the neural substrate for attractor memory network. In the light of attractor hypothesis, cortical memory representations correspond to attractor states supported by recurrent excitatory connections. Attractor networks have several dynamical attractors, to which similar activity patterns in terms of a combination of specific active and inactive units are attracted. These attractors can be stored by means of synaptic learning. The attractor dynamics lends the memory system several attractive features. First of all, such memory networks are noise resistant and fault tolerant in the sense that a noisy, corrupted or incomplete stimulus can still activate a full corresponding memory pattern – the effect known as pattern completion. Furthermore, when conflicting stimuli are provided the phenomenon of pattern rivalry occurs. In addition, the use of local, synaptic learning rules are sufficient to form global memory patterns using highly parallel processing. Despite this locality, an attractor network trained with a Bayesian-Hebbian learning rule [21] retrieves the pattern provided with the stronger evidence based on the statistics of the input and previous learning examples. In addition, storage capacity in large-scale attractor networks appears to meet biological needs [22].

Despite a high degree of compatibility between the functionality of attractor networks and that of cortical memory, it is relevant to study the actual anatomical substrate of attractor dynamics in cortex. As mentioned in the beginning, we hypothetically designate layer 2/3 to be the main driver of such dynamics; mostly due to the predominant presence of dense recurrent connections, necessary to support attractor function. From a neurodynamical perspective, these layers

seem to be the main source of excitatory drive in the cortical circuitry [23, 24]. In addition, the phylogenetically oldest parts of the cerebral cortex only contain the superficial layers so if the attractor functionality is central to cortical processing, it should be harbored there. The deeper layers, which emerged later in evolution, could still be directly involved in or supporting attractor function. Then however they would be likely to rather address the needs arising from the expanding cortex size such as readout and output to subcortical structures [25, 26], or participate in the selection and modulation of both task-relevant and task-irrelevant cortical modalities. The latter notion is bolstered by the fact that layer 5 seems critical in the regulation of cortical up and down states [27, 28], i.e. in the regulation of global excitability of entire cortical areas.

In further support of attractor dynamics in the superficial layers, stimulus evoked neural activity exhibited in layers 2/3 is also sparser with lower average firing rate and is more selective to input statistics compared to the deeper layers [29, 30]. These characteristics are congruent with sparse and distributed memory patterns stored in attractor networks. However, a consequence of this relatively sparse activity is that it is likely to be obscured by deep layer activity when large quantities of spiking data is collected, which hinders the acquisition of direct neural evidence for attractor-like dynamics. It is not surprising therefore that the most direct in vivo evidence of attractor dynamics comes from olfactory cortex ([31, 32] and references found therein), and hippocampus [12, 33, 34], i.e. cortical structures that lack the deeper layers. There is also evidence for self-sustained and input specific activity from inferotemporal [35, 36] and prefrontal cortex [37-39], which are late in the processing stream and therefore should be more strongly influenced by the intrinsic connectivity. In addition, two-photon calcium imaging studies have produced relevant insights into the attractor hypothesis since the imaging method can reveal calcium current traces with good temporal resolution in tens to hundreds of neurons simultaneously within a small cortical volume of the superficial layers in vivo [40-43]. This technique was recently used to demonstrate non-linear attractor-like activity in auditory cortex [42]. In particular, spatially organized neuronal sub-groups were shown to respond discretely in time to specific auditory cortex input [42]. Here, groups of stimuli evoked all-or-nothing responses in distinct neural sub-groups. These discrete activities were however partly obscured by a large trial-to-trial variability.

Finally, there is evidence for attractor dynamics sustained by the recurrent connectivity in striate cortex [44, 45]. Using voltage-sensitive dye imaging, Kenet et al. [45] found that the superficial layers switched spontaneously and in a coordinated fashion between re-occurring states spanning several cortical columns. These spontaneous states showed strong correlation to visually evoked patterns of activity and have later also been reported to match the structured, horizontal long-range connections in layer 2/3 [46]. It thus seem likely that visually evoked states are strongly related to self-sustained attractor states supported by recurrent connectivity in superficial layers.

However, it is not clear whether such switches between stable activity patterns are indeed compatible with the dynamics of computational networks as for such models, unlike biology, full connectivity between units is often used. Further, single units in attractor networks display very high firing rates with low variability while superficial activity in vivo has low rate and is

highly variable. From the modeling perspective, the implications of the questionable assumption about all-to-all cortical connectivity adopted in theoretical studies (mathematically it ensures convergence to stable states) have hardly been investigated in the context of biological plausibility of attractor dynamics and function. Nor have the very low firing rates reported in vivo been reproduced. Our approach relying on a biophysically detailed attractor network model of cortex with a spatial scale spanning several hypercolumns [16], which draws from known anatomy and connectivity, allows for addressing some of these questions.

3. The network model

The network contains two types of neurons, excitatory pyramidal and inhibitory basket cells, composed of several compartments modeled by Hodgkin-Huxley equations. The basic functional units of the network are however minicolumns, each containing 30 recurrently connected pyramidal cells (Figure 1), inspired by the columnar structure of sensory cortex [41, 47-49]. These should not necessarily be seen as anatomical columns but rather functional columns consisting of subgroups of more tightly connected neurons, as found throughout cortex [40, 42, 43, 50-55].
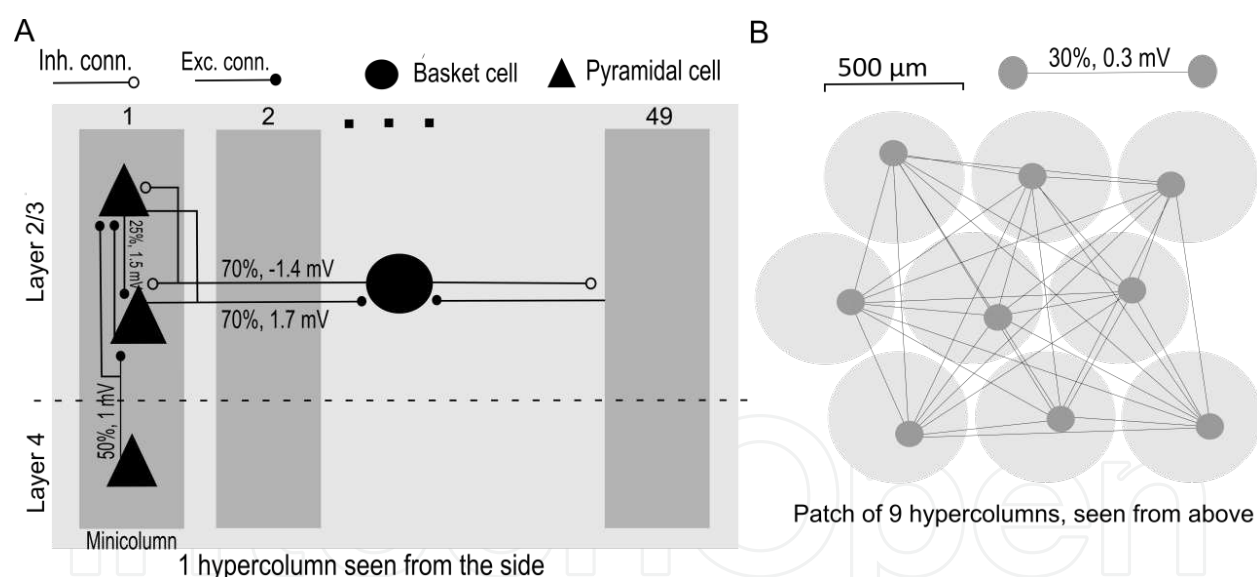

Figure 1. Network setup and connectivity. A: A detailed connectivity of a single hypercolumn, containing 49 minicolumns. B: A sketch of the long-range connectivity within a cortical patch, consisting of several hypercolumns (9 in a full patch). The numbers on the arrows give the connectivity and post synaptic potential (PSP) size at resting potential of the post-synaptic cell.

A cluster of minicolumns, spanning a few hundred microns, constitutes a hypercolumn in the network. Since the minicolumns within each cluster are coupled through a pool of basket cells, a hypercolumn can be defined by the extent of non-specific feedback inhibition [52] (Figure 1). In earlier studies [16] we used down-scaled hypercolumns containing 8 minicolumns, but in the subsequent work hypercolumns contained at least 49 minicolumns [17-20]. The feedback

from the basket cells has several functions. It normalizes activity in the network, provides the means for mutual competition that implements winner-take-all (WTA) dynamics within a hypercolumn and finally produce oscillations, which in turn add several interesting dynamical features to the network. Similar local WTA dynamics, on the scale of ~200 microns, was recently observed in auditory cortex in vivo [42].

We have typically modeled a cortical patch of about 1.5x1.5 mm using a 9-hypercolumn network. Distributed, retrievable and sparse patterns of activity are stored as attractors in this network. This is achieved by long-range interactions between pyramidal cells in minicolumns across different hypercolumns (Figure 1). Such structured, horizontal connections had originally been adopted in the model as an assumption but later on they received increased experimental support from studies of layer 2/3 connectivity [46, 56, 57]. In the work presented here, only orthogonal attractor patterns are stored, i.e. each minicolumn only participates in one global pattern. Although overlapping patterns, where each minicolumn participates in several patterns, increase memory storage capacity, they lead to similar results [58]. Data from in vivo paired recordings are used to bring connectivity and synaptic weights as close to biology as possible [59], but assumptions regarding long-rage connectivity have to be made.

4. Attractor properties, low firing rates and nested oscillations

We have found that stable attractor activity can indeed be maintained for plausible synaptic weights and very low firing rates, if the network is operating in an oscillatory regime (Figure 2). These oscillations, in the range of 25-40 Hz, correspond to upper beta [60] and gamma-like [61] oscillations in vivo, which have been correlated with active stimulus processing and memory recall [60-68]. In our network, the oscillations are generated by the strong feedback inhibition from basket cells (pyramidal interneuron gamma (PING) network; [69, 70]). This feedback inhibition also effectively underlies the selection of a winning population in the WTA circuit within a hypercolumn and controls firing rates in this winning cell assembly.

The oscillatory regime is also interesting for other computational reasons. Due to the gamma-cycle dynamics, an attractor cell assembly could maintain its activity and suppress the activity of competing assemblies already at an average firing rate of 3 s^{-1} per pyramidal cell [17]. This can be explained by the dynamics of the gamma cycle, which has a phase dominated by excitation where pyramidal cells have an opportunity to fire, and followed by a phase where the innervated basket cells shut down the activity in the network. As this inhibitions wears off, there is a race between populations of pyramidal cells to reach the firing threshold before recruited basket cells shut down the activity again [65]. As a result, only a small bias (low firing rates) to one of the competing populations is needed to activate or maintain a given attractor.

Since the network is highly dependent on the activity in the distant recurrently connected hypercolumns an intrinsic bias is mediated by long-range excitation, which arrives out of phase with respect to local excitatory inputs (Figure 2A), often in the inhibition-dominated part of the local gamma rhythm. This reflects an integration of global evidence for a given memory pattern on the gamma time-scale and implies that the resulting decision to either maintain

current activity pattern in the network or not is made within a short temporal window in each gamma cycle, serving as a discrete time unit. Consequently, transitions from and to active attractor states are globally sharp. The inter-hypercolumnar connections underlying these computations in a modular network help to stabilize the oscillatory regime since the global excitation arrives out of phase with respect to the local firing [17].
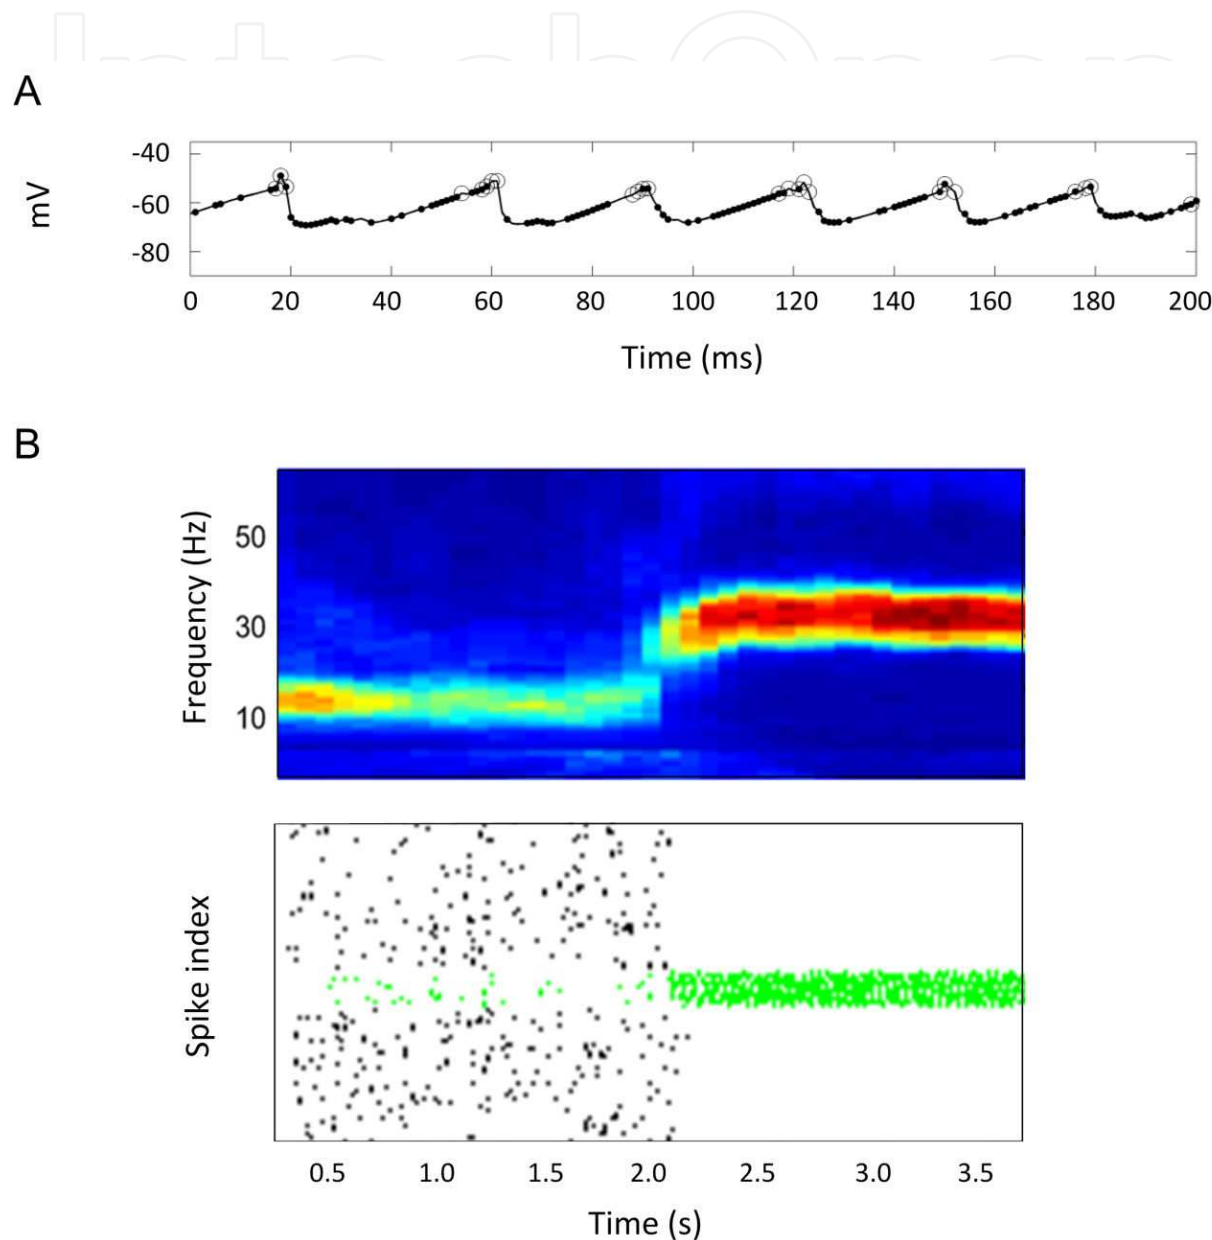

Figure 2. Oscillatory activity in the various network states. A: During attractor retrieval each minicolumn in the active assembly oscillates at gamma frequency (25-40 Hz). All pyramidal spiking within a minicolumn is concentrated to the peak of each oscillation (circles) while the incoming spikes from distant minicolumns are evenly distributed across the whole oscillatory cycle, stabilizing activity within the assembly. B: Bistable network receiving stimulation of one of its coding attractors at t = 2s. This time point marks a transition from alpha like (ground state) to gamma like (attractor state) oscillations (top) and a simultaneous transition from diffuse low rate firing to the concentrated higher rate spiking (bottom) in a specific cell assembly. Spiking from pyramidal cells in this assembly is shown as green dots while all other spikes are depicted as black dots.

The fast state transitions in the network can also be understood from the perspective of balanced excitation and inhibition [11, 71, 72]. Since spiking of individual neurons in balanced networks is driven by input fluctuations rather than average net excitation, with membrane potential close to the firing threshold, rapid state transitions can occur [71, 72]. The balance, which also results in highly irregular firing on the single cell level [72], is roughly preserved in a large parametric region of the oscillatory regime [17, 69]. The model can therefore operate in this regime without a need for fine-tuning or plasticity-induced synaptic changes, which are otherwise necessary in memory networks [9, 73]. The oscillatory regime thus results in fast transitions and irregular firing [17, 69] with a CV_2 close to one (during attractor activation), as often reported in vivo during delay match-to-sample tasks [39, 74, 75].

Classical attractor networks remain in an attractor state once they fall into one as long as there is no external input forcing a transition. One of the biological mechanisms that can cause such global transitions out of active attractor states is neural fatigue, implemented in our modeling work by the inclusion of cellular adaptation and synaptic depression in the model [76]. Together they render the attractor lifetime finite and the level of adaptation has a direct effect on the attractor duration (Figure 3A). The dynamics of activation and deactivation of attractors with finite life-times result in an increase in theta/delta-band power of the synthesized local field potentials (LFPs) (Figure 4A). The peak frequency of this rhythm corresponds roughly to the inverse of the attractor's dwell time. In consequence, the co-emerging gamma and theta/delta rhythms are coupled, i.e. the phase of slower theta/delta wave modulates the amplitude of faster gamma activity (Figure 4B). Such nested oscillations have been widely reported as a neural correlate of various memory paradigms [62, 77-81]. Theta oscillations by themselves have also been connected to both encoding, learning and retrieval of memory objects [68, 82-87]. In addition, theta phase modulations of firing rates observed in vivo [85] can also be found in the model (Figure 4C).

