

In search of common foundations for cortical computation

William A. Phillips

Center for Cognitive and Computational Neuroscience, Department of Psychology, University of Stirling, FK9 4LA, Scotland, United Kingdom
Electronic mail: wap1@forth.stir.ac.uk www-psych.stir.ac.uk/~wap

Wolf Singer

Max Planck Institute for Brain Research, D-60496, Frankfurt, Germany
Electronic mail: singer@mpih-frankfurt.mpg.de

Abstract: It is worthwhile to search for forms of coding, processing, and learning common to various cortical regions and cognitive functions. Local cortical processors may coordinate their activity by maximizing the transmission of information coherently related to the context in which it occurs, thus forming synchronized population codes. This coordination involves contextual field (CF) connections that link processors within and between cortical regions. The effects of CF connections are distinguished from those mediating receptive field (RF) input; it is shown how CFs can guide both learning and processing without becoming confused with the transmission of RF information. Simulations explore the capabilities of networks built from local processors with both RF and CF connections. Physiological evidence for synchronization, CFs, and plasticity of the RF and CF connections is described. Coordination via CFs is related to perceptual grouping, the effects of context on contrast sensitivity, amblyopia, implicit influences of color in achromatopsia, object and word perception, and the discovery of distal environmental variables and their interactions through self-organization. Cortical computation could thus involve the flexible evaluation of relations between input signals by locally specialized but adaptive processors whose activity is dynamically associated and coordinated within and between regions through specialized contextual connections.

Keywords: cell assemblies; cerebral cortex; context; coordination; dynamic binding; epistemology; functional specialization; learning; neural coding; neural computation; neuropsychology; object recognition; perception; reading; self-organization; synaptic plasticity; synchronization

1. Introduction

The possibility of common foundations for cortical computation was first discussed by the current authors (Phillips is a psychologist, Singer a neurophysiologist) in 1980. We had collaborated in the early 1970s, comparing single unit activity in cat lateral geniculate nucleus with the ability of humans to detect the appearance and disappearance of elements in random dot patterns (Phillips & Singer 1974; Singer & Phillips 1974). This helped convince us that psychophysics and neurophysiology could combine fruitfully and in detail. Since then we had not met for some years and the neurophysiologist asked what the psychologist's current interests were. The ensuing conversation went roughly as follows:

Psychologist: Well, I am mainly interested in the fundamental differences between different cognitive domains. For example, what are the basic differences between sensorimotor systems and the higher conceptual systems? Then, within the conceptual systems, what are the basic differences between visuospatial processing and verbal processing?

Neurophysiologist: But why are you emphasizing differences? The cortical algorithm is everywhere the same.

Psychologist: Well, if that is so it is very interesting, but from the psychological point of view there certainly seem to

be some major differences. Consider learning and memory, for example. Information storage in the sensory systems is of very short duration, less than a second in visual sensory storage, whereas once it is put into a schematic conceptual form information can be voluntarily maintained for many seconds in short-term memory (STM) and can be learned and stored indefinitely in long-term memory (LTM).

Neurophysiologist: But there is also long-term plasticity in sensory systems, both during development and later. The receptive fields of cells in primary sensory cortex depend upon the stimulation they get during development. These use-dependent modifications of synaptic transmission can also occur in adults.

Psychologist: Yes, of course, such effects are well established, but that is a quite different kind of learning.

Neurophysiologist: Well, is it? Why do you suppose that learning and processing in the sensory cortex are fundamentally different from learning and processing elsewhere in the cortex? Perhaps they are very similar, and from the neurophysiological point of view that's how it seems.

We are still searching for answers to the questions raised by this discussion. Are there information processing operations that are common to different cortical regions and different cognitive subsystems, and if so, what are these operations? Why are they useful? How are they implemented by cortical processes? Of course different cognitive

functions are performed by different cortical regions and at different levels of organization, but all regions of the neocortex share a common basic internal organization; it is because of this predominant homogeneity that it is also called isocortex. Computational capabilities of general utility may therefore arise from this common design. This target article is concerned with what those capabilities might be and how they arise from cortical structures and processes.

The organization of cognition into distinct subsystems is even more firmly established now than it was twenty years ago. This does not imply differences in the information processing operations that they perform, however, because subsystems may differ in the information upon which they operate, but not in the operations that they perform upon that information. Many cognitive subsystems are distinguished from each other by reference to the information on which they operate, but it is also likely that some cognitive functions require special information processing capabilities. These include episodic memory and working memory, intentional representation (i.e., processes that distinguish between representation and referent), and the creative aspects of language and long-range strategic planning. Higher cognitive functions such as these are central to human mental life and depend to a large extent upon cortical activity. These functions may not arise in any simple way from basic capabilities common to cortex in general, however, because (1) intentional representation and language are not characteristic of mammals in general but are restricted to one or at most a few; (2) in contrast to skills, episodic memories cannot be acquired in the absence of the hippocampus (Squire 1992) and may require special computational capabilities (McClelland et al. 1995); and (3) the ability to dynamically create more than one level of grouping within the same set of units, such as [(AB)(CD)], may involve special computational problems (Fodor & Pylyshyn 1988; Hummel & Holyoak 1993). Thus, our working assumption is that some cognitive functions require special capabilities in addition to those that are common to cortex in general. Furthermore, although we take the abilities that are provided by the common foundations for granted, they are crucial to the sensory, perceptual, and motor skills on which our daily lives depend.

1.1. Functional specialization and contextual integration as basic aspects of cortical computation

A vast body of findings shows that different cortical regions and different cells within regions transmit information about different things. Discussions of how the activity of these specialized processors can be coordinated have an equally long history, but this aspect remains much less well understood. The particular form of integration with which this target article is concerned is the one that arises from a myriad of local coordinating interactions between pyramidal cells within and between cortical regions. This does not deny that the music of the hemispheres might be guided by some kind of conductor, but it does imply that integration can be achieved, at least in part, through local interactions between the players themselves. Musicians have two different sources of information that they normally use in two different ways. They have the score to tell them what to play, but they also watch and listen to each other to

determine exactly when and how loudly to play it. The local processors that we postulate also have two classes of input. One is the receptive field input that tells them what features to signal, and the other is contextual input from the concurrent activity of other processors that is used to determine exactly when and how confidently to signal the features for which they have evidence.

A simple, general, and precise framework for describing functional specialization in neural systems is provided by the adaptive filter formalism (Carpenter 1989). The basic idea is: the strengths of the synapses that mediate receptive field input perform a selective filtering operation that can be adapted through experience to better match the environment and tasks to which the system is exposed. Filtering is necessary for at least two reasons: (1) the amount of sensory data to be processed is so great that predictive relationships can only be found after dimensionality reduction; and (2) different information is relevant to different purposes.

We now need to add contextual integration to such formalisms. Filtering is useful because it contributes to the more general goal of making good predictions. Predictive relationships of varying degrees of complexity are richly embedded within the input to the cortex, across both space and time, and the discovery and use of these relationships is a major goal of cortical computation at all stages and levels of processing. The integrative interactions that we hypothesize can be thought of as using these predictive relationships to produce patterns of activity that are coherent both within and between various streams and levels of processing. There is evidence that this involves synchronizing the activity of dynamically specified subsets of cells, using special synchronizing contextual connections that influence the probability that the target cells fire at any moment (see, e.g., Engel et al. 1992; Singer 1990; 1993; 1994a; 1995; Singer & Gray 1995). A crucial aspect of this form of integration is that context affects activity without corrupting the information transmitted by that activity about the cell's receptive field input. Here we summarize the evidence for synchronization and for contextual connections, analyzing computational capabilities that arise when cortical processors receive local contextual inputs that can be used to guide both learning and processing.

In the remainder of section 1, we outline the issues and hypotheses to be discussed. Section 1.2 gives an informal outline of the possibilities that arise when local cortical processors coordinate their activities by using specialized contextual inputs to form synchronized population codes and to guide learning. Section 1.3 relates the codes and processes we propose to other aspects of cortical function. Section 1.4 reviews prior proposals using synchronized population codes and contextual guidance. Section 1.5 summarizes the main hypotheses that we expect to be controversial. Section 2 outlines arguments for and against the hypothesis of common foundations for cortical computation. Section 3 specifies the goals of contextual guidance using formal concepts from information theory and multivariate statistics. Computational studies showing the basic capabilities of simple networks built from local processors with contextual guidance are described. Section 4 outlines evidence from neurobiology for contextual guidance and synchronized population codes. Section 5 posits the relevance of these hypotheses to psychological issues, describing evidence that is already available from behavioral

studies as well as ways in which our suggestions can be further tested and developed by such studies. Finally, section 6 discusses a few of the issues that arise from these hypotheses.

1.2. Coordinating activity through the local contextual guidance of learning and processing

A starting point for our approach is the hypothesis that although cortical circuits are constrained to operate only upon the information that is locally available to them, coordination of their activity with what is going on elsewhere is central to their computational role. This is possible because they receive locally specific contextual input from other processors (but directly only from a tiny fraction of the other processors in the cortex as a whole). The contextual input is used to enhance selectively the transmission of that information in the processor's receptive field (RF) input that is coherently related to the context. Networks of such processors therefore tend to transmit sets of signals that as far as possible maximize their mutual coherence. As a signal transmits information only if it varies, we call this capability the "maximization of coherent variation." Useful consequences that follow from it are discussed in section 1.2.3.

The usefulness of the ability to organize distributed patterns of activity into coherent groups is widely acknowledged in discussions of the "binding problem." This problem would be solved if cells currently forming a coherent group synchronized their spike trains to within a few milliseconds. This possibility was proposed by Milner (1974) and has long been advocated on both theoretical and biophysical grounds (von der Malsburg 1981). Neurophysiological evidence (outlined in sect. 4) now suggests that the spiking activity of cortical neurons can be synchronized to within a few milliseconds in a way that is appropriate to the prevailing context and includes synchronization between neurons in different streams of processing and between neurons at different stages of processing.

Synchronization would be an effective signal for grouping because inputs to pyramidal neurons are summed much more effectively if they are synchronized (Abeles 1982, 1991; Bernander et al. 1991). Synchronization is an inherently relational signal because it depends upon temporal relations between inputs from separate sources. Thus, unlike the more commonly studied rate and place codes, it is not defined upon the signals produced by individual cells, and will not be revealed by studies of single cell activity.

A major feature of the work on synchronization is that it suggests the existence of specialized cortico-cortical synchronizing connections that modulate postsynaptic activity without corrupting the information transmitted about the receptive field features to which the cell is selectively sensitive (Engel et al. 1991b; König et al. 1993; Löwel & Singer 1992; Munk et al. 1992). In other words, these connections help determine exactly when a cell fires but they do not change the feature signaled by that activity. To explain how this is possible, it is often suggested that the synchronizing connections influence the phase but not the amplitude of oscillatory outputs produced by the local processors (e.g., Hummel & Biederman 1992; Schillen & König 1994; Shastri & Ajjanagadde 1993). Though useful, this may not be a general solution, however, to the problem of combining both feature and grouping information within

the same signal (see, e.g., Nelson 1995). First, although oscillations are likely to play a major role in synchronization, they are not necessary because even single impulses can be synchronized. Second, there is evidence that synchronization can occur without oscillations (König et al. 1995). Third, there is doubt as to the generality of oscillations in the normal functioning of cortex (e.g., Bair et al. 1994; Tovee & Rolls 1992; Young et al. 1992). The computational studies described in section 3 show how contextual inputs can guide processing without corrupting the transmission of RF information, and they do this in a way that does not require oscillations even though it is compatible with them.

Another major focus of this paper is the possibilities that arise for learning when processors receive local contextual inputs. The computational studies outlined in section 3 show that idealized local processors with contextual inputs can discover those receptive field features that are predictably related to the context within which they occur and can discover the predictive relations between them: In addition to learning the associations between features, the local processors can preferentially discover those features that are associated.

1.2.1. Local processors with contextual guidance. Local processors with contextual guidance receive a set of receptive field (RF) and contextual field (CF) inputs (Fig. 1a). These processors are intended to be loosely analogous to local cortical circuits. In relation to the RF input, they act as filters that transmit information about the RF features to which they are selectively sensitive. This selective sensitivity is specified by the strengths, W , of the synaptic connections that mediate the RF input. In addition, the probability that they transmit information about any RF feature at any moment is increased if that feature is as predicted by the context and decreased if it is incompatible with that prediction. The predictions are specified by the CF inputs as mediated by the strengths, V , of their synaptic connections. A crucial aspect of this form of processing is that the predictions are not confused with the RF evidence. The role of context is not to impose its predictions upon the processor, but to emphasize those outputs for which there is RF evidence and which are also coherently related to the context. This can be done by using context to influence the confidence with which decisions are made on the basis of the RF evidence and to synchronize coherent outputs.

The strengths of the synapses (W and V) are not permanently fixed, but can change so as to better adapt the detailed operations performed by the processors to the statistical structure of the inputs they receive. We hypothesize that providing local processors with contextual input enhances the learning of which they are capable. Major issues to be discussed accordingly concern what the goals of this learning could be, and by what synaptic modifications they may be achieved.

As illustrated by the width of the arrows in Figure 1, processors are assumed to have fewer outputs than RF inputs. This reflects the long-standing hypothesis that a major goal of sensory and perceptual processes is recoding to reduce redundancy (Atick 1992; Atick & Redlich 1990; 1993; Attneave 1954; Baddeley & Hancock 1991; Barlow 1959; 1961; 1972; 1989; Barlow & Foldiak 1989; Foldiak 1990; Li & Atick 1994; Linsker 1988; Redlich 1993). The underlying idea is that the flood of data to be processed can be reduced to more manageable amounts by using the

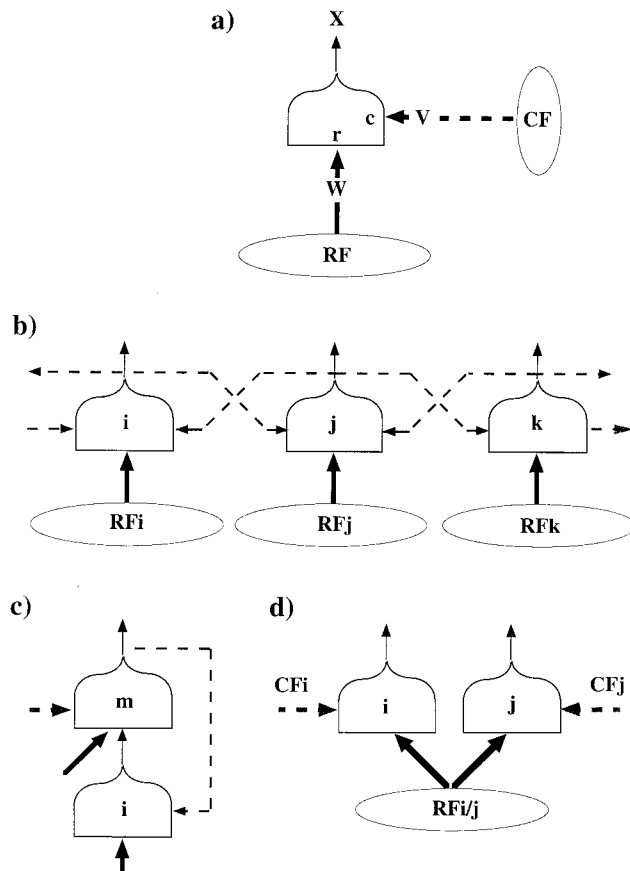


Figure 1. Local processors with contextual guidance and some of the network architectures that can be built from them. Solid black lines show receptive field (RF) inputs; dashed grey lines show contextual field (CF) inputs. (a) A local processor receives input from a receptive field vector, RF, via a vector of synaptic strengths, W , and from a contextual field vector, CF, via a vector of synaptic strengths, V . Each element of the two input vectors is multiplied by its particular synaptic strength. These are summed to give the integrated RF and CF inputs r and c , which are then used as specified by the transfer function to determine output probability. For simplicity, the synaptic strength vectors and the integrated inputs are omitted in the examples of possible architectures shown in (b), (c), and (d). (b) Three processors, i , j , and k , with nonoverlapping RFs are linked by contextual connections. Large arrays of processors could be linked by such lateral contextual connections within and between cortical regions. (c) A processor, i , receives contextual input from a processor, m , to which it contributes RF input. The processor m also receives both RF and CF input from other sources, however, so its activity is only partially determined by the input it receives from i . Descending feedback pathways in the cortex could include such contextual connections. (d) Two processors, i and j , receive RF input from a common source, but receive contextual input from different sources. As a result of these differing contextual inputs, the processors i and j could learn to extract different aspects of their common RF input, thus establishing appropriate functional specializations.

statistical structure in the data to recode the information it contains, with frequent input patterns being translated into codes that contain much less data than the patterns themselves. If the computational goals can be clearly specified, such as by using information theory, then rules for learning can be derived from those goals (Intrator & Cooper 1995a). An important limitation of recoding as a goal for cortical

computation is that ultimately it is subordinate to the goal of associative learning. There would be no point in recoding information about variables that have no relation to anything else known to the system. Proponents of recoding to reduce redundancy therefore usually see it as being preparatory to associative learning (e.g., Barlow 1993). A distinctive advantage of the processors that we propose here is that they are not forced to transmit whatever RF variables carry the most information, but can selectively discover those that are associatively related to the context in which the processor operates.

1.2.2. Architectures that can be built from local processors with contextual guidance. Many different network architectures can be built from such processors. Felleman and Van Essen (1991) review a great deal of evidence concerning the overall system architecture within which local cortical processors operate. They distinguish three broad classes of connections between cortical regions: ascending feedforward connections, descending feedback connections, and lateral connections between regions that are at approximately the same stage of processing. The ascending projections from one stage to the next are localized so that neurons receive their primary ascending inputs from a small subset of neurons at the preceding stage, with different local groups of neurons receiving input from different subsets. Thus, many distinct streams of processing project through a few stages in converging and diverging ways, with primary feedforward connections distinguishable from lateral and descending connections.

The local processors hypothesized here are broadly compatible with such an architecture. The ascending connections could provide much of the RF input, and both the descending and the lateral connections could include CF input. Furthermore, within cortical regions there are many distinct streams of processing; these are linked by long-range horizontal collaterals that could transmit synchronizing contextual information. Mutual contextual guidance of this sort is shown in Figure 1b; it could link distinct streams of processing both within and across cortical regions. Another possibility is that contextual inputs could be received from processors to which the processor concerned contributes RF input. This is shown in Figure 1c and would allow activity to be coordinated across different stages of processing, as well as within stages. Finally, Figure 1d shows that one set of RF signals can be transmitted to separate processors with different contextual fields. Each processor will emphasize the RF information that is relevant to its context, thus enabling different processors to extract different aspects of the same RF activity.

The patterns of connectivity shown in Figure 1 are not mutually exclusive and can be combined in various ways. Furthermore, separate modules with recurrent internal RF connectivity could be linked by CFs that coordinate their activities. We assume that genetic constraints play a major role in determining patterns of RF and CF connectivity, and that the CF inputs, like RF inputs, are specific to the role of each local processor. For example, at early stages of the analysis of a visual scene, other parts of the scene might provide a useful context, whereas at later stages of processing, information from other modalities might provide a more appropriate context. If local cortical processors do receive specialized contextual inputs as proposed, then principles or heuristics for determining where they should

come from will be an important issue. Tononi et al. (1994; 1996) have shown that the overall pattern of cortical connectivity balances functional integration, produced here by the contextual inputs, against functional segregation, produced here by the RF filter functions. This produces a system with high complexity, high computing power, and the ability to use context to “go beyond the sensory information given” in an appropriate way.

1.2.3. Computational capabilities of networks built from local processors with contextual guidance. (1) The most important capability that arises in relation to processing is that context will increase the probability that mutually coherent subsets of units will be active at any moment. That is, they will tend to produce synchronized population codes. It has already been argued in detail elsewhere (Singer 1990; 1993; 1994) that such codes have important advantages: They are flexible because they are created dynamically; they are fast because in the limit all that needs to be synchronized are single spikes from each of the cells to be grouped; they can signal many different patterns because at different times each cell can be part of different groups; they do not compromise the meaning of the signals to be grouped; and, finally, they transmit appropriately structured patterns of activity rather than just arbitrary or unstructured labels (Phillips 1997; Phillips et al. 1995a).

A layer built from local processors with contextual guidance therefore produces patterns of activity where the RF filter functions ensure that the individual signals are justified by the RF input and the contextual connections maximize their coherence as a group. This implies that in the case of perceptual grouping, for example, the Gestalt criteria are embodied in the CF connections between the entities that are grouped. This predicts that synchronization of active cells in the visual system should reflect Gestalt principles of grouping. Evidence that this is so will be outlined in section 4.

One way to see what is implied by dynamic flexibility is to note that if these processes are common to cortex then the inputs to each area will themselves be organized by the grouping processes operating in the areas generating those inputs. The RF filter functions will therefore not operate upon rigidly fixed data bases but upon ones that are already organized so as to emphasize coherent subsets of data within the RF. As those coherent subsets can occur within each RF in a very large number of different ways, this enables the receiving cells to respond appropriately to many more inputs than they could do without the dynamic grouping. In addition, there is evidence that cortex receives inputs that are dynamically grouped in the thalamus (Sillito et al. 1994) and retina (Neuenschwander et al. 1996).¹

(2) In addition to grouping related features that are clearly evident in the RF input, local contextual information could improve the perception of features that are weak or ambiguous.

(3) An important capability that arises in relation to learning is that by using contextual input local processors can become selectively sensitive to those variables within their RF input that are predictably related to that context.²

The possibility of using relations between separate data-sets as a basis for self-organization is illustrated in Figure 2. As an example of this approach, Becker and Hinton (1992) show how stereo-depth can be discovered by using the mutual information between separate streams of process-




Figure 2. The aim of learning (based upon an unpublished diagram by Geoff Hinton). Data from an input image are processed by many parallel streams, of which only two are shown here. One part of the input space is represented by the oval on the left, with just one data point of many being shown. Another part of the input space is represented by the oval on the right, with more of the data points being shown. Predictive relationships exist between separate input data-sets but these relations will often be between higher-order functions rather than between pairs of individual data points. The aim of learning is to find mappings into new spaces that make the predictive relationships easier to compute. Predictions are here shown going in only one direction but they may also be reciprocal. In our approach, but not in that of Becker and Hinton (1992), the predictions are used to guide the short-term processing as well as the learning.

ing that receive inputs from neighboring patches of the image that are independent except for having the same disparity. Stone and Bray (1995) and Stone (1996) have also shown how coherence across time can be used to learn invariances and other salient visual parameters.³ A general epistemological argument for this approach is that predictive relationships between diverse data-sets must depend in some way upon their distal origins. Discovering those relationships will therefore reveal distal variables and interactions within the proximal data.

(4) Processors with contextual inputs from different sources, as shown in Figure 1d, will become selectively sensitive to those aspects of their inputs that are relevant to those contexts. Thus this will help create appropriate functional specializations and will entail appropriate generalizations because the outputs of each processor will generalize across irrelevant dimensions of RF variation. We must recognize facial expressions despite variations in personal identity; we must recognize individual faces despite variations in facial expressions. Variables that are crucial to one goal may be irrelevant sources of noise to another. This problem would be solved if different cortical regions have a selective sensitivity to just those variables that are relevant to their role, and the evidence suggests that this is how face perception is organized (Bruce 1988); but how can local

processors know what is relevant? Genetic specification cannot be the whole answer because some functional specialization is established through learning. Context could contribute to such learning by guiding RF selectivity to the relevant variables. For example, regions receiving RF input from visually perceived faces and CF input from regions concerned with evaluating emotional states could then learn to become selectively sensitive to just those variables in face images that are predictably related to emotional expression.

1.3. Relation to other aspects of cortical computation

The CF inputs that we postulate are not equivalent to inputs from “beyond the classical receptive field” in general. Many investigations show that the response of cortical cells to their preferred stimulus is suppressed by the presence of similar stimuli in the surround (e.g., Allman et al. 1985; Blakemore & Tobin 1972; Knierem & Van Essen 1992; Nelson & Frost 1978). These effects are not owing to the CF inputs that we propose because they are not concerned with producing coherent patterns of activity across multiple feature detectors, but with using information about the surround to suppress responses that do not differ from that surround. This is quite different from the use of contextual predictions to increase the probability of signals that agree with those predictions. It is more appropriate to view the subtraction of activity that is summed over some surrounding region as being included in the mechanisms that determine RF selectivity, rather than as being part of the mechanism for coordinating the activity of many simultaneously active feature detectors.

Receptive fields that emphasize contrast with the surround show how the maximization of coherence is compatible with evidence for processes that emphasize the unexpected. If a single element differs from the others in an array on some simple variable, such as color or orientation, then the odd one out is very noticeable. This is evidence for processes that emphasize what is *not* predicted by the surround; we account for that evidence by noting that RFs usually develop so as to detect such differences. Enhanced transmission of the unexpected is also proposed by some theories to be a major role for the descending feedback projections from higher to lower stages of processing (e.g., Mumford 1992; Pece 1992). In contrast to these theories, Sillito et al. (1994) provide evidence that feedback from the primary visual cortex (V1) to the lateral geniculate nucleus (LGN) synchronizes the activity of those LGN cells that agree with the interpretation at the higher level. Furthermore, psychological experiments show that context often supports what is consistent with that context (e.g., Biederman 1972; McClelland 1978; McClelland et al. 1986; Palmer 1975).

We must distinguish the population codes proposed here from the population vector codes proposed by Georgopoulos (1990). The latter is a proposal about how a single vector could be signaled by the activity of a group of cells. Synchronization specifies which cells are in the same group. Synchronized population codes are compatible with but do not require population vector coding.

Finally, we need to relate the local contextual guidance that we hypothesize to spatial attention, arousal, and other strategic control processes. Local contextual guidance arises automatically from the interactions of local pro-

cessors and does not require specialized circuitry such as that hypothesized to be involved in spatial attention (e.g., Posner & Rothbart 1994; Van Essen et al. 1994). Furthermore, the contextual inputs that we postulate are highly specific in relation to both timing and the features they predict. Spatial attention seems to operate on a longer time scale and simply enhances the processing of whatever features are present at the locations and spatial scales attended, rather than enhancing some features and opposing others on a locally specific basis (Kröse & Julesz 1989; Nakayama & Mackeben 1989). Nevertheless, if there are local contextual interactions of the kind that we postulate, then attentional control processes will operate upon the synchronized population codes that those interactions produce, possibly themselves using mechanisms that increase the synchronicity of attended items (see Tiitinen et al., 1993, for empirical evidence on this from EEG recordings, and Goebel, 1993, for computational studies).

1.4. Prior theories using synchronized population codes and contextual guidance

The idea of population coding has a long history, with Hebb's (1949) notion of the cell-assembly serving as the leading representative. The further possibility that synchronized activity on a fine time scale specifies which subset of neurons is grouped to form the population code at any moment also has a long history (e.g., Abeles et al. 1993a; 1993b; Bienenstock 1995; Crick & Koch 1990; Eckhorn et al. 1991a; 1991b; Engel et al. 1992; Goebel 1993; Tononi et al. 1992a; 1992b; von der Malsburg 1981; Wang et al. 1990; Yamaguchi & Shimizu 1994). It has also been shown how synchronization can play a major role in neural network architectures based upon adaptive resonance (artmaps) and the boundary contour system (BCS) (Grossberg 1993; Grossberg & Somers 1991).

Various versions of the distinction between RFs and CFs also occur in prior theories. The coupling connections referred to in previous discussions of the substrate of synchronization (e.g., Singer 1990; 1993; 1994) are an example of what are referred to here as CFs. Further examples are the linking connections proposed by Eckhorn et al. (1991a) and the fast enabling links proposed by Hummel and Biederman (1992). Some aspects of the distinction also occur in theories that do not rely upon synchronization. Ullman (1994), for example, proposes that there is a bidirectional, bottom-up and top-down flow of information in which activity in one stream “primes” activity in the equivalent units in the reverse stream. This priming acts to increase the probability of transmission of any RF information with which it agrees. It uses mechanisms that differ from those of CFs but has similar effects: it modulates the probability of signals being transmitted in very locally specific ways and does so without corrupting the transmission of RF information.

One of the earliest theories with a distinction that is analogous to that between RFs and CFs is Edelman's (1978; 1989). This approach has been developed using highly detailed synthetic modeling (Reeke et al. 1990), which use large simulations (from about 1 to about 3.5 million connections) with many biological features, from intracellular processing to the overall principles of connectivity being built in. The models developed achieve perceptual group-

ing (Sporns et al. 1991) and form recognition that is independent of color and position (Tononi et al. 1992b); they also account for several other perceptual phenomena. Detailed analysis of their fine-grained internal temporal dynamics fits well with that observed in cortex (Sporns et al. 1989; Tononi et al. 1992b). Phasic reentrant signaling is crucial to the success of these simulations. Its functional role and mechanisms are closely analogous to that proposed here for contextual input. Our approach differs from theirs in minor differences of emphasis, however. For example, we put more emphasis upon simplifying computational studies, formal descriptions of the underlying computational goals, and the possibility that information supplied by the contextual connections could guide RF learning, thereby helping to establish some of the functional specialization that studies of reentrance have so far built in.

The use of predictive relationships between separate streams of processing to guide learning within streams has also been studied previously (e.g., Becker 1996; Becker & Hinton 1992; de Sa 1994a; 1994b; Schmidhuber & Prelinger 1993; Stone 1996; Stone & Bray 1995). Although there is basic agreement between our work and that of Becker and Hinton (1992), there are also some important differences. One difference is that we emphasize the use of context to coordinate ongoing activity and to form synchronized population codes, whereas Becker and Hinton communicate information between streams of processing for the purposes of learning only. Their reasoning was to ensure that distinct streams of processing could not increase the mutual information in their outputs simply by driving each other. Thus, in the approach of Becker and Hinton, local processors receive inputs that are used to change synaptic strength without directly affecting post-synaptic activity. We know of no biological evidence for such a process. Furthermore, a second major difference that follows from the first is that in Becker and Hinton's approach there are no cross-stream predictions to learn, whereas in our approach the CF predictions do play a major role because they embody the knowledge used to integrate ongoing activity and form synchronized population codes. A third difference is that in our approach a single parameter specifies the balance between maximizing information transmission within streams and maximizing coherence between streams. No such parameter exists in their approach.

1.5. Summary of some controversial hypotheses

The above considerations lead to hypotheses that we expect to be controversial. The first hypothesis is that there are basic computational capabilities common to many different cortical regions and to many different species. Second, in relation to the general functional role of any such capabilities, our hypothesis is that they include processes that gradually adapt them to the general statistical structure of the world in which the cortex finds itself, and they do so by maximizing the transmission of information that is predictably related to the context in which it occurs. Third, in relation to coding, we argue for synchronized population codes. Such codes differ from single-cell codes in that they convey information about internal structure; they differ from the more usual form of distributed code in that stored knowledge is used to group the elements into coherent subsets. Fourth, in relation to the short-term processing

dynamics, we hypothesize that local processors use contextual predictions to guide processing but without confounding those predictions with the information that they transmit about their receptive field inputs. This contrasts with the assumption, common to many connectionist theories of cognition, that local processors treat all of their specific informative inputs in essentially the same kind of way. Fifth, in relation to learning, we propose that RF features that are predictably related to the context in which they occur can be discovered – along with the predictive relationships between them. This contrasts with the common assumption that feature discovery is independent of associative learning. Finally, in relation to epistemology, we suggest that by discovering latent variables in diverse data-sets and the relations between them, the local processors are in effect discovering distal variables and relationships. As a consequence, they lay foundations for representation and meaning. Nevertheless, we will argue that these foundations do not constitute intentional representation proper because such local processors do not distinguish between the signals they receive and the distal causes from which those signals arise.

2. Arguments for and against common foundations for cortical computation

2.1. Specialized cortical regions and their common intrinsic organization

It is well established that the cortex contains many specialized regions, but our central concern is with their internal organization. In what ways is it common, and in what ways does it vary? Although differences exist, there is a widespread belief in commonalities: "It is easy to recognize a histological (e.g., Golgi) preparation as being cortex rather than cerebellum or tectum. It is much more difficult to tell whether it is human or bovine, motor, sensory, or associative cortex" (Braitenberg 1978, p. 444). "Laminations and vertical connections between laminae are hallmarks of all cortical systems, the morphological and physiological characteristics of cortical neurons are equivalent in different species, as are the kinds of synaptic interactions involving cortical neurons. This similarity in the organization of the cerebral cortex extends even to the specific details of cortical circuitry" (White 1989, p. 179). "Despite the many detailed properties that can be used to differentiate among the various cortical areas, the common properties of all the cortical areas are overwhelming. The same cell types, the same types of connections, and the same distributions of cells and connections across the cortical depth are found in all parts of the isocortex. These properties of the cortex are markedly different from those found in the other parts of the brain" (Abeles 1991, p. 33).

If there are commonalities, it is crucial to find out what they are. For extensive reviews of this issue, see Edelman and Mountcastle (1978), Rakic and Singer (1988), Martin (1988), White (1989), Shepherd (1990), Braitenberg and Schüz (1991), and Abeles (1991). Commonalities may exist at a number of different levels of organization and with respect to various aspects of function. Some may arise from small populations of pyramidal cells and their associated local circuit neurons, such as in the "canonical circuit" of Douglas and Martin (1990), or the "basic circuit" of Shepherd and Koch (1990). Others may arise at lower levels,

such as the morphology and physiology common to pyramidal cells. Hence, a common multicellular circuit is not necessarily implied by the hypothesis of common foundations for cortical computation because they could also arise at other levels of organization.

The basic homogeneity of the neocortex is widely thought to imply common information processing operations: "The typical wiring of the cortex, which is invariant irrespective of local functional specialization, must be the substrate of a special kind of operation which is typical for the cortical level" (Braitenberg 1978, p. 444). "It is taken as an article of faith that there is an information processing algorithm unique to cortex that is a product of the regularities of its architecture" (Stryker 1988, p. 133). "For many anatomists, it seems perverse to regard the visual cortex as an ad hoc collection of specialist circuits, rather than a set of basic circuits adapted to perform many different tasks. . . . For the neocortex, an unconventional class of models needs to be developed – models that are neural networks, but based directly on the biology; derived from visual cortex, but not designed to solve a particular problem in visual processing" (Douglas & Martin 1991, pp. 291–92). Such views have a long history (e.g., Edelman & Mountcastle 1978; Lorentó de No 1949; Rockel et al. 1980).

The belief in commonalities is supported by evidence that cortex contains some generalized learning algorithm that adapts each region to the input it receives. For example, it has been shown that sensitivity to visual features can be induced in the primary auditory cortex of neonatal ferrets by replacing its normal auditory input with visual projections (Sur et al. 1988). Similarly, it has been shown that visual cortex has the potential to develop an array of functional units that is appropriate to the somatosensory input (Schlaggar & O'Leary 1991).

Although these arguments for commonalities have force, they are not conclusive. Noting similarities will not be convincing until we can see clearly how they provide capabilities that are useful to many different cortical regions. Differences that are critical from a computational point of view may not be obvious from an anatomical or physiological point of view. Furthermore, the suggestion that some form of columnar organization is common to the whole of cortex can be criticized (e.g., Purves et al. 1992; Swindale 1990). Therefore, it is important to note that although "cortical columns" are not central to the hypotheses developed here, criticism of this idea suggests limitations upon anatomical arguments for commonalities.

2.2. Specialized cognitive functions and their common computational requirements

Functional specialization is also a major feature of cognitive organization. This has been established by studies of both normal and brain-damaged subjects, with cognitive neuropsychology providing a rich source of data and theory (Ellis & Young 1988; McCarthy & Warrington 1990; Shallice 1988). Functional specialization is most firmly established for perceptual and motor functions. Its existence and nature in the highest level functions such as strategic control is less firmly established, but there is some evidence for it even there (Shallice 1988; 1991).

The inferences drawn from cognitive and neuropsychological investigations are often shown in diagrams of func-

tional specialization and information flow. Our primary concern here is not with this system level of organization, but with the operations that are performed by the various cognitive subsystems. What are these operations, and which, if any, are common to different subsystems? Mapping the cognitive architecture is a complex and important task, but adding or deleting subsystems and routes between them will only be crucial to the search for commonalities to the extent that it changes the set of basic computational capabilities required. What those capabilities are is not obvious. This issue needs wider discussion.

One simple aspect of cognitive neuropsychological theory that suggests common operations is that subsystems are often distinguished by the content of the information with which they are concerned. This suggests that they differ primarily in what they operate upon, rather than in the operations that they are required to perform upon that information.

In contrast to our emphasis upon commonalities, studies of basic cognitive processes can give rise to skepticism about the value of a search for general principles. Crick reports the view to which Ramachandran has been led by his elegant and ingenious psychophysical studies, as follows: "It may not be too farfetched to suggest that the visual system uses a bewildering array of special-purpose tailor-made tricks and rules-of-thumb to solve its problems. If this pessimistic view of perception is correct, then the task of vision researchers ought to be to uncover these rules rather than to attribute to the system a degree of sophistication that it simply doesn't possess. Seeking overarching principles may be an exercise in futility" (Crick 1988, p. 156). Even with respect to Ramachandran's argument, however, Crick then adds: "It is, of course, possible that underlying all the various tricks there are just a few basic learning algorithms that, building on the crude structures produced by genetics, produces this complicated variety of mechanisms" (p. 156).

2.3. Evolutionary arguments

The cerebral neocortex evolved as an add-on to preexisting neural systems and has expanded rapidly at various stages of mammalian evolution (Jerison 1973). The speed of this evolution has been used to support the view that it embodies a multipurpose form of computation: "Neocortex has expanded rapidly in phylogeny by creating multiple new areas. While mammals with very small cortices have behavioral capacities no more impressive than noncorticate animals, the capacity for rapid phylogenetic change may be the most important feature of cortex" (Stryker 1988, p. 133).

There is a separate and important evolutionary function that a *generic* principle for the development of a perceptual network layer – whether it be infomax or some other principle – can serve. Suppose that an evolutionary mutation produces a modified eye, or merges the auditory signals into the visual pathway at some new point. If there were no generic principle for layer development, we might imagine that mutations would have to occur simultaneously in the processing function of several layers, for those layers to be able to use the novel input properly. But if there is such a generic principle – one that applies to each layer regardless of what type of input reaches it – then the novel input will automatically be processed in accordance with that principle. This suggests that the existence of a generic principle may greatly increase the likelihood of a mutation being adaptive. (Linsker 1988, pp. 116–17)

Another evolutionary argument for commonalities arises from the comparative study of learning. After an extensive search for basic differences in learning abilities across various species, Macphail (1987) concluded that all the problem solving abilities of nonhuman animals arise directly from a common basic associative process. Furthermore, the common process that he inferred from these comparisons is one that learns the causal links between events, that is, one that learns what predicts what.

Evolutionary arguments can also be used to oppose the hypothesis of commonalities. Tooby and Cosmides (1995) argue that the evolutionary perspective entails the functional analysis of niche-differentiated cognitive and neural machinery that is unique to the species: "the human cognitive architecture is far more likely to resemble a confederation of hundreds or thousands of functionally dedicated computers, designed to solve problems endemic to the Pleistocene, than it is to resemble a single general-purpose computer equipped with a small number of general-purpose procedures such as association formation, categorization, or production-rule formation" (Tooby & Cosmides 1995, p. 1189). The denial of a small number of general-purpose learning procedures is a crucial part of this perspective. Gallistel (1995) concludes that the catalog of special-purpose learning procedures, such as the ability of birds to learn the position of the celestial pole, could be enlarged indefinitely.⁴

These arguments do not settle the issue, however. Neither the presence of highly specific abilities nor the absence of a single all-powerful ability implies that there are no abilities common to many different species and to many different cortical regions. To rebut the view that classical and operant conditioning are general purpose procedures, Gallistel (1995) proposes that they are specialized for the solution of problems in multivariate, nonstationary time series analysis. This enables them to figure out what predicts what. Such a capability may not be all-purpose but it is far less specialized than an ability that can learn only the position of the celestial pole.

2.4. Computational arguments

Impressive advances in the theory and technology of neural computation since 1980 have greatly encouraged our search for commonalities. This is because they show that powerful multipurpose capabilities can be implemented in neural systems (e.g., Amit 1989; Gluck & Rumelhart 1990; Rumelhart & McClelland 1986). They suggest that these capabilities are likely to contrast with those of conventional von Neumann computation and it has often been proposed that this contrast depends upon the use of distributed representations or population codes:

Distributed representations give rise to some powerful and unexpected emergent properties. . . . For example, distributed representations are good for content-addressable memory, automatic generalization, and the selection of the rule that best fits the current situation. . . . Thus, the contribution that an analysis of distributed representations can make to these higher-level formalisms is to legitimize certain powerful, primitive operations which would otherwise appear to be an appeal to magic. (Hinton et al. 1986, p. 79)

This viewpoint is important because it suggests that cognition may be based upon computational primitives that are not obvious a priori, but which are of general utility.

3. Computational studies of the contextual guidance of learning and processing

This section is concerned with what Marr (1982) calls computational theory. If we are ever to understand how the cortex works, we must understand the work it does. What that work might be at the level of local cortical circuits is far from obvious. The hypothesis being examined here is that it includes the maximization of coherent variation, that is, transmitting as much information as possible while keeping it coherently related to what is going on elsewhere, and thus keeping it "meaningful." These studies are designed as simplifying abstractions, not as detailed models of biological systems. Their goal is to make the underlying computational task and strategy clear (Marr 1982; Phillips 1997; Sejnowski et al. 1988). This will make it easier to build and interpret detailed models that embody that strategy, and to design experimental paradigms to determine whether it is used by real biological systems.

3.1. Transfer functions

If the role of context is to modulate transmission through local processors so as to emphasize coherent outputs but without corrupting the information that is transmitted about the RF input, then we need a transfer function with the following properties. If there is no RF input then output should remain at the neutral level; if there is no CF input then the output should be monotonically related to RF input in some standard nonlinear and biologically plausible way; if RF and CF inputs agree then the gain of the function relating output to RF input should be increased; if RF and CF inputs disagree then the gain of the function relating output to RF input should be decreased; CF input should affect the *confidence* with which decisions are made but only the RF input should determine *what* decisions are made. Physiological studies outlined in section 4.2 show that neurons do indeed receive two classes of input that differ in approximately the way this suggests. In addition to the classical forms of excitatory and inhibitory input, they also receive inputs, such as those mediated by NMDA receptors, whose effects depend upon the prevailing state of activation and could therefore fulfill the gain-controlling role of CFs (e.g., Fox et al. 1990). A function, $A(r,c)$, giving the internal activation of probabilistic bipolar $(-1,1)$ units has been derived from these computational and physiological considerations (Kay & Phillips 1994; 1997; Phillips et al. 1995b), such that $A(r,c) = 0.5r [1 + \exp(2rc)]$ where r = summed weighted RF inputs including any bias input and c = summed weighted CF inputs including any bias input. An equivalent activation function can be given for processors that produce binary $(0,1)$ outputs. To compute the output probability in the simulations we apply the standard logistic squashing function to the internal activation, so the transfer function as a whole is composed of the activation function followed by the squashing function. The neutral level is given by an output probability of 0.5. The continuous value transmitted between units is the expected value of outputs with this probability, which ranges from -1 to 1 . As Figure 3 shows, this transfer function has the properties required. It is not unique, but it is a clear and simple representative of the limited class of functions with these properties (Kay & Phillips 1994; 1997). A natural interpretation of the output given by this transfer function is that

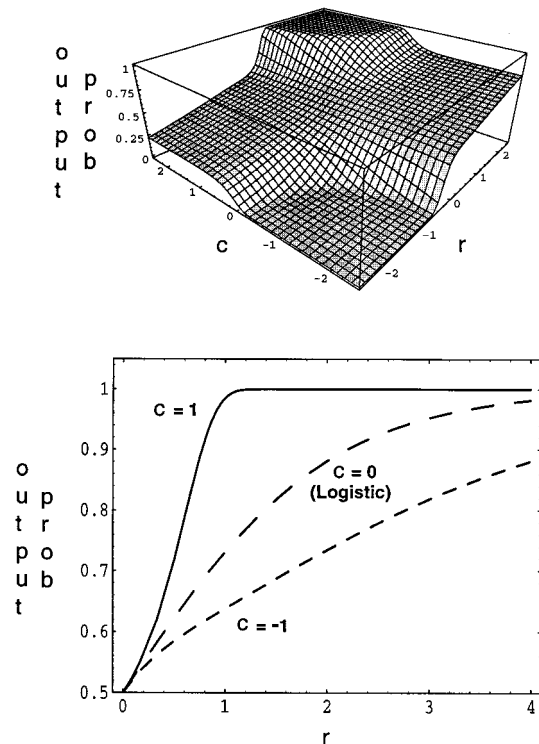


Figure 3. A transfer function for local processors with contextual guidance showing how output probability is related to the integrated receptive field (RF) input, r , and to the integrated contextual field (CF) input, c . When the RF input provides no evidence in favor of either a positive or a negative output decision, then $r = 0$; output remains at the neutral probability level of 0.5 and is unaffected by CF input. The lower panel shows the effect of context in the case where the RF input supports a positive output decision, for just three specific values of c . When $c = 0$, the probability of a positive output increases from 0.5 as specified by the logistic function. When the context predicts a positive output, $c = 1$; the probability of a positive output increases more rapidly as a function of positive RF input. When the context predicts a negative output, $c = -1$, the probability of a positive output increases less rapidly as a function of positive RF input. The equivalent effects also occur when the RF input favors a negative decision.

it gives the probability of a discrete event such as an action potential.

Our computational and empirical studies both emphasize three closely interrelated but distinct forms of neural signaling: relative timing, place, and firing rate. The possibility of using synchronization, or relative timing, to signal grouping using the transfer function just defined follows from the way in which the CFs influence output probability. They increase the probability that outputs from different processors will be produced at the same time if they are mutually predictive, and they reduce this probability if the outputs are opposed. Section 3.4.3 shows that this produces coherent groupings.

Place coding is the transmission of information about different features or variables by different cells. It allows for the possibility that a number of different cells could all transmit information about the same feature. This form of signaling is preserved in the computational studies through the use of outputs from different units to signal different variables. Our emphasis upon self-organization implies that

this coding is not fully prespecified, but may change as the system adapts to its inputs through learning.

The classical form assumed for rate coding is the transmission of information through the firing rate of single cells measured over a time period that is long relative to the duration of individual action potentials. This is one way to transmit information about continuous variables such as the output probabilities generated by the above transfer function. It is not the only way, however. Imagine a set of cells that produces action potentials with a probability that is essentially the same for all cells, for example, some form of neuronal group as proposed by Edelman (1978; 1989). The crudest estimate of that probability is given by sampling a single cell for a single brief interval that is long enough for just one binary output, for example, 1 or 2 msec. A better estimate can be obtained by sampling many of the cells for this brief interval. If the output probability remains approximately constant for a time of more than 1 or 2 msec, then an even better estimate can be obtained by sampling many of the cells over that longer time. Thus in this simple case these measures give different estimates, with varying amounts of precision and bias, of the same underlying quantity. This suggests that much of classical single-unit neurophysiology has been developed so as to exploit situations in which the relevant underlying quantity remains constant for long enough to allow adequate estimates to be obtained by sampling a single cell over longer intervals and by averaging across trials. The success of this enterprise does not imply that there are no sets of cells signaling the same underlying quantity in those situations, nor does it imply that there are no situations where that underlying quantity changes rapidly. The only way to obtain an accurate estimate in the latter case would be by averaging the outputs of a set of cells over a brief interval.

Continuous values were transmitted between processors in most of the simulations summarized below. However, a few simulations have been run in which only binary values were transmitted, that is, single units were used for each output probability and at each iteration of the computation of the short-term dynamics each unit transmitted just a single binary output with that probability. Performance did not seem to be very sensitive to this change, so our working assumption is that high precision in transmitting the output probabilities is not a necessary requirement of the computational approach being developed here.

3.2. Information-theoretic objective functions

The goal of maximizing the transmission of coherent information can be specified in a precise but general way by using the concepts of Shannon entropy, mutual information, and conditional information (Kay & Phillips 1994; 1997; Phillips et al. 1995a; 1995b). The Shannon entropy, $H(X)$, is the average amount of information in any variable X with a given probability distribution. Mutual information, $I(X;Y)$, is a measure of the average amount of information that is shared by the probability distributions of two variables, for example, X and Y . It is a measure of the extent to which uncertainty about one variable is reduced by observing the other, and it is a commonly used measure of information transmission. If two variables are independent then their mutual information will be zero. Conditional information, $H(X|Y)$, is a measure of how much uncertainty is left about one variable, for example, X , given that we

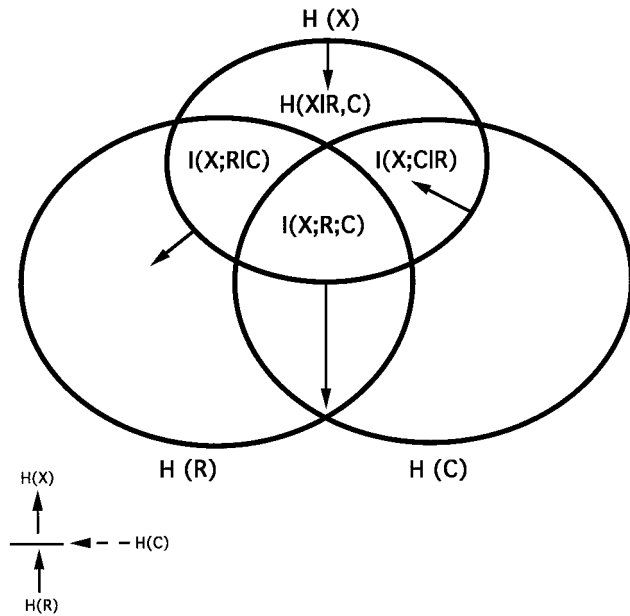


Figure 4. A Venn diagram illustrating the decomposition of the information in the output of a local processor, X , into four disjoint components (in the case where $I(X;R;C)$ is positive). A goal for processing can be specified as the relative importance attached to increasing or decreasing each component. $H(X)$ is the total information (Shannon entropy) in X , etc. The length and directions of the arrows indicate the goal of increasing or decreasing each component of coherent information, $I(X;R;C)$, as much as possible, increasing the transmission of the information in the RF that is independent of context, $I(X;R|C)$, but to a lesser extent, and reducing the transmission of any information in the context that is independent of the RF, $I(X;C|R)$. The fourth component of $H(X)$, i.e., $H(X|R,C)$, is information in the output that was in neither the RF nor the CF input. This is noise generated by the processor itself. Reduction in each of these last two components can either be specified as part of the goal or can be left to occur by default as a result of dedicating output channel capacity to other components. The information flow is shown by the icon bottom left.

already know another variable, for example, Y . From these definitions it follows that $H(X) = I(X;Y) + H(X|Y)$. For a lucid introduction to these concepts see Hamming (1980).

Consider a local processor to have input vectors R and C constituting the RF and CF inputs respectively, and to produce an output vector X . The Shannon entropy in X can be decomposed as follows

$$H(X) = I(X;R;C) + I(X;R|C) + I(X;C|R) + H(X|R,C)$$

where the first term on the right is a measure of the information that is common to X , R , and C ; the second is that common to X and R but not to C ; the third is that common to X and C but not to R ; and the fourth is information in X that is in neither R nor C . Figure 4 illustrates this decomposition in the case where all components are positive.

A goal for any local processor, X , can now be specified in terms of these four components. Each processor must adapt on the basis of just the information that is locally accessible to it. We therefore specify how X should adapt taking R and C as givens, but allowing for the possibility that any connections upon which R and C depend may themselves be adapting in the same way. We require X to convey information about major sources of variation in R , and in

particular those that are predictably related to C . Data compression, as argued for in sections 1.1 and 1.2.1, will be ensured by constraining processors to have fewer outputs than inputs.

Discovering major sources of variation in R requires the maximization of the mutual information between the output and the RF input, $I(X;R)$, which consists of two components, that is, $I(X;R;C) + I(X;R|C)$. Consider first the transmitted information that is common to the RFs and CFs, that is, $I(X;R;C)$. This is the RF information that is coherently related to the context, and we require the local processor to transmit as much of it as possible. If the RFs and CFs arise from separate data-sets then any information they share must reflect some common distal influences upon those data-sets, and the more diverse the data-sets the more distal those common influences are likely to be. Maximizing this component is therefore likely to transmit information about variables with relevance to the environment within which the system operates. Variables in the RF input that are unrelated to the context, that is, $I(X;R|C)$, may also be useful at some later layer of processing or stage of learning, however; thus, this component could also be increased, though with a lower priority than the information that is meaningfully related to its local context. Information that is shared by X and C but not by R , that is, $I(X;C|R)$, should be decreased because the role of X is to transmit information about R , but not about C . Finally, $H(X|R,C)$ denotes variation in X that is due to neither R nor C , that is, intrinsic noise that is added by the processor to its output. We would normally wish this component to be reduced.

We can now formulate the general class of objective functions

$$F = \phi_0 I(X;R;C) + \phi_1 I(X;R|C) + \phi_2 I(X;C|R) + \phi_3 H(X|R,C)$$

where F is the objective to be maximized, and ϕ_0 , ϕ_1 , ϕ_2 and ϕ_3 are parameters in the range 1 to -1 . The ϕ_i parameters weight the various components of the transmitted information, $H(X)$, with positive values for the components that we wish to increase and negative values for components that we wish to actively decrease. Different objectives can therefore be given as different values of these parameters. The goal of maximizing information transmission within streams requires an objective function $F = I(X;R)$, which is given by setting $\phi_0 = \phi_1 = 1$ and $\phi_2 = \phi_3 = 0$. This is the goal studied by Linsker (1988) and many others; it is called Infomax. The goal of maximizing the transmission of information that is predictably related to the context requires an objective function $F = I(X;R;C)$. This is given by setting $\phi_0 = 1$ and $\phi_1 = \phi_2 = \phi_3 = 0$. We call this goal Coherent Infomax. It is equivalent to that intended by Becker and Hinton (1992) but has a different form because we make explicit the requirement to maximize the transmission of RF information.⁵

3.3. Learning rules

Learning rules can be derived by performing gradient ascent on the objective function, F , relative to the strengths of the RF and CF connections (Kay & Phillips 1994; 1997; Phillips et al. 1995b). The dependence of change in connection strength upon postsynaptic activity as specified by these rules is shown in Figure 5 (from Smyth 1994). The learning rules have the same general form for the RFs and for the CFs. The change in synaptic strength is proportional




Figure 5. Change in the strength of the RF and CF connections (weight change) as a function of postsynaptic activity (output probability). The threshold above which connection strengths are increased is a function of prior activity (from Smyth 1994).

to presynaptic activity, but it is nonmonotonically related to postsynaptic activity. The nonmonotonicity required is similar to the computationally powerful BCM learning rule proposed by Bienenstock et al. (1982), and to a simpler version, the ABS rule, that has been shown to have biological plausibility (Artola et al. 1990; Hancock et al. 1991a; 1991b). Other learning rules have also been developed within this general approach, including one that maximizes the covariance between the integrated RF and CF inputs (Der & Smyth, in press; Smyth & Der 1995). The latter rule may be easier to implement than the one shown in Figure 5, but has the same overall form of dependence on the level of postsynaptic activity, which it shares with the BCM and ABS rules.

A major feature of the learning rule shown in Figure 5 is the threshold of postsynaptic activity below which connection strengths are decreased and above which they are increased. In the rules derived by Kay and Phillips this depends upon three specific dynamic conditional averages of prior activity, namely, the average prior output probability of the unit taken over all RF and CF inputs; the average output probability for the current RF input taken over all CF inputs; and the average output probability for the current CF input taken over all RF inputs.⁶ Only the first of these three is used by the BCM rule; its role is to make it harder to increase the strengths of connections to units that have already been active too frequently, and vice versa.

3.4. Simulations of simple networks of local processors with contextual guidance

3.4.1. Learning A simulation performed by Darragh Smyth at Stirling shows how context can guide RF learning. Two streams of single-unit processors were linked by CFs; as a consequence, they were able to discover variables that were correlated across streams (Fig. 6). Pairs of vertical or horizontal bars were presented that were both vertical on 70% of occasions and both horizontal on 30% at random. The bars were bright on dark or vice versa at random. The signs of the vertical bars were uncorrelated across streams,




Figure 6. The architecture and stimuli used to compare the goals of maximizing coherence across streams, that is, Coherent Infomax, with that of maximizing information transmission within streams, that is, Infomax. Within streams the single unit in each processor received RF input from all nine receptors. It also received CF input from the unit in the other stream. The stimuli consisted of the six pairs of inputs, each being presented at random with the probabilities shown. The sign of the horizontal bars was not the most informative variable within streams but it was correlated across streams.

but the signs of the horizontal bars were perfectly correlated across streams. The sign of the vertical bar therefore carries more information within streams, but the sign of the horizontal bar is more relevant to the correlation across streams.

The course of learning and the receptive fields found after learning are shown in Figure 7. When the goal of learning was specified as the maximization of information transmission within streams, both local processors became sensitive to the sign of the vertical bar. When the goal of learning was specified to be the maximization of coherence across streams, both local processors became sensitive to the sign of the horizontal bar. The CF connection strengths were then also learned correctly, thus embodying the cross-stream predictions.

Other simulations show that these networks have a rich array of possible behaviors depending upon the goal specified, the activation function used, the learning rate, the starting weights, and the correlations within and between the RF and CF inputs. When the goal is to maximize the RF input that is predictably related to the context, this is done

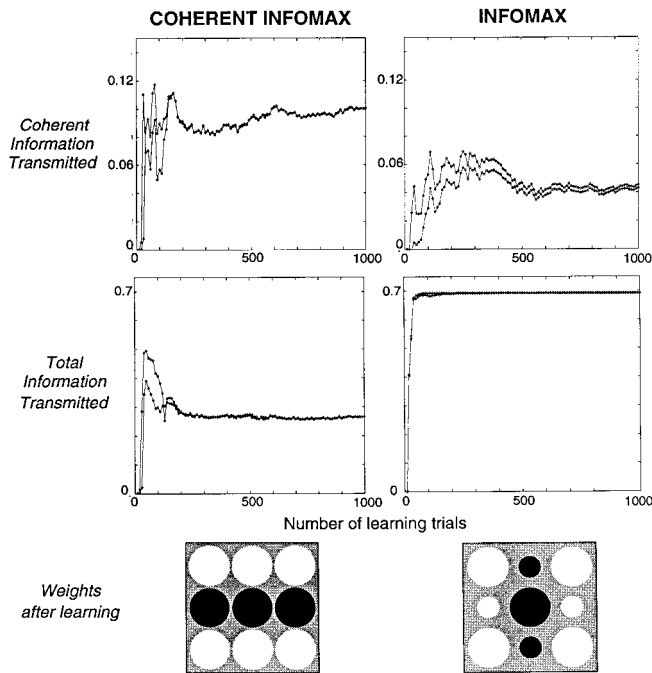


Figure 7. Results of the simulations using the architecture and stimuli shown in Figure 6: Learning specified by the goal of maximizing coherence across streams is shown on the left; that specified by maximizing information transmission within streams is shown on the right. Each of the four panels shows the information transmitted (\log_e) by each of the two output units at different stages of learning. The synaptic strengths produced by each form of learning are shown at the bottom. White is positive, black is negative, and diameter is proportional to the absolute value. Each of the two sets of synaptic strengths is representative of what was found in both streams.

irrespective of whether or not it maximizes transmission of information about the RF. When the goal is to maximize transmission of information about the RF, that is achieved irrespective of contextual predictability. Transition between these two goals is specified by varying a single parameter, ϕ_1 , from 0 to 1.

The ability to discover the relevant RF variables and to ignore the irrelevant has been shown when (1) the relevant variables are the most informative within streams, when (2) they are not the most informative within streams, and even when (3) there is no evidence within streams as to the existence of these particular variables (Kay & Phillips 1994; 1997; Phillips et al. 1995b; Smyth 1994; Smyth et al. 1994). These abilities have been shown within a variety of network architectures including networks with multiple streams and contextual connections between streams; multistream networks with two layers of processing and contextual connections within streams and from higher to lower layers; multistream networks in which the RF fields of different streams overlap with each other; and multistream networks with contextual connections between neighboring streams only (Kay & Phillips 1994; 1997; Phillips et al. 1995b; Smyth 1994; Smyth et al. 1994). In all cases the correct contextual predictions are learned together with the discovery of the RF features that they relate. The networks learn faster with more streams and are sensitive to small correlations between streams (Phillips et al. 1995b).

In the example shown in Figures 6 and 7, the features that were correlated across streams were the same, that is, the sign of a horizontal bar, and in the case of visual input the different streams are most easily thought of as arising from different places in the image. Neither of these aspects is crucial, however. The features that are mutually predictable could arise from different features in different streams, from the same cues to a common underlying variable at different places in the image, or from different cues to a common underlying variable at the same place in the image; they can also include cross-modal contextual input.

3.4.2. The codes produced by these learning rules. The approach has been developed to include local processors with multiple output units, as shown in Figure 8 (Floreano et al. 1995; Kay et al., in press). Different units within a processor adapt their RF weights in order to transmit information about different variables, thus increasing the amount of information that each processor can transmit. These networks thus have four levels of organization: units, processors, streams, and layers. Local codes were produced for the relevant variables in the simulation shown in Figures 6 and 7, where there was little room for any other form of coding. When more than one variable is relevant within streams, and when multiunit processors are used, a greater range of possible codings exists (Floreano et al. 1995; Kay et al., in press). The codes produced vary but are not reliably related to the input variables in any simple way. Simple input variables are sometimes signaled by single units and

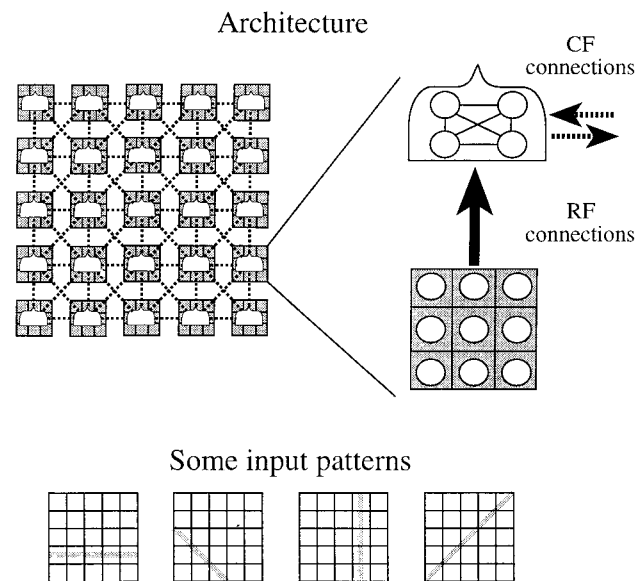


Figure 8. The architecture of the net used to study coherent grouping. Twenty-five streams of processing with nonoverlapping RFs were arranged as a 5×5 array with contextual connections between neighboring streams as shown by the lines joining the local processors. The enlargement on the right shows that in each stream the processors were composed of four units, with each unit receiving RF input from each of nine input units. Units receive CF input from each of the units in the neighboring streams, as well as from all other units within the same processor. Input patterns were continuous, positive or negative, horizontal, vertical, and diagonal lines as shown by the examples.

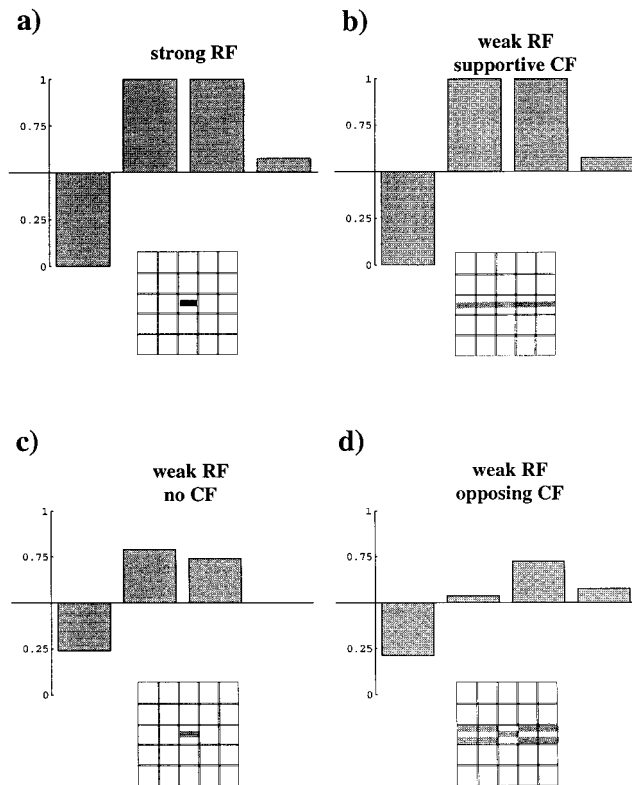


Figure 9. The effects of opposing or supportive context on the output probabilities of the four units of the local processor at the centre of the 5×5 array in the architecture shown in Figure 8. The four output probabilities are shown as increases or decreases from the neutral level of 0.5.

sometimes they are not. In short, single units within multi-unit processors do develop selective tuning functions; yet these are generally not related to the input in an intuitively obvious way, and they vary not only across streams but also in different instances of the same network architecture and input pattern set. However, it matters most that the relevant information is transmitted; how that information is distributed across the available output signals is not crucial. In relation to the study of receptive field selectivity in cortical neurons this suggests that the information conveyed by a population of cells may be more important than the exact way in which it is distributed across the individual cells, and this is consistent with the rich variation in detailed selectivity that is often observed in cortical neurons.

3.4.3. Dynamic grouping. To show how contextual connections can produce coherent grouping, Kay et al. (in press) simulated large arrays of multiunit processors and studied the effects of the CF input on the short-term processing dynamics. In one simulation, 25 streams of four-unit processors were arranged as a 5×5 array, with each stream receiving RF input from a 3×3 array of units. All four units in each stream received contextual input from all units in their neighboring streams (Fig. 8). The training input was comprised of collinear horizontal, vertical, and diagonal bars displayed upon the 15×15 input array. The learning algorithm was set to maximize coherence across streams (i.e., Coherent Infomax, see sect. 3.2). Learning was found to scale up successfully to this case, with all streams tending

to discover the relevant input variables at around the same time.

The influence of the CFs after learning can be seen in Figure 9, which shows the effects produced on one stream of processing by supportive or opposing activity in other streams. The output probabilities for the four units of the processing stream at the centre of the 5×5 array when presented with a horizontal bar on its 3×3 receptive field are shown for four cases: (1) when the RF input is strong, (2) when the RF input is weak and there is no contextual input, (3) when the RF input is weak and there is supportive contextual input, and (4) when the RF input is weak and there is opposing contextual input. Note that the only way the context can influence output in this architecture is via the contextual connections. The results show that these contextual inputs increase the probability of outputs coherently related to that context and decrease the probability of opposing outputs. These effects are produced rapidly, within just one or two iterations of the computations that update the outputs.

A second simulation, run by Floreano (1994, unpublished), is analogous to demonstrations such as the Rubin vase and the Necker cube, where two different perceptual organizations are possible but only one occurs at a time. Such phenomena might reflect the effects of mutual contextual guidance in cases where the input provides evidence for both of two internally coherent feature sets that are mutually incompatible; in order to study the short-term processing dynamics in such a case, a net with 100 single-unit streams was simulated. For the sake of this demonstration, two subsets of nine units were specified so that each had positive CF input from all other units in the same subset and negative CF input from all of the units in the other subset. All streams received inputs that varied randomly across iterations with the constraint that inputs to both subsets of nine units were positive.

The output probabilities produced at each iteration of the computations updating the activity of the network are shown in Figure 10. As these iterations involve just one synaptic delay, each iteration corresponds to just a few milliseconds. A coherent subset of features emerges from the background within 3 or 4 iterations, but only one of the two alternative organizations emerges at any one instant. Within a cooperative subset, all outputs emerge from the background simultaneously and are less affected by random variation in their inputs than are the responses to the background. These effects are similar to the retrieval of memories in recurrent auto-associative attractor networks (e.g., Amit 1989; Hopfield 1982), but with the important difference that contextual connections just organize the input data into coherent subsets without adding features for which there is no evidence in the input. The short-term dynamics of nets with contextual guidance and a feedforward RF connectivity is therefore constrained to remain close to the input.⁷

3.5. Assessment of the computational studies

The simulations described above used very simple nets, but models of large and complex nets that combine contextual integration with many biological details show that they can preserve the capabilities required of the short-term dynamics.⁸ As far as the contextual guidance of learning is concerned, theoretical considerations and our own simula-




Figure 10. Coherent grouping in a net with 100 single-unit streams. Each unit receives just a single RF input. Two groups of nine units, each arranged as a 3×3 rectangle, have positive CF connections with all other members of their group and negative CF connections with all members of the other group. Input and output strengths for each unit are shown as the absolute deviation from an output probability of 0.5; zero is black and 0.05 is white. Outputs are calculated over 12 iterations. Input strengths vary randomly from iteration to iteration, but the inputs to both sets of nine units are always positive. Output probabilities increase simultaneously for all nine units of a group within a few iterations, but only one group emerges from the background at a time.

tions suggest that this may become easier rather than harder in larger systems, because more streams can provide better guidance.

We have so far had only limited success in using the learning rule outlined in section 3.3 to discover arbitrary nonlinear functions in nets with two layers of feedforward RF weights. Nets with feedback contextual guidance from higher to lower layers can sometimes discover such functions; they do so more often with more streams of processing (Phillips et al. 1995b). They do not solve such problems reliably, however; one reason for this is that when units in higher layers compute nonlinear functions then the feedback predictions conflict with what the units in the intermediate layers are able to compute. Others have shown that learning by maximizing coherence across streams can discover useful higher-order functions when applied to real-world problems (de Sa 1994a; 1994b; Stone 1996; Stone & Bray 1995; for a review and further applications, see Becker 1996). These algorithms do seem more limited in what they can learn than supervised algorithms such as error back-propagation, but, as we will argue in section 6.4, this does not necessarily make the algorithms less plausible as analogies to self-organization in the cortex.

Important outcomes of these computational studies are as follows: First, the goal of maximizing the transmission of

contextually relevant information can be specified precisely within the framework of information theory. This can be done despite the antithesis between information and meaning that is described so clearly by Hamming (1980, p. 103), and which has limited the usefulness of information theory to psychology and neurobiology (Horgan 1995). The approach being developed here may therefore help extend the application of information theory to brain function beyond the sensory systems. Second, feature discovery and associative learning can cooperate in such a way as to discover variables that are predictably related across diverse data-sets without needing a supervisor that already knows about those variables. Third, it is possible for the output of a local processor to be affected by contextual input while still transmitting unambiguous information about the RF input. Fourth, the form of learning derived analytically from the information-theoretic goals adds further support to the hypothesis that changes in synaptic strength depend nonmonotonically upon postsynaptic activity in approximately the way proposed for the BCM and ABS rules (Artola et al. 1990; Bienenstock et al. 1982; Hancock et al. 1991a; 1991b).

4. Physiological evidence for contextual integration and synchronized population codes

Here we outline evidence for context-dependent synchronization of activity in the cortex, for corticocortical contextual connections that are involved in this synchronization, and for plasticity of the receptive field and contextual field connections. For more detailed reviews see Singer (1990; 1993; 1994; 1995) and Singer and Gray (1995).

4.1. Context-dependent synchronization

Intracolumnar interactions are shown by simultaneously recording the activity of cells within a small region of cortex. Synchronization of neighboring cells (< 200 μm apart) has been observed in many different species and cortical regions of awake and anaesthetized animals; it can also be observed in the local field potential (LFP) as well as in the multiunit and paired single-unit recordings (e.g., Gray & Singer 1989; Gray & Viana Di Prisco 1993; Kreiter & Singer 1992; Michalski et al. 1983; Toyama et al. 1981; Ts'o et al. 1986). Synchronization of neighboring cells with overlapping RFs and feature selectivity sometimes reflects common thalamic input, but it is more often characterized by dynamic properties that can only be accounted for by reciprocal interactions via local intracortical connections. Overall, the evidence suggests that the activity of local neuronal groups of cells is often closely synchronized.

Intercolumnar interactions are shown by simultaneously recording the activity of cells in different parts of the cortex, and synchronization has been observed between cells that are far apart (e.g., > 2 mm). In that case it occurs predominantly between cells with similar receptive field selectivity, and it reduces with distance (e.g., Gray et al. 1989; Michalski et al. 1983; Schwartz & Bolz 1991; Ts'o et al. 1986). Its occurrence within and between visual areas depends upon whether the cells that are observed are stimulated by single or separate objects. For example, synchronization is strong when two cells in V1 with non-overlapping but collinear preferred orientations are stimu-

lated by a single long bar moving across their RFs (Gray et al. 1989). It is weaker when they are stimulated by two short bars moving in the same direction, and it is abolished altogether when the two short bars move in opposite directions. These and many other results support the view that the synchronization of distributed activity in the visual system implements the well established Gestalt principles of perceptual grouping.

The prediction that cells can be part of different groupings at different times depending upon the stimulating conditions has been tested in the primary visual cortex of the cat (Engel et al. 1991b) and of the awake behaving monkey (Kreiter & Singer 1994). These experiments show that when two cells with different orientation and direction preferences are stimulated by a single moving bar that is suboptimal for both, then they synchronize, but when they are stimulated by two separate bars, each being optimal for one of the cells, then they do not. Synchronization occurs within the secondary visual area MT of the awake behaving monkey, and depends upon whether the cells are activated by a single common stimulus or by two different stimuli (Kreiter & Singer 1996). Synchronization has also been observed within and between a variety of other cortical regions, including olfactory, somatosensory, and motor regions, as well as across hemispheres (Singer & Gray 1995).

4.2. Local contextual fields and their role in synchronization

The specific thalamic afferents to primary visual cortex, V1, provide examples of RF inputs, and the excitatory long-range horizontal collaterals connecting pyramidal cells in V1 with nonoverlapping RFs provide an anatomical basis for the CF inputs (Gilbert 1995). Long-range horizontal connections are common in V1 (Gilbert & Wiesel 1983; 1989; McGuire et al. 1991; Rockland & Lund 1983) and in other cortical regions (Gilbert 1992), and these connections have a synchronizing action (König et al. 1993; Löwel & Singer 1992). It has also been shown that inter-hemispheric connections have a specific role in synchronizing activity (Engel et al. 1991a). The descending connections from higher stages may also include signals that have a synchronizing role. Such connections are ubiquitous but do not seem to play a primary role in driving the cells to which they project. In accordance with this suggestion it has been found that the activity of cells at different stages of visual processing can be synchronized (Engel et al. 1991c), and that cells at later stages of processing in the visual system can synchronize the activity of relevant subsets of cells at earlier stages of processing (Bullier et al. 1992; Sillito et al. 1994).

The anatomical and physiological evidence therefore suggests that the contextual connections within and between regions of the visual cortex are organized as shown in Figure 11. These connections are not distinguished solely by their source, but also by the effect that they have on the processors to which they project, because they have a modulatory rather than a primary driving role. One way in which they could fulfill this role is through voltage-gated receptors. Synaptic receptors that are both ligand and voltage-gated have become known as NMDA receptors, and they are widely distributed on pyramidal cells throughout the cortex. These receptors provide a mechanism for voltage dependence because there is a magnesium block on them that is reduced by depolarizing the cell (e.g., Ascher et al. 1988). These channels therefore contribute more effectively to further depolarization when the cell is already partially depolarized; they thereby provide a mechanism for gain control. Fox et al. (1990) show that cells in cat visual cortex have one class of receptor channel that pro-

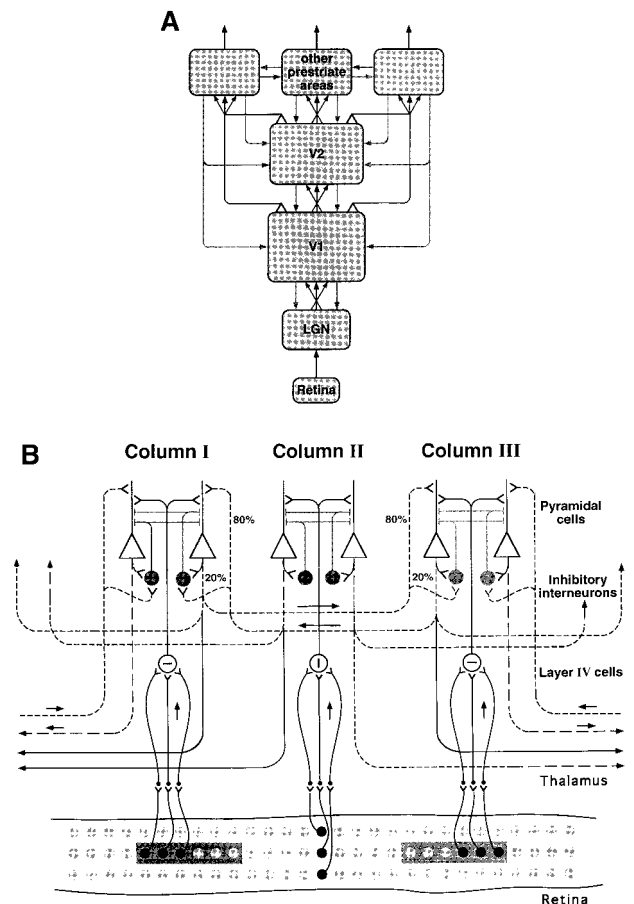


Figure 11. Organization of the synchronizing connections between (A) and within (B) regions of the visual cortex (Singer 1996). RF connections are shown as solid lines; synchronizing connections are shown as grey lines in (A) and as dotted lines in (B). (A) shows feedforward RF connections projecting from the lateral geniculate nucleus (LGN) to primary and secondary visual cortex (V1 and V2) and then on to other visual areas. CF connections link areas occupying the same level in the processing hierarchy; they are also included in the descending back-projections from higher to lower processing stages. (B) shows three processing columns within V1, two tuned to horizontal contours (Columns I and III), and one to vertical contours (Column II). Ganglion cells in the retina converge appropriately onto cells in Layer IV of the cortex after relay through the LGN in the thalamus. The output of the orientation selective layer IV cells is then relayed to pyramidal cells in other layers of V1, and these in turn project to higher stages of processing as indicated in (A). CF connections originate from pyramidal cells and link columns with the same or similar orientation specificity, here Columns I and III. Of these intraregional CF connections, about 80% terminate directly on pyramidal cells, and therefore provide excitatory input. The remaining 20% terminate on inhibitory interneurons (shown here as filled circles) which in turn synapse on pyramidal cells. Note that this wiring diagram shows only the essentials of the CF connections linking processors within a cortical area, omitting much sophisticated intracortical circuitry.

nism for voltage dependence because there is a magnesium block on them that is reduced by depolarizing the cell (e.g., Ascher et al. 1988). These channels therefore contribute more effectively to further depolarization when the cell is already partially depolarized; they thereby provide a mechanism for gain control. Fox et al. (1990) show that cells in cat visual cortex have one class of receptor channel that pro-

vides the primary drive and summates linearly, and a second class that provides amplifying gain-control (see Fox & Daw, 1992, for computational studies of possible mechanisms for these effects). If the long-range horizontal collaterals do provide synchronizing contextual input as hypothesized here, then their synaptic inputs should be predominantly modulatory rather than driving. The available evidence suggests that this is so (e.g., Hirsch & Gilbert 1991). Furthermore, if these long-ranging intraregional connections did contribute to the structure of the receptive field proper, then the receptive fields of cortical neurons would be much larger and more broadly tuned than they actually are.

Note that the hypotheses developed here do not imply that all voltage-dependent channels mediate CF rather than RF inputs. If any such distinction is relevant to cortex it is more likely that RF inputs produce strong activation of both voltage-dependent and non-voltage-dependent channels (Armstrong-James et al. 1993), whereas CF inputs produce strong activation only of voltage dependent channels. Note also that the absolute division of inputs into either RFs or CFs is a simplifying idealization. In cortex individual inputs may contribute to both roles but to varying degrees. Furthermore, there is evidence that although the long-range horizontal input is usually modulatory it can become more effective in generating spiking activity itself when the primary RF input is removed for many weeks (Das & Gilbert 1995).

4.3 Plasticity of the receptive field connections

It is now well established that the activity-dependent self-organization of synaptic connections could provide a substrate for learning in the cortex (Singer 1987; 1990). This is likely to involve long-term potentiation (LTP) and long-term depression (LTD), as well as control by global gating systems, and it applies to mature as well as to developing cortex (Singer & Artola 1994).

The learning rules formally derived from the information-theoretic objectives in section 3 require synaptic strength on active inputs to remain unchanged when postsynaptic activity is very low, to decrease when it is at intermediate levels, and to increase when it is high. The plasticity observed in slices of adult rat neocortex by Artola et al. (1990) supports these three specific predictions. Furthermore, it has been shown that much of the data on activity-dependent self-organization in the visual cortex can be explained by the BCM learning rule (Clothiaux et al. 1991), which makes the same three predictions.

Given the contextual input to local processors the possibility arises for this input to affect RF learning. Indeed, it is unlikely to have no effect, and we have argued above that it could have effects with far-reaching computational consequences. We know of no empirical studies explicitly designed to explore these possibilities, but results reported by Gilbert and Wiesel (1990) may be relevant. Their studies were mostly concerned with the effects of concurrent context upon the response to RF stimulation. Such effects could therefore be due to context modulating postsynaptic activity, but without having any effects upon the strengths of the synapses that carry RF input. However, it was also observed that, under some conditions, prior contextual stimulation altered the orientation tuning function that was later obtained in response to RF stimulation alone. This

suggests that the prior contextual stimulation played a role in changing the strengths of RF synapses, but as it could also have been due to other adaptation effects, studies using a modification of their paradigm to address this issue more directly would be worthwhile.

4.4. Plasticity of the contextual field connections

Although the basic organization of the CF connections could be genetically specified, shaping by experience is also necessary because RF feature selectivity depends upon experience. In keeping with this it has been shown that long-range horizontal collaterals undergo activity-dependent changes in synaptic strength (Hirsch & Gilbert 1993; Löwel & Singer 1992). Furthermore, there is evidence that the selection of these connections follows a correlation rule establishing preferential coupling between cells exhibiting correlated activity (Löwel & Singer 1992), in agreement with the predictive role proposed here for the CF connections. These ensemble-forming connections remain susceptible to use-dependent modifications in the adult (Singer & Artola 1994). Indeed, most of the experiments demonstrating LTP and LTD in neocortical synapses have been performed on corticocortical connections terminating on pyramidal cells in layers II/III or V (Artola & Singer 1993; Singer 1995), so they could predominantly reflect the plasticity of CF connections.

5. Psychological implications and evidence

This section shows how our hypotheses concerning local contextual integration can be tested and developed by behavioral methods, including cases where these are combined with physiological methods. We will argue that: (1) studies of the detection and grouping of simple stimulus elements provide behavioral evidence for contextual integration of the kind we propose; (2) it is reasonable to search for such processes at the higher levels of cognition, such as the perception of objects and words, because they are implemented by mechanisms that are widely distributed and of general utility; and (3) there is already theoretical and empirical support for the view that contextual integration at the level of local processors is relevant to these higher levels of cognition.

Our focus is on the visual perception of simple line element displays and on the perception of words. Theories using contextual integration and synchronization have been applied to a wide variety of other tasks with psychological relevance, however. Examples include the cocktail-party problem (von der Malsburg & Schneider 1986), perceptual grouping within and between multiple visual feature domains (Eckhorn et al. 1991a; Schillen & König 1994; Tononi et al. 1992b; Wang et al. 1990), form from motion and motion capture (Tononi et al. 1992b), object recognition (Hummel & Biederman 1992; Neven & Aertsen 1992; Tononi et al. 1992b), selective attention and scene perception (Goebel 1993), the binding of events across widely distributed cortical zones (Damasio 1989), reasoning (Shastri & Ajanagadde 1993), and consciousness (Crick & Koch 1990). Although these theories differ in detail they all suggest ways in which contextual integration at the level of local circuits can produce useful cognitive capabilities.

5.1. The effects of strabismus on synchronization and behavior

If synchronization is relevant to behavior then stimuli should be more perceptible if they produce synchronized activity, and conditions that reduce synchronization should impair perception. These predictions have been tested by comparing the effects of induced strabismus (squint) in cats on both synchronization and behavior. In strabismic cats neurons driven by different eyes lose the long-range horizontal intracortical connections that initially connect them (Löwel & Singer 1992). As a consequence these neurons do not synchronize, and the cats cannot fuse images from the two eyes (König et al. 1993). Furthermore, strabismus often leads to impaired perception in one eye (amblyopia), but the cortical activity evoked by input to that eye has so far seemed to be normal. The discrimination of gratings by cats using either their normal or their amblyopic eye has now been shown to be closely related to the extent to which the gratings produce synchronized activity (Engel et al. 1994; König et al. 1994; Roelfsma 1994a; 1994b). Both discrimination and synchronization are reduced in the amblyopic eye, and in both cases this reduction is greater at higher spatial frequencies. Recordings from the primary visual cortex of awake strabismic cats show that the amount of synchronization is directly related to perception and motor control. Under conditions of stimulation that lead to binocular rivalry, neurons connected to the eye that dominates perception and oculomotor response show increased synchronization and neurons connected to the suppressed eye show decreased synchronization (Fries et al. 1996). Changes in perceptual dominance are unrelated to changes in firing rate, however. These results show that response selection is closely related to synchronization; they thus support the view that internal grouping through synchronization on a fine time scale is important for the selection of perceptually or behaviorally relevant signals.

5.2. Contour integration in human vision through local "association" fields

Psychophysical evidence for dynamic grouping through a network of local linking connections between feature detectors is reported by Field et al. (1993). Subjects were shown arrays of 256 oriented band-pass line elements (Gabor patches) and had to detect a path of 12 elements with gradually changing orientation that was embedded within the random background. They found that performance was impaired by increases in the distance between the line elements and with the deviation of their orientations from collinearity, but it did not depend upon their relative "phases" (i.e., black on white or vice versa). Field et al. (1993) conclude that the ability to detect such paths is due to local "association" fields that link feature detectors in an organized way that depends upon their relative RF selectivities (Fig. 12). Their results show that these connections link feature detectors over distances that are large compared with the sizes of their receptive fields, and do so in such a way as to implement Gestalt grouping principles of proximity and continuity.

The dynamic grouping observed by Field et al. (1993) supports the hypothesis of locally specific processes of integration, and the "association" fields they infer from their findings are much the same as the CFs we propose. Field et al. (1993) note the similarities between the condi-

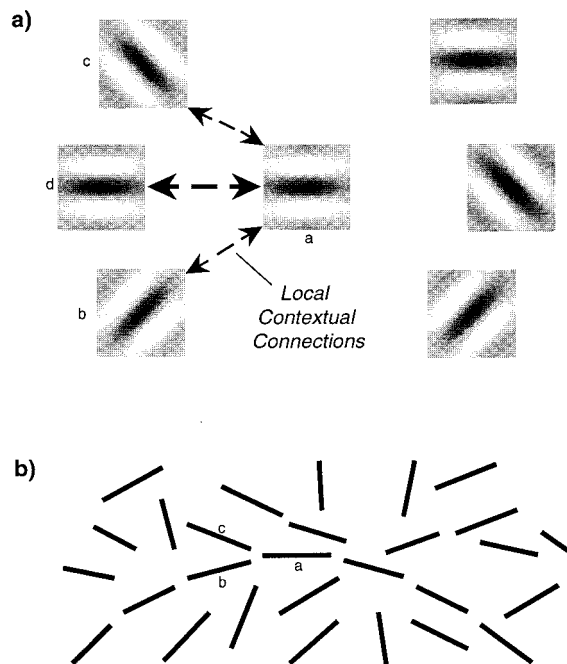


Figure 12. (a) Local contextual, or "association," fields as inferred from psychophysical studies of contour integration in human vision (Field et al. 1993): the square patches stand for oriented RFs at a particular spatial scale and in the relative positions shown. A sample of the contextual fields through which the output of cell *a* can be grouped with the outputs of other cells is shown. The arrows on the left show that they connect cells whose preferred orientations form first-order curves, with a strength that reduces with deviation from collinearity. Cell *a* will therefore not be connected to cells whose RFs have orientations and positions as shown on the right. This joint dependence on orientation and position is similar to that for the synchronizing connections shown in Figure 11. (b) An illustration of the perceptual effects assumed to arise from such contextual grouping connections. A long upward-facing curve can be seen running through the whole array. In this example, element *a* is more likely to be seen as being grouped with *b* than with *c*, even though it could equally well be paired with either on the basis of the pairwise connections. Such grouping processes may involve cooperative population effects, but an additional possibility is that the activity of cells operating at spatial scales that distinguish the individual elements is grouped by contextual inputs from cells that operate at a coarser scale and detect the long curve without distinguishing its elements.

tions under which they find good perceptual grouping and the conditions producing synchronization in V1. They also note that these conditions seem well matched to those determining the extent to which pyramidal cells are linked by long-range horizontal collaterals. These similarities are further strengthened by detailed comparisons between perceptual grouping criteria and the anatomy of the tangential intracortical connections which show them to be closely matched (Schmidt et al. 1997).

Finally, Field et al. (1993) argue that detection of the path in their studies must have been mediated by the grouped activity of the set of detectors activated rather than by the activity of a single high-level detector because a new path was formed randomly on each trial, and because the bandpass nature of the stimuli precluded their detection by cells with classically defined RFs covering the whole of the path. This argument is relevant to the issue of the relative

roles of local and distributed codes to which we will return below.

5.3. Effects of local context on the detection of low-contrast line elements

The psychophysical evidence just discussed provides evidence for local contextual fields in the grouping of easily detected stimuli. If such fields exist then they could also mediate locally specific effects of context on the detection of faint or ambiguous stimuli. Several psychophysical experiments indicate that this is so (Polat & Sagi 1993; 1994a; 1994b). Such studies show that targets are surrounded by a small region within which additional stimuli suppress target detection, and then by a larger region within which they facilitate detection provided that they are coherently related to the target such as by being collinear or near collinear. They thus suggest that within local streams of processing, inhibitory mechanisms force a choice between alternative features, whereas between streams, contextual interactions facilitate the detection of coherent features.

A detailed comparison of psychophysical and physiological evidence for such facilitatory effects of context on the detection of target elements is reported by Kapadia et al. (1995). They show that human visual contrast sensitivity is improved by a neighboring suprathreshold line element in a way that is reduced by increases in their spatial separation and with the deviation of their orientation selectivities from collinearity. Using equivalent stimuli in electrophysiological studies, Kapadia et al. (1995) also show that the response of superficial layer complex cells to low-contrast stimuli in V1 of awake attending Rhesus monkeys, as measured by summing spikes over 200 msec, depends upon the local relations between target and context in a way that is very similar to that seen in the psychophysical experiments. The physiological experiments also showed that these effects were not due to the context encroaching within the RF of the recorded cell, but were due to modulatory interactions between cells with nonoverlapping RFs.

5.4. Common contextual interactions revealed by psychophysical, physiological, and anatomical methods

The contextual conditions producing enhanced detection of low contrast stimuli in both the psychophysical and the electrophysiological studies of Kapadia et al. (1995) are very similar to those producing grouping of high-contrast stimuli in the psychophysical studies of Field et al. (1993) and synchronization in electrophysiological studies (e.g., Singer & Gray 1995). All four sets of findings show similar effects of spatial separation, spatial frequency, orientation, and collinearity. They may therefore be of great importance in reflecting common underlying mechanisms for local contextual integration. This view is strengthened further by their close match to the anatomy of long-range horizontal collaterals (e.g., Schmidt et al. 1997). If such mechanisms indeed do exist then they will provide an obvious candidate for explaining many other locally specific effects of context. The hypothesis that grouping and contextual facilitation of element perception involve common processes will be discussed further in section 5.7.2 in relation to theoretical and experimental studies of contextual integration in word perception.

5.5. Contextual modulation by variables that cannot themselves be perceived

One of our main hypotheses is that context can modulate the transmission of information about something other than the context. Thus, if some input variable is used only for contextual guidance then no explicit information will be transmitted about that variable in its own right.

5.5.1. Modulation by implicit color boundaries in achromotopic patients. A dramatic demonstration of this possibility comes from two neuropsychological patients who no longer see the world in color but only in black and white. The first, HJA, has been studied for many years (Humphreys & Riddoch 1987), but it has now been discovered that color does have implicit influences on his detection of luminance contrasts (Humphreys et al. 1992). In one task he had to say whether the top and bottom halves of a rectangular display differed in the second or the first of two intervals. Sometimes the two halves differed only in luminance, sometimes only in color, and sometimes in both. When the color difference was presented without any luminance difference, performance was at chance. This shows his achromotopia. When the two halves differed only in luminance, performance improved as luminance contrast increased. When a color difference was added, performance improved more rapidly than it did with the luminance contrasts on their own. The second patient, WM, was also tested on these tasks, and was compared in detail with HJA on a variety of other tasks. He has a different form of achromotopia, but shows the same kind of implicit modulatory effects of color differences upon the detection of luminance differences (Troscianko et al. 1993; 1996).

A simple interpretation of these results is that, in these patients, color streams continue to modulate luminance streams, but with their own feedforward outputs no longer functioning properly. Color differences can therefore still influence the detection of luminance differences, but without themselves being perceived. Color differences may influence the detection of luminance differences because they are highly correlated in natural visual images, and this is reflected in the connectivity between color and luminance channels. The correlations are not used to conflate the two variables, however, but to provide contextual guidance. Further evidence for facilitatory effects of color can also be found in psychophysical studies of normal subjects (Gur & Akri 1992; Troscianko 1994). The above findings therefore illustrate a key difference between RFs and CFs. Outputs are "seen" by later stages of processing not as conveying information about the CF input, in this case color, but as conveying information about the RF input, in this case luminance. The introspective reports of the two patients described above suggest that in this case those later stages involve conscious awareness.

5.5.2. Modulation by implicit context in visual texture segregation by normal subjects. To show how analogous effects can be sought using normal subjects we outline experiments that are currently being run at Stirling by Craven et al. on the interaction of target and contextual cues to texture segregation. A 20×20 array of small line elements, divided into two halves differing in mean line length, is displayed for 1 sec, and Ss decide whether the display is divided into long and short elements by a vertical

or a horizontal boundary. Contextual input is provided by also dividing the array into two halves that differ in the mean orientation of the elements. This boundary is coincident with the length boundary on 70% of trials and is orthogonal to it on 30%. If modulatory interactions are occurring then (1) the effect of context will increase as target strength increases; (2) this will occur just for a low range of target strengths; and (3) this range will be at higher values of target strength for weak than for strong contexts (Smyth et al. 1996).⁹ We also predict that the perception of the target boundary will be facilitated by a context with which it is coherent. Results obtained so far support these predictions and include modulatory effects of weak contextual cues that have no direct influence on response themselves.

5.6. Psychophysical evidence for RF and CF plasticity in sensory systems

Until the early 1980s plasticity in sensory systems was widely thought to be restricted to a critical period during development. There is now abundant anatomical, physiological, and psychophysical evidence for such plasticity in adults. Plasticity has been shown in visual, auditory, somesthetic, and motor systems. It is particularly dramatic at the cortical level and has been shown by studies of reorganization following deafferentation, nerve section, and cortical lesions, and from studies of the effects of more subtle changes in the patterns of sensory input received (Gilbert 1995; Kaas 1995). These effects occur on various time scales and include psychophysical evidence for fast perceptual learning in the visual sensory system of adult humans (e.g., Karni & Sagi 1991; Poggio et al. 1992; Polat & Sagi 1994b). These psychophysical effects are interpreted as being due to changes at sensory stages of processing because they are specific for eye, orientation, and spatial frequency, as well as for spatial position.

Some of these findings can be interpreted as changes in feedforward RF selectivity, such as when major reorganizations of sensory or motor maps occur. In other cases they are more likely to involve changes in the connections mediating intracortical contextual integration (Gilbert 1995). Consider the psychophysical findings reported by Polat and Sagi (1994b): prior to practice, the spatial range within which contextual stimuli facilitate target detection is up to about six times the spatial period of the target (Polat & Sagi 1993). After a few hours of practice, covering the whole range of separations to be spanned, the range of facilitation was increased by at least a factor of three (Polat & Sagi 1994b), probably by strengthening chains of local facilitatory interactions between filters with nearby but nonoverlapping RFs.

There are also other ways in which psychophysical experiments could study the learning of contextual predictions. For example, consider the techniques for studying local contextual integration described in sections 5.2, 5.3, and 5.5.2. Each of these could be used in paradigms where the predictive relationships between context and target are manipulated to see whether the effects of local context depend upon experience. There are already some findings suggesting that such studies would be worthwhile, for example, (1) the effects of strabismus outlined above show that experience affects both the CF connections and the probability of synchronization; (2) further studies of the

achromotopic patient WM suggest that the color-luminance interactions adapt to correlations that are experimentally induced between them (Troscianko et al. 1995); and (3) there are large practice effects in the texture segregation task outlined in section 5.5.2.

An example of the effect of learning on cross-modal integration is provided by Durgin (1995) who presented random dot patterns that had a greater dot density on either the left or the right. A tone was presented simultaneously and its pitch was perfectly correlated with the side of greater density. After 180 such pairings, a staircase procedure was used to measure perceived equivalence between left and right at an intermediate density. Simultaneous presentation of a tone affected matching such that the side that had been denser when that tone was presented during training was seen as being denser than it should have been in order to give an accurate match. The extent to which such cross-modal contextual learning affects discrimination within modalities when later tested in the absence of the cross-modal stimulus is not yet clear, but it is clearly amenable to further psychophysical study. Physiological evidence for effects of auditory stimulation on activity in the visual cortex (e.g., Fishman & Michael 1973; Spinelli et al. 1968) encourages such experiments. Psychophysical studies of learned cross-modal effects upon grouping would be of particular interest given the evidence suggesting that contextual interactions play a major role in grouping.

In sum, there are large effects of learning in both mature and immature sensory systems. These effects often include changes in RF connectivity, but changes in CF connectivity are also likely, as expected on the grounds that the contextual predictions would otherwise be invalidated by RF changes. Finally, psychophysical experiments designed specifically to study changes in contextual integration due to learning are both possible and worthwhile.

5.7. Local contextual integration in the perception of objects and words

5.7.1. Dynamic grouping in the perception of objects.

Ways in which binding through synchronization could help produce object recognition performance similar to that of humans has already been discussed in detail elsewhere (e.g., Hummel & Biederman 1992; Mozer et al. 1992), so we restrict ourselves to a brief outline. One central idea is that shape recognition could generalize well across irrelevant dimensions, such as position, if shape descriptors are insensitive to those dimensions. Synchronization is used to bind shapes to positions in order to show what shapes are where. The Hummel and Biederman (1992) model has seven layers through which image features are combined into parts, and then into structural descriptions of objects in terms of their parts and relationships. Synchronization is used to group image features into volumetric parts, and to bind parts to relationships, and is achieved through a network of fast enabling links that are similar to our CFs. Crucial aspects of human performance displayed by the model include recognition that generalizes well across position, size, left-right reflection, and rotation in depth, but poorly across rotation in the picture plane. Goebel (1993) developed a similar model that has more flexible synchronizing connections and which also incorporates mechanisms that produce performance consistent with psychophysical evidence for selective spatial attention. As

in human performance, such systems are highly dependent upon internal grouping processes so that the same image grouped in different ways can give rise to very different outputs. These demonstrations of computational feasibility and similarity to human performance are encouraging, but direct tests of the hypothesis that timing on a fast time scale is important for grouping in object perception are also needed, and this is a major goal for further research.

5.7.2. Contextual integration in the perception of words.

Word perception is a major focus for studies of contextual integration within cognitive psychology, so here we discuss several ways in which the two areas of research can be related. We note similarities between context effects in word perception and in the perception of simple line elements; we conclude that as the latter can be studied by both psychophysical and physiological methods, this may bring a rich new source of evidence to bear on cognitive conceptions of contextual integration. We relate this to dynamic grouping and to synchronized population codes, and outline neuropsychological evidence for such codes in word perception. Finally, we note a possible role for contextual guidance in learning to perceive words.

Various paradigms have shown that both letter and phoneme perception depend upon local context. Cattell (1886) showed that letters are recognized more accurately in the context of a familiar word. The many subsequent studies of such phenomena include demonstrations that forced-choice discrimination between a pair of pre-specified letters is better if the test letter appears within the context of a familiar word or pronounceable nonword (e.g., Johnston & McClelland 1973; Reicher 1969; Rumelhart & McClelland 1982). Strong effects of local context also occur within speech perception. For example, Massaro and Cohen (1983) presented computer generated syllables such as /sli/, /tri/, /sri/, and /tli/, to subjects and asked them to classify the middle phoneme as an /l/ or an /r/. This phoneme was presented at seven different levels on a continuum from being very /l/-like to being very /r/-like by varying the frequency from which the third formant (F3) began at the phoneme's onset. Each was factorially combined with one of four different leading consonants, /s/, /t/, /p/, and /v/. Perception of the central phoneme was predominantly determined by the direct acoustic cue to that phoneme as given by F3 frequency, but it was also affected by the preceding consonant, particularly when the direct cue was most ambiguous. Similar effects have been shown in studies of reading. For example, when an ambiguous lower-case letter that can be read as either an *e* or a *c* is placed in contexts that support one or the other alternative then identification is biased towards the contextually appropriate alternative (Massaro 1979).

These effects are similar to those in the perception of simple line elements displays (e.g., Kapadia et al. 1995; Polat & Sagi 1993, 1994a; 1994b) in several ways. Although they occur at a higher level of analysis, context effects in word perception also occur rapidly and automatically. The effects are locally specific in that they depend upon the particular properties of the entities that are interacting. For example, the context /s/?/i/ supports some target letters and not others, just as the context of two collinear oriented line elements supports some intervening oriented line elements and not others. Furthermore, in both cases the interactions are such as to emphasize things that are ex-

pected in that context, rather than things that are unexpected. In both domains the effects of context on the perception of individual elements seem to be greatest when the direct stimulus evidence is most ambiguous. Yet another similarity is that contextual interactions in both domains are affected by learning. Finally, there are effects of object-specific knowledge on the detection of line segments (e.g., McClelland 1978; Weisstein & Harris 1974) that may be analogous to the effects of word-specific knowledge on the forced-choice discrimination between letters. At present there seems to be enough similarity, therefore, to justify the search for a common explanation for these various effects of context.

Theories of contextual interaction in word perception have been undergoing vigorous development for many years and can now account for many details of performance with impressive precision (e.g., Massaro 1989a; McClelland & Elman 1986; McClelland & Rumelhart 1981; Richman & Simon 1989; Rumelhart & McClelland 1982). The problem is not that no explanation is available but that each of several diverse theories can fit the data so well that it is difficult to know what inferences to draw. Thus we need to find common elements, or fundamental mechanisms, that may appear in different forms in the different theories and which are crucial to the cognitive functions with which they are concerned (Richman & Simon 1989). As the theories all aim at generality, one way to do this is to widen the range of phenomena to which they are applied, including psychological tasks that can be related to known physiological mechanisms if possible. It is particularly appropriate to relate Interactive Activation and Competition (IAC) theories (e.g., McClelland & Rumelhart 1981) to anatomy and physiology because they explicitly use a neural style of computation and played a major role in promoting the use of that style in cognitive theory. Basic aspects of these models that are broadly compatible with what is known about cortical physiology include: having just a few different levels of analysis; having many replications of processors spanning feature space at different input positions within a level; and having local inhibitory relations to force a choice between incompatible alternatives within levels. For further in-depth discussions of these theories and related issues, see McClelland (1991), Massaro and Cohen (1991), and Movellan and McClelland (1995).

We now discuss five major unresolved issues from the perspective of our general approach and of the analogy between context effects in the perception of words and of simple line elements: (1) What is the architecture of information flow, and in particular to what extent do the effects of context depend upon the feedback of information from higher to lower levels of analysis? (2) How does contextual information affect processing, and in particular does it do so in essentially the same way as direct evidence from the target? (3) Do the mechanisms producing the effects of context on the perception of ambiguous or just detectable stimulus elements also play a role in dynamically grouping those elements? (4) To what extent does each level use local as opposed to population codes for those entities with which it is concerned? (5) Is the goal of maximizing coherence between distinct streams of processing relevant to learning within streams?¹⁰

1. What is the architecture of information flow, and in particular to what extent do the effects of context depend upon the feedback of information from higher to lower

Levels of analysis? The detailed properties of context effects in the perception of line elements strongly suggest that they are due in part to long-range horizontal collaterals that directly link distinct entities within the same level of analysis (e.g., Kapadia et al. 1995; Singer 1995). If there are connections that are specialized to mediate contextual interactions within the segmental (i.e., phonemic or letter) levels of word perception, for example, then this could contribute to the superiority of regular over irregular non-words and may help explain some of the more rapidly occurring components of contextual interaction. Contextual connections within segmental levels may not be the best way to explain the effects of word-specific knowledge, however, and these are often assumed to involve activity at a higher level of analysis that is specialized to distinguish words. The question that then arises is whether activity at this level influences processing at the segmental levels. This issue is unresolved but an analogous issue can be studied at lower levels of visual processing using physiological techniques. Studies of the role of feedback from V1 to LGN provide evidence that it synchronizes the firing of those LGN cells that combine to form higher level entities in V1 (Sillito et al. 1994). This suggests that higher levels do influence processing at lower levels but in a way that is distinct from the ascending RF input and more like the process of contextual integration within levels. As feedback connections are ubiquitous within the cortex, this could be a general feature of cortical processing, so it would be worthwhile obtaining further evidence on it by combining psychophysical measures with electrophysiological measures in V1 as in Kapadia et al. (1995), but measuring both rate and relative timing in multiple single unit recordings, and using stimulus elements that either do or do not combine together to form a single familiar entity at a higher level. The possibility that a major role of feedback is to group activity at lower levels will be examined further under point 3 below, which discusses grouping processes in word perception.

2. *How does contextual information affect processing, and in particular does it do so in essentially the same way as direct evidence from the target?* If not then contextual integration is not a fundamental issue within that domain. Note that as contextual interactions may be reciprocal the question is not whether some parts of a word are processed just as targets and others just as contexts, but whether words are composed of distinct but interacting parts such that the processes by which they interact differ from those by which they are kept distinct.

One of the clearest ways of distinguishing between contextual and direct stimulus effects in word perception is that it is the direct stimulus input and not the context that determines the alternatives between which choice is made; that is, context influences the choice between those competing alternatives for which there is direct but ambiguous stimulus support.¹¹ As evidence for this view Massaro (1989b) notes that context does not by itself produce the phoneme restoration effect (Warren 1970) because some bottom-up support for the presence of a missing phoneme is required, even if only in the form of a brief noise burst. This possible asymmetry in the roles of context and target is not apparent in the experiment of Massaro and Cohen (1983) because target support was always available. It should therefore be possible to make it more apparent by presenting stimuli such as /li/, /ri/, /si/, /ti/, /sli/, and /tri/

with various levels of ambiguity for the /l/ or /r/ phoneme and asking subjects to decide whether an /l/ or /r/ or neither is present. The distinction proposed here predicts that context will have a large effect in the presence of an ambiguous target while having little or no effect in its absence. In contrast, the direct stimulus cues should have large effects whether the context is present or not. Such a result would not be surprising but it would show how the asymmetry in the roles of RFs and CFs can be reflected in performance. Furthermore, this asymmetry in dependence would contrast with the various ways in which information from different sources can be combined such that they all influence decision in essentially the same way (e.g. Massaro & Friedman 1990), as do different sources of information from within the RF. An important aspect of this latter form of combination is that although the different sources may be independent prior to combination, their individual contributions are not kept distinct in the output decisions produced. Thus, in contrast with the effects of CF inputs, all RF inputs help determine the meaning of the decisions to which they contribute.

The distinction between the roles of RFs and CFs may also be relevant to the long-running debate between theories that emphasize the use of internal knowledge to go beyond the input from external stimuli and theories that emphasize remaining faithful to that input so as to avoid hallucination. The latter danger is often noted in discussions of the effects of context in word perception (e.g., Massaro 1989a; Massaro & Cohen 1991) and can be used as an argument for assuming that context and target do not interact. If context has distinct effects upon processing, then instead of having to choose between avoiding hallucinations and allowing contextual interaction we can have both, including direct contextual interactions within levels, which might otherwise overwhelm stimulus processing with hallucination.

3. *Do the mechanisms producing the effects of context on the perception of ambiguous or just detectable stimulus elements also play a role in dynamically grouping those elements?* Sections 5.2 and 5.3 outlined evidence suggesting that for simple line element displays the grouping of elements into coherent wholes depends upon the same knowledge embodied in the same mechanisms as do the effects of context on the perception of the individual elements. If this is so and if contextual integration is achieved in the same way in different regions, then it will also apply to word perception. We take it for granted that grouping processes are a crucial part of word perception at both lexical and sublexical levels, and this is easily demonstrated. At the lexical level, for example:

this must be grouped dynamically using knowledge of specific words.

Internal grouping processes also occur at sublexical levels. PIGHAM, for example, will be pronounced either with or without consonants in the middle depending upon whether or not IGH is grouped to form one grapheme. If grouping is computed dynamically, as it must be if it is signaled by relative timing, then such groupings could change rapidly from moment to moment, thus making various possible alternative groupings successively available. Studies of a neurological patient, who will be discussed further below, illustrate the relevance of such dynamic grouping processes to word perception. She was quite unable to read PIGHAM

as a single unfamiliar nonword but could read it easily when she saw it as two familiar words (Goodall & Phillips 1994). She made it clear in various ways (e.g., by drawing a pencil line between the appropriate letters) that this involved feedback of grouping information from a lexical level to a level containing a precise topographic map of the individual letters.

There is also evidence from normal subjects that the effects of word familiarity on perception involve internal grouping processes. For example, familiarity reduces asymmetrical left-to-right letter position effects that can be explained as being due to processing letters separately in unfamiliar stimuli (Phillips 1971). Many other effects can also be explained as being due to processing familiar items as a single coherent whole, or "chunk," but processing unfamiliar items as a number of separate chunks (Richman & Simon 1989).

One implication of these considerations is that in order to understand the role of feedback in word perception it may be worthwhile emphasizing tasks that reveal the effects of grouping processes. For example, if grouping and disambiguating involve a common mechanism then disambiguating interactions between elements should depend upon whether those elements are grouped together or not. Consider the disambiguating effects of local context in the experiments of Massaro and Cohen (1983), for example. In those experiments the phonemes that interacted were always part of a single word. If more than one word were presented then it would be possible to test whether or not the interaction between neighboring phonemes depends upon their being perceived as parts of the same word or phrase.

4. *To what extent does each level use local as opposed to population codes for those entities with which it is concerned?* A fundamental difference between these two possibilities is that population codes can transmit information about inner structure but local codes cannot. This difference was used to provide evidence on this issue by studies of two neuropsychological patients whose ability to read and write is very largely restricted to words with which they are familiar (Goodall 1994; Goodall & Phillips 1994; Phillips & Goodall 1994). Their reading and writing therefore provides a direct window on the contribution of lexical knowledge when isolated from sublexical processes that treat the input as a string of separate letters or phonemes. Several experiments were run with the general format of giving the patients visual discrimination training on a set of nonwords and then testing their ability to write those nonwords to dictation as compared with matched but unfamiliar nonwords.¹² These experiments show that the patients can write visually familiar nonwords to dictation accurately and fluently on first hearing them spoken aloud even though they have never written them before and cannot write any of the unfamiliar nonwords. This is a surprising result for theories proposing that the role of word recognition is to produce a local code that indicates which word has been recognized, because that cannot explain how familiarity makes a description of internal structure available for matching with input from other modalities and for generating appropriately structured output. These results do not show that local codes do not exist for familiar words, but they cannot be explained by proposing such codes.¹³ Furthermore, if the output of word recognition systems is a structured description then this could help explain how

output lexicons could obtain knowledge of word structure. If it is assumed that they just receive local codes from input lexicons and then generate the appropriately structured output then it is not obvious how they acquire knowledge of those structures.

These findings therefore support theories that use distributed or population codes for familiar words, but other evidence outlined above suggests that they are processed as single whole items or chunks. How can these two conclusions be reconciled? If familiar words are given a local code then we have difficulty explaining the transmission of information about their structure. If we assume that they are given a classical distributed code then we have difficulty explaining the effects of familiarity adequately, because such codes can transmit descriptions of novel items as well as of familiar items.¹⁴ Theories using distributed codes that are formed by internal grouping processes may be able to overcome both difficulties because then the way in which elements of a distributed pattern of activity are combined depends upon internal grouping processes that embody knowledge of familiar or coherent combinations. This assumes a highly distributed input to all stages of cortical processing, and emphasizes the importance of being able to separate that input into distinct subsets on the basis of their internal coherence. It also suggests the possibility that information that is conveyed as a number of separate groupings at one level or stage of processing may be conveyed as a single coherent grouping at later levels or stages.

5. *Is the goal of maximizing coherence between distinct streams of processing relevant to learning within streams?* The effects of context on disambiguation and grouping in word perception imply that the contextual predictions are learned, but do contextual inputs from other streams also influence feature discovery within streams? Possible examples of such an influence are the effect of learning to read on phonological awareness (Bentin et al. 1991) and of phonological awareness on learning to read (Bentin 1992). Another example is provided by Kanevsky (1989), who gave subjects with impaired hearing supplementary tactile input generated from the speech signals that they were lip-reading. Their performance when lip-reading alone was significantly improved by several months of training with the coherent cross-modal inputs. Thus processing of features within one stream was affected by their correlations with features in another stream.

The possibility that coherence between distinct streams could be of relevance to speech processing has been studied computationally by Becker (1996) and de Sa (1994a; 1994b). Both emphasize the need for an algorithm that can discover the features that distinguish words in the absence of any external supervisor that tells the network what those features are, and both propose that such an algorithm could use the statistical dependency between distinct streams of processing. Becker applied the I_{max} algorithm to the Peterson-Barney data consisting of the first and second formant frequencies of 10 vowels spoken by many different speakers. She showed that from just these two low-level acoustic features alone the I_{max} algorithm could discover higher-order features that could be used to classify the vowels with better than 75% accuracy.¹⁵ Cross-modal interactions have been studied by de Sa and Ballard using the task of learning to categorize consonant-vowel syllables that were both heard and seen as a pattern of mouth move-

ments. Tightly synchronized audio and video recordings of 5 speakers saying 5 vowels were fed into two separate streams of processing, which learned to classify their inputs by minimizing disagreements between the two streams. When then tested on each stream of input alone, syllable categorization accuracy was better than 90% for the audio input and better than 75% for the video input (de Sa & Ballard 1997). These demonstrations are valuable because they show how learning by discovering statistical dependencies between separate streams of processing can discover useful features within streams when applied to real-world speech processing tasks.

6. Issues arising

If there are common foundations for cortical computation then a central goal for neurobiology must be to discover what they are and how they are embodied in cortical structures and processes. The primary issues that arise therefore concern the justification of the search for common foundations and the promise of the various means by which common cortical organization, common computational primitives, and common requirements of cognitive subsystems may be revealed. We hope to encourage discussion of these general issues, as well as of our more specific findings and hypotheses. Many issues arise in regard to the latter, such as ways in which contextual integration and synchronization can be related to attention (e.g., Tiitinen et al. 1993), imagery (e.g., Ishai & Sagi 1995), and consciousness (e.g., Crick & Koch 1990). Here we note just those most closely related to the aspects that we have emphasized above.

6.1. Does the distinction between RFs and CFs have biological relevance?

The possibility of distinguishing between RF and CF inputs raises a number of specific questions: (1) Can RF and CF inputs be distinguished in the cortex on the basis of the synaptic receptor channels that they activate? One possibility is that CF inputs are more dependent upon voltage-dependent channels than are RF inputs, but more research is required on this issue. (2) Can RF and CF inputs be distinguished on the basis of the anatomical distribution of their input sites? For example, at least one source of contextual input, the long-range intraregional tangential connections, seems to preferentially contact the apical dendrites of their target pyramidal cells (Gilbert & Wiesel 1983; Kisvarday et al. 1986). (3) If there are contextual connections then do they affect only timing, without having any effect upon the total number of spikes that are produced by the cell or local group of cells, or can they affect both (König et al. 1995)? One possibility is that the answer depends on the circumstances, with effects upon the number of spikes, for example, being more likely with weak RF input.

6.2. Does the distinction between RFs and CFs have psychological relevance?

Section 5 suggested several paradigms through which this distinction might be reflected in behavior. Some of these ask whether the effects of appropriately selected contextual

variables depend upon the target to a greater extent than the effects of the target depend upon them. From an information-theoretic perspective they seek conditions under which behavior transmits information about the target rather than about the context, even though it is influenced by context. One specific test for this is to see how the effects of context depend upon the strength of the target evidence. If the putative contextual variable does not modulate transmission of information about the target but just contributes directly to the response decision itself then the effects of context will *decrease* as the strength of the target variable increases (Massaro 1989a; 1989b; Massaro & Friedman 1990). If it does have a modulatory influence, however, then the effects of context will *increase* as the strength of the target variable increases from low to medium values (Smyth et al. 1996). We have also suggested that some aspects of contextual integration in higher cognitive functions, such as word perception, might be related to cortical neurobiology by analogy to studies of the perception of simple line element displays that can combine both psychological and physiological techniques. Whether the research that this suggests will reward the investment required remains to be seen. However this turns out, it is hard to see how the distinction between RFs and CFs could fail to have psychological relevance if it is biologically valid.

6.3. Does the cortex use synchronized population codes?

Various doubts have been raised concerning the functional role of synchronization, but biophysical and theoretical arguments suggesting that it has a major role have already been noted (e.g., Abeles 1982; 1991; von der Malsburg 1981), and sections 4 and 5 outlined further support for this, including evidence that synchronization and behavior are closely related and that the stimulus conditions producing synchronization have detailed similarities to those producing grouping in perception. Further evidence could be obtained on this issue by experimentally manipulating the synchronization of spikes on a fine time scale to see whether this then has consequences for behavior. This is technically difficult but one way in which it might be done is by studying how the modulatory effects of context depend upon the precise temporal relationships between target and context. For example, a paradigm for distinguishing modulatory from direct effects such as the one described by Smyth et al. (1996) could be used in a texture segregation task in which there are both target and contextual cues to segregation. The stimulus parameters used to specify the target and contextual boundaries could both change rapidly, and with various phase relations. The central question would then be whether the context preferentially modulates the detection of target boundaries with which it is synchronized to within a few milliseconds. To improve the chances that the external timing relations imposed upon the cues will give adequate control of the timing of internal activity the stimuli should be presented under conditions that maximize the use of fast high temporal resolution pathways (e.g., low spatial frequencies, dark adapted Ss, etc.). Such studies are encouraged by results showing that in texture segmentation tasks temporal differences in stimulus onset of as little as 10 msec can be used to segregate figure from ground (Leonards et al. 1996).

6.4. What form or forms of self-organization occur within the cortex?

As noted in section 3.5, algorithms that learn by maximizing coherence across streams can discover useful higher-order functions (Becker 1996) but may nevertheless be more limited in what they can learn than supervised algorithms such as error backpropagation (Rumelhart et al. 1986). This does not necessarily make them less biologically plausible, however, and would make them more plausible if their limitations were shared with the processes of cortical self-organization. Unfortunately, those limitations are not yet clear. Some studies suggest that the human visual system has great difficulty learning arbitrary nonlinear functions such as the XOR (Thorpe et al. 1989), but more evidence is required. Furthermore, in cases where such difficult learning problems are solved, this might be because the input is transformed in a way that makes the problem easier (Clark & Thornton 1997), so we need to know how any input data is recoded internally in order to know what internal learning problems have been solved. Major tasks for future research are therefore to determine what functions can and cannot be learned by cortical self-organization and to compare them with the capabilities and limitations of algorithms derived from computational theory.

Four well-specified learning rules that have been proposed as a basis for cortical self-organization turn out to have a close family resemblance to each other even though they were derived in different ways (i.e., those of Bienenstock et al. 1982; Der & Smyth, in press; Hancock et al. 1991; Kay & Phillips 1994; 1997). They all relate the change in synaptic strength to postsynaptic activity in approximately the nonmonotonic way shown in Figure 5. Furthermore, there is direct physiological evidence for a dependence of this general form (Artola et al. 1990; Dudek & Bear 1992; Kirkwood et al. 1996; Singer & Artola 1994). Intrator and Cooper (1995b) provide detailed arguments for the computational value and biological plausibility of such rules, and for the view that they are common to the hippocampus as well as to both mature and immature cortex. A key feature of all of these rules is the threshold of postsynaptic activity above which connections are strengthened, and they differ in the way that this is specified. Some keep it fixed, others move it dynamically as a function of prior activity, and the latter differ in what aspects of prior activity are used and in how they are used. It is possible that various versions of this rule exist, with simpler and less computationally powerful ways of specifying the threshold being found in some species and/or neural subsystems, and more complex and powerful ways being found in others. There is physiological evidence for a dynamically moving threshold (Huang et al. 1992; Kirkwood et al. 1996), and investigations of possible molecular bases for it have already begun (Mayford et al., in press). The detailed properties of this threshold and its relation to theories of learning is clearly a major issue for neurobiology.

Does context affect RF learning? Gilbert and Wiesel's (1990) observation that prior contextual stimulation altered the orientation tuning function that was later obtained in response to receptive field stimulation alone suggests that it can, but these effects were not studied systematically and their causes are uncertain. More direct study of this issue is therefore required, and ways in which

this can be done were suggested in sections 4.3, 5.6, and 5.7.2 (point 5).

6.5. To what extent do these processes provide a representation of the external world?

The ability to form representations of the external world is often thought to be fundamental to cortical computation. We have argued that networks of local processors with contextual guidance can in effect discover distal variables and relationships by discovering mutual information in diverse data-sets. This may not be the same as forming explicitly intentional representations of the external world, however (Phillips et al. 1995a). "Representation" is commonly used to describe what sensory and perceptual systems do, and in that general sense it is synonymous with "the transmission of information." It also has a more specific meaning, however, such that the distinction between representation and referent is critical. Intentional representation is using one thing, the representation, in the place of another thing, the referent. It implies a user that knows about and distinguishes both the representation and the referent. The relation between representation and referent can be iconic, symbolic, or both, and the relationship is asymmetrical. This asymmetry does not apply to the notion of information transmission, however, which is defined as the mutual information that is shared between input and output.

Representation proper plays an important role in human cognition, but much cortical function may proceed without it, nevertheless. Phylogenetic studies of such representational abilities show them to be rare or absent outside the primate line and to emerge progressively within it (e.g., Byrne & Whiten 1992; Chavalier-Skolinkoff 1983). In children, different aspects of this skill develop during years 1 to 6, gradually becoming more flexible, internalized, and differentiated (e.g., Campbell & Olson 1990; De Loach 1987; Piaget 1954; Wimmer & Perner 1983; Zaitchik 1990). We know of no physiological observations showing that local cortical processors treat their inputs as standing for something other than themselves. Furthermore, we know of no neural network model that puts intentional representation into the dynamics of the network, rather than just leaving it in the mind of the designer. Such models often have considerable computational power, nevertheless, even though they do not incorporate any proper intentional representations of an external world.

We conclude that representation proper is not a common feature of cortical computation, but arises late in both phylogenetic and ontogenetic development, and perhaps upon foundations such as the implicit form of realism that has been hypothesized here to result from the maximization of coherent variation by local cortical processors. If so, a better understanding of the capabilities and limitations of those foundations will help us understand how higher cognitive functions are possible, and why they are needed.

ACKNOWLEDGMENTS

This work is supported by a Network grant from the Human Capital and Mobility Program of the European Community. Special thanks to Jim Kay, Dario Floreano, and Darragh Smyth for the computational studies outlined in section 3. We thank our many colleagues at Stirling and Frankfurt for their crucial contributions to this work. We also thank Moshe Abeles, Elie Bi-

enstock, Ralf Der, John Hertz, Geoffrey Hinton, Nathan Intrator, Michael Jordan, Ralph Linsker, Tim Shallice, Olaf Sporns, Tom Troscianko, John Taylor, Christoph von der Malsburg, and Rolf Würtz for valuable discussions of these issues. We also thank Sue Becker and two anonymous referees for their helpful and insightful suggestions.

NOTES

1. For further discussions of the advantages of such “constructive” effects of contextual guidance, or reentrance, upon RF selectivity, see Finkel and Edelman (1989) and Tononi et al. (1992b; 1996).

2. From a statistical point of view these learning abilities combine the descriptive aims of techniques such as principal component analysis with the predictive aims of multiple regression. They are therefore related to techniques of latent structure analysis, such as canonical correlation (Giffins 1985; Hotelling 1936), which seek functions defined upon separate data-sets such that the correlation between them is as large as possible. Canonical correlation can in principle be implemented within a neural network (Kay 1992), but it is concerned with linear functions of just two data-sets. The computational studies in section 3 extend it to the multistage analysis of multiple data-sets.

3. Further examples of this approach are provided by Zemel and Hinton (1991), who show how it can discover viewpoint-invariant relationships that characterize objects, and by Becker (1993), who shows how it can learn to compute translation-invariant object categories using just the continuity of an object's identity across time.

4. One of the many examples that Gallistel cites is that migratory birds learn the position of the celestial pole while they are fledglings (Able & Bingham 1987). He argues that this requires a procedure that is specially designed to compute and store the values of variables that specify the position of the celestial pole in the pattern of dots defined by the circumpolar stars. He concludes that “we should no more expect to find a general-purpose learning mechanism than we should expect to find a general-purpose sensory organ” (Gallistel 1995, p. 1266).

5. The goals of Infomax and of Coherent Infomax differ only in the value given to ϕ_1 . This therefore suggests an important role for this parameter. It specifies the balance between increasing information transmission within streams of processing and increasing predictability across streams. It is analogous to the parameter “eta” in the algorithm for discovering predictable classifications developed by Schmidhuber and Prelinger (1993). An advantage of specifying the relative priority to be given to these two goals by a single parameter is that it provides for a continuous transition between them. Appropriate values for this parameter might then vary with the layer of processing and stage of learning.

6. How these two conditional probabilities could be estimated in a biological system is not yet clear, so a simpler rule such as that proposed by Der and Smyth (in press) or by de Sa (1994a; 1994b) may be more plausible. On the other hand, it can be argued that we often underestimate the computational power of individual neurons, so an approximation to the more complex rule may be feasible.

7. These demonstrations show how CF input increases the probability of at any moment producing outputs that are mutually supportive of each other. They thus generate grouping through synchronization. These effects are compatible with temporal structure in the outputs arising from any of a variety of sources, including periodicity arising at the level of single cells, local circuits, or larger populations. Any temporal variations in the outputs will be transmitted through the CF connections in such a way as to increase the probability that coherent outputs will be active simultaneously.

8. Examples of such nets are those incorporating reentrant connections (e.g., Sporns et al. 1991; Tononi et al. 1992b).

9. Tests for modulatory effects can be specified in terms of the mutual and conditional information between target, context, and

response (Smyth et al. 1996). These tests can be used in cases where context and target variables are correlated, and where the context either does or does not also have a direct influence on response. One distinctive sign of modulatory effects is provided by an *increase* in $I(X;C|R)$ as target strength increases from low to medium values. $I(X;C|R)$ represents the information transmitted about the context in addition to any information transmitted about the target; this would be intuitively expected to decrease with increases in target strength. To test for an increase in $I(X;C|R)$ we estimate the probability of a correct response in a two-alternative forced choice task in the presence of either supporting or nonsupporting context at two different levels of target strength. If the effect of context *increases* with target strength then the context is modulating transmission of information about the target.

10. Other major issues that need to be resolved in order to provide an adequate account of context effects in word perception include the following: (1) What are the effects of noise in the input and noise generated from within the system itself (Movellan & McClelland 1995)? (2) How is overt responding related to the different levels of analysis? In particular, to what extent do segmental decisions (i.e., decisions concerning letters or phonemes) reflect the activity at segmental levels and to what extent at word or morphemic levels? (3) What decision processes are used to generate overt responses from the information made available by the speech perception system (Massaro & Friedman 1990)? (4) What is the time course of context effects upon processing? In particular, how soon do these effects become apparent after stimulus onset?

11. Another way of distinguishing contextual influences may be by proposing that target information sets limits upon discriminability whereas context just adds bias. This is not identical to our emphasis upon contextual interactions that do not corrupt the meaning of the signals that interact because information transmission depends upon both sensitivity and bias but it is similar, and much of the evidence from studies of word perception can be seen as supporting it (e.g., Krueger & Shapiro 1979; Massaro 1989a; Massaro & Cohen 1983; Morton 1969). The evidence is not wholly unambiguous on this issue, however, because accuracy in forced-choice discrimination between letters depends upon context (e.g., Johnston & McClelland 1973; Reicher 1969), and because target discriminability, as measured by d' , is also affected by context in some other paradigms (e.g., Phillips 1971; Samuel 1981; 1996). Furthermore, it is difficult to draw strong inferences about internal processing from such phenomena alone because whether they are compatible with certain aspects of a model (e.g., the architecture) depends upon other factors (e.g., the role of noise; Movellan & McClelland 1995).

12. Patient AN had a stroke in 1979 when she was 47; she has been studied at Stirling since 1985 (Goodall & Phillips 1994, Phillips & Goodall 1994). Patient AM is her nephew; he had a stroke in 1990 at the age of 35, leaving him with very similar deficits (Goodall 1994). The mental status of both patients is good; they are cooperative and insightful during testing. In one study (Goodall 1994), five 4-letter nonwords were created and printed on a sheet that was placed in front of the patient, who had to say whether each of a number of separately presented test items was on the sheet. To enforce accurate visual discrimination the test items included nonwords differing from the training items by just one letter. This training proceeded until performance was accurate, which took about 10 minutes. At no time did either the patient or the experimenter speak or write the nonwords; indeed, the patients could not accurately read them aloud, either before or after training. After a 5-minute break the five familiar nonwords and five unfamiliar nonwords were read for writing to dictation. The familiar nonwords were all written correctly and fluently, but none of the unfamiliar nonwords were written either correctly or fluently. The familiar nonwords were read to 42 normal subjects for writing to dictation, but without any prior presentation of their written form. The spellings produced by these subjects were quite different from those given by our two patients, because unlikely but plausible spellings were chosen when selecting the stimuli to

be used with the patients. This confirms that the patients were using the knowledge obtained from the visual discrimination training when writing the familiar nonwords to dictation.

13. It may be possible to preserve the view that the normal processes of word perception use local codes for words by proposing that these results reflect the use of voluntary visual imagery or episodic memory strategies. Several aspects of the results greatly weaken such interpretations, however. First, neither patient ever claimed to be visualizing. Retrieval of the learned information was fast, fluent, and automatic in the sense that the patients showed no sign of cognitive mediation at retrieval. Indeed, both were surprised at their ability to write the nonwords and on one occasion AM, quite unprompted, said "Now how did I do that?" Second, the information transmitted from the visual input to the written output was not at the level of a visual image because the patients transcribed printed text into their own cursive script. Third, knowledge of the new items was acquired gradually and retained over periods of months. This is not characteristic of episodic memory functions. Fourth, writing to dictation is a typical implicit memory task and was used as such here. No reference was made to previous presentations. The patients were simply asked to write down what was said to them. This is exactly the kind of task that relies on procedural skills and is consequently unimpaired in patients with deficiencies in episodic memory (McCarthy & Warrington 1990; Shallice 1988). Fifth, familiar English words occurred as error responses to unfamiliar nonwords, but the familiar nonwords did not. This shows that nonwords were not being processed in a way that kept them separate from knowledge of familiar English words.

14. This difficulty is made clear by studies comparing patient AN's ability to read and write familiar items with her ability to read and write closely related but unfamiliar items. She can read familiar words such as NAPKIN and MUSKET but not unfamiliar nonwords such as NAPKET and MUSKIN made by recombining their component syllables. This is unlikely to be due to the absence of semantic mediation because she can learn to read meaningless nonwords with which she is familiarized (Goodall & Phillips 1994). She can also learn to copy meaningless nonwords such as BONSED and MUNIZE, but this ability does not transfer to the copying of BONIZE and MUNSED (Phillips & Goodall 1994). If both kinds of item are processed as distributed codes that are treated as a single group, whether familiar or not, then analysis and simulation both suggest that there will be good transfer to novel items that share components with familiar items. This applies to both feedforward networks (e.g., Baldi & Hornik 1989; Brousse & Smolensky 1989; Phillips et al. 1993) and recurrent architectures with attractors (e.g., Plaut & McClelland 1993). Thus, these studies of AN suggest that the letter strings were processed as distributed codes to which internal processes were applied to form familiar or coherent groupings. Computational studies of coding through synchronization in the dynamic link architecture demonstrate the possibility of mapping structured descriptions as familiar wholes while specifically avoiding false combinations of their component subsets (e.g., Lades et al. 1993; von der Malsburg 1988), but whether this approach can be developed to account for word perception in detail remains to be seen. A simpler alternative possibility is that familiar items are signaled by distributed codes of the kind used in most connectionist theories but with the addition of internal grouping processes that segregate unfamiliar combinations into distinct subsets that are processed separately.

15. Becker (1996) applied the Imax algorithm to the task of learning to categorize phonemes by presenting each of two streams with different versions of the same vowel. We assume that this design could easily be modified to explore the use of sequential dependencies between successive acoustic inputs of different phonemes to learn how to categorize them. Guiding learning within streams by coherence across streams does not require either that the coherence be perfect or that the different streams process different instances of the same thing.

Open Peer Commentary

Commentary submitted by the qualified professional readership of this journal will be considered for publication in a later issue as Continuing Commentary on this article. Integrative overviews and syntheses are especially encouraged.

Is synchronization necessary and is it sufficient?

Daniel J. Amit

*Istituto di Fisica, Università di Roma, La Sapienza, Rome and
Racah Institute of Physics, Hebrew University, Jerusalem, Israel;
amit@jupiter.romal.infn.it*

Abstract: The strong coupling of binding to cross-correlations is methodologically problematic. A completely unstructured network of neurons can produce cross-correlations very similar to the measured ones, and yet they have little dynamic effect.

If stimuli in separate receptive fields actually integrate to an instantaneous whole, one is faced with a complex dynamical problem at the neurophysiological level. Less speculative is the fact that in primary visual cortex, neuronal spiking activity exhibits a certain level of synchrony just when the separate stimuli can be considered parts of a whole. We need little convincing and no neurophysiology to agree that context must be part of perception (as well as of more elaborate mentation) to supplement and give sense to the fragmented way in which we receive the "world."

The issue that Phillips & Singer's (P&S's) target article addresses seems to be that a concatenation of the view that context must function and that the Singer-Gray-Koenig-Kreiter synchronization has been observed, brings new insight, either to cognitive science or to neurobiology. I do not find this very convincing, despite the weighty evidence represented by some 260 learned references.

Clearly if a well-defined cognitive problem can be solved by a specific neurophysiological phenomenon, then its solution must be known *a priori*. The discovery of the physiological fact can at best demonstrate that a device exists in the biological brain that may be the vehicle for the acknowledged solution. To argue that a given physiological fact is either required for the solution of the cognitive problem, or that it is the physiological fact that has given birth to the solution, is a rather risky undertaking.

First, arguing that a physiological fact is necessary for a solution exposes one to the vicissitudes of complex dynamics and treacherous numbers: it is argued that synchronization makes neural responses more effective in dealing with the dynamical question of how synchronizations, when they arise, are propagated from where they are observed to where they are needed as a solution. The target article quotes Koch and collaborators (Bernander et al. 1994) to argue that synchronized signals are effectively propagated. But this study is rather ambiguous on the issue, and in different regimes of parameters the system may make synchronization either more or less effective than unsynchronized inputs of the same spike rate.

Should one conclude from the fact that synchronization seems necessary for a solution to a cognitive problem, that the parameters in the cortex are in the right regime? Or conversely, that if the parameters of cortex are unfavorable, the cognitive problem has no solution?!

Is the cognitive problem solved for any level of synchronization (as measured by the peak in the cross-correlations of neurons (see, e.g., Kreiter & Singer 1996) or is some low threshold value required? I will try to indicate below that there are reasons to believe that the synchronizations observed in these remarkable experiments can be considered small.

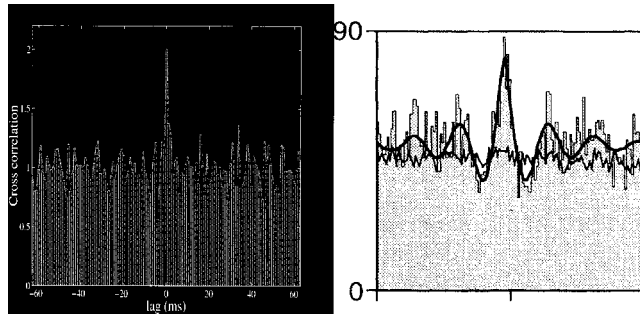


Figure 1 (Amit). Right (experiment): averaged cross-correlation from Figure 7 of Kreiter and Singer 1996. Left (simulation): cross-correlation (averaged over eight 5-second runs, the same two cells recorded simultaneously and normalized to the product of the two rates) from a simulation of a network of 7,500 cells in spontaneous activity. In both, the coincidence bin is 1 msec. The rates on the left are 50–60Hz, on right 30–35Hz.

P&S's position seems to be that the fact that synchronization solves an outstanding problem in cognitive psychology is not sufficient reason for its existence. It is supplemented by an argument of maximum information transmission. This again seems to raise I think more problems than it solves:

1. Is binding related to maximum information transmission? Or is it a coincidence?
2. What is the space in which this maximization is affected?
3. Is information transmission maximized at any level of synchronization and any level of background noise?
4. To maximize information rate, some notion of what is the relevant information to be transmitted is required. One can define an enormous amount of information in the spatial variation of the depolarization on a single cell that is unlikely to be relevant to cognitive science.

But wouldn't the identification of the relevant information that can be transmitted with profit for cognitive phenomena be itself a major success in a bottom-up approach of cognition? Do we know what it is?

Finally, concerning the magnitude of synchronization: in Figure 1 we reproduce a cross-correlation curve, expressing stimulus selective synchronization from Kreiter and Singer (1996). To its right is a cross-correlation from a simulation of spontaneous activity of a totally structureless, noisy network of spiking integrate and fire neurons (Amit & Brunel 1996). One notes first that the two curves are qualitatively identical. The peak at zero delay is about the same factor above the background value corresponding to no correlation between the spike trains. The remarkable fact about the experimental result is that it depends on the stimuli presented to the two nonoverlapping receptive fields. This is of course absent in the simulation. What is significant about the simulation is that despite the presence of this level of "synchronization," all of its properties other than the cross-correlation itself can be theoretically deduced quantitatively, assuming that the spike trains are independent. This would imply that correlations at this level and of such duration are not noticed up-stream. [See also Amit: "The Hebbian Paradigm Reintegrated" *BBS* 18(4) 1995.]

ACKNOWLEDGMENT

This work is supported in part by a grant of the Basic Research Fund of the Israeli Academy of Science and a Human Capital and Mobility grant of the EEC.

Do the biological details matter?

James M. Bower

Computation and Neural Systems Program, California Institute of Technology, Pasadena, CA 91125;

jbower@smaug.bbb.caltech.edu www.bbb.caltech.edu/bowerlab

Abstract: Phillips & Singer (P&S) extend ideas derived from the observation eight years ago that the coherence (synchronization) of cortical oscillations can be modulated by the structure of visual stimuli. As described in the target article, a large part of the continued interest in this finding is related to independent theoretical work suggesting that synchronized cell firing could help solve the problem of binding together within cortex neuronal activity associated with different attributes of visual stimuli. The authors present an abstract "proof of concept" model describing how their cortical processing scheme could work, but our biologically realistic models of cortical relationships suggest that the proposal is biologically implausible. Our realistic models lead to a very different interpretation of the significance of cortical oscillations.

Abstract models as existence proofs. Our own more realistic models of cerebral cortical networks and neurons suggest that Phillips & Singer's (P&S's) ideas are biologically implausible. Shortly after the first reports of oscillations in visual cortex (Gray 1989), we replicated the results by slightly modifying our existing models of oscillations in olfactory cortex (Wilson & Bower 1988; 1992) to reflect neocortical structure (Wilson & Bower 1991). These simulations demonstrated the likely importance of horizontal connections in sustaining the physiologically described phase relationships and suggested that the mechanisms of oscillations in these two very different sensory cortices could reflect "a very general property of cerebral cortical function" (Wilson & Bower 1991). At the same time, however, we pointed out that there were likely to be several fundamental problems with implementing phase-dependent neuronal grouping in a real biological system, a result supported by our subsequent single cell modeling (see Bower 1995 for review).

Are cortical oscillations phase constant? For phase to group neurons, the phase relationships between the neurons should, in principle, be stable during the entire grouping period. Our model predicted that "zero phase relationships between particular neurons should exist . . . only over multiple trials," and stated that "the absence of within-trial coherence . . . (would) seriously confound the interpretation of fine phase differences in higher cortical areas." We concluded that "phase relationships reflect the structure of intercellular communication within (cortical) networks" and not neural coding, per se. In other words, phase relationships were not sufficiently instantaneous, continuous, or regular to support the binding hypothesis.

Are temporally coincident presynaptic effects more effective postsynaptically? At the base of the proposal concerning synchronous oscillations is the idea that post synaptic cells are more likely to fire with synchronous inputs. However, recent modeling and physiological investigations suggest that this basic assumption is not necessarily true. For example, the timing of network generated inhibition is equally or more important than excitatory presynaptic coincidence. Further, with proper timing, non-synchronous synaptic input patterns can result in larger somatic EPSPs than synchronized inputs (c.f. Protopapas & Bower 1994).

Is it likely that the timing of spike occurrence is independent of the information content in the spike sequence? In the abstract model presented, the receptive field (RF) and contextual field (CF) signals are simply declared to be independent through the built in transfer function: $A(r,c) = 0.5r[1 + \exp(2rc)]$. However, in real neurons it is likely that the precise timing of the spikes itself carries information about the stimulus (Bialik et al. 1991). Therefore, anything that changes the timing of spikes will alter the signal being transmitted to the next stage.

What then is the significance of cortical oscillations? As a result of our biological based simulations, we proposed many years ago that "phase may not be a robust strategy for encoding stimulus specific information" and instead suggested that cortical oscilla-

tions reflected the structure of inter-cellular communication within and between networks (for review, see Bower 1995). Our more recent modeling of single pyramidal cells suggest that these cells are very sensitive to the specific spatial and temporal (and not just synchronous) patterns of synaptic input. Furthermore, we have suggested for many years that oscillatory activity patterns also indirectly reflect the timing and organization of computational cycles within cortical circuits (Bower 1995). Finally, whatever the functional role of oscillations in cortical networks, I believe it is more useful to develop theories based on causal relationships between neural components manifest in realistic simulations than to attempt to squeeze ad hoc suggestions from machine vision into brain circuits.

Binding by synchronisation: A task-dependence hypothesis

Guido Bugmann

School of Computing, University of Plymouth, Plymouth PL4 8AA, United Kingdom; gbugmann@soc.plym.ac.uk
www.tech.plym.ac.uk/soc/staff/guidbugm/bugmann.htm

Abstract: Binding needs to be task dependent, and cannot usefully be driven by properties of the stimulus alone. However, task dependent binding can only take place after the patterns in a stimulus have been identified. Thus pattern recognition needs to be done prior to binding. Synchronisation may be a consequence of pattern recognition and can be used to localise the pattern and tag its attributes at different levels of information processing.

It is frequently proposed that the function of synchronisation is to bind features belonging to the same stimulus. Let us assume that this is the case. The working hypotheses in early papers on synchronisation were that binding is driven by properties of the stimulus and is mediated by lateral connections within a cortical area (or between areas). These hypotheses suffer from several weaknesses:

(1) The data most frequently cited as evidence that lateral connections cause synchronisation do not provide such evidence. In strabismic cats there are no lateral connections between neurons in V1 responding to inputs from different eyes. The reduced synchronisation between these neurons is interpreted as indicating a synchronising role of the lateral connections (König et al. 1993). However, strabismic cats only analyse the visual information of one eye at a time. So, the data can also be interpreted as indicating that (lack of) synchronisation results from the (lack of) processing of the information by higher visual layers. Another set of data shows that severing the callosal link between the visual areas in the two hemispheres eliminates the synchronisation between neurons in these areas (Engel et al. 1991b). This is also taken as evidence for the role of lateral connections for synchronisation. However, callosal fibres also comprise the feedback fibres from higher areas in the contralateral hemisphere. Thus the data could also indicate that synchronisation is mediated by *feedback projections* from the area that uses the information. This is consistent with findings by Sillito et al. (1994) and Munk et al. (1995).

(2) The data used to indicate that synchronisation is necessary for perception do not give such evidence. For instance, in cat and humans with strabismic amblyopia, only one eye is normally used. If forced to use the other eye, subjects reveal deficits in connecting contours and discriminating details near contours (see references in Roelfsma et al. 1994a). Only weak synchronisation is observed between neurons responding to this other eye (Roelfsma et al. 1994a). This is interpreted as indicating that synchronisation is needed for correct perception, but it could just as well indicate the opposite, namely, that correct *perception causes synchronisation*.

(3) The hypothesis of stimulus-driven binding does not explain how neurons know what they should bind. For instance when you

are observing your grandmother's face, you may wish to determine her identity, if she has her eyes open, or if her skin is smooth. In the three cases, different elements of the visual input need to be grouped. Thus, if binding by synchronisation takes place, it cannot be stimulus-driven. External inputs are needed to control binding in a *task-dependent* way.

The context field (CF) connections postulated in the paper by Phillips & Singer (P&S) open the door to a task dependent binding. However, they are not used for that purpose. Instead, they help in setting up a connectivity that maximises the transfer of relevant information about the input, where the "relevance" is determined by properties of the input. Thus binding remains input-driven rather than task-driven.

Unfortunately, task-driven binding is problematic too. It raises a new question as a counterpart to (3) above: How can a high-level process know which parts of the input image to group before it knows what is in the image?

One way to solve this problem is to abandon the notion that a binding mechanism is necessary for pattern recognition. The feature detectors at various stages of visual information processing may exhibit enough positional dependence to make further binding unnecessary. Also, if a visual scene contains numerous objects or interpretations of objects, why should these not all be recognised by their respective groups of neurons? The problem may be more to *localise* an object or pattern of interest and that is where synchronisation may be useful. In a neural network described by Bugmann and Taylor (1993; 1994a; 1994b), when a neuron in a higher layer starts firing, it induces the synchronisation of all neurons in the lower level layers contributing to its response by virtue of a cascade of feedback projections. The resulting area of synchronised activity in a low layer (e.g., corresponding to V1 or LGN) indicates the location of the pattern. By enabling or disabling high level neurons in a task dependent way, only the relevant visual features are synchronised. In such a scheme,

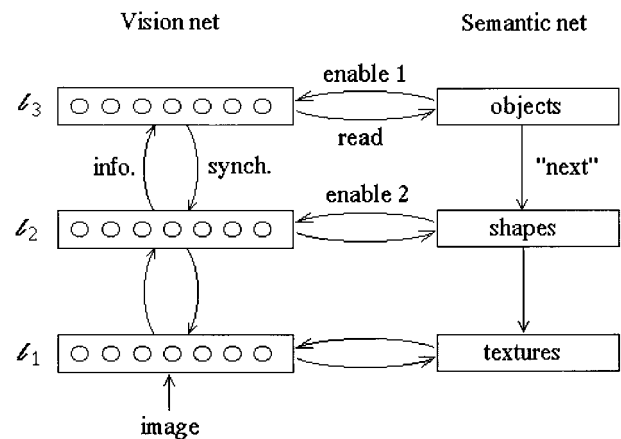


Figure 1 (Bugmann). Principle of multi-resolution binding. A question such as "How smooth is the skin of my grandmother's face?" is answered after the following sequence of events: First, the grandmother cell in layer l_3 is enabled by the semantic "objects" layer. When the grandmother is recognised, neurons in layer l_2 and l_1 representing attributes of the face are synchronised by feedback projections. Finally the semantic "texture" layer can use the group of synchronised neurons in layer l_1 to evaluate the smoothness of the skin. If the texture of the eyelids is of interest, one more processing step is needed: After recognition in layer l_3 is acknowledged by the semantic "object" layer, signals from the semantic "shape" layer restricts the set of active neurons in layer l_2 to those coding for the eyes. This in turn limits the population of synchronised neurons in l_1 to those corresponding to the eyes. Anatomical data in (Pandya & Yeterian 1985) may support such a model.

synchronisation and binding is a consequence of the recognition of a pattern, not its cause.

Task-dependent synchronisation predicts synchronisation between distant neurons which represent an object relevant to the task at hand (Bugmann & Taylor 1994b). This differs from the familiarity factor suggested by P&S.

Another possible role of top-down synchronisation is multi-resolution binding. For instance, in the grandmother example given above, the texture of the skin may be represented at a relatively low level, the shape of her eyes at a higher level, and her identity at an even higher level. Thus different tasks require information to be extracted from different levels. By allowing a high level neuron representing the concept "my grandmother" to fire, the subset of low level texture neurons corresponding to her skin would be synchronised. This sub-set could then be used to produce the response to the question about the smoothness of the skin (Fig. 1). Thus, a function of synchronisation may be the tagging of attributes of objects at different levels of resolution. Sensory information processing can then be viewed as a *presentation* process in which information is organised for other systems that need it.

Support for grouping-by-synchronization, the context-field, and its mechanisms, but doubt in the use of information theory by the cortex

Reinhard Eckhorn

Department of Neurophysics, Philipps-University, Renthof 7, D-35032 Marburg, Germany; eckhorn@bio.physik.uni-marburg.de

Abstract: Our work supports synchronization for binding within Phillips & Singer's "contextual field" (CF) as well as the type of its lateral interaction they propose. Both firmly agree with our "association field" (AF) and its modulatory influences (Eckhorn et al. 1990). However, the CF connections seem to produce anticorrelation among assemblies representing unrelated structures, whereas experimental evidence indicates decoupling. Finally, it is unclear how the cortex can have access to the logistic function used in the "coherent infomax" approach.

Evidence for stimulus-related and perception-related synchronization in the cortex. We can provide additional support for Phillips & Singer's (P&S's) hypothesis of context-related feature grouping from observations of synchronization among fast oscillations (35–90 Hz) in the visual cortex of anesthetized cats and awake monkeys; among neurons within a single vertical column; between separate columns; and between assemblies in different cortical areas (e.g., Eckhorn et al. 1988; 1993). A general finding was *zero correlation delay* even between signals from serially arranged cortical areas like V1 and V2 of the monkey (Frien et al. 1994). This is indicative of simultaneous processing in coupled assemblies despite transmission delays – which requires special types of symmetric mutual connectivity (e.g., Eckhorn 1997). Last, and most important, we have found first evidence of *perception related synchronization* in an intact behaving monkey during a binocular rivalry task. Here, the monkey's perception was correlated with the synchronization of oscillations between those neurons representing the features of the currently perceived stimulus (Eckhorn et al. 1996; Kottmann et al. 1996).

The cortex may use externally or internally generated synchronization for context-related processing. The "coherent infomax" learning procedure requires covariant changes of signals in neighboring channels in order to assess "Gestalt principles" in sensory systems. During the first learning phase, these covariant signals can be supplied only by stimulus-locked cortical activities. Fast processing of feature grouping is required at later stages of development as well as in the absence of stimulus-locked activations. The context connections hence should induce specific

temporal correlation structure on their own. We found both types of specific synchronization, often simultaneously (e.g., Kruse & Eckhorn 1996). Fast-changing stimuli dominated cortical synchronization by direct stimulus-locked components and generally suppressed synchronized oscillations. However, fast oscillations were generated during phases of slowly changing retinal images (like ocular fixation and smooth pursuit), that is, in difficult stages of signal processing. Hence, both types of synchronization may play a role in contextual grouping operations.

Anti-correlation among separately synchronized groups was not found in the cortex. The P&S target article concentrates on context-related grouping arising from facilitative lateral coupling. However, separation is equally important when assemblies representing similar features do not belong together. The "contextual" connections take this into account by inhibiting uncorrelated signals in neighboring channels; this probably produces negative temporal correlations (troughs in cross-correlogram). However, to our knowledge, this effect has not been observed in the cortex. Instead, incoherent stimuli generally induced uncorrelated activities among their cortical representations (flat cross-correlogram). We have recently demonstrated such stimulus-related decoupling in a particularly sensitive test of figure-ground segregation in the monkey's visual cortex (Guettler et al. 1997).

Shape of orientation detectors' association-field (AF) in primary visual cortex. In the past we used a definition slightly different from the CF of P&S, and called it association- or linking-field (AF; e.g., Eckhorn et al. 1990). While CFs and AFs in models only supported coupling among coaxially aligned orientation detectors, our recordings of synchronized oscillations from cats' and monkeys' visual cortex revealed additional Gestalt principles. Besides coaxially aligned synchronization ("good continuation," example noted by P&S) coupling among orientation detectors with different RF positions and orientations exhibited "proximity" and "similarity" rules (e.g., Frien et al. 1996).

Two types of connections are necessary for receptive fields (RF) and association fields (AF). We can strengthen P&S's emphasis put on the difference in connections for establishing RF and CF actions by showing that feed-forward connections establish RF properties and are of a "standard" type while AF inputs modulate RF inputs (e.g., "offset multiplicative") and hence do not degrade the RF properties while sensitive synchronization is possible (Eckhorn 1997; Eckhorn et al. 1990). In addition, our results from models confirmed quantitatively that such modulatory connections are more powerful than additive coupling in enhancing or reducing spatial "signal correlation contrast" (Juergens & Eckhorn 1979).

Can the cortex use information-theoretic measures for learning and function? Our answer is "no" at the present state. We will not disagree that "information" reduction is a main goal of the sensory cortex, but no one knows how "relevant information" is defined and how it may be selectively extracted while irrelevant information is suppressed. The use of Shannon information for "mutual infomax" cannot solve this problem in principle because it is restricted to structural information (in space and time) and does not include the concept of relevance or meaning. In addition, it does not provide rules for coding but allows only estimations of upper bounds of what might be achieved maximally with a given code (e.g., Eckhorn et al. 1975b). Moreover, in the P&S review it is not evident how the logistic function defined by information measures is made available to the cortex (during development and function). In addition, information theory requires stationary conditions to determine the probabilities of symbols carrying the information, whereas property of cortical function is changes in its state. Such instationarities are present either on long time scales due to learning or on fast scales during sensory processing.

Finally, the P&S do not discuss explicitly the often disputed question of whether and how the cortex extracts and utilizes sensory information that is redundant (as it seems to be in the "mutual infomax") or whether it suppresses it to obtain sparse independent code symbols (Barlow & Földiak 1989). The results

of the “mutual infomax” approach seem interesting, however, and future work may show whether something similar is really used by the cortex.

ACKNOWLEDGMENT

The authors thank A. Frien, E. Juergens, K. Kopecz, and T. Schanze for their comments.

An internal teacher for neural computation

Dario Floreano

Department of Computer Science, Mantra Center for Neural Computation, Swiss Federal Institute of Technology, LAMI-INF-EPFL CH-1015 Lausanne, Switzerland; floreano@di.epfl.ch diwww.epfl.ch/lami/team/floreano

Abstract: Contextual signals might supervise the discovery of coherently varying information between cortical modules computing different functions of their receptive field input. This hypothesis is explored in two sets of computational experiments, one studying the effects on learning of long-range unidirectional contextual signals mediated by intervening processors, and the other showing contextually supervised discovery of a high-order variable in a multilayer network.

1. Supervision and biological plausibility. Supervised models represent a successful paradigm of learning for artificial neural networks. Not only have such models been applied to a wide range of engineering problems, but they have also been used to infer computational properties of living brains (e.g., Zipser & Andersen 1988). However, the requirement of corrective feedback for output neurons (provided by an “external teacher”), often coupled with backpropagation of error signals along the same synaptic connections (Rumelhart et al. 1986), raises serious doubts about the biological plausibility of neural models of supervised learning (Crick 1989).

The theoretical framework and the computational model described by Phillips & Singer (P&S) might now provide a conciliatory solution which combines intermodule supervision with a biologically plausible learning scheme. Contextual signals between different processors computing different functions of their inputs might play the role of an “intrinsic teacher.” The idea is that spatial or temporal correlation of events could be exploited by the neural system to guide acquisition of a novel processing ability.

To test this hypothesis, at least three constraints must be satisfied: (1) If supervision takes place between processors computing different functions (which might thus be located far apart in the cortical tissue), it is necessary to check whether contextual guidance can have effects at long distance (longer than that implied by the average extension of lateral intra-cortical connections), for example, through the mediation of other intervening processors. (2) Since supervision makes sense only when at least one of the interconnected cortical processors cannot discover the appropriate function, the “intrinsic teacher” hypothesis should be tested on a computationally “difficult” function such as a nonlinear transformation of the input. (3) Although one could imagine a situation where an already developed neural module supervises the learning of another module, the model should also account for the situation when both modules learn together.

2. Long-range effects of contextual signals. The hypothesis that contextual signals might function as an “intrinsic teacher” was tested in a set of computational experiments guided by the above constraints and based on the mathematical model outlined in P&S’s target article. An important feature of these experiments is that contextual connections between different processors (or modules, if they are composed of more than one processor) were unidirectional, whereas in all the other experiments described in the target article contextual connections were reciprocal and symmetrical. The first experiment was aimed at testing the first constraint, namely, whether a contextual signal traveling through several processors with nonoverlapping receptive fields could guide discovery of the same feature in all processors.




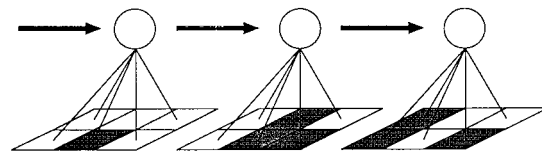


Figure 1 (Floreano). (a) Top: contextually supervised edge-contrast detection (black = -1 ; white = $+1$). Thick arrows are contextual connections. Bottom: Coherent Information transmitted by each processor during learning (solid line: processor 1; dashed line: processor 2; dotted line: processor 3) and final CF and RF structure (black is inhibitory, white is excitatory, size is proportional to strength). (b) Top: contextually supervised stereo-depth. Each output unit of the edge-contrast module sends a contextual signal to each unit of the stereo-depth network (thick arrows). Input to stereo network consists of a right (top line) and left (bottom line) retinal image. Left image is shifted one pixel to the right. Bottom: Coherent Information transmitted by edge-contrast network (graph on the left) and by output unit of stereo-depth network (graph on the right).

The architecture was composed of three processors that were presented with a horizontal edge contrast whose sign was correlated across processors (see top of Fig. 1). The structure of the input for each processor can be easily understood if we visualize the receptive field of each module as a 2×3 matrix whose entries can take random bipolar values -1 or $+1$. The sign of the contrast edge was given by the difference between the two sums of the row components. For every pattern presentation, each processor received a new random input, but the sign of the edge contrast was correlated across processors. Although this arrangement might look a bit artificial, it guaranteed that a simple Infomax approach could not discover the information correlated across processors. The three processors were arranged on a line, each receiving a contextual signal through a unidirectional connection from the preceding processor; the contextual signal to the first processor, which was in this experiment supposed to come from an already developed neural module, was a bit whose value alternated together with the sign of the edge contrast presented as input to the network.

Starting from initially random weights for both receptive field (RF) and contextual fields (CF) connections, in less than 500 learning cycles all processors learned to signal the sign of the edge contrast even though the contextual signal to the last two processors was mediated by intervening processors (the graph at the top of the figure plots Coherent Infomax values for each processor

during learning). The final structure of weight strengths for each module reflected the edge variable (top of Fig. 1 represents weight strengths using the same format as Fig. 7 of P&S's target article). Similar results were also obtained with longer chains of processors, with different input features (lines against background), and also with partial overlap of receptive fields (25% of the total surface area) between adjacent processors.

3. Discovery of surface depth from multiple cues. Having shown that contextual signals could drive discovery of correlated functions even at long distances, a further set of experiments was run where a neural module attempted to drive another module to extract depth from stereo images. Depth perception from stereo images requires a nonlinear transformation which could be learned by a multi-layer perceptron, but not with self-organized Hebbian learning (unless some specific constraints are imposed on the network architecture). This seemed a suitable problem to test whether contextual signals arising from another neural module could guide the learning of a computationally difficult function. Also, stereo-depth had already been used as a test problem in a similar framework by Becker and Hinton (1992) (see sect. 1.4, para. 4, of the target article for a discussion of important differences).

The neural network architecture was composed of two modules, each one presented with input variables coherently varying across modules (see bottom of Fig. 1). One module, which was composed of two processors linked by reciprocal contextual connections, was presented with the same edge-contrast input described above. The other module, composed of a single multi-layer processor, was presented with two horizontal "slices of pixels," each one extracted from the images forming a random-dot stereogram. During learning, variation of the sign of the edge was correlated with a left or right shift in one of the random-dot images. The hidden unit and the output unit of the stereo-depth module received contextual signals from the output units of the edge-contrast network through unidirectional connections. Although these input patterns might be seen as multiple retinal cues correlated with depth (such as direction of contrast between foreground and background, and corresponding direction of binocular disparity), the choice of this particular setup was motivated mainly by coherence with previous experiments described in section 2 and by constraints (2) and (3) suggested in section 1.

Several experiments were run where both modules simultaneously learned with a new set of initial random strengths for RF and CF connections. On average, the stereo-depth module successfully learned in 75% of the experiments to signal correctly the direction of retinal shift (the edge-contrast module always learned successfully), as can be seen by the value of Coherent Infomax (graphs at bottom of Fig. 1).

Several types of random-dot stereograms were used in these experiments, varying the size of the image, using a Gaussian image-centered shift probability for each pixel, and adding various levels of uniform noise. Although these variations did not have an influence on whether the stereo-depth network learned successfully or not, the cases of failure took place when the edge-contrast network learned very quickly. As a control experiment, we observed that a stereo-depth module composed of two or more interconnected processors without contextual guidance from other modules could never discover the sign of the shift.

As expected, the strength of the contextual connections to the hidden unit of the stereo-depth network was smaller than that of the connections to the output unit, because the hidden unit activation cannot be perfectly correlated with the direction of shift (if that was the case, a simple perceptron would be sufficient). Nonetheless, from a set of other experiments, it was clear that both contextual connections to the hidden unit and to the output unit were necessary for the stereo-depth network to learn correctly.

4. Conclusion. The experiments described here do not rule out the possibility that the discovery of nonlinear high-order coherent variables could simply arise from bottom-up information processing and reciprocal lateral interactions. Rather, they are intended to

suggest a further powerful functionality that unidirectional contextual connections might serve: that of an intrinsic teacher. Not only does this possibility help bridge the gap between a powerful learning paradigm and biological plausibility, but it makes Phillips & Singer's theory even more appealing as a common foundation of cortical computation.

ACKNOWLEDGMENTS

Part of this work was supported by a Network grant from the Human Capital and Mobility Program of the European Community. Special thanks to Darragh Smyth for several interesting discussions.

Nonlinear computation and dynamic cognitive generalities

Robert A. M. Gregson

Division of Psychology, School of Life Sciences, Australian National University, Canberra, ACT 0200, Australia; robert.gregson@anu.edu.au

Abstract: Although one can endorse the complexity of the data and processes that Phillips & Singer (P&S) review, their mathematical suggestions can be compared critically with cases in nonlinear psychophysics, where the theoretician is faced with analogous problems. Owing to P&S's failure adequately to recognise both the intricate properties of nonlinear dynamics in networks and the constraints of metabolic demands on the temporal generation of patterns in biological nets their conclusions fail to meet the problems they properly address.

Phillips & Singer (P&S) have generously provided us with a challenge that is subtle, complex, and erudite. It would be pretentious to attempt a global critique; I restrict myself to the smaller area of multivariate psychophysics where there are parallels and agreements, and perhaps profound differences. Their distinction (sect. 1) between subsystems differing in information upon which they operate, but not in some of the fundamental operations they have available as processors, is one I completely endorse. Indeed, it is fundamental to the application of nonlinear dynamics in connected channels and networks to the creation of a theoretical psychophysics independent of any one specific sensory modality (Gregson 1988; 1992; 1995; 1996).

Do we need the sorts of mathematics at which P&S hint, and which they describe here obliquely in words and pictures, to exhibit the intricate dynamics to which the data surveys point? We have repeated mention of nonlinearities, levels of operation, local coupling, multivariate statistics, synchronization of vectors and spikes, recoding and condensation of information, and dynamic grouping subject to central plastic revision. I concur with all of these, and more, but I would add, having tried to write theory without it, some consideration of yet another contextual field (what they call contextual field [CF] on receptive field [RF] processes) and that is metabolism. The intrinsic delays and pacings required by a system that can, because of its energy consumption, run temporarily out of steam, impose boundary conditions on modeling. The brain is in real time and is a dissipative system both biologically and mathematically.

The hippocampus has a central role in information processing; its operations include recursive loops, with different delays and cortical mappings. For the dynamic theorist it is the properties of such loops in an environment where the system floats in its phase space between stabilities of various types, including chaos, that can carry the encoding and the temporal pacing of patterns of information.

Using mathematics quite different from the transfer function (sect. 3.1) and constructing bivariate lattices and cascades of such lattices, I obtained simulations of many of the sensory and perceptual properties that P&S want to derive (sect. 3.4.2) from their units, processors, streams, and layers. What is interesting is that the coherence or synchronicity effects between local regions can emerge (Gregson 1995) without any CF to induce it, but in lattices of cross-coupled nonlinear complex trajectories.

If we use complex nonlinear dynamics, however, we do not need some of the induction effects of CF on RF which P&S bring in. Hence we can question the metatheoretical status of the CF/RF distinction. I agree that to have the system learn as well as to sense and to perceive, some other upstream layers are required. One way is to use series which only transmit at local dynamic windows (Gregson 1993). But, interestingly, if we nest nonlinear dynamics in slower paced loops which themselves encode boundary conditions restricting the rate of local processing (which, not too fancifully, are metabolic constraints), then trajectories with a diversity of periodicities and (apparently) superimposed local perturbations will arise. Some of these are strictly periodic with a long inter-signal interval on a constant base and can thus serve as a biological clock; others have spikes which are aperiodic on a chaotic base and are reminiscent of epileptic EEG records. One can generate this diversity by using a collection of attractors (Gregson & Harvey 1992), but one can equally well do it by nesting one structurally rich attractor within one or two hierarchically nested recursive loops. The brain is not short of such structures.

One must be very reluctant to postulate special mechanisms to induce information-encoding patterns when the biology of the situation could induce their appearance naturally with quite tiny parameter changes. The basic nonlinear dynamics, pushed recursively back into stable regions of its phase space, will do many things spontaneously. It needs compatibility with the underlying neurophysiology, but not creating seven layers! (sect. 5.7.1).

The answer to P&S's question (sect. 5.7.2, para. 5) "do the mechanisms producing the effects of context on the perception . . . of stimulus elements also play a role in dynamically grouping those elements?" probably is that the question is back to front. If we start with nonlinear dynamics evolving in time, then the grouping can come before the perception and perhaps before the emergence of consciousness (Ruhnau 1995).

The comment (sect. 5.7.2, Number 3 of major unresolved issues) that "if grouping is computed dynamically . . . signaled by relative timing, then such groupings could change rapidly from moment to moment" is indeed insightful. The serious question is, what sort of dynamical computing: there is an infinity of possibilities of which I have merely scratched the surface in using Γ cubic complex polynomial recursions. Some physicists make use of \tanh mappings for $-\infty, +\infty \rightarrow -1, +1$ for similar purposes (Amit 1989). Their $A(r,c)$ (sect. 3.1) is not adequate, and if it leads to Figure 3 then I am sure it is wrong.

There are no adequate dynamic formalizations in P&S's paper but tantalising hints as to which alternatives might be compared. There is some sort of consensus as to what fundamentals are needed, and what phenomena need to be accommodated. Phillips & Singer are to be commended on listing so many hurdles we have to leap, but I don't accept that their running shoes have the right spikes.

Principles of cortical synchronization

Stephen Grossberg

Department of Cognitive and Neural Systems, Boston University,
Boston, MA 02215; steve@cns.bu.edu
cns-web@bu.edu/Profiles/Grossberg.html

Abstract: Functional roles for cortical synchronization in self-organizing neural networks are described. These properties are best understood by models that link brain to behavior. Synchrony can express itself differently in cortical circuits that perform different behavioral tasks. Cortical temporal properties that seem inexplicable by synchrony are also mentioned.

Phillips & Singer (P&S) nicely summarize key cortical properties without disclosing the cortical design principles from which they arise. I will discuss some missing principles.

Let me start with background remarks: "Behavioral success drives brain evolution." Properly designed neurons must properly

interact in networks whose emergent properties control behaviors that autonomously adapt to changing environments. Thus, one needs to simultaneously link three levels: neuron, network, and behavior, or else the functional significance of brain data cannot be tested. P&S do not establish this linkage.

One needs models to do so. P&S emphasize models that explain data on one level. Models that bridge all three levels obey nonlinear feedback laws. One thus needs to keep neurons, networks, behavior, and nonlinear feedback models simultaneously in mind. Most neuroscientists have not been trained to bridge these levels.

"Neo-classical" theories abound instead. For example, information theory works in stationary worlds with unchanging states, but brains adapt to nonstationary worlds and learn new states. Connectionist learning models are mentioned that cannot self-organize, but brains are self-organizing organs par excellence. None of these models links realistic brain mechanisms to parametric behavioral data to support its biological relevance. Models that do bridge all four levels are hardly mentioned, even though many data that P&S review were anticipated by them. Nor have the authors surveyed the model ideas that generated these predictions. I will briefly review some of them now.

Why does synchrony occur on both circuit and system levels? Concerning circuits: I showed in 1968 that nerve cells avoid losing sensitivity to fluctuating input patterns (viz., contextual information) if they obey membrane equations and interact via on-center off-surround networks. This result demonstrated that "the unit of short-term memory is a spatial pattern," or that these networks synchronize their activities to process distributed analog patterns of contextual information (see Grossberg 1982 for discussion and references).

In neural networks derived from behavioral postulates about associative learning, cell activities, or short-term memory traces, interact with adaptive weights or long-term memory traces. These learning laws are like those of Artola and Singer (1993). I proved, under general conditions of cell signal processing, feedback, and modulation of learning that "the unit of long-term memory is a spatial pattern"; that is, associative networks learn contextually informative synchronized patterns of activation (Grossberg 1968; 1971; 1982).

These results proved mathematically that synchrony is an emergent property of ubiquitously occurring on-center off-surround and associative networks, and that synchrony could occur with or without oscillations. Much of the field of neural networks has grown from these results, including self-organizing feature maps and the so-called Hopfield networks.

Concerning systems: Grossberg and Somers (1991) showed that two types of networks could synchronize quickly: (1) reciprocally interacting bottom-up and top-down adaptive filters for perceptual tuning and categorization, and (2) horizontal interactions for perceptual grouping. However, these synchronized states have qualitatively different functional properties. A top-down adaptive filter cannot activate a target cell unless it also gets a bottom-up input. This is the well-known ART (adaptive resonance theory) matching law (for reviews, see Carpenter & Grossberg 1991; Grossberg 1995), which suggests how top-down feedback attentionally primes and selects consistent bottom-up signals, as in phonemic restoration and the Sillito et al. (1994) corticothalamic data. Horizontal grouping interactions can activate target cells that do not get bottom-up input, as in illusory contour formation. The association fields of Field et al. (1993) satisfy the constraints on horizontal interaction that Grossberg and Mingolla (1985) predicted to explain perceptual grouping data.

P&S overlooked this key difference between adaptive filters and horizontal interactions, even though it illustrates different functional roles: bottom-up and top-down filtering and matching mechanisms enable the cortex to learn stably throughout life. Horizontal interactions for perceptual grouping complete missing information, as over blind spots and retinal veins. Grouping laws would destabilize learning if they were used for top-down matching.

A recent model combines these two types of processes. Grossberg et al. (1997b) simulated how the laminar, columnar, map, and network properties of interblob and interstripe cortical areas V1 and V2 interact reciprocally with LGN to explain perceptual grouping data. This model describes functionally different and experimentally testable roles for adaptive filters and horizontal interactions.

P&S compare horizontal visual interactions with word recognition. I think they make the same mistake here. A large modeling literature uses ART to explain speech and word recognition data as well as visual filtering and categorization data (e.g., Carpenter & Grossberg 1993; Gove et al. 1995; Grossberg 1987; 1995). Horizontal visual interactions (e.g., Kapadia et al. 1995) have a different explanation (Grossberg et al. 1997b).

Various brain timing properties have been modeled using these distinct mechanisms: groupings persist for variable durations in a context-sensitive manner (Francis et al. 1994). Top-down matching during word perception can reorganize which word strings are temporally fused or segregated (Grossberg et al. 1997a). Phillips & Singer's synchrony analysis cannot handle such data.

Synchronizing oscillations: Coding by concurrence and by sequence

V. G. Haase and L. F. M. Diniz

Departamento de Psicologia, Universidade Federal de Minas Gerais, Belo Horizonte, Brazil 31270901; haase@fafich.ufmg.br

Abstract: Synchronizing oscillations may be just one case of integration and/or coding, one which explains associations by concurrence. Understanding the sequencing of neural/behavioral events requires a clock mechanism that imposes structure behind mere associations, and may be best served by dissociating oscillations and synchronization in terms of physiologic and computational mechanisms.

Discussing the linguistic typology of aphasias, Roman Jakobson quoted Karl Pribram as having said that "the two linguistic axes find their correspondence in the brain: viz. decoding/encoding is posterior/frontal in the brain; concurrence/sequence (or simultaneity/successiveness) is arranged as dorsolateral/mediobasal in the brain" (Jakobson 1964, p. 39). It is worth quoting these time honored speculations to make us aware of the risks of oversimplifying. It also helps us remember that, in addition to concurrence, sequencing represents an important "association" mechanism. Phillips & Singer (P&S) make an impressive argument regarding the cellular mechanisms of associations along the concurrence or simultaneity axis. Although this may be an important "common computation" in the neocortex, it is not the only one. To be sure, P&S do not claim it is. However, a full account of the sequencing of neural and behavioral events may not be possible in terms of associations alone. Some structure is also required (Lashley 1951).

Structure may be provided by a temporal quantizing mechanism in the brain. According to one research tradition, synchronizing neural oscillations in the 30–40 Hz range is a possible integrative mechanism (Harter 1967; Pöppel 1970; 1978; 1994). Time is processed discontinuously in the brain. Each time packet is defined by the period of oscillatory activity. Inside a given time quantum, temporal relations are not consciously discernible and everything that happens during such a moment of time is integrated into a neural/psychological event. Mental activity consists of sequentially associating such events up to an upper limit of around 3 seconds. In other words, oscillations may work as a clock mechanism, a *Zeitgeber*, and their period may define an integration time window. Dynamic properties are assured so far as this mechanism consists of an instantaneously resettable clock with the properties of a relaxation oscillator (Soutif 1970; Wever 1965).

To integrate neurophysiological observations on gamma oscillations with the perceptual moment hypothesis, a differentiation

between the physiological mechanisms and computational functions of both oscillations and synchronization is in order. Recent theorizing about the role of synchronizing gamma oscillations in neuronal assembly formation seems somehow ambiguous (Singer 1990; 1993; 1994b). In earlier papers by Singer and coworkers, oscillatory mechanisms played a major role as a synchronizing mechanism. This emphasis diminished gradually as evidence accumulated that oscillatory activity is not necessary for synchronization (Tovee & Rolls 1992; Young et al. 1992) and that in some cases there is no correlation between oscillatory activity and specific parameters of sensation or movement (Ghose & Freeman 1992; Murthy & Fetz 1992; Sanes & Donoghue 1993). Matters could be cleared if associations and dissociations between gamma oscillations and distributed encoding through coherent neural assemblies and their respective functional and computational roles were more thoroughly explored.

Theunissen and Miller (1995) called attention to the theoretical importance of differentiating between an encoding and an integration time window. Experimental work on the locust olfactory system also suggests a mechanism through which temporal quantizing of neuronal events may contribute to sequencing mechanisms (see review in Laurent 1996). On a cycle-by-cycle basis, different neuronal ensembles are defined by their coherent activity, and information is also contained in the precise temporal sequence by which these assemblies are updated along successive periods. This last finding opens the possibility of integrating in one coherent framework not only oscillations and synchronization, but also the temporal patterning of discharge trains which has recently been the subject of increasing interest (Bialek & Rieke 1992; Krüger & Becker 1991; McClurkin et al. 1991).

Word recognition in the split brain and PET studies of spatial stimulus-response compatibility support contextual integration

Marco Iacoboni

Brain Mapping Division, Department of Neurology and Psychology, University of California, Los Angeles, Los Angeles CA 90095; iacoboni@loni.ucla.edu brainmapping.loni.ucla.edu

Abstract: The neural substrates of context effects in word perception are still largely unclear. Interhemispheric priming phenomena in word recognition, typically observed in normal subjects, are absent in commissurotomy patients. This suggests that callosal fibers may provide contextual integration. In addition, certain characteristics of human frontal cortical fields subserving sensorimotor learning, as investigated by positron emission tomography, provide evidence for contextual integration not confined to the visual system. This supports the notion of common aspects of cortical computations in different cerebral areas.

Phillips & Singer (P&S) propose (1) synchronization as a neurophysiological mechanism of contextual integration, and (2) context effects in word perception as behavioral evidence supporting contextual guidance in cortical computations. Their proposal would be strengthened by empirical evidence that provides some possible links between these two phenomena.

That "the cortical algorithm is everywhere the same" (sect. 1) is obviously an appealing and relevant aspect of P&S's proposal. However, the neurophysiological evidence cited in support of contextual integration is largely confined to the visual system. If the cortical algorithm is everywhere the same, then there must be some biological evidence for contextual integration in other domains, motor, cognitive, attentional, and so on.

Interhemispheric lexicality priming in the normal and the split brain. In lateralized lexical decision between words and orthographically legal pseudowords, responses are faster and more accurate for letter strings in the right visual field than in the left visual field. In normal subjects, responses are more accurate and faster for two copies of a word simultaneously presented in both

visual fields, than for a single copy presented in the right visual field. This phenomenon is absent in the split brain (Mohr et al. 1994).

In normals, when two different letter strings (one cued as the target, the other uncued as the distractor) are simultaneously presented in the two visual fields (bilateral presentation), responses are less accurate and slower, compared to responses to a single letter string (unilateral presentation). When bilateral presentations are considered, congruent pairs (both target and distractor are words or pseudowords) are processed faster and more accurately than noncongruent pairs (word-nonword and nonword-word) (Iacoboni & Zaidel 1996). Once again, the split brain seems not to show these interhemispheric "context" effects (Iacoboni et al. 1996a; in preparation).

Taken together, these findings seem to support the role of callosal fibers in contextual integration and may provide a link between the observed role of callosal connections in synchronizing activity (Engel et al. 1991) and behavioral contextual effects in word recognition.

Premotor cortex, spatial stimulus-response compatibility, and sensorimotor learning. In a positron emission tomography (PET) experiment on sensorimotor integration and sensorimotor learning, we used a task that requires the coding of stimuli in extrapersonal space and of responses in personal space (spatial stimulus-response compatibility). In this experiment we manipulated the mapping of extrapersonal and personal space. This was done by asking subjects to respond in half of the PET scans with their hands in the uncrossed position, that is, each hand in its homonymous space, and to respond in the other half of the PET scans with hands in crossed response position, with each hand in its heteronymous hemispaces, that is, left hand in right hemispaces and right hand in left hemispaces (Iacoboni et al. 1997). During the imaging session comprising twelve PET scans, four subjects ("switchers") switched very frequently 10 times from the uncrossed to the crossed response position. The other four subjects ("non-switchers") switched less frequently. Frequent switching of uncrossed and crossed response position is known to affect the sensorimotor learning in this task (contextual interference) (Proctor & Dutta 1993).

As expected, switchers showed a smaller learning effect (as measured by shortening in reaction times) than nonswitchers. In spatial stimulus-response compatibility, learning is generally associated with blood flow increases in the caudal sector of dorsal premotor cortex (PMdc) (Iacoboni et al. 1996b). In this region, learning-related blood flow increases were bigger in nonswitchers than in switchers. Moreover, nonswitchers presented a significant activation of the caudal sector of the supplementary motor area (SMA-proper) (Picard & Strick 1996), compared to switchers. This suggests that SMA-proper provided contextual reinforcement to PMdc to facilitate learning acquisition in non-switchers. This contextual effect would have been subserved by the robust corticocortical connection between SMA-proper and PMdc (Matelli et al. 1991).

The rostral sector of dorsal premotor cortex (PMdr) is also strongly connected with PMdc (Matelli et al. 1991). PMdr is the premotor region that codes the explicit stimulus-response association rule in spatial stimulus-response compatibility (Iacoboni et al. 1997). Thus, PMdc fits the model of local processor with contextual guidance outlined by Phillips & Singer in section 1.2.1 of the target article. PMdc would receive receptive field (RF) inputs from PMdr and contextual field (CF) inputs from SMA-proper. PMdc would use contextual inputs from SMA-proper "to influence the confidence with which decisions are made on the basis of the RF evidence," as provided by PMdr. Thus, these PET data seem (1) to fit with the model, and (2) to provide biological evidence for contextual guidance that is not restricted to the visual system.

Internal context and top-down processing

Peter König, Carl Chiang, and Astrid von Stein

The Neurosciences Institute, San Diego, CA 92121; peterk@nsi.edu; cchiang@nsi.edu; astein@nsi.edu www.nsi.edu/users/konig

Abstract: Recent experimental work suggests that the concept of contextual fields should be generalized to allow the modulation of local information extraction by both external and internal context. The external context relates to the coherent information of the stimulus; the internal context refers to the parts of this information which are relevant for behavior. This dual interaction, present at every level of the hierarchy, requires a fundamental unit of processing more complex than a single neuron appears today. We argue that the cortical column supplies the required mechanisms for a separate top-down and bottom-up processing and allows for interaction without destructive cross-talk. This flexibility might explain why the same basic architecture of a cortical column is used in the whole isocortex.

In Phillips & Singer's target article, the "context field," which is used to extract coherent information, is defined primarily by the properties of the external stimulus. Examples of "context fields" are features of the visual stimulus outside the classical receptive field or surrounding letters/words in a lexical task. These paradigms place the emphasis on the passive processing of external stimuli. However, humans and animals alike cannot be described simply as stimulus-response machines. For a human or an animal, even more important than the detection of coherent information in the environment is the extraction of those parts which are relevant to its behavior. The reaction to stimuli depends very much on internal variables like motivation, attention, and intent. In our view, in order to serve as a general framework for the cortical dynamics, the concept of "context fields" must be generalized to include the internal state of the animal.

This generalization leads to the prediction that if the modulation of feature extraction by the "context field" is mediated by synchronized activity, as is the central claim of the target article, the behavioral relevance of the stimulus and the expectancy of the animal should influence the synchronization patterns in sensory areas. Indeed, recent experimental evidence supports this notion (Chiang et al. 1996; von Stein et al. 1996). To determine the role of synchronization in top-down guided binding we simultaneously recorded field potentials from different cortical layers in primary visual cortex (area 17) and sensorimotor association cortex (area 7) in behaving cats. They were presented with two different visual stimuli and trained to respond by either watching one (no-go stimulus) or by pressing a lever after seeing the other (go stimulus). We found that: (1) compared to the no-go situation, the go stimulus induced a much stronger synchronization between the two areas; (2) this increase in synchronization between area 17 and area 7 was correlated with the correct performance of the task; (3) contrastingly, a highly salient but irrelevant stimulus failed to induce synchronized activity. Thus, the coupling between these two areas was dependent on the behavioral context. The synchronization pattern not only reflects Gestalt laws – corresponding to bottom-up synchronization – but can be influenced decisively by the internal state – corresponding to top-down processing.

The interaction between external stimuli and internal variables as described in the previous paragraph places a considerable computational burden on the processing units. The generation and evaluation of synchronous events must be kept separate for multiple sources and targets at each level. From present knowledge, this seems a hard task for a single neuron to accomplish. It has been argued before (Edelman 1987), that groups of closely coupled neurons are the basic cortical processing units, which then can form larger assemblies (Hebb 1949). For a number of reasons we think that a cortical column implements such a group and performs the tasks described above:

1. Under physiological conditions, coincidence detection is the dominant processing mode of cortical neurons (König et al. 1996a) and can produce precisely timed output (Mainen & Sejnowski 1995).

2. Due to the termination of afferents from feedforward – relating to the external context – and reentrant connections – relating to the internal context – onto neurons in different layers (van Essen & Maunsell 1984), synchronous events originating from different levels of the processing hierarchy can be evaluated independently.

3. Along similar lines, the projection neurons with different targets are segregated in different layers. Thus binding can be specific for targets on lower or higher levels of the hierarchy, avoiding global synchronization.

4. *In vitro* experiments demonstrate the complex fast dynamics within a column upon electrical stimulation (Langdon & Sur 1990).

5. In the experiments described above (Chiang et al. 1996; von Stein et al. 1996), a layer-specific synchronization between areas at different levels of the visual processing hierarchy has been found. Furthermore, during presentation of expected stimuli, the interaction between the different areas occurred in particular in the alpha range, while during perception of novel stimuli, synchronization in the gamma range was predominant.

These experiments demonstrate that the cortical column is capable of operating on a millisecond time scale, and that the temporal dynamics depends on the match between expectation and stimulus. Thus, any description of cortical microcircuitry (Douglas & Martin 1991a) must be extended to include fast temporal processes.

In sum, the cortical column appears to be a rather flexible device, able to extract information which is dependent on internal as well as external context. It is this unique functional feature that may explain why the processing unit is used with so few changes in all areas of the neocortex, and may make it a true “common foundation for cortical computation,” which the target article sets out to find.

ACKNOWLEDGMENT

This work was supported by the Neurosciences Research Foundation.

Glossing over too much

Gin McCollum

R. S. Dow Neurological Sciences Institute, Portland, OR 97209;
mccollum@ohsu.edu
www.ohsu.edu/som-NeuroScience/sys/mccollum.html

Abstract: Although Phillips & Singer’s proposal of commonalities seems sound, information theory and artificial neural network modeling omit important detail. An example is given of a distributed neural transformation that has been characterized mathematically and found to have both overall commonalities and differences of detail in different regions. P&S’s contextual field is compared to inclusive regions in a formalism relevant for modeling bodily-kinaesthetic intelligence.

Phillips & Singer’s (P&S’s) idea that context changes the way we register incoming information has to be true. For example, sometimes we read for comprehension (skipping over typographical errors) and sometimes we proofread. Also, it is plausible that there are basic operations common to all cortical and cognitive processes. However, their methodology – information theory and artificial neural network modeling – is unequal to the task of finding those common, basic operations. Their methods may apply to micro-information, but not to macro-information: Who wrote that symphony? How do I ski down this wooded slope?

Precision in modeling. More formal structure is needed at intermediate levels to allow conclusions about basic operations, both cortical and cognitive. It is true that information theory can be used in both the cognitive and the neural domains, but information theory is too broad-brush to yield specific conclusions. There is mathematical structure in the ensemble-to-ensemble transformations of neural activity communicating to,

within, and from the cerebral cortex. Once that structure is characterized mathematically, the Shannon information can be derived, if it is needed, from the more accurate formalism.

For example, we have characterized mathematical structure in cerebellar climbing fiber tactile receptive fields (Castelfranco et al. 1993). Within the particular mathematical structure, one can determine what specification is conveyed out of what range of possibilities: the Shannon information. The mathematical structures of sets of receptive fields differ from one part of the nervous system to another. (See Schwartz, 1977, for a mathematical characterization of other receptive fields, including some in the cerebral cortex.) P&S propose that there are commonalities across the cerebral cortex with differences of detail. The example of cerebellar receptive fields has that feature also; it shows up clearly when the receptive fields are characterized mathematically (Castelfranco et al. 1993).

P&S propose synchrony as an important mechanism, with or without rhythmicity. Synchrony and rhythmicity occur also in the cerebellar cortex (Llinás & Sasaki 1989; Simpson et al. 1996). The mathematical characterization of receptive fields suggests a subtle relationship between synchrony and rhythmicity (McCollum 1997).

These remarks hold also for the output of a neural ensemble, for movements, and for cognitive processes. Linguistics already has a formal characterization which should be used. For example, rather than describing the /l/-versus-/r/ difference verbally (sect. 5.7.2), it can be formalized in a number of ways, including phonetically, geometrically, algebraically, or from a movement point of view. Mathematical characterizations at those levels will yield the cognitive information flow in various cognitive systems.

Managing degrees of freedom in learning. Even though some learning can be described as approaching some target performance, a more general description is elaborating a set of behavioral or imaging strategies (e.g., McCollum et al. 1995). [See also Nashner & McCollum: “The Organization of Human Postural Movements” *BBS* 8(1) 1985.] The cerebral cortex’s role of grasping connections and envisioning possibilities can then be expressed in terms of a more general form of P&S’s CFs and RFs: forming CFs from RFs and relating RFs to (already existing) CFs. We have formalized movement design problems in analogous but mathematically more specific terms (McCollum 1994). In an abstract space including all the degrees of freedom required for movement, a particular movement pattern (for example, speaking English) structures the space in a particular way. Another movement pattern (for example, speaking German) structures the space in a different way. Potential movements are grouped into regions of movement space that specify a particular movement. For example, pronouncing /r/ involves a certain region of a movement space, which is specified by the overall phonetic structure of the language, as expressed in movement. (Similar remarks could be made about the perceptual aspect.)

More precision is needed in specifying what is lost and what is gained in isolating movement patterns (P&S’s CFs), such as the overall phonetic structure of a language or the proper pronunciation of /r/. P&S refer to reduction in connection numbers in terms of “data compression” (sect. 3.2), “reduce redundancy” (sect. 1.2.1), and “discover the relevant” (sect. 3.4.1), as if no information will be lost. In contrast, Bernstein (1967) acknowledged the loss of information in reducing degrees of freedom when producing movement. When the variety of activity is reduced in structuring the flow of perception or movement, both redundancy and information are lost – a complex situation.

Even though P&S’s idea of common operations across the cerebral cortex and cognitive processes is probably correct, it would be better pursued with a bigger toolbox. We need to have more respect for the complexity of the processes involved.

Topologic organization of context fields for sensorimotor coordination

Pietro Morasso, Vittorio Sanguineti, and Francesco Frisone

Department of Informatics, Systems and Telecommunications, University of Genova, 16145 Genova, Italy; morasso@dist.unige.it
www.laboratorium.dist.unige.it/STAFF/motasso.html

Abstract: In field computing a topologic organization of CFs is necessary to support sensorimotor planning. A simple model of cortical dynamics can exploit such topologic organization.

We agree with Phillips & Singer (P&S) that “networks of local processors with contextual guidance can in effect discover distal variables and relationships by discovering mutual information in diverse data-sets” (sect. 6.5, para. 1). However, although this may be enough for understanding a good many sensory/perceptual systems, it is probably ill-suited for dealing with the guidance of action in sensorimotor coordination. What is needed there is a stronger notion than mutual information, which is essentially *passive*: something which captures the causality and circularity of purposive action.

Learning target-directed movements implicitly requires a Piagetian self-supervised paradigm, although the local learning rule may appear unsupervised (Morasso & Sanguineti 1996). The essential feature for enabling cortical maps to handle the Piagetian *reaction circulaire* is a capacity for topologic organization which allows an internal representation of external, distal space to emerge in a coordinate-free way (Sanger 1994). A good deal of work has been carried out on the ways in which simple, self-organizing heuristics can produce such internal representations, from the early work of Kohonen (1982) to the Martinetz & Schulten’s (1994) more recent Topology Representing Network and its application to modeling cortical map dynamics (Sanguineti et al. 1997). Although it might be argued that self-organizing maps are just a variant of the CF formalism, our opinion is that there is a subtle but critical gain in embedding CFs in a topological structure. Such extended CFs can do a lot more than optimally modulating the information transmitted by the RFs: they can perform field operations like diffusion, dispersion, convection, and so on, which allow the cortical areas to behave like a field computer (Amari 1983; MacLennan 1997).

To illustrate the concept, consider the following generic dynamic model of a cortical region

$$\frac{dV_i^x}{dt} + \gamma_i V_i^x = f(h_i^{lat} + h_i^{ext}) \quad (1)$$

where the i -th local processor (or cell assembly, or cortical column) is characterized by an activity level V_i and two kinds of inputs (h_i^{lat} , h_i^{ext}), which correspond to intra-cortical and external connections, respectively. The equation simply says that V_i evolves under the action of three competing influences: (1) self-inhibition (weighted by γ_i), (2) net input h_i^{lat} coming from the set of lateral connections inside the same cortical region, and (3) net input h_i^{ext} determined by thalamo-cortical or cortico-cortical connections across different cortical areas. The h_i^{lat} term is intended to express the massive lateral connections, which are symmetric and excitatory. This is a recurrent input which triggers the coupled dynamics of the whole network, according to global processes such as diffusion, relaxation, and so on. The h_i^{ext} term refers to the set of input neurons projecting to the given column (e.g., thalamic neurons or output neurons from other cortical areas). These connections are unidirectional, which means they are asymmetric and cannot support recurrent dynamics: as a consequence, what h_i^{ext} expresses is a smoothed version of the pattern of activity present in the projecting region.

As shown by Sanguineti et al. (1997), if sufficient gating inhibition is embodied by the nonlinearity on the right-hand side of the equation, the qualitative behavior of this kind of system can be

characterized as follows: (1) stable states of the system are identified by focused islands of activation (or population codes), in register with a local maximum in the input distribution; (2) if the input distribution shifts or jumps to a new pattern, a diffusion process is initiated which spreads around the initial cluster of activity, in a wave-like manner; (3) this activity automatically starts coalescing again as soon as it finds some kind of ridge or peak in the input distribution; (4) if such a distribution has many peaks, that is, several possible targets, the system automatically selects the *closest* one, with the notion of distance being dictated by the topologic organization of the lateral connections.

This is a different mechanism for assuring that the CF connections do not corrupt the nature of the RF of each processor; it incorporates a geometrically meaningful navigation capability in the cortical representation, which can explain invariant features of motion planning such as the near-straightness of reaching movements (Morasso 1981). Interesting and complex dynamic behaviors can be obtained if we couple two such maps, allowing the output of one to become the input of another, via a suitable set of cross-connections.

Testable predictions about the nature and shape of wave-like phenomena on the cortical area in sensorimotor planning can be derived from this scheme. It remains to be seen if this mechanism of navigation in cortical maps is also meaningful for grouping, dynamic binding, and other perceptual processes.

’Tis all in pieces (separate RFs and CFs), all coherence gone

Ernst Neibur^a and Marius Usher^b

^aZanvyl Krieger Mind/Brain Institute and Department of Neuroscience, The Johns Hopkins University, Baltimore, MD 21218;

niebur@jhu.edu russell.mb.jhu.edu/cns/cnslab.html

^bDepartment of Psychology, University of Kent at Canterbury, Canterbury, CT2 7NP, United Kingdom; m.usher@ukc.ac.uk

snipe.ukc.ac.uk/psychology/people/usher/index.html

Abstract: We argue that the separation between CF (contextual field) and RF (receptive field) in relation to the NMDA (N-methyl-D-aspartate) system is empirically questionable and that it is functionally unnecessary. In addition, the proposed suppression of unexpected information will in many cases be counterproductive.

Phillips & Singer (P&S) are to be commended for their attempt to put the concept of context onto a more physiologically concrete basis than is customary in this field. Understandably, they are more successful in this attempt when they talk about visual perception than in the second major field they tackle, human language. Nevertheless, we take issue with their specific implementation and with its computational rationale, as summarized below.

NMDA cannot mediate synchrony at short time scales. The authors assume that unlike RF connections, CF connections are based on receptors of the NMDA type. It is difficult to see how such connections could mediate synchrony on the time scale of “few milliseconds” (sect. 1.2, para. 2) since the fastest time constants for the closing of NMDA receptors are 35–66 msec (Carmignoto & Vicini 1992), and hundreds of milliseconds are more typical values (Lester et al. 1990). A “simultaneous” increase in firing probability, on this time scale, due to NMDA-mediated CF connections is not suitable for mediating the proposed synchronizing effect since it will only contribute to increase the *baseline* of the cross-correlation function but will not generate narrow peaks (of width less than, say, 25 msec).

Inconsistent context is not always suppressive. P&S assume throughout the target article that the function of the CF is to confirm the hypotheses formed on the basis of previous input. Sensory input from the RF is enhanced when it is consistent with these hypotheses and suppressed otherwise (sect. 1.2.1). We agree that this is the case in many situations, and P&S do provide

additional evidence in support of this function, but this cannot be the whole truth.

To use a standard example: if a dangerous predator appears unexpectedly, P&S's model would predict that information about this unexpected event would be suppressed, since it is inconsistent with the present context and expectations. Obviously, the opposite is the case; such information is processed with a *higher* priority.

It is possible to reconcile these two apparently contradictory capabilities without invoking a separation between the processing of context and content information. We have suggested recently (Stemmler et al. 1995) that the difference between situations in which the context enhances a sensory input and those in which it suppresses it may be the intensity of the input in question. We hypothesized that weak input will be enhanced if it is consistent with its context; strong input, on the other hand, does not benefit from such an enhancing influence and may be suppressed by similar surround stimuli. This is consistent with psychophysical data on threshold facilitation (Polat & Sagi 1993) and apparent contrast (Cannon & Fullenkamp 1993). Our functional argument is that this strategy yields maximal advantage in both cases: if the input is too weak to "stand on its own," the situation may be saved by taking contextual information into account. If the signal is strong, there is no reason to ask for confirmation from the background. On the contrary, a strong signal can be dealt with more efficiently if redundant information (provided by the context) is suppressed.

Multiplicative CFs are unnecessary. We agree with the authors that synchrony is an excellent marker for relational information provided by the context. In previous work (Stemmler et al. 1995; Usher & Niebur 1995), we demonstrated that this can be achieved by regular (non-NMDA, thus fast) lateral connections which are both excitatory and inhibitory. If such connections balance excitation with inhibition, they modify the synchronization of discharges without much affecting the response strength (both inhibitory and excitatory connections can synchronize neural activity; see, e.g., Van Vreeswijk et al. 1994). Finally, we emphasize that context does not need to affect processing in a multiplicative way if context effects are small relative to the stimulus (Usher & Niebur 1996). In fact, evidence from letter perception contradicts the multiplicative assumption (Massaro 1989; McClelland 1991). These data show parallel lines (for z-score of choice probability versus the perceptual continuum) implying additive and independent effects of context and stimulus.

To conclude, we believe that introducing a new distinction between different kinds of processing (RF versus CF) is not a productive strategy in the search of common principles of information processing.

ACKNOWLEDGMENT

The work of EN is supported by a grant from the Markey Foundation.

Synchronization, binding, multiscale dynamic processing, and neuron sociology

Paul L. Nunez

Brain Physics Group, Department of Biomedical Engineering, Tulane University, New Orleans LA 70118; pnunez@mailhost.tcs.tulane.edu
tatanka.bmen.tulane.edu/res/brain.html

Abstract: Well-posed questions about information processing may require physiologically based, quantitative models of large scale neocortical dynamic function. "Synchronization" of this dynamics can be viewed in different contexts of the binding problem.

Dynamical systems and cortical information processing. Phillips & Singer (P&S) illustrate that, given our meager understanding of the behavior of even relatively small neural "circuits," forming well-posed questions about cortical information processing

is challenging. For example, exactly what is meant by "common foundations" in the heading of section 2, "Arguments for and against common foundations for cortical computation." Or consider the cited conclusion of Macphail (1987), "all problem-solving abilities of nonhuman animals arise directly from a common basic associative process." These phrases contain strong tautological facets. In fact, when reading the cognitive science literature, this physical/neuro scientist is often reminded of some of the writings of Newton's contemporaries in which concepts for which we now have distinct mathematical definitions (e.g., force, momentum, and energy) were used almost interchangeably, thereby severely limiting communication between scientists.

How can we forge stronger connections between the subfields of neuroscience and pose questions about cortical information processing in a more precise and broader scientific context? P&S cite literature on artificial neural networks and information theory, which may provide interesting metaphors and suggest useful conjectures. But these fields appear far too limited to provide any kind of general framework in which to view cortical function. Alternately, we may view neocortex as a complex, multiscale, dynamical system. We might first search for agreement between theoretical models of neocortical interaction and physical data (EEG, MEG, PET, MRI, etc.) which can be partly predicted with combined physics/physiology/anatomy approaches, independent of cognitive data. Next, we should identify the most robust correlations between cognitive and physical data. Thus, we can infer correlations between the theoretical models and cognitive events. Once this has been accomplished, we will be in a far better position to form well-posed questions about, for example, cortical information processing, questions that suggest unambiguous, new experiments.

"Synchronization" of neural populations. P&S make a compelling case for the importance of synchronized population codes and cite experimental evidence for synchronized potentials recorded from intracranial electrodes. In complex systems, however, the general idea of "synchronization" requires much closer scrutiny because it is both scale and frequency dependent.

One other measure of synchrony is coherence, which measures phase consistencies at different frequencies between two signals. For example, scalp-recorded EEG coherence typically decreases in the alpha band during transitions from resting to states of mental calculation, while in the same data, theta band coherence increases (Nunez et al. 1997). Furthermore, coherence between two neural populations can depend critically on population size (Nunez 1995). Very large scale coherence between cortical locations separated by 10 or 20 cm (measured with scalp electrodes) can be high, whereas small scale coherence (measured with intracranial electrodes) between the same locations are likely to be small (Bressler 1995; Silberstein 1995). For example, subdural coherence measured with 2 mm diameter electrodes typically falls to zero at all frequencies for electrode separations greater than about 2 cm (Bullock et al. 1995). Thus, some electrophysiologists suspect that moderate to high EEG scalp coherence is due only to passive spread of currents in the head volume conductor. We now have conclusive evidence, however, that long range EEG scale coherence can be large even when volume conduction effects are removed (Nunez 1995a).

In these scalp experiments, coherence between neural populations of roughly 10^7 (high resolution EEG or MEG) to 10^9 (conventional EEG) is measured. By contrast, intracranial recordings may involve tissue containing anything between one and perhaps 10^6 neurons, depending on electrode size and location. "Synchronization" (as well as other dynamic measures) between cortical regions depends on both spatial scale and EEG oscillation frequency. This point apparently accounts for misunderstandings between clinical or cognitive electroencephalographers and animal electrophysiologists. One conjecture is that cortical information processing can occur simultaneously at multiple scales, and the rules governing such processing may be scale dependent. For

example, when expressed in manageable form, the rules (e.g., equations) that govern interactions between fluid elements (e.g., in weather systems) are quite different from the rules that determine molecular interactions within the fluid, even though the former are partly dependent on the latter. It would appear that any successful neocortical dynamic theory must also be closely tied to its appropriate measurement scale, and that theories formulated at different scales are likely to look very different.

Mechanisms of cortical information processing may perhaps be illuminated by considering other complex systems. However, such physical systems typically lack properties critical to neocortical dynamics, which severely limits their use as metaphors. From a dynamic viewpoint, the most important neocortical properties appear to include: competition between functional integration and functional segregation; multiscale interactions involving neurons, minicolumns, macrocolumns, and so on; multiple interaction lengths, including the size of synaptic cleft through the length of the longest corticocortical fibers; and multiple time scales, including synaptic delays and anterior/posterior corticocortical propagation times. These basic physiologic time scales must somehow be compatible with much longer scales associated with conscious experience, short-term memory, and so on.

A sociological metaphor is appropriate to illustrate the dynamics underlying neocortical behaviors. For example, person-person interactions are both local and nonlocal. Important interactions take place across spatial scales, as when information aggregated by a TV station is transmitted to individuals. Communities at different locations and scales have varying degrees of isolation so that the rules of sociological dynamics are inhomogeneous and scale dependent. Adopting P&S's parlance, one may ask if "common foundations" underlie information processing in a marriage, a football game; a revival meeting, political convention, a city, and so on. There are certainly many common features, but at what point do "features" become "foundations," and how important is spatial scale to the rules governing information processing?

Neocortical dynamic theory. Various approaches to neocortical dynamic theory at the large scales appropriate for comparisons with both scalp EEG and cognitive data are reviewed in Nunez (1995). Such theories attempt to integrate data obtained at small scales (e.g., single neuron behavior) with theoretical tools and macroscopic anatomy/physiology so as to predict generic aspects of large scale dynamic behavior, that is, general expected behavior which is relatively independent of assumptions about poorly known neurophysiological parameters. The methods are complementary and include a statistical mechanics of neocortical interactions (Ingber 1985; 1995), local "circuit" approaches in which PSP rise and decay times are critical (Freeman 1975; van Rotterdam et al. 1982; Wilson & Cowan 1973; Zhadin 1984), and global theory in which delays in corticocortical fibers play a dominant role (Katznelson 1981; Nunez 1974). More recently, local and global theories have been shown to combine naturally (Jirsa & Haken 1996; Nunez 1989; 1995; Wright & Liley 1996), so some brain states may be more strongly influenced by local dynamics (functional segregation) and other states more dependent on global dynamics (functional integration). Also, plausible connections between local/global theory and the statistical mechanical theory have been advanced (Ingber 1995; Nunez 1995a).

In contrast to artificial neural networks (which require free parameters), such physiologically based theories have made correct qualitative, and in some cases, quantitative predictions of generic neocortical dynamic behavior. For example, zero phase lag of field potentials over remote regions having no direct connections and simultaneous local/global resonances near 10 and 40 Hz occur naturally in such global systems (Andrew & Pfurtscheller 1996; Nunez 1995a; 1995b; Silberstein 1995). The occurrence of multiple global oscillation modes (identical resonant frequencies at all locations) suggests a possible mechanism for "binding by synchrony" (Singer 1993). For example, oscillations generated in local cortical regions A, B, and so on (internal "pacemakers") can

drive multiple global modes. These enhanced global oscillations can then facilitate remote cell assemblies C, D, and so on, even when they involve overlapping tissue (refer to pp. 694–702 of Nunez, 1995a, for mathematical support).

In summary, I find several of P&S's qualitative arguments compelling, but I suggest that many are better addressed in the context of physiologically based, dynamic theories, even though such theories can provide only very crude representations of actual neocortical dynamic processes.

ACKNOWLEDGMENT

Thanks to Lester Ingber for helping to make his neocortical statistical mechanics appear less sadistical.

Synchronicity and its use in the brain

Guenther Palm and Thomas Wennekers

Department of Neural Information Processing, University of Ulm, D-89069 Ulm, Germany; palm@neuro.informatik.uni-ulm.de
www.informatik.uni-ulm.de/ni

Abstract: We briefly review the long-standing ideas about the use of synchronicity in the brain, which rely on Donald Hebb's views on cell assemblies and synaptic plasticity. More recently the distinction among several timescales in the description of neural activity has become a focus of theoretical discussion. Phillips & Singer's target article is criticized mainly because it does not distinguish these timescales properly and hence does not really address the questions so intensely debated today.

The idea that temporal coincidence of neural activity is translated into togetherness has a long history (see Palm 1982). In sensory areas of the brain, it means that neurons that are activated simultaneously signal co-occurring events in the outside world. Neurons that signal events which often occur together should also belong together internally, for example, by strengthening mutual excitatory connections through Hebbian synaptic plasticity (Hebb 1949). In areas further away from the direct sensory projection this coincidence of activity still means togetherness. It is efficiently detected by a typical neuron because synchronous activation of different afferents of a neuron is much stronger than asynchronous activation (Abeles 1982); it may be further enhanced by Hebbian synaptic plasticity. These ideas, of course, call for mechanisms of synchronization. Here we have to rely on theoretical arguments and in order to relate them to experimental observations we have to specify the timescale on which synchronization has to be achieved.

As a matter of fact, there are several timescales that have to be specified: (1) The time window for synchronicity; (2) the amount by which mechanisms of synchronization can shift time differences between neural events; (3) the time it takes to establish synchronization in the ongoing activity, and (4) the time it takes to change the network connectivity (Hebbian learning) due to previous synchronization.

One would ordinarily expect at least one order of magnitude between one timescale and the next. For example, if we take the synchronicity window in the range of 1–2 msec, then (2) should be at least 10 msec, (3) 100–200 msec, and (4) at least several seconds. These timescales are probably the ones that most researchers would agree on today (see also Palm 1993; Wennekers & Palm 1997).

Theoretical ideas concerning synchronization of neural activity have been put forward in two contexts: in the context of coupled oscillations, well-established theories of synchronization have been developed either for abstract oscillator networks (Kuramoto 1984) or more realistic spiking neurons (review in Gerstner 1995); in the context of synfire-chains as a simplified description of the physiological spreading of neural activity, the issue of synchronization has been one of the starting points (Abeles 1991).

In the experimental literature, the idea of internally generated synchronization lies at the heart of concept pairs such as “stimulus-locked versus stimulus-induced” synchronization (Eckhorn 1994) or “raw” versus “stimulus corrected” correlation (cf. Aertsen et al. 1989). A peak in the raw correlation between two neurons can still be explained by the two neurons being locked to the stimulus, whereas a peak in the stimulus corrected correlation can only be generated internally even if it is stimulus-induced, that is, if it occurs only for certain stimulus situations.

In the context of the ideas outlined above the target article is extremely vague exactly at the following interesting points.

(1) The different time-scales are not well distinguished in the theory. New stimuli are sampled at every time increment or every few in the simulations (e.g., P&S's Fig. 10). This suggests a step size of a large fraction of a stimulus epoch, which is clearly inappropriate for revealing any spike timing effects in the msec range. Nevertheless, the outcome of the simulations is explained on a short scale, intentionally scale (1), since it is discussed in relation to experimentally observed spike synchronization.

A consistent interpretation with a step size of 1 msec is not possible, because stationary firing probabilities close to one (reached after an initial transient phase) would imply rates of 1000 Hz. Physiological rates of about 50 Hz would imply a step size of about 20 msec, which is scale (2). But then again spike timing on scale (1) cannot be resolved. Lowering the stationary probabilities to .05 by rescaling the neural squashing function would also lead to 50 Hz for a step size of 1 msec, but does not solve the problem. Due to the probabilistic spike generation, cells then fire – on average – every 20 timesteps, but not necessarily synchronously.

(2) The target article does not propose a mechanism for synchronization. In response to a sudden stimulus onset, local processors linked by “coherent” context fields raise their firing probabilities simultaneously. This network response is taken as evidence for synchronization. However, since it is input driven, it can at best be related to stimulus locked signals in experiments (Eckhorn 1994). Moreover, it does not imply internal synchronization: in response to a stimulus consisting of two widely separated bars, subsets of cells driven by each of the bars simultaneously increase their rates. This implies neither that cells between those subsets are synchronized (which is not the case in experiments), nor that they are synchronized inside any of the sets (which they appear to be). Similarly, the stationary state achieved in the simulations after the stimulus locked response also fails to reveal synchronization. Cross-correlograms between any two “coherent” cells during this phase are flat, neither peaked nor oscillatory, because the neurons fire in every time-step. Thus stimulus induced synchronization seems to be completely absent.

(3) In fact P&S's simulations show that contextual guidance is possible without oscillation or internal synchronization, since the external synchronicity (on a firing rate level) provided by the stimulus is sufficient. This observation accords well with P&S's note that the simulations work the same way whether spiking or graded response neurons are used (sect. 3.1, last paragraph).

In summary, the target article takes up and discusses some very interesting issues concerning synchronicity and synchronization. However, in the computer simulations it neither provides mechanisms for such synchronization nor shows the effects of it. It only considers synchronicity provided by the external stimuli and does not clearly distinguish the time-scales. Thus the theoretical level of the discussion is brought back to the time of Hebb's original books.

Schizophrenia as a model of context-deficient cortical computation

Steven M. Silverstein and Lindsay S. Schenkel

Department of Psychiatry, University of Rochester Medical Center, Rochester, NY 14642; sslv@cvs.rochester.edu

Abstract: Phillips & Singer's compelling presentation is weakest in its demonstration of commonalities between sensory plasticity and higher forms of learning and behavior. We propose that available data on schizophrenia can provide such evidence, because of the presence of impairments in a number of functions central to their model, and strong relationships between these dysfunctions and behavior.

Schizophrenia as a model of context-deficient cortical computation. Phillips & Singer's (P&S's) compelling presentation is weakest in demonstrating commonalities between basic and higher level learning mechanisms. We argue below that data on information processing in schizophrenia can provide such evidence. In particular: (1) schizophrenia is characterized by abnormalities in a number of functions (e.g., perceptual organization) that Phillips & Singer argue are manifestations of a common cortical processing algorithm; (2) relationships exist between these impairments that would be predicted by the model; (3) the proposed neurophysiological bases of these impairments are consistent with the model; and (4) the behavioral and symptomatic features of the disorder that are associated with these impairments suggest important extensions of the range of learning phenomena that may be incorporated into the model.

Perceptual organization appears to be a clear manifestation of the computational process they discuss. Conditions under which it fails to occur may therefore provide useful data regarding its underlying mechanisms. In several studies, schizophrenic patients have demonstrated impaired perceptual organization of visual (Silverstein et al. 1996a) and auditory (Silverstein et al. 1996b) stimuli. Schizophrenic patients are also impaired in their ability to alter perceptual organization of ambiguous stimuli based on current context (Silverstein et al. 1996b), or to develop perceptual organization for unstructured patterns after repeated exposure (Silverstein et al. 1996b). In another case, schizophrenic patients' perceptual organization impairment was reversed after a contextual enhancement, suggesting that a greater than normal degree of contextual input is required for perceptual organization to occur in schizophrenia (Silverstein et al. 1996a). Comparing the cortical activity of schizophrenic and control subjects during perceptual organization tasks may thus clarify the mechanisms involved in perceptual grouping, and the hypothesized ability of these mechanisms to adapt to current context.

Schizophrenic patients are also characterized by reduced contextual influences in functions such as memory (Calev et al. 1983), attention, and language processing (Cohen & Servan-Schreiber 1992). Thus, schizophrenia may involve a generalized disruption of the proposed cortical processing algorithm (possibly the result of a pervasive cytoarchitectural and/or neurotransmitter abnormality). Further evidence for this comes from an examination of the relationships between perceptual organization ability and language production in schizophrenia (Knight & Silverstein, in press). Consistent with predictions, degree of perceptual organization impairment was significantly correlated with levels of disorganized and associative thought disturbance, the two forms of thought disorder that best reflect the idea of a loss of context during speech production. Schizophrenia thus provides multiple opportunities for exploring specific and general aspects of a cortical computation dysfunction. In addition, examining the normalization of these impairments during pharmacologic treatment may reveal important information about the physiology of normal cortical computation.

P&S hypothesize that the physiological mechanisms underlying context-dependent changes involve NMDA receptor functioning. In this regard, it is interesting that a recent model of schizophrenia points to NMDA receptor hypofunction as an etiological factor

(Olney & Farber 1995). This model is based, in part, on the ability of NMDA antagonists to produce the positive (i.e., hallucinations, delusions) and cognitive/disorganized (i.e., thought disorder, disorientation) symptoms of schizophrenia in healthy controls. In two studies, we found that both context-related impairments in perceptual organization and a failure to develop perceptual organization after repeated exposure were significantly associated with positive and cognitive/disorganized symptoms in samples of schizophrenia patients (Knight & Silverstein, in press; 1996c; 1996d). These data support the hypothesis that normal perceptual organization and other context-related phenomena involve NMDA receptor activity.

Evidence that more complex forms of learning such as social skill acquisition are related to basic computational processes comes from consistent findings that perceptual organization impairment is found almost exclusively in patients with histories of poor premorbid (i.e., prior to the onset of symptoms) social functioning. Moreover, the perceptual impairment is most pronounced among these patients when they are experiencing an acute exacerbation of their illness (Silverstein et al. 1996a). These data suggest that a large subgroup of persons who develop schizophrenia may be characterized by: (1) a developmental failure to benefit from context, as manifested in impaired social cue recognition and social competence; and (2) a more severe disruption in cortical computation associated with illness onset and later periods of acute psychosis (characterized by glutamatergic and dopaminergic dysregulation), which leads to disorganization in basic functions such as perceptual organization.

P&S state that the adaptive phenomena they discuss are implemented by the ability of processors and networks to conform their activity to the statistical structure inherent in the external world by maximizing the transmission of information that is predictably related to the context in which it occurs. We suggest that this function may be the basis of reality testing, the absence of which is a defining feature of the acute phase of schizophrenia. If acute schizophrenia is characterized by a general failure to flexibly adapt distributed processing on the basis of external context, this could explain the occurrence of symptoms such as hallucinations, which appear to be the result of internal context dominating perception. The onset of hallucinations in acute schizophrenia (and attenuation during pharmacologic treatment), and the relationship between hallucinations and perceptual organization again suggests that schizophrenia is a fertile environment for exploring a number of relevant issues, including the avoidance of hallucinations in normal word processing noted by Phillips & Singer. We hope these comments highlight the usefulness of integrating their model with the rich data on information processing in schizophrenia.

On the normalization of coherent contrast and the semantics of synchronization

Darragh Smyth

University Laboratory of Physiology, Oxford University, Oxford OX1 3PT, United Kingdom;

darragh.smyth@physiol.ox.ac.uk www.physiol.ox.ac.uk/rdms

Abstract: This commentary describes some extensions of the Phillips & Singer model of contextual interactions to cater for the contrast- and extent-dependency of surround facilitation and suppression. I also comment briefly on some semantic problems with what exactly the role of synchronization might be: input preprocessing or output postprocessing?

Phillips & Singer (P&S) present an excellent review of the evidence for facilitating influences on processing and learning from outside the classical receptive field (RF). I will comment on some extensions to their model of contextual effects on the rate of spikes while pondering some of the claims about effects on the timing of spikes.

There are two features about contextual stimuli in vision that have been neglected in the model of contextual guidance de-

scribed. First, there is evidence for suppression when the surround is stimulated by large-field textures (references in sect. 1.3, paras. 1 and 2; also Kastner et al. 1997). Second, there is plenty of evidence for contrast-dependent long-range facilitation and suppression (Levitt & Lund 1996; Mizobe et al. 1996; Polat & Norcia 1996; Weliky et al. 1995). P&S choose to attribute the former to the classical receptive field (RF) while the latter is not catered for at all with the monotonic model described in section 3 of the target article.

Attributing the effects of texture suppression and pop-out to the classical RF may be purely a semantic point but I feel it is more appropriate to explain as much of the data on contextual interactions as possible within the same model. Excitatory and inhibitory long-range interactions are mediated via layer 2/3 cells. This is reasonably homogeneous throughout cortex. Recently it has been shown that although large surround textures induce response suppression, the amount of suppression depends on the contrast between the center and surround in the orientation and direction domains (Kastner et al. 1997), as distinct from the luminance domain. Thus surrounds with orthogonal orientations or opposite directions of motion induce larger firing rates than homogeneous textures. This may suggest that there is a sort of receptive field structure in the orientation-direction-position space. I feel that it is more convenient to interpret contextual fields as having different subfields of ON or OFF polarity across the orientation, direction and topographic maps in area 17. There may be similar subfield segregation of lateral connections in other parts of the brain if the model of contextual guidance is a common cortical algorithm. Thus, in effect, contextual fields are performing exactly the same type of integration as classical RFs, except that they have a modulatory influence on the local response, instead of driving.

The contrast-dependency has been explained by differential effects of excitatory and inhibitory inputs via lateral connections (Stemmler et al. 1995). However, this model depends on different levels of noise in excitatory and inhibitory cells and thus may not be so applicable. An alternative model is to apply contrast normalization to cell responses. Cell responses are half-squared and normalized by the overall activity of the other cells within the hypercolumn (and beyond, as suggested by this extension). This model has been successfully applied to cells within hypercolumns to explain the effects of superimposed stimuli of different orientations (Heeger 1993). It may be necessary to extend this model to lateral connections. This may explain why large-field textures give lower responses. It must be pointed out that these two models are not mutually exclusive as there is some evidence for both occurring in the brain.

To summarize this point, P&S's view of contextual interactions can be extended to include all such interactions, facilitative and suppressive, while preserving the concept of maximizing coherent variation. This is achieved by considering modulatory contextual fields (CFs) to be integrating across ON and OFF subfields in the output feature space. In addition, some form of contrast-dependent normalization is required to explain the non-monotonicity of the experimental results. Effectively, lateral interactions are modulating and normalizing signals according to the presence of coherence in the feature space, in contrast to their background.

My second point concerns P&S's description of the effects of context on the timing of spikes. The argument appears to be a little fuzzy so I would like to clarify some issues. It is continuously argued that synchronization is a dynamic grouping mechanism. However in section 1.2.3, paragraph 3, P&S imply that RF filters can signal more than one feature using different dynamically grouped subsets of their inputs. Does this not undermine the notion that a RF is a filter for one particular feature that is hard-wired into its anatomical configuration of excitatory and inhibitory inputs? In section 1.2, paragraph 3, they rightly argue that synchronization improves integration and therefore signaling. This is probably necessary for responding to weak stimuli and/or in the presence of noise. However is there any general evidence that the

filter of a RF is not hard-wired into the anatomy, or could it actually be something dynamic as P&S suggest?

It is too easy to view synchronization as solving the “binding” problem. One could argue that this implies the role of a homunculus in perception. If RFs respond to singular features only then synchronization plays an important role in the integration of weak or noisy signals at the next stage. This is not mutually exclusive with the possibility of synchronized population codes; it simply depends on the level of abstraction at which the researcher decides to view the code. We need to think of experiments to test the function of synchronization as a presynaptic mechanism rather than a post-synaptic one. The brain is a multi-layered system and so with synchronization occurring at all levels (V1, V2, MT, etc.), it is not simply some function of the “outputs” of a neural network.

ACKNOWLEDGMENT

Special thanks to Dario Floreano for many constructively critical conversations deciphering the role of context in the brain.

Information theory: The Holy Grail of cortical computation?

James. V. Stone

Wellcome Research Fellow, Psychology Department, Sheffield University, Sheffield, S10 2UR, England;

stone@aivru.sheffield.ac.uk www.shef.ac.uk/psychology/stone

Abstract: Simple hypotheses are intrinsically attractive, and, for this reason, need to be formulated with utmost precision if they are to be testable. Unfortunately, it is hard to see how Phillips & Singer's hypothesis might be unambiguously refuted. Despite this, the authors have provided much evidence consistent with the hypothesis, and have proposed a natural and powerful extension for information theoretic approaches to learning.

Characterising the “Holy Grail” of cortical computation. The idea of a common foundation for cortical computation was explored as early as 1970 by Marr (Marr 1970), and has been discussed by a number of authors since (Creutzfeldt 1978; Douglas & Martin 1994; Ebdon 1993; Stone 1996a; Szenatgoathai 1978). To date, the idea has remained an intriguing possibility, rather than a testable hypothesis. Despite the efforts of Phillips & Singer (P&S) to place this hypothesis on a firm footing, however, it is not obvious how one would set about falsifying the hypothesis in the form provided in this paper. The authors have not indicated what type of experimental findings would suffice to refute the hypothesis. For example, if it were possible to demonstrate that two cortical regions used qualitatively different operations, would this constitute a refutation? For a number of reasons, I think it would not; but this is due to my own interpretation of the hypothesis, and not a logical imperative of the hypothesis as expressed in the paper. The issue is not whether such a finding would constitute a refutation, but rather, whether *any* finding is capable of unambiguously refuting the hypothesis.

For example, P&S note that, “although ‘cortical columns’ are not central to the hypotheses developed here, criticism of this idea suggests limitations upon anatomical arguments for commonalities” (sect. 2.1, last para.). This is a fair point, but there is no clear indication of what types of arguments could not also suggest limitations on their use. Surely a necessary first step in searching for the “Holy Grail” of cortical computation is to establish what form it might take, and what forms it could not take. Whilst the search for common foundations for cortical computation is undoubtedly a worthwhile endeavour, it would be helpful to establish conditions under which to call off the search, and conditions under which we might shout, “Eureka!”

Is maximising information enough? Several authors have devised neural network models which learn by maximising information theoretic measures (Becker 1996; Linsker 1988; Stone

1996b). As noted by Phillips & Singer (P&S), such rules operate in the absence of any ethological considerations, so that even variables which are irrelevant to an animal's behaviour would be extracted by such rules. This appears to be a fundamental limitation of information theoretic models, which either extract any variables (Linsker 1988), or only variables which conform to certain assumptions implicit in the learning algorithm (Becker 1996; Stone 1996a). By introducing the idea of cross-stream contextual inputs as constraints on learning, P&S have effectively overcome an important limitation on conventional information theoretic approaches. Whilst others have argued that information maximisation methods are appropriate only for low level sensory processing (presumably because they tend to extract all variables in the input data), P&S argue that contextual information can be used to extract selectively only those variables directly relevant to a particular set of behaviours. This important insight opens up the possibility that a single principle can be used to account for learning of low level perceptual invariances, as well as “higher order” variables (such as the association between the shape and taste of a fruit). Moreover, it suggests a natural extension to a whole class of neural network models which learn by explicitly maximising information theoretic quantities (Becker 1996; Linsker 1988; Stone 1996b).

Maximising Shannon entropy is hard. On a more technical level, the general approach adopted for generating learning rules shares with others the assumption that information theoretic quantities are suited to extracting variables implicit in “sensory” data streams. Whilst information theory provides a principled method of deriving learning rules, it is not the only, nor necessarily the best, means to this end.

In its “raw” form, Shannon entropy takes no account of the temporal ordering of inputs. A number of authors (Barlow 1996; Becker 1996; Foldiak 1991; Stone 1996a; 1996b) have argued for the use of learning rules based on the tendency of distal variables to vary smoothly over time. Indeed the BCM learning rule discussed by P&S (sect. 3.5, para. 3) is important precisely because it takes advantage of the temporal sequence of inputs. More recently, Becker (1996) has used temporal smoothness as an explicit assumption in deriving an information theoretic learning rule. Whereas P&S's learning method ignores temporal contiguity of inputs, Becker's method ignores cross-stream constraints. It remains to be seen if these two approaches can be profitably combined, and if the resultant learning method can be related to learning in the neocortex.

ACKNOWLEDGMENTS

Thanks to Nikki Hunkin and Kevin Gurney for comments on this document.

Information: In the stimulus or in the context?

Giulio Tononi and Gerald M. Edelman

The Neurosciences Institute, San Diego, CA 92121; tononi@nsi.edu

Abstract: The distinction between receptive field and conceptual field is appealing and heuristically useful. Conceptually, it is more satisfactory to distinguish between information from the environment and from the brain. We emphasize here a selectionist view that considers information transmission within the brain as modulated by a stimulus, rather than information transmission from a stimulus as modulated by the context.

Contextual field and reentry. The key distinction between receptive field (RF) and contextual field (CF) has several appealing features, not the least of which is its simplicity. As long as RF input is identified with that relayed through forward, voltage-independent connections, and CF input with that relayed through backward or lateral, voltage-dependent connections, one has simply the physiological translation of an anatomical or biochemical

distinction. The heuristic value of this distinction is supported by large-scale simulations which, in much the same way as envisioned by Phillips and Singer (P&S), explicitly introduced backward and lateral voltage-dependent connections alongside the usual forward connections capable of driving neurons (e.g., Lumer et al. 1997; Sporns et al. 1991; Tononi et al. 1992). These simulations illustrated many of the “contextual” effects of reentry, the process of ongoing, recursive neural signaling occurring via lateral and backward connections (Edelman 1978; 1987). These effects include its correlative, constructive, and associative properties: reentrant signaling can modulate the timing of neural response across neuronal groups, generate new response properties, and influence plastic changes by making them sensitive to signals from many brain regions (Tononi et al. 1992; 1996). As P&S point out, the notions of reentrant signaling and of contextual input have much in common.

At the experimental and conceptual levels, the demarcation between RF and CF is less clear. How reliably can one differentiate experimentally between RF and CF for units in higher sensory areas, or in motor and premotor areas? Conceptually, the distinction between the firing of such units, which according to P&S would indicate the presence of specific features in the external environment, and the so-called “confidence” for that firing, which would indicate its coherence with the local context, is problematic: the assumption that the firing of a given unit represents specific features in the environment, even “distal features,” has been severely questioned. Discriminating between firing and “confidence” of firing is an idealization that could only work if there were an unequivocal distinction, over a short time scale, between active and inactive neurons. Rather than attempting to distinguish between direct and contextual input at the conceptual level, it seems much safer to distinguish simply between the extrinsic mutual information between a brain unit and the sensory input, and the intrinsic mutual information between a unit and the rest of the brain (Tononi et al. 1996).

The role of context in perception and action: “Beyond the information given.” P&S’s account has the considerable merit of making their formulations explicit by using general information-theoretical measures. They suggest that at the time scale of perception and action (mutual) information transmitted by the RF is modulated by information transmitted by the CF. This is a clear improvement over typical information-processing approaches which consider RF inputs exclusively.

There are other reasons, however, to abandon altogether the metaphor of the brain as an information processing device (Edelman 1987), while still using information-theoretical measures. Anatomically, most neurons receive signals from other neurons rather than directly from sensory inputs. Physiologically, the brain is spontaneously active, and is so already *in utero*. Phenomenologically, dreaming and imagery are striking demonstrations that the adult brain can produce meaning without any direct input or information from the periphery. Turning the traditional emphasis on sensory input on its head, we have considered the information that is exchanged within the brain in the absence of external inputs (Tononi et al. 1994). This is measured by neural complexity, which is the sum of the average mutual information between each particular subset of brain units and the rest of the brain. Complexity is high only if individual brain units are not only specialized to signal different events, but also distribute these signals efficiently to the rest of the system. Thus, a complex brain is a collection of specialists that talk to each other, and complexity strikes an optimal balance between segregation and integration of function. It seems to us that the distinction between segregation and integration of function, and their reconciliation in terms of neural complexity, captures the conceptual gist of the distinction between RF and CF in a way that is perhaps more flexible and theoretically better founded.

How should one consider, then, the fate of the signals transmitted for the sensory sheets during perception? In our view, such extrinsic signals convey information not so much in themselves,

but by virtue of how they modulate the intrinsic signals that are constantly being exchanged within the brain along reentrant pathways. These constitute an intrinsic informational context which dominates the response to the extrinsic signals. In information-theoretic terms, for a small value of the extrinsic mutual information between an individual stimulus and the brain there will be a large change in the intrinsic mutual information among subsets of units within the brain (Tononi et al. 1996). Thus, it is not so much CF input that modulates RF input, but extrinsic information from a stimulus that modulates intrinsic information within the brain. We have shown that the extent to which this modulatory action is successful reflects the experience that the brain has of a set of related stimuli. This can be measured by a quantity, complexity matching, which is the change in neural complexity as a result of the encounter with the stimulus (Tononi et al. 1996). According to this analysis, the distinction between transmission and storage of information in the brain vanishes and, at every instant, the brain goes far “beyond the information given” (Bruner 1972).

The role of context in learning: “Adjustment of inner to outer relations.” At the time scale of learning, P&S propose that brain units should increase the mutual information relayed by their RF input that is coherent with the mutual information relayed by their CF input, a principle they call “coherent infomax.” This is another step forward with respect to proposals that deal exclusively with maximizing extrinsic mutual information. It is also, however, only as valid as the distinction between RF and CF. Since most neurons receive most of their input signals through reentrant circuits, and since the brain is spontaneously active, it may be more appropriate again to turn matters on their head and take a selectionist as opposed to an information-processing view of learning (Edelman 1987; 1993). Such a selectionist view is perfectly compatible with information-theoretical analyses: learning is conceptualized as an overall increase in the intrinsic mutual information between subsets of brain units, modulated by information from the external world, rather than *vice versa*. This corresponds to an increase in neural complexity in response to extrinsic stimuli, that is, to an overall increase in complexity matching (Tononi et al. 1996), thus representing the information-theoretic equivalent of an “adjustment of inner to outer relations” (James 1890). As an increase in complexity satisfies the joint requirements of segregation and integration at the shorter time scales of perception and action, an increase in matching satisfies the joint requirements of categorization and association at the longer time scales of learning. In practice, an increase of matching over time will occur if there are neural mechanisms such that, *on average*, the mutual information between subsets of brain units increases more when extrinsic input is present than when it is absent. A biochemical difference between backward (lateral) and forward connections is a way to approximate this goal. Another possibility, not considered by P&S, is that the activity of certain neuromodulatory systems with diffuse projections, which is high during waking and low during sleep, influences the mechanisms mediating plastic changes in the brain (Cirelli et al. 1996).

Global context and diffuse ascending systems. A key reason for the usefulness of contextual influences in learning is what is known as the bias/variance dilemma (German et al. 1992): given a limited amount of time, it is difficult, in a completely unsupervised manner (without some bias), to categorize or cluster input signals with high variance purely on the basis of their statistical properties. Local context, mediated by reentrant input, provides part of this bias: what is learned preferentially is what fits with the local context. But other sources of bias are also important. A key source is provided by the aforementioned activity of diffuse ascending systems or value systems, which signal the occurrence of events that are salient for the entire organism and thereby represent a form of global context. The noradrenergic and serotonergic systems, for instance, diffusely innervate most of the brain, fire tonically during waking but much less or not at all during sleep, fire phasically during salient events, and release neuromodulators

that can influence both synaptic activity and plasticity. We have suggested that, in a process called value-dependent learning, synaptic changes should depend on both local correlations (whether mediated by direct or reentrant connections) and on correlations with the firing of diffuse ascending systems signaling salient events. This form of learning has been shown in simulations to account for adaptive learning and for map registration in the brain (Friston et al. 1994; Reeke et al. 1990; Rucci et al. 1997; Tononi et al. 1992b). While this is just another example of the pervasive influence of local and global context on neural activity and plasticity, it clearly indicates how essential it is to go beyond a strict stimulus-driven, information-processing approach to brain function. As demonstrated by P&S's admirable effort, there is no alternative but to consider, both theoretically and experimentally, the context provided by the entire brain and by its history.

Synthesizing synchrony versus dissecting dissonance

Alessandro Treves

Programme in Neuroscience, SISSA, 34013 Trieste, Italy;
ale@imbo.sissa.it www.sissa.it

Abstract: A series of intriguing experiments has emphasized the need to understand synchronization effects. Information theory is proving to be an important tool for understanding the computations performed by cortical circuits. Nevertheless, Phillips & Singer's target article fails to provide a convincing correlation between synchronous firing and information transmission, which could well prove to be anticorrelated.

Synchronization, be it in neuronal firing in epileptic tissue or in goose-stepping by parading soldiers, rarely indicates intelligent behavior. In contrast, available quantitative evidence suggests that the tremendous computational power of the brain arises out of the functional diversity, or dissonance, among its constituent units. Phillips & Singer (P&S), participating in everyone's search for basic principles in cortical function, Solomonically assign half the cortex's connections to the generation of independent neuronal responses, and half to their linkage. Should we concede that half of our brain's precious connections are devoted to goose-stepping?

P&S concisely review a series of very interesting experiments from Singer's and other labs, which raise the issue of the interaction between the amount of zero-lag synchronization among sensory cortical neurons and the nature of the input being perceived. They also sketch an information-theoretic formalism, by Phillips and coworkers, that makes it possible to derive learning rules in which weight modification depends to a varying degree on a subset of the inputs to a postsynaptic unit, the so-called contextual inputs. Nowhere in the target article, however, do they attempt to establish evidence for a relation between synchronization and the amount of information transmitted by units which learn through such modification rules.

Perhaps aware of the weakness of their central claim, which is the one direct (RF) input they provide to the debate on basic principles of cortical operation, P&S surround it with distantly related (CF) arguments, from a variety of other scientific streams, which together take up most of the target article. They thus set the stage for a first experimental falsification of their claim: if the commentaries on the target article turn out not to be fully synchronized, it means that contextual effects can indeed affect the population response (a trivial observation), but by means other than synchronizing it!

Two cornerstones are still missing for the construction P&S attempt. The first is the use of the correct information-theoretic measures not just to derive efficient modification rules applicable to formal network models in simulations, but to extract quantitative analyses of the operation of real neurons in actual experiments. This is now a blooming enterprise (see, e.g., Golomb et al. 1997; Panzeri et al. 1996; Rolls et al. 1997), of which P&S appear

curiously unaware. The second, a cornerstone which sits directly upon the first, is understanding how to extract information measures from the activity of population of cells, in a way that directly probes the specific contribution of synchronization or, in general, precise spike timing, to information transmission. One very small piece of anecdotal evidence I can provide on this issue is from quantifying the amount of information that populations of 30–60 simultaneously recorded hippocampal cells transmit about the position of the rat, in experiments done in Tucson (see, e.g., Treves et al. 1996). When I randomly shuffled, cell by cell, time bins which corresponded to the same position of the rat, thereby simulating nonsimultaneous recording, little changed, quantitatively, in the very robust (and informative) population code, expressed thus both by simultaneous and nonsimultaneous activity vectors. A great deal of serious work obviously has to be done before evidence can emerge on the actual contribution of precise spike timing. Smart new procedures are necessary, both to address the limited sampling problem (Treves & Panzeri 1995) and to manipulate external correlates, in order to dissect the differential contribution to mutual information of the enormous number of variables available in the firing of populations of cells. Until then, P&S can only claim to entertain a hope that binding up cells into synchronization will prove to squeeze more out of them than letting them free, a hope which I am afraid is unlikely to be fulfilled.

ACKNOWLEDGMENTS

I am greatly indebted to Bruce McNaughton who, like Edmund Rolls, has granted me access to data recorded in his lab.

On the computational basis of synchronized codes

DeLiang L. Wang

Department of Computer and Information Science and Center for Cognitive Science, The Ohio State University, Columbus, OH 43210-1277;
dwang@cis.ohio-state.edu www.cis.ohio-state.edu/~dwang

Abstract: For scene analysis, it is important to ask the question of how synchronized population codes – the basic representation employed in the target article – are generated. Recent computational advances have resolved the critical challenges of rapid synchronization with local coupling and rapid desynchronization. The synchronized codes make real contributions to tackling the problem of scene analysis.

A fundamental aspect of perception is its ability to group elements of a perceived scene into coherent clusters (objects). This ability underlies perceptual processes such as perceptual organization, figure/ground segregation, and separation of multiple objects, and it is generally known as scene analysis (segmentation). Regarding this problem, Phillips & Singer (P&S), in their general framework of cortical computation, have answered the following two basic questions explicitly. (1) How are coherent clusters represented in the brain? (2) What is the neurobiological substrate for the representation? Regarding the first question (the “binding problem”), P&S argue for synchronized population codes. That is, an object is represented by the synchronized firing activity of the scattered neurons coding different features of the object (Abeles 1982; Milner 1974; von der Malsburg 1981). As to the second question, P&S distinguish between receptive fields (RF) and contextual fields (CF), and argue that lateral CF connections linking RFs are the neurobiological substrate for the synchronized codes.

There is another major question that needs to be addressed in the same context: (3) Given CF connections, *how* can the synchronized codes be generated? The question is a computational one, and is important for two reasons. First, it must be answered in order to explain how the brain analyzes various scenes in P&S's framework. Second, addressing this question would amount to

constructing a computational system that does automatic scene analysis: an objective of great significance in its own right. Automatic scene analysis is a fundamental task of machine perception and a tremendously challenging engineering problem. This question, however, is not treated in any depth in the target article; dynamic grouping is touched upon in section 3.4.3, but, from the computational perspective, the functionality is simple and it is unclear how the networks generalize to handle realistic input patterns. On the other hand, there is a large body of literature that deals exclusively with this question, and major progress has been made. My commentary is focused on this issue, as it has important implications on the computational foundation of the target article.

The discovery of long-range synchronous oscillations in the visual cortex (Eckhorn et al. 1988; Gray et al. 1989) triggered many computational studies on the synchronized codes. There are two major computational challenges. The first concerns how to achieve rapid synchronization within a population of locally coupled oscillators.¹ Earlier models for achieving phase synchrony generally rely on all-to-all connections. However, a network with full connectivity indiscriminately connects and synchronizes all activated oscillators; it lacks topological information (Sporns et al. 1991; Wang 1993b). The indiscriminate grouping problem hinders this class of models from addressing real images. The second challenge concerns how to achieve fast desynchronization among different groups of oscillators representing distinct objects. We regard the rapidity of synchrony and desynchrony as particularly important, not only because speed is critical for real-time scene analysis but also because perception is very rapid. It is known that human subjects can segment and identify an object in a small fraction of a second (Biederman 1987; I. Biederman, personal communication, 1994). This suggests that both synchrony and desynchrony must be achieved in just a few cycles if the synchronized codes of 40-γ Hz rhythms (Eckhorn et al. 1988; Gray et al. 1989) are the underlying mechanisms. Because of these two challenges, the synchronized codes have not contributed much to building successful artificial neural systems for analyzing real images.

Somers and Kopell (1993) and Wang (1993a; 1995) have independently recognized the severe limitations of widely used sinusoidal oscillators in generating global synchrony based on local coupling; they proposed using different oscillator models to overcome the problem. More recently, Terman and Wang have proposed and analyzed locally excitatory globally inhibitory oscillator networks (l_{egion}) (Terman & Wang 1995; Wang & Terman 1995). In a l_{egion} network, each oscillator is modeled as a standard relaxation oscillator with two time scales (see also Somers & Kopell 1993). Local excitation is implemented by lateral coupling and global inhibition is realized by a global inhibitor. Whether an oscillator can oscillate is determined by the external stimulus to the oscillator; the connections in l_{egion} modify only the phases of oscillators. Thus, the l_{egion} network is fully compatible with the general RF/CF framework of the target article. The network exhibits a mechanism of *selective gating*, whereby an oscillator jumping up to the active phase rapidly recruits the oscillators stimulated by the same pattern, while preventing other oscillators from jumping up. We have proven that, with selective gating, the network rapidly achieves both synchronization within groups of oscillators that are stimulated by connected regions and desynchronization between different groups. In sum, the l_{egion} network provides an elegant solution to the two challenges outlined above.

The ability of l_{egion} in producing the synchronized codes presents a unique approach to addressing scene analysis. Wang and Terman (1997) applied l_{egion} to segmenting gray-level images, and reported very promising results. For gray-level images, each oscillator corresponds to a pixel, and two neighboring oscillators are connected with a weight proportional to corresponding pixel similarity. As an example, Figure 1A shows one gray-level image, and Figure 1B shows the result of segmentation. The entire image is segmented into 23 regions. Each region

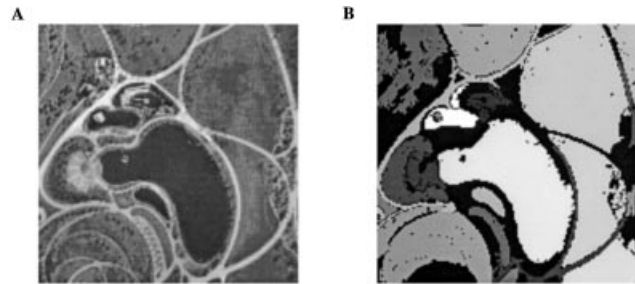


Figure 1 (Wang). A. A gray-level image consisting of 160×160 pixels. B. The result of segmentation by a l_{egion} network with 160×160 oscillators. From Wang & Terman, 1997.

corresponds to a different density in the figure, indicating the phases of oscillators. In the simulation, different regions “pop out” alternately. As can be seen from Figure 1B, almost all major regions in Figure 1A are segmented. The black scattered regions in the figure represent the background that always remains inactive (for details see Wang & Terman 1997). Other images, including MRI (magnetic resonance imaging) images and texture images, have also been successfully segmented.

In summary, I think it is important to ask the question of how the synchronized codes can be generated. Major advances have been made in addressing this question. In particular, the critical challenges of rapid synchrony based on lateral coupling and rapid desynchrony have been successfully resolved and these synchronized codes make real contributions to tackling the hard engineering problem of scene analysis. These advances lay a sound computational ground for the theoretical framework of Phillips & Singer.

ACKNOWLEDGMENT

Thanks to R. Lewis for helpful comments on an earlier version.

NOTE

1. An “oscillator” here is used as a mathematical notion, and this does not necessarily imply an oscillatory (periodic) outcome. An oscillator is compatible with a spiking neuron.

Local attractor dynamics will introduce further information to synchronous neuronal fields

J. J. Wright

Brain Dynamics Laboratory, Mental Health Research Institute of Victoria, Parkville, Melbourne, Victoria 3052, Australia; jjw@cortex.mhri.edu.au www.mhri.edu.au/bdl

Abstract: Simulations and analytic considerations show that synchronisation occurs in delay neural networks at the surrounds of externally driven sites. In the synchronous fields, network transmission has a static transfer function and $H(X|R,C)$ is minimal. But when autonomous local states with attractor dynamics develop in the network, $H(X|R,C)$ is not minimized. Physiological realism may therefore require some modifications in application of coherent infomax.

Basic assumptions. Phillips & Singer’s (P&S’s) arguments, while possessed of considerable explanatory power, assume that small domains of cortex can be described as filters. In the sensory neurophysiological experiments in which synchronous oscillation is commonly observed (e.g., Singer & Gray 1995), CF inputs do not corrupt transmission of RF data (sect. 1.2) and this justifies their attribution to small cortical locales of a transfer function incapable of introducing significant fresh information to the on-ward data-stream. But is this always justified?

Mechanisms of synchrony. P&S are careful to separate dynamic mechanisms of synchronous oscillation from the information theoretic implications of the phenomena. However, their assumption of a static network transfer itself forces constraints on possible mechanisms. These mechanisms are both controversial and (probably) diverse, and include input of synchronous subcortical signals (Steriade et al. 1993), intrinsic oscillations (Llinas et al. 1991) including “chattering cells” (Gray & McCormick 1996), inhibitory-inhibitory oscillations (Traub et al. 1996), entrained nonlinear limit cycles (Tass & Haken 1996), and so on.

A mechanism which requires minimal assumptions and which appears to be a ubiquitous property of coupled neuron-like delay elements has recently been reported. This finding emerged in relation to numerical models of EEG activity (Wright & Liley 1996) and is reported in detail in Wright (1997). Essentially, when any two points of a matrix of elements are driven with independent, asynchronous noise inputs, in approximation of RF inputs, zero-lag synchrony in the neighbourhoods of both sites of input rapidly develops. The extent of the fields of cosynchronous activity reflects dynamic gain factors analogous to inputs of nonspecific cortical activation, as well as structural (cortico-cortical) connectivity. Correlation functions derived from the simulation's state equations (Robinson et al. 1997a; 1997b) confirm that the phenomenon depends only on dendritic and axonal delays and gains. Findings qualitatively resemble experimental findings summarised in Singer & Gray (1995), although approximation to the complex conditions of experiment is as yet crude.

In the conditions in which this simulation has been studied, pulse densities remain at a low level and any small cortical locale is described by a static transfer function. But the simulation also exhibits extreme instability of pulse firing rate, moving from a low pulse density to one of runaway excitation with minimal perturbation. Thus, if the model is adequately parameterised and otherwise valid, then local or global mechanisms must exist in the real physiological situation, to regulate stability much more tightly. Likely stabilising mechanisms were discussed in the *BBS* commentary in Wright and Liley (1996) and have been partially incorporated in simulations in continuing work.

It now appears that the operation of such stabilising mechanisms can lead to localised autonomous, self-organising states. When locales in the network are sufficiently excited to enter autonomous activity they generate outputs which are not describable as a static transform of their inputs.

Local autonomous states and choice of objective function: When is coherent infomax appropriate? P&S's coherent infomax (sect. 3.2) depends upon the minimisation of the Shannon information at a cortical site not accounted for by RF and CF inputs. But what if activity at the site in question cannot be described in terms of a transfer function, because local activity has entered some self-organising state? Such states are very likely in real cortex, as described in Amit's (1995) models of neurophysiological findings in delayed response. The limit cycle and/or chaotic dynamics of these more complex conditions imply that $H(X|R,C)$ is likely to be high, and not accountable as noise.

In future developments of P&S's models, perhaps autonomously active sites ought to be included within the definition of RF inputs. “RF” would then describe both intrinsic and extrinsic sources of information, while coherent infomax would still apply to transmission of information throughout the wider, lower firing rate, domains of the cortex. This generalisation might help avoid problems when experimental data obtained in more active cognitive processes is considered. Moreover an important putative property of coherent infomax – the discovery of distal variables and relationships in diverse data sets (sect. 6.5) – might also apply to interactions between patches of autonomous, self-organising cortical activity, as well as to features from the environment. Such a future model might offer the possibility of introducing the decision-making capacity of attractor dynamics to the association capacity of synchronous fields.

Context dependent feature groups, a proposal for object representation

Rolf P. Würtz

*Institut für Neuroinformatik Ruhr-Universität Bochum D-44780 Bochum, Germany; rolf.wuertz@neuroinformatik.ruhr-uni-bochum.de
www.neuroinformatik.ruhr-uni-bochum.de/ini/people/rolf/top.html*

Abstract: The usefulness of contextually guided processors is investigated a little further. A more general use for binding V1 cell responses than the one in Phillips & Singer's target article is proposed, which takes into account that strong responses of these cells can mean more than the presence of lines and edges. The possibility of different grouping depending on the activities of neighboring cells is essential to the approach.

Phillips & Singer's (P&S's) target article is good news for a computational modeler for several reasons. First, it is rare luck to have a computational concept which is approved by a psychologist and a neurobiologist. Second, this concept is presented at a level of abstraction and simplicity which is not usually adopted by biologists, although it is a crucial prerequisite for reasonable computational models. Such a paradigm can spare (for the time being) the complicated (and computationally expensive) simulation of the nitty-gritty of synchronization. Instead, the proposed processors can be used to show that relatively simple networks can do something useful.

Now, why should new cortical processors be necessary at all? After all, it is clear that the brain consists of neurons, and computational models of neurons abound on all levels. Still, from an esthetic point of view, the system properties observed in the brain should arise “naturally” from the properties of the simple processing elements used. Most prominent among these system properties is invariant recognition in the various sensory modalities. Although much research has been done, there seems to be no convincing way to get this property from conventional neural networks.

Multiplicative synapses (see the exponent of the transfer function, first paragraph in sect. 3.1), have been popular in computational neuroscience for a long time. What is interesting in P&S's paper is that multiplication occurs only between the CF and RF inputs. The proposed processors are ideal for invariant recognition by dynamic routing (Olshausen et al. 1993; 1995). In these articles, contextual guidance is used to test hypotheses about translation and scale of simple objects, with special cells being responsible for the mapping between input and stored model. A drawback of these models is that they operate only on very simple images with good contrast and few ambiguities. The RF properties of V1 cells are not part of these models.

Let me propose a model that creates robust object descriptions from the responses of simple and orientation-selective complex cells. An important property is that these responses must mean different things depending on the response distribution of neighboring cells. Concretely, they can stand for the presence of line elements, texture elements, or local pieces of structure.

In the first case, cells sharing the same RF center and orientation but of different RF size must be grouped. These groups code for the presence of a line/edge element and, as the authors mention, they in turn must group with neighboring groups of the same or slightly different orientation.

In the other cases, all cells sharing the same center must be grouped to code for a local texture or structure element (Lades et al. 1993; Würtz 1995; 1997). If those groups mean a texture element, they should group with neighboring ones that have the same activity distribution to yield areas of constant texture. For structure elements, top down information about the local structure of known objects must be used in order to bind the right elements to an object description. Such a description usually also contains boundary lines, that is, groups of the first type in a suitable spatial arrangement.

Such an approach at object description can be more robust, because constraints like the need for a *closed* boundary can be

Table R1. Views on four basic issues

relaxed. An object can be also be described by an imperfectly closed boundary together with constant texture or known structure.

Now, what is required for this? First, a possibility to group V1 cells without destroying their filtering properties, which is clearly provided by P&S. Second, the process must be cascaded, that is, groups of groups must be possible. This is a bit more problematic, because the number of *different* synchronized groups at a time is very limited. For the system proposed here, this is not a serious problem, because the processors keep their RF properties intact, so the relevant cells for a whole object may be synchronized into one group without confusion. For implementing, for example, a graph structure like the one used by Lades et al. (1993) it might constitute a serious limitation.

If such object descriptions develop in a self-organized way based on the input image it should be relatively easy to match them with stored models of objects, although the integration of flexible matching and associative memory remains an open problem. Anyway, the target article opens a couple of interesting new routes of investigation.

Authors' Response

Progress toward an understanding of cortical computation

W. A. Phillips^a and W. Singer^b

^aCenter for Cognitive and Computational Neuroscience, Department of Psychology, University of Stirling, FK9-4LA, Scotland, United Kingdom; wap@forth.stir.ac.uk www-psych.stir.ac.uk/~wap; ^bMax Planck Institute for Brain Research, D-60496, Frankfurt/Main, Germany; singer@mpih-frankfurt.mpg.de

Abstract: The additional data, perspectives, questions, and criticisms contributed by the commentaries strengthen our view that local cortical processors coordinate their activity with the context in which it occurs using contextual fields and synchronized population codes. We therefore predict that whereas the specialization of function has been the keynote of this century the coordination of function will be the keynote of the next.

We thank all of our commentators for producing thoughtful and insightful commentaries in the few weeks available. The commentaries come from five broad areas of expertise: neuroscience and neurology (13); computing, neural computing, and informatics (10); psychology (5); theoretical brain physics (2); and psychiatry (2). This diversity is a testament to the importance of trying to understand cortical computation. Anything common to such diverse sources of knowledge must reflect a deeper reality. We do not suppose, however, that all these perspectives can be boiled down to any single perspective, including that emphasized in the target article. Just as we suppose that cortical cells can cooperate so as to emphasize coherence while nevertheless maintaining their own distinctive contribution, so we suppose that different scientific perspectives can all shed light on a common underlying reality while nevertheless remaining distinct.

We first present a summary of the views taken on four central issues, and then discuss two of the most basic: the working hypothesis of common foundations, and the role of contextual coordination. The main body of the response is divided into nine sections. R1 to R5 discuss the five issues

Are common foundations for cortical computation likely?

YES Amit, Bower, Bugmann, Eckhorn, Floreano, Gregson, Grossberg, Haase & Diniz, Iacoboni, König et al., McCollum, Morasso et al., Niebur & Usher, Palm & Wennekers, Silverstein & Schenkel, Smyth, Tononi & Edelman, Treves, Wang, Wright, and Würtz.

NO
? Nunez, Stone.

Is the coordination of activity a fundamental goal?

YES Amit, Bugmann, Eckhorn, Floreano, Gregson, Grossberg, Haase & Diniz, Iacoboni, König et al., McCollum, Morasso et al., Niebur & Usher, Nunez, Silverstein & Schenkel, Smyth, Stone, Tononi & Edelman, Wang, Wright, and Würtz.

NO Treves.
? Bower, Palm & Wennekers.

Is the distinction between CFs and RFs useful?

YES Bugmann, Eckhorn, Floreano, Grossberg, Iacoboni, König et al., McCollum, Morasso et al., Silverstein & Schenkel, Smyth, Stone, Tononi & Edelman, Wang, Wright, and Würtz.

NO Niebur & Usher.
? Amit, Bower, Gregson, Haase & Diniz, Nunez, Palm & Wennekers, and Treves.

Does the cortex use synchronized population codes?

YES Bower, Bugmann, Eckhorn, Floreano, Gregson, Grossberg, Haase & Diniz, Iacoboni, König et al., McCollum, Niebur & Usher, Nunez, Palm & Wennekers, Silverstein & Schenkel, Tononi & Edelman, Wang, Wright, and Würtz.

NO Amit, Treves.
? Morasso et al., Smyth, and Stone.

listed in sections 6.1 to 6.5, and in the same order. R6 discusses possible additional foundations proposed by commentators. R7 discusses further issues raised by commentators. R8 takes up the invitation of one of the commentators to pursue the analogy between the scientific enterprise and cortical computation. Finally, R9 notes promising directions for future progress.

Table R1 classifies commentaries in relation to the views taken on four basic issues: (1) Are common foundations for cortical computation likely? (2) Is the coordination of activity a fundamental goal? (3) Is the distinction between contextual fields (CFs) and receptive fields (RFs) useful? (4) Does the cortex use synchronized population codes? Each commentary is classified as "yes," "no," or "agnostic" on each question. Most classifications are based on what is said explicitly, but in a few cases we inferred what seemed implied. Nevertheless, the classifications given are ours. There is no guarantee that commentators would agree.

On each of the four questions Table R1 classifies the majority of commentators as giving the same answers we

do. Few argue explicitly against the positions we take on these issues. They were not short of criticisms, of course, with the most prominent being as follows. **Amit** and **Treves** doubt any role for synchronization. **Bower** says that his biological models suggest a role for it, but a role that differs from ours. **Niebur & Usher** agree with us on the role for synchronization, but doubt CFs. **Palmer & Wennekers** think that synchronization is important, but criticize the computational studies for not modeling activity at the level of individual spikes. Some commentators see no role for information theory. There is thus no agreement between these criticisms, and much disagreement. Several commentaries call for more emphasis upon top-down processes. This is well within the spirit of our approach, and we agree.

All commentaries agree that it is worth searching for common foundations. This consensus does not show that view to be correct. There is no reason to suppose that those whose skeptical views were noted in sections 2.1, 2.2, and 2.3 are no longer skeptical. It is simply that their views are not represented here. It would help if they were, because their perspective suggests aspects of cognition that are unlikely to be directly comprehensible in terms of common foundations. Our search will be led seriously astray if it attempts to explain, without additions, aspects that require further capabilities. With respect to language, at least, the best supported hypothesis is that it does involve special capabilities (Pinker 1994). Whether these depend upon the common foundations, as our perspective suggests, remains to be seen.

Stone and **Nunez** both note that the hypothesis of common foundations is not well specified. We have not attempted any formal specification, simply using the word “common” to mean most or many, and not necessarily all. We do not see how this can be formalized. It is an empirical rather than a formal issue, so we keep our minds open as to what constitutes the range of cortical and cognitive processes to be understood under a common rubric. The more the better.

We now turn to views on the role of contextual coordination. We argued that functional specialization and contextual coordination together form two basic and mutually constraining principles of cortical organization (sect. 1.1). Functional specialization is undisputed. The issue is whether it is balanced against a need for coordination. **Tononi & Edelman** have argued elsewhere (Tononi et al. 1996), both formally and forcefully, that this is so, and von der Malsburg and Singer (1988) have shown how the emergence of global order from local cooperation is a very general property of physical systems composed of many distinct but interacting elements. The forms of contextual coordination emphasized in the target article include the selection of contextually relevant signals, their grouping into coherent subsets, and contextual disambiguation. Although most commentaries seem to agree with the need for such coordination, two raise important doubts. **Treves** argues for the value of dissonance and of letting cells be free rather than having them goose-stepping in synchrony. **Niebur & Usher** argue that as unexpected events can be both important for survival and highly salient, they constitute cases where our emphasis upon coherence is exactly the opposite of what is required.

Treves stressed the possibility that information transmission may be reduced rather than increased by synchrony. We will discuss evidence for synchrony as a cortical

code in section R3.1. Here we focus upon conceptions of the goals of cortical processing. Underlying Treves’s concern for squeezing as much information as possible out of cortical cells is the long and strong tradition that emphasizes the goal of recoding to reduce redundancy (sect. 1.2.1). If the goal is to transmit as much information as possible using a limited and noiseless channel, an optimal recoding will map activity into statistically independent variables. We agree that this perspective makes an important contribution to our understanding of cortical computation, but taken to the extreme it would lead to something worse than chaos. First, neither the external world nor the cortex are noiseless, so redundancy is necessary. More important, the transmission of as much information as possible seems to be a highly implausible goal for cortex. It is more plausible that it discovers what predicts what, including various aspects of reinforcement, and uses that to transmit information about just those variables that matter. This will emphasize signals that are coherently related, and thus reduce total information transmission.

Niebur & Usher note the sudden appearance of a predator as an example of the salience of “unexpected” events; they view this as evidence for the opposite of our emphasis upon coherence. We are well aware of the salience of abrupt changes. They were central to our studies of the detection of appearance and disappearance of single elements in random-dot arrays using psychophysical and physiological techniques (Phillips & Singer 1974; Singer & Phillips 1974). Gradients of change in both space and time are so crucial that they are central to RF specificity at all levels. This has long been clear, and it explains why contrast in either time or space (including higher feature spaces) is so noticeable. Does this imply that incoherent signals are more salient than coherent signals? Not at all. Consider a random-dot array in which two subsets of elements appear, one in some structured or familiar arrangement, the other at random. It is the structured or familiar subset that will be most salient, not the random one. The sudden appearance of a predator will produce activity across a wide array of cells in the visual cortex. Camouflage evolves so as to reduce the coherence of that activity, not so as to increase it. Thus, Niebur & Usher provide an excellent example that expresses at the intuitive and behavioral level our view that coherent patterns of activity, not incoherent ones, are most effective at the cellular level.

R1. Does the distinction between RFs and CFs have biological relevance?

R1.1. Niebur & Usher’s doubts

Most commentaries agree that this distinction is useful. **Niebur & Usher**, however, argue explicitly that it is not. Their grounds for this are that lateral interactions in the cortex can be explained by excitatory and inhibitory interactions alone, making no use of gain-control mechanisms. **Smyth** discusses the same phenomena as Niebur & Usher, but, in contrast, concludes that such lateral interactions support the RF/CF distinction and should be included in the CFs. Who is right? In section R1.7 we will outline a scenario in which they could both be partly right. Here we emphasize two more general points. First, CFs are primarily distinguished by the effects they have, that is, modulatory and synchronizing, not by the source from which they

come. They should not be identified with lateral interactions. Some long-range interactions may be concerned with computing RF features that require comparison with a large surround. RF features defined by comparing a center with a large surround are perfectly compatible with the contextual coordination of those features through CFs, however. Second, CFs are primarily distinguished by the effects they have, not by the mechanism producing those effects. Furthermore, although Niebur & Usher suggest a different mechanism for lateral interactions, they acknowledge elsewhere that voltage-dependent channels (e.g., NMDA) are common (Usher & Niebur 1996), and may play a role in "top-down attentional modulation." Thus, in our terms, they do assume a role for voltage-dependent CF mechanisms, suggesting that it is specifically intraregional voltage-dependent CFs that they doubt. If so, their view seems similar to that of **Grossberg** in this respect (sect. R1.7). The absence of such CFs is not yet clearly established. It cannot be established by showing that some long-range intraregional connections contribute to RFs, nor by showing that they may in part use mechanisms different from those we emphasized.

R1.2. How are RFs and CFs distinguished?

Given that this distinction is useful, how should it be drawn? At least five possible contrasts must be considered: (1) RFs have a driving effect (driving activity either up through excitation or down through inhibition); CFs are modulatory (producing either facilitative or suppressive gain-control); (2) RFs determine what information signals transmit; CFs help determine exactly when they are sent; (3) RFs determine what information signals transmit; CFs specify with which other signals they are to be grouped. (4) RFs determine what decision signals transmit whereas CFs help determine the confidence with which that decision is made; (5) RFs tend to be feedforward; CFs tend to be lateral and feedback.

Tononi & Edelman and **Eckhorn** cite further evidence that most of those contrasts are relevant. Our view is close to theirs but with a few differences of emphasis. **Tononi & Edelman** doubt the usefulness of distinguishing between decisions and confidence in those decisions. They rightly note the difficulty of making a categorical decision by distinguishing between active and inactive neurons. Our working assumption is that different outcomes of categorical decisions are represented by different groups of neurons with reciprocal inhibition between groups. The categorical decision could then be conveyed by the group transmitting more spikes or more synchronized ones, than the others, and the confidence by how much more. **Tononi & Edelman** put more emphasis on contrast 5 than we do. We propose that CFs are distinguished by the effects they have, not by the source from which they come. Nevertheless we agree that CFs are more likely to come from lateral and feedback sources. **Eckhorn** notes that there is little or no physiological evidence for suppressive effects between uncorrelated signals. This is a simple misunderstanding. The computational framework outlined in section 3 implies that CF connections will adapt so that they have no effect if they link streams between which activity is uncorrelated.

R1.3. Terminology for CFs

Given the similarities between our view of CFs and those of **Tononi & Edelman** and of **Eckhorn**, we explain briefly

why we use the terms we do, rather than those they propose. Calling the CFs reentrant connections puts the emphasis upon the source from which they come rather than upon the effects that they have. If it turns out that some bottom-up connections mediate contextual modulation, for example, to speed-up coordination, we could refer to them as feedforward CFs. Calling them feedforward reentrant connections would sound less coherent. **Eckhorn** calls the CFs association fields, as do **Field et al.** (1993). Although this has some appropriate connotations, it seems more appropriate to the retrieval of concepts. We therefore prefer to use the term associative connections to refer to the internal connections in attractor neural networks, such as those analyzed in detail by **Amit** (1989), **Hopfield** (1982), and many others.

R1.4. How could CFs produce synchronization?

It is noted by **Palm & Wennekers** that the computational theory outlined in section 3 does not show in detail how biological CFs could produce synchronization. Previous studies by many others have studied this in detail (sect. R1.9), however. An excellent example of the use of CFs in technological implementations is given here by **Wang**. His success in processing real images shows that CFs would be useful if they had biological reality. Further computational support for synchronizing CFs is provided by **Würtz** who notes that their use may overcome the problem of distinguishing multiple simultaneously active groups, which is a difficulty for the dynamic link architecture (DLA) (**Lades et al.** 1993). Thus, a system using CFs may be less flexible than the DLA, but closer to the biology.

R1.5. Is there a role for inter-regional feedback CFs?

There is an explicit call by **Tononi & Edelman, König et al., Grossberg,** and **Bugmann** for more emphasis upon top-down processes. We agree. Top-down sources for CFs are clearly shown in Figure 1c. Physiological evidence for such effects was cited in section 4, and their relevance to word perception was discussed in section 5.7.2. Nevertheless we understand why several commentators were misled about our views on this, because we did put more emphasis upon bottom-up stimulus-driven contextual coordination. This was not because we doubted the importance of top-down influences. Two classes of top-down effect must be distinguished, stimulus driven and strategically driven. There is clear physiological evidence for top-down contextual modulation that is stimulus driven (**Zipser et al.** 1996), and our computational studies have included such effects (e.g., **Phillips et al.** 1995a; 1995b). **König et al.** now provide evidence for strategic top-down coordination by showing that synchronization depends upon the current task and not just upon the current stimulus. We agree that such effects are important and could be mediated by CFs.

R1.6. Is there a role for internal sources of RF drive?

The flow of activity that arises from within the cortex itself is emphasized by **Tononi & Edelman,** and **Wright** suggests that intrinsic attractor dynamics be incorporated into the concept of a local-processor, and in such a way as to be seen as part of the RF. We see our approach as supporting both proposals. They rightly remind us that internal systems, for

example those concerned with plans and strategies, can provide a source of driving as well as of coordinating activity. Coordination of this internally generated activity is also crucial, however, because our thoughts and plans also need to be coherent (sect. R7.4).

R1.7. How are CFs related to long-range interactions in visual cortex?

The distinction between RFs and CFs is unnecessary, **Niebur & Usher** argue, because the effects of context can be explained without invoking a separation between the processing of context and content. It is not clear whether they are denying the functional distinction or proposing a different mechanism for it, however. Having denied the distinction, they then use it so that the sign of the effect of “context” depends upon the strength of the “input” (i.e., the RF input). If there were no difference between the processing of context and content then the sign of the effect of RF input should depend upon the strength of the context, and they do not claim that to be the case. It thus seems more likely that they are proposing a mechanism for the distinct effects of context. This interpretation is strengthened by their description of the physiological and psychophysical data to be explained as showing contextual modulation (Stemmler et al. 1995). The mechanism they propose is based upon the phenomenon of “stochastic resonance,” in which noise enhances detection of weak signals by increasing the gain at weak signal strengths (Stemmler et al. 1995). Bezrukov & Vodyanov (1997) have now proved that stochastic resonance can occur in voltage-dependent ion channels, however, so the mechanism that Niebur & Usher propose may be more compatible than they thought with the one we emphasized. If so, that would enable them to preserve their account of the relevance of stochastic resonance without having to explain away evidence that long-range intra-regional connections activate voltage-dependent channels (Hirsch & Gilbert 1991). Even that is unlikely to provide a complete account of long-range lateral interactions, however, because it is hard to see how enhancing the detectability of weak elements explains the figural grouping of high-contrast elements (e.g., as in Field et al. 1993).

Smyth, in clear contrast to **Niebur & Usher**, interprets the data on long-range interactions in V1 as evidence for CFs. Kastner et al. (1997) provide evidence for modulatory effects upon cells in visual cortex (V1). Their Figure 4a shows a cell with narrow orientation tuning. It is inhibited by a single central bar orthogonal to its preferred orientation. There is no sign of inhibition being produced by a surrounding array of parallel elements alone, but when they are presented along with an optimally oriented central bar they reduce firing rate from about 20 to about 5 spikes per sec above spontaneous. This is clear evidence for modulatory rather than driving effects, and Kastner et al. describe their findings in terms of effects upon “saliency.” This implies that they view the effects of the surround as not changing what is signaled, but as changing how effective that signal is at subsequent levels. This is in accord with what Smyth proposes.

The view just outlined allows for the dependence of context effects on the strength of the central stimulus. The hypothetical activation function used in section 3.1 of the target article produces smaller effects of context as target

strength increases, and in addition sect. 6.1 suggests that effects of context upon spike rate in cortex may reduce as target strength increases. Studies of context using visual evoked potentials show this to be so (Polat & Norcia 1996). Thus this aspect of “contrast-dependence” is catered for by the simple assumptions made in section 3, contrary to what **Smyth** may be taken as implying. The aspect that he intends to emphasize, however, is the change in *sign* of the context effects with changes in the strength of the central RF stimulus (personal communication with Smyth). We agree that this is not catered for in section 3.

For a broader view of long-range interactions it is crucial to distinguish several levels of processing. Two levels of analysis can be distinguished in the perception of textured stimuli such as those used in many of these experiments (Sagi 1995). At the lower level local features are computed, for example by anisotropic Gabor-like filters. At the higher level, further center-surround comparisons are used to compute gradients of change in the lower-level feature space (e.g., by isotropic Gabor-like filters with larger RFs than at the lower level). The model outlined by Sagi (1995) is concerned only with the basic feedforward RF architecture, so we need to add contextual coordination within levels and from the higher to the lower level. Detectors of distinct stimulus elements could therefore interact through modulatory CF effects at the lower level, but with their outputs being combined in RF computations at the higher level. At the higher level of analysis, however, long-range comparisons are used to compute gradients in feature space, and these are best thought of as being part of RF computation because information about the separate stimulus elements is lost (Sagi 1995). This is in keeping with the view expressed in section 1.3, except that it occurs in secondary rather than primary visual cortex.

This scenario may provide a reply to the second question that **Smyth** raises, that is, it shows how synchronization can have a role in input pre-processing. It allows for the dynamic sculpting of the exact extents of the effective RF centres and surrounds of cells in secondary visual cortex. These would otherwise be too rigid to reconcile, for example, the conflicting demands of segregation and integration (Braddick 1993). There is evidence that this involves top-down modulatory inputs (Salin & Bullier 1995), including those that are stimulus-driven (Zipser et al. 1996), as well as those that are task-dependent (Maunsell 1995). Zipser et al. (1996) found evidence for top-down contextual influences on V1 cells by showing that firing rate is increased by any of a variety of higher-level cues that made the features being signaled part of the figure rather than the ground. Synchronization of firing was not studied. Given the evidence for diffuse axonal bifurcation of the feedback connections (Salin & Bullier 1995), and given that pyramidal cells respond to synchronization over the relevant timescales (König et al. 1996), we predict that the outputs of V1 cells sensitive to different elements within the figure will tend to become more synchronized by the common top-down inputs that they receive.

An important difference between **Grossberg’s** perspective and the one we outlined is that in his models top-down and lateral connections have very different roles, with only the top-down being analogous to our CFs. He states his conviction that the function of horizontal intraregional connections is to complete missing information, as over blind spots and retinal veins, but cites no evidence that

convinces us of this. It is hard to see how it could be the main part of the story, because lateral intraregional connections in V1 predominantly link cells with normal retinal input.

R1.8. What is the distribution of CF inputs?

Given the usefulness of distinguishing CFs from RFs, **Eckhorn** and **Morasso et al.** note that a major issue that arises is to map the distribution of facilitatory and suppressive CF inputs. We agree. Much relevant data concerning their distribution in V1 already exists (e.g., Kastner et al. 1997; Levitt & Lund 1996; Polat & Norcia 1996; Schmidt et al. 1997; Weliky et al. 1995; Zipser et al. 1996), but even there important issues await further clarification. In other regions there is much new territory to be explored, as Morasso et al. emphasize.

CFs are local in that each processor has its own particular set of direct CF inputs. All-to-all CF connectivity is neither possible nor desirable (**Wang**). If CF influences can be mediated through other processors, however, then each processor could be influenced indirectly by a much larger set. **Floreano's** simulation of mediated long-range contextual influences clearly demonstrates this possibility. He shows how activity in distant parts of the network can influence RF learning given a chain of intermediate CF connections. An important constraint on this possibility, however, is that the mediating processors in the chain must themselves simultaneously receive RF input. This is in striking agreement with the capabilities and constraints observed by Polat and Sagi (1994b) in their psychophysical studies of perceptual learning in human vision.

R1.9. What synaptic receptor channels do CFs use?

Doubts are expressed by **Niebur & Usher** about whether NMDA-receptor channels can mediate synchrony at short-time scales. NMDA channels have long been thought to have time constants of more than 150 msec, but it is now known that there are also fast components to the NMDA response with time constants closer to 15–50 msec (D'Angelo et al. 1990; Hestrin et al. 1990; Monyer et al. 1994). Although the fast components of the NMDA response are not well known, they have been used in models of STM (Jensen et al. 1996; Lisman & Idiart 1995). None of this implies that there are no other mechanisms for synchronization, and there is evidence for others (e.g., Cobb et al. 1995; Traub et al. 1996).¹

R2. Does the distinction between RFs and CFs (and/or synchronization) have psychological relevance?

In section 5 we argued that cognitive and neuropsychological evidence from a wide variety of paradigms supports our approach. Few of those arguments were questioned. Overall, commentaries did not subtract much from the psychological case presented, but, with one notable exception, neither did they add much. This suggests that a great deal remains to be done to foster facilitatory interactions between psychological and neurobiological studies of contextual coordination. The encouraging exception is provided by **Silverstein & Schenkel**. They have long used a variety of cognitive psychological paradigms to study perceptual

grouping and other forms of contextual coordination in schizophrenia and conclude that it is selectively impaired in a distinct subset of patients. If so, studies of such states could provide new insights into the broader psychological significance of contextual coordination and of its pharmacological bases. Their suggestion will be considered closely in section R7.4.

R2.1. What is the role of top-down contextual control in cognition?

Internal sources of RF drive are dominant over external sources, **Tononi & Edelman** argue. This contrasts with the greater emphasis in the target article upon stimulus-driven effects. We agree that this is a major issue, and that the contextual modulation provided by strategic control may have much import for psychology. Perhaps the balance between internal and external sources varies, with the balance shifting too far toward external sources in some cases (e.g., attention deficit hyperactivity disorder, frontal lobe damage), and too far toward internal sources in others (e.g., psychotic delusions, dream states). They also argue for the importance of global context as provided by diffuse ascending systems. We agree, but emphasize that they do not carry semantic information and differ anatomically, pharmacologically, and functionally from mechanisms for local contextual coordination.

König et al. also emphasize the importance of internal strategic control variables, and show that they are amenable to combined electrophysiological and behavioral study. Their results indicate that such control uses similar CF synchronizing mechanisms as do stimulus-driven effects, and thus suggest promising new directions for future research. Section 1.4 in the target article notes several prior theories with similarities to our approach, including that of **Grossberg**. We did not specifically relate his work to these issues, however, and he has now done that better than we could have done. Instead, we put more emphasis upon relations to other theories that are closer to ours in spirit, and to the neurobiology in detail. Nevertheless, it is encouraging that Grossberg also sees important roles for synchronization and contextual modulation, for example in the top-down matching process of ART. **Bugmann** also makes a good case for top-down task-dependent contextual effects. We agree, but note that this does not imply the absence of bottom-up stimulus-dependent contextual effects. We await with interest for further developments in the use of CFs in computational models of strategic control.

R2.2. Are there modulatory interactions in word perception?

The relevance of contextual modulation to word perception is called into question by **Niebur & Usher** on the grounds that the effects of context and target are independent until additively combined at the final stage of response selection (Massaro 1989a; sect. R7.1.1). There are four replies to these doubts. First, Movellan and McClelland (1995) show how the data suggesting independence is compatible with interactions between streams even in the case where those interactions are mediated by driving excitatory and inhibitory connections. Second, even Massaro acknowledges an asymmetry in the effects of context and target, because it is only the target that determines the alternatives between

which choice is made. Third, Samuel (1996) shows that there are conditions where the signal detection measures used as evidence for independence do show effects of context upon signal detectability, even though these effects are small and fragile. Fourth, and most importantly, measures of signal detection are not equivalent to measures of information transmission. Changes in bias can have large effects on information transmission, so other measures are needed to test for the interactions that we propose. The appropriate conditional mutual information measures are presented in Smyth et al. (1996). We do not know of any studies applying these measures to context effects in word perception, but possible paradigms for doing so were outlined in section 5.7.2.

Iacoboni summarizes data showing that interhemispheric priming effects are mediated by the corpus callosum, as most theories would predict. This could reflect interhemispheric contextual modulation, but more direct evidence is required. In particular, measures that distinguish modulatory from driving effects (e.g., Smyth et al. 1996) would have to be used in order to show that the interhemispheric priming is indeed modulatory. Until then, a *prima facie* case for the hypothesis can be made on the grounds that response selection is determined by the target, with cross-hemispheric context predominantly influencing the RT.

R.3. Does the cortex use synchronized population codes?

R3.1. The doubts of Amit and Treves

Both **Treves** and **Amit** doubt that synchronization is used. This question can only be resolved experimentally. Spurious synchronization in Amit's models in no way detracts from the many observations of stimulus-dependent synchronization in real animals. Treves's use of information measures may well help analyze this data, but it is not just the transmission of information in general that must be measured, but information about the particular stimulus relations that affect synchronization. Stimulation paradigms have to be used that require response selection and grouping, and evidence has to be obtained that the result of these operations is signaled by changes in synchronization rather than by changes in discharge rate. Such evidence is available. Whether neurons with spatially segregated RFs are activated with a single continuous contour or with two different stimuli is reflected by the degree of synchronization rather than by rate changes in the cat retina and LGN (Neuenschwander et al. 1996), cat area 17 (Engel et al. 1991; Freiwald et al. 1995; Gray et al. 1989), for groups of neurons distributed across area 17 of the two hemispheres (Engel et al. 1991; Munk et al. 1996), or across areas 17 and PMLS (Engel et al. 1991), in area MT of awake monkeys (Kreiter & Singer 1996), and in primary auditory cortex (deCharms & Merzenich 1996). In amblyopia, a developmental disorder associated with deficits in contextual grouping, the only detectable abnormality in primary visual cortex was drastically reduced synchronization among responses conveyed by the amblyopic eye. There were no differences in discharge rates between neurons driven by the normal or the amblyopic eye even when these responses were evoked by gratings whose spatial frequency was so high that they could no longer be resolved by the cats when

viewing with the amblyopic eye alone (Roelfsema et al. 1994a; 1994b). Finally, dynamic selection of responses, such as occurs during binocular rivalry, is associated with changes in synchrony but not discharge rate in area 17 of awake cats experiencing rivalry. Upon introduction of a rivalrous stimulus, neurons responding to the eye which won the competition increased their synchrony while neurons responding to the losing eye did the reverse (Fries et al. 1996).

R3.2. Should external (stimulus-locked) and internal synchronization be distinguished?

Stimulus-locked and internal synchronization are distinguished by **Tononi & Edelman, Eckhorn, and Wright**. We agree. In both cases synchronization is exploited for binding. Psychophysical evidence suggests that stimulus-locked synchronization is used for figure-ground segmentation if figure elements are presented with temporal offset relative to background elements. The temporal precision of this segmentation is better than 8 msec and the temporal information is provided by the magnocellular pathway (Leonards et al. 1996; Leonards & Singer 1997; but see Kiper et al. 1996 for a different result). This grouping by external timing can be overridden by internal grouping, however. If the figure is defined by textual coherence of its elements, detection is not impeded by false temporal conjunctions. If two different figures are defined, one by temporal cues and the other by texture cues, there is rivalry and the figure defined by the more salient cues is perceived.

R3.3. What are the time-scales of synchronization?

Questions concerning the relevant time scales of synchronization are raised by **Palm & Wennekers** and **Nunez**. One of the main advantages of synchronization is that it can operate on a much faster time scale than selecting responses by increasing their rate – provided that it is possible to synchronize responses rapidly. If responses are selected by rate increases, a cell has to integrate over a sufficiently large number of incoming EPSPs before emitting a response in order to assume that the response is actually generated by the selected input and not by accidentally arriving EPSPs or nonselected inputs. The required duration of this safety interval depends on average firing level and can be set by modulating the membrane potential (i.e., the distance to threshold and the membrane time constants). Assuming response frequencies of cortical neurons in the range of 50 Hz and postulating summation over at least 4–5 EPSPs before reaching a decision amounts to transmission times per processing step of at least 80 to 100 msec. This is far too slow. Input selection by spike synchronization is much faster because it does not rely on temporal summation. Coincident EPSPs can reach firing level within a few msec, thus reducing transmission times to near the synaptic delay. The speed of selection is thus primarily constrained by the time required to obtain synchrony. Contrary to earlier models, which generated synchrony by coupling harmonic oscillators, more recent models based on spiking neurons (i.e., relaxation “oscillators”) indicate that synchrony can be obtained very rapidly (**Wang; Somers & Kopell** 1993). This is supported by experimental evidence. In the retina, responses to coherent stimuli are

synchronized within milliseconds after their onset (Neuenschwander et al. 1996).

In the cortex, it is much more difficult to determine experimentally the time required for the alignment of spikes by cross-correlation analysis because the windows must have a minimal duration to obtain sufficient entries, and because response onset is not well defined when moving stimuli are used. The data suggest that newly induced oscillatory responses synchronize right away once oscillatory patterning develops (Singer & Gray 1995). As cortical activity has a complex oscillatory structure (Arieli et al. 1996) and as oscillatory modulation of the membrane potential shifts spike timing, feature specific synchronization may occur very rapidly. The rate at which different assemblies can be organized successively without merging depends critically on the integration time constants of cortical neurons. Here views diverge, ranging from estimate of less than a millisecond (König et al. 1995; Softky 1995) to tens of milliseconds (Shadlen & Newsome 1994). Data on echo location and spatial location of sound sources show that neurons are capable of very precise coincidence detection (Yan & Suga 1996). Whether this is also true in the cortex remains to be seen, but recent evidence on action potentials backpropagating into dendrites (Stuart & Sakmann 1994) and other regenerative events in dendrites (Connors & Gutnick 1990) may lead to a drastic change in views on the timescale of dendritic integration.

R3.4. What is the role of oscillation?

Questions are raised by **McCollum**, and **Haase & Diniz** concerning the relation between oscillation and synchronization. Experimental observations indicate a close relation between oscillation and the precision of synchrony (König et al. 1995). Also, there is evidence that cortical neurons engage in synchronous oscillations in the beta and gamma frequency range across visual, association, somatosensory, and motor areas when the animal focuses its attention in the preparation of a visuomotor reaction; this synchronization increases in a task-specific way during execution of the task but collapses upon completion (Roelfsema et al. 1997). This suggests that oscillatory modulation helps synchronize responses, facilitating handshaking among neurons that need to be bound together for the execution of the task. Whether this frame-setting oscillatory modulation of membrane potential is due to oscillatory input from inhibitory or excitatory sources is not crucial for the principle and requires experimental testing.

Oscillatory patterning may be advantageous for rapid synchronization of discharges because it allows for temporal shifting of responses to synaptic input in both directions. The duration of the possible shift intervals is determined by oscillation frequency, and in the range of about a half-cycle. An assembly defined by synchronized discharges need not necessarily extend over many oscillation cycles. In some cases it may be sufficient to organize a single volley of synchronized discharges and have a new assembly at the next processing stage as proposed by Abeles (1991). In that case different assemblies can be organized on successive cycles. However, oscillatory patterning can also serve to generate the same synchronous volleys over several cycles, thereby refreshing assemblies that need to be stabilized for a while. The advantage over stabilizing assemblies by simple rate increases would be that it reduces the possibility

of getting false conjunctions if several assemblies need to be maintained within the same array of neurons, as false bindings are avoided if unrelated assemblies do not fire in synchrony. This does not require fixed phase shifts, as it can be achieved by interactions that modulate the regularity of the respective oscillations.

R4. What forms of learning occur within the cortex?

R4.1. Can the abstract learning rules discover higher-order variables?

There is a recognition by **Stone** of the potential of using context to guide learning to variables that are statistically related to variables of significance to the organism, including, but not limited to, reinforcement. We agree that this can include temporal as well as spatial structure. The algorithms that use temporal constraints to discover statistical structure (e.g., Becker 1996; de Sa & Ballard 1997; Stone 1996a; 1996b) are sufficiently similar to those outlined in section 3 that ways in which the two constraints can be used together merit further study.

Floreano shows that the learning rules outlined in section 3.3 can discover nonlinear transforms. The example used is stereo depth. He studies the possibility that this is guided by contextual input from other cues to depth. As there are usually multiple cues it may be that those that are either genetically specified or more easily learned guide the discovery of others. Floreano assumes that different depth cues are processed in different streams, and this is also plausible because different cues are not always in agreement (e.g., depth of a picture plane versus depth in the picture). The success of Floreano's simulation in sometimes discovering stereo depth using such an internal teacher is therefore encouraging. It did not do so reliably, however, and although his suggestion as to why that was so may be part of the story, the capabilities and limitations of this whole class of learning algorithms remain unclear.

R4.2. Does unsupervised learning in the cortex discover higher-order variables?

In section 6.4 we asked whether there is any evidence that self-organization in the cortex can discover nonlinear variables such as XOR. No such evidence was offered in the commentaries, nor have we yet found any from other sources. The continued failure of such evidence to appear suggests that reliable discovery of such nonlinear variables may not be a fundamental capability of cortex.

R4.3. Is there a floating threshold for LTP, and if so what are its determinants?

Sections 3.3 and 6.4 note converging evidence for learning rules of the BCM/ABS type (Artola & Singer 1993; Bienenstock et al. 1982; Hancock et al. 1991a). A central feature of these learning rules is a non-monotonic dependence of synaptic change upon post-synaptic activity with a threshold (that may be movable) below which synapses are weakened (LTD) and above which they are strengthened (LTP). **Grossberg** notes that some of his learning rules have a similar form, and **Stone** also notes their importance. A major goal for the neurobiology of learning is therefore to

find out whether the threshold for LTP does move and if so what it depends upon.²

R4.4. Does SMA modulate sensorimotor learning in PMdc?

Neuroimaging evidence is outlined by **Iacoboni** showing that activity in the supplementary motor area (SMA) provides a contextual input that guides the learning of a sensorimotor mapping from the rostral sector of dorsal premotor cortex (PMdr) to the caudal sector of dorsal premotor cortex (PMdc). We agree that this may be so but other explanations of the results are possible. One way forward may be to show that activity in both PMdr and PMdc is necessary for learning, but that activity in SMA is not, even though it has a guiding effect when present. Experiments of this sort are now under way (Iacoboni, personal communication).

R5. To what extent do these processes provide a representation of the external world?

One of our central suggestions is that local processors could lay foundations for representation and meaning by discovering commonalities across diverse data sets. This is because CFs provide a means whereby local processors can select the relevant information. **Morasso et al.**, **Stone**, **Tononi & Edelman**, and **Wright** all supported this suggestion and no one argued against it. **Amit**, **Eckhorn**, and **McCollum** all overlooked this suggestion, however, calling for a notion of relevance as though we had not addressed the issue. We must thus reiterate our suggestion, which is essentially a simple one. If local processors receive only RF input then we can see no way in which they could distinguish the relevant from the irrelevant. If they also receive CF input, then they can do so by, for example, selectively transmitting only the RF information that is statistically related to the context. Thus, in clear contrast to what **McCollum** assumes, this does involve information loss, and crucially so.

In section 6.5 we asked whether the processors we propose could help form proper intentional representations of the external world. We suggested that while they might contribute they could not be sufficient. We have so far not been very precise as to what “external input” means. In relation to the simulations it simply means external to the network being simulated. For most local processors in the cortex this is analogous to input from another cortical subsystem, and not to input from a world external to the system as a whole. To make this distinction the system must have some notion of “itself.” **Wright’s** insightful suggestion that the ability to discover coherence in multiple sources of activity could apply to sources from within the organism may be relevant here, but philosophical clarification of the issues involved would also be helpful.

R6. Additional computational foundations suggested by commentators

We proposed functional specialization and contextual coordination as two fundamental and mutually constraining foundations of cortical computation, expecting thereby to

provoke proposals of additional or alternative foundations. Few others were forthcoming. **Grossberg** refers us to his neural network models. **Tononi & Edelman** suggest a perspective from which the predominant information flow is between sources internal to the cortex, with the main role of external input being to modulate that flow. We agree that the distinction between internal and external sources is important, and that the balance between them is a topic of fundamental importance (sect. R2.1). It is clear that at the level of local circuits **Tononi & Edelman** emphasize just the same two fundamental principles as we do.

In keeping with an emphasis upon internal sources, **Wright** notes that local processors may be organized so that they are autonomously active and with an intrinsic dynamic. In this case they cannot be adequately described only in terms of a transfer function. We agree that in this case such sources should be seen as being included within the total set of primary RF inputs, and that developments of the kind he suggests are well worthwhile. **Haase & Diniz** and **Stone** note that structure in time is as important as structure in space, so an ability to process temporal sequence is also crucial. This seems compatible with our emphasis upon both coherence and synchrony. Indeed, as **Haase & Diniz** note, precision in temporal sequencing may be well served by the use of synchronized population codes (**Abeles** 1991). Overall, functional specialization and contextual coordination still appear most prominent as fundamental principles of cortical computation.

R7. Further issues raised by commentators

R7.1. What mathematical formalisms are useful?

We take it for granted that if there are common foundations then it must be possible to describe them in abstract terms that are independent of any specific content.

R7.1.1. Is information theory useful? We can ask, echoing **Amit’s** questions concerning synchronization “Is information theory necessary to an understanding of cortical function?” and “Is it sufficient?” To both we answer with an unequivocal “No.” **McCollum** identified our methods with information theory and neural network simulations. Given our backgrounds in experimental psychology and neurophysiology, and the allocation of only 20% of the target article to computational studies we do not see that as being accurate. Nevertheless, it is justified to ask “Is information theory useful to an understanding of cortical function?” and “Are other formalisms useful?” To both we answer with an unequivocal “Yes.” Of the commentators, **Stone**, **Tononi & Edelman**, **Treves**, and **Wright** also find information theory useful, but **Gregson**, **Eckhorn**, **Nunez**, and **McCollum** do not. These differences may be due to a misunderstanding of the uses proposed for information theory, so we will try again to make them clear.

Three uses are exemplified in the target article and commentaries. First, there is its use for job specification, that is, to provide a conception of the essential information processing operations performed by local cortical circuits. In clear contrast to **Grossberg’s** perception of the target article, its central concern is with general principles. In using information theory to formulate a conception of the goals of cortical computation we are in a long (e.g., **Attneave** 1954; **Barlow** 1959) and still strong (e.g., **Intrator & Cooper**

1995a; Olshausen & Field 1996) tradition. The goals of information compression, data reduction, and sparsification are relevant to the concerns rightly raised by **Gregson** and by **Eckhorn** in relation to metabolic constraints. In accord with the work of **Stone** and of **Tononi & Edelman**, a central aim of the target article is to extend these conceptions to include the challenging task of information *selection* (sects. 1.2.1, R5).

Second, and also in accord with the work of **Stone**, we have used the information theoretic job specification to derive a learning algorithm that changes connection strengths to better approximate the goal, given the statistical structure within the input received. Our description of the goal as maximizing the objective function specified (sect. 3.2) may mislead some commentators, for example, **Amit**. There is no reason to suppose that in any realistic situation perfect optimization is either necessary or possible. All that is required for the learning to be useful is that it more closely approximates the system to such a state. Although it may not be necessary to use information theory to understand learning, should synaptic plasticity be found to have a dynamic threshold for LTP that moves as specified by one of the abstractly derived learning rules (sect. R4.3), then many will wish to understand better how that rule was derived.

Third, information theory can be used for data analysis, for example, to measure information transmitted, as emphasized by **Treves**. We agree that such measures are useful (sect. R3.1). In addition, we use information theory to tackle a problem that no commentator discussed, that is, to provide an adequate specification of what should count as “context.” There was agreement that context is important, but this means little unless we can give an adequate account of what can count as context.

Consider a psychophysical experiment in which there are two cues to texture segregation and that both influence performance, although subjects are asked to base their judgments on just one of them. We could call the specified cue the target, and the other the “context,” thereby hoping to provide evidence for the relevance of context. Though common, this conception of context is wholly inadequate. On that conception “target” and “context” are distinguished simply by what experimenter and subject call them. It may be that both cues contribute in essentially the same way to the observed responses, so that dividing them into target and context is arbitrary.

This is exactly what is implied by Massaro’s (1989a) fuzzy logic model of perception. Perceptual channels are independent until they all contribute in essentially the same way to response decision. In such a system there is no need to distinguish target from context. Either an input variable contributes to response or it does not. The same can also apply to the cellular level; either an input effects output or it does not. If it does, it is part of the cell’s RF, and if not, it isn’t. From this viewpoint there is no need to distinguish RF influences from CF influences (Gilbert, personal communication). Thus, in contrast to what **Amit** assumes, the case for “context” does have to be made. Our use of information theory to specify what can count as context is developed further in Smyth et al. (1996), and is being used at Stirling to analyze the interaction of cues to texture segregation by human subjects (thus addressing **Gregson**’s doubts about the empirical validity of the transfer function used in the

computational theory). We hope to put the study of contextual coordination on a sounder basis than is provided either by showing that variables other than those that the experimenter calls the target can affect subjects’ responses, or by showing that cellular activity is affected by more than what the experimenter calls the classical receptive field.

R7.1.2. What other formalisms are useful? We do not expect to squeeze all of cortical function into a single mathematical formalism. We agree with **Gregson**, **Nunez**, **McCollum**, and **Morasso et al.** that others are needed, and non-linear dynamic systems analysis in particular. They complement rather than supplant the uses just listed for information theory, however. The challenge for such formalisms is to give rise to revealing experimental paradigms and to make their findings more comprehensible than they would otherwise have been. Having tried to do that for what are perhaps some of the more basic and intrinsically comprehensible aspects of cortical function, we know how big a challenge that is.

R7.2. What is the role of modeling?

None of the uses listed in section R7.1.1 for information theory can be described as “modeling.” Simulations are outlined in section 3.4 but they are not models of biological systems. As **Bower** indicates, they are “proof of concept” simulations, being designed to explore the abstract goals, transfer functions, and learning rules. **Würtz** also understood this, and thought our focus on simple computational concepts a distinct advantage. Nevertheless, **Palm & Wennekers**’s attempt to interpret the simulations as models of spiking neurons, and then criticize them for not being so. Their attempt to interpret the simple demonstration whose results are shown in Figure 10 of the target article shows that they misunderstood it in several ways.³ Palm & Wennekers ask questions that are not pertinent to any of the uses made of the computational theory, but they could be asked of a biological model, or more pertinently still of the biological system itself (sect. R3.3).

Bower continues to argue for the relevance of detailed biological models. We agree that they are useful, but their limitations as a primary focus for research are well revealed by Bower’s commentary. Necessary responses to the question posed in his title are “What details?” and “To what is their relevance being assessed?” If they matter, they are not details. If we were trying to discover how birds stay up in the air it would be counterproductive to try to include all details of bird structure and physiology. The “details” that matter are those essential to aerodynamic lift. Bower implies that the details that matter in the case of cortical computation are those in his models. **Grossberg** also describes his models as being biologically realistic, but they are not the same as Bower’s and are used to support different conclusions. The most convincing way to test the relevance of a theory to biology is to test it on real organisms, not on models that try to replicate them. We therefore direct our efforts towards seeking a conceptual understanding of cortical computation that can be tested and developed through investigations of real brains and real people, as exemplified here by **König et al.** and by **Silverstein & Schenkel**. As the role of computational theory is still so widely and deeply misunderstood Table R2 compares it

Table R2. *Relations between theory and biology*

	Bird flight	Cortical computation
Goal ¹	Stay up	Coordinate activity ²
Strategy	Use aerodynamic lift	Use contextual information
Formalization	Aerodynamic equations for lift, drag, moment, etc.	Transfer functions and information theoretic objective functions, etc.
Mechanism	Wings ³	Contextual fields ³
Testing the strategy	Wings in wind tunnels ^{4,5}	Neural net "simulations" ^{4,5}
Biological testing	Do bird's wings have the required aerofoil section? Is air velocity adequate?	Do cortical cells receive contextual input? Does it affect output as required?

1. Only one goal is considered in each case. Both bird flight and cortical computation have other goals but separate goals are best considered separately.

2. The goals of cortical computation are far less obvious than those of flight. Discovering what they are is part of the problem we have to solve.

3. Wings and CFs are neither necessary nor sufficient for their goals.

4. The theories are best tested using versions of the mechanism that are as simple as possible yet consistent with achieving the goal in accordance with the theory.

5. The mechanisms tested are real. The wing in the wind tunnel generates lift. The "simulated" net computes. Thinking of them as models of something else does not contribute anything to their role in testing the strategy.

with the role of aerodynamics in the study of flight, in the hope that that will clarify our views on this issue.

R7.3. What is the role of CFs and/or synchronization in motor control?

The possible relevance of CFs and synchronization to motor control is discussed by **Iacoboni, McCollum, and Morasso et al.** The selection of particular motor responses raises combinatorial problems similar to those raised by scene segmentation. The number of different constellations of muscle contractions needed for the execution of different movements is very large. As the same muscles are used in different combinations for different movements, the commands for their contraction have to be bound in ever different constellations. Representing every possible movement by a command unit that distributes its output to the appropriate constellation of motor neurons leads to the same combinatorial explosion as representing every distinguishable object by a pontifical cell. The analogy with scene segmentation and object representation is obvious if one equates motor units with elementary features and a motor program with an object. This predicts that one should observe oscillations and synchronization during the initial selection phases of motor programming and a sequence of successively structured synchronized assemblies of the type ABCD etc. during execution. The findings of Murthy and Fetz (1996) are compatible with such a view.

R7.4. Is contextual coordination deficient in schizophrenia?

Contextual coordination, **Silverstein & Schenkel** suggest, may be impaired in particular schizophrenic states. *Prima facie*, this has an intuitive appeal. If there are specialized mechanisms for contextual coordination, fragmentation will result from their impairment. If these mechanisms are widely distributed across cortex, disorganized perceptions, thoughts, and actions are all possible consequences. Many studies of cognition in schizophrenia can be seen as evidence for such disorganization. Silverstein & Schenkel's suggestion is further strengthened by evidence for under-activity of NMDA-receptor channels in schizophrenic states (Olney & Farber 1995) and for the psychotomimetic effects of ketamine (Krystal et al. 1994), which blocks NMDA-channels. This is in keeping with our hypothesis that contextual coordination is achieved via CFs that exert gain control, for example, via NMDA-channels. It thus becomes of paramount importance to understand the basic pharmacology of coordinating gain-control channels, because that may advance our understanding of schizophrenia in a way that directly links molecular mechanisms with their cognitive consequences.

Attempts to understand schizophrenia have been plagued by finding a task on which patients are impaired, and then building a grand explanatory model on that (Ronan O'Carroll, personal communication). A particular problem that arises here is that the "contextual coordination" that is impaired in schizophrenia may have little or nothing to do with the "contextual coordination" for which the target article presents evidence. We therefore first consider studies of perceptual grouping in schizophrenia, then note unresolved issues that arise.

Silverstein et al. (1996a) review many studies of perceptual grouping in schizophrenia beginning with those of Cox and Leventhal (1978) and Place and Gilmore (1980), and present further evidence using the Banks and Prinzmetal (1976) paradigm to study the effect of visual grouping on visual search. All these studies suggest that in severe schizophrenic states perception is more fragmented than normal, with reduced effects of processes that group the stimulus elements into larger units. The evidence suggests that both stimulus-induced and top-down attentional processes may be involved in this impairment. The impairment is indicative of schizophrenia (Knight 1992; Silverstein et al. 1992), and correlates with the severity of other psychotic symptoms (Silverstein et al. 1996c). Using Navon's (1977) terminology, these patients have a tendency to see the trees before the forest, whereas control subjects have a tendency to see the forest before the trees. The evidence cannot be dismissed as being due to a general impairment of performance for at least two reasons: (1) The altered processes involve a reversal in the relative difficulty of conditions as compared with controls (Silverstein et al. 1996a); (2) This patient group performs *better* in tasks where it is advantageous to see the trees, not the forest (Place & Gilmore 1980; Rabinowicz et al. 1996). These perceptual disorders thus fit well with those to be expected, given impairments to the processes of contextual coordination hypothesized in the target article, however several unresolved issues then arise.

1. This patient group shows no deficits in grouping under conditions where there are either strong stimulus cues to

grouping (Knight 1992) or a strong push to allocate attention to global properties (Silverstein et al. 1996a). The hypothesis must therefore be expressed in terms of a reduced *tendency* to use grouping, and this needs further conceptual clarification and empirical testing. Gestalt grouping processes may not be at all impaired, either because they are genetically specified, or because if they are acquired or refined by visual input then the input correlations upon which they depend are so overwhelmingly present in natural input that they are learned to asymptotic levels even by an impaired learning mechanism. The acquisition of new assemblies at higher levels is not based upon such strong statistical structure, however, so impairments in acquiring and using those may be associated with impaired binding and reduced synchronization both within and between regions.

2. As **Silverstein & Schenkel** note, the evidence suggests impairments of perceptual learning as well as of perceptual processing, and this is in keeping with other perspectives on this disorder (e.g., Gray et al. 1991). An unresolved issue that arises here is whether it is only CF learning (i.e., assembly formation) that is impaired and, if so, how this can be reconciled with the hypothesis that CFs guide RF learning.

3. Some of the top-down attentional effects on perceptual grouping that have been shown to be impaired in schizophrenia clearly involve episodic and working memory processes (Silverstein et al. 1996b). The framework as presented in the target article would therefore have to be extended to show how it is related to higher level processes such as attention and episodic and working memory, if it is to be applicable to such findings. The value of such an extension is also emphasized by **Bugmann, König et al., and Tononi & Edelman** (sect. R3).

4. It is easy to see how an impairment in contextual coordination could produce symptoms of disorganization, but it may at first seem less easy to see how it could account for the positive symptoms, such as hearing voices. One approach would be to seek a separate explanation for these positive symptoms. Alternatively, we could speculate that they may reflect: (1) reduced modulatory control by external input of the flow of activity arising from internal thoughts; and (2) reduced knowledge of the source of internal speech and percepts because of reduced inter-regional binding.⁴ The pharmacological evidence encourages the search for an integrated explanation because drug-induced psychoses that reduce NMDA activity mimic both positive and negative symptoms (Krystal et al. 1994).

Silverstein & Schenkel's hypothesis thus survives close scrutiny well and, as they argue, it opens important directions for further work. One possibility may be to study the effect of drug-induced psychoses on perceptual learning and inter-regional synchrony using local field potentials. More broadly, this hypothesis should encourage physiological, psychological, and neuroimaging studies of contextual coordination in schizophrenic and drug-induced psychotic states using both human and animal subjects.

R8. Science as a search for coherence

It is suggested by **Treves** that if the commentaries "turn out not to be fully synchronized" (by which we assume he means not in full agreement) then they will contribute to the population response but by means other than synchronizing

it. We agree that it may be both revealing and entertaining to draw an analogy between the outputs of local cortical processors and scientific writings. Each paper written tries to convey something that is both distinctive and relevant. Relevance does not imply agreement, however. Different members of a group can have opposing effects at a higher level of analysis. The primary code used by *BBS* to signal mutual relevance is simultaneous publication of target articles and associated commentaries, that is, grouping through synchronous transmission. This is not because synchrony is necessary to grouping, but because it makes it a lot easier for readers to process them as a whole.

This analogy is also useful for another reason. By seeing science as a continuation of the search for coherence, by means beyond the reach of any algorithm possessed by mammals in general, we are reminded that local processors should be seen as moving toward that goal, not as reaching it (sect. R7.1.1).

R9. Promising directions for future progress

Specialization of function has been the keynote of this century. Coordination of function may be the keynote of the next. This will require more attention to contextual interactions, synchronized population codes, and fast dynamics. We believe that future research must emphasize relational codes rather than focusing only upon response properties of single cells. The relevant relations can be detected only by analyzing simultaneously recorded activity. We shall have to look at the outputs of any processing level in the same comprehensive way as it is looked at by the subsequent processing stages to which it projects. This requires analysis of mutual dependencies among the responses of a large number of neurons, and hence the study of ensemble dynamics. There is thus a logical progression from early studies of the topology of brain functions, to the functional analysis of individual neurons located in different compartments, and now to the attempt to understand the spatio-temporal patterns of coordinated neuronal activity. In conclusion, then, our prediction is that the temporal organization and coordination of brain activity is as sophisticated and subtle as is its topological organization, the former being an emergent property of the latter.

ACKNOWLEDGMENTS

We thank Peter Hancock, Jim Kay, Steve Silverstein, Leslie Smith, and Darragh Smyth for comments on a draft of this response.

NOTES

1. Indeed, if NMDA receptor activation is combined with oscillatory modulation of the membrane potential, latency of responses is shifted as a function of oscillation phase and output phase is locked to input phase (Volgushev et al., submitted). This can synchronize discharges within a single oscillation cycle with high temporal precision because of the voltage-dependent gating of the MG^{2+} -block, which allows for very rapid activation and inactivation of NMDA-receptor mediated input.

2. Evidence is becoming available that it is not the absolute level of postsynaptic depolarization or the absolute level of the postsynaptic Ca^{2+} -surge alone that determines whether a particular synapse undergoes LTP or LTD. The actual state of the synapse, that is, the actual release probability with which it operates (Volgushev et al. 1997) and the rate with which it is activated during induction of the modification also matter (Hansel et al., in preparation). Thus synaptic changes are affected by information stored locally at each synapse, and this makes complex

learning rules such as those of Phillips et al. (1995b) and Kay et al. (in press) less implausible.

3. The simulations assume four levels of organization: cells, units, local processors, and layers. They operate at the level of units, but assume them to be composed of stochastic binary elements. The simulation shown in Figure 10 is intended to be analogous to the perception of an ambiguous figure, such as the Rubin vase, where alternative perceptions emerge successively, even though the external stimulus does not change. Thus, in contrast to what Palm & Wennekers assumed, the input cannot represent direct input from the external stimulus. Instead, the processing layer to which the simulation relates is assumed to be some layers distant from the external stimulus. Its inputs from the previous layer fluctuate because processing is stochastic. Palm & Wennekers suggest that the simulation in Figure 10 would work in the same way if the graded response units were replaced by single binary units. This is not so. If outputs are restricted to their extremal values, the CFs have no effect, in which case the outputs in at each iteration in Figure 10 would simply be identical to the inputs. We should therefore have made in clear in section 3.1 that the apparent insensitivity of performance to such a replacement refers only to the learning, with which most of the simulations were concerned. The two sets of 3×3 units with internal coherence in the study of grouping shown in Figure 10 represent the two alternative percepts. CFs were therefore positive within and negative between these two alternative groupings. The background elements received no CF input. The first six iterations show the simultaneous emergence from the background of all elements in one of the two alternative groupings. Thus, contrary to what Palm & Wennekers conclude, this demonstrates the role of internal grouping processes, because the coherent sets of elements are distinguished from the background only by their CFs. To show the successive emergence of alternative percepts activity in the units was stopped and restarted after six simulations. The cortex is not short of mechanisms that would have such an effect, so there is no need for either it or a simulation to become stuck in the stationary state to which Palm & Wennekers refer. The choice of six iterations was arbitrary, but by varying it in an appropriate way periodic or other temporal structure could easily be introduced to the simulation. That would not change the effects with which it was concerned, however.

4. The latter possibility was suggested to us by Frith's (1992) interpretation of psychotic hallucinations as being due to a failure of self-monitoring. Words activated within the articulatory-loop can arise from either internal or external sources. Their source may therefore be unknown if they are not bound to activity in the other cortical modules from which they arise.

References

Letters "a" and "r" appearing before authors' initials refer to target article and response, respectively.

- Abeles, M. (1982) *Local cortical circuits: Studies of brain function*, vol 6. Springer-Verlag. [GP, aWAP, DLW]
 (1991) *Corticoids: Neural circuits of cerebral cortex*. Cambridge University Press. [arWAP]
- Abeles, M., Bergman, H., Margalit, E. & Vaadia, E. (1993b) Spatiotemporal firing patterns in the frontal cortex of behaving monkeys. *Journal of Neurophysiology* 70:1629–38. [aWAP]
- Abeles, M., Vaadia, E., Bergman, H., Prut, Y., Haalman, I. & Slovlin, H. (1993a) Dynamics of neuronal interactions in the frontal cortex of behaving monkeys. *Concepts in Neuroscience* 4:131–58. [aWAP]
- Able, K. P. & Bingham, V. P. (1987) The development of orientation and navigation behavior in birds. *Quarterly Review of Biology* 62:1–29. [aWAP]
- Aertsen, A. M. H. J., Gerstein, G. L., Habib, M. K. & Palm, G. (1989) Dynamics of neuronal firing correlation: Modulation of "effective connectivity." *Journal of Neurophysiology*, vol. 61, 5:900–17. [GP]
- Allman, J., Miezen, F. & McGuinness, E. (1985) Stimulus specific responses from beyond the classical receptive field: Neurophysiological mechanisms for local-global comparisons in visual neurons. *Annual Review of Neuroscience* 8:407–30. [aWAP]
- Amari, S. (1983) Field theory of self-organizing neural nets. *IEEE Transactions on Systems, Man, and Cybernetics* SMC-13:741–48. [PM]
- Amit, D. J. (1989) *Modelling brain function: The world of attractor neural networks*. Cambridge University Press. [RAMG, arWAP]
 (1995) The Hebbian paradigm reintegrated: Local reverberations as internal representations. *Behavioral and Brain Sciences* 18:617–57. [JJW]
- Amit, D. J. & Brunel, N. (submitted) Dynamics of a recurrent network of spiking neurons before and following learning. (See e.g. www.fiz.huji.ac.il/staff/acc/faculty/damita.) [DJA]
- Andrew, C. & Pfurtscheller, G. (1996) Event-related coherence as a tool for studying dynamic interaction of brain regions. *Electroencephalography and Clinical Neurophysiology* 98:144–48. [PLN]
- Arieli, A., Sterkin, A., Grinvald, A. & Aertsen, A. (1996) Dynamics of ongoing activity: Explanation of the large variability in evoked cortical responses. *Science* 273:1868–71. [rWAP]
- Armstrong-James, M., Welker, E. & Callahan, C. A. (1993) The contribution of NMDA and non-NMDA receptors to fast and slow transmission of sensory information in the rat SI barrel cortex. *Journal of Neuroscience* 16:480–87. [aWAP]
- Artola, A., Brocher, S. & Singer, W. (1990) Different voltage-dependent thresholds for the induction of long-term depression and long-term potentiation in slices of the rat visual cortex. *Nature* 347:69–72. [aWAP]
- Artola, A. & Singer, W. (1993) Long-term depression of excitatory synaptic transmission and its relationship to long-term potentiation. *Trends in the Neurosciences* 15:218–26. [SG, arWAP]
- Ascher, P., Bregestrovski, P. & Nowak, L. (1988) N-methyl-D-aspartate activated channels of mouse central neurones in magnesium-free solutions. *Journal of Physiology* 339:207–26. [aWAP]
- Atick, J. J. (1992) Could information theory provide an ecological theory of sensory processing? *Network* 3:213–51. [aWAP]
- Atick, J. J. & Redlich, A. N. (1990) Predicting ganglion and simple cell receptive field organizations. *International Journal of Neural Systems* 1:305. [aWAP]
 (1993) Convergent algorithm for sensory receptive field development. *Neural Computation* 5:45–60. [aWAP]
- Attneave, F. (1954) Informational aspects of visual perception. *Psychological Review* 61:183–93. [arWAP]
- Baddeley, R. J. & Hancock, P. J. B. (1991) A statistical analysis of natural images matches psychophysically derived orientation tuning curves. *Proceedings of the Royal Society of London B* 246:219–23. [aWAP]
- Bair, W., Koch, C., Newsome, W., Britten, K. & Niebur, E. (1994) Power spectrum analysis of bursting cells in area MT in the behaving monkey. *Journal of Neuroscience* 14:870–92. [aWAP]
- Baldi, B. & Hornik, K. (1989) Neural networks and principal component analysis: Learning from examples without local minima. *Neural Networks* 2:5358. [aWAP]
- Banks, W. P. & Prinzmetal, W. (1976) Configural effects in visual information processing. *Perception & Psychophysics* 19:361–67. [rWAP]
- Barlow, H. (1996) Intraneuronal information-processing, directional selectivity and memory for spatiotemporal sequences. *Network: Computation in neural systems* 7(2):251–59. [JVS]
- Barlow, H. B. (1959) Sensory mechanisms, the reduction of redundancy, and intelligence. In: *The mechanisation of thought processes*. Her Majesty's Stationery Office. [arWAP]
 (1961) Stationary principles underlying the transformations of sensory messages. In *Sensory communication*, ed. W. A. Rosenblith. MIT Press. [aWAP]
 (1972) Single units and sensation: A neuron doctrine for perceptual psychology? *Perception* 1:371–94. [aWAP]
 (1989) Unsupervised learning. *Neural Computation* 1:295–311. [aWAP]
 (1993) The biological role of neocortex. In: *Information processing in the cortex*, ed. A. Aertsen & V. Braitenberg. Springer-Verlag. [aWAP]
- Barlow, H. B. & Foldiak, P. (1989) Adaptation and decorrelation in the cortex. In: *The computing neuron*, ed. R. Durbin, C. Miall & G. Mitchison. Addison-Wesley. [RE, aWAP]
- Becker, S. (1993) Learning to categorize objects using temporal coherence. In: *Advances in neural information processing systems 5*. Morgan Kaufmann. [aWAP]
 (1996) Mutual information maximization: Models of cortical self-organization. *Network* 7:7–31. [arWAP, JVS]
- Becker, S. & Hinton, G. E. (1992). A self-organizing neural network that discovers surfaces in random-dot stereograms. *Nature* 355:161–63. [DF, aWAP]
- Bentin, S. (1992) Phonological awareness, reading, and reading acquisition: A survey and appraisal of current research. In: *Orthography, phonology, morphology, and meaning*, ed. R. Frost & L. Katz. Elsevier. [aWAP]
- Bentin, S., Hammer, R. & Cahan, S. (1991) The effects of aging and first grade

- schooling on the development of phonological awareness. *Psychological Science* 2:271–74. [aWAP]
- Bernander, O., Douglas, R. J., Martin, K. A. C. & Koch, C. (1991) Synaptic background activity influences spatiotemporal integration in single pyramidal cells. *Proceedings of the National Academy of Sciences USA* 88:11569–73. [aWAP]
- Bernander, O., Koch, C. & Usher, M. (1994) The effect of synchronized inputs at the single neuron level. *Neural Computation* 6:622–41. [DJA]
- Bernstein, N. (1967) *The co-ordination and regulation of movements*. Pergamon Press. [GM]
- Bezrukov, B. M. & Vodyanov, I. (1997) Stochastic resonance in non-dynamical systems without response thresholds. *Nature* 385:319–12. [rWAP]
- Bialek, W., Rieke, F., Vansteveninck, R. R. D. & Warland, D. (1991) Reading a neural code. *Science* 252:1854–57. [JMB]
- Bialek, W. & Rieke, F. (1992) Reliability of information transmission in spiking neurons. *Trends in Neurosciences* 15: 428–34. [VGH]
- Biederman, I. (1972) Perceiving real-world scenes. *Science* 177:77–80. [aWAP]
- (1987) Recognition-by-component: A theory of human image understanding. *Psychological Review* 94:115–47. [DLW]
- Bienenstock, E. (1995) A model of neocortex. *Network* 6:179–224. [aWAP]
- Bienenstock, E., Cooper L. N. & Munro P. W. (1982) Theory for the development of neuron selectivity: Orientation specificity and binocular interaction in visual cortex. *Journal of Neuroscience* 2:32–48. [arWAP]
- Blakemore, C. & Tobin, E. A. (1972) Lateral inhibition between orientation detectors in the cat's visual cortex. *Experimental Brain Research* 15:439–40. [aWAP]
- Bower, J. M. (1995) Reverse engineering the nervous system: An in vivo, in vitro, and in computo approach to understanding the mammalian olfactory system. In: *An introduction to neural and electronic networks, second edition*, ed. S. Zornetzer, J. Davis and C. Lau. Academic Press. [JMB]
- Braddick, O. (1993) Segmentation versus integration in visual motion processing. *Trends in the Neurosciences* 16:263–68. [rWAP]
- Braitenberg, V. (1978) Cortical architectonics: General and areal. In: *Architectonics of the cerebral cortex*, ed. M. A. B. Brazier & H. Petsch. Raven. [aWAP]
- Braitenberg, V. & Schüz, A. (1991) *Anatomy of the cortex*. Springer-Verlag. [aWAP]
- Bressler, S. L. (1995) Large scale cortical networks and cognition. *Brain Research Reviews* 20:288–304. [PLN]
- Brousse, O. & Smolensky, P. (1989) Virtual memories and massive generalization in connectionist combinatorial learning. In: *Proceedings of the 11th Annual Conference of the Cognitive Science Society*. Erlbaum. [aWAP]
- Bruce, V. (1988) *Recognising faces*. Erlbaum. [aWAP]
- Bruner, J. S. (1972) *Beyond the information given*. Norton. [GT]
- Bugmann, G. & Taylor, J. G. (1993) A model for latencies in the visual system. In: *Proceedings of the third conference on artificial neural networks (ICANN'93, Amsterdam)*, ed. S. Gielen & B. Kappen. Springer-Verlag. [GB]
- (1994a) Role of short-term memory in neural information propagation. In: *Extended abstract book of the International Symposium on Dynamics of Neural Processing*, Washington. [GB]
- (1994b) *A top-down model for neuronal synchronization*. Research Report NRG-94-02, School of Computing, University of Plymouth. [GB]
- Bullier, J., Munk, M. H. J. & Nowak, L. G. (1992) Synchronization of neuronal firing in areas V1 and V2 of the monkey. *Society of Neurosciences Abstracts* 18:11.7. [aWAP]
- Bullock, T. H., McClune, M. C., Achimowicz, J. Z., Iragui-Madoz, V. J., Duckrow, R. B. & Spencer, S. S. (1995) EEG coherence has structure in the millimeter domain: Subdural and hippocampal recordings from epileptic patients. *Electroencephalography and Clinical Neurophysiology* 95:161–77. [PLN]
- Byrne, R. W. & Whiten, A. (1992) Cognitive evolution in primates: Evidence from tactical deception. *Man* 27:609–27. [aWAP]
- Calev, A., Venables, P. H. & Monk, A. F. (1983) Evidence for distinct verbal memory pathologies in severely and mildly disturbed schizophrenics. *Schizophrenia Bulletin* 9:247–64. [SMS]
- Campbell, R. N. & Olson, D. R. (1990) Children's thinking. In: *Understanding children: Essays in honour of Margaret Donaldson*, ed. R. Grieve & M. Hughes. Blackwell. [aWAP]
- Carpenter, G. A. (1989) Neural network models for pattern recognition and associative memory. *Neural Networks* 2:243–57. [aWAP]
- Carpenter, G. A. & Grossberg, S. (1991) *Pattern recognition by self-organizing neural networks*. MIT Press. [SG]
- (1993) Normal and amnesic learning, recognition, and memory by a neural model of cortico-hippocampal interactions. *Trends in Neurosciences* 16:131–37. [SG]
- Castelfranco, A.M., Robertson, L. T. & McCollum, G. (1993) Detail, proportion, and foci among face receptive fields of climbing fiber responses in the cat cerebellum. *Somatosensory and Motor Research* 11:27–46. [GM]
- Cattell, J. M. (1886) The time taken up by the cerebral operations. *Mind* 11:377–92. [aWAP]
- Chevalier-Skolnikoff, S. (1983) Sensori-motor development in orangutans and other primates. *Journal of Human Evolution* 12:545–46. [aWAP]
- Chiang, C., von Stein, A. & König, P. (1996) Synchronous activity between primary visual and sensorimotor cortex in awake behaving cat. *Society for Neuroscience Abstracts* 22:255.2. [PK]
- Cirelli, C., Pompeiano, M. & Tononi, G. (1996) Neuronal gene expression in the waking state: A role for the Locus Coeruleus. *Science* 274:1211–15. [GT]
- Clark, A. & Thornton, C. (1997) Trading spaces: Computation, representation and the limits of uninformed learning. *Behavioral and Brain Sciences* 20(1) 57–90. [aWAP]
- Clothiaux, E. E., Cooper, L. N. & Bear, M. F. (1991) Synaptic plasticity in visual cortex: Comparison of theory with experiment. *Journal of Neurophysiology* 66: 1785–1804. [aWAP]
- Cobb, S. R., Buhl, E. H., Halasy, K., Paulsen, O. & Somogyi, P. (1995) Synchronization of neuronal activity in hippocampus by individual GABAergic interneurons. *Nature* 378:75–78. [rWAP]
- Cohen, J. D. & Servan-Schreiber, D. (1992) Context, cortex, and dopamine: A connectionist approach to behavior and biology in schizophrenia. *Psychological Review* 99:45–77. [SMS]
- Connors, B. W. & Gutnick, M. J. (1990) Intrinsic firing patterns of diverse neocortical neurons. *Trends in the Neurosciences* 13:99–104. [rWAP]
- Cox, M. D. & Leventhal, D. N. (1978) A multivariate analysis and modification of a preattentive perceptual dysfunction in schizophrenia. *Journal of Nervous and Mental Disease* 166:709–18. [rWAP]
- Creutzfeldt, O. D. (1978) The neocortical link: Thoughts on the generality of structure and function of the neocortex. In: *Architectonics of the cerebral cortex*, ed. M. A. B. Brazier & H. Petsche. Raven Press. [JVS]
- Crick, F. (1988) *What mad pursuit*. Penguin. [aWAP]
- (1989) The recent excitement about neural networks. *Nature* 337:129–32. [DF]
- Crick, F. & Koch, C. (1990) Towards a neurobiological theory of consciousness. *Seminars in the Neurosciences* 2:263–75. [aWAP]
- Damasio, A. R. (1989) The brain binds entities and events by multiregional activity from convergence zones. *Neural Computation* 1:123–32. [aWAP]
- D'Angelo, E., Rossi, P. & Garthwaite, J. (1990) Dual-component NMDA receptor currents at a single central synapse. *Nature* 346:467–69. [rWAP]
- Das, A. & Gilbert, C. D. (1995) Long-range horizontal connections and their role in cortical reorganization revealed by optical recording of cat primary visual cortex. *Nature* 375:780–84. [aWAP]
- deCharms, R. C. & Merzenich, M. M. (1996) Primary cortical representation of sounds by the coordination of action-potential timing. *Nature* 381:610–13. [rWAP]
- De Loach, J. (1987) Rapid change in the symbolic functioning of very young children. *Science* 238:1556–57. [aWAP]
- Der, R. & Smyth, D. (in press) Local online learning of coherent information. *Neural Networks*. [aWAP]
- de Sa, V. (1994a) *Unsupervised classification learning from cross-modal environmental structure*. PhD thesis, University of Rochester, NY. [aWAP]
- (1994b) Learning classification with unlabeled data. In: *Advances in neural information processing systems* 6. Morgan Kaufmann. [aWAP]
- de Sa, V. & Ballard, D. H. (1997) Category learning through multi-modality sensing. *Neural Computation* (in press). [arWAP]
- Douglas, R. J. & Martin, K. A. C. (1990) Neocortex. In: *The synaptic organization of the brain*, ed. G. M. Shepherd. Oxford University Press. [aWAP]
- Douglas, R. J. & Martin, K. A. C. (1991a) A functional microcircuit for cat visual cortex. *Journal of Physiology* 440:735–69. [PK]
- (1991b) Opening the grey box. *Trends in the Neurosciences* 14:286–93. [aWAP]
- (1994) The canonical microcircuit: A co-operative neuronal network for neocortex. In: *Structural and functional organization of the neocortex*, ed. B. Albowitz, K. Albus, U. Kuhnt, H.-Ch. Nothdurft & P. Wahle. Springer-Verlag. [JVS]
- Dudek, S. M. & Bear, M. F. (1992) Homosynaptic long-term depression in area CA1 of hippocampus and the effects on NMDA receptor blockade. *Proceedings of the National Academy of Sciences USA* 89:4363–67. [aWAP]
- Durgin, F. H. (1995) *Contingent aftereffects of texture density: Perceptual learning and contingency*. PhD thesis, Department of Psychology, University of Virginia. [aWAP]
- Ebdon, M. (1993) Is the cerebral neocortex a uniform cognitive architecture? *Mind and Language* 8(3):369–403. [JVS]
- Eckhorn, H. (1997, in press) Neural mechanisms of scene segmentation. *Journal of Artificial Neural Networks*. [RE]

- Eckhorn, R. (1994) Oscillatory and non-oscillatory synchronizations in the visual cortex and their possible roles in associations of visual features. *Progress in Brain Research* 102:405–26. [GP]
- Eckhorn, R., Bauer, R., Jordan, W., Brosch, M., Kruse, W., Munk, M. & Reitboeck, H. J. (1988) Coherent oscillations: A mechanism of feature linking in the visual cortex? *Biological Cybernetics* 60:121–130. [RE, DLW]
- Eckhorn, R., Dicke, P., Arndt, M., & Reirboeck, H. (1991a) Flexible linking of visual features by stimulus-related synchronizations of model neurons. In: *Induced rhythms in the brain*, ed. E. Basar & T. H. Bullock. Birkhauser. [aWAP]
- Eckhorn, R., Frien, A., Bauer, R., Woelbern, T. & Kehr, H. (1993) High frequency (60–90 Hz) oscillations in primary visual cortex of awake monkey. *NeuroReport* 4:243–46. [RE]
- Eckhorn, R. & Gruesser, O. -J., Kroeller, J., Pellnitz, K. & Poepel, B. (1976) Efficiency of different neural codes. *Biological Cybernetics* 22:40–60. [RE]
- Eckhorn, R., Reitboeck, H. J., Arndt, M. & Dicke, P. (1990) Feature linking among distributed assemblies. *Neural Computation* 2:293–306. [RE]
- Eckhorn, R., Schanze, T., Brosch, M., Salem, W. & Bauer, R. (1991b) Stimulus-specific synchronizations in cat visual cortex: Multiple microelectrode and correlation studies from several cortical areas. In: *Induced rhythms in the brain*, ed. E. Basar & T. H. Bullock. Birkhauser. [aWAP]
- Edelman, G. M. (1978) Group selection and phasic re-entrant signalling: A theory of higher brain function. In: *The mindful brain*, ed. G. M. Edelman & V. B. Mountcastle. MIT Press. [aWAP, GT]
- (1987) *Neural Darwinism: The theory of neuronal group selection*. Basic Books. [PK, GT]
- (1989) *The remembered present: Biological theory of consciousness*. Basic Books. [aWAP]
- (1993) Neural Darwinism: Selection and re-entrant signalling in higher brain function. *Neuron* 10:1–20. [GT]
- Edelman, G. M. & Mountcastle, V. B. (1978) *The mindful brain*. MIT Press. [aWAP]
- Ellis, A. W. & Young, A. W. (1988) *Human cognitive neuropsychology*. Erlbaum. [aWAP]
- Engel, A. K., König, P., Kreiter, A. K., Schillen, T. B. & Singer, W. (1992) Temporal coding in the visual cortex: New vistas on integration in the nervous system. *Trends in the Neurosciences* 15:218–26. [aWAP]
- Engel, A. K., König, P., Kreiter, A. K. & Singer, W. (1991a) Interhemispheric synchronization of oscillatory neuronal responses in cat visual cortex. *Science* 252:1177–79. [MI, arWAP]
- Engel, A. K., König, P. & Singer, W. (1991b) Direct physiological evidence for scene segmentation by temporal coding. *Proceedings of the National Academy of Sciences USA* 88:9136–40. [GB, arWAP]
- Engel, A. K., Kreiter, A. K., König, P. & Singer, W. (1991c) Synchronization of oscillatory neuronal response between striate and extrastriate visual cortical areas of the cat. *Proceedings of the National Academy of Sciences USA* 88:6048–52. [arWAP]
- Felleman, D. J. & Van Essen, D. C. (1991) Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex* 1:1–47. [aWAP]
- Field, D. J., Hayes, A. & Hess, R. F. (1993) Contour integration by the human visual system: Evidence for a local “association” field. *Vision Research* 33:173–79. [SG, arWAP]
- Finkel, L. H. & Edelman, G. M. (1989) The integration of distributed cortical systems by reentry: A computer simulation of interactive functionally segregated visual areas. *Journal of Neuroscience* 9:3188–208. [aWAP]
- Fishman, M. C. & Michael, C. R. (1973) Integration of auditory information in the cat's visual cortex. *Vision Research* 13:1415–19. [aWAP]
- Floreano, D., Phillips, W. A. & Kay J. (1995) A computational theory of learning visual features via contextual guidance. *Perception* 24(suppl):22. [aWAP]
- Fodor, J. A. & Pylyshyn, Z. W. (1988) Connectionism and cognitive architecture: A critical analysis. *Cognition* 28:3–71. [aWAP]
- Foldiak, P. (1990) Forming sparse representations by local anti-Hebbian learning. *Biological Cybernetics* 64:165–70. [aWAP]
- (1991) Learning invariance from transformation sequences. *Neural Computation* 3(2):194–200. [JVS]
- Fox, K. & Daw, N. (1992) A model for the action of NMDA conductances in the visual cortex. *Neural Computation* 4:59–83. [aWAP]
- Fox, K., Sato, H. & Daw, N. (1990) The effect of varying stimulus intensity on NMDA-receptor activity in cat visual cortex. *Journal of Neurophysiology* 64:1413–28. [aWAP]
- Francis, G., Grossberg, S. & Mingolla, E. (1994) Cortical dynamics of feature binding and reset: Control of visual persistence. *Vision Research* 34:1089–1104. [SG]
- Freiwald, W. A., Kreiter, A. K. & Singer, W. (1995) Stimulus dependent intercolumnar synchronization of single unit responses in cat area 17. *NeuroReport* 6:2348–52. [rWAP]
- Frien, A., Eckhorn, R., Bauer, R., Woelbern, T. & Kehr, H. (1994) Stimulus-specific fast oscillations at zero phase between visual areas V1 and V2 of awake monkey. *NeuroReport* 5:2273–77. [RE]
- Frien, A., Eckhorn, R. & Reitboeck, H. J. (1996) Fast oscillations in V1 of awake monkey. *Society for Neuroscience Abstracts* 22:255.5. [RE]
- Fries, P., Roelfsma, P. R., Engel, A. K., König, P. & Singer, W. (1996) Synchronized gamma frequency oscillations correlate with perception during binocular rivalry in awake squinting cats. *Society of Neuroscience Abstracts* 22(1):117.3 [arWAP]
- Friston, K. J., Tononi, G., Reeke, G. N., Jr., Sporns, O. & Edelman, G. M. (1994) Value-dependent selection in the brain: Simulation in a synthetic neural model. *Neuroscience* 59:229–43. [GT]
- Frith, C. D. (1992) *The Cognitive neuropsychology of schizophrenia*. Erlbaum. [rWAP]
- Gallistel, C. R. (1995) The replacement of general-purpose theories with adaptive specializations. In: *The cognitive neurosciences*, ed. M. S. Gazzaniga. MIT Press. [aWAP]
- Geman, S., Bienenstock, E. & Doursat, R. (1992) Neural networks and the bias/variance dilemma. *Neural Computation* 4:1–58. [GT]
- Georgopoulos, A. P. (1990) Neural coding of the direction of reaching and a comparison with saccadic eye-movements. *Cold Spring Harbour Symposium on Quantitative Biology* 55:849–59. [aWAP]
- Gerstner, W. (1995) Time structure of the activity in neural network models. *Physical Review E* 51(6):738–58. [GP]
- Ghose, G. M. & Freeman, R. D. (1992) Oscillatory discharge in the visual system: Does it have a functional role? *Journal of Neurophysiology* 68:1558–74. [VGH]
- Giffins, R. (1985) *Canonical analysis: A review with applications in ecology (Biomathematics 12)*. Springer-Verlag. [aWAP]
- Gilbert, C. D. (1992) Horizontal integration and cortical dynamics. *Neuron* 9:1–13. [aWAP]
- (1995) Dynamic properties of adult visual cortex. In: *The cognitive neurosciences*, ed. M. S. Gazzaniga. MIT Press. [aWAP]
- Gilbert, C. D. & Wiesel, T. N. (1983) Clustered intrinsic connections in cat visual cortex. *Journal of Neuroscience* 3:1116–33. [aWAP]
- (1989) Columnar specificity of intrinsic horizontal and cortico-cortical connections in cat visual cortex. *Journal of Neuroscience* 9:2432–42. [aWAP]
- (1990) The influence of contextual stimuli on the orientation selectivity of cells in primary visual cortex of the cat. *Vision Research* 11:1689–1701. [aWAP]
- Gluck, M. A. & Rumelhart, D. E. (1990) *Neuroscience and connectionist theory*. Erlbaum. [aWAP]
- Goebel, R. (1993) Perceiving complex visual scenes: An oscillator neural network model that integrates selective attention, perceptual organisation, and invariant recognition. In: *Advances in neural information processing systems* 5, ed. S. J. Hanson, J. D. Cowan & C. L. Giles. Morgan Kaufmann. [aWAP]
- Golomb, D., Hertz, J., Panzeri, S., Treves, A. & Richmond, B. J. (1997) How well can we estimate mutual information from limited samples of neuronal responses? *Neural Computation* 9:649–65. [AT]
- Goodall, W. C. (1994) Neuropsychological studies of reading and writing. PhD thesis, University of Stirling, Scotland, UK. [aWAP]
- Goodall, W. C. & Phillips, W. A. (1994) Three routes from print to sound: Evidence from a case of acquired dyslexia. *Cognitive Neuropsychology* 12:113–47. [aWAP]
- Gove, A., Grossberg, S. & Mingolla, E. (1995) Brightness perception, illusory contours, and corticogeniculate feedback. *Visual Neuroscience* 12:1027–52. [SG]
- Gray, C. M., König, P., Engel, A. K. & Singer, W. (1989) Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Nature* 338:334–37. [JMB, arWAP, DLW]
- Gray, C. M. & McCormick, D. A. (1996) Chattering cells - superficial pyramidal neurons contributing to the generation of synchronous oscillations in the visual cortex. *Science* 274:109–13. [JJW]
- Gray, C. M. & Singer, W. (1989) Stimulus-specific neuronal oscillations in orientation columns of cat visual cortex. *Proceedings of the National Academy of Sciences USA* 86:1698–1702. [aWAP]
- Gray, C. M. & Viana Di Prisco, G. (1993) Properties of stimulus-dependent rhythmic activity of visual cortical neurons in the alert cat. *Society of Neuroscience Abstracts* 19:359.8. [aWAP]
- Gregson, R. A. M. (1988) *Nonlinear psychophysical dynamics*. Erlbaum Associates. [RAMG]
- (1992) *n-Dimensional nonlinear psychophysics*. Erlbaum Associates. [RAMG]
- (1993) Learning in the context of nonlinear psychophysics: The Gamma Zak Embedding. *British Journal of Mathematical and Statistical Psychology* 46:31–48. [RAMG]
- (1995) *Cascades and fields in perceptual psychophysics*. World Scientific. [RAMG]
- (1996) n-Dimensional nonlinear psychophysics: Intersensory interaction as a

- network at the edge of chaos. In: *Fractals of brain, fractals of mind*, ed. E. MacCormac & M. I. Stamenov. John Benjamins. [RAMG]
- Gregson, R. A. M. & Harvey, J. P. (1992) Similarities of low-dimensional chaotic auditory attractor sequences to quasirandom noise. *Perception and Psychophysics* 51:267–78. [RAMG]
- Grossberg, S. (1968) Some nonlinear networks capable of learning a spatial pattern of arbitrary complexity. *Proceedings of the National Academy of Sciences* 60:758–65. [SG]
- (1971) Pavlovian pattern learning by nonlinear neural networks. *Proceedings of the National Academy of Sciences* 68:828–31. [SG]
- (1982) *Studies of mind and brain*. Reidel. [SG]
- (1987) *The adaptive brain, vol. II*. Elsevier. [SG]
- (1993) Self-organizing neural models of categorization, inference and synchrony. *Behavioral and Brain Sciences* 16:460–61. [aWAP]
- (1995) The attentive brain. *American Scientist* 83:438–49. [SG]
- Grossberg, S., Boardman, L. & Cohen, M. (1997a) Neural dynamics of variable-rate speech categorization. *Journal of Experimental Psychology: Human Perception and Performance* 23:481–503. [SG]
- Grossberg, S. & Mingolla, E. (1985) Neural dynamics of perceptual grouping: Textures, boundaries, and emergent segmentations. *Perception and Psychophysics* 38:141–71. [SG]
- Grossberg, S., Mingolla, E. & Ross, W. (1997b) Visual brain and visual perception: How does the cortex do perceptual grouping? *Trends in Neurosciences* 20:106–111. [SG]
- Grossberg, S. & Somers, D. (1991) Synchronized oscillations during cooperative feature linking in a cortical model of visual perception. *Neural Networks* 4:453–66. [SG, aWAP]
- Guettler, A., Eckhorn, R., Juegens, E. & Frien, A. (1997) Neural correlation contrast in visual cortex of monkey changes with stimulus contrast across an object-background-border. In: *From membrane to mind*. Thieme. [RE]
- Gur, M. & Akri, V. (1992) Isoluminant stimuli may not expose the full contribution of color to visual functioning: Spatial contrast sensitivity measurements indicate interaction between color and luminance processing. *Vision Research* 32:1253–62. [aWAP]
- Hamming, R. W. (1980) *Coding and information theory*. Prentice-Hall. [aWAP]
- Hancock, P. J. B., Smith, L. S. & Phillips W. A. (1991a) A biologically supported error-correcting learning rule. *Neural Computation* 3:201–12. [aWAP]
- (1991b) A biologically supported error-correcting learning rule. In: *Proceedings of the International Conference on Artificial Neural Networks*, ed O. Simula. Elsevier. [aWAP]
- Hansel, C., Artola, A. & Singer, W. (in preparation) Relation between dendritic Ca²⁺ levels and the polarity of synaptic long-term modifications in rat visual cortex neurons. [rWAP]
- Harter, M. R. (1967) Excitability cycles and cortical scanning: A review of two hypotheses of central intermittency in perception. *Psychological Bulletin* 68:47–58. [VGH]
- Hebb, D. O. (1949) *The organization of behaviour*. Wiley. [PK, GP, aWAP]
- Heeger, D. J. (1993) Modeling single cell direction selectivity with normalized, half-squared, linear operators. *Journal of Neurophysiology* 70(5):1885–98. [DS]
- Hestrin, S., Sah, P. & Nicoll, R. A. (1990) Mechanisms generating the time course of dual component excitatory synaptic currents recorded in hippocampal slices. *Neuron* 5:247–53. [rWAP]
- Hinton, G. E., McClelland, J. L. & Rumelhart, D. E. (1986) Distributed representations. In: *Parallel distributed processing. Explorations in the microstructure of cognition*, vol. 1, ed. D. E. Rumelhart. & J. L. McClelland. MIT Press. [aWAP]
- Hirsch, J. A. & Gilbert, C. D. (1991) Synaptic physiology of horizontal connections in cat's visual cortex. *Journal of Neuroscience* 11:1800–9. [aWAP]
- (1993) Long-term changes in synaptic strength along specific intrinsic pathways in the cat visual cortex. *Journal of Physiology* 461:247–62. [aWAP]
- Hopfield, J. J. (1982) Neural networks and physical systems with emergent collective computational capabilities. *Proceedings of the National Academy of Sciences USA* 79:2554–58. [aWAP]
- Horgan, J. (1995) From complexity to perplexity. *Scientific American* June:74–79. [aWAP]
- Hotelling, H. (1936) Relations between two sets of variables. *Biometrika* 28:321–77. [aWAP]
- Huang, Y.-Y., Colino, A., Selig, D. K. & Malenka, R. C. (1992) The influence of prior synaptic activity on the induction of long-term potentiation. *Science* 255:730–33. [aWAP]
- Hummel, J. E. & Biederman, I. (1992) Dynamic binding in a neural network for shape recognition. *Psychological Review* 99:480–517. [aWAP]
- Hummel, J. E. & Holyoak, K. J. (1993) Distributing structure over time. *Behavioral and Brain Sciences* 16:464. [aWAP]
- Humphreys, G. W. & Riddoch, M. J. (1987) *To see but not to see: A case study of visual agnosia*. Erlbaum. [aWAP]
- Humphreys, G. W., Troscianko, T., Riddoch, M. J., Boucart, M., Donnelly, N. & Harding, G. F. A. (1992) Covert processing in different visual recognition systems. In: *The neuropsychology of consciousness*, ed. A. D. Milner & M. D. Rugg. Academic Press. [aWAP]
- Iacoboni, M., Rayman, J. & Zaidel, E. (1996a) Left brain says yes, right brain says no: Normative duality in the split brain. In: *Toward a science of consciousness: The first Tucson discussions and debates*, eds. S. R. Hameroff, A. W. Kaszniak, A. C. Scott. MIT Press. [MI]
- (in preparation) Lateralized lexical decision with unilateral and bilateral presentations: Evidence from the split brain. [MI]
- Iacoboni, M., Woods, R. P., Lenzi, G. L. & Mazziotta, J. C. (1997) Merging of oculomotor and somatomotor space coding in the human right precentral gyrus. *Brain* 120(3). [MI]
- Iacoboni, M., Woods, R. P. & Mazziotta, J. C. (1996b) Brain-behavior relationships: Evidence from practice effects in spatial stimulus-response compatibility. *Journal of Neurophysiology* 76:321–31. [MI]
- Iacoboni, M. & Zaidel, E. (1996) Hemispheric independence in word recognition: Evidence from unilateral and bilateral presentations. *Brain and Language* 53:121–40. [MI]
- Ingber, L. (1985) Statistical mechanics of neocortical interactions: Stability and duration of the 7+/- rule of short-term memory capacity. *Physical Review A*, 31:1183–86. [PLN]
- (1995) Statistical mechanics of multiple scales of neocortical interactions. In: *Neocortical dynamics and human EEG rhythms*, ed. P. L. Nunez. Oxford University Press. [PLN]
- Intrator, N. & Cooper, L. N. (1995a) Information theory and visual plasticity. In: *The handbook of brain theory and neural networks*, ed. M. A. Arbib. MIT Press. [aWAP]
- (1995b) BCM theory of visual cortical plasticity. In: *The handbook of brain theory and neural networks*, ed. M. A. Arbib. MIT Press. [aWAP]
- Ishai, A. & Sagı, D. (1995) Common mechanisms of visual perception and imagery. *Science* 268:1772–74. [aWAP]
- Jakobson, R. (1964) Towards a linguistic typology of aphasic impairments. In: *Disorders of language. Cyba Foundation Symposium*, ed. A. V. S. de Reuck & M. O'Connor. Little, Brown & Company. [VGH]
- James, W. (1890) *The Principles of Psychology*. Reprint. Dover. [GT]
- Jensen, O., Idiart, M. A. P. & Lisman, J. E. (1996) Physiologically realistic formation of autoassociative memory in networks with theta/gamma oscillations: Role of fast NMDA channels. *Learning and Memory* 3:243–56. [rWAP]
- Jerison, H. J. (1973) *Evolution of the brain and intelligence*. Academic Press. [aWAP]
- Jirsa, V. K. & Haken, H. (1996) Field theory of electromagnetic brain activity. *Physical Review Letters* 77: 960–63. [PLN]
- Johnston, J. C. & McClelland, J. L. (1973) Visual factors in word perception. *Perception and Psychophysics* 14:365–70. [aWAP]
- Juegens, E. & Eckhorn, R. (1997) Parallel processing by a homogenous group of coupled model neurons can enhance, reduce and generate signal correlations. *Biological Cybernetics* 76:217–27. [RE]
- Kaas, J. H. (1995) The reorganization of sensory and motor maps in adult mammals. In: *The cognitive neurosciences*, ed. M. S. Gazzaniga. MIT Press. [aWAP]
- Kanevsky, D. (1989) A multiple source, or, is a striped apple more than a striped orange? *Behavioral and Brain Sciences* 12:767–69. [aWAP]
- Kapadia, M. K., Ito, M., Gilbert, C. D. & Westheimer, G. (1995) Improvement in visual sensitivity by changes in local context: Parallel studies in human observers and in V1 of alert monkeys. *Neuron* 15:843–56. [SG, aWAP]
- Karni, A. & Sagı, D. (1991) Where practice makes perfect in texture discrimination: Evidence for primary visual cortex plasticity. *Proceedings of the National Academy of Sciences USA* 88:4966–70. [aWAP]
- Kastner, S., Nothdurft, H. -C. & Pigarev, I. N. (1997) Neuronal correlates of pop-out in cat striate cortex. *Vision Research* 37(4):371–76. [DS]
- Katznelson, R. D. (1981) Normal modes of the brain: Neuroanatomic basis and a physiologic theoretical model. In: *Electric fields of the brain: The neurophysics of EEG*, ed. P. L. Nunez. Oxford University Press. [PLN]
- Kay, J. (1992) Feature discovery under contextual supervision using mutual information. *Proceedings of the International Joint Conference on Neural Networks* Book 4:79–84. Baltimore. [aWAP]
- Kay, J., Floreano, D. & Phillips, W. A. (in press) Contextually guided unsupervised learning using local multivariate binary processors. *Neural Networks*. [aWAP]
- Kay, J. & Phillips, W. A. (1994) *Activation functions, computational goals and learning rules for local processors with contextual guidance (Technical Report CCCN-15)*. Centre for Cognitive and Computational Neuroscience, University of Stirling, Scotland, UK. [aWAP]
- Kay, J. & Phillips, W. A. (1997) Activation functions, computational goals and learning rules for local processors with contextual guidance. *Neural Computation* 9:763–78. [aWAP]

- Kiper, D. C., Gegenfurtner, K. R. & Movshon, J. A. (1996) Cortical oscillatory responses do not affect visual segmentation. *Vision Research* 36(4):539–44. [rWAP]
- Kirkwood, A., Rioult, M. G. & Bear, M. F. (1996) Experience-dependent modification of synaptic plasticity in visual cortex. *Nature* 381:526–28. [aWAP]
- Kisvárdy, Z. F., Martin, K. A. C., Freund, T. F., Maglóczky, Z., Whitteridge, D. & Somogyi, P. (1986) Synaptic targets of HRP-filled layer III pyramidal cells in the cat striate cortex. *Experimental Brain Research* 64:541–52. [aWAP]
- Knierem, J. J. & Van Essen, D. C. (1992) Neuronal responses to static texture patterns in area V1 of the alert macaque monkey. *Journal of Neurophysiology* 67:961–80. [aWAP]
- Knight, R. A. (1992) Specifying cognitive deficiencies in poor premorbid schizophrenics. In: *Progress in experimental personality and psychopathology research: Vol. 15*, ed. E. F. Walker, R. H. Dworkin & B. A. Cornblatt. Springer.
- Knight, R. A. & Silverstein, S. M. (In press) The role of cognitive psychology in guiding research on cognitive deficits in schizophrenia. In: *Experimental psychopathology and the pathogenesis of schizophrenia*, ed. M. F. Lenzenweger & R. H. Dworkin. APA Press. [SMS]
- Kohonen, T. (1982) Self-organizing formation of topologically correct feature maps. *Biological Cybernetics* 43:59–69. [PM]
- König, P., Engel, A. K., Löwel, S. & Singer, W. (1993) Squint affects synchronization of oscillatory response in cat visual cortex. *European Journal of Neuroscience* 5:501–8. [GB, aWAP]
- König, P., Engel, A. K., Roelfsema, P. R. & Singer, W. (1996a) Coincidence detection or temporal integration. The role of the cortical neuron revisited. *TINS* 19:130–37. [PK]
- König, P., Engel, A. K. & Singer, W. (1995) Relation between oscillatory activity and long-range synchronization in cat visual cortex. *Proceedings of the National Academy of Sciences USA* 92:290–94. [arWAP]
- (1996b) Integrator or coincidence detector? The role of the cortical neuron revisited. *Trends in the Neurosciences* 19:130–37. [rWAP]
- Kottmann, H., Eckhorn, R., Woelber, T., Frien, A. & Bauer, R. (1996) Synchronized fast oscillations in striate cortex correlate with perception of grating orientation in a binocular rivalry task in monkey. In: *Brain and evolution*, ed. N. Elsner & H.-U. Schnitzler. Thieme. [RE]
- Kreiter, A. K. & Singer, W. (1992) Oscillatory neuronal responses in the visual cortex of the awake macaque monkey. *European Journal of Neuroscience* 4:369–75. [aWAP]
- (1994) Global stimulus arrangement determines synchronization of neuronal activity in the awake macaque monkey. *Supplement European Journal of Neuroscience* 7:153. [aWAP]
- (1996) Stimulus-dependent synchronization of neuronal responses in the visual cortex of the awake macaque monkey. *Journal of Neuroscience* 16:2381–96. [DJA, arWAP]
- Kröse, B. J. A. & Julesz, B. (1989) The control and speed of shifts of attention. *Vision Research* 29:1607–9. [aWAP]
- Krüger, J. & Becker, J. D. (1991) Recognizing the visual stimulus from neuronal discharges. *Trends in Neurosciences* 14:282–86. [VGH]
- Krueger, L. E. & Shapiro, R. G. (1979) Letter detection with rapid serial presentation: Evidence against word superiority at feature extraction. *Journal of Experimental Psychology: Human Perception and Performance* 5:657–73. [aWAP]
- Kruse, W. & Eckhorn, R. (1996) Inhibition of sustained gamma oscillations (35–80 Hz) by fast transient responses in cat visual cortex. *Proceedings of the National Academy of Sciences USA* 93:6112–17. [RE]
- Krystal, J. H. et al (1994) Subanesthetic effects of the noncompetitive NMDA antagonist ketamine in humans. *Archives of General Psychiatry* 51:199–214. [rWAP]
- Kuramoto, Y. (1984) *Chemical oscillations, waves and turbulence*. Springer. [GP]
- Lades, M., Vorbrüggen, J. C., Buhmann, J., Lange, J., von der Malsburg, C., Würtz, R. P. & Konen, W. (1993) Distortion invariant recognition in the dynamic link architecture. *IEEE Transactions on Computers* 42:300–11. [arWAP, RPW]
- Langdon, R. B. & Sur, M. (1990) Components of field potentials evoked by white matter stimulation in isolated slices of primary visual cortex: Spatial distributions and synaptic order. *Journal of Neurophysiology* 64:1484–1501. [PK]
- Lashley, K. (1951) The problem of serial order in behavior. In: *Cerebral Mechanisms in Behavior. The Hixon Symposium*, ed. L. A. Jeffress. Wiley. [VGH]
- Laurent, G. (1996) Dynamical representation of odors by oscillating and evolving neural assemblies. *Trends in Neurosciences* 19:489–96. [VGH]
- Leonards, U. & Singer, W. (1997) Interactions between segmentation mechanisms with differential sensitivity for color and luminance contrast. Submitted to *Vision Research*. [rWAP]
- Leonards, U., Singer, W. & Fahle, M. (1996) The influence of temporal phase differences on texture segmentation. *Vision Research* 36:2689–97. [arWAP]
- Levitt, J. B. & Lund, J. S. (1996) Contrast dependence of modulatory surround effects in macaque striate neurons. *Investigative Ophthalmology & Visual Science* 37(3):S482. [DS]
- Li, Z. & Atick, J. J. (1994) Efficient stereo coding in the multiscale representation. *Network* 5:157–74. [aWAP]
- Linsker, R. (1988) Self-organization in a perceptual network. *Computer* 21:105–17. [aWAP, JVS]
- Lisman, J. E. & Idiart, M. A. P. (1995) Storage of $7+/-2$ short-term memories in oscillatory subcycles. *Science* 267:1512–14. [rWAP]
- Llinás, R. & Sasaki, K. (1989) The functional organization of the olivo-cerebellar system as examined by multiple Purkinje cell recordings. *European Journal of Neuroscience* 1:587–602. [GM]
- Llinas, R. R., Grace, A. A. & Yarom, Y. (1991) In vitro neurons in cortical layer 4 exhibit intrinsic oscillatory activity in the 10 to 50 Hz frequency range. *Proceedings of the National Academy of Sciences USA* 88:897–901. [JJW]
- Lorentó de No, R. (1949) Cerebral cortex: Architecture, intracortical connections, motor projections. In: *Physiology of the nervous system*, ed. J. F. Fulton. Oxford University Press. [aWAP]
- Löwel, S. & Singer, W. (1992) Selection of intrinsic horizontal connections in the visual cortex by correlated neuronal activity. *Science* 255:209–12. [aWAP]
- Lumer, E. D., Edelman, G. M. & Tononi, G. (1997) Neural dynamics in a model of the thalamocortical system. I. Layers, loops, and the emergence of fast synchronous oscillations. *Cerebral Cortex* 7:207–27. [GT]
- MacLennan, B. (1997) Field computation in motor control. In: *Self-organization, computational maps, and motor control*, ed. P. Morasso & V. Sanguineti. Elsevier. [PM]
- Macphail, E. M. (1987) The comparative psychology of intelligence. *Behavioral and Brain Sciences* 10:645–95. [aWAP]
- Mainen, Z. F. & Sejnowski, T. J. (1995) Reliability of spike timing in neocortical neurons. *Science* 268:1503–06. [PK]
- Marr, D. (1970) A theory for cerebral neocortex. *Proceedings of the Royal Society of London (B)* 176:161–234. [JVS]
- (1982) *Vision*. Freeman. [aWAP]
- Martin, K. A. C. (1988) The Wellcome Prize lecture: From single cells to simple circuits in the cerebral cortex. *Quarterly Journal of Experimental Physiology* 73:637–702. [aWAP]
- Martinetz, T. & Schulten, K. (1994) Topology representing networks. *Neural Networks* 7:507–22. [PM]
- Massaro, D. W. (1979) Letter information and orthographic context in word perception. *Journal of Experimental Psychology: Human Perception and Performance* 5:595–609. [aWAP]
- (1989a) Testing between the TRACE model and the fuzzy logic model of speech perception. *Cognitive Psychology* 21:398–21. [arWAP]
- (1989b) Multiple book review of *Speech perception by ear and eye: A paradigm for psychological inquiry*. *Behavioral and Brain Sciences* 12:741–94. [aWAP]
- Massaro, D. W. & Cohen, M. M. (1983) Phonological constraints in speech perception. *Perception and Psychophysics* 34:338–48. [aWAP]
- (1991) Integration versus interactive activation: The joint influence of stimulus and context in perception. *Cognitive Psychology* 23:558–614. [aWAP]
- Massaro, D. W. & Friedman, D. (1990) Models of integration given multiple sources of information. *Psychological Review* 97:225–52. [aWAP]
- Matelli, M., Luppino, G. & Rizzolatti, G. (1991) Architecture of superior and mesial area 6 and the adjacent cingulate cortex in the macaque monkey. *Journal of Comparative Neurology* 311:445–62. [MI]
- Maunsell, J. H. R. (1995) The brain's visual world: Representations of visual targets in cerebral cortex. *Science* 270:764–68. [rWAP]
- Mayford, M., Wang, J., Kandel, E. R. & O'Dell, T. J. (in press) CaMKII regulates the frequency response function of hippocampal synapses for the production of both LTD and LTP. *Cell*. [aWAP]
- McCarthy, R. A. & Warrington, E. K. (1990) *Cognitive neuropsychology. A clinical introduction*. Academic Press. [aWAP]
- McClelland, J. L. (1978) Perception and masking of wholes and parts. *Journal of Experimental Psychology: Human Perception and Performance* 4:210–23. [aWAP]
- (1991) Stochastic interactive processes and the effect of context on perception. *Cognitive Psychology* 23:1–44. [aWAP]
- McClelland, J. L. & Elman, J. L. (1986) The TRACE model of speech perception. *Cognitive Psychology* 18:1–86. [aWAP]
- McClelland, J. L., McNaughton, B. L. & O'Reilly, R. C. (1995) Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review* 102:419–507. [aWAP]
- McClelland, J. L. & Rumelhart, D. E. (1981) An interactive activation model of context effects in letter perception: I. An account of basic findings. *Psychological Review* 88:375–407. [aWAP]

- McClelland, J. L., Rumelhart, D. E. & Hinton, G. E. (1986) The appeal of parallel distributed processing. In: *Parallel distributed processing: Explorations in the microstructure of cognition*, vol. 1, ed. D. E. Rumelhart & J. L. McClelland. MIT Press. [aWAP]
- McClurkin, J. W., Optican, L. M., Richmond, B. J. & Gawne, T. J. (1991) Concurrent processing and complexity of temporally encoded messages in visual perception. *Science* 253:675–77. [VGH]
- McCollum, G. (1994) Navigating a set of discrete regions in body position space. *Journal of Theoretical Biology* 167:263–71. [GM]
- (1997) Dynamics of cerebellar neurons. *International Journal of Theoretical Physics* (in press). [GM]
- McCollum, G., Holroyd, C. & Castelfranco, A. M. (1995) Forms of early walking. *Journal of Theoretical Biology* 176:373–90. [GM]
- McGuire, B. A., Gilbert, C. D., Rivlin, P. K. & Wiesel, T. N. (1991) Targets of horizontal connections in macaque primary visual cortex. *Journal of Comparative Neurology* 305:370–92. [aWAP]
- Michalski, A., Gerstein, G. L., Czarkowska, J. & Tarnecki, R. (1983) Interactions between cat striate cortex neurons. *Experimental Brain Research* 51:97–107. [aWAP]
- Milner, P. M. (1974) A model for visual shape recognition. *Psychological Review* 81:521–35. [aWAP, DLW]
- Mizobe, K., Polat, U., Kasamatsu, T. & Norcia, A. M. (1996) Lateral masking reveals facilitation and suppression from the same single cells in cat area 17. *Investigative Ophthalmology & Visual Science* 37(3):S483. [DS]
- Mohr, B., Pulvermüller, F., Rayman, J. & Zaidel, E. (1994) Interhemispheric cooperation during lexical processing is mediated by the corpus callosum: Evidence from the split-brain. *Neuroscience Letters* 181:17–21. [MI]
- Monyer, H., Burnashev, N., Laurie, D. J., Sakmann, B. & Seeburg, P. H. (1994) Developmental and regional expression in the rat brain and functional properties of four NMDA receptors. *Neuron* 12:529–40. [rWAP]
- Morasso, P. (1981) Spatial control of arm movements. *Experimental Brain Research* 42:223–27. [PM]
- Morasso, P. & Sanguineti, V. (1996) How the brain can discover the existence of external egocentric space. *Neurocomputing* 12:289–310. [PM]
- Morton, J. (1969) Interaction of information in word recognition. *Psychological Review* 76:165–78. [aWAP]
- Movellan, J. R. & McClelland, J. L. (1995) *Stochastic interactive processing, channel separability, and optimal perceptual inference: An examination of Morton's law (Technical Report PDP.CNS.95.4)*. Department of Psychology, Carnegie Mellon University, Pittsburgh, PA. [aWAP]
- Mozer, M. C., Zemel, R. S., Behrmann, M. & Williams, C. K. I. (1992) Learning to segment images using dynamic feature binding. *Neural Computation* 4:650–65. [aWAP]
- Mumford, D. (1992) On the computational architecture of the neocortex: 2. The role of cortico-cortical loops. *Biological Cybernetics* 66:241–51. [aWAP]
- Munk, M. H. J., Nowak, L. G., Chouvet, G., Nelson, J. I. & Bullier, J. (1992) The structural basis of cortical synchronization. *European Journal of Neuroscience Supplement* 5:21. [aWAP]
- Munk, M. H. J., Nowak, L. G., Nelson, J. I. & Bullier, J. (1995) Structural basis of cortical synchronization. 2. Effects of cortical-lesions. *Journal of Neurophysiology* 74:2401–14. [GB]
- Munk, M. H. J., Roelfsema, P. R., König, P., Engel, A. K. & Singer, W. (1996) Role of reticular activation in the modulation of intracortical synchronization. *Science* 272:271–74. [rWAP]
- Murthy, V. N. & Fetz, E. E. (1992) Coherent 25- to 35-Hz oscillations in the sensorimotor cortex of awake behaving monkeys. *Proceedings of the National Academy of Sciences (USA)* 89:5670–74. [VGH]
- (1996) Synchronization of neurons during local field potential oscillations in sensorimotor cortex of awake monkeys. *Journal of Neurophysiology* 76:3968–82. [rWAP]
- Nakayama, K. & Mackeben, M. (1989) Sustained and transient components of focal visual attention. *Vision Research* 31:1221–36. [aWAP]
- Navon, D. (1977) Forest before trees: The precedence of global features in visual perception. *Cognitive Psychology* 9:353–83. [rWAP]
- Nelson, J. I. (1995) Binding in the visual system. In: *The handbook of brain theory and neural networks*, ed. M. A. Arbib. MIT Press. [aWAP]
- Nelson, J. I. & Frost, B. J. (1978) Orientation-selective inhibition from beyond the classic visual receptive field. *Brain Research* 139:359–65. [aWAP]
- Neuenschwander, S., Engel, A. K., König, P., Singer, W. & Varela, F. J. (1996) Long-range synchronization of oscillatory light responses in the cat retina and lateral geniculate nucleus. *Nature* 379:728–33. [aWAP]
- Neven, H. & Aertsen, A. Rate coherence and event coherence in the visual cortex: A neuronal model of object recognition. *Biological Cybernetics* 67: 309–22. [aWAP]
- Nunez, P. L. (1974) The brain wave equation: A model for the EEG. *Math Biosciences* 21:279–97. [PLN]
- (1981) *Electric fields of the brain: The neurophysics of EEG*. Oxford University Press. [PLN]
- (1989) Generation of human EEG rhythms by a combination of long and short-range neocortical interactions. *Brain Topography* 1:199–215. [PLN]
- (1995a) *Neocortical dynamics and human EEG rhythms*. Oxford University Press. [PLN]
- (1995b) Multiscale neocortical dynamics, experimental EEG measures, and global facilitation of local cell assemblies. *Behavioral and Brain Sciences* 19:305–306. [PLN]
- Nunez, P. L., Silberstein, R. B., Carpentier, M. R., Srinivasan, R., Wijesinghe, R. S., Tucker, D. M., Cadusch, P. J. & Wu, X. (1997) EEG coherency II: Comparisons of multiple measures of EEG coherency differences between resting and cognitive states. *Electroencephalography and Clinical Neurophysiology*, submitted. [PLN]
- Olney, J. W. & Farber, N. B. (1995) Glutamate receptor dysfunction and schizophrenia. *Archives of General Psychiatry* 52:998–1007. [SMS]
- Olshausen, B. A., Anderson, C. H. & van Essen, D. C. (1993) A neurobiological model of visual attention and invariant pattern recognition based on dynamic routing of information. *Journal of Neuroscience* 13(11):4700–19. [RPW]
- (1995) A multiscale dynamic routing circuit for framing size- and position-invariant object representations. *Journal of Computational Neuroscience* 2:45–62. [RPW]
- Olshausen, B. & Field, D. J. (1996) Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature* 381:607–09. [rWAP]
- Palm, G. (1982) *Neural assemblies. An alternative approach to artificial intelligence*. Springer. [GP]
- (1993) On the internal structure of cell assemblies. In: *Brain theory*, ed. A. Aertsen. Elsevier Science Publishers. [GP]
- Palmer, S. E. (1975) The effects of contextual scenes on the identification of objects. *Memory and Cognition* 3: 519–26. [aWAP]
- Pandya, D. N. & Yeterian, E. H. (1985) Architecture and connