

1

Dynamic Coordination in Brain and Mind

William A. Phillips, Christoph von der
Malsburg, and Wolf Singer

Abstract

This chapter discusses the concept of dynamic coordination and the major issues that it raises for the cognitive neurosciences. 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. Dynamic coordination is the coordination that is created on a moment-by-moment basis so as to deal effectively with unpredictable aspects of the current situation. Different computational goals for dynamic coordination are distinguished and issues discussed that arise concerning local cortical circuits, brain systems, cognition, and evolution. Primary focus is on dynamic coordination resulting from widely distributed processes of self-organization, but the role of central executive processes is also discussed.

Introduction: Basic Concepts, Hypotheses, and Issues

How Is Flexibility Combined with Reliability?

The universe is lawful but unpredictable. Regularities make life possible, but unpredictability requires it to be flexible, so, biological systems must combine reliability with flexibility. Neural activity must reliably convey sensory information, cognitive contents, and motor commands, but it must do so flexibly and creatively if it is to generate novel but useful percepts, thoughts, and actions in novel circumstances. Neural activity, however, is widely distributed, which suggests that activity is dynamically coordinated so as to produce coherent patterns of macroscopic activity that are adapted to the current context, without corrupting the information that is transmitted by the local signals.

It is clear that, both within and between brain regions, different cells or groups of cells convey information about different things. New techniques and findings continue to add to our knowledge of this local selectivity and

its gradual adaptation to different environments and tasks. Much of cognitive neuroscience is therefore primarily concerned with finding out what is dealt with where; that is, with localizing brain functions. We need now to understand better how these diverse activities are coordinated. Most locally specialized streams of activity can affect most others, either directly or indirectly. In that sense, the brain operates as an integrated whole. It is also clear that percepts, thoughts, and actions are normally, in some sense, both “coherent” and “relevant” to the current context, and psychological studies provide much information on the underlying processes. However, we do not yet have widely agreed conceptions of “coherence” and “relevance,” nor of the dynamic coordinating interactions by which they are achieved. A central goal of this Forum was therefore to review the relevant evidence, to assess the underlying concepts, and to evaluate possible ways of exploring the major issues that arise.

Prima facie, the need to combine flexibility with reliability raises difficulties because of contrasting requirements. To be reliable, neural codes must code for the same thing when used at different times and in different contexts; however, to be flexible, codes must be used in different ways at different times and in different contexts. The notion of dynamic coordination directly raises this issue because it emphasizes the need to combine and sequence activities in a context-sensitive way that changes from moment to moment, while implying that the activities coordinated maintain their distinct identities.

Thus, fundamental issues arise for neurobiology, neurocomputational theory, cognitive psychology, and psychopathology. Does dynamic coordination depend on specific synaptic or local circuit mechanisms? Does it depend on distinct neural pathways? Can the capabilities and requirements of dynamic coordination be clarified by neurocomputational theory? What are the cognitive consequences of dynamic coordination, and what are the psychopathological consequences of its malfunction?

How Is Holism Combined with Localism?

Localism and holism have been contrasting themes throughout the history of neuroscience, with the evidence for localism being clearly dominant (Finger 1994). Emphasis upon holistic organization has waxed and waned within both neuroscience and psychology. The form of holism known as Gestalt psychology was strong for a few decades from the 1920s, but has had limited impact since then. It is now ready for a renaissance in which its phenomenological methods are complemented by rigorous psychophysical methods, and its speculations concerning electrophysiological fields are replaced by better founded neural network theories (Kovács 1996; Watt and Phillips 2000).

The apparent conflict between localism and holism is clearly shown by studies of the visual cortex. The classic studies of Kuffler, Hubel, Wiesel, and many others provide strong support for localism. They show that individual cells or small populations of cells convey information about very specific

aspects of visual input. Higher levels within the hierarchy of processing compute more abstract descriptions, but local selectivity is maintained up to at least object recognition levels, and probably beyond. Analogous forms of local selectivity have been found throughout all perceptual and motor modalities. In contrast, many other experiments clearly show that activity in the visual cortex is strongly affected by a surrounding context that extends far beyond the “classical receptive field” (Gilbert 1992; Kovács 1996). There is now evidence for many forms of contextual modulation, dynamic organization, and task dependence throughout all levels of the visual system (Albright and Stoner 2002; Lamme and Spekreijse 2000). Such findings provide clear support for the holistic tradition and have led some to doubt the value of the classical localist tradition, emphasizing instead the rich nonlinear dynamics of highly connected networks. As it is clear that the arguments and evidence are very strong for both localism and holism, a better understanding is needed of how they are combined.

What Is Dynamic Coordination?

Though various aspects of the notion of dynamic coordination have long been implicit in neuroscience, it is still in need of a clear explicit expression. Here, we provide an informal perspective and briefly discuss more formal views. For further discussion of coordination and its functions, see von der Malsburg (this volume).

Simply put, dynamic coordination refers jointly to fundamental neuro-computational functions such as contextual gain modulation (Phillips and Silverstein 2003; Tiesinga et al. 2005), dynamic grouping (e.g., Singer 2004), dynamic linking (e.g., Wolfrum et al. 2008), and dynamic routing (e.g., Fries 2005). All imply that novel context-sensitive patterns of macroscopic activity can be created by modulating the strength and precise timing of local neural signals without corrupting the information that these local signals transmit. For example, consider attention, which is a paradigmatic case of gain modulation. There is ample evidence that attention amplifies or suppresses signals conveyed by local processors, but with little or no effect on the receptive field selectivity of the cells generating those signals (Deco and Thiele 2009; Reynolds and Heeger 2009; Maunsell 2004). There is also ample evidence that concurrent contextual inputs from far beyond the classical receptive field can produce similar modulatory effects on the strength and timing of local signals (Gilbert 1992; Lamme 2004). Evidence concerning contextual modulation, grouping, linking, and routing was discussed in depth at the Forum and is reviewed in the following chapters.

The proposed distinction between coding and coordinating interactions can be clarified by considering coordination at a molecular level. The genetic code that was first discovered in the structure of DNA concerns the mapping of codons (i.e., sequences of three bases along the DNA chain) to amino acids.

This coding has remained essentially constant throughout the history of life on Earth. We now know that expression of those codes is controlled in a highly context-sensitive way by noncoding sequences and epigenetic mechanisms to ensure effective adaptation to current circumstances. Furthermore, we know that only a small percentage of codons code for amino acids and that the majority are involved in coordinating their expression. Thus, the effects of coding sequences, noncoding sequences, and epigenetic mechanisms are seamlessly intertwined, but coding and coordinating mechanisms have nevertheless been clearly distinguished.

Our working assumption is that the distinction between coding and coordinating interactions can also be applied to brain function. At each level of hierarchical processing, single-unit or population codes are used to represent many diverse entities. At each level, expression of those codes is flexibly controlled in a highly context-sensitive way by various coordinating interactions to ensure effective adaptation to current circumstances. The effects of coding and coordinating interactions are seamlessly intertwined in all neuronal activity and behavior, but they are clearly distinguishable. Under natural conditions the selection, timing, and amplification or suppression of signals is coordinated with concurrent stimuli and current goals. The contribution of these coordinating interactions can be greatly reduced by presenting isolated stimuli to anesthetized animals. When this is done we find a highly selective and hierarchical functional architecture, as revealed by much of neuroscience since the 1950s. The selectivity thus discovered has led to fundamental concepts such as that of receptive field selectivity and regional specialization. In stark contrast to the genetic code, however, receptive field coding is not fixed; it adapts to the environment in which the organism finds itself. This adaptation is usually gradual, as when visual cortex adapts to the statistical structure of visual input, but it may also operate more rapidly. When it does, the new codes must remain reliable over the time span of their use, and their rate of creation must not outrun the ability of projective sites to interpret them correctly.

As our primary focus is on dynamic coordination rather than on local selectivity, we must first make clear that local selectivity does make a crucial contribution to neural activity in awake-behaving animals in natural environments. First, any system that learns from examples must find some way to overcome the “curse of dimensionality” (Bellman 1961; Edelman 2008a); that is, an exponential increase in the number of samples required with dimensionality. This can be done by first mapping the input into low-dimensional subspaces or manifolds. Fortunately, neural systems can do this because the natural input they receive has a hierarchical structure that can be exploited to reduce dimensionality. Second, there is plenty of evidence showing that neural systems discover and exploit this hierarchical structure. For example, at the macroscopic level of neuroimaging, cognitive neuroscience is replete with evidence for reliable local selectivity within and between cortical regions in awake-behaving subjects, as well as across subjects within and between species. Much remains to

be discovered, however, concerning this local selectivity. *Prima facie* evidence that “context” modifies local selectivity may therefore be more easily interpreted by refining notions of the local selectivity involved.

Thus, in arguing for a central role for coordinating interactions, we are not arguing against a role for local selectivity. Rather, we take that as a given and then argue that the selectivity is, and must be, so great that coordination is inevitably required. Figure 1.1 demonstrates the modulatory role of context in vision to emphasize our central concern with the Gestalt level of organization created by these coordinating interactions. Contextual modulation is common in perception and can be seen as using Bayesian inference or prior probabilities to interpret input (Schwartz et al. 2007). In response to discussion of the concept of dynamic coordination (e.g., Moser et al., this volume), we must



Figure 1.1 Demonstration of the strength and validity of contextual disambiguation, based on a computer-generated demonstration of lightness constancy by Edward Adelson (MIT) and on a painting by René Magritte. The upper panel is a photograph (taken by Ben Craven, <http://bencraven.org.uk/>) of a real checkerboard and pipe. The central light-gray square of the checkerboard actually transmits less light to the camera (or eye) than the arrowed dark-gray square because of the pipe’s shadow. The lower panel, where all but the arrowed black square and the central white square of the checkerboard are deleted, demonstrates this; cues that the visual system uses to disambiguate reflectance from incident light intensity have been removed. Most people will see the central square as being much lighter in the upper panel than in the lower. The central squares are, however, identical and are in both panels darker (brightness about 27, as measured by Photoshop) than the arrowed square (brightness about 33). Such contextual disambiguation is ubiquitous throughout perception, demonstrating that the capabilities of attention are highly constrained. We cannot voluntarily ignore the surrounding context (e.g., when looking at the central square in the upper panel). We think that we can attend to independent elements of sensory awareness, but we cannot. This is but one of several common, but false, beliefs about the nature of consciousness.

emphasize that the concept does not imply that coordinating interactions have no effect on local circuit activities. On the contrary, they can have significant effects. The point is that these effects preserve local selectivity. This obviously applies to dynamic grouping, linking, and routing, and in relation to contextual disambiguation by gain modulation. Tiesinga, Fellous, and Sejnowski (2008:106) state that “multiplicative gain modulation is important because it increases or decreases the overall strength of the neuron’s response while preserving the stimulus preference of the neuron.”

To some it might seem that, in principle, coding and coordinating interactions cannot be distinguished, because signals necessarily transmit information about everything that affects them. This intuition is misleading. The information that is transmitted specifically about a modulatory or coordinating input given the coding, or receptive field, tends to be negligible, even when that modulating input has a large effect on the transmission of receptive field information (Kay et al. 1998). Therefore, conditional mutual information measures can be used to distinguish coordinating from coding interactions (Smyth et al. 1996). By applying these measures to two alternative forced-choice data obtained in a texture segregation task with multiple cues, it was found—as predicted—that cue fusion is a coding rather than coordinating interaction, whereas attention involves coordinating rather than coding interactions (Phillips and Craven 2000).

Three formal conceptions are discussed by Engel et al. (this volume): coherent infomax (e.g., Kay et al. 1998), coordination dynamics (e.g., Kelso 1995), and predictive coding under the free energy principle (e.g., Friston 2009; Friston and Keibel 2009). They were found to have much in common, including emphasis upon the necessity for dimensionality reduction, correlations between dimensions, and the distinction between driving and modulating interactions. A major difference in emphasis between coordination dynamics and coherent infomax is that the former is concerned with coordination in general, whereas the latter has been predominantly used in relation to dynamic coordination in particular. Another difference is that the formal analyses of coordination dynamics have been applied mainly to sensorimotor coordination, whereas coherent infomax has been applied mainly to perception. Optimization under the free energy principle is the most general and extends beyond coherent infomax by formally including a loop back from motor output to sensory input, thus enabling the system to achieve its objectives both by adapting itself to its world, and its world to itself.

Above we used the notion of dynamic coordination to mean coordination that is generated on a moment-by-moment basis so as to deal effectively with the current situation. It contrasts with prespecified coordination, i.e., those activity patterns that are explicitly specified and available before the current tasks and stimulus conditions are known. Highly practiced or stereotypical patterns of coordinated activity do not have to be created dynamically as they are already available. For example, the coordinated motions of the limbs that

produce normal locomotion are so well specified prior to their use that they do not need to be created anew each time. Such skills may be acquired or refined by learning, but once acquired they function as prespecified components from which appropriate patterns of activity can be built. Stereotypical patterns of coordination can be highly stable and can act as strong attractors in the dynamic landscape (Kelso 1995). Dynamic coordination, by contrast, is that which cannot be prespecified because the relevant stimulus conditions or tasks are not known until they arise. For example, in figure–ground segregation, the particular set of low-level features that are to be grouped to form a Gestalt can often be determined only after the stimulus is available (Watt and Phillips 2000). Prespecified and dynamic coordination may be seen as the contrasting ends of a spectrum, rather than as a dichotomy, because novel patterns of coordination are usually built from familiar components. In some cases the novelty may be great; in others small. Possible ways of characterizing this spectrum and of placing various cognition functions on it, such as Gestalt perception, attention, and cognitive control, were major topics at this Forum.

Finally, we wish to emphasize three points concerning population coding, hierarchy, and downward causation. First, we assume that coding is more reliable at the level of the local population, than at the level of single cells. Second, interpretations computed at higher levels of the hierarchy depend upon coordinating interactions at lower levels, and higher-level interpretations serve as a context that modulates activity at lower levels. Third, as holistic Gestalt organization can have large effects on local activity, it provides clear examples of downward causation from macroscopic to microscopic levels.

What Neural Mechanisms Express Dynamic Coordination?

The dilemma mentioned earlier—how to guarantee stability in the representation of local information transmission while permitting flexibility in building situation-dependent representations—can only be solved by clearly discriminating between local information transmission and relational Gestalt organization. One possibility is temporal binding, according to which the neurons to be bound into a group synchronize their signals in time, such that neuronal identity and rate of firing in a given interval signals local meaning, whereas the temporal fine structure expresses relations to other neurons. Neurons would thus always convey two orthogonal messages in parallel: they would first signal whether the feature for which they serve as a symbol is present and then they would communicate with which other neurons they cooperate in this very moment to form a holistic representation.

One common view proposes that feature identity and salience are conveyed by labeled line and rate codes. The binding-by-synchrony hypothesis proposes that, in addition, grouping information is conveyed by the temporal fine structure of the discharge sequence which allows for the definition of precise temporal relations between simultaneously active neurons, especially if they

engage in oscillatory activity (Singer, this volume). In this way, graded relations can be encoded by varying degrees of synchrony and/or phase shifts and nested relations by phase locking among oscillations of different frequencies, known as $n:m$ locking. Such dynamically organized spatiotemporal patterns self-organize on the backbone of networks with fixed anatomical connections and change rapidly in a context-dependent way without requiring changes of the backbone, because effective coupling among anatomically connected neurons is not only determined by synaptic efficiency but also by the temporal relations between the respective discharges, and in dynamic systems these can change rapidly (Womelsdorf et al. 2007). A second hypothetical mechanism would be based on rapidly and reversibly switching synapses, such that groups of neurons that are to express integration into a pattern do so by momentarily activating a network of mutual connections. These two mechanisms are not mutually exclusive and may indeed be like two sides of the same coin: signal correlations switching synaptic strengths, synaptic strengths modulating signal correlations (von der Malsburg 1981/1994). Alternately, synapses could be actively controlled by dedicated “control units” (Lücke et al. 2008). A third mechanism would require that for each local feature to be signaled (such as a visual edge at a given position and orientation), there must be several alternate neurons, each carrying their own distinct connection patterns. By selecting among these neurons, the system could express integration of the signaled fact into this or that network. Figure–ground separation could thus be expressed by representing each possible visual element by two neurons: an *f*-neuron and a *g*-neuron. Within the figure, *f*-neurons would be active; within the ground, *g*-neurons would be.

Mechanisms for contextual gain modulation, including those underlying attention, play a major role in dynamic coordination; they amplify activity relevant to current stimulus and task contexts while suppressing activity that is irrelevant. One mechanism discussed at this Forum concerns the “windows of opportunity” for pyramidal cell spiking that are provided by synchronizing the activity of fast-spiking inhibitory interneurons (see Whittington et al., this volume). This mechanism is clearly modulatory rather than driving, and relates dynamic coordination to synchrony, rhythms, and psychopathology.

To What Extent Does Dynamic Coordination Depend on Distributed Self-organization and to What Extent on Executive “Control”?

A pervasive idea holds that dynamic coordination is the result of interpretation and strategic planning by a central executive housed in prefrontal cortex (PFC). Conceptual difficulties with this idea and evidence from several disciplines suggest, however, that coordination is primarily distributed. It arises predominantly from many local coordinating interactions between the local processors themselves, rather than simply being imposed by commands from the PFC. This does not contradict the idea that PFC may function to formulate,

From "Dynamic Coordination in the Brain: From Neurons to Mind,"

C. von der Malsburg, W. A. Phillips, and W. Singer, eds. 2010. Strüngmann Forum Report, vol. 5, series ed. J. Lupp. Cambridge, MA: MIT Press. ISBN 978-0-262-01471-7.

initiate, and monitor strategic plans, but rather emphasizes the dynamic coordination that arises through self-organization in addition to any contribution from the PFC.

As it is now well established that different areas within the PFC have distinct roles, PFC activities may themselves need to be dynamically coordinated. Thus, it is possible that basic mechanisms of dynamic coordination also operate within the PFC. The classic concept of self-organization does not include goal orientation, which is a hallmark of strategic planning. The extent to which goal orientation can arise within a highly distributed self-organized system is therefore a major issue that remains to be addressed in future research.

These issues can be seen as analogous to those concerning the organization of society. The assumption that the PFC is necessary to coordinate activity is equivalent to the assumption that without commands from a central government, humankind is left only with anarchy. Originally, “anarchy” referred to the order that arises without a ruler. That it has now come to mean the absence of order clearly reveals the common but false assumption that without a ruler there cannot be order.

Themes and Variations

In seeking a better understanding of dynamic coordination, we begin with the assumption that it takes several different forms, all sharing some basic properties. Throughout, our search for common themes is not intended to imply that there are no variations. On the contrary, one of the central goals is to understand both the themes and the variations. For example, when we ask whether there is a canonical cortical microcircuit that includes mechanisms for dynamic coordination, we assume that both themes and variations will be important.

Computational Goals of Dynamic Coordination

Coordination is a major issue within distributed computing technologies (Singh and Huhns 2005) and must often be solved dynamically at run-time, rather than being hard-wired (Moyaux et al. 2006). The same applies to perceptual organization in computer vision, and the great amount of research devoted to this in many labs clearly shows that it is a major, and, as yet, unresolved issue. Within the brain, two basic goals can be distinguished: contextual disambiguation, which uses context to disambiguate local signals, and dynamic grouping, which organizes them into coherent subsets. These goals require knowledge of what predicts what, and thus some form of learning analogous to latent structure analysis (e.g., Becker and Hinton 1992; Kay et al. 1998). At least three other goals can also be distinguished: dynamic linking, dynamic routing, and dynamic embedding. All contribute to the flexibility and power of

cognition, as well as to Gestalt perception, attention, working memory, and thought, in particular.

Contextual Disambiguation

Local activity can be ambiguous in several ways. How is it to be interpreted? Is it relevant to the current task? These ambiguities can be greatly reduced by using information from the broader context within which the signals occur. Neurophysiology, psychophysics, and neuropsychology all show context sensitivity to be very widespread, and examples will be given in later sections.

Long ago work on computer vision demonstrated a crucial role for contextual disambiguation at higher levels of object and scene perception (e.g., Clowes 1971). Many recent computational studies also show the effectiveness of context-sensitive disambiguation at the level of local feature processing (Kruger and Wörgötter 2002). Bayesian techniques or, more generally, probabilistic techniques such as factor graphs are highly relevant, as context, given in the form of conditional probabilities, can be used to reduce uncertainty of interpretation of local data.

Dynamic Grouping

At each level of processing, activity must be dynamically grouped into coherent subsets so that it can have appropriate effects at subsequent levels. Without such grouping, storing sensory patterns in their entirety in memory would be of little use, because they are unique and have no direct connection to each other. It is necessary to extract significant patterns from a given input, patterns that generalize from one situation to others. These significant patterns are to be found by grouping mechanisms. The principles of Gestalt psychology provide criteria by which this grouping may be achieved. They have been much studied and discussed in relation to what has been referred to as the “binding problem.” Previous discussions of binding use inconsistent terminology, however, so that our intended meaning must be clarified: “Grouping” means organizing a set of data into subsets upon which some operations are to be performed. Dynamic grouping and prespecified grouping must be distinguished, as explained above, because they have different requirements and constraints. Dynamic grouping must then be divided into attentive and preattentive, because some forms of dynamic grouping occur rapidly and in parallel across the visual field, whereas others are slower, serial, and require attention (Watt and Phillips 2000).

Dynamic Routing

The evidence for functional localization and automatic processing up to high levels of analysis shows clearly that much inter-regional communication is prespecified. Many cognitive functions require flexibility in communication

between brain regions (Fries 2005), however, so the fundamental capabilities and constraints of such flexible routing, its neural implementation, and cognitive consequences were major issues of discussion at this Forum.

Dynamic Embedding

Groups within groups can be described as being “embedded” or “nested.” This embedding is dynamic when it is created as required by novel circumstances, rather than being given by a prespecified hierarchy. It is most obviously relevant to linguistic syntax and metacognition (e.g., thoughts about thoughts that incorporate novel syntactic structures within novel syntactic structures).

To demonstrate more precisely what we mean by dynamic embedding, let $ABCD$ stand for variables and $abcd$ represent their realized values in a particular instance. Prespecified grouping processes specify a subset of items on which to perform some operation: grouping $ABCD$ as (AB) (CD) organizes four variables into two subsets in the same way for all values of the variables. Dynamic grouping processes also organize the items into subsets, but can do so only after the realized values are known: grouping $abcd$ as (ab) (cd) . Most models of hierarchical processing assume prespecified levels of embedding: grouping $ABCD$ as $((AB)$ $(CD))$, i.e., as two groups at a lower level and as one group at a higher level. Such networks have considerable computational power, but cannot create novel hierarchies of embedding. An important issue is therefore whether meta-cognitive processes, such as analogies, syntax, and thoughts about thoughts, require dynamically embedded groupings (Hummel and Holyoak 2005). If so, some phylogenetic and ontogenetic differences may be due to the evolution and development of that specific capacity (Penn et al. 2008). The neuronal dynamics that could embody any such capability have yet to be discovered. Using a simple marker that signals whether two elements are members of the same group, such as whether the spike trains are synchronized or not, would not be adequate (Hummel and Holyoak 1993) because different markers would be required for different levels of grouping. One suggestion is that dynamic embedding may involve synchronizing activity within different frequency bands; however, this hypothesis now requires further examination.

Dynamic Coordination at Four Levels of Organization

Corresponding to the framework of this Forum, we outline the specific issues that were relevant to each discussion group.

Evolution of Dynamic Coordination

Working on the assumption that dynamic and prespecified coordination cannot be disentangled, Balaban et al. (this volume) focused extensively on

comparative anatomy and coordination in general. Clearly, their discussions point out the need for a better understanding of the distinction between pre-specified and dynamic coordination.

If processes of dynamic coordination make major contributions to Gestalt perception, attention, and working memory, then they should be found in many, if not all, species. If, in addition, processes of dynamic coordination play a central role in higher visuospatial cognition, relational reasoning, language, and fluid intelligence in general, then some of those processes may vary greatly across species, while some may be found in only a few. Important evolutionary issues arise, therefore, at the level of microcircuits, system architectures, cognitive capabilities, and subcortical structures.

Microcircuits

What microcircuit commonalities exist across species? What role do any such commonalities play in dynamic coordination? Is there evidence of differences in the microcircuitry across species that could be relevant to differences in their ability to coordinate activity? In particular, we need comparative studies of the local circuit mechanisms identified below as making major contributions to dynamic coordination.

System Architectures

System architectures vary greatly across species, but are any of these differences of particular relevance to dynamic coordination? For example, as long-range lateral and descending connections play a major role in dynamic coordination, comparative studies of their size or organization may reveal relevant evolutionary changes.

Cognitive Capabilities

To what extent does the evolution of cognitive capabilities arise from the evolution of new forms of dynamic coordination? Are there evolutionary quantum leaps in the ability to coordinate behavior? How do Gestalt perception, attention, working memory, and executive control evolve? Are some forms of thought or relational reasoning found in only a few species, such as humans? Do they depend upon dynamic mappings between novel relational structures? To what extent does human language depend upon capabilities already present in higher forms of visuospatial cognition, and to what extent upon underlying capabilities unique to language? Does consciousness involve dynamic coordination, and do new forms of consciousness evolve as new forms of dynamic coordination emerge? Do episodic memory and insightful problem solving require any special form of coordination? Though there is relevant comparative

data on such cognitive capabilities (Grodzinski and Clayton, this volume), much more is needed.

Subcortical Structures

Do subcortical structures, such as the basal ganglia, play a major role in dynamic coordination, and if so, do their roles change in the course of evolutionary development? Such questions are of particular importance because there has long been evidence that tonically active neurons in the striatum play a central role in coordinating the distributed modular circuitry of corticobasal channels (Graybiel et al. 1994; Aosaki et al. 1995), thus providing an analogy in adaptive motor control to “binding” processes hypothesized to contribute to the coordination of perceptual processes.

Dynamic Coordination in Local Cortical Microcircuits

Much is known about the anatomical and physiological properties of local cortical microcircuits. Thus, Mehta et al. (this volume) focused on the following issues.

Canonical Cortical Microcircuits

Is there a canonical cortical microcircuit (e.g., as discussed by Douglas and Martin 2007), and, if so, what are its major variants? Commenting on evidence for the basic nonuniformity of the cerebral cortex, Rakic (2008) concludes that although cortex is generally organized into vertical columns, the size, cellular composition, synaptic organization, and expression of signaling molecules varies dramatically across both regions and species. Our working hypothesis is that, although there is variation, it is not so great as to remove all functional commonalities.

Elementary Computational Operations and Mechanisms

What elementary computational operations are performed by cortical microcircuits, and, in particular, are there mechanisms that make a special contribution to dynamic coordination? Possible mechanisms include those underlying “modulatory” as contrasted with “driving” interactions (Sherman and Guillery 1998), “contextual fields” as contrasted with “receptive fields” (Kay et al. 1998; Phillips and Singer 1997), local circuit mechanisms for gain control (Tiesinga and Sejnowski 2004), and local circuit mechanisms for coordinating phase relations between rhythmic activities (Whittington and Traub 2003). Rhythms were thought to be highly relevant because they have a special role in dynamic grouping, attention, and other relevant cognitive processes (Buzsáki 2006).

Is Feedforward Transmission Driving?

How and why does feedforward transmission (e.g., from thalamus to primary sensory regions) have such a dominant effect on activity, when it constitutes less than 10% of synaptic input (e.g., to layer IV cells)? From a functional point of view, it seems likely that this occurs because feedforward transmission is primarily driving, whereas lateral and descending connections are primarily coordinating.

Glutamate Receptor Subtypes

Are some glutamate receptor subtypes of particular relevance to dynamic coordination? *Prima facie*, NMDA receptors seem to function as highly selective gain controllers and thus could help mediate coordination within the cognitive system (Phillips and Singer 1997). This hypothesis is supported by evidence that NMDA receptor malfunction or dysregulation constitutes a crucial part of the pathophysiology of cognitive disorganization in psychosis (Phillips and Silverstein 2003).

It is usually said that NMDA receptors are relevant to cognition because of their role in synaptic plasticity and learning. This view, however, neglects all of the evidence which shows that they have large and immediate effects on ongoing activity. NMDA receptors amplify activity that is relevant to the current context and suppress that which is irrelevant. Therefore, they have a major impact on current information processing—not simply on learning. Indeed, it can be argued that their role in learning is secondary to their role in current processing (i.e., synaptic plasticity tends to record those patterns of neural activity that were selectively amplified because of their coherence and relevance to the context within which they occurred). Differences between NMDA receptor subtypes may be of particular importance because they have different temporal dynamics. For example, the 2A subtype has faster deactivation kinetics than the 2B subtype, and thus may play a special role in rapid and precise cognitive processes.

Inhibitory Interneurons

What do inhibitory local circuit interneurons and the various subtypes of GABA receptors contribute? They have a central role in generating and coordinating rhythmic activities (Whittington and Traub 2003), in enhancing attended activities (Tiensinga and Sejnowski 2004), and in suppressing those that are irrelevant. Thus, in combination with NMDA receptors, they may play a central role in dynamic coordination. This hypothesis is supported by evidence implicating GABAergic neurotransmission in the pathophysiology of disorganized cognition (e.g., as in schizophrenia; Lewis et al. 2005).

Apical and Distal Dendritic Compartments

Are some coordinating interactions predominantly mediated by apical and distal dendritic compartments? *Prima facie*, it seems that basal and proximal synapses are well placed to have a central role in driving postsynaptic activity, and that apical and distal compartments are well placed to receive and integrate inputs from the modulatory context. Modeling studies support this hypothesis (e.g., Körding and König 2000; Spratling and Johnson 2006).

Windows of Opportunity Created by Rhythmically Synchronized Disinhibition

Pyramidal cells receive strong perisomatic inhibitory input from fast-spiking basket cells, which temporarily prohibits spiking. “Windows of opportunity” for pyramidal cell spiking are provided by the periods of recovery from this inhibition. The effects of excitatory inputs to principal cells can therefore be modulated by controlling these windows of opportunity. In part, this depends upon how synchronized the inhibitory inputs are, because when they are synchronized so are the periods of recovery from inhibition. Models of gain modulation through synchronized disinhibition show that it could play a major role in attention, coordinate transformation, the perceptual constancies, and many other cases of contextual disambiguation (Salinas and Sejnowski 2001; Tiesinga et al. 2008). These models show that such gain modulation is particularly effective at gamma frequencies. Furthermore, as synchronization of the phases of these windows of opportunity within and between cortical regions could contribute to dynamic grouping, linking, and routing, they could play a central role in all four main neurocomputational functions referred to jointly as dynamic coordination. Thus, this aspect of local circuit function may play a pivotal role in future studies of dynamic coordination. Furthermore, it is also closely related to NMDA and GABAergic function, gamma and beta rhythms, and the pathophysiology of cognitive disorganization in disorders such as schizophrenia (Roopun et al. 2008).

Synaptic Assemblies

The notion of “synaptic assemblies” is outlined by Mehta et al. (this volume). It potentially provides a far richer and more dynamic concept than does that of Hebbian “cell assemblies,” and may be more centrally involved in dynamic coordination. Most importantly, it is inherently a relational concept, as it is concerned with connections between cells, rather than simply with the activity of the cells themselves. This resonates with the idea of the rapid formation, matching, and dissolution of effective network architectures as used within dynamic link architectures (Wolfgramm et al. 2008).

Dynamic Coordination in Brain Systems

Lateral and Descending Connections

Is coordination in perceptual hierarchies predominately achieved through lateral and descending connections? Our working assumption is that this is so, at least to a first approximation. If feedforward drive is the primary determinant of receptive field selectivity, then lateral and descending connections may modulate the effects of that drive so as to increase overall coherence within and between different levels of the hierarchy (Lamme and Roelfsema 2000). An analogy with Bayesian techniques may be of relevance here because feedforward pathways can be seen as transmitting the data to be interpreted, whereas lateral and descending pathways carry information about conditional probabilities that are used to help disambiguate the perceptual decisions (Körding and Wolpert 2004). As some lateral and descending connections may actively fill in missing data, these connections will need to be distinguished from those that are purely modulatory.

Contextual Modulation

What does the neurophysiological evidence for contextual modulation tell us about dynamic coordination? Many such studies have been conducted and all show that activity in perceptual areas is modulated by contextual input from far beyond the classical receptive field (Gilbert 1992; Lamme and Spekreijse 2000). Contextual disambiguation is therefore a widespread and common process in perceptual systems. Similar forms of contextual modulation may also apply to the neural activities underlying working memory and thought.

Neural Bases of Gestalt Grouping

How does neural activity discover and signal perceptual Gestalts? It is clear that temporal proximity plays a central role in determining what events are grouped, but one version of this is of particular relevance here: the hypothesis that dynamically created groupings are signaled by synchronizing the spike trains of activities that are to be processed as a whole, and by desynchronizing those that are to be segregated. This hypothesis has led to many empirical studies and much debate.

It is possible that information about different properties of specified regions of space (e.g., color, texture, depth, and motion) is linked via their reference to common spatial locations. This cannot apply, however, to the linking of information that arises from different regions of space as required to distinguish figure from ground. Psychophysical studies of contour integration and of lateral interactions between distinct spatial channels suggest that, in addition to the inputs specifying their receptive field selectivity, cells in visual cortex have

“association fields” that amplify coherent groupings. Neurophysiological studies of this possibility are therefore of central importance (Seriés et al. 2003), and a central goal for future research must be to assess the extent to which the Gestalt criteria for dynamic grouping are implemented by such connections.

Temporal Structure and Synchrony

What is the role of synchrony and high frequency rhythms in neuronal signaling? One possibility is that spike rate and spike synchronization operate in a complementary way such that salience can be enhanced by increasing either or both. Evidence for this is provided by physiological studies of brightness induction (Biederlack et al. 2006) and by psychophysical studies which suggest a central role for synchronized rate codes in figure-ground segregation (Hancock et al. 2008).

Both synchrony and high frequency (gamma and beta) rhythms have been implicated in attention, short-term memory, associative memory, sensorimotor coordination, consciousness, and learning (Melloni et al. 2007; Singer 2004; Wespatat et al. 2004). Both were major foci for discussion at this Forum (see Singer and Moser et al., both this volume). In conscious perception, for example, selective attention aids the formation of coherent object representations by coordinating the resolution of many widely distributed local competitions, and physiological studies (e.g., Fries et al. 2001) suggest that synchrony plays a role in attentional coordination. Furthermore, computational studies (e.g., Tiesinga et al. 2005) show that synchronized disinhibition could play a fundamental role in contextual gain modulation, thus extending the potential role of rhythmic synchrony to cover all four fundamental neurocomputational functions, referred to jointly as dynamic coordination.

Dynamic Linking

How is dynamic linking in object recognition and graph matching generally achieved? As mentioned earlier, pattern recognition is best based on assessment not only of the presence of figural elements (feature-based recognition) but also of the correctness of their spatial or temporal relations (correspondence-based recognition). Such recognition amounts to graph matching; that is, to the comparison of a stored model of the pattern and the potential instance to be recognized, where both have the form of a graph with feature-labeled nodes and links to represent neighborhood relationships. During graph matching or recognition, dynamic links between corresponding nodes must be established, as modeled in Wolfrum et al. (2008).

Neuronal Bases of Attention

What are the neuronal bases of attention? This issue is of central importance because selective enhancement of the relevant and suppression of the irrelevant

are major forms of dynamic coordination. Many specific issues concerning attention arise: How is attention related to synchrony? Do stimulus-based and task-based influences on salience have essentially similar kinds of effect? What determines the balance between them? What is the current status of the neurobiological evidence for biased-competition models, and what further evidence is needed? Why do attention and working memory have such limited capacity? Did evolution get stuck in a dead end somehow, or do these limitations reflect fundamental underlying computational constraints? What can we learn from neurological disorders of attention? Do Bayesian models of salience and visuospatial attention clarify these issues?

One major conclusion was that biased competition, in the form of normalization models (Reynolds and Heeger 2009), is currently best able to explain the modulatory effects of attention at the local circuit level.

Dynamic Routing

How flexible is the routing of information flow between cortical areas, and how is that flexibility achieved? Communications between brain regions are largely pre-specified by the cortical architecture. For example, V2 responds to activity in V1 and affects activity in V3 and V4, regardless of the current task. Some communication must be flexible, however, as some tasks and circumstances require novel routing of information (e.g., from sensory to motor areas). The central proposal discussed at the Forum was that such dynamic routing is achieved by coordinating the phases of transmission and receptivity across brain regions (Fries 2005; Womelsdorf et al. 2007), and much evidence was thought to support this proposal.

The Role of Prefrontal Cortex

In what ways does the PFC have a special role in coordination, and how well is activity coordinated by distributed processes of self-organization independently of any contribution from PFC? Does activity in the various components of PFC need to be coordinated, and, if so, how is that achieved?

Dynamic Coordination in Cognition and Behavior

The computational capabilities provided by dynamic coordination at the neuronal level have major consequences for Gestalt perception, attention, working memory and thought, and thus for learning as well. We assume that they are crucial to language, but have focused more on their relevance to visuospatial cognition.

Gestalt Perception

What do psychophysical studies tell us about Gestalt organization in perception? Does it operate within all perceptual pathways? How does it develop?

Evidence indicates a major role for Gestalt organization in dimensionality reduction and contextual modulation (Kovács, this volume). However, although knowledge of contextual relationships in image structure begins to be acquired in the first few months of life, evidence shows that it continues to develop over several years in human vision (Doherty et al. 2010; Kovács 2000).

Contextual Disambiguation

To what extent does cognition use context to reduce local uncertainty? A vast literature is available showing context sensitivity in a wide range of cognitive domains. Our focus was on consistency constraints in the perception of edges, surfaces and objects, because they provide paradigmatic examples of such coordinating interactions.

Structured descriptions are fundamental to cognition. It is obvious that language depends on descriptions generated by a grammar, which, though finite, enables us to produce and understand infinitely many novel sentences. Analogous “grammars” have been used to interpret line drawings of opaque polyhedra (solid objects with straight edges) in computer vision (e.g., Clowes 1971). Clowes demonstrated formally how local ambiguity can be reduced by using “co-occurrence restrictions,” and showed their relevance to human vision using demonstrations of ambiguity, paraphrase, and anomaly, just as Chomsky did for language.

Figure 1.2a is ambiguous in several ways: it can be seen as a tent or as a book standing up (i.e., the central vertical edge can be seen as folding either outwards or inwards). Changes of viewing angle reduce this ambiguity (Figure 1.2a, b), thus demonstrating how context resolves local ambiguity; the arrow and fork junctions that are ambiguous in Figure 1.2a are disambiguated by the other junctions in Figure 1.2b and Figure 1.2c.

The equivalence of one interpretation of Figure 1.2a with Figure 1.2b and of another with Figure 1.2c shows that perception distinguishes between scenes and images, just as language comprehension distinguishes between deep and surface structures. Figure 1.2d is evidence that objects are seen as structures, and not merely as members of a prespecified category. Though you had never seen a tent with such a hole before, you have now, and that required you to create an internal description of a new distal structure and to map image fragments to parts of that structure.

The local ambiguities in such pictures and their resolution using contextual constraints are made explicit in Figure 1.3, which is based on the computer vision system developed by Clowes (1971). His analysis is deterministic, but would be more relevant to human perception if formulated in Bayesian or other statistical terms.

The arcs linking compatible local interpretations in Figure 1.3 are emphasized here because they clearly show essential properties of coordinating interactions. They resolve local ambiguity by applying co-occurrence constraints

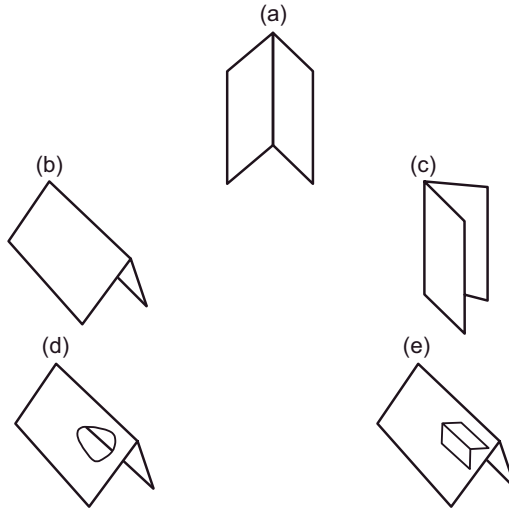


Figure 1.2 Local ambiguities disambiguated by context. The central edge of (a) can be seen as folding either inwards or outwards, because the arrow junction at the top and the fork junction at the bottom of that edge are ambiguous. This ambiguity is reduced in (b) and (c), where the context provided by the other junctions makes one interpretation more probable than the others. Drawing (d) shows that interpretation is not merely categorization, but relates structured descriptions in the picture and scene domains. Drawing (e) shows how context can reveal another possible (though improbable) interpretation of (a), i.e., as a chevron-shaped hole through which part of a more distant edge can be seen.

to select coherent subsets of the local interpretations that are made possible by the locally available data, thereby playing a major role in determining the Gestalt created. Assuming that different local interpretations are conveyed by different neuronal populations, this clearly demonstrates our assumption that macroscopic organization can have large effects on local microcircuit activity. Eleven vertices can be seen in the object in Figure 1.3. If each has at least three alternative interpretations, when considered independently, then there would be at least 3^{11} possible combinations. From that host of possibilities, coordinating interactions select the few that are coherent when taken as a whole. Our assumption is that this is done by selectively amplifying and grouping the neuronal activity representing local interpretations that are coherent.

These demonstrations do not imply that global coherence is necessary for perception. The importance of global coherence is often thought to be demonstrated by impossible figures. Take, for example, Figure 1.4, which is known as the devil's pitchfork because it has three round prongs at one end but two square prongs at the other. Such drawings do bring global inconsistencies to our attention, but they also show that contextual constraints are primarily local. We clearly see the various possible local interpretations, even though they are not globally coherent. Indeed, such drawings are only perceived as impossible

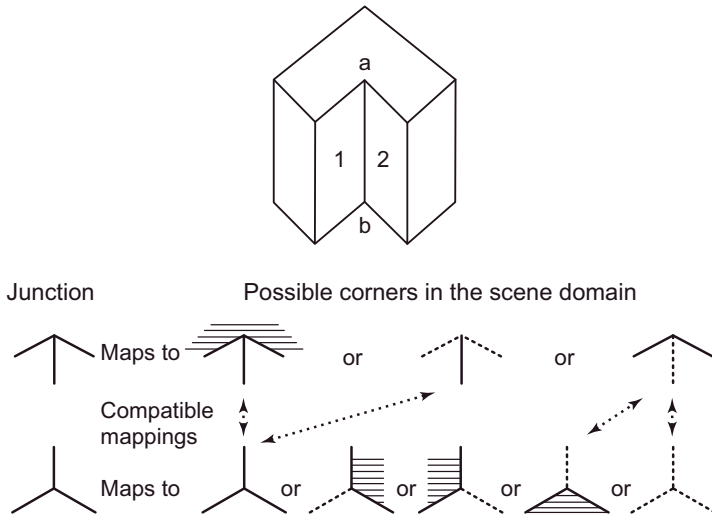


Figure 1.3 Example of how the Clowes' (1971) algorithm used contextual constraints to interpret line drawings of solid objects with plane surfaces and straight edges. The drawing at the top shows two further interpretations of junctions in Figure 1.2a. Here, regions 1 and 2 can be seen as being on two separate limbs of a block with a square column cut into one corner. The two upper edges of junction “a” would then be convex, and the edge joining “a” and “b” would be concave. Regions 1 and 2 can also be seen as the two sides of a square column rising into the top corner of a ceiling. The two upper edges would then be concave, and the edge joining “a” and “b” would be convex. Locally possible interpretations of the arrow junction at “a” and the fork junction at “b” as corners are shown below, with hatching indicating a surface that disappears behind an edge, solid lines indicating convex edges, and dotted lines concave edges. When junctions share an edge, as do “a” and “b,” only certain combinations are compatible. Surfaces are regions surrounded by a set of compatible mappings from junctions to corners. Objects are a compatible set of surfaces.

because the local consistencies map so strongly on to partial three-dimensional structures. Without that, the drawing would simply look like an arbitrary collection of lines. Thus, such impossible figures can be thought of as showing the primacy of local over global coherence. An alternative perspective is to view this demonstration as showing that all elements of the whole figure may be bound if attended to as a whole, even though those elements are not seen as forming a single coherent structure.

Attention and Working Memory

What do cognitive studies of attention tell us about dynamic coordination? They seem to be of central importance, and show the conditions under which relevant signals can be enhanced and irrelevant suppressed. Many issues remain to be resolved, however, before we understand why and how this occurs. Cognitive studies support biased competition models of attention (Duncan

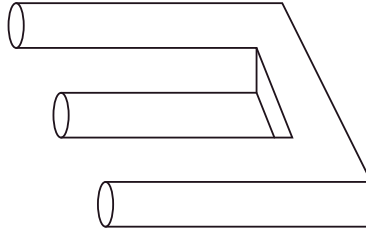


Figure 1.4 The devil's pitchfork. Although there is no globally coherent interpretation as a three-dimensional object, local consistencies are strong enough to produce partial perception of three-dimensional structure. This can be thought of as showing the primacy of local over global coherence.

2006), and their more recent formulation as normalization models (Reynolds and Heeger 2009), but how can that support be clarified and strengthened? Are competitions resolved by a single central executive controller or locally, as suggested by independence of their resolution across modalities (Hupé et al. 2008)? What is the role of attention in creating coherent percepts? How are conflicts between stimulus-driven effects on salience and task-driven effects on relevance resolved? Attention is needed to bind features that are not already bound pre-attentively (Treisman 1999). However, some dynamic grouping occurs pre-attentively (Watt and Phillips 2000). So, what determines whether or not attention is required? What do cognitive studies tell us about how and why the capacities of attention and working memory are so limited? Is working memory a major source of contextual influence on current processing, and how convincing are the neuronal models proposed as explanations of those influences (e.g., Kerns et al. 2004)? Some working memory studies have been interpreted as indicating the importance of dynamic grouping that is implemented by neuronal synchrony (e.g., Luck and Vogel 1999), but how strong is that evidence? Can theories of attentional selection based on dynamic link architectures (e.g., von der Malsburg and Schneider 1986) be tested psychophysically? Although these issues were discussed in depth at this Forum (see Engel et al., this volume), it is clear that they provide major goals for future research.

Reasoning

Are there any theories of relational thinking that are both psychologically and neurally plausible? Does relational reasoning require special forms of dynamic coordination? Though Hummel and Holyoak (2005) imply that it does, they propose a model in which the critical binding of roles to fillers is achieved using the same form of synchronization thought to underlie dynamic grouping in general.

The use of context to resolve local ambiguities, as shown in Figure 1.3, is an example of an approach to constraint satisfaction that can generally be applied to reasoning and problem solving. As the network of constraints grows,

the possibility of checking for global consistency becomes less feasible; thus it is unlikely that fully global consistency can be established for rich knowledge structures, such as those in human cognition. The best that can then be done is to maximize local consistencies and minimize local inconsistencies (Dechter 1992).

Central Executive Functions

Do cognitive studies imply the existence of a central executive, and, if so, what are its specific responsibilities and capabilities? What do they tell us about its internal organization, and about the way in which the activities of its separate components are coordinated? Do these cognitive studies lead to the same conclusions as physiological investigations of PFC? Among the participants at this Forum, there was wide agreement that executive processes (e.g., cognitive control) play an important role in coordinating other activities, and are themselves highly dependent on the dynamic coordination of their various components.

The Distinction between Multimodal Coordination and Multimodal Sensor Fusion

How is multimodal coordination related to sensor fusion? Multimodal “integration” has been extensively studied, but the term “integration” is ambiguous. It often refers to sensor fusion, where inputs from different modalities are combined so as to produce an output signal within which the separate contributions are not distinguished. This is not the same as multimodal coordination, however, because in that case the separation between modalities is maintained even though the signals that they generate are coordinated.

Bayesian Analyses

Are Bayesian analyses relevant? Discussed in depth at this Forum (see Triesch et al., this volume), we think that they are useful because they account for much behavioral data. In addition, they compute posterior probabilities using data and priors in fundamentally different ways that match our distinction between driving and modulatory interactions, particularly if the Bayesian inference is implemented so as to amplify transmission of predicted signals (e.g., Kay and Phillips, submitted; Spratling 2008).

The Relation between Cognition and the Precise Temporal Structure of Neural Activity

What is the status of the evidence relating synchrony and brain rhythms to cognition? Addressed by both Tallon-Baudry and Engel et al. (this volume), it was agreed that such studies may be pivotal, potentially relating fundamental

cognitive functions to neural activity at the population and local circuit levels. Furthermore, as noted next, they also provide a link to psychopathology.

Disorders of Dynamic Coordination

What do failures of dynamic coordination in various forms of psychopathology tell us about its role in normal brain function? What do studies of the pathophysiology of those conditions reveal about its neuronal mechanisms? Silverstein (this volume) shows that psychopathology can provide a rich source of evidence linking behavior and cognition to neural activity at population and synaptic levels. In support of this view, consider the findings of Roopun et al. (2008). They show how the windows of opportunity for pyramidal cell spiking at gamma and beta frequencies depend on GABAergic and NMDA receptor activity, and they relate those mechanisms to changes in cortical dynamics that occur in schizophrenia. Furthermore, they provide evidence that there are variations in those rhythms and mechanisms across cortical regions, so we now need to discover the computational, cognitive, and psychopathological consequences of such variations. Finally, it was agreed that the theories of Friston (1999) and of Phillips and Silverstein (2003) have much in common. Friston suggested that the “disconnection” theory of schizophrenia could be referred to as the “dysconnection” theory. This label would then apply to both theories if it is understood that it is not connections in general that are malfunctional, but coordinating connections in particular.

Conclusion

Dynamic coordination makes fundamental contributions to brain function and cognition. Much is already known about why and how it does so, but much more remains to be discovered. Many concretely specified issues and predictions have now been identified that can be investigated using well-established or recently developed techniques. Over the coming decades, we expect substantial advances in our understanding of dynamic coordination in the brain, thereby building reliable bridges between neurobiological and psychological perspectives on mental life.