From the functional perspective, the network is capable of memory completion and pattern rivalry (Figure 5). Memory completion was tested by providing the network with partial stimuli of the stored patterns and examining whether full activation of the stored activity pattern was achieved via the lateral long-range connections. This occurs when roughly one third of the minicolumns in a pattern receives brief stimulation (Figure 5A, B). Pattern rivalry reflects the network's ability to resolve ambiguities in the input. When two patterns are simultaneously stimulated and their relative strengths vary, it turns out that small differences between stimuli can have a decisive impact on which pattern is activated and which one is extinguished [16]. Lundqvist et al. [16] demonstrated that relative differences in input strength of 25% consistently selected the more strongly stimulated assembly. This is by no means the lower limit though and here we used 10% differences (Figure 5B). Once the activity of the winning pattern is terminated due to adaptation and the same conflicting stimuli is applied again, the weakly stimulated pattern typically gets activated (Figure 5B).
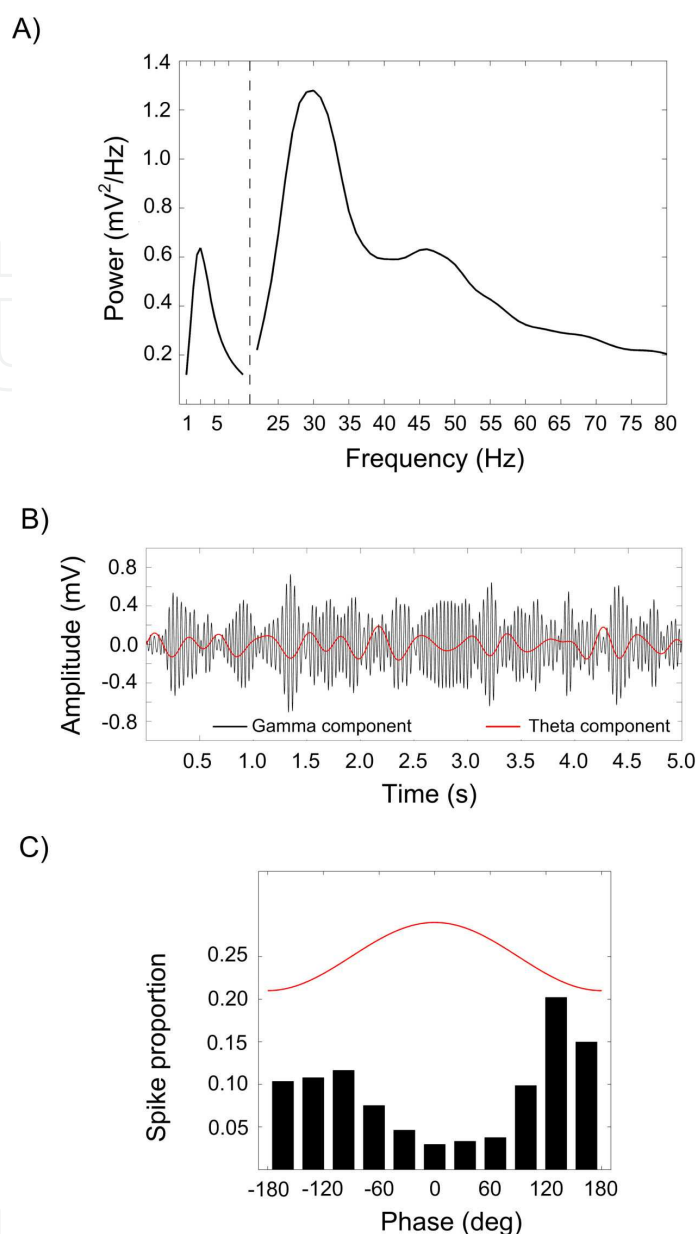
Figure 3. Cartoon of energy landscapes in various regimes. Solid lines depict the energy of various states at time t=0, and broken lines at a later time point (specified in A-C). The ball indicates what state the network is likely to be in, i.e. one of the states with the lowest energy. A: Network with attractors of limited life-time, t_{lt} (200-800 ms depending on parameters). The network will quickly end up in one of its attractor states at the onset of each simulation. At t = t_{lt} (broken lines) neural fatigue and synaptic depression has increased the energy of this attractor such that noise will bump it to another attractor state. If there is no neural fatigue the attractor states will be persistently active until the network is deterred into a new state by external stimulation. B: Bistable network with one default state and several coding attractors. This bistability is achieved by either scaling up the network or increasing mutal inhibition between cell assemblies. At the onset of simulation we have here stimulated a specific coding attractor. At t=t_{lt} (broken line) the network will again exit this state but now jump into the ground state. The network will remain in this state until one of the coding attractors are stimulated. C: Bistable network with added synaptic augmentation. Solid lines show the network state just after the stimulated attractor has terminated due to neural fatigue, and the network has retreated to its ground state. After some time t, larger than the fast decay of neural fatigue but smaller than the decay of the more long-lasting synaptic augmentation, the energy landscape is altered (broken lines). During this time window the network is likely to jump back in to the previously active attractor spontaneously.
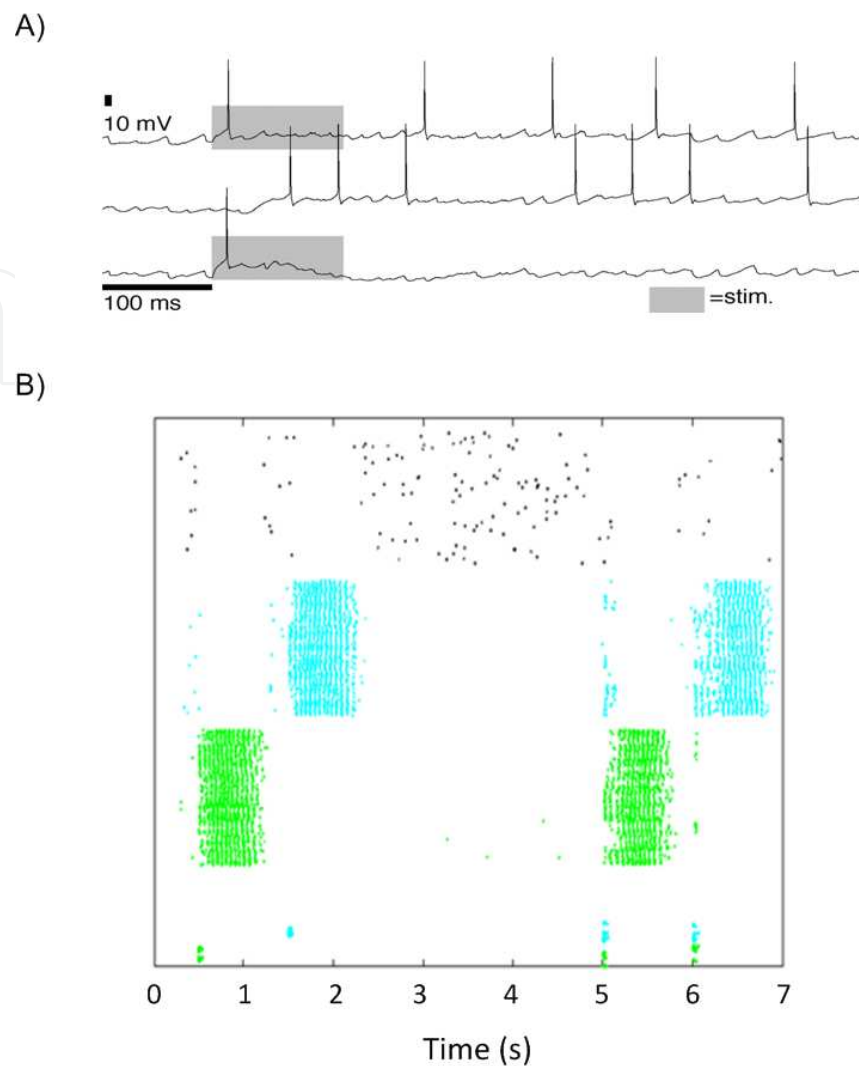
Figure 4. Gamma and theta phase locking. Gamma and theta power (A) and gamma-filtered (black) and theta-filtered (red) components during active sequential retrieval of attractors (B). C: Histogram of spike events in relation to the theta phase.

