

Cortical Dynamics

Wolf Singer

Abstract

A hallmark of cortical organization is the coexistence of serial feedforward with re-entrant processing. The latter is based on feedback projections from higher to lower processing levels and massive reciprocal excitatory projections which link neurons located within the same cortical areas as well as cortical areas occupying the same level in the processing hierarchy. These reentrant connections, together with local negative feedback loops, give rise to exceedingly complex dynamics that are characterized by oscillations in a broad range of frequencies, synchronization of discharges, and cross-frequency coupling. Evidence is reviewed which suggests that these dynamic properties support specific computations: the flexible binding of distributed neurons into functionally coherent assemblies, the attention-dependent selection of sensory signals, the conversion of semantic relations into temporal relations, the comparison of stored priors with sensory evidence, the selective routing of signals in densely interconnected networks, the definition of relations in the context of learning, and the dynamic formation of functional networks. Arguments challenging a functional role of oscillations and synchrony, due to their volatile nature, are discussed in relation to recent evidence that highlights the advantages of volatility.

The Encoding of Relations

Organisms have evolved cognitive functions that enable them to construct an internal model of the world. This model permits interpretation of sparse and noisy sensory information, the generation of predictions, and the planning of well-adapted responses. To generate and exploit such models, organisms have developed efficient mechanisms to extract relevant features selectively from the plethora of available physical and chemical signals, to detect consistent spatial and temporal relations between these features, to memorize the corresponding relational constructs, to compare sensory evidence with stored knowledge, to evaluate the behavioral relevance of the actual conditions, and to select the appropriate response. In this chapter, I focus on mechanisms that underlie the detection and encoding of relations, as relations define the virtually infinite space populated by cognitive objects that can be generated

by combining a limited set of elementary features. With a limited number of symbols, for instance, alphabets sufficed throughout history to compose vast bodies of literature due to the richness and the combinatorial power of relational codes.

One common strategy for the detection of relations and their encoding in neuronal responses, realized in virtually all neuronal systems, is based on convergent feedforward circuits. Neurons tuned to respond to particular, frequently co-occurring, and hence related features of the environment are selectively connected to higher-order target cells. By adjusting the gain of these converging connections and the threshold of the target cell, it is assured that the latter responds preferentially to only a particular conjunction of features (Barlow 1972). In this way consistent relations among features become represented by the activity of conjunction-specific neurons. Iterating this strategy across multiple layers in hierarchically structured feedforward architectures then leads to the representation of complex relational constructs by conjunction-specific neurons of higher order. This basic principle for the evaluation and encoding of relational constructs has been realized independently during evolution in the nervous systems of different phyla (molluscs, insects, vertebrates) and reached the highest degree of sophistication in the hierarchical arrangement of processing levels in the cerebral cortex of mammals and, in particular, primates. It has also been implemented in numerous versions of artificial neural networks (Rosenblatt 1958; Hopfield 1987; DiCarlo and Cox 2007; LeCun et al. 2015). The highly successful, recent developments in the field of “deep learning” (LeCun et al. 2015) capitalize on the scaling of this principle in large multilayer architectures.

There are, however, marked differences between these artificial neuronal networks and the processing architectures found in the nervous systems of highly evolved species, suggesting that these exploit additional principles of information processing. This is particularly true for brains endowed with cortical structures, such as the hyperstriatum in birds and the hippocampus and neocortex in higher vertebrates. These structures possess feedback connections between processing levels; in addition, neurons within the same layer are reciprocally coupled by myriads of recurrent lateral connections. These reentry connections are missing in most feedforward artificial systems, but in natural brains they are more abundant than feedforward connections (Markov et al. 2014; Bastos et al. 2015). Differences between natural and artificial systems exist also with respect to the learning mechanisms. In technical systems, the supervised adjustment of the synaptic gain of feedforward connections is achieved by the so-called “backpropagation algorithm,” which is biologically implausible and differs from the various learning mechanisms implemented in natural brains. This latter difference is not crucial, because the results of training procedures are ultimately similar; however, the differences in processing architectures are consequential as they permit computations that go beyond those realizable in feedforward architectures. The latter possess no short-term

memory and therefore have difficulties to process and classify temporal relations. Recurrent networks, by contrast, exhibit fading memory and hysteresis because of the nonlinear dynamics that evolves on the backbone of reverberating circuits (see below). Although it is argued that any multilayered feedforward network can be unrolled so as to function like a recurrent network, the number of layers required to cope with behaviorally relevant temporal relations (e.g., in speech processing) is prohibitive. Recent designs of artificial systems, therefore, incorporate recurrent networks that realize a function characterized as long short-term memory (Hochreiter and Schmidhuber 1997; Silver et al. 2017; Banino et al. 2018).

Recurrent networks are not only better suited for the processing of temporal relations, they also offer complementary solutions for the encoding of spatial relations. As proposed by Donald Hebb (1949), consistent relations among features can be encoded by forming functionally coherent assemblies that bind individual, feature-coding neurons together. Neurons coding for features that typically co-occur, and thereby define a particular cognitive object, get bound together into an assembly which, as a whole, represents a concrete perceptual object or a more abstract cognitive content (e.g., category, concept, action plan). In this case, the *binding* of specific features is not achieved through the convergence of feedforward connections onto conjunction-specific neurons but by reciprocal connections between feature-selective neurons. These connections are strengthened by correlation-dependent synaptic plasticity mechanisms (Hebbian synapses, see below) and by selectively increasing the mutual interactions between nodes encoding a related feature, which enhance the vigor and/or coherence of the responses of the respective nodes. In this way, consistent relations among features characteristic for perceptual objects or, at higher processing levels, among more abstract contents, are translated into the weight distributions of reciprocal connections among the nodes of the network. Accordingly, the information about the presence of a particular constellation of features is not represented by the discharge of a single conjunction-specific neuron but by the amplified or more coherent responses of a distributed assembly of neurons. If a particular feature constellation matches the weight distributions of the recurrent connections linking the respective feature-sensitive nodes, the assembly will self-amplify its responses through reverberation and then assume the same function as a conjunction-specific neuron, except that now the coding unit is an assembly of cells.

Both relation-encoding strategies have advantages and disadvantages, and evolution has apparently opted for a combination of the two. Feedforward architectures are well suited to evaluate relations between simultaneously present features (e.g., spatial relations), and they allow for fast processing because they rely exclusively on a series of simple summation and thresholding operations. Moreover, they are easy to implement and are robust because interactions are essentially linear, well-controllable, and show no runaway dynamics. However, feedforward architectures are less apt to handle relations

among temporally segregated events because they lack memory functions. In addition, they are costly in terms of hardware requirements. All information about the statistical contingencies of features must be stored in the weights of the feedforward connections; because the dynamic range of neurons limits the number of converging driving connections, processing hierarchies require a large number of levels to cope with the combinatorial complexity of possible feature constellations. Since the number of required output units scales linearly with the number of relations that can be analyzed and encoded (combinatorial explosion), biological systems that rely exclusively on feedforward architectures can only afford representation of a limited number of behaviorally relevant relational constructs. Finally, feedforward architectures cannot easily cope with entirely new constellations of features because they lack the associative capacities of recurrent networks.

By contrast, assemblies of recurrently coupled, mutually interacting neurons can cope very well with the encoding of temporal relations (sequences) because such networks exhibit fading memory due to reverberation. Assembly codes are also much less costly in terms of hardware requirements, because individual feature-specific neurons can be recombined flexibly to yield a very large number of different assemblies, each representing a different cognitive content (combinatorial code). In addition, coding space increases dramatically because information about the statistical contingencies of features can be stored in the synaptic weights of feedforward connections as well as in the weights of the recurrent and feedback connections. Finally, the encoding of entirely new or the completion of incomplete relational constructs is facilitated by the nonlinear dynamics of recurrently coupled networks; their dynamics allows for self-organization and hence the completion of patterns and the generation of novel associations (generative creativity).

These advantages, however, come with a price. Processing may be slower than in purely feedforward architectures because assembly formation depends on time-consuming self-organizing processes based on the nonlinear dynamics of recurrent networks that exhibit reverberation, hysteresis, and attractor dynamics. Moreover, implementation of this distributed combinatorial coding strategy is not trivial; additional mechanisms are required to assure stability of network dynamics. To exploit the advantages of recurrent networks fully, their dynamical range needs to be well controlled because

- fast formation of assemblies requires a delicately regulated level of resting activity, and
- if global excitation drops below a critical level, recurrent networks may cease to operate or, if a critical level of excitation is reached, engage in runaway dynamics and become epileptic.

For cortical structures, this problem is taken care of by a number of cooperating, self-regulating mechanisms (involving inhibitory interneurons, excitation–inhibition balance, and ascending modulatory systems) that keep the

network within a narrow working range just below criticality, beyond which it would enter into a chaotic regime.

Another particularly challenging problem, known as the superposition catastrophe, involves the segregation of simultaneously active assemblies, in particular if they comprise spatially interleaved neurons or have to share some of the feature-selective neurons. If assemblies were solely distinguished, as proposed by Hebb, by enhanced activity of the constituting neurons, it becomes difficult to distinguish which of the more active neurons actually belong to which assembly (the “binding problem”). The option to multiplex in time coexisting rate-coded assemblies is also problematic because readout of enhanced discharge rate requires temporal integration in downstream structures. Given the low discharge rate of cortical neurons, it might take several hundreds of milliseconds before the more active members of assemblies become distinguishable from less active neurons, and hence multiplexing is achievable only on a slow timescale. Moreover the inability to configure temporally overlapping representations would jeopardize the associative capacities of assembly coding. It has thus been proposed that the salience of the responses of neurons temporarily bound into assemblies should not be enhanced solely by an increase of their discharge rate but by the precise synchronization of their action potentials (Gray et al. 1989; Gray and Singer 1989; Singer 1999). There is ample evidence that synchronous inputs are particularly effective in driving neurons above threshold. Thus, activation of target cells at the subsequent processing stage can be assured by increasing either the rate or the synchronicity of discharges in the afferents converging from the respective lower level. The advantage of increasing salience by synchronization is that integration intervals for synchronous inputs are very short, thus allowing for instantaneous detection of enhanced salience. Hence, information about the relatedness of responses can be read out very rapidly. *In extremis*, single spikes can be labeled as salient and belonging to a particular assembly if synchronized with a precision in the millisecond range. Thus, assemblies defined by synchrony rather than rate increases can be multiplexed at a much faster rate than rate-coded assemblies without becoming confounded.

At this point, one might argue that the readout of assemblies again requires conjunction-specific neurons that convert relational information into discharge rates and that assembly coding offers no advantages. This argument would probably hold if the output of the system would solely be conjunction-specific grandmother cells at the top of the sensory processing streams, or individual command neurons at the top of the inverse hierarchy of motor output. However, this is not the case. As far as we understand the system’s coding strategy, information remains distributed over many nodes: from the very early sensory structures all the way to the motor output. Thus, there is no bottleneck requiring condensation of relational information in the discharges of individual cells (grandmother cells), and hence assembly coding can be maintained throughout the whole processing stream. Assemblies formed at

a lower level can ignite corresponding assemblies at the respective next processing stage. The ability to cope with the “combinatorial explosion” (i.e., the virtually infinite number of possible feature constellations) is one of the advantages of assembly coding: it permits hardware-efficient, flexible binding of features represented by the reciprocally coupled network nodes at all levels of processing. However, the provision of an anatomical backbone for assembly coding is not the only advantage of recurrent networks. In addition to the ability to cope with the encoding of temporal relations mentioned above, it is likely that the complex dynamics which evolve in recurrent networks provide additional computational options. These, as well as more complex dynamic properties of cortical networks, are considered later in the chapter.

Temporal Relations as Code for Semantic Relations

Consistently (i.e., frequently and stereotypically) occurring temporal relations between real world events signal relatedness. Simultaneously occurring events usually have a common cause or are interdependent because of interactions. If one event consistently precedes the other, the first is likely the cause of the latter; if there are no temporal correlations between the events, they are most likely unrelated. Nervous systems exploit this fact in the detection and encoding of relations. Accordingly, learning rules adopted by evolution are sensitive to such temporal relations, thereby permitting the generation of internal models of relational constructs that have considerable predictive power—a likely reason for the striking conservation of the mechanisms supporting use-dependent modifications of synaptic transmission. Without exception, the associative mechanisms of synaptic plasticity evaluate correlations with a precision in the millisecond range. For the evaluation of correlations over longer time spans, required for the detection and encoding of contingencies separated by long intervals, additional mechanisms have been implemented. These involve memory functions at different timescales, ranging from fading memory in recurrent networks over short- to long-term memory mechanisms in devoted structures of the brain.

Expressing semantic relations in temporal relations does not, of course, exclude the common encoding of relations in the responses of conjunction-specific neurons (labeled line codes). However, unless the connections converging on such neurons are genetically determined, the input connections to the future conjunction-specific neurons have to be selected by experience. This selection, in turn, is again based on the time-sensitive mechanisms of synaptic plasticity, both during development and learning. Thus, the implementation of nongenetically specified, conjunction-specific “binding” neurons also requires that semantic relations become recoded in temporal correlations.

The exquisite sensitivity for temporal relations of associative synaptic plasticity mechanisms constrains the strategies used by cortical networks for

the encoding of relations. Therefore, the respective mechanisms will now be briefly reviewed.

Time-Sensitive, Associative Synaptic Modification Rules in a Nutshell

Interestingly, the mechanisms that support activity-dependent shaping of neuronal architectures during development and those that mediate use-dependent long-term modifications of synaptic gain, thought to underlie learning in the adult, share numerous similarities (for a review, see Singer 2018b). Major differences are that during development, functionally weakened synaptic connections eventually get physically and irreversibly removed while the pool of connections available for selection is permanently replenished through newly formed connections. The initial steps, however, that serve the evaluation of temporal relations are based on very similar molecular mechanisms and correspond to the rules proposed by Hebb (1949) for the strengthening of interactions between neurons that exhibit correlated activity: neurons wire together if they fire together. Hence statistical contingencies get translated into coupling strength. The first experimental confirmation of Hebb's hypothesis was the seminal discovery by Bliss and Lomo (1973) that tetanic stimulation of excitatory pathways in the hippocampus causes a long-term potentiation (LTP) of synaptic transmission. Subsequent studies demonstrated that LTP induction required a critical level of postsynaptic depolarization. If postsynaptic cells are prevented from responding to excitatory input by concomitant inhibition or hyperpolarizing current injection, modifications either do not occur or change to long-term depression (LTD) (Artola et al. 1990). With the advent of Ca imaging it became clear that the initial trigger for both LTP and LTD is a surge of calcium in the postsynaptic dendrites and that the polarity of the modifications depends on the rate of rise, amplitude, and sources of this Ca increase. Fast and strong increases lead to LTP, whereas slow and smaller increases trigger LTD (Bröcher et al. 1992; Hansel et al. 1996, 1997). Accordingly, both modifications can be obtained by raising intracellular Ca concentrations through the liberation of caged Ca in a concentration-dependent manner (Neveu and Zucker 1996). Moreover, the source of the Ca increase is of importance. Calcium entering through N-methyl-D-aspartate (NMDA) receptor-associated channels favors the induction of LTP, whereas Ca entering solely through voltage-dependent Ca channels is more likely to trigger LTD. The complex molecular cascades leading to these changes in synaptic gain have been thoroughly studied, and evidence indicates that both pre- and postsynaptic modifications are involved (for a review, see Morishita et al. 2005).

The evidence that both experience-dependent circuit modifications during development as well as use-dependent changes of synaptic efficiency in the

adult depend on correlations between pre- and postsynaptic activation found a mechanistic explanation when it was discovered that both involve activation of NMDA receptors. These function as coincidence detectors because they are permeable for Ca ions only if glutamate is bound to the receptor while the postsynaptic cell is sufficiently depolarized to remove the magnesium block (Nowak et al. 1984; Artola and Singer 1987; Kleinschmidt et al. 1987; Bear et al. 1990; for a review, see Collingridge and Singer 1990). Since the level of depolarization of the postsynaptic membrane does not only depend on the activity of the local excitatory synapses, but also on all the other excitatory and inhibitory inputs, this mechanism also accounts for the cooperativity that characterizes use-dependent synaptic modifications. Even weak inputs can increase their gain if they are active in synchrony with other nearby excitatory inputs that contribute to depolarization and the removal of the magnesium block. With the advent of two-photon imaging technology it became possible to demonstrate *in vivo* that contingent activation of weak inputs converging onto the same dendritic branch could induce sufficient depolarization to activate regenerative dendritic responses (Na and Ca spikes) and to induce LTP (Grienberger et al. 2015). Conversely, concomitant activation of inhibitory inputs can prevent even strongly activated inputs from depolarizing the postsynaptic dendrite above LTP threshold. In this case, presynaptic activity that would normally induce LTP may either induce LTD or induce no change at all (Artola et al. 1990).

The discovery that spikes can backpropagate into dendrites and contribute to the postsynaptic depolarization that gates synaptic plasticity revealed yet another aspect of time-sensitive plasticity (Markram et al. 1997; Bi and Poo 1998; Stuart and Häusser 2001). Varying the timing between a single excitatory postsynaptic potential (EPSP) and the backpropagating spike showed that small changes in the *temporal relations* have a massive impact on synaptic modifications. When the EPSP precedes the backpropagating action potential by less than 50 ms, the synapse potentiates and the strength of potentiation increases with decreasing delay. However, once the EPSP occurs after the backpropagating spike, there is a sharp transition toward LTD. The underlying mechanism is the same as detailed above. If the backpropagating spike occurs shortly *before* the EPSP, it can contribute to lifting the Mg block, thus allowing LTP to occur; if it arrives *after* the EPSP, the repolarizing currents prevent NMDA receptor activation, and LTD is the likely result. This special case of a use-dependent synaptic modification, known as spike timing-dependent plasticity (STDP), has an important implication: use-dependent synaptic modifications are not only sensitive to the coherence of converging activity but also to causal relations. The gain of excitatory connections increases if their activity can be causally related to the activation of the postsynaptic neuron and weakens when this is not the case.

These empirical results have been formalized in rules addressed as the BCM (Bienenstock et al. 1982), ABS (Artola et al. 1990), and STDP (Markram et

al. 1997; Bi and Poo 1998) rules. In conclusion, the net effects of these use-dependent synaptic modifications of excitatory connections are

- a strengthening of (reciprocal) connections among pairs of cells that are frequently activated in temporal contiguity or when one cell successfully drives its target neuron,
- a strengthening of the gain of converging inputs that are frequently active in temporal contiguity,
- a weakening of connections among pairs of cells whose activity is un- or anti-correlated or when one cell discharges shortly after its target neuron,
- a weakening of inputs active in contiguity with inhibition of the postsynaptic cell, and
- a weakening of connections that are inactive while the postsynaptic cell is strongly activated by other inputs (heterosynaptic depression).

Thus, the crucial variable that determines the occurrence and polarity of synaptic gain changes is the *temporal relation* (contiguity) between discharges in converging presynaptic inputs and/or between the discharges of presynaptic afferents and the depolarization of the postsynaptic neuron. These mechanisms evaluate correlation patterns with a precision in the range of tens of milliseconds and, in the case of STDP, even in the millisecond range. Moreover, the dependence of synaptic modifications on cooperativity between pre- and postsynaptic activity predicts that synchronous oscillations should provide a particularly favorable condition for the induction of use-dependent synaptic modification. This prediction is supported by the finding that use-dependent modification of orientation maps in the visual cortex is facilitated by entrainment of cortical networks in synchronous gamma oscillations (Galuske et al. 2019).

However, a learning mechanism has recently been discovered (Bittner et al. 2017) that operates on much longer timescales and is associative but does not seem to evaluate causal relations. If a postsynaptic neuron is sufficiently depolarized to generate a plateau potential—an active, Ca^{2+} -dependent dendritic response—then EPSPs arriving before *or* after this plateau potential undergo potentiation. This potentiation is maximal for EPSPs arriving in temporal contiguity with the plateau potential and decays with increasing temporal distance from the plateau potential. Thus, this mechanism is also sensitive to temporal coherence but allows for a gradation of potentiation as function of the temporal offset between the postsynaptic “reward” signal and incoming EPSPs. Obviously, similar to STDP, this mechanism is ideally suited for the learning of sequences as it converts sequence order into graded changes of synaptic efficiency. At first sight this mechanism seems to relax the constraint of precise timing relations between pre- and postsynaptic events. However, if information about sequence order were to be stored by the gradation of synaptic weights, the temporal relations between input sequences and the postsynaptic event, the plateau potential, need again to be precise and reproducible. Precise temporal sequences in population responses of cortical neurons are commonly

observed, and there is evidence from the visual cortex that the sequence order contains as much stimulus-specific information as the firing rate of individual neurons. When stimulated with a moving grating, cortical neurons discharge in a fixed sequence, as revealed by systematic offsets of cross-correlation peaks in the range of ± 10 ms. The order of the neurons in this sequence changes with stimulus orientation and allows determination of stimulus orientation and direction of motion with the same precision as measurements of discharge rates (Havenith et al. 2011).

Use-dependent gain changes have also been described for inhibitory connections onto excitatory neurons (I/E) and connections among inhibitory neurons (I/I). These modifications are also sensitive to the relative timing of pre- and postsynaptic activity but the modification rules appear to be more heterogeneous than for E/E connections, which matches the large diversity of interneuron types. The database for I/E and I/I plasticity is comparatively sparse but both Hebbian and anti-Hebbian modifications have been observed (Moore et al. 2010).

In conclusion, the time-sensitive mechanisms of synaptic plasticity, especially those identified for E/E connections, are well suited to convert information about the relatedness of events encoded in the precise timing relations of neuronal discharges in lasting changes of functional architectures.

Stimulus-Locked versus Internally Generated Temporal Relations

The fact that the learning rules are exquisitely sensitive to the precise temporal relations between individual discharges of connected neurons has far-reaching implications for the way nervous systems capture and encode information about relations. First, it has to be assured, if temporal relations between external events are to be evaluated, that the timing relations between events in the environment are reliably encoded in spike timing to permit learning of correct associations. This requirement is met in all sensory modalities by the implementation of transmission chains, commonly referred to as “phasic systems,” which operate with high temporal resolution and accuracy. *In vivo* recordings from higher visual areas as well as the auditory and the somatosensory cortex have revealed that the discharge latencies of individual neurons signal the temporal structure of stimuli with extreme precision in the millisecond range. This proves that precise timing of discharges can be preserved despite numerous intervening synaptic transmission steps (Buracas et al. 1998; Reinagel and Reid 2002). Simulation studies, partly based on the concept of synfire chains proposed by Moshe Abeles (1991), confirmed that conventional integrate-and-fire neurons are capable of transmitting temporal information with the required precision (Mainen and Sejnowski 1995; Diesmann et al. 1999).

Additional mechanisms are required, however, when selective associations have to be established between neuronal responses that lack precise temporal structure. Such binding operations are most likely required for the association of sensory responses to stimuli lacking a temporal dimension or for the association of internally generated activity. A parsimonious solution would be to implement intrinsic mechanisms that impose temporal structure on neuronal responses that satisfies the contingency requirements of the classical plasticity rules and to utilize the existing plasticity mechanisms also for the association and segregation of signals that initially lack temporal structure.

