

Coordination in Brain Systems

Edvard I. Moser, Maurizio Corbetta,
Robert Desimone, Yves Frégnac, Pascal Fries,
Ann M. Graybiel, John-Dylan Haynes, Laurent Itti,
Lucia Melloni, Hannah Monyer, Wolf Singer,
Christoph von der Malsburg, and Matthew A. Wilson

Abstract

This chapter reviews the concept of dynamic coordination, its mechanistic implementation in brain circuits, and the extent to which dynamic coordination, and specific manifestations of it, have the power to account for functions performed by interacting brain systems. In our discussions, we addressed how on-the-fly changes in coupling between neural subpopulations might enable the brain to handle the fast-changing recombination of processing elements thought to underlie cognition. Such changes in coupling should be apparent, first and foremost, in the statistical relationship between activity in interconnected brain systems, rather than in the individual firing patterns of each subsystem. Dynamic coordination may manifest itself through a variety of mechanisms, of which oscillation-based synchronization is likely to play an important but not exclusive role. Also discussed is how modulation of phase relationships of oscillations in different brain systems, in neocortex and hippocampus of the mammalian brain, may change functional coupling, and how such changes may play a role in routing of signals at cross sections between cortical areas and hippocampal subdivisions. Possible mechanisms for oscillation-based synchronization, particularly in the gamma frequency range, are explored. It is acknowledged that the brain is likely capable of producing zero-phase lag between spatially dispersed cell populations by way of rather simple coupling mechanisms, primarily when neuronal groups are coupled symmetrically. Synchronization with remote areas may be most efficient with phase differences that match the conduction delays. Fast-conducting, long-range projecting interneurons are identified as a potential substrate for synchronizing one neural circuit with another. A number of research strategies are identified to enhance our understanding of dynamic coordination of brain systems and how it might contribute to the implementation of the functions of those systems.

Introduction

In its simplest form, coordination is defined by the multiple interrelations that can be drawn between elements of any given assembly, and its phenomenological expression is signaled by the reconfiguration of elementary dynamics. The potential relations are viewed as functions of an externally defined context or an internally self-generated goal. In their introductory chapter, Phillips et al. (this volume) further constrain the issue by adding: "In general, coordinating interactions are those that produce coherent and relevant overall patterns of activity, while preserving the essential individual identities and functions of the activities coordinated."

We tried to define the process of dynamic coordination by contrasting it with non-vacillating alterations, ever aware of the fact that virtually all neuronal processes are coordinated in one way or another. Much of the coordination of activity required for the formation of specific response properties, such as the receptive fields of neurons in sensory systems or the generation of sequences necessary for the execution of movements, can be achieved by appropriate neuronal architectures that allow for the recombination and sequencing of signals in processing cascades based on divergence, selective convergence, and feedback. Such architectures allow for highly complex coordination and association of signals, even if these originate in separate processing streams, provided that there are adequate connections between the various stages of these processing cascades. To define relations and support selective grouping, it would be sufficient to increase jointly the rate of the responses that are to be associated with each other. Sparse coding and topographic coding would further enhance the salience of rate changes and reduce the risk of grouping unrelated but simultaneously enhanced responses. The fact that a number of phenomena can be predicted by firing rate-based models raises the question whether the fine-grained temporal structure of neural activity in different brain regions has additional explanatory power. Thus, we discussed whether more dynamic mechanisms are required to allow the flexibility, robustness, and speed at which cognition operates in the performing brain.

Dynamic coordination is required when the results of computations achieved in different processing cascades need to be recombined and associated in a flexible, nonstereotyped way. A paradigmatic case is working memory, where ever-changing items have to be temporarily associated with each other. It is unlikely that fixed neuronal architectures would be sufficient to anticipate and cope with the virtually infinite number of possible constellations of associable contents. For each possible constellation, one would need a devoted set of neurons receiving the appropriately selected converging inputs. Because of this limitation, the mechanisms required would need to be capable of establishing transitorily, and in a highly flexible way, relations among signals originating in different processing cascades. These mechanisms would need to be able to select in a dynamic, task- or goal-directed way signals from different, spatially

segregated processing streams and assure selective interactions between, and further joint processing of, these signals. In our deliberations, we considered the possibility that if several such signals are to be coordinated for the first time, that is, before anything is known about their relatedness, a possible mechanism for such relatedness (presumably due to common cause) might be in the form of temporal correlation. This transient coordination may be necessary to establish new linking paths, although its subsequent recruitment might not always be required.

Furthermore, we recognized that most behaviors, and relations between different processing streams, are not generated *ex novo* but are highly predictive based on the history of activation and learning of the system. We rarely reach for a visual object by allowing an arm to cross the midline; we rarely raise a leg and an arm simultaneously unless we are dancing. These relationships cannot be coded just in the anatomy, but must be implemented in the statistical properties of the intrinsic (not task-driven) functional connectivity, which may also require mechanisms of temporal correlation. Dynamic coordination thus does not arise from noise, but from a patterned baseline landscape that may constrain to some extent its flexibility.

The question then arises: How can the presence of dynamic coordination be diagnosed? We agreed that dynamic coordination should be apparent in the patterns of interactions between defined neuronal populations. For all coordinated processes it should be the case that more information can be retrieved by considering the joint activity of neurons belonging to different processing cascades than evaluating the activities of the respective neurons in isolation. In nondynamically coordinated processes, the relations among the respective firing sequences will be stereotyped across trials, if stimulus conditions remain constant. The additional information contained in the relations between the respective firing patterns can therefore be retrieved in sequential recordings from these neurons and with averaging across trials. Such analysis has been applied, for example, to the motor cortex and has led to the discovery of population codes for movement trajectories (see Georgopoulos et al. 1986; Frégnac et al., this volume). However, this approach may be insensitive to more fine-grained temporal relations between the discharges of dynamically grouped neurons. These relations can only be determined by simultaneously recording from the neurons whose activity is suspected to be coordinated. What one might observe in these cases is that the individual responses change only little, if at all, in terms of average rate, while measures of relations between the responses change in a systematic way.

A further constraint is that, in dynamic coordination, the information containing relations must change in a context-, task-, and goal-dependent way. For example, temporal structures may depend on whether a trial was successful or an error trial, and they may change as a function of stimulus context or shifts in attention or goal definitions. To understand the functions of dynamic coordination on the basis of extracellular recordings, it is imperative to perform

multisite recordings and to search for real-time correspondences between temporal structure and task demands. Such approaches may require analytical methods that extract activity patterns in many cells at the same time and take their trial-by-trial covariation into account. Note, however, that it is possible to infer coordination-related processes from single site recordings by appropriate multiscale analysis of intracellular membrane recordings, which reflect the impact of long-distance feedback and lateral connectivity (see Frégnac et al. 2009; also Frégnac et al., this volume).

Experimental evidence for such conditions of dynamic coordination is available. In neuronal recordings during binocular rivalry, for example, whether the considered neuron is part of the processing cascade leading to conscious experience cannot be predicted from the rates of individual neurons in V1. In contrast, a measure of synchrony, in this case precise synchronization of periodic activity in the gamma frequency range, predicts correctly whether or not a given pair of neurons participates in the processing cascade which, in this very moment, conveys the activity that reaches consciousness and is perceived (Fries et al. 1997). In inferotemporal cortex, this information can be retrieved from the rates of individual neurons, suggesting that synchronization (dynamic coordination) at early processing stages was used to select responses for further joint processing by enhancing their saliency, which then facilitated transmission of this synchronized activity to higher levels where it then evoked increased rate responses.

Examples for highly specific inter-regional coordination can also be found in human neuroimaging. Functional neuroimaging is particularly suitable to reveal large-scale dynamic coordination processes that span the entire brain. Several methods are available that allow researchers to investigate how cognitive processes change the interactions between remote brain regions. These include psychophysiological interactions (Friston et al. 1997), dynamic causal modeling (Friston et al. 2003), and Granger causality mapping (Roebroeck et al. 2005). Haynes et al. (2005), for example, have investigated the effects of attention on the connectivity between representations of attended locations. They let subjects attend to two out of four spiral stimuli and report whether they had the same or different handedness. Thereafter they measured the functional connectivity between the individual representations of these stimuli in retinotopic visual cortex. They found that functional connectivity was increased between the retinotopic representations of jointly attended stimuli, both within regions (i.e., V1–V1, V2–V2) and between regions (V1–V2, V2–V1).

A topic of our discussion was whether the definition of dynamic coordination as structured temporal relationships between neuronal populations excludes simultaneous changes in the individual populations to be synchronized. A strict application of the criteria implied for dynamic coordination by Phillips et al. (this volume) would suggest that local processes, such as the discharge fields of V1 neurons, would remain unchanged under associative stimulation protocols. The underlying assumption is that the relational information should

remain separable from the information carried initially by each stimulus component taken in isolation. It was agreed that a strict application of the definition may not be valid as there are numbers of examples which suggest that changes in dynamic coordination give rise to, or are associated with, changes in individual subcircuits. Three examples illustrate the diversity of interactions between coordinated processes and activity in single cells and single populations.

The first example was taken from the stomatogastric ganglion of the lobster. This particular sensorimotor network constitutes a striking example of assembly dynamic reconfiguration correlated with changes in behavior where, during the coordination, the electrical input–output properties of individual elements are not preserved. Coordination is controlled by “orchestra leaders” (PS cell), whose activity triggers the widespread broadcast of neuromodulators. This neuromodulation impacts on the intrinsic reactivity of the other cells by changing reversibly the expressed repertoire of membrane conductances. Consequently, the individual excitability patterns of any given cell will change depending on the context (before, during, or after the orchestra leader cell has fired). Note that in this paucineuronal biological network where all partners are known, the coordinator is identified and the causal link between temporal assembly motifs and the behavioral actions, as well as their functional significance (food swallowing, crunching, and expulsion), are clearly defined (Meyrand et al. 1991).

In a second example, the coordinating agent was part of the high-order statistical features present in the sensory input stream. Changing the statistical regularities of the environment may produce drastic reorganization of ensemble activity patterns and their stimulus-locked reliability. For instance, it is well known that repeated presentation of drifting luminance gratings in V1 receptive fields evokes strong but highly unreliable responses, both at the spiking and subthreshold levels. In contrast, in the same cells, virtual eye-movement animation of natural scenes evokes temporally precise sparse spike responses and stimulus-locked membrane potential dynamics which are highly reproducible from one trial to the next (Frégnac et al. 2005). In this second example, coordination is unrelated to the behavioral outcome or neuromodulation since it is observed in the anesthetized and paralyzed preparation as well as in the attentive-behaving monkey (Vinje and Gallant 2000). This self-organized process adapts the temporal precision of the sensory code to the statistics of the input. However, in contrast to the first example, this adaptive form of temporal coordination is done in the absence of internal executive or supervision units. As demonstrated in the first example, the full field “whole” condition will affect the functional identity of the recorded unit (i.e., the individual receptive fields of the V1 cells).

These two examples illustrate conditions where properties of the individual units of a circuit clearly change in parallel with coordinating processes; however, the literature also contains illustrations where the information that can be stored or recalled on the basis of coordinated activity is separable from the

rate responses of single neurons. This can be seen in recordings of responses to long-duration single frequency tones from the auditory cortex, where highly transient burst responses are detected at the onset or offset of the tone whereas the mean activity is unchanged during the tonic phase of the stimulation. The presence of the stimulus is here signaled by a dramatic and tonic elevation in the correlation between cortical units coding for the sound frequency, without any apparent change of firing rate (deCharms and Merzenich 1996).

It remains a challenge to define a taxonomy of coordination where the underlying mechanisms of each phenomenological form can be clearly separated. Nonetheless, we agreed that dynamic coordination is apparent in a number of studies that show changes in the temporal structure of the joint activity of two or more neuronal populations that differ in character to those taking place at the single-population level.

A Possible Need for Fast-changing Neuronal Architectures

All coordination requires the definition of relations. In the nervous system, relations are established by the anatomical connections among neurons, the anatomical architecture of the networks, and the patterns of inter-regional spontaneous activity, the baseline or intrinsic functional architecture of the networks. Work in nonhuman and human primates indicates that the anatomical connectivity matrix has small world properties allowing for the coexistence of local processing and long-range integration (Kotter 2004; Hagmann et al. 2008). This small world architecture also gives rise to space-time structures of coupling and time delays, which in the presence of noise defines a dynamic framework for the emergence of spontaneous and task-driven cortical dynamics at different temporal scales (minutes, seconds, hundreds of milliseconds) and could support both long- and short-term changes in functional connectivity. To allow for dynamic coordination in behavior (task-dependent selection of responses for joint processing, selective association of subsystems to be engaged, etc.), the functional architecture must be modifiable at the same rapid pace as cognitive and executive processes can change. This requires fast changes in effective coupling among neurons; that is, the gain or the efficiency of a connection must be modifiable. The brain is likely to have a number of mechanisms for achieving such changes in coupling, operating at different timescales.

Coordination by Gain Modulation

Dynamic coupling can to some extent be accomplished by well-characterized gain-modulation mechanisms. Synaptic gain changes can be induced within tens of milliseconds, they may be (but do not have to be) associative and can

last from a few tens of milliseconds (e.g., during frequency-dependent changes in transmitter release) to many decades (e.g., when activity is stored in long-term memory). Effective coupling can also be changed by purely activity-dependent gating, such as when dendritic segments are switched off by shunting inhibition, or when the sequence of activated synapses along a dendrite is changed so that excitatory postsynaptic potentials (EPSPs) either summate effectively or shunt each other, or when the nonlinear amplifying effect of NMDA receptors is enabled or vetoed by local or global adjustments of membrane potential. The question raised in our discussions was whether such gain-modulating mechanisms would be sufficient to account for the speed and flexibility of cognitive operations.

Coordination by Synchronization of Oscillation Patterns

A candidate mechanism for effective change of the coupling among neurons involves rhythmic modulation of discharge activity (neuronal oscillations). Oscillating networks facilitate the establishment of synchrony because they can capitalize on the effects of entrainment and resonance. Oscillators that are tuned to similar frequencies have the tendency to engage in synchronous oscillations if reciprocally coupled. This is the case even if coupling is very weak and even if their frequency tuning is broad and the preferred frequencies are not identical.

An oscillatory modulation of membrane potential, such as occurs in oscillating cell assemblies, confines spiking to the rising slope of the depolarizing phase. Thus, spikes emitted by networks engaged in synchronous oscillations become synchronized. The temporal precision of this synchronization increases with oscillation frequency. In the case of gamma oscillations, output spikes can be synchronized with a precision in the range of a few milliseconds. Because of the coincidence sensitivity of neurons, this synchronization greatly increases the impact that the output of synchronized cell assemblies has on subsequent target neurons.

Another virtue of oscillations is that they allow the exploitation of phase (relative timing) for coding (see discussions on phase precession in the hippocampus in Mehta et al., this volume). In oscillating, synchronized cell populations, responses to strong excitatory inputs will occur earlier on the rising phase of the oscillation than responses to weak inputs. Thus, intensity can be encoded in the time of spiking relative to the oscillation phase. This is a convenient way of coding since the latency of first spikes already contains all information about the amplitude of the driving input. Early studies on retinal coding by Kuffler (1953) showed that relative intensities of visual stimuli can readily be assessed from the relative latencies of the first spikes of ganglion cells. Later studies showed that image reconstruction from first spike latencies is as good as counting rates over several hundred milliseconds (VanRullen and

Thorpe 2001). Thus, readout time for this temporal code is much faster than for the rate code. In the case of the retina, these intensity-dependent differences of spike latencies are of course caused by receptor kinetics. In central processing, the same conversion of an amplitude code into a temporal code can be achieved, in principle, by oscillatory modulation of cell assemblies.

These considerations provide answers to the question: Under which circumstances are oscillations needed? They are needed or at least highly advantageous if (a) spikes have to be synchronized with high precision to support their propagation in sparsely connected networks (see synfire chains of Abeles 1991); (b) spike timing has to be adjusted with high precision for the definition of relations in learning processes such as spike timing-dependent plasticity (STDP); or (c) phase is used as coding space, i.e., if timing relations between spikes or between spikes and the phase of a population oscillation convey information about input amplitude or the relatedness of distributed processes.

There was consensus in our group that several pieces of experimental data are consistent with a role for oscillation-based synchronization in cognitive processes. For example, researchers have shown that attention during visual search correlates with increases in coherence between local field potentials (LFPs) from the frontal cortex and the parietal cortex (Buschman and Miller 2007, 2009). Around the time when monkeys find and shift attention to a visual target, there is an increase in coherence in two different frequency bands: an upper frequency band (35–55 Hz) for bottom-up attention (pop-out), and a lower frequency band (22–34 Hz) for top-down attention (conjunction search). During search for conjunctions, the monkeys shift the location of their attention every 40 ms. The attention-related shifts in frontal eye field spiking activity were correlated with increased power in the lower frequency band, suggesting that the oscillations act as a “clocking” signal that controls when attention is shifted (Fries 2009). The study suggests that serial covert shifts of attention are directed by the frontal eye field and that synchronization between cortical systems may regulate the timing of cognitive processing. Task-induced changes in synchronization or coherence have been reported at the level of individual regions during sensory integration (Roelfsema et al. 1997), selective attention (Fries, Reynolds et al. 2001), working memory (Pesaran et al. 2002; Howard et al. 2003), and motor control (Crone et al. 1998). Between distant cortical regions they have been reported during object recognition (Varela et al. 2001), working memory (Jones and Wilson 2005), long-term memory encoding (Fell et al. 2001), visual attention (Gregoriou et al. 2009), and sensorimotor integration (Roelfsema et al. 1997).

Oscillations and Dynamic Routing

Oscillations may influence routes of communication within structurally constrained brain networks. Consider two groups (*A* and *B*) of neurons that provide

converging synaptic input to a common target group (C) and compete for influence on this target group. If there is rhythmic synchronization among the neurons in group A and among the neurons in group B , but not between A and B , then C will most likely synchronize to either A or B , but not to both at the same time (Börgers and Kopell 2008). The locking of C to either A or B implements a winner-takes-all between the competing inputs of A and B and establishes an exclusive communication link between the target C and the more strongly synchronized input (Fries 2005; Fries et al. 2008).