Conceptually, we consider minicolumns, rather than single cells, as the basic functional units of the network. This means that information processing does not rely on single cells but on recurrently connected neuronal populations. This perspective recently obtained additional experimental support [42] and has several important implications. Firstly, the connectivity within the network on the unit level can be increased without affecting the biologically realistic connectivity on the single cell level [22]. Since a pyramidal cell receives roughly 10 000 synapses, full cell-to-cell connectivity is not possible, even within a small cortical volume. With minicolumns acting as computational units, a closer approximation of the full connectivity, assumed in theoretical studies of attractor networks, can be obtained (another factor that reduces the need for full cell-to-cell connectivity is the dense local inhibition implementing disynaptic connections between a vast number of pyramidal cells). Secondly, since the average output of each minicolumn rather than that of a single cell reflects the activation of a distributed

memory pattern, memory retrieval is more robust to cellular variability or synaptic failures. In this light, irregular and rare firing of individual pyramidal cells does not undermine the stability of the retrieval process. On the contrary, this irregular firing is instead the manifestation of a dynamically regulated network state where the population activity does not depend on spike timing and firing rates of individual cells. In consequence, the network function is robust to cell death and synaptic loss [88]. Without adjusting the synaptic weights, more than 50% of cells could be removed with no detrimental effects to the attractor retrieval dynamics (Figure 6). As regards the removal of synapses, it can be performed in two different ways. First, if connections are removed from one cell at a time, a similar effect can be obtained by simply removing cells. Second, in the scenario where individual connections are removed at random the network becomes slightly more sensitive, but still tolerates a synaptic loss of roughly 40%. This number can be increased to 60% if the loss is compensated by increasing the conductance of the remaining synapses [88].
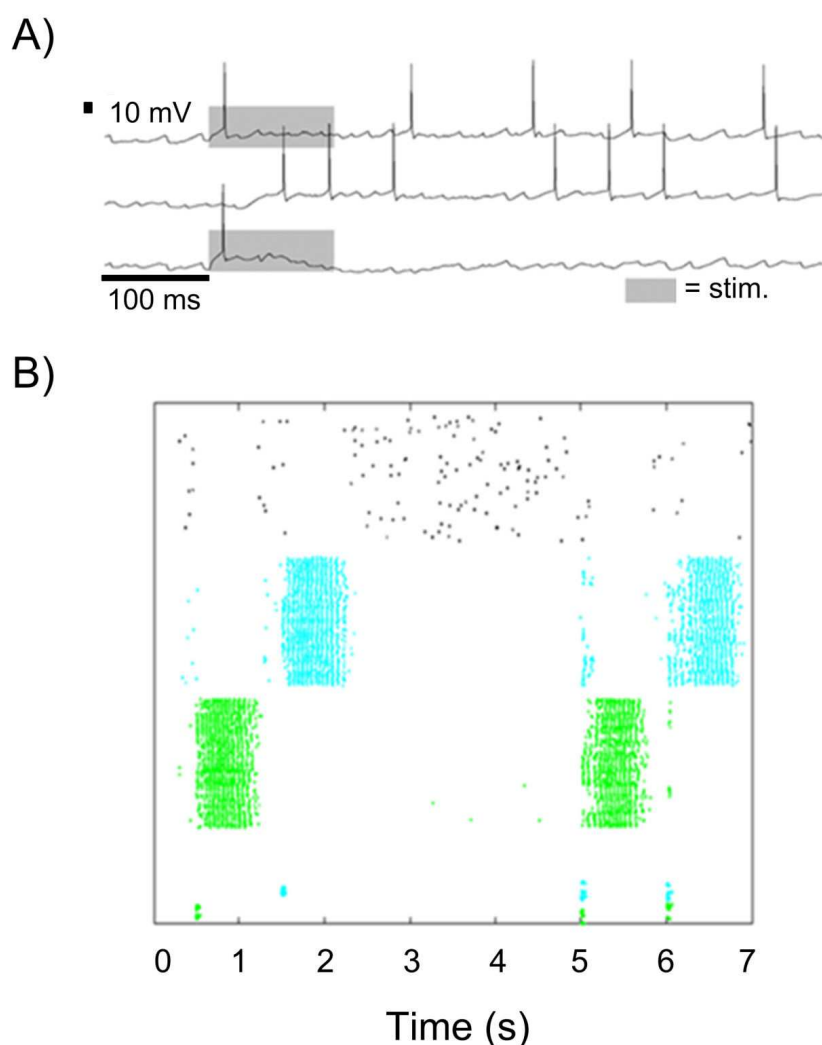

Figure 5. Pattern rivalry and completion. A: Single cell dynamics during pattern rivalry and completion. Two cells (top and bottom respectively), part of two distinct assemblies, receive input. The cell at the top is part of an assembly that

receives slightly stronger excitation and will sustain its activity while the cell at the bottom will be suppressed (pattern rivalry). The cell in the middle didn´t receive any stimulation but belonged to the winning assembly and becomes active (pattern completion). B: Global spiking dynamics demonstrating completion and rivalry. Two assemblies receive brief, partial input to 1/3 of their pyramidal cells at t=0.5 s (green) and t= 1.5 s (cyan). These inputs quickly spreads so that the full patterns are activated (pattern completion). At t=5 s both patterns receive stimulation simultaneously, but the green pattern receives 10% stronger input. This pattern quickly activates at the cost of the cyan pattern. At t=6 s the green pattern again receives 10% stronger input but due to the recent activation it is partly fatigued and the cyan pattern prevails.
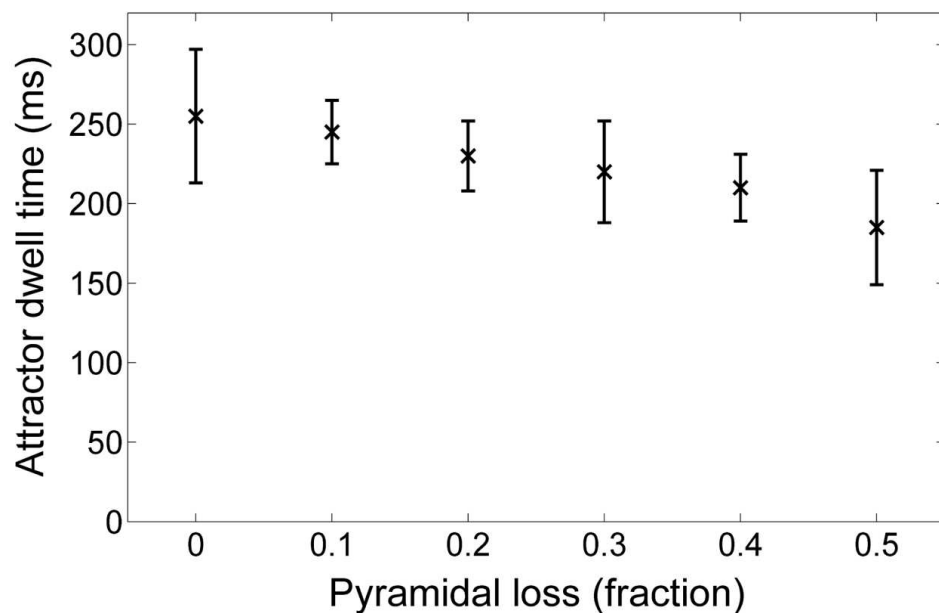

Figure 6. Tolerance to cellular death. The stability of attractor dynamics, measured as the dwell time (y-axis) of stimulated cell assemblies receiving brief stimulation. Cells are removed at random from the network (x-axis) without adjusting connectivity or synaptic weights.

5. Scaling the network and the emergence of bistability and alpha oscillations

Since the scale of the original model was small relative to a cortical area in terms of the number of hypercolumns and minicolumns (while the number of cells within each minicolumn was consistent with biological evidence), it becomes relevant to investigate whether biologically plausible neural dynamics and attractor function can be maintained at much larger simulated scales. For instance, the question as to whether a large distribution of axonal delays can co-exist with stable and coherent activations of cell assemblies should be addressed. In addition, it is important to show that the relatively few connections that each pyramidal cell can form are sufficient for stable memory retrieval even at cortical scales. In order to handle these questions, we scaled the network considerably, up to the size of mouse cortex containing 22 million neurons and spanning 16 cm^2 [20].

Due to the modular structure of the network, and arguably cortex, it is indeed scalable with largely preserved dynamics. Once hypercolumns are scaled to realistic size, only the density in the connections across them has to be re-scaled in order to maintain the dynamical regime as the network grows. As the number of hypercolumns in the network was increased, we kept the number of long-range (cross-hypercolumnar) connections terminating on each pyramidal cell constant, progressively diluting the probability that two distant neurons connect. Since biological neurons have limited physical space to make connections on their dendrites, an equivalent process seems likely as in vivo systems are scaled up. As a result, a single cell sees roughly the same amount of excitatory and inhibitory input once an attractor state is entered regardless of network size. Dynamics and function during attractor retrieval were maintained even for the largest simulations without any parameter changes [20]. The transitions to and from attractor states turned out surprisingly coherent, even though the slowest time delays within each assembly were 50-60 ms. This effect was again, as described above, obtained due to the interdependence of minicolumns in each pattern mediated by the gamma cycle dynamics and the network operating in a balanced regime, where only small changes in excitation are needed for state transitions to occur.

Despite largely preserved attractor retrieval dynamics there are functional and dynamical consequences of scaling up the network. Most importantly, another dynamical state of the network emerges [11, 20] in addition to the aforementioned active attractor coding state (Figure 2B) once each hypercolumn has more than 25 minicolumns. Since this new state becomes the default condition of the bistable (Figure 3B) network in the absence of any external stimulation, it is referred to as the ground state. It is in our network manifested by global alpha-band (~10-20 Hz) oscillations (Figure 2B) and is characterized by very low levels of activity in all minicolumns without a dominance of any patterns. This state is facilitated by the mutual competition between attractor patterns [11], stabilized by feed-back inhibition growing with the network size. In the smaller network, noise fluctuations quickly activated one pattern at the expense of the others leading to a sequential recall of the patterns in a random order. In the larger network, on the other hand, it is possible to maintain the state of competition between attractor patterns as long as there is no sufficient bias to one of them, thus the emergence of a new stable state. This bias could be either in the form of external stimulation of a specific pattern or internal mechanisms such as synaptic facilitation, which we used to store a subset of patterns in working memory ([18, 19]; see section *Multi-item working memory*).

