

Neocortical Rhythms

An Overview

Wolf Singer

Abstract

Information processing systems need to be able to identify and encode relations. Relations can be defined in space and time. Nervous systems exploit both dimensions for the handling of relations. Their anatomical layout is characterized by selective convergence of connections on target cells, allowing the establishment of relations among signals of different origin. In addition, relations can be expressed more dynamically by adjusting the temporal rather than the spatial contiguity of signals. It is proposed that this is achieved by rhythmic modulation of neuronal activity and context as well as task-dependent modulation of oscillation frequencies and phases.

The Encoding of Relations

Analyses of neuronal responses in sensory systems reveal a characteristic sequence of processing steps. At subcortical levels, neurons encode local properties of perceptual objects and barely any relations among these elementary features. As one proceeds along the hierarchically arranged processing areas of the neocortex, responses become increasingly selective for complex constellations of elementary features. This sensitivity is the result of interactions mediated by intracortical tangential connections and iterative recombination of feedforward connections from lower- to higher-order neurons. Thus, responses of higher-order neurons signal not only the presence of certain sets of component features but also prespecified relations among these features.

This binding of features by recombination of feedforward connections (labeled-line coding) needs to be complemented by mechanisms that permit flexible encoding of relations to resolve the ambiguities inherent in the responses to natural environments. Natural scenes, for example, usually contain a large number of different objects, the contours of which may be overlapping, partially occluded, or superimposed on those of the background. A neuron may thus

be stimulated by contours belonging to different objects. To avoid the encoding of false relations (conjunctions), these ambiguities need to be resolved before signals can be subject to grouping in the respective convergent pathways. This sorting of appropriately groupable responses must occur in a context-dependent way at each level of processing. At low levels, it is indispensable for scene segmentation; at high levels, it is required for the disambiguation of simultaneously configured distributed representations (assemblies) (Wang 2005).

Experimental data and theoretical considerations suggest that perceptual objects are represented not only by individual neurons that respond selectively to single objects but also by distributed assemblies of cells (Singer 1999; Tsunoda et al. 2001). First, object-specific neurons are rare and seem to exist only for highly overlearned objects or for objects of particular behavioral relevance (Quiroga et al. 2005; Logothetis et al. 1994). Second, novel objects cannot be represented by preestablished neurons because the required feedforward architectures would have to be specified a priori to support formation of the appropriate conjunctions. Third, objects that are simultaneously encoded in different sensory modalities elicit responses from several different sensory systems, and these need to be interrelated to arrive at a comprehensive polymodal description of a particular object. Here, the space of possible conjunctions is so large that it cannot be covered by applying the “labeled line” strategy alone. These considerations suggest that objects which cannot be represented by individual conjunction-sensitive neurons are encoded by assemblies of distributed neurons, each of which represents only a particular component of the object. In assembly coding, however, a relation-defining mechanism is required that tags responses which are evoked by the components of the same object as related. In essence, neurons have to convey two orthogonal messages in parallel: (a) they must signal whether the feature or conjunction of features for which they stand as a “labeled line” symbol is present, and (b) they must indicate, from instance to instance, with which subset of the myriads of simultaneously active neurons they are actually forming a coherent assembly. There is consensus that the first message is conveyed by a rate code: the higher the discharge frequency, the higher the probability that the neuron’s preferred feature is present. The second message, so the proposal defended in the chapter, is conveyed by precise temporal relations among the discharges of neurons and, thus, by a relational temporal code.

Synchrony as a Tag of Relatedness

Although some of the Gestalt grouping occurs preattentively, it is commonly held that attentional mechanisms also play an important role in relation-defining grouping operations, both at low levels, where scene segmentation is accomplished, and at higher levels, where objects are thought to be represented (Treisman 1999). However, the mechanisms underlying these functions are

still poorly understood. One proposal states that responses are selected for further joint processing by joint increases in saliency, i.e., in discharge rate (Cook and Maunsell 2002). Neurons recruited into the same assembly would be distinguished from all others by their higher discharge frequency. This interpretation has been challenged by the argument that response amplitude may be an ambiguous signature of relatedness for two reasons: (a) because it depends on stimulus variables, such as intensity, that are inappropriate grouping cues, and (b) because it requires long readout times, which makes it difficult to segregate assemblies from one another that are simultaneously configured within the same neuronal network (Singer 1999). Therefore, it has been proposed that neurons should be able to signal independently of their actual discharge rate to which of the other simultaneously active neurons their responses are related. This latter code should ensure that responses tagged as related are processed jointly at subsequent stages, i.e., are routed together into the appropriate channels and/or are unambiguously recognizable as originating from cells of the same assembly.

