

Olaf Sporns

Neural Dynamics and Cognitive Synthesis

Abstract

Nervous systems facing complex environments have to balance two seemingly opposing requirements. First, there is a need quickly and reliably to extract important features from sensory inputs. This is accomplished by functionally segregated (specialized) sets of neurons, e.g. those found in different cortical areas. Second, there is a need to generate coherent perceptual and cognitive states allowing an organism to respond to objects and events, representing conjunctions of numerous individual features. This is accomplished by functional integration of the activity of specialized neurons through their dynamic interactions. These interactions produce patterns of temporal correlations or functional connectivity involving distributed neuronal populations, both within and across cortical areas. Empirical and computational studies suggesting that changes in functional connectivity may underlie specific perceptual and cognitive states involving the integration of information across specialized areas of the brain offer new insights into the linkage between neural dynamics and cognitive synthesis.

Introduction

Humans experience a world that is composed of coherent objects and events. For example, when viewing a visual scene, its constituent objects appear both unitary and differentiated with respect to each other. This is the case, despite the fact that objects are characterized by conjunctions of numerous features, often shared by other objects present within the visual field. Scene segmentation, the unambiguous assignment of object boundaries within a visual image, is known to be a hard problem in computer vision and artificial intelligence; yet, biological nervous systems seem to solve this problem exceedingly well. In a sense, scene segmentation is just one among a set of related problems encountered throughout

perception and cognition, all involving the merging of often disparate information into a single unified construct. How is perceptual or cognitive synthesis achieved within the networks of the brain?

Modern neuroscience has made tremendous progress in identifying neural correlates for many aspects of visual perception. Much of the neural activity underlying visual perception takes place within the visual cortex. As numerous studies have shown, the visual cortex (as well as every other subdivision of the cortex) is organized as a mosaic of functionally segregated areas (Zeki, 1993; Mountcastle, 1998), each specialized to deal with a particular aspect of the visual scene. For example, there are separate visual areas dealing with shape, movement, or color. This poses a problem. Neurally, numerous areas of the brain are specialized to deal with only a small subset of visual features. However, perceptually, visual objects are both coherent (i.e. defined as conjunctive sets of numerous features) as well as differentiable components within a visual scene. This “binding problem” (von der Malsburg, 1995; Treisman, 1996) extends over several problem domains, including the binding of features to form objects, the binding of parts of a spatially extended object and its differentiation from a background, as well as the binding of different views of an object across time.

This brief essay examines the binding problem in the context of the more general issue of how information distributed throughout the brain can be integrated to yield coherent mental or cognitive states (cognitive synthesis). Empirical as well as computational studies suggest that neural dynamics, specifically the coordinated firing of large sets of neuronal populations, can provide a possible neural basis for perceptual and cognitive synthesis. This emerging view has important consequences for our theoretical understanding of the relationship between neural and cognitive states.

Possible Solutions to the Binding Problem

Neurons communicate through trains of discrete impulses (action potentials, or spikes), each lasting less than a millisecond. There are two main “coding strategies” which allow neuronal activity to carry information (Perkel and Bullock, 1968; Theunissen and Miller, 1995). First, a neuron can signal by turning on or off. For example, a neuron can convey information by changing its mean firing rate, evaluated as the number of spikes per unit time (often taken to be on the order of hundreds of milliseconds). This “*rate coding*” is often contrasted with the second main strategy, called “*temporal coding*”. Here, a neuron can convey information by the timing of its spikes, often considered relative to spikes emitted by other neurons. In one version of temporal coding, two (or more) neurons can carry information by virtue of whether or not their spikes are temporally correlated

or synchronous (a special case of correlation with zero phase lag). Whether neurons exhibit rate or temporal coding has profound consequences for how they act as a population (deCharms, 1998). Rate coding implies that neurons signal primarily through their firing level and, when forming populations, their signals can be considered to carry independent pieces of information. In other words, the population activity can be viewed as a superposition of essentially *independent* elements. In contrast, if information is carried in the relative timing of spikes within a population, the population as a whole can no longer be considered as consisting of independent elements; rather, its *coordinated* pattern of activity must be taken into account.

Different possible solutions of the binding problem make different use of rate and temporal coding. One possibility involves the activity of neurons that integrate frequently co-occurring sets of features by combining convergent inputs from other parts of the brain. Such neurons form specialized “detectors” for complicated objects (e.g. the infamous “grandmother cell”; Barlow, 1972). They signal the presence of a particular object by increasing their firing rate (i.e. rate coding). However, such a mechanism is unlikely to be the predominant mechanism for integration. Highly specific feature detectors have been found only for a handful of complex objects, most of them neurons specialized for the recognition of faces (Perrett et al., 1987). To date, specialized cortical areas whose activity represents entire perceptual or mental states have not been identified. It has also been noted (von der Malsburg, 1995) that the vast number of possible objects and events far exceeds the number of available neuronal groups (or even single neurons). In addition, convergence along specialized connections does not allow for dynamic (“on-the-fly”) conjunctions in response to unanticipated or novel stimuli. These empirical and theoretical considerations would suggest that other mechanisms must be at work, in order to account for the capacity of the brain rapidly to integrate information and yield perceptual and cognitive coherence.

A closer look at the anatomical and functional organization of the cerebral cortex provides important clues for a more general mechanism for neural integration. Within a cortical area, distant neuronal groups are linked by long- range intrinsic connections (also called tangential or horizontal connections), forming a dense patchy network (Gilbert and Wiesel, 1989). Between cortical areas, detailed anatomical studies have shown the existence of a meshwork of anatomical pathways (Felleman and van Essen, 1991; Scanned et al., 1995; 1999), comprising bundles of axons linking vast numbers of neurons. Almost all of these pathways are reciprocal. These reciprocal pathways provide the structural substrate for dynamic interactions between cell populations over large distances, within and between cortical areas (reentry; Edelman, 1987).

In the late eighties, several reports were published indicating that orientation-selective neurons located within the cat primary visual cortex showed oscillatory

responses around 40 Hz when stimulated (Eckhom et al., 1988; Gray and Singer, 1989). When recorded from simultaneously, widely separated neurons with nonoverlapping receptive fields synchronized their responses when a coherently moving spatially extended contour was presented, while no synchrony was detected in response to an incoherent stimulus (Gray et al., 1989). These results represented the first direct neurophysiological evidence of synchronization among neurons in a manner suggestive of a role for synchrony in perception. They sparked a vigorous debate on the role of temporal correlations in solving the binding problem that has continued to the present day. Computational models have been an important component in this debate and have made significant contributions in the search for possible neural mechanisms underlying perceptual synthesis (for reviews of such models see von der Malsburg, 1995; 1999; Gawne, 1999).

Modeling Binding and Perceptual Synthesis

Almost as soon as experimental evidence on synchronicity in the visual cortex became available, detailed computer simulations were conducted (Sporns et al., 1989) showing that dynamic reentrant interactions between functionally specialized groups of neurons can lead to patterns of short-term correlations. Individual groups of neurons, composed of sparsely interconnected sets of excitatory and inhibitory neurons, readily engaged in oscillatory activity when exposed to a visual input. The oscillation was clearly evident at the population level, while the firing of individual modeled neurons rarely followed the oscillatory pattern with great accuracy. Thus, neuronal groups acted as population oscillators, exerting effects on other groups due to local cooperative interactions. These computational studies gave rise to the proposal (Sporns et al., 1991) that basic visual perceptual processes such as perceptual grouping (according to the Gestalt laws; see e.g. Kanisza, 1979) and figure-ground-segregation may have a neural basis in patterns of synchronicity. These dynamic patterns are constrained by the anatomical patterns of intra-areal connections linking groups or columns of neurons within the primary visual cortex. In agreement with experimental results (Gray et al., 1989), correlations in the model emerged quickly after stimulus onset and linked groups of neurons for periods of up to several hundred milliseconds ("short-term correlations"). Correlations were found to benefit from incorporating a novel mechanism of short-term plasticity within the intra-areal connections linking groups of neurons over long distances. According to this mechanism, connections had larger effects on postsynaptic neurons that were already somewhat depolarized, rendering these inputs voltage-dependent. Thus, groups that tended to fire synchronously were connected with greater efficacy,

which in turn tended to boost their synchronization. Evidence for such voltage-dependence of some types of cortico-cortical connections was found in primary visual cortex (Hirsch and Gilbert, 1991).

These computer simulations were subsequently extended, from considering dynamic interactions within just one visual area to considering both intra-areal and inter-areal interactions across multiple visual areas. A computer simulation of nine segregated visual cortical areas (Tononi et al., 1992), modeled on the anatomy and physiology of the macaque monkey visual cortex, demonstrated that short-term correlations can be generated over an extended network of pathways, despite conduction delays, and in a rapid and stimulus-dependent manner. Using a visual choice task as a simple instantiation of the binding problem, it was shown that short-term correlations can lead to coherent perceptual performance and behavior in the absence of a master area. The model was able to select a visual stimulus, defined as a conjunction of three independent features, without using or generating specialized feature detectors; instead binding occurred “on the fly”, through a network of dynamic interactions involving almost all areas of the model. When these interactions were disrupted (e.g. by cutting anatomical pathways, or by eliminating their reciprocity) the behavioral performance was significantly impaired. This simulation study demonstrated that a distributed pattern of synchronized activity can have a distinct effect on a behavioral output.

Large-scale computer simulations of the mammalian thalamocortical system have provided important insights into the structural bases for the generation of different kinds of rhythmic and synchronized activity (Lumer et al., 1997a). In a model comprising a total of 65,000 spiking neurons organized into multiple cortical and thalamic areas, it was found that the thalamocortical system consists of several classes of ‘macro-circuits’, e.g. the network of horizontal patchy connections within each cortical layer, interlaminar cortical loops, interareal reciprocal pathways, and thalamocortical loops. High-frequency rhythmic activity within extended populations of neurons depended on the dynamics of these (often polysynaptic) loops and on the lateral spread of oscillations due to intra-areal and inter-areal connectivity. Several anatomical and physiological manipulations were performed that would be difficult, if not impossible, to perform in a real animal. First, selectively eliminating backward connectivity (i.e. connection pathways originating from hierarchically higher areas of the cortex and terminating on lower areas) resulted in systematic time lags across areas, disrupting synchronous activity with zero time lag, a hallmark of actual cortical dynamics (e.g. Engel et al., 1991). Second, when spike timing was selectively perturbed by “jittering” (i.e. shifting in time) of individual action potentials, significant effects on patterns of synchrony as well as on firing rates were observed throughout the model (Lumer et al., 1997b). This computational result points to

an important interaction between firing rates and patterns of synchrony or correlations. Evidently, for neurons to be correlated they must also be active, but levels of activity are in turn affected by patterns of correlations within a network. This is due to the sensitivity of cortical neurons, individually and in populations, to the timing of their inputs, rendering them cellular “coincidence detectors” (Konig et al., 1996). This important property of cortical neurons provides a conceptual link between rate coding and temporal coding.

Empirical Evidence for Synchronized Neural Activity

Over the past decade, a large number of neurophysiological experiments have shown the existence of stimulus-dependent temporally correlated activity within areas of the visual (Gray et al., 1989), auditory (deCharms and Merzenich, 1996), frontal (Vaadia et al., 1995) and motor cortex (Murthy and Fetz, 1996; Donoghue et al. 1998). Robust synchronized activity has been found in a variety of animal species, most notably in the cortex of awake primates (e.g. Kreiter and Singer, 1996; Gray, 1999). Thus, it appears that the existence of temporally correlated neuronal activity is a widespread phenomenon occurring throughout the cerebral cortex of most mammalian species.

What is the structural substrate of the extended patterns of temporal synchrony observed both within and between cortical areas? In principle, synchronized activity of two cell groups could be due to common input, or to dynamic coupling mediated by direct (usually reciprocal) connections. In the cerebral cortex, synchronization depends on the integrity of reciprocal pathways, as was demonstrated by the abolishment of cross-hemispheric correlations in visual cortex after transection of the corpus callosum (Engel et al., 1991). Further evidence for a relationship between patterns of cortico-cortical connections and patterns of synchrony was provided by experiments on strabismic cats. A misalignment of the optical axes of the two eyes (squint) resulted in changes in the anatomical pattern of intra-areal connections in primary visual cortex (Lowel and Singer, 1992), such that only groups of cells with same ocular dominance were linked. When recording neuronal population activity, temporal correlations were readily observed between cells of same ocular dominance, while correlations between groups of cells with different ocular dominance were reported absent (Konig et al., 1993). Taken together, these studies indicated that temporal correlations are predominantly due to reentrant connections both within and between cortical areas, and that changes in anatomical connections result in corresponding changes in temporal correlations.

It has always been critical to show that synchronized activity is related to changes in behavioral or perceptual states. Recently, some evidence points to

a role of synchronized neural activity in perception during binocular rivalry in awake strabismic cats (Fries et al., 1997). Binocular rivalry refers to the alternating perception of two different stimuli presented separately to the two eyes. It was found that only neuronal responses to a stimulus that was actually perceived were strongly synchronized. Responses to a stimulus that was not perceived were not synchronized, although their mean firing rate remained unchanged. Outside of primary visual cortex, synchronized neuronal activity has also been observed between sensory and motor areas (Bressler et al., 1993; Roelfsema et al., 1997). Interestingly, sensorimotor synchronization occurred with zero phase lag despite the existence of significant transmission delays between these widely separated regions of the cerebral cortex. The presence of sensorimotor synchrony was related to the execution of behavioral tasks that require sensorimotor integration. Electrophysiological studies in awake monkeys have shown significant modulations of temporal correlations between neurons in motor (Riehle et al., 1997) and frontal cortex (Vaadia et al., 1995) occurring during the performance of behavioral tasks. Episodes of increased synchrony between cortical neurons were found to coincide with behaviorally significant time intervals of several hundred milliseconds, often in the absence of corresponding changes in firing levels.

The hypothesis that temporal correlations may be involved in solving the binding problem is receiving additional support by a recent set of psychophysical studies. These studies showed that temporal correlations present in visual stimuli can produce distinct perceptual effects. In particular, temporally co-varying changes within spatially coherent features of a stimulus array produced percepts of high saliency (Blake and Yang, 1997). Perceptual grouping, one of the simplest examples of binding in vision, was aided by correlated changes in the contrast of local visual features (Alais et al., 1998). This relationship between (externally generated) temporal modulations present within the stimulus and perceptual coherence generally supports the idea that (internally generated) temporal correlations in neural activity may underlie different perceptual states.

From Synchronized Activity to Functional Connectivity

The experimental and computational studies briefly reviewed here are consistent with the view that the cortex is composed of functionally specialized local populations of neurons that are interacting dynamically along anatomical loops and pathways, both within and between segregated areas. Fundamentally similar dynamic processes operate at multiple levels of scale and across widely separated regions of the brain. The large-scale patterns of temporal correlations generated by the dynamics of neuronal interactions across the brain are often

referred to as functional connectivity (Friston, 1994). What is the evidence that changes in functional connectivity are associated with higher cognitive processes, and how can functional connectivity be experimentally measured and computationally evaluated?

In the past, neuroimaging studies have typically attempted to identify which brain regions show increased or decreased activity in the course of particular perceptual or cognitive tasks. Increasingly, perceptual and cognitive states of the human brain are investigated in the context of changes in functional connectivity, i.e. in terms of which brain regions are co-active or correlated. Numerous examples of changes in functional connectivity exist, most of them obtained from awake and conscious human subjects. Performance of a visuomotor task requiring the integration of visual and motor signals was found to be accompanied by increases in interregional coherence between visual and motor cortex (Classen et al., 1998). Induced coherent (synchronized) activity in several brain regions was observed while a previously presented visual stimulus was held in short-term memory (Tallon-Baudry et al., 1997). Changes in functional connectivity involving lateral prefrontal cortex were reported to reflect whether experimental subjects were aware of an association between different sensory stimuli or not (McIntosh et al., 1999). Finally, changes in synchronization between distinct brain areas were found to reflect changes in perceptual dominance during binocular rivalry (Srinivasan et al., 1999).

By its nature, functional connectivity involves statistical relationships between potentially large numbers of segregated elements. There is a growing need to describe such patterns efficiently and in ways that are meaningful in the context of brain function. Several computational and statistical techniques addressing this need have been introduced, for example techniques using multivariate statistics, including multidimensional scaling, path analysis and cluster analysis. One way of approaching the problem of analyzing patterns of functional connectivity is through a consideration of functional segregation and integration in terms of information theory.

Complexity

Functional segregation and integration present seemingly opposing requirements to the cortical architecture. On the one hand, functional segregation requires that the firing of specialized groups of neurons be mutually independent, i.e. such groups should represent essentially orthogonal stimulus dimensions. On the other hand, functional integration requires that their joint activities be highly coherent, i.e. sets of groups should represent conjunctive sets of features, such as higher order objects or concepts. Using two basic and general concepts

from information theory, entropy and mutual information, we have developed a theoretical measure that captures the interplay of functional segregation and integration within a given system (Tononi et al., 1994; 1998). The system is composed of a number of discrete but interacting elements, such as neurons or neuronal groups. The behavior of these elements over time is described by a set of random variables. Entropy and mutual information capture the degree to which subsets of the system exhibit deviations from statistical independence. We consider subsets of all sizes within the system and define a measure that captures the degree to which a system incorporates both local and global structure, i.e. deviations from statistical independence at small or large scales. We called this measure complexity, a term that reflects the recognized tendency of complex systems to exhibit a high degree of spatiotemporal order (contain large amounts of “interesting structure”), while keeping a balance between the extremes of randomness and total regularity (e.g. Huberman and Hogg, 1986). Such systems are not only found in the brain, but throughout biology, and the dichotomy of local cooperativity and global interactivity has been recognized even in economic (Simon and Ando, 1961) and socio-cultural (Flannery, 1972) contexts.

Complexity is a descriptor of dynamics, or, more precisely, of the overall pattern of statistical deviations (e.g. temporal correlations) generated by a system’s activity. Clearly, one of its key determinants must be found in the underlying structure of the network. In the case of the brain, this structure is equivalent to its neuroanatomy. Are there specific structural motifs that are uniquely associated with high levels of complexity? We investigated this question (Sporns et al., 2000a; 2000b) by generating large numbers of simple connectivities (graphs), running them as dynamical systems, and computing their complexity. An iterative evolutionary strategy was applied to select graphs that produced dynamics with high complexity; this strategy allowed to efficiently search the vast space of possible connectivities. After selecting for high complexity, anatomical patterns emerged which consistently comprised densely connected clusters of elements that were sparsely interconnected between each other. The patterns were analyzed using structural measures derived from graph theory, and the results were compared to those obtained from an analysis of real cortical connection matrices, describing the interconnectivity of the macaque monkey visual cortex (Felleman and Van Essen, 1991) or the cat cortex (Scanned et al., 1995; 1999). It was found that real cortical connection matrices share most structural features of graphs selected for high complexity. In fact, when implemented as dynamical systems, the resulting pattern of functional connectivity was found to be highly complex. These structural and functional analyses indicated that the actual organization of cortico-cortical connection pathways is highly conducive to a near-optimal balance between functional segregation and functional integration.

Tononi and Edelman (1998) have proposed a testable hypothesis concerning the neural substrate of conscious experience, the dynamic core hypothesis. The hypothesis addresses two key properties of conscious experience, the fact that it is integrated (i.e. the conscious scene is unified and coherent) as well as differentiated (allowing an individual to experience any of a huge number of different conscious states). Rapid and effective integration of neural activity in distributed populations forms the basis for the coherent quality of consciousness, while the availability of a large repertoire of neural states accounts for the differentiated quality of consciousness. Integrated neural activity is thought to occur in a dynamic core capable of traversing a huge space of possible dynamic states. Information-theoretic measures of functional clustering and complexity provide empirical tools to identify the neural bases of consciousness in actual recordings of brain activity. The dynamic core hypothesis differs from other proposals in that it attempts to characterize a *neural process* underlying conscious experience, avoiding the category error of ascribing correlates of consciousness to particular neurons or brain regions. This neural process is a manifestation of neural dynamics, the coordinated activity of large populations of neurons distributed throughout the brain. In a sense, this is the most profound and ultimate link between neural dynamics and cognitive synthesis.

Dynamical Systems and Cognition

In this brief essay, it was argued that cognitive synthesis depends upon the interplay between functional segregation and integration in the brain. Essentially, its underpinnings are dynamic phenomena generated by the activity of large populations of neurons organized in distinct anatomical patterns. The last ten years of debate over the binding problem (a special case of cognitive synthesis) have led to a reappraisal and resurgence of approaches to brain function based on the dynamics of large populations of neurons. More and more evidence suggests that the rapid integration of information in the brain does not occur in a particular location, but constitutes a unified neural process, that can only be characterized in terms of neural dynamics.

This realization constitutes a step away from the “classical” notions of neural computation as rule-driven transformation of discrete representations, a view that seems oddly inappropriate in light of the rich dynamics exhibited by even the simplest of nervous systems. The key to understanding perception and cognition may lie in the realization that brains are complex dynamical systems, delicately balanced between the opposing forces of functional segregation and functional integration.

Acknowledgement. The work reviewed in this paper was carried out as part of the theoretical neurobiology program at The Neurosciences Institute, which is supported by Neurosciences Research Foundation. The Foundation receives major support for this work from Novartis and the W.M. Keck Foundation.

References

- Alais, D., Blake, R., and Lee, S.-H. (1998) Visual features that vary together over time group together over space. *Nature Neurosci.* 1:160^a164.
- Barlow, H. B. (1972) Single units and cognition: a neurone doctrine for perceptual psychology. *Perception* 1:371-394.
- Bressler, S. L., Coppola, R. and Nakamura, R. (1993) Episodic multiregional cortical coherence at multiple frequencies during visual task performance. *Nature* 366:153-156.
- Classen, J., Gerloff, C., Honda, M., and Hallett, M. (1998) Integrative visuomotor behavior is associated with interregionally coherent oscillations in the human brain. *J. Neurophysiol.* 79:1567-1573.
- deCharms, R. C., and Merzenich, M.M. (1996) Primary cortical representation of sounds by the coordination of action-potential timing. *Nature* 381:610-613.
- deCharms, R. C. (1998) Information coding in the cortex by independent or coordinated populations. *Proc. Natl. Acad. Sci. USA* 95:15166-15168.
- Donoghue, J. P., Sanes, J.N., Hatsopoulos, N.G. and Gaal, G. (1998) Neural discharge and local field potential oscillations in primate motor cortex during voluntary movements. *J. Neurophysiol.* 79:159-173.
- Eckhom, R., Bauer, R., Jordan, W., Brosch, M., Kruse, W., Munk, M., and Reitboeck, H. J. (1988) Coherent oscillations: a mechanism for feature linking in the visual cortex? *Biol. Cybern.* 60:121-130.
- Edelman, G. M. (1987) *Neural Darwinism*. Basic Books, New York.
- Engel, A. K., König, P., Kreiter, A. K. and Singer, W. (1991) Interhemispheric synchronization of oscillatory neuronal responses in cat visual cortex. *Science* 252:1177- -1179.
- Felleman, D. J., and D. C. Van Essen (1991) Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex* 1:1-47.
- Flannery, K. V. (1972) Cultural evolution of civilizations. *Ann. Rev. Ecology Systematics* 3: 399-426.
- Fries, P, Roelfsema, P. R., Engel, A. K., König, P, and Singer, W. (1997) Synchronization of oscillatory responses in visual cortex correlates with perception in interocular rivalry. *Proc. Natl. Acad. Sci. USA* 94:12699-12704.
- Friston, K. J. (1994) Functional and effective connectivity in neuroimaging: A synthesis. *Human Brain Mapping* 2:56-78.
- Gawne, T. J. (1999) Temporal coding as a means of information transfer in the primate visual system. *Crit. Rev. Neurosci.* 13:83-101.
- Gilbert, C. D., and Wiesel, T.N. (1989) Columnar specificity of intrinsic horizontal and cortico-cortical connections in cat visual cortex. *J. Neurosci.* 9:2432-2442.

- Gray, C. M., Kónig, P., Engel, A. K. and Singer, W. (1989) Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Nature* 338:334-337.
- Hirsch, J., and Gilbert, C. (1991) Synaptic physiology of horizontal connections in the cat visual cortex. *J. Neurosci.* 11:1800-1809.
- Huberman, B. A., Hogg, T. (1986) Complexity and adaptation. *Physica* 22D:376-384.
- Kanizsa, G. (1979) *The Organization of Vision*. Praeger, New York.
- Kónig, P, Engel, A. K., Lówel, S., and Singer, W. (1993) Squint affects synchronization of oscillatory responses in cat visual cortex. *Eur. J. Neurosci.* 5:501-508.
- Kónig, P, Engel, A. K., and Singer W. (1996) Integrator or coincidence detector? The role of the cortical neuron revisited. *Trends Neurosci.* 19:130-137.
- Kreiter, A. K., and Singer, W. (1996) Stimulus-dependent synchronization of neuronal responses in the visual cortex of the awake macaque monkey. *J. Neurosci.* 16:2381- - 2396.
- Lówel, S., and Singer, W. (1992) Selection of intrinsic horizontal connections in the visual cortex by correlated neuronal activity. *Science* 255:209-212.
- Lumer, E. D., Edelman, G. M., and Tononi, G. (1997a) Neural dynamics in a model of the thalamocortical system. I. Layers, loops and the emergence of fast synchronous rhythms. *Cerebral Cortex* 7:207-227.
- Lumer, E. D., Edelman, G. M., and Tononi, G. (1997b) Neural dynamics in a model of the thalamocortical system. II. The role of neural synchrony tested through perturbations of spike timing. *Cerebral Cortex* 7:228-236
- McIntosh, A. R., Rajah, M. N., and Lobaugh, N. J. (1999) Interactions of prefrontal cortex in relation to awareness in sensory learning. *Science* 284:1531-1533.
- Mountcastle, V. B. (1998) *Perceptual Neuroscience: The Cerebral Cortex*. Harvard University Press.
- Murthy, V. N. and Fetz, E. E. (1996) Synchronization of neurons during local field potential oscillations in sensorimotor cortex of awake monkeys. *J. Neurophysiol.* 76:3968-3982.
- Perkel, D. H., and Bullock, T. H. (1968) Neural coding. *Neurosci. Res. Progr. Bull.* 63:221-348.
- Perrett, D. I., Mistlin, A. J., and Chitty, A. J. (1987) Visual neurones responsive to faces. *Trends Neurosci.* 10:358-364.
- Roelfsema, P. R., Engel, A. K., Kónig, P, and Singer, W. (1997) Visuomotor integration is associated with zero time-lag synchronization among cortical areas. *Nature* 385:157-161.
- Scannell, J. W., Blakemore, C., and Young, M. P. (1995) Analysis of connectivity in the cat cerebral cortex. *J. Neurosci.* 15:1463-1483.
- Scannell, J. W., Bums, G. A. P. C., Hilgetag, C. C., O'Neil, M. A., and Young, M. P (1999) The connectional organization of the cortico-thalamic system of the cat. *Cerebral Cortex* 9:277-299.
- Simon, H. A., Ando, A. (1961) Aggregation of variables in dynamic systems. *Econometrica* 29:111-138.
- Singer, W., and Gray, C. M. (1995) Visual feature integration and the temporal correlation hypothesis. *Annu. Rev. Neurosci.* 18:555-586.

- Spoms, O., Gaily, J.A., Reeke, Jr., G. N., and Edelman, G. M. (1989) Reentrant signaling among simulated neuronal groups leads to coherency in their oscillatory activity. *Proc. Natl. Acad. Sci. USA* 86:7265-7269.
- Spoms, O., Tononi, G., and Edelman, G. M. (1991) Modeling perceptual grouping and figure-ground segregation by means of active reentrant circuits. *Proc. Natl. Acad. Sci. USA* 88:129-133.
- Spoms, O., Tononi, G., and Edelman, G. M. (2000a) Theoretical neuroanatomy: Relating anatomical and functional connectivity in graphs and cortical connection matrices. *Cerebral Cortex* 10:127-141.
- Spoms, O., Tononi, G., and Edelman, G. M. (2000b) Connectivity and complexity: the relationship between neuroanatomy and brain dynamics. *Neural Network* 13: 909- 922.
- Srinivasan, R., Russell, D. P., Edelman, G. M., and Tononi, G. (1999) Increased synchronization of neuromagnetic responses during conscious perception. *J. Neurosci.* 19:5435-5448.
- Tallon-Baudry, C., Bertrand, O., Peronnet, F., and Pemier, J. (1998) Induced g-band activity during the delay of a visual short-term memory task in humans. *J. Neurosci.* 18:4244-4254.
- Theunissen, F. and Miller, J. P. (1995) Temporal encoding in nervous systems: a rigorous definition. *J. Comput. Neurosci.* 2:149-162.
- Treisman, A. (1996) The binding problem. *Curr. Opin. Neurobiol.* 6:171-178.
- Tononi, G., Spoms, O. and G. M. Edelman (1992) Reentry and the problem of integrating multiple cortical areas: Simulation of dynamic integration in the visual system. *Cerebral Cortex* 2:310-335.
- Tononi, G., Spoms, O. and G. M. Edelman (1994) A measure for brain complexity: Relating functional segregation and integration in the nervous system. *Proc. Natl. Acad. Sci. USA* 91:5033-5037.
- Tononi, G., Edelman, G. M., and Spoms, O. (1998) Complexity and coherency: Integrating information in the brain. *Trends Cogn. Sci.* 2:474—484.
- Tononi, G., and Edelman, G. M. (1998) Consciousness and complexity. *Science* 282:1846- 1851.
- Vaadia, E., Haalman, I., Abeles, M., Bergman, H., Prut, Y., Slovin, H., and Aertsen, A. (1995) Dynamics of neuronal interactions in monkey cortex in relation to behavioral events. *Nature* 373:515-518.
- Von der Malsburg, C. (1995) Binding in models of perception and brain function. *Curr. Opin. Neurobiol.* 5:520-526.
- Von der Malsburg, C. (1999) The what and why of binding: The modeler's perspective. *Neuron* 24:95-104.
- Zeki, S. (1993) *A Vision of the Brain*. Blackwell, Oxford.