The described constellation of two neuronal groups converging onto one target group is a fundamental motif in cortex. While this motif renders the postsynaptic neurons selective to diagnostic features of the learned input pattern, it also renders them nonselective or invariant to nondiagnostic accidental features. This invariance is an advantage, because it might provide the basis for object recognition in the face of changes to irrelevant stimulus aspects; however it is also a curse, because a given stimulus will never cover the complete input space of a given neuron, leaving room for competing stimuli. It would be beneficial if the effective input of a given neuron at a given moment in time were limited to functional subsets corresponding to one actual object. This selective efficacy of subsets of a neuron's input might be implemented through the above mentioned exclusive communication link, possibly by synchronization in the gamma frequency band. For this solution to work, two conditions have to be met simultaneously: First, inputs driven by a given stimulus need to be rhythmically synchronized to each other, but not to inputs driven by other stimuli. This corresponds to the binding-by-synchronization hypothesis (Singer and Gray 1995; von der Malsburg 1981/1994). Second, one of the input segments has to be given a competitive advantage over the other through an enhancement. This corresponds to the hypothesis of biased competition through enhanced synchronization (Fries 2005). Thus, the way to use structural convergence in order to harvest both selectivity and invariance seems to lie in the interplay between structural neuronal connectivity and dynamic neuronal synchronization.

In proposing a role for oscillatory activity in dynamic coordination of neuronal populations, our group agreed that one should not forget that oscillatory activity, which may certainly be considered a signature or a manifestation of dynamic coordination, does not necessarily explain the causes or mechanisms by which such coordination arises. Consider a simple dynamic routing example like the one described above, where information in a low-level sensory processing area could be routed toward one of two possible targets in a higher area, with the choice of direction being endogenous, i.e., not dictated by the stimulus. When information is routed one way (choice A), some neurons may oscillate in one manner; when routed the other way (choice B), the same or other neurons may oscillate in a different manner. The mere existence of these oscillations does not explain how the selection was implemented in spatially specific synchronization patterns. The result, signature, or manifestation of

choosing between *A* and *B*, and communicating it to lower areas, may be what is expressed in patterns of oscillations in the functionally connected areas, but how the decision is made and what mechanisms and pathways are employed to communicate it remain pressing questions.

Oscillations and Phase Relationships

Coupling between cell populations is heavily dependent on the phase relationship of the cells in the groups to be linked. By adjusting phase angles, coupling can be modified over the whole range from ineffective to maximal. Controlled changes in firing phases can also be used for dynamic routing if sender and receiver are oscillating at a similar frequency and phases are adjusted. Because oscillations can occur over a wide frequency range, many routes can be specified at the same time without interference. Finally, because coupling in oscillatory networks depends on phase and because we observe the coexistence of oscillations in different frequency bands (theta, beta, gamma), many different and graded adjustments of couplings can be structured, providing opportunities for establishing dynamically graded and nested relations, which could be advantageous for the encoding of compositionality.

Consistent with a role for oscillations in routing of information, experimental data suggest that the phase of the ongoing oscillation can establish preferential windows for information processing. Inputs that arrive in the “good phase” of the ongoing oscillation will be processed preferentially, whereas those arriving at the “bad phase” will be suppressed. For a long time it has been known that the ability to perceive weak signals fluctuates slowly over several seconds (streaking effect). A recent study showed that infra-slow (0.01–0.1 Hz) fluctuations of ongoing brain activity correlate with this behavioral dynamics. In this study (Monto et al. 2008), the probability of detecting a tactile target at threshold was 55% more likely in the rising phase of the fluctuation, and strongly correlated with the power amplitude of higher frequency (1–40 Hz) EEG fluctuations. Support for the same hypothesis comes from two more recent studies showing a relationship between visual detection and phase in the theta–alpha range (Busch et al. 2009; Mathewson et al. 2009).

Finally, we discussed the potential impact of precise phase relationships on learning mechanisms. This seems important because processing architectures have to be adjusted to the requirements of mechanisms establishing durable relations (e.g., in associative learning); that is, they have to transform the (semantic) relations defined during processing into permanent changes in coupling that represent these relations. If any of the mechanisms of associative synaptic plasticity known to date (LTP, LTD, STDP) have anything to do with learning, it would seem that processing architectures need to be capable of defining relations in the temporal domain and that they will have to do so by adjusting the timing of individual spikes with a precision of a few milliseconds.

In STDP, for example, it matters whether an EPSP arrives just a few milliseconds before or after a spike to increase or decrease the efficacy of a connection. It would seem, therefore, that in signal processing and in dynamical coordination, relations should be specified with a similar temporal resolution and precision. Synchronization and phase adjustments in the gamma frequency range could provide the time frame for the precise adjustment of spike timing required for STDP. One should note, however, that the evidence for STDP *in vivo* is relatively scarce and that it is easier to demonstrate the negative (depression) than the positive (potentiation) parts of the STDP curve in the adult cortex (Jacob et al. 2007).

A Wider Repertoire of Coordination Mechanisms

Oscillations cannot be the sole mechanism of dynamic coordination. The neuroscience literature contains a number of examples of dynamic coordination where brain circuits communicate using precise temporal codes not expressed as lasting synchronization in oscillatory patterns in the LFP. The diversity of mechanisms can be illustrated by patterns of hippocampal–neocortical interactions in slow-wave sleep. Slow-wave associated transitions in excitability from low firing rate (putative down-state) to high firing rate (putative up-state) exhibit a systematic timing relationship in which the neocortex leads the hippocampus. During the elevated firing rate period that follows that transition, the hippocampus expresses a series of sharp wave-ripple burst events that replay sequential spatial memory information in which the timing relationship is reversed, with the hippocampus thought to be leading the neocortex. The dialog that may be reflected in shifting timing relationships may reflect the dynamic coordination of oscillatory modes during memory processing.

The possibility of a wider repertoire of mechanisms was further illustrated by discussion of the mechanisms for rapid object recognition in the visual cortex. A given scene may be analyzed in terms of complex arrays of relations by focal attention visiting here and there, and even if the relations thus identified are indeed represented in terms of correlated oscillations, a more permanent trace of these relations must be left behind to be available at later visits by focal attention. If, for instance, there is a number of objects and a number of persons present in a scene, then sequential focal attention may discover which object belongs to which person, one by one, as a result of some inference; when coming back to one of the persons or objects, this result should be available immediately without the necessity of going through the process of inference from scratch. In addition, there is the necessity of maintaining ongoing relations or links between different neuronal ensembles over longer timescales. Many patterns of behavior are predictable, although not necessarily across individuals, repeated over and over with little variation, while at the same time novel behaviors can be stabilized with learning.

What could the mechanism of such short- and longer-term storage of relations be? It was proposed that in addition to the elementary symbols represented by neurons (or groups of neurons), there might be a large network of dynamic links. These links correspond to permanent neural connections, which can, however, be modified (e.g., made ineffective) temporarily. In this view the brain's representations would not have the form of vectors of activity, or neural signals, as in classical conceptualizations, but they would have the form of dynamical graphs. There are various mechanisms by which the efficiency of connections can be rapidly modified. There is synaptic plasticity on a continuous range of timescales, starting from a few milliseconds, and the effective connectivity (Aertsen et al. 1989) of a network can be changed by a variety of presynaptic and postsynaptic influences.

Recent work in human brain imaging shows that spontaneous activity, as measured by fluctuations of the blood oxygenation level dependent (BOLD) signal, is not random but organized in specific spatiotemporal patterns (Deco and Corbetta 2010; Fox and Raichle 2007) that resemble functional networks recruited during active behavior. These correlations occur at a very slow temporal scale (<0.1 Hz), which correspond to fluctuations of slow cortical potential (0.1–4 Hz) and band limited power fluctuations of the gamma band (He et al. 2007). These patterns of spatiotemporal correlation at rest reflect not only the underlying anatomy, but are gated by their recruitment during tasks. The leading hypothesis, supported by studies showing changes with learning and lesions, is that these patterns of spontaneous activity code for relations in the cortex that are related to the history of network activation and learning. They may represent attractor states that constrain and potentially bias the recruitment of brain networks during active behavior.

If these views are correct, then neuroscience is currently ignoring a large part of the representational machinery of the brain—very large indeed, as there are many more connections than there are neurons in the brain.

If coordination is expressed largely by dynamic connections, then what is the importance of signal correlations? We agreed that signal correlations are likely to be indispensable when a set of neurons are to be coordinated for the first time; that is, when the downstream circuits have not yet encoded this relatedness in their link structure. In this way, neural oscillations could play a vanguard role, appearing only early in some learning task, disappearing as soon as the coordination pattern is encoded in some connectivity structure.

Computer modeling work will be particularly useful in shaping our thoughts about neural operations if the model can be related in a convincing way to neural operations, instead of just using the brute force of high-speed computers, and if the performance of the model can be proven superior in public benchmark tests. Such tests are available for face recognition (e.g., FRVT 2002; Messer et al. 2004; Phillips et al. 2005). The consistently winning systems were all correspondence-based; that is, they are based on representations of faces in terms of two-dimensional arrays of local features (mostly of Gabor

type, i.e., modeled after receptive field types found in primary visual cortex) and on finding correspondences between local feature sets in model and image. As a large number of such correspondences are to be found for a given match, temporal coding is bound to be rather time-consuming. A model of correspondence-finding by temporal coding (Wiskott and von der Malsburg 1996) would, if implemented with realistic neurons, have needed more than 10 s to recognize a face, two orders of magnitude slower than human performance. If, however, connectivity patterns were installed that allowed for the fast dynamic installation of topographic fiber maps with the help of control units (introduced by Olshausen et al. 1993), high-performance face recognition by correspondence-finding within physiological times of 100 ms was feasible (Wolfrum et al. 2008). In a related study (Bergmann and von der Malsburg, submitted), it is shown that the necessary control unit circuitry can be developed on the basis of synaptic plasticity controlled by synchrony coding.

Correspondence-based object recognition models have explicit representation of shape. It was argued, however, that pure feedforward models (such as Serre et al. 2007; Poggio and Edelman 1990), which do not make use of dynamic coordination, are also able to represent shape. The Chorus of Prototypes algorithm (Duvdevani-Bar and Edelman 1999) and the Chorus of Fragments model (Edelman and Intrator 2003) may represent shape if endowed with a mechanism for relating together the responses of the ensemble of neurons that represent, in a distributed yet low-dimensional manner, the current input. Temporal binding by synchrony may be just such a mechanism (as was proposed in Hummel and Biederman 1992).

The Mechanisms for Synchronization between Neural Populations

Although there was consensus that the brain has a wide repertoire of mechanisms for achieving dynamic coordination, we chose to discuss in more detail coordination by synchronization of oscillatory activity across neuronal populations. This form of coordination has support in the experimental literature, as suggested above, and there is now a considerable literature exploring mechanisms of synchronization at the level of cellular assemblies.

We began the discussion by reviewing models for synchronization between cell populations. Several models have been proposed as mechanisms for achieving zero-phase lag between same-frequency oscillatory activity in different populations, which by definition might be seen as the ultimate expression of synchronization. Evidence indicates that zero-phase lag synchronization is ubiquitous and can occur over surprisingly large distances, such as between the hemispheres (Engel, König, Kreiter et al. 1991), despite the rather considerable conduction delays of pathways connecting the synchronized assemblies. At first it may seem that such synchronization between widely dispersed populations could be achieved only by common input from a central

oscillator that slaves the respective synchronized assemblies. However, since callosotomy abolishes interhemispheric synchrony, it may instead rely on interactions between the synchronized assemblies. Several models for such interactions were considered.

One class of models relies on spike doublets of inhibitory interneurons (Traub et al. 1996; Ermentrout and Kopell 1998). When neurons synchronize locally in the gamma band, there is a characteristic interaction between excitatory and inhibitory neurons: excitatory neurons spike first and trigger inhibitory neuron spiking with short delays. The ensuing inhibition shuts down the local network until inhibition fades and the cycle starts again with the firing of excitatory neurons. Long-range synchronization between two such gamma oscillatory groups can occur when excitatory neurons of group *A* excite interneurons of group *B*, even if this entails a conduction delay of a few milliseconds between *A* and *B*. Essentially, the excitatory input from *A* to *B* triggers a second inhibitory spike in *B* and thereby prolongs the inhibition inside *B* by the conduction delay. The two interneuron spikes in rapid succession gave this model the name “spike doublet mechanism.” One prediction of this model is that local gamma band synchronizations decrease in frequency when coupled across long distances and the frequency decrease is proportional to the conduction delay.

Synchronization across long distances might also be supported by other configurations of reciprocal interaction between the subcircuits. Evidence is now available which shows that zero-phase synchrony can be established despite conduction delays in the coupling connections both from experiments with coupled lasers (Fischer et al. 2006) and modeling of networks with spiking neurons (Vicente et al. 2008). As long as at least three reciprocally coupled systems are allowed to interact (triangular configurations), zero-phase synchrony is easily established and very robust against scatter in conduction times of coupling connections.

A useful mathematical perspective on the phenomenon of zero-phase synchronization comes from the study of coupled map lattices and globally coupled maps. These are systems of coupled nonlinear dynamical systems, whose long-term (ergodic) behavior can show some universal properties under some simplifying assumptions (Tsuda 2001). One of these assumptions is that the system is globally and symmetrically coupled with a single coupling strength. Under these constraints, it can be shown that the states of every coupled dynamical system come eventually to occupy a synchronization manifold. Crucially, because of the symmetry constraints on the dynamic equations, the set of all solutions must obey the same symmetry. Zero-phase synchronization represents a symmetrical solution. Due to spontaneous symmetry breaking, however, individual solutions might violate symmetry (i.e., exhibit non-zero-phase synchronization). A simple example is a ball sitting on top of a hill in a completely symmetric state. However, as this state is unstable, the slightest perturbation will cause the ball to roll downhill. This movement will not occur

symmetrically in all directions but in each case in some particular direction. Each single solution (i.e., rolling in a specific direction) breaks the symmetry of the initial problem. Only the set of all solutions and the probability for the ball to roll in a specific direction is symmetric. This has been used in a model investigating oscillatory interactions in primary visual cortex (Schillen and König 1991). Here a specific type of excitatory tangential connection avoids the trivial solution of global synchronization. With the additional assumption of ergodicity (i.e., the system that evolves over a long timescale and visits all regions of state space), individual solutions have to obey symmetry constraints. For systems with more than two coupled oscillators, this reduced the solution for the entire system to global synchrony with zero-lag quasi periodic or chaotic oscillations. At low but nonnegligible coupling strength, the synchronization manifold is “riddled” with unstable points that “eject” the trajectory away from the synchronization manifold to produce intermittent bursts of localized activity (Breakspear et al. 2009).

The main message from these theoretical treatments is that there is nothing mysterious about zero-phase lag synchronization among three or more populations. Indeed, under the constraints of the model, it is impossible to get any synchronization other than zero lag. To get consistent (nonzero) phase coupling, one has to break the symmetry, in terms of the intrinsic parameters of the system or its coupling parameters. This basic phenomenon has been illustrated using neuronally plausible simulations by Chawla et al. (2001), where it proved difficult to break the symmetry provided by three or more neuronal systems that are interconnected in a roughly symmetrical fashion.

Although zero-phase synchronization could serve as a useful guide for understanding the mechanisms underlying long-range synchronization between neural circuits, it was argued that the current models for producing zero-phase lag have relied on unrealistic architectures and that the physiological properties of the model neurons do not match those of the performing brain (e.g., neurons do not regularly fire in doublets during gamma oscillations). It was proposed that zero-phase lags may not even be desirable for synchronization when information is communicated over long distances. It may often be advantageous to introduce systematic phase shifts to coordinate convergence of distributed information from different sources or to enforce timing relationships that would establish specific patterns of dynamic routing. The actual phase lags between oscillating populations in two regions may vary across task conditions and network states. One example for modulating the efficiency of interareal coupling by systematic phase shifts between oscillatory activity is cortico-tectal communication (Brecht et al. 1998). It was also recognized that regulation of spike timing through systematic phase locking can be used to encode temporal relationships (such as spatial behavioral sequence encoding through theta phase precession).

An example of time-shifted synchronization across brain areas was recently reported in a study of frontal eye field (FEF)–V4 interactions in an attention

task (Gregoriou et al. 2009). There is considerable evidence that FEF plays an important role in the top-down control of attention in visual cortex, including V4. In the Gregoriou et al. study, spikes and LFPs were recorded simultaneously from FEF and V4, in monkeys trained in a covert attention task. One stimulus always appeared inside the shared receptive field and two others appeared outside; the monkey was cued to attend to a different stimulus on each trial. Spike-field coherence in the gamma band increased with attention in V4 and FEF. The effect was particularly strong when cells in the two areas had overlapping receptive fields. However, there was almost a 180° phase lag in synchrony in the gamma frequency band between FEF and V4, corresponding to a time delay of about 10 ms. The same 10 ms time shift was found in other frequency bands of the V4–FEF synchronous activity, suggesting that there is a constant 10 ms time shift between the time while cells spike in one area and cells are maximally depolarized in the other. It was suggested that this time shift may be accounted for by conduction and synaptic delays between the two areas. If so, then spikes from one area would actually arrive in the connected area at a time when the receiving cells were most prepared (depolarized) to receive them, which is consistent with the strong effects of FEF activity on the top-down attentional modulation of V4 responses. The study illustrates the potential role of time-shifted synchrony between areas as a common mechanism for functional interactions between cortical areas and raises the possibility that zero-lag synchronization may be implemented primarily in local circuits.

To add to the complexity, a neuronal population may have different phase lags to different subsets of a population with which it interacts. The recent description of traveling theta frequency waves in the hippocampus (Lubenov and Siapas 2009) suggests that neurons in regions that communicate with the hippocampus may be synchronized with a subset of the hippocampal population across a wide range of the oscillatory cycle, but the identity of the neurons with which synchronization occurs may change with phase. These phase lags may influence the wider patterns of coherence between the hippocampus and other structures, such as the striatum, for which phase angle changes with task and with learning (Tort et al. 2008; DeCoteau et al. 2007).

There was general consensus that the mechanisms enabling synchronous firing across widespread brain regions are poorly understood, especially for the higher frequency (e.g., gamma), and that alternative solutions should be considered. One possibility considered involves long-range axonal collaterals. Synchronization between two oscillating populations might be achieved if the collaterals of gamma-modulated pyramidal cells in one location phase reset the basket neurons in the other location. A fundamental problem with this model is the limited axonal conduction velocities of pyramidal cells. An alternative fast-conducting conduit between distant sites may instead be provided by axon collaterals of so-called “long-range” interneurons. The anatomical “short cuts” provided by the long-range interneurons may offer the interarea fast transmission that is required to phase-synchronize gamma oscillations

between distant cortical regions. Such far-projecting fast-conducting interneurons have been described in the hippocampus (Sik et al. 1994). The axons of this interneuron family innervate multiple regions of the hippocampus and can project to multiple external regions, including medial septum, subiculum, pre-subiculum, entorhinal cortex, induseum griseum, and possibly other cortical regions. Similarly, GABAergic interneurons in the medial septum project to the hippocampus, preferentially to GABAergic interneurons (Freund and Antal 1988), hippocampal GABAergic neurons provide long-range projection back to the septum (Gulyás et al. 2003), the basal forebrain has GABAergic neurons that project widely across the cortex (Sarter and Bruno 2002), and long-range GABAergic interneurons are known to connect remote areas in the ipsilateral and contralateral cortex (Buhl and Singer 1989; Gonchar et al. 1995; Kimura and Baughman 1997; Tomioka et al. 2005). A common property of many of the long-range interneurons is that their axon caliber is nearly twice as large as that of parallel conduits from pyramidal cells connecting the same regions and the diameter of the surrounding myelin is three times thicker (Jinno et al. 2007). The estimated volume of the total axon arbor of a long-range interneuron is several times larger than the volume occupied by the axon tree of pyramidal cells, suggesting that only few such neurons may be needed to establish coherence between regions. There was consensus in the group about the need for further investigation of the potential role of long-range fast-transmitting inhibitory interneurons in fast inter-area cortical synchronization.

We also discussed the potential role of ascending neuromodulatory systems in synchronization of activity across brain regions. The broad terminal fields of axonal projections from monoaminergic and cholinergic cell groups generally speak against a role in controlling dynamic changes in specific subsets of interacting cell clusters, as does the slowness of many receptors for such transmitters (e.g., dopamine) and the long time that it takes to clear the transmitter from the synaptic cleft. These facts do not, however, exclude a key permissive function for ascending neuromodulatory systems in providing necessary conditions for inducing oscillatory activity. The discharge patterns of cholinergic as well as monoaminergic cell groups change radically during transitions between brain states (e.g., when subjects switch between awake states and sleep), and such changes are temporally correlated with massive changes in the oscillatory properties of cortical networks. Observations suggest that, although terminal fields are broad, subtypes of intermingled interneurons are innervated by different neuromodulatory systems (e.g., 5-HT axons terminate on CCK-expressing interneurons but not parvalbumin-expressing cells, whereas cholinergic projections primarily terminate on basket cells). The specificity of the neuromodulatory innervation, as well as the specific combinations of receptor subtypes expressed by different classes of interneurons, and the ability of neuromodulators to change the time constants of GABA receptor potentials are likely to have significant impact on the generation of oscillations

and synchrony across brain regions, although the exact mechanisms remain to be determined.

In our discussions, we briefly straddled the issue as to whether areas in the neocortex only exchange information once they have finished their respective computations and then transmit the result (discontinuous communication) or whether they permanently interact (continuous processing) until they converge to a collective result. We felt that the latter scenario is more realistic, although some ERP studies seem to suggest that information is transmitted in discrete packages.

Steps into the Future

Computational Models

How might neuroscientists improve their understanding of the brain's mechanisms for dynamic coordination? Our discussion of models and experiments will be presented sequentially, although the consensus is that advances require an integrated approach.

Models will play a critical role in interpreting the many disparate empirical findings regarding coordination in neural systems. Cortical models for dynamic coordination across brain systems can be roughly categorized according to whether they are focused on the role of large population influences on single neuron properties versus models centered on the nature of interactions across two or more specific cortical structures or layers.

There are numerous examples of models that examine the effects of attentional feedback or task demands on single neuron properties. The feedback in these models comes from unspecified sources, and in most cases the models only consider the effects of feedback on average firing rates. In the field of attention, for example, biased competition (Desimone and Duncan 1995; now described as normalization models, Reynolds and Heeger 2009), feature-similarity gain (Maunsell and Treue 2006), and response gain models all attempt to explain how attentional feedback cause the enhancement of responses to attended targets and the suppression of responses to unattended distracters. Normalization models explain and predict the large majority of attentional effects that have been reported on single neuron properties.

In contrast to these attentional models, based on average firing rates, some models also address the role of spike timing and synchrony in neural populations. It is claimed that only spiking neuron models that incorporate gamma synchrony can explain the effects of attention on competing stimuli within the same receptive field (Börgers et al. 2008), although direct tests of competing models on these data are missing. In the future, it will be critical to make differential predictions from models based on static firing rates versus synchrony

and population dynamics, which can then be tested empirically in neurophysiological studies.