Following the discovery that neurons in the primary visual cortex can synchronize their spike discharges in a context-dependent way and with a precision in the millisecond range (Gray and Singer 1989), it was proposed that the synchronization of responses could serve as the required tag for relatedness (Gray et al. 1989) and thus, for binding them together. Precise temporal synchronization of spike discharges is as effective as joint rate increases for selectively raising the saliency of neuronal responses (Biederlack et al. 2006). The reason is that synchronized input to target neurons has a stronger impact than temporally uncoordinated input. For example, simultaneously arriving excitatory postsynaptic potentials (EPSPs) summate more effectively than temporally dispersed EPSPs; active dendritic conductances amplify fast-rising depolarizations of large amplitude (Ariav et al. 2003); the frequency adaptation of synaptic release and postsynaptic receptors attenuates the effects of rate increases (Markram and Tsodyks 1996); and the dependence of firing threshold on the rising slope of depolarizations favors responses to synchronized inputs (Azouz and Gray 2003). Because synchronization capitalizes on spatial rather than temporal summation, it modulates the efficiency of individual synaptic events and therefore operates with a temporal resolution in the millisecond range. Thus, relations can be defined within narrow time windows (<10 ms), and different relations can be encoded in rapid succession.

The notion that precise temporal synchrony serves as a tag of relatedness agrees well with the temporal sensitivity of correlation-based learning mechanisms. One mechanism, known as spike timing-dependent plasticity (STDP), classifies discharges as related (unrelated) and causes synapses to strengthen (or weaken) as a function of the precise temporal relations among pre- and postsynaptic activity patterns. This mechanism operates with a temporal resolution in the range of milliseconds; it also operates with millisecond precision (Markram et al. 1997; Zhang et al. 1998; Wespstat et al. 2004). Thus, there is a

perfect match between the signatures of relatedness used in signal processing and Hebbian learning. This cannot be otherwise because both processes must rely on the same relation-defining code to avoid learning false conjunctions.

Oscillations as a Timing Mechanism

Precise synchronization of discharges is often associated with an oscillatory patterning of the neuronal responses (Gray and Singer 1989). Because individual cells tend to skip cycles, these oscillations are rarely detectable in the spike trains of single cells, but they are readily seen in data representing the responses of large populations of neurons, as in multiunit recordings or recordings of local field potentials (LFPs). The periodic patterning of these responses is the result of oscillations generated within the various pools of inhibitory interneurons. These interneurons are coupled through chemical and electrical synapses and are capable of sustaining oscillatory activity patterns (Kopell et al. 2000; Whittington et al. 2001; Cardin et al. 2009; Sohal et al. 2009). Oscillatory inhibitory inputs to pyramidal cells veto the latter's discharges during the inhibitory troughs and favor discharges at the depolarizing peaks, thus causing synchrony in firing. Surprisingly, these locally synchronized oscillatory responses can become phase-locked, with zero delay over large distances, despite considerable conduction delays in the reciprocal excitatory corticocortical connections that mediate long-range synchrony (Engel, König, Kreiter et al. 1991).

Several mechanisms have been proposed that are capable of establishing zero-phase lag synchronization despite conduction delays. One mechanism relies on the fact that interneurons discharge with spike doublets when the networks engage in beta and gamma oscillations (Kopell et al. 2000). Another mechanism exploits the special topology of coupling connections and the non-linear properties of networks of coupled oscillators (Vicente et al. 2008). The precision with which spike timing can be adjusted increases with oscillation frequency (Volgushev et al. 1998) and, often, one observes a relation between oscillation frequency and the distance over which synchronization is maintained. Synchronization among remote groups of neurons, or among large assemblies of neurons, tends to occur at oscillations in the theta or beta frequency range, whereas the highly precise synchronization of local clusters of cells is carried out by gamma oscillations. Often, oscillations in different frequency bands coexist and exhibit complex phase relations (Roopun, Kramer et al. 2008). This concatenation of rhythms offers the attractive option of establishing graded correlations between neuronal assemblies of different size, thereby encoding nested relations. Such encoding is required for the representation of both composite perceptual objects and composite movement trajectories. This hypothesis, however, awaits further experimental testing.