Mechanisms for the Generation of Temporally Structured Activity

Neuronal mechanisms capable of generating temporally structured activity are diverse, abundant, and evolutionarily ancient. A common and highly conserved strategy to generate temporally structured activity is the oscillatory patterning of activity, the basic principle of parsing time used in virtually all clocks. Neuronal networks have a high propensity to engage in oscillatory activity. These oscillations cover a broad frequency range, from below 0.1 to more than 200 Hz, and they tend to occur in typical frequency bands that are characteristic for particular brain structures and brain states. As reviewed recently (Buzsáki et al. 2013), these frequency bands are surprisingly well conserved across different species and even across different phyla. This suggests that they reflect some basic dynamics of nerve cells and/or circuits and are adapted to serve particular cognitive and/or executive functions. Spectral decomposition of global measures of brain activity, such as EEG or MEG or local field potential (LFP) recordings, usually reveals a continuous distribution of oscillation frequencies, the power of the oscillations decreasing with increasing frequency (the $1/f$ rule). To be addressed as such, an oscillatory process should appear in the power spectrum as a narrow-band “bump” of increased power. This distinction is important because with certain recording techniques (e.g., electrocorticography) one usually observes a task-associated broadband increase of power in the high-frequency range (from ~ 80 – 120 Hz) that is often addressed as “enhanced activity in the gamma-frequency range.” This broadband activity reflects synaptic currents and action potentials and is a good measure of neuronal activity but it must not be confounded with narrow-band gamma oscillations.

Common to all oscillatory processes is that the neurons engaged in an oscillation undergo periodic changes of excitability. Phases of increased excitability, often associated with action-potential generation, alter with phases of low excitability (Fries et al. 2007). The mechanisms causing these cyclic changes in excitability are heterogeneous. In certain cells, often addressed as pacemaker neurons or clock cells, these cyclic oscillations of excitability are caused by interactions among voltage-gated ion channels that have antagonistic effects

on the membrane potential (Heyer and Lux 1976). Here oscillation frequency depends on channel kinetics, membrane time constants, and driving forces. Individual pacemaker neurons usually operate in a characteristic frequency band and are typically found in pattern generator circuits that control rhythmic motion. Swimming, locomotion, respiration, heartbeat, and peristaltic movements are prime examples (Marder and Buchner 2001; Grillner 2006). However, cells with such properties are also found in structures traditionally not considered to be involved in pattern generation, such as the visual cortex. The chattering cells that engage in gamma oscillations when depolarized by current injection are one example (Gray and McCormick 1996). Recently it was discovered that such cells have narrow spikes and constitute a large fraction of the excitatory neurons in monkey V1 (Vinck, pers. comm.). Finally, even standard neurons can be considered as (relaxation) oscillators because action-potential firing is followed by refractory periods imposing cyclic alterations of excitability.

Another prominent mechanism for the generation of oscillations are circuit motifs giving rise to an oscillatory patterning of responses, the most common involving negative feedback loops, also addressed as recurrent inhibition. Excitatory neurons drive inhibitory neurons that inhibit the very same excitatory cells and these reciprocal antagonistic interactions naturally lead to an oscillatory patterning of the responses of both cell populations, the discharges of the inhibitory cells lagging slightly behind the discharges of the excitatory neurons (Whittington et al. 2000; Börgers and Kopell 2008; Buzsáki and Wang 2012). Here the frequency and regularity of the oscillations depend on a host of variables such as (a) the time constants of EPSPs, IPSPs, and dendritic integration, (b) the conduction delays of the feedback loops, (c) the excitatory drive, and (d) the embedding of the oscillator circuits in the network. Circuits tend to have their characteristic preferred oscillation frequency that can, however, be modulated over wide ranges by excitatory drive (Lowet et al. 2017).

Prominent examples are the septohippocampal circuits which generate the theta rhythm (Buzsáki 2006), the thalamocortical interactions responsible for the alpha rhythm (Steriade et al. 1993), and the cortical microcircuits which generate the beta and gamma oscillations known as ING and PING circuits (Kopell et al. 2000; Börgers and Kopell 2008; for a review, see Buzsáki et al. 2013). The excitability of cells participating in such oscillating circuits also undergoes a periodic modulation. At the peak of increased excitability, incoming EPSPs have a high probability to summate effectively and to generate action potentials while they are barely effective when arriving during the subsequent phase of enhanced inhibition. In this phase, EPSPs cannot summate effectively because of shunting inhibition and are less likely to reach firing threshold because of hyperpolarization. Thus, when cells oscillate, irrespective of whether the oscillations are due to pacemaker currents or circuit interactions, their ability to relay signals is modulated periodically, whereby the duration of the windows of opportunity decreases with oscillation frequency.

In conclusion, nervous systems are endowed with diverse and highly conserved mechanisms that are capable of imposing temporal structure on neuronal activity. How these temporal patterns can be used to establish precise temporal relations between the discharges of different neurons will be discussed in the following section.

Mechanisms for the Establishment of Temporal Relations

Consistent temporal relations among the discharges of pairs of neurons are commonly addressed as correlations. These can have several causes, some of which are trivial. Discharges can become correlated by external events. This is a frequent phenomenon, addressed as stimulus locking, and can be distinguished from internally generated correlations by trial shuffling (shift predictor). Simple causes for internally generated temporal correlations are (a) common input from bifurcating axons, in which case one observes sharp and single peaks in the correlogram with close to zero phase lag; (b) a direct connection from a feeder cell to the target cell, in which case the correlogram also exhibits a single sharp peak, but now with a consistent offset; (c) reciprocal E/E connections among excitatory neurons, in which case correlograms tend to have broader peaks centered around zero; and (d) common fluctuations of global excitability, usually caused by descending or ascending modulatory systems, that regulate excitability in a global and state-dependent way. In this case one observes joint fluctuations of discharge rates that lead to very broad peaks in the correlograms. These coordinated rate fluctuations are addressed as “noise correlations.” The term suggests that they are considered detrimental for information transmission, in particular for population coding (Averbeck et al. 2006). The controversies regarding the consequences of noise correlations will be discussed in detail later in this review.

A particularly versatile mechanism for the dynamic generation of temporal relations exploits the nonlinear interactions among *coupled oscillators* and therefore warrants an in-depth discussion.

The Establishment of Temporal Relations in Coupled Oscillator Networks

An important feature of oscillatory circuits is their propensity to resonate and be entrainable by periodically modulated inputs. As observed as early as 1665 by Christiaan Huygens, a Dutch watchmaker, very weak interactions suffice to synchronize coupled oscillators if their preferred frequencies are similar. Huygens noticed that pendulum clocks synchronized their beat when fixed to the same timber due to the weak interactions caused by mechanical coupling. If the difference between preferred frequencies increases, stronger coupling is required to assure synchrony with stable phase locking, and if the frequency

difference increases beyond a critical point, synchronization becomes unstable. Phase offset gradually increases and this may lead to intermittent phase resetting or a complete breakdown of synchrony. These complex and highly nonlinear relations have been analyzed in numerous theoretical studies (e.g., Winfree 1967; Aronson et al. 1990; Kuramoto 1990) and are summarized in the so-called Arnold tongue regime (Glass and Sun 1994). A graphical representation of synchronization behavior relating the difference in preferred frequency to increasing coupling strength leads to a “tongue”-shaped surface of possible synchronization regimes: the Arnold tongue.

Such reciprocal coupling between oscillatory circuits is a common motif in recurrently coupled neuronal networks and the cerebral cortex is a prime example. Thus, the Arnold tongue formalism can be applied to describe the relation between coupling strength and synchronization probability. Applied to the cerebral cortex, this predicts that the probability of two coupled columns synchronizing should increase with the strength of their reciprocal coupling. Experimental evidence from the visual cortex indicates that this is indeed the case. The network of the tangential intracortical connections that reciprocally couple neurons located in different functional columns is anisotropic. Columns responding to features that have a high probability to co-occur in natural scenes are more strongly coupled than columns tuned to features that are rarely contiguous (Gilbert and Wiesel 1989; Bosking et al. 1997; Stettler et al. 2002; Pecka et al. 2014). This selectivity of coupling is to a large extent due to experience-dependent pruning processes (Singer and Treter 1976; Löwel and Singer 1992; Smith et al. 2015) whereby the statistical contingencies of features in the outer world are translated into the functional architecture of the network of intracortical recurrent connections (Iacaruso et al. 2017).

In addition to coupling strength and preferred oscillation frequency, synchronization probability also depends critically on the conduction velocity of the coupling connections that varies over a wide range in neuronal networks. If conduction delays exceed a critical value, synchronization breaks down and interactions may lead to drifting phase behavior or a complete shutdown of the oscillations of one or all of the coupled oscillators (Aronson et al. 1990; Niebur et al. 1991; Reddy et al. 1998; Vicente et al. 2008; for a review, see Pajevic et al. 2014). Thus, synchronization by reciprocal coupling can only be achieved over larger distances if coupling connections are fast conducting or if oscillation frequency is reduced. This agrees with the evidence that, in general, long-distance synchronization is more common in the beta and alpha bands than in the gamma band (Buzsáki et al. 2013). However, there is evidence that high-frequency gamma oscillations can become synchronized between the two hemispheres and that this synchronization is mediated by the reciprocal callosal connections rather than a common synchronizing input from subcortical projections (Engel et al. 1991). Thus, even high-frequency oscillations can be synchronized over fairly large distances. Gamma oscillations can also synchronize

between distant cortical areas of the same hemisphere (Buschman and Miller 2007; Gregoriou et al. 2009), and analysis of phase lags and Granger causality suggests that this long-range synchronization results from direct (reciprocal) interactions. In this case, however, it cannot be excluded that the respective distant oscillators are synchronized in addition by common oscillatory input from a third source, which could either be another cortical area or a subcortical projection system (e.g., Saalman et al. 2012).

Still another possibility to coordinate high-frequency oscillations is phase locking to a slower oscillatory process. Evidence that fast oscillations can be coupled to slow oscillations is available and addressed as cross-frequency coupling (for review and critical discussion, see Palva et al. 2005; Canolty et al. 2006; Montgomery et al. 2008; Belluscio et al. 2012; Aru et al. 2015). The most commonly observed form of coupling is that the amplitude (power) of the fast oscillations is modulated periodically by the phase of the slow oscillation. However, it is not yet clear whether this coupling can actually assure phase synchronization of concomitantly modulated fast oscillators, as this would require phase locking or coordinated phase resetting of the fast oscillations rather than solely amplitude modulation.

Finally, an important variable determining synchronization probability is the connectivity motif of the coupling connections. It matters whether the coupling connections are excitatory, inhibitory or both, whether they impinge only on excitatory or inhibitory elements of the (oscillatory) circuits or on both, and whether more than two oscillators are coupled with one another (Vicente et al. 2008; Pérez et al. 2011). As the most frequent motif for the coupling of oscillatory circuits, anatomical evidence suggests reciprocal interactions between the inhibitory elements of oscillatory circuits (I-I connections) and excitatory connections that impinge both on the excitatory and inhibitory elements (E-E/E-I connections). Prominent examples are the septohippocampal circuits supporting the theta rhythm (Buzsáki 2006), the thalamocortical interactions responsible for the alpha rhythm (Steriade et al. 1993), and the cortical microcircuits generating the gamma oscillations, known as ING and PING circuits (Kopell et al. 2000; Börgers and Kopell 2008; Veit et al. 2017; for a review, see Buzsáki et al. 2013). In general, long-range interactions are mediated by glutamatergic projections but recent tracing studies indicate that inhibitory projections also span large distances. Examples are putative GABAergic neurons in the cerebral cortex that project across the corpus callosum (Buhl and Singer 1989), inhibitory cells in the septum and entorhinal cortex that innervate the hippocampus, and GABAergic cells in the basal forebrain that project to the cerebral cortex (for a review, see Caputi et al. 2013). These GABAergic long-range projections tend to innervate selectively inhibitory neurons in the respective target structures and are therefore in a good position to influence, via I-I interactions, the respective oscillatory circuits.

In conclusion, the probability that reciprocally coupled oscillators synchronize, irrespective of whether they consist of pacemaker neurons or oscillatory

microcircuits, increases with the similarity of their preferred oscillation frequencies, the strength of mutual coupling, and the conduction velocity of the reciprocal connections. In addition, state variables such as arousal, expectation, and readiness to act play an important role. They modulate the excitability of the nodes and thereby the entrainability of the coupled oscillator networks. An example is the cholinergic facilitation of induced gamma oscillations (Herculano-Houzel et al. 1999; Rodriguez et al. 2004).

So far only pairwise interactions have been considered but in reciprocally coupled networks, such as the cerebral cortex, interactions among multiple nodes lead to much more complex temporal patterns. These have been identified by multisite recordings and appear as stereotyped sequences of discharges distributed across nodes (songs, synfire chains) (Abeles 1991; Diesmann et al. 1999; Yuste et al. 2005; Gansel and Singer 2012) and traveling waves (Ermentrout and Kleinfeld 2001).

Altogether these mechanisms provide ample opportunities to impose temporal structure on neuronal activity and to establish precise temporal relations among discharges. These relations can then be converted by the established mechanisms of time-sensitive synaptic plasticity into modifications of functional architectures, whereby the selectivity with which the gain of synaptic connections between oscillating nodes can be modified depends critically on the oscillation frequency. For fast oscillations, for example, in the gamma range, the window of opportunity for the strengthening of synapses (i.e., the phase of heightened excitability) is very short, in the range of 10 ms or less (Wespatat et al. 2004), whereas for theta oscillations it is in the range of 100 ms. Thus, the faster the oscillations are, the greater the temporal selectivity is with which synchronous firing can be converted into lasting gain changes of synaptic connections.

In conclusion, the oscillatory patterning of neuronal responses and their synchronization can endow neuronal responses with the temporal structure required to make use of the established time-sensitive mechanisms of synaptic plasticity. This permits the selective association of neuronal populations whose responses are not time locked to stimuli but generated internally (e.g., during imagery or recall of memories). Evidence that this option is likely exploited has been provided by investigations on memory consolidation in human subjects (Miltner et al. 1999; Axmacher et al. 2008; Fell et al. 2011) and animals (for additional references, see Yamamoto et al. 2014; Singer 2017).

The Role of Oscillations and Synchrony in Information Processing and Coding

An undisputed function of oscillations in neuronal systems is the generation of rhythmic movements (see above). Furthermore, the cyclic changes

of excitability associated with oscillations play a crucial role in the gating of signal transmission and the routing of activity. Thus, transmission of sensory signals is modulated as a function of their timing relative to the phase of self-generated or stimulus-evoked oscillations (for a review, see Lakatos et al. 2008, 2013; Van Rullen 2016), and this modulation likely plays a role in the discontinuous sampling of information (Sergent et al. 2005; Landau and Fries 2012; Ni et al. 2017). Engaging in an oscillation can also increase the saliency of neuronal signals because it often results in burst firing (Gray and McCormick 1996). Bursts, in turn, enhance the impact of excitatory input on target cells without increasing average discharge rate, especially in sparsely connected networks (Larkum 2013).

Particularly consequential functions emerge once several nodes of a network engage in oscillations and synchronize. The high propensity of coupled oscillators to synchronize (see above) can be exploited to establish precise and highly specific temporal relations between the discharges of the neurons participating in the respective oscillatory circuits. These temporal relations can range from zero phase lag synchrony, in which case the engaged neurons discharge simultaneously, to synchronization with various phase lags, in which case one obtains precisely timed *discharge sequences* distributed across populations of neurons (Havenith et al. 2011). This option is exploited for a host of executive functions. The control of bird song, speech, and composite movements are but a few examples (Marder and Buchner 2001; Suthers and Margoliash 2002; Grillner 2006).

It is likely, however, that these dynamics may also play a role in cognitive processes. Provided that information can be encoded in and read out from temporal patterns in the discharges of distributed neurons, for which there is evidence (Masquelier et al. 2009), coding space can be considerably enlarged by exploiting an additional dimension of fine-grained temporal relations.

The Role of Oscillations and Synchrony in the Dynamic Gating of Neuronal Interactions

Another important function supported by the synchronization of nodes is the dynamic gating of interactions which permits flexible routing of signals across the backbone of fixed anatomical connections without requiring synaptic gain changes. If two anatomically coupled nodes engage in oscillatory activity, phase and frequency shifting can be used to selectively and dynamically modulate the strength of functional coupling without changing the gain of synapses. If phase relations are adjusted such that the EPSPs of one node arrive at the other during its phase of heightened excitability, transmission is facilitated, whereas it is blocked when phase is shifted by 180° (Fries 2005, 2009). Direct experimental evidence for such a phase-dependent gating mechanism has been

obtained in the visual system (Womelsdorf et al. 2007; Bosman et al. 2009; Bastos et al. 2015).

This dynamic gating mechanism has been shown to serve multiple functions. Recent evidence indicates that it is exploited for attention-dependent signal selection. If two oscillatory signals compete for transmission in an upstream target area, the signal wins that succeeds to phase lock with oscillations in the target area (for a review, see Maris et al. 2016).

Synchronization has also been shown to support dynamic configuration of functional networks. Accordingly, measures of coherence among oscillatory signals such as phase locking indices, paired phase consistency, Granger causality, and transfer entropy are commonly applied to identify functional networks. As expected, these measures reveal networks composed of nodes linked by strong anatomical connections because, as discussed above, these nodes have a higher probability to engage in synchronous activity than weakly coupled nodes. An example of such preconfigured functional subsystems is the default network or the language network. However, coherence analysis also revealed networks that get configured on the fly in a context- and task-dependent way.

Cats trained to perform a motor response to the change of a visual stimulus synchronized, in anticipation, those cortical areas that needed to be engaged to accomplish the task. In response to an auditory cue announcing the next trial, neurons in the visual, parietal, somatosensory, and motor cortex of the hemisphere controlling the motor response engaged in beta oscillations that were synchronized with close to zero phase lag *before* the visual stimulus appeared (Roelfsema et al. 1997). Thus, communication among processing stages required to accomplish the task is likely to be facilitated in anticipation by entrainment in coherent oscillations. By enhancing coherence of oscillatory activity, similar task-dependent formation of functional networks has been observed for diverse executive and cognitive functions that require dynamic coordination of distributed processing nodes (Buschman and Miller 2007; Siegel et al. 2008; Gregoriou et al. 2009; Grothe et al. 2012; Salazar et al. 2012; Dotson et al. 2014; for a review, see von der Malsburg et al. 2010).

Finally, as proposed following the discovery of synchronized gamma oscillations in the visual cortex, this gating mechanism also supports the dynamic binding of neurons driven by perceptually groupable (semantically related) features into functionally coherent assemblies (distributed coding). As reviewed above, the architecture and coupling strength of intercolumnar horizontal connections reflects the statistical contingencies of feature constellations. Neurons tuned to features that have a high probability to co-occur, and hence to be related, are more strongly coupled than neurons coding for unrelated features. Consequently, and as shown in numerous studies, neurons responding to related features—those that tend to be bound perceptually according to the well-established Gestalt rules for perceptual grouping and scene segmentation—have a high probability to synchronize their responses if their

preferred stimuli are presented together. Thus, synchronization probability reflects the Gestalt rules for perceptual binding. Nearby columns responding to closely spaced contours are all strongly coupled, which corresponds to the Gestalt rule of vicinity. More widely separated columns exhibit strong coupling only if they respond to the segments of continuous elongated contours or contours with similar orientation, especially when these are aligned collinearly and/or move with the same speed in the same direction. This reflects the Gestalt criteria of continuity, similarity, collinearity, and common fate. As shown in numerous experimental and simulation studies, neurons in these strongly coupled columns have a high probability of synchronizing their oscillatory discharges with close to zero phase lag when co-activated by their preferred features (Gray et al. 1989; Gray and Singer 1989; König and Schillen 1993; Schillen and König 1994; Kohn and Smith 2005; for reviews, see Singer 1999 and Uhlhaas et al. 2009).

Thus, temporal contiguity of discharges appears again to serve as a code of relatedness. However, in this case it is not used for developmental pruning of connections or the association of neuronal assemblies in the context of learning; it serves to define relations during the processing of visual patterns as proposed by the “binding by synchrony” hypothesis.

In conclusion, precise timing relations among neuronal discharges are used during development and learning to translate consistent correlations among features in lasting changes of coupling strength, and these anisotropies in coupling, in turn, determine the probability of synchronization once such “trained” networks are presented with sensory stimuli. Thus, synchronization is used in an instructive way during development and learning to store “knowledge” about statistical contingencies in the functional architecture of networks, and it is then used again for the readout of stored knowledge during information processing.

This gating of neuronal interactions by synchronization of oscillatory activity likely plays a more important role than is suggested by the firing statistics of neurons considered in isolation. The reason for this is that neurons participate considerably more often in oscillatory processes than is suggested by extracellular recordings of discharges. Even when the population of interacting neurons engages in regular oscillations—as indicated by oscillatory field potential fluctuations and periodic fluctuations of the cells’ membrane potential—the discharges of individual cells may still exhibit Poisson statistics because of cycle skipping (Zeitler et al. 2008; Palmigiano et al. 2017).

While there is common agreement that precise correlations among neuronal discharges play a crucial role in synaptic plasticity during developmental pruning and learning, the proposal that such correlations also play a role in signal processing has given rise to substantial controversies. Thus a brief summary of the most pertinent arguments for and against a functional role of synchrony in signal processing seems warranted when discussing the role of cortical dynamics.

Counterarguments

It has been argued that synchronized oscillations cannot have a functional role because they are too volatile and nonstationary to support reliable processing. To appreciate this criticism, it is necessary to consider the constraints of the commonly applied methods for the detection of oscillations and synchrony, such as Fourier and wavelet analysis, auto- and cross-correlation techniques, or various coherence measures. These techniques are applicable only if oscillation frequencies are reasonably stable and if oscillatory episodes last for a substantial number of cycles. Even for the analysis of the high-frequency gamma oscillations, measurement windows typically range from 200 ms to 1 s. As a consequence, most studies investigating oscillations and synchrony have been performed under conditions that favor the occurrence of sustained, frequency stable oscillations. Typical cases are the induction of gamma oscillations in the visual cortex with drifting grating stimuli or isolated contours. These studies generated a host of interesting insights into the dependence of oscillations and synchrony on stimulus features, central states, attention, and behavioral goals. However, because of the applied stimulation and analysis techniques, these studies also nurtured the notion that functionally relevant synchronization can only be established under very restricted stimulation conditions and that it is absent or too spurious to serve a function under more natural conditions, especially when information processing has to occur on fast timescales, as is the case during visual exploration. Primates perform on average four saccades per second, which implies that new visual information is sampled every 250 ms (Maldonado et al. 2008; Ito et al. 2011), and psychophysical evidence indicates that familiar scenes can indeed be segmented and recognized within such short intervals. Thus, it has been argued that there is not enough time to rely on synchronization of sustained oscillations for feature binding and the formation of object-specific neuronal assemblies. Even more problematic is that synchronous oscillations, when assessed with conventional Fourier or wavelet analyses, decrease in amplitude or become undetectable when complex stimuli are presented, such as cluttered scenes or images that lack clear high-contrast boundaries (Lima et al. 2010). Moreover, it has been argued that the dependence of oscillation frequency and power on contrast, motion speed, eccentricity, and complexity of stimuli is incompatible with the idea that spike synchronization can be used to encode semantic relations, as postulated in the binding by synchrony hypothesis (Singer 1993; Singer and Gray 1995), and to gate communication by coherence, as postulated by the CTC hypothesis (Fries 2005). These arguments have been made explicit by Atallah and Scanziani (2009), Burns et al. (2010, 2011), Ray and Maunsell (2010), Jia et al. (2013a, b); for reviews, see Ray and Maunsell (2015) and Palmigiano et al. (2017).

Before reviewing more recent evidence that counters these arguments, it is worth noting that the constraint of processing speed also poses a problem

for coding strategies that rely on joint rate increases for the definition of relations, because discharge rates of cortical neurons are low and can carry only little information when integrated over short intervals. This problem could, in principle, be solved by averaging over large populations of cells. However, considerations on sparse coding and noise correlation make this solution appear suboptimal as well (Averbeck et al. 2006). For these very reasons, it has been proposed that encoding relatedness by spike synchronization, rather than by joint rate increases, accommodates the superposition problem inherent in assembly coding (see above) as well as the constraint of processing speed (Singer 1999; Van Rullen et al. 2001, 2005; for a review, see Korndörfer et al. 2017).

Supporting Evidence

Recent experimental studies that combine multisite recordings with time-resolved analysis of coherent oscillations and spike synchronization in awake, behaviorally trained monkeys have shown that high-frequency oscillations and their synchronization also occur under natural conditions. In addition, they have confirmed the volatile features of synchronization phenomena: nonstationarity, frequency variability, short duration, rapid phase shifts, and fast formation and dissolution of coherent states. These studies also revealed, however, that the brief bouts of coherent activity contain information about the contents of working memory (Lundqvist et al. 2016), the communication between cortical areas (Siegel et al. 2008), the dynamic formation of functional networks (Buschman and Miller 2007), and the direction of information flow (Bastos et al. 2015; Lowet et al. 2017).

The indications that even short periods of synchronization are functionally relevant and informative are in perfect agreement with the dynamics of simulated recurrently coupled networks of spiking neurons. Korndörfer et al. (2017) recently demonstrated that neurons in reciprocally coupled networks engage very rapidly in synchronous discharges when activated by structured input, and that the synchronization probability is determined by the strength of coupling. This confirms the experimental evidence that neurons responding to features that need to be bound together for perceptual grouping and scene segmentation synchronize their responses if activated by patterns containing groupable features (Gray et al. 1989). In the study by Korndörfer et al. (2017), no explicit oscillatory properties of the nodes were implemented, but the spiking neurons had the usual refractory period and hence shared features of relaxation oscillators. The additional implementation of inhibition influenced network dynamics, as expected, but did not interfere with the fast synchronization of spikes.

Another comprehensive simulation study by Palmigiano et al. (2017) investigated the effect of synchronization on information transfer in coupled oscillator networks. They simulated *delay-coupled* recurrent networks with

spiking excitatory and inhibitory neurons that shared essential connectivity motifs of supragranular cortical layers. In this study, two regimes of network dynamics were examined. One, addressed as the transient synchrony regime, was characterized by low coherence, as is typically observed for resting-state activity in awake animals or when cortical networks are challenged with complex stimuli. Spike statistics were stochastic but field potentials computed from averaged membrane potential fluctuations exhibited bursts of synchronous gamma oscillations which only lasted a few cycles. When the strength of recurrent inhibition was increased, together with an increase in background drive, oscillation frequency and discharge rate remained in the same range, but oscillations became more sustained and synchrony more robust. This second regime resembles *in vivo* conditions when responses are evoked with regularly structured, high-contrast stimuli that induce strong excitation and inhibition such as drifting gratings. The interesting and unexpected outcome of measurements of information transfer in this delay-coupled oscillatory network was that phase-dependent gating of information transfer was as effective and even more flexible in the transient synchrony regime than in the regime exhibiting more sustained oscillations and synchrony. The short duration of the oscillatory bursts allowed for fast opening and closing of transmission channels by frequency tracking and the direction of information flow rapidly reversed if phase relations changed. Palmigiano et al. concluded that “features that at first sight appear to be noncompliant with information routing may actually provide the brain with a particularly flexible routing mechanism.” This conclusion received strong support from the recent experimental study by Lowet et al. (2017), who investigated the dynamics of synchronized gamma oscillations in V1 of awake monkeys using massive parallel recordings. The observed dynamics resemble in great detail those reproduced by the simulated network and confirm essential predictions of the hypothesis that supragranular layers of the cerebral cortex can be considered as a delay-coupled recurrent oscillator network. In this study, all predictions derived from the Arnold tongue formalism could be confirmed experimentally.

Taken together, both the results of electrophysiological experiments and simulation studies indicate that synchronization of spike discharges can be fast enough to serve feature binding/perceptual grouping within the short intersaccadic fixation intervals (see also Lowet et al. 2016). Moreover, the transient and variable nature of synchronized gamma oscillations, characteristic of free viewing conditions and the processing of complex scenes, does not compromise the many other functions assigned to synchronization. Rather, the fast fluctuations between synchronized and uncorrelated states are advantageous for the rapid and flexible definition of relations between distributed neuronal responses, as is required for attention-dependent input selection, the flexible and selective routing of signals on the backbone of the fixed connectome, the task-dependent formation of functional networks, and the flexible binding of distributed feature detectors into functionally coherent assemblies. Recent

experimental results obtained in awake monkeys by Lowet et al. (2016), Bosman et al. (2009) on the effect of microsaccades, and Brunet et al. (2015) on oscillations during free viewing of natural scenes are fully compatible with this view.

Dynamics beyond Oscillations and Synchrony

The nonlinear interactions within delay-coupled oscillator networks give rise to very complex dynamics that provide a very high-dimensional space that, in principle, can be exploited for flexible and efficient computation. In the following, some of these options are discussed and substantiated with recently obtained experimental evidence. I argue that such networks are particularly suited to cope with a number of hitherto poorly understood functions: the encoding of temporal sequences, the storage of vast amounts of information about the environment in the networks of sensory cortices, the ultrafast retrieval of information in processes requiring comparison between input signals and stored knowledge, and the fast and effective classification of complex spatiotemporal input patterns.

Theories of perception, formulated more than a hundred years ago (von Helmholtz 1896), and a plethora of experimental evidence indicate that the brain interprets sparse and impoverished input signals on the basis of an internal model of the world. The information provided by this model is used to reduce redundancy, to facilitate segregation of figures from background, to bind signals evoked by features constituting a perceptual object, and to enable classification and identification. Given the daunting complexity of the visual world, the store containing such an elaborate model must have an immense capacity to accommodate the vast number of statistical contingencies required for the interpretation of ever-changing sensory input patterns. Moreover, this massive amount of prior knowledge needs to be arranged in a configuration that renders it accessible within fractions of a second to meet the constraint of known processing speed.

The proposal is that these requirements can be met if encoding, storage, and processing of information take place in the high-dimensional state space provided by complex systems with nonlinear dynamics.

As reviewed above, information about contingencies in the outer world is stored in the synaptic weight distributions of the tangential intracortical connections that reciprocally couple feature selective nodes. In low visual areas, the nodes are selective for elementary features, such as the orientation of contour borders, whereas in higher visual areas, the nodes code for increasingly complex constellations of elementary features. Accordingly, one expects that the weight distributions of the tangential coupling connections in these areas reflect contingencies of higher order. However, because multisite recordings and correlation analysis have not been performed in higher visual

areas, we know little about the interactions between network nodes in higher visual areas.

Resting- and Stimulus-Induced States

The brain's spontaneous activity is constrained by its functional architecture. Hence, the dynamics of resting activity must reflect the weight distributions of the structured network that harbors the entirety of latent internal priors. This predicts that resting activity is high dimensional and represents a vast but constrained manifold inside the universe of all theoretically possible dynamical states. If input signals become available, they are likely to trigger a cascade of effects: they drive in a graded way the feature-sensitive nodes, thereby constraining the network dynamics. If the evidence provided by the input patterns matches well the priors stored in the network architecture, the network dynamics will collapse to a specific substate that provides the best match with the corresponding sensory evidence. Such a substate is expected to have a lower dimensionality than the resting activity, to exhibit specific correlation structures, and to be metastable due to reverberation among nodes supporting the respective substate. Because these processes occur within a very high-dimensional state space, substates induced by different input patterns are likely to be well segregated and therefore easy to classify. They can then serve either as input to the next cortical processing stage, where higher-order priors are implemented, permitting emergence of more abstract interpretations, or they can be classified by local readout units that directly feed into executive centers.