Fewer quantitative models take on the daunting task of modeling the interactions among two or more cortical areas. Efforts are ongoing to collect data on a large number of individuals (upward of 2,000 healthy subjects) to characterize the anatomical, functional, and electrophysiological neuromatrix of the human brain (The Human Connectome Project). The goal of this project is to provide the neuroscience community with a public data set, which will hopefully describe for the first time the entire array of cortical areas, as well as their anatomical and functional links. This will allow quantitative mathematical modeling of their properties and exploration of the range of dynamics and interactions that are possible within these networks both in healthy and damaged brains. Presently, more limited systems models are being considered. In attention, Hamker (2005) proposed a model that incorporates interactions among a large number of visual areas and the “attentional control” system that provides feedback. Quantitative models of object recognition typically incorporate the receptive field properties of neurons located along the ventral stream. Examples of these types of models are ones developed by Poggio, Edelman and colleagues (Serre et al. 2007; Edelman and Duvdevani-Bar 1997; Duvdevani-Bar and Edelman 1999). These models are strictly feedforward, based on findings that inferotemporal neurons show object-selective responses at times so short that they seem to preclude multiple recursive cycling up and down the visual pathways. When trained on a large database of images, these models are able to achieve recognition performance of human observers who classify images based on very brief stimulus presentation times. For more complex, cluttered scenes that require more recognition time, the latest version of the Poggio model incorporates attentional feedback (Chikkerur et al. 2009). By contrast, the face recognition model of von der Malsburg incorporates feedback to visual cortex from neurons holding stored representations of faces (see above). This feedback model achieves a high level of performance on published databases of faces. However, it was argued that in all of these system models, only average firing rates are considered and the timescale of the feedback is still relatively slow. A critical goal for the future is to find out whether the proven success of object classification and face recognition models are only first steps and that models based on binding mechanisms can be expanded into a broad range of functional models for dynamically coordinated perception.

It was agreed that an essential element for evaluating models is their performance on large, publically available image databases. Although some databases exist, there is a need for more realistic conditions in the databases, including the recognition of objects at different scales and embedded in complex scenes. Furthermore, beyond simple recognition, there is a need for models that can answer at least basic questions about the objects, such as shape, size, and location. The development of such models will help in understanding how and why synchronous interactions may be important for perception and memory.

Experiments

We considered a number of experimental approaches to the testing of the role of dynamic coordination in cognitive performance. Because oscillations and synchronization are currently the best-explored mechanistic paradigms, our discussion focused on possible ways to test whether such phenomena are necessary and sufficient for the cognitive functions performed by those brain regions where synchrony is observed. There was consensus that such experiments must monitor activity from two or more cell populations at the same time; as discussed in the introductory section, changes in the joint activity of two or more cell ensembles can be seen as a defining criterion for dynamic coordination.

We agreed that much of the current evidence linking synchronization of oscillatory patterns to coordination functions is correlation-based—a concern that is shared with most other fields of study in systems neuroscience—and that results thus, in principle, might be explained by other models, including those based solely on rate changes. However, the literature does contain some interventional studies which at least partly address the question of whether synchronization between cell populations is necessary for behavioral functions relying on the synchronized assemblies. In a study with multisite recordings from the frog retina, for example, activity was recorded from cells that respond to changes in shadows on the retina. Interventions that disrupted the synchrony of firing across the recording electrodes disrupted escape behavior elicited by shadow stimuli under conditions that did not change the average rates of the cells (Ishikane et al. 2005). Other experiments, performed in the hippocampus of the rat, have shown that using cannabinoids or other approaches to disrupt temporal order in hippocampal place cells, in a manner that does not change the average firing rates of the neurons, is sufficient to disrupt navigational performance in a spatial memory task (e.g., Robbe et al. 2006). In awake-behaving monkeys and healthy human subjects, some experiments have modified activity in visual cortex during stimulus detection by stimulating putative attention control regions in frontal cortex (Ruff et al. 2006; Ekstrom et al. 2008). Interference with frontal or parietal regions by TMS has been shown to alter, in behaviorally significant ways, anticipatory alpha rhythms in occipital visual cortex (Capotosto et al. 2009). The invention of optogenetic tools for selective stimulation or silencing of genetically defined cell populations is likely to result in a number of experiments along these lines within the next few years. It is clear that synchrony can be interrupted experimentally, and those data that exist so far suggest that such interventions may disrupt the functions performed by the affected cell populations.

Although interventional approaches represent the gold standard for studies of causal relations between coordination and brain function, we agreed that the caveats of such studies should not be forgotten. Interventions such as stimulation or inhibition of target cell populations may have additional effects

on top of the intended ones; for example, disrupting synchronization between brain areas may also affect the proximal activity of each subpopulation, such as firing rates or precise local phase relationships. We also spent some time discussing strategies for gaining insight about coordination mechanisms under circumstances where physiological variables cannot be manipulated directly. One possible approach exploits the fact that human subjects often confuse the color and shape of different objects. Such “illusory conjunctions” can be used as a diagnostic tool to investigate which neural mechanism breaks down during binding errors. To investigate the role of response synchronization in feature binding, one could ask patients with intracranial electrodes to report on the color and shape of multiple objects in the visual field under conditions that lead to occasional misbindings (e.g., when stimuli are presented very briefly). One would need to record from cells that encode two distinct properties of two different objects in the visual field (e.g., one could choose color and motion as features and then record from color-sensitive cells in V4 and from motion-sensitive cells in MT). If synchronization is indeed the neural mechanism for feature binding, one would expect that the action potentials of cells belonging to the same object are synchronized when perception is successful, and that synchronization reflects illusory conjunctions when they occur. The same recordings could be used as well to test a different model, where the positional information encoded in V4 and MT signals maps corresponding features together. In this case, the positional information might be disrupted or shifted in either of these populations, thus providing a potential alternative account for the misassignment of features and spatial positions. If intrinsic dynamic connectivity turns out to be an important mechanism to code relations, especially for behaviors that are predictable or well-learned, then new investigations should be directed toward manipulating the ongoing intrinsic connectivity, either through behavioral paradigms or interventions like stimulation or disruption, and then correlate these changes to behavioral performance or task-driven activity.

We concluded that a variety of experimental approaches and systems are available to explore the function of oscillation-based synchronization and other possible mechanisms of dynamic coordination between neuronal populations. A common factor of all experiments that aim to test these functions should be the recording of activity from two or more brain regions at the same time; this is the only way to study changes in inter-regional temporal structure that may or may not be accompanied by activity changes in each of the areas locally. A number of brain systems, each with their unique advantages, should be used to extract the mechanisms of coordination. The study of temporal structure in large dispersed neuronal populations is likely to require an arsenal of new analytical and statistical techniques. Finally, there is a strong need for interaction between computational models and experimental testing; models should make clear predictions about activity changes in realistic neuronal architectures, and

experimental strategies should be developed to test specific predictions from the models.