The Duration of Synchronized Events

Early studies were based mostly on conventional cross-correlation analysis of cell discharges and/or LFPs. This method reliably detects synchronous firing if it is sustained over prolonged periods, but it fails if synchronous events are concentrated within narrow temporal windows. Novel measures have thus been developed to allow brief events of coincident firing to be assessed. One of these methods—the unitary event analysis—uses statistical methods to identify single, nonaccidental incidences of coincident firing (Pipa et al. 2007, 2008); the other evaluates consistent phase relations between the discharges of individual neurons and LFP oscillations (spike-field coherence; Fries et al. 2002). Applying these methods to data obtained from the visual cortex of monkeys exploring natural scenes has revealed that episodes of excess synchronized firing occur shortly after the onset of visual fixation and are restricted to epochs as short as a few tens of milliseconds (Maldonado et al. 2008). This finding agrees with the evidence that scene segmentation and object identification can be accomplished within less than 200 ms, leaving only 10–20 ms per processing stage to accomplish the required grouping operations (Thorpe et al. 1996; VanRullen and Thorpe 2001). Because information encoded in variations of discharge rates is limited (since cells can generate only a few spikes within such short time windows), a substantial amount of information is likely encoded in the precise timing relations between individual discharges of distributed neurons.

Functions Attributed to Synchronization

Binding

Evidence from studies of the visual system suggests that response synchronization may be used throughout all processing stages, from the retina to the highest cortical areas. Its purpose is to establish relations among distributed responses (i.e., to bias grouping of responses for subsequent joint processing) and to tag responses of assembly members as related (Kreiter and Singer 1996; Neuenschwander and Singer 1996; Castelo-Branco et al. 1998; Castelo-Branco et al. 2000). In all cases, synchronization probability reflects some of the Gestalt criteria that are used for scene segmentation and perceptual grouping. In the retina, ganglion cell responses synchronize with millisecond precision if evoked by continuous contours (Neuenschwander and Singer 1996). This synchronization is associated with high frequency oscillations (up to 90 Hz) and is based on horizontal interactions within the network of coupled amacrine cells. In the visual cortex, synchrony is often associated, especially when it is observed over larger distances, with an oscillatory pattern of spike discharges in the beta and gamma frequency range (30–60 Hz).

Synchronization correlates well with elementary Gestalt rules such as continuity, colinearity, and common fate (Engel, König, and Singer 1991; Engel et al. 2001; Castelo-Branco et al. 2000; Samonds et al. 2006). The substrate for this context-dependent synchronization is made up of tangential intracortical connections that preferentially link columns encoding features, which, according to common Gestalt criteria, tend to be grouped (Löwel and Singer 1992). In the inferior temporal cortex of the primate brain (the likely site for the generation of representations of visual objects), synchronization probability appears to reflect the formation of object-representing assemblies (Tsunoda et al. 2001). Neurons responding to the components of faces (e.g., eyes, nose, mouth) synchronized their responses when the arrangement of these components was such that the animals signaled having recognized a face; however, they did not synchronize when the components were scrambled or presented in a way the animal deemed incompatible with the appearance of a normal face. Interestingly, the distinction between face and nonface could not be derived from changes in the discharge rate. These findings are compatible with the interpretation that discharge rate signals the presence of particular features, whereas the correlations among the discharges indicate the relatedness of these features.

Attention and Stimulus Selection

Grouping operations based on elementary Gestalt rules and the binding of the stereotyped feature constellations of highly familiar objects can occur preattentively and are observable in anesthetized preparations. This automatic, attention-independent grouping is likely based on binding by convergence in fixed feedforward architectures and on the synchronizing effects of the intrareal tangential fiber systems (discussed above).

In addition to this evidence for attention-independent grouping by synchrony, more recent results clearly indicate that synchronization is highly susceptible to top-down, attention-dependent influences. They also show that it plays an important role in attention and expectancy-dependent response selection (Fries, Reynolds et al. 2001; Fries et al. 2002). An in-depth discussion of the mechanisms involved is provided by Whittington et al. (this volume) and Börgers and Kopell (2008).

Various measures have been used to assess the influence of selective attention on neuronal synchrony: correlations among spike discharges, spike-field coherence, correlations in phase locking between oscillatory field potentials, and, finally, the amplitude and phase locking of oscillatory responses as seen in MEG and EEG recordings. Because the amplitude of these latter signals depends not only on the number of active neurons but also on the degree of synchronicity of the captured activity, both phase locking and the power of oscillations can be taken as measures of synchrony. At all levels of analysis, evidence indicates that focusing attention on a particular stimulus or modality

increases the synchrony of responses in the gamma and beta frequency range in the neuronal networks that are devoted to the processing of the attended stimulus (Roelfsema et al. 1997; Schoffelen et al. 2005). Interestingly, in these and many other cases, synchronization of oscillatory activity is not necessarily associated with major changes in the discharge activity of neurons. This observation supports the notion that synchronization and rate of discharges can be adjusted independently, and that precise synchronization can be used to raise the saliency of responses independently of discharge rate (Fries, Neuenschwander et al. 2001; Fries et al. 2007). Recent results from multisite recordings in cats and monkeys take this proposal one step further and suggest that the (often anticipatory) induction of coherent oscillations across distributed cortical areas and executive structures facilitates selective routing of activity and rapid handshaking among the involved processing stages (Womelsdorf et al. 2007). However, at present, it is not known which mechanisms coordinate these preparatory phase adjustments of oscillatory activity.

Association with Consciousness

Probability increases that signals become part of consciousness when they are attended to, either because they are salient and attract attention, or because they are selected by focused attention. Moreover, contents appearing in consciousness are usually interrelated (unity of consciousness). Because synchronization in the gamma range enhances saliency (Biederlack et al. 2006), supports selection by attention (Fries, Reynolds et al. 2001), and establishes relations (Gray et al. 1989), it is a prime candidate for being a neural correlate of consciousness. Evidence from studies on binocular rivalry in cats (Fries et al. 1997; Fries et al. 2002) and from masking experiments in human subjects (Melloni et al. 2007) indicates that there is indeed a close correlation between gamma synchronization and conscious processing. In cat primary visual cortex, responses to the respective perceived stimulus differed from those to the suppressed stimulus because the former were more synchronized, not because they were more vigorous. In human subjects, conscious processing of stimuli has been associated with precise phase locking of gamma oscillations across widely distributed cortical areas, whereas unconsciously processed stimuli evoked only local gamma oscillations (Melloni et al. 2007).

Abnormal Synchrony and Mental Disorders

Several clinical conditions (e.g., schizophrenia, autism, and Alzheimer's disease) are associated with cognitive impairments that suggest disturbed coordination of distributed brain processes. This concept has received support from recent EEG and MEG studies (Uhlhaas and Singer 2006; Vierling-Claassen

et al. 2008), as well as from *in vitro* pharmacology (Roopun, Cunningham et al. 2008). When challenged with cognitive tasks that require feature binding (Uhlhaas, Linden et al. 2006; Uhlhaas and Singer 2010) or storage of visual information in working memory (Haenschel et al. 2009), schizophrenic patients exhibited several deficits: reduced synchronization of early evoked responses, reduced power of evoked and induced gamma oscillations, and a dramatic loss of the ability to synchronize gamma oscillations across distant cortical areas. These abnormalities were also seen (albeit in a less pronounced manner) in unmedicated, first admission patients. Impaired synchronization of oscillatory activity in the beta and gamma frequency range was also found in patients with autism (Wilson et al. 2007) and Alzheimer's disease (Koenig et al. 2005; Stam et al. 2007).

These findings suggest that some of the dissociative symptoms characteristic of these disorders could result from abnormalities in the precise temporal coordination and binding of distributed cortical processes. This insight offers researchers the option to use synchronization of oscillatory activity as an endophenotype for further investigations into the pathophysiology of these severe and hitherto incurable brain diseases.

The Stance of a Sceptic

No contemporary neurobiologist would deny the existence and necessity of "dynamic coordination" of distributed processes in the brain. Such coordination requires fast, temporary, and mostly reversible modification of neuronal interactions within the fixed or only slowly changing anatomical architecture. This implies that effective connectivity must be modifiable on the fly in a context-, attention-, or task-dependent way. This can be achieved by a host of well-established mechanisms, raising the question as to why oscillations, synchrony, and temporal codes should matter at all. The gain of connections is effectively modified by changes of firing rate, since this increases or decreases the saliency of responses. This suffices to select signals for further processing, especially if saliency maps enhance contrast between selected and non-selected responses. Experiments on attention-dependent rate enhancements suggest such a mechanism. To define relations and support selective grouping, it is sufficient to increase jointly the rate of the responses that are to be associated with each other. Since coding is sparse, activity is low, and since there is topological (spatial) organization, the risk of confusion and grouping of unrelated but simultaneously enhanced responses is small. At least it is not greater than misinterpreting spurious correlations among simultaneously active neurons as meaningful. The flow and routing of activity can further be controlled very effectively by short-lasting and reversible changes in synaptic gain, by active dendritic conductances that introduce nonlinearities in the summation of synaptic inputs, by shunting inhibition which can switch off entire

dendritic segments, by modulatory inputs that alter time and length constants, and by top-down influences which, in principle, can gate all these processes. The problem of feeding learning mechanisms with activity that is sufficiently structured in time can probably be solved by statistics, assuming simply that the conditions required for gain increases occur more often, if coupled neurons are more active. There have been negative findings. Some labs have failed to find oscillatory activity and response synchronization in brain areas where others had observed them, or found no relation between the occurrence of synchrony and perceptual functions that should involve dynamic grouping of responses. Some of the conflicts between controversial findings could be resolved. Thus, it is well established that gamma oscillations and the associated synchronization of spikes are extremely state and task dependent. In light anesthesia they occur only during states of activated EEG (Herculano-Houzel et al. 1999), whereas in the awake brain they are strongly dependent on attentional mechanisms (Fries, Reynolds et al. 2001). Moreover, cell discharges, even if synchronized to an oscillatory process, may fail to exhibit an oscillatory firing pattern because of irregular cycle skipping. Finally, some of the negative findings may have to do with the fact that different cortical areas accomplish different functions. Just as one would obtain negative evidence for a rate code if one searched for face-specific responses in area MT rather than IT, one would obtain negative evidence for synchrony as tag of relatedness if one searched for higher-order binding functions in V1. Still, there are negative findings that require a continuous and critical evaluation of results supporting the temporal coding hypothesis.

One could argue that we have come a long way by assessing the rate variations of individual neurons without coming across explananda that would require us to search for something additional. Why then should we look for fine-grained temporal relations between the discharges of distributed neurons, since this necessitates technically much more challenging multisite recordings? The fact that oscillations, synchrony, and fine-grained temporal relations are observed and interesting functions can be associated with such phenomena is not sufficient to assign a function to them. What is the argument behind the view that conventional mechanisms do not suffice to account for what we observe and need to understand? What evidence exists to show that all these temporal dynamics are not just an epiphenomenon of the conventional, purely rate-based processes?

Obviously, to address these questions in an ultimate fashion, causal rather than correlative evidence is required. Such evidence is equally difficult to obtain for conventional rate codes and oscillation-based temporal codes, because both are constitutive attributes of neuronal processing. Interfering with rates most often entrains changes in oscillations and synchrony. Moreover, manipulating rates interferes in an often trivial way with the cell's sole signaling mechanism. Conversely, blocking oscillations and synchrony often leads to changes in discharge rate. However, studies are available that can claim having

obtained causal evidence for either rate or temporal codes. Assuming that high frequency microstimulation acts by increasing rates rather than by rendering discharges more synchronous, the many studies that show relations between stimulation frequency and the vigor of perception or motor responses can be taken as support of rate codes (Salzman et al. 1992). Support for a coding role of synchrony comes from studies in which synchronization has been reduced or abolished without interfering with discharge rate. In frogs, an escape response was abolished when the synchronization of retinal ganglionic cells was prevented (Ishikane et al. 2005), and in locusts the discrimination of odor mixtures was impaired when synchrony was abolished among projection cells in the olfactory lobes (Stopfer et al. 1997). Finally, optogenetic methods have recently been applied to enhance synchronization in the gamma frequency range, and this led to predicted and functionally relevant changes in the performance of neuronal networks (Cardin et al. 2009; Sohal et al. 2009).

In conclusion, it appears that there is ample and equally convincing correlative evidence that the brain uses rates as well as temporal relations between spike trains in parallel to encode complementary information. However, if the brain does indeed exploit temporal relations among the firing sequences of large assemblies of neurons, new challenges must be overcome before we can decipher these temporal codes. Multisite recordings will be imperative. Moreover, analysis of relations must go beyond pairwise correlations, comprise all frequency bands, and consider all phase relations in case of oscillatory activity or relative delays between spike times of non-oscillatory activity. This leads to a dramatic expansion of search space, just as exploitation of temporal relations for information processing—if the brain applied that strategy—would dramatically increase coding space.

Conclusion

Temporal relations among distributed neuronal responses can be assessed only with multisite recordings. Because this approach has a relatively short history, we are just beginning to understand coding strategies based on the dynamic interactions among large numbers of neurons. It may turn out that precise synchronization is only one, albeit very important signature of the many potentially significant dynamical states. Precisely timed phase offsets between oscillating cell assemblies, concatenations of different rhythms, and sequences of patterns defined by specific temporal relations are likely to play an equally important role (Fries et al. 2007). To analyze these more complex patterns, and to examine whether they contain information that can be related to behavior, remains one of the great challenges for future systems neuroscience.