Analogies from Computational Studies

In a much simplified version, the nonlinear dynamics characteristic of recurrent networks are exploited for computation in certain AI systems, the respective strategies being addressed as “echo state, reservoir, or liquid computing” (Buonomano and Maass 2009; Lukoševičius and Jaeger 2009; D’Huys et al. 2012). The following Gedanken experiment illustrates the principle.

Objects impact at different intervals and locations in a pond of water and generate propagating waves whose parameters reflect the size, impact speed, and location of the object. The wave patterns fade with a time constant determined by the viscosity of the liquid, interfere with one another, and create a complex dynamic state. This state can be analyzed by measuring, at several locations in the pond, the amplitude, frequency, and phase of the respective oscillations, and from these variables a trained classifier can subsequently reconstruct the exact sequence and nature of the impacting “stimuli.”

Fernando and Sojakka (2003) put these ideas into practice using an actual bucket of water and showed that the interference between waves on the water surface allowed a simple perceptron to solve the XOR problem. The advantages of this computational strategy are as follows:

1. Low-dimensional stimulus events are projected into a high-dimensional state space, where nonlinearly separable stimuli become linearly separable.
2. The high dimensionality of the state space can allow for the mapping of more complicated output functions (like the XOR) through simple classifiers.
3. Information about sequentially presented stimuli persists for some time in the medium (fading memory). As such, information about multiple stimuli can be integrated over time, allowing for the mapping of complicated temporal functions.

These properties make artificial recurrent networks extremely effective for complex sequence processing as demonstrated by a number of applications: grammar learning, automatic driving systems, generation of handwritten text, image captioning, and so forth. Recent simulation studies have actually shown that performance of an artificial recurrent network is substantially improved if the recurrent connections are made adaptive and can “learn” about the feature contingencies of the processed patterns (Lazar et al. 2009; Hartmann et al. 2015). Interestingly, the local plasticity mechanisms used for learning and adaptation in these models have lasting consequences on the spatiotemporal structure of spontaneous activity. Changes in the structure of spontaneous activity are consistent with the statistical properties of the previously observed stimuli and have functional consequences for stimulus disambiguation and interpretation (Hartmann et al. 2015).

Experimental Evidence

The structural similarities between these artificial recurrent networks and supragranular layers in the cerebral cortex suggest that some of the computational strategies applied in reservoir computing might actually be used by the cerebral cortex. Experimental studies testing this hypothesis are still rare and have become possible only through the advent of massive parallel recordings from the network nodes. So far, however, the few predictions that have been subject to experimental testing appear to be confirmed. The covariance structure of resting activity does indeed reflect the anisotropic layout of the tangential reciprocal connections (Bosking et al. 1997; Fries et al. 2001a; Kenet et al. 2003), is modified by learning (Lewis et al. 2009; Kundu et al. 2013), and reveals hallmarks of an internal model of the environment (Berkes et al.

2011). Evidence is also available that cortical computations exploit the high-dimensional dynamics of recurrent networks to encode information about stimulus sequences. Responses to successively presented visual stimuli (letters and numbers) were recorded with matrix electrodes simultaneously from a random sample of ~ 60 neurons in cat primary visual cortex; linear classifiers were trained on short segments (5–100 ms) of the activity vectors of a training set of responses; these classifiers were then used to identify the nature of the presented stimuli in a test set (Nikolic et al. 2009). These experiments revealed the following:

1. Information about a particular stimulus persists in the activity of the network for up to a second after the end of the stimulus (fading memory).
2. Information about sequentially presented stimuli superimposes so that two subsequent stimuli can be correctly classified some time after the end of the second stimulus.
3. Information about stimulus identity is distributed across many neurons (> 30) and encoded in the correlation structure of the responses.

Evidence has also been obtained that sensory signals matching the priors stored in the network cause a collapse of high-dimensional network dynamics into metastable subregions of the state space that are stimulus specific and distinguished by enhanced coherence (covariance, synchrony), reduced variability, and lower dimensionality (Churchland et al. 2010; Banyai et al. 2018; Klein et al., in prep.; Lazar et al., in prep.). Stimuli that do not match the stored priors (e.g., “unnatural” stimuli) evoke substates that are less coherent, less stimulus specific (Banyai et al. 2018), and more difficult to classify (Lazar et al., pers. comm.). This could explain why grating stimuli elicit particularly strong synchronized gamma oscillations. A grating matches several priors (Gestalt criteria, grouping criteria), namely those of continuity, collinearity, similarity (in this case in the orientation domain), regularity, and (if the grating drifts) common fate. As discussed by Vinck and Bosman (2016), stimuli that match predictions stored in the cortical network appear to give rise to particularly well-synchronized gamma oscillations. Evidence also indicates that the cortical network “learns” about stimulus statistics and exploits this knowledge to optimally segregate the representations of different stimuli. Repeated presentation of stimuli has been shown to cause unsupervised changes in the network with the consequence that familiar stimuli evoke substates that are better classifiable than those evoked by less familiar stimuli, because they are better segregated in high-dimensional dynamic space (Lazar et al., in prep.). These experience-dependent modifications of the network have an impact also on the structure of resting-state activity: the vectors specific for highly familiar stimuli are spontaneously replayed (Lazar et al., pers. comm.).

There are also indications that top-down mechanisms related to attention and expectancy constrain the dynamic space in anticipation of having to

respond to the cued stimulus. This top-down effect is associated with a massive enhancement of gamma oscillations induced by the attended (Fries et al. 2001b) or expected stimulus (Lima et al. 2011). Also, one observes a ramp up of the power of gamma oscillations and a change in dimensionality of the network dynamics following presentation of a cue that instructs the animal about the sequence of future events (Klein et al., in prep.). These top-down influences change the correlation structure of the activity vectors and could contribute to speeding up the formation of specific substates once sensory evidence becomes available.

Concluding Remarks

Despite considerable effort, there is still no unifying theory of cortical processing. As a result, numerous experimentally identified phenomena lack a cohesive theoretical framework. This is particularly true for the dynamic phenomena reviewed here, because they cannot easily be accommodated in the prevailing concepts which emphasize serial feedforward processing and labeled line codes. However, with its preponderance of reciprocal connections and the rich dynamics that result from these reciprocal interactions, the cortical connectome suggests that additional processing strategies are implemented. Here I proposed concepts that assign specific functions to oscillations, synchrony, and the more complex dynamics that emerge from a delay-coupled recurrent network. These concepts are fully compatible with the robust evidence for the encoding of relations by anatomical convergence (labeled line codes) but complement this mechanism through a scenario in which the precise temporal relations among the discharges of coupled neurons serve as complementary code for the definition of semantic relations, both in signal processing and learning. In addition, I introduced a computational strategy that capitalizes on the high-dimensional coding space offered by reciprocally coupled networks. In this conceptual framework, information is distributed and encoded in the discharge rate of individual nodes (labeled lines) as well as, to a substantial degree, in the precise temporal relations among the discharge sequences of distributed nodes. The core of the hypothesis is that the dynamic interactions within delay-coupled recurrent oscillator networks (a) endow responses with the precise temporal structure required for the encoding and learning of semantic relations, (b) exhibit complex, high-dimensional correlation structures that reflect the weight distributions of the coupling connections and serve as an internal model, and (c) permit fast convergence toward stimulus-specific substates that are easy to classify because they occupy well-segregated loci in the high-dimensional state space. Analysis of the correlation structure of these high-dimensional response vectors is still at the very beginning. However, methods are now available for massive parallel recording from large numbers of network nodes in behaving

animals. It is to be expected, therefore, that many of the predictions derivable from this novel concept will be amenable to experimental testing in the near future.