In the scaled-up bistable network, successful pattern activation by an external cue is coupled to a transition in the oscillatory dynamics from the alpha to gamma rhythm (Figure 2B). Similar stimulus induced transitions have been reported in layer 2/3 of the visual cortex in vivo [66, 89]. In the context of extensive experimental work on neural oscillations, our two distinct network states correspond with a general view that alpha reflects idling or pre-stimulus readiness (for a review see [90, 91]) and gamma is a correlate of active processing ([61]; for review see [64, 65]).

What are the mechanisms underlying these rhythms, and, more importantly, the transition between them in our network? In balanced networks with oscillatory population activity and irregular firing, the oscillatory frequency is dependent, among other factors, on the level of

overall excitation in the network [92]. Comparing spiking populations in the two stable network states, the excitation level and firing rates are higher during active memory retrieval, thereby increasing the oscillatory frequency relative to the ground state. At the limit where the recurrent excitation within cell assemblies is just strong enough to promote stable attractor states, the switch from the ground state to one of the coding attractors is associated with a minimal increase in excitation and oscillatory frequency. Although the total amount of spikes elicited from the pyramidal cell population as a whole remains the same, after stimulation all spikes are elicited from the active cell assembly, i.e. the combination of single minicolumns across all the hypercolumns, instead of being spread out between all pyramidal cells as in the ground state [17]. This effect occurs since cells in the active assembly climb slightly faster to firing threshold in each oscillatory cycle, thereby shutting down competitors before they get a chance to spike and influence the network dynamics. It illustrates how a very small bias can have a strong impact on the spiking in oscillatory, balanced networks. As recurrent excitation is increased, the gap between the oscillatory frequencies in the two states also widens towards a clear distinction between alpha and gamma rhythms, hence reflecting the gradual stabiliza‐ tion of the active state. To maintain attractor-coding activity, a cell assembly has to oscillate faster than the ground state frequency. Towards the end of the attractor's lifetime the oscilla‐ tory frequency drops due to adaptation and the network consequently falls back into the ground state.

As with the balanced regime, the bistable regime with two simultaneously stable states exists also in non-oscillatory networks [11]. However, the advantage of oscillatory networks amounts to the fact that the parametric range of the bistable regime becomes much wider and less sensitive to perturbations in excitation [17]. The strong feedback inhibition needed for a stable ground state does not destabilize the active attractor states. On the contrary, it has relevant functional and dynamical implications for the network during memory retrieval, as discussed in the previous section.

In general, neural oscillations as a population phenomenon occur due to strong feedback inhibition that periodically shuts down activity in a network, and therefore typically desta‐ bilizes persistent activity in a cell assembly [93]. However, if this cell subset is biased in any way, in our case by the long-range excitation out of phase with respect to the local oscillations, the persistent activity in the oscillatory regime becomes extremely stable instead. Once the network can tolerate periodic hyperpolarization without terminating the activity permanently, strong feedback inhibition can be used to dynamically balance fast changes in excitation. Then, as long as the inhibition is strong enough to periodically shut down the network, it remains roughly balanced.

6. Multi-item working memory

Attractor networks have been proposed as a modeling framework for a working memory system, which temporarily maintains a small subset of memory items. Models of spatial working memory have for instance used persistent activity in bump-attractor networks to

preserve a trace of a specific direction [8, 93]. We can obtain a similar effect in our network when the adaptation mechanisms are subdued. Then a stimulated attractor will remain persistently active as a cued memory over several seconds [17]. The persistent activity approach is limited however since only one item or direction can be stored at any given time due to the mutual inhibition between attractors, whereas working memory is reported to contain up to seven items simultaneously [94, 95]. In this section we discuss an alternative approach to working memory maintenance known as periodic replay [10, 14, 18, 19, 84, 96], which allows for storing multiple items.
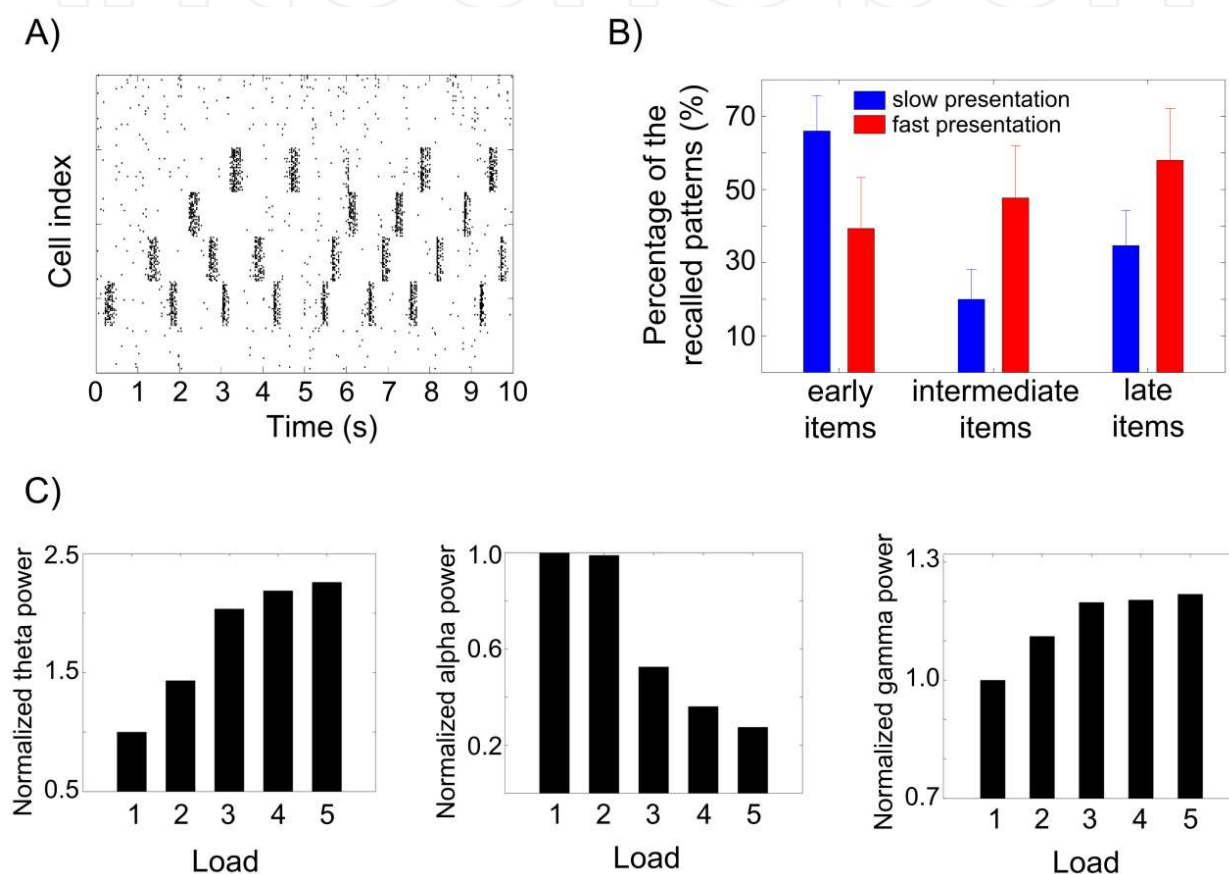
Figure 7. Multi-item working memory through synaptic augmentation. Synaptic augmentation causes attractors to spontaneously reactive some time after they have been terminated. This augmentation is then refreshed upon each reactivation, and the attractor is held in working memory in a cyclic fashion. A: Here four items are stimulated (at 0.2, 1.2, 2.2, and 3.2 sec), and the space between active recall is filled by ground state activity. The items presented early start their re-activations already during the presentation period (0-3.5 s). This can explain the bias for items presented early in the list to be remembered as seen in (B). Here 10 items are presented followed by a recall phase where we test which items that are replayed. Early (1, 2) and late (9, 10) items have a higher probability to be remembered than intermediate (4-6) items (blue bars). If the list is presented at the rate of 2 s⁻¹ (red bars), the tendency for early items to be remembered is removed. C: Frequency modulations by memory load. Bars show integrated power in the three different power bands (2–6, 10–18, and 28–40 Hz) and five different load conditions. Bars are normalized relative to the power in Load 1 condition (one memory item), such that power in Load 1 is 1.

Although in both working memory models only one attractor can be active at any given time, in the periodic replay paradigm it has a brief lifetime instead of being persistently active. The

encoded items are then retrieved in a sequence one after another and get periodically reactivated. In computational networks this effect can be achieved by incorporating either cellular [96] or synaptic [10, 14, 18, 19] mechanisms that adjust the excitability of activated neurons dynamically. In the latter case, it can be achieved by adding synaptic augmentation, observed in prefrontal neuronal subgroups [97], on top of faster synaptic depression in a bistable attractor network. On the single synapse level, this makes the conductance vary dynamically over time. During a brief pre-synaptic spike train the amplitude of excitatory post-synaptic potentials (EPSPs) remains static or slightly decreases over time due to the combined effect of synaptic augmentation and synaptic depression. However, due to the slower decay of the augmentation, a new spike arriving roughly one second after the initial burst elicits a significantly magnified EPSP. On the cell assembly level, this implies that an attractor that has been activated by stimulation is temporarily more excitable than the ground state some time after its termination (Figure 3C). During this window it has a high chance to spontaneously reactivate and in the process refresh the synaptic augmentation. This way, a pattern stimulated initially becomes periodically reactivated. During silent periods there is an opportunity for other assemblies to be replayed (Figure 7a). Due to the decay of augmentation, the subset of memory patterns selected for replay need to be reactivated within the decay time window following their last deactivation in order to maintain their elevated excitability. As a consequence, a limited number of items can be stored. In particular, up to ~6 attractor memories can be simultaneously augmented and hence periodically reactivated [18, 19] for biologically realistic levels of synaptic augmentation.

The notion that individual memory objects are replayed at a theta time-scale during working memory maintenance has support from human MEG recordings [84]. The model can also explain the widely reported finding that alpha-band power decreases [98, 99] while gamma- and theta-band power increase [67, 98-100] with working memory load. We obtain this effect (Figure 7C) since for each additional memory item encoded in working memory, the network spends on average shorter time in the alpha-dominated ground state and longer time in its active retrieval state, correlated with nested theta-gamma oscillations [18]. The effect saturates at the full memory capacity of the network.

The notion of theta-coupled replay of memory items with accompanying theta-gamma phase-amplitude coupling is also consistent with single-cell spike statistics obtained from recordings in prefrontal areas and superficial layers of cortex, where a relative abundance of cells displaying clumpy-bursty behavior with Lv [101] and CV_2 [102] well above 1 was observed [103, 104]. This clumpy-bursty behavior can be reproduced when single cells burst in specific theta periods and are silent in the other ones as is seen in the periodic replay paradigm [19]. Although the estimated variability during the active theta periods results in Lv close to 1, the inclusion of long inter-spike intervals (ISIs) introduced by the silent theta periods boosts Lv to 1.5 (Figure 8), as reported for clumpy-bursty cells in vivo. This effect occurs for firing rates within a certain range, overlapping with the ones observed in our network model [19].

Finally, we would also like to present unpublished results from a study aimed at reproducing the phenomenon of recency and primacy effects [105] in list-learning paradigms. When a list of items exceeding the capacity of working memory is to be remembered by a subject, there is

a marked tendency for objects from the beginning (primacy) and the end (recency) of the list to be recalled with a greater likelihood. To simulate this, the network is presented with 10 memory items at the rate of 1 s^{-1} followed by a 10 s period corresponding to a free recall phase. On average in 100 trials, 5.0±0.7 (mean ± standard deviation) items are maintained such that they are replayed in the recall phase. In addition, memory items in the beginning and at the end of the list are more frequently encoded than those presented in the middle (Figure 7B). The simulated recency effect can be explained by the fact that augmentation in the assemblies activated towards the end of the presentation period is relatively high when the free recall period starts. The primacy effect, on the other hand, can be explained by the fact that the network has time in between presentations to replay these items already in the presentation phase, and thus re-enforce their increased excitability. If the network is largely denied this opportunity by presenting the list of items in quicker succession (at the rate of 2s^{-1}), around five items are again maintained in working memory (4.9±0.8), but the first items now have the smallest chance of being remembered (Figure 7B). At their cost, the last items instead have an even elevated chance of being replayed during the free recall period.
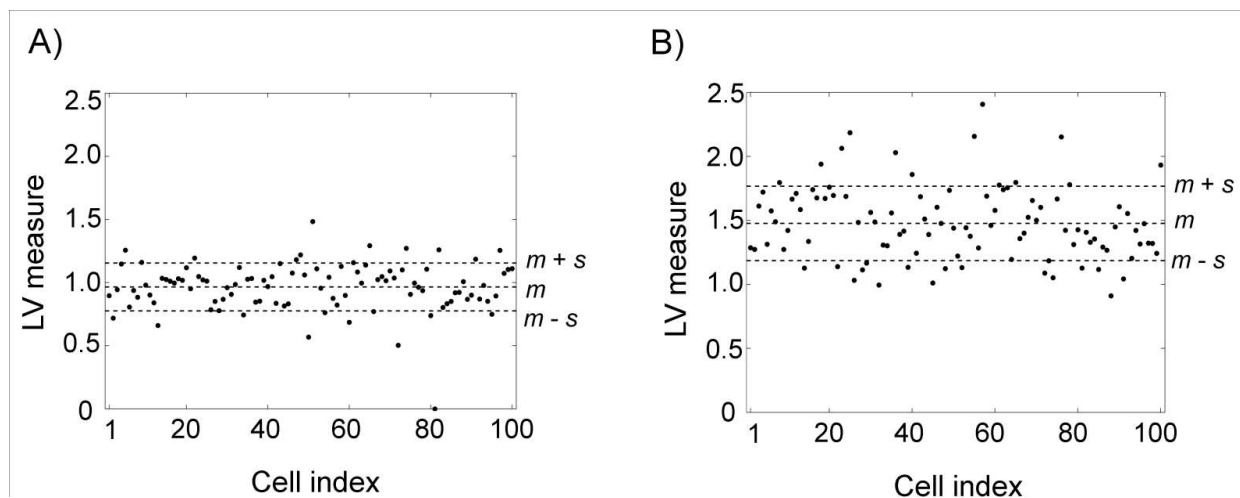

Figure 8. Scatter plot of Lv for 100 cells drawn from the persistently active network (A) and the replay network (B). The dotted lines mark the range of Lv values within one standard deviation from the mean.

7. Attentional blink

Attractor networks also allow us to study attentional mechanisms and their functional consequences. Attentional effects can be incorporated into such models in several different ways. For instance, it has been studied how top-down activity can bias certain attractors at the cost of others and thus serve as a model for top-down attention [13, 106]. Generally, in our work we rather focus on the potential neural manifestations of attention and examine how they correlate with the network's capability to retrieve weakly stimulated memory pattern. In that vein, we are currently studying the effects of both phase and power modulations of ongoing alpha oscillations on the network's performance. Here, however, we want to discuss

results related to the attentional blink phenomenon [107-110]. It is concerned with an inability to detect and process two relevant stimuli presented in quick succession by humans; the first item masks the perception of the second one even if they are presented equally long. This masking effect is not maximal when the visual targets are shown immediately one after another, but instead when the relative delay is around 300 ms [108, 109]. The attentional blink phenomenon was correlated with the P300 component [108] and evoked gamma oscillations in the electroencephalography signals [109].


Figure 9. Attentional blink. In the period of time closely after activation of one stored pattern in the network, triggering another pattern requires more stimulation than otherwise. Here, one pattern was first activated at t = 0. After some delay, a second stimulus was applied to a different pattern, attempting to trigger its activation. Two data series (rings and crosses) are shown, corresponding to separate experiments (random seeds). Each data point shows the minimum number of minicolums in a pattern that have to be stimulated in order to activate the second pattern after a given delay. Third degree polynomia have been fitted to the data points.

In the network, we obtain qualitatively similar time-dependent attentional blink effect [16] related to evoked gamma oscillations. This is due to the fact that an activated cell assembly attains a peak in firing rates after some delay relative to the stimulation. The delay corresponds to the time needed for the recurrent network to build up activity before the adaptation causes the reduction in rates again. Other competing assemblies are maximally suppressed and thus harder to activate at this peak of activity. We associate this effect with the impaired ability to detect the subsequent stimuli in consistence with psychophysical data ([108, 109], Figure 9). This phenomenon was recently studied in more detail using the same network model [58].

8. Summary and conclusions

We have reviewed evidence that the neural activity of superficial cortical layers is to a large extent compatible with the non-linear dynamics displayed in recurrent attractor networks. Research in the field of computational neuroscience has touched upon various aspects of the attractor theory with emphasis on its biological relevance and functional implications. In our modeling work, where we have used a biophysically detailed attractor network inspired by cortical connectivity, we have demonstrated how novel features such as modular structure and oscillatory dynamics render the model more robust and consistent with biological findings. In addition, we have shown how our mesoscopic network model can be utilized to link lower-level neural substrate with higher-order cognitive or behavioral phenomena. In particular, we have conceptually replicated recency, primacy and attentional blink effects. In the light of the network's dynamics we have also motivated the limited capacity of working memory. The model can be perceived as a crude model of the superficial layers of associative cortex taking the form of a large distributed network of attractor networks. Future work is intended to follow the direction of diverging individual cortical areas with respect to connectivity and function. We envisage that this work will be accelerated by the concerted effort in computational neuroscience to study cortical function from a bottom-up perspective.

Author details

Mikael Lundqvist[*], Pawel Herman and Anders Lansner

KTH and Stockholm University, Stockholm, Sweden

References

[1] Hebb DO (1949) The Organization of Behaviour. New York: Wiley.

[2] Cossart R, Aronov D, Yuste R (2003) Attractor dynamics of network UP states in the neocortex. Nature 423: 283-288.

[3] Plenz D, Thiagarajan (2007) The organizing principles of neural avalanches: cell assemblies in the cortex? TINS (30): 101-110.

[4] Hopfield JJ (1982) Neural networks and physical systems with emergent collective computational properties. PNAS 81:3088–3092.

[5] Palm G (1982) Neural assemblies. An alternative approach to artificial intelligence. Berlin: Springer.

[6] Lansner A, Fransén E (1992) Modeling Hebbian cell assemblies comprised of cortical neurons. Network: Comput Neural Systems 3: 105–119.

[7] Fransén E, Lansner A (1998) A model of cortical associative memory based on a horizontal network of connected columns. Network: computations in neural systems 9(2): 235-264.

[8] Compte A, Brunel N, Goldman-Rakic PS, Wang XJ (2000) Synaptic mechanisms and network dynamics underlaying spatial working memory in a cortical network model. Cereb Cortex 10: 910–23.

[9] Barbieri F, Brunel N (2008) Can attractor network models account for the statistics of firing during persistent activity in prefrontal cortex? Front Neurosci 2: 114-122.

[10] Sandberg A, Tegnér J, Lansner A (2003) A working memory model based on fast Hebbian learning. Network 14: 789-802.

[11] Amit DJ, Brunel N (1997) Model of global spontaneous activity and local structured activity during delay periods in cerebral cortex. Cereb Cortex 7 (3): 237-252.

[12] Treves A, Rolls ET (1994) Computational Analysis of the Role of the Hippocampus in Memory. Hippocampus 4: 374–391.

[13] Deco G, Thiele A (2009) Attention - oscillations and neuropharmacology. Eur J Neurosci 30: 347-354.

[14] Mongillo G, Barak O, Tsodyks M (2008) Synaptic theory of working memory. Science 319: 1543-1546.

[15] Lansner A (2009) Associative memory models: from the cell-assembly theory to biophysically detailed cortex simulations. Trends Neurosci 32: 178–186.

[16] Lundqvist M, Rehn M, Djurfeldt M, Lansner A (2006) Attractor dynamics in a modular network model of neocortex. Network 17: 253-276.

[17] Lundqvist M, Compte A, Lansner A (2010) Bistable, Irregular Firing and Population Oscillations in a Modular Attractor Memory Network. PLoS Comput Biol 6(6): e1000803. doi:10.1371/journal.pcbi.1000803.

[18] Lundqvist M, Herman P, Lansner A (2011) Theta and gamma power increases and alpha/beta power decreases with memory load in an attractor network model. J Cog Neurosci 23: 3008–3020.

[19] Lundqvist M, Herman P, Lansner A (2012) Variability of spike firing during theta-coupled replay of memories in a simulated attractor network. Brain Res 1434: 152–161.

[20] Djurfeldt M, Lundqvist M, Johansson C, Rehn M, Ekeberg Ö, Lansner A (2008) Brain-scale simulation of the neocortex on the IBM Blue Gene/L supercomputer. IBM Journal of Research and Development 52: 31-41.

[21] Sandberg A, Lansner A, Petersson KM, Ekeberg Ö (2002) A bayesian attractor network wth incremental learning. Network, 13: 179-184.

[22] Johansson C, Lansner A (2007) Imposing Biological Constraints onto an Abstract Neocortical Attractor Network Model. Neural Comput 19(7): 1871–1896.

[23] Weiler N, Wood L, Yu J, Solla SA, Shepherd GMG (2008) Top-down laminar organization of the excitatory network in motor cortex. Nat Neurosci. 11 (3): 360-366.

[24] Binzegger T, Douglas RJ, Martin KAC (2009) Topology and dynamics of the canonical circuit of cat V1. Neural Networks 22: 1071-1078.

[25] Wise SP, Jones EG (1977) Cells of origin and terminal distribution of descending projections of the rat somatic sensory cortex. J Comp Neurol 175: 129–157.

[26] Deschenes M, Bourassa J, Pinault D (1994) Corticothalamic projections from layer V cells in rat are collaterals of long-range corticofugal axons. Brain Res 664: 215–219.

[27] Shu Y, Hasenstaub A, McCormick DA (2003) Turning on and off recurrent balanced cortical activity. Nature 423: 288–293.

[28] Wester JC, Contreras D (2012) Columnar Interactions Determine Horizontal Propagation of Recurrent Network Activity in Neocortex. J Neurosci. 32 (16): 5454-5471.

[29] Sakata S, Harris KD (2009) Laminar Structure of Spontaneous and Sensory-Evoked Population Activity in auditory cortex. Neuron 64: 404-418.

[30] de Kock CPJ, Sakmann B (2009) Spiking in primary somatosensory cortex during natural whisking in awake head-restrained rats is cell-type specific. PNAS 106 (38): 16446-16450.

[31] Haberly LB, Bower, JM (1989) Olfactory Cortex: Model Circuit for Study of Associative Memory? Trends Neurosci. 12: 258–264.

[32] Wilson DA, Sullivan RM (2011) Cortical processing of odor objects. Neuron 72: 506-519.

[33] Wills TJ, Lever C, Cacucci F, Burgess N, O'Keefe J (2005) Attractor dynamics in the hippocampal representation of the local environment. Science 308: 873–876 .

[34] Niessing J, Friedrich RW (2010) Olfactory pattern classification by discrete neuronal network states. Nature 465, 47–52.

[35] Rolls ET, Tovee MJ (1994) Processing speed in the cerebral cortex and the neurophysiology of visual masking. Proceedings of the Royal Society, 257: 9-15.

[36] Desimone R (1996) Neural mechanisms for visual memory and their role in attention. PNAS 93(24): 13494-13499.

[37] Fuster JM, Alexander GE (1971) Neuron activity related to short-term memory. Science 173: 652–4.

[38] Goldman-Rakic PS (1995) Cellular basis of working memory. Neuron 14: 477-485.

[39] Compte A, Constaninidis C, Tegnér J, Raghavachari S, Chafee MV, Goldman-Rakic PS, Wang XJ (2003) Temporally irregular mnemonic persistent activity in prefrontal neurons of monkeys during a delayed response task. J Neurophysiol 90: 3441–3454.

[40] Kerr JND, de Kock CPJ, Greenberg DS, Bruno RM, Sakmann B, Helmchen F (2007) Spatial Organization of Neuronal Population Responses in Layer 2/3 of Rat Barrel Cortex. J Neuroscience 27(48): 13316-13328.

[41] Ohki K, Chung S, Ching YH, Kara P, Reid RC (2005) Functional imaging with cellular resolution reveals precise micro-architecture in visual cortex. Nature 433: 597-603.

[42] Bathellier B, Ushakova L, Rumpel S (2012) Discrete Neocortical Dynamics Predict Behavioral Categorization of Sounds. Neuron 76: 435-449.

[43] Rothschild, G, Nelken I, Mizrahi A (2010) Functional organization and population dynamics in the mouse primary auditory cortex. Nat Neurosci 13: 353–360.

[44] Arieli A, Sterkin A, Grinvald A, Aertsen A (1996) Dynamics of Ongoing Activity: Explanation of the Large Variability in Evoked Cortical Responses. Science 273: 1868-1871.

[45] Kenet T, Bibitchkov D, Tsodyks M, Grinvald A (2003) Spontaneously emerging cortical representations of visual attributes. Nature 425: 954-956.

[46] Muir DR, Da Costa NMA, Girardin CC, Naaman S, Omer DB, Ruesch E, Grinvald A, Douglas RJ (2011) Embedding of Cortical Representations by the Superficial Patch System. Cerebral Cortex, doi:10.1093/cercor/bhq290.

[47] Mountcastle V (1957) Modality and topographic properties of single neurons of cat's somatic sensory cortex. J Neurophysiol 20: 408-434.

[48] Hubel DH, Wiesel TN (1963) Shape and arrangement of columns in cat's striate cortex. J Physiol 165: 559-568

[49] Peters A, Yilmaz E (1993) Neurological organization in area 17 of cat visual cortex. Cereb Cortex 3: 49-68.

[50] Sato TR, Gray NW, Mainen ZF, Svoboda K (2007) The functional microarchitecture of the mouse barrel cortex. PloS Biol. 5, e189.

[51] Song S, Sjöström PJ, Reigl M, Nelson S, Chklovskii DB (2005) Highly Nonrandom Features of Synaptic Connectivity in Local Cortical Circuits. PLOS Biology 3 (39): 507-519.

[52] Yoshimura Y, Dantzker JL, Callaway EM (2005) Excitatory cortical neurons form fine-scale functional networks. Nature 433: 868-873.

[53] Kampa BM, Letzkus JJ, Stuart GJ (2006) Cortical feed-forward networks for binding different streams of sensory information. Nature Neuroscience 9 (12): 1472-1473.

[54] Perin R, Berger TK, Markram H (2011) A synaptic organizing principle for cortical neuronal groups. PNAS 108(13): 5419-5424.

[55] Rao SG, Williams GV, Goldman-Rakic PS (1999) Isodirectional Tuning of Adjacent Interneurons and Pyramidal Cells During Working Memory: Evidence for Microcolumnar Organization in PFC. J Neurophysiol 81: 1903-1916.

[56] Stettler D, Das A, Bennett J, Gilbert C (2002) Lateral connectivity and contextual interactions in macaque primary visual cortex. Neuron 14: 739–750.

[57] Binzegger T, Douglas RJ, Martin KAC (2009) Topology and dynamics of the canonical circuit of cat V1. Neural Networks, 1071-1078.

[58] Silverstein D, Lansner A (2011) Is Attentional Blink a Byproduct of Neocortical Attractors? Front Comput Neurosci 5 (2011): 1-14.

[59] Thomson AM, West DC, Wang Y, Bannister AP (2002) Synaptic connections and small circuits involving excitatory and inhibitory neurons in layers 2-5 of adult rat and cat neocortex: triple intracellular recordings and biocytin labelling in vitro. Cereb Cortex 12: 936-953.

[60] Pesaran B, Pezaris JS., Sahani M, Mitra PP, Andersen RA (2002) Temporal structure in neuronal activity during working memory in macaque parietal cortex. Nat Neurosci 5: 805-811.

[61] Gray CM, Singer W (1989) Stimulus-Specific Neuronal Oscillations in Orientation Columns of Cat Visual Cortex. PNAS 86: 1698-1702.

[62] Ward L.M (2003) Synchronous neural oscillations and cognitive processes. Trends Cogn. Sci. 7: 553-559.

[63] Tallon-Baudry C, Bertrand O, Peronnet F, Pernier J (1998) Induced gamma-band activity during the delay of a visual short-term memory task in humans. J Neurosci. 18: 4244-4254.

[64] Tallon-Baudry C, Bertrand O (1999) Oscillatory gamma activity in humans and its role in object representation. Trends in Cognitive Sciences 3: 151-162

[65] Fries P, Nikolic D, Singer W (2007) The gamma cycle. Trends Neurosci 30: 309–316.

[66] Fries P, Womelsdorf T, Oostenveld R, Desimone R (2008) The Effects of Visual Stimulation and Selective Visual Attention on Rhythmic Neuronal Synchronization in Macaque Area V4. J Neurosci 28: 4823-4835.

[67] Howard MW, Rizzuto DS, Caplan JB, Madsen JR, Lisman J., Aschenbrenner-Schreibe R, Schulze-Bonhage A, Kahana MJ (2003) Gamma Oscillations Correlate with Working Memory Load in Humans. Cereb Cortex 13: 1369-1374.

[68] Jacobs J, Kahana M (2009) Neural Representations of Individual Stimuli in Humans Revealed by Gamma-Band Electrocorticographic Activity. J Neurosci 29: 10203-10214.

[69] Brunel N, Hakim V (1999) Fast global oscillations in networks of integrate-and fire neurons with low firing rates. Neural Comput 11(7): 1621–71.

[70] Whittington MA, Traub RD, Kopell N, Ermentrout B, Buhl EH (2000) Inhibition-based rhythms: experimental and mathematical observations on network dynamics. Int J Psychophys 38: 315-36.

[71] Tsodyks MV, Sejnowski T (1995) Rapid state switching in balanced cortical network models. Network: Comput Neural Systems 6: 111–124.

[72] van Vreeswijk C, Sompolinsky H (1996) Chaos in neuronal networks with balanced excitatory and inhibitory activity. Science 274: 1724–1726.

[73] Vogels TP, Sprekeler H, Zenke F, Clopath C, Gerstner W (2011) Inhibitory Plasticity Balances Excitation and Inhibition in Sensory Pathways and Memory Networks. Science 334: 1569-1573.

[74] Joelving FC, Compte A, Constantinidis C (2007) Temporal properties of posterior parietal neuron discharges during working memory and passive viewing. J Neurophysiol 97: 2254–2266.

[75] Shafi M, Zhou Y, Quintana J, Chow C, Fuster J, et al. (2007) Variability in neuronal activity in primate cortex during working memory tasks. Neuroscience 146: 1082–1108.

[76] Tsodyks M, Pawelzik K, Markram H (1998) Neural networks with dynamical synapses. Neural Comp 10(4): 821-835.

[77] Axmacher N, Henseler MM, Jensen O, Weinreich I, Elger CE, Fell J (2010) Cross-frequency coupling supports multi-item working memory in the human hippocampus. PNAS 107 3228-3233.

[78] Canolty RT, Edwards E, Dalai SS, Soltani M, Nagarajan SS, Kirsch HE, Berger MS, Barbaro NM, Knight RT (2006) High Gamma Power Is Phase-Locked to Theta Oscillations in Human Neocortex. Science 313: 1626-1628.

[79] Jensen O, Colgin LL (2007) Cross-frequency coupling between neuronal oscillations. Trends Cogn Sci 11: 267-269.

[80] Palva JM, Palva S, Kaila K (2005) Phase synchrony among neural oscillations in the human cortex. J Neurosci. 25(15): 3962-3972.

[81] Kendrick KM, Zhan Y, Fisher H, Nicol AU, Zhang X, Feng J (2011) Learning alters the theta amplitude, theta-gamma coupling and neural synchronization in inferotemporal cortex. BMC, Neuroscience 12: 55.

[82] Raghavachari S, Kahana MJ, Rizzuto DS, Caplan JB, Kirschen MP, Bourgeois B (2001) Gating of human theta oscillations by a working memory task. J Neurosci 21: 3175–3183.

[83] Sederberg PB, Kahana MJ, Howard MW, Donner EJ, Madsen JR (2003) Theta and gamma oscillations during encoding predict subsequent recall. J Neurosci 23: 10809-10814.

[84] Fuentemilla L, Penny WD, Cashdollar N, Bunzeck N, Duzel E (2010) Theta-coupled periodic replay in working memory. Curr Biol 20: 1-7.

[85] Sirota A, Montgomery S, Fujisawa S, Isomura Y, Zugaro M, Buzsáki G (2008) Entrainment of neocortical neurons and gamma oscillations by the hippocampal theta rhythm. Neuron 60: 683–697.

[86] Huerta PT, Lisman JE (1993) Heightened synaptic plasticity of hippocampal CA1 neurons during a cholinergically induced rhythmic state. Nature 364: 723–725.

[87] Rutishauser U, Ross IB, Mamelak AN, Schuman EM (2010) Human memory strength is predicted by theta-frequency phase-locking of single neurons. Nature 464: 903–907.

[88] Bruederle et al. (2011) A comprehensive workflow for general-purpose neural modeling with highly configurable neuromorphic hardware systems. Biol Cybern 104: 263–296 DOI 10.1007/s00422-011-0435-9.

[89] Buffalo AE, Fries P, Landman R, Buschman TJ, Desimone R (2011) Laminar differences in gamma and alpha coherence in the ventral stream. PNAS 108 (27): 11262-11267.

[90] Pfurtscheller G, Stancak A, Neuper C (1996) Event-Related Synchronization (ERS) in the Alpha Band—An Electrophysiological Correlate of Cortical Idling: A Review. Int. J. Psychophysiol. 24: 39–46.

[91] Klimesch W, Sauseng P, Hanslmayr S (2007) EEG alpha oscillations: the inhibition-timing hypothesis. Brain Res Rev 53: 63-88.

[92] Brunel N, Wang XJ (2003) What determines the frequency of fast network oscillations with irregular neural discharges? I. Synaptic dynamics and excitation-inhibition balance. J Neurophysiol 90: 415-430.

[93] Wang XJ (1999) Synaptic basis of cortical persistent activity: the importance of NMDA receptors to working memory. J Neurosci 19: 9587–9603.

[94] Miller G (1956) The magical number seven, plus or minus two: some limits to our capacity to process information. Psychol Rev 63: 81-97.

[95] Cowan N (2001) Meta theory of storage capacity limits. Behavioral and Brain Sciences 24: 154-185.

[96] Lisman J, Idiart M (1995) Storage of 7 +/- 2 short-term memories in oscillatory subcycles. Science 267: 1512-1515.

[97] Wang Y, Markram H, Goodman PH, Berger TK, Ma J, Goldman-Rakic PS (2006) Heterogeneity in the pyramidal network of the medial prefrontal cortex. Nat Neurosci 9: 534-542.

[98] Meltzer JA, Zaveri HP, Goncharova II, Distasio MM, Papademetris X, Spencer SS, Spencer DD, and Constable RT (2008) Effects of working memory load on oscillatory power in human intracranial EEG. Cereb Cortex 18: 1843-1855.

[99] Mölle M, Marshall L, Fehm HL, Born J (2002) EEG theta synchronization conjoined with alpha desynchronization indicate intentional encoding. Eur J Neurosci 15: 923-928.

[100] Jensen O, Tesche CD (2002) Frontal theta activity in humans increases with memory load in a working memory task. Eur J Neurosci 15: 1395-1399.

[101] Shinomoto S, Shima K, Tanji J (2003) Differences in Spiking Patterns Among Cortical Neurons. Neural Comput 15(12): 2823–2842.

[102] Holt GR, Softky WR, Koch C, Douglas RJ (1996) Comparison of discharge variability in vitro and in vivo in cat visual cortex neurons. J Neurophysiol 75: 1806–1814.

[103] Shinomoto S, Miyazaki Y, Tamura H, Fujita I (2005) Regional and laminar differences in in vivo firing patterns of primate cortical neurons. J Neurophysiol 94: 567–575.

[104] Shinomoto S, Kim H, Shimokawa T, Matsuno N, Funahashi S, et al. (2009) Relating Neuronal Firing Patterns to Functional Differentiation of Cerebral Cortex. PLoS Comput Biol 5(7), e1000433.

[105] Tzeng OJL (1973) Positive recency effect in a delayed free recall. Journal of Verbal Learning and Verbal Behavior (12): 436–439.

[106] Beck DB, Kastner S (2009) Top-down and bottom-up mechanisms in biasing competition in human brain. Vision Res 49(10): 1154-1165.

[107] Broadbent DE, Broadbent MH (1987) From detection to identification: Response to multiple targets in rapid serial visual presentation. Perception & Psychophysics 42: 105–113.

[108] McArthur G, Budd T, Michie P (1999) The attentional blink and P300. Neuroreport 26: 3691–3695.

[109] Fell J, Klaver P, Elger CE, Fernández G (2002) Suppression of EEG Gamma Activity May Cause the Attentional Blink. Consciousness and Cognition 11: 114-122.

[110] Shapiro KL, Raymond JE, Arnell KM (1994) Attention to visual pattern information produces the attentional blink in rapid serial visual presentation. J Exp Psychol Hum Percept Perform 20: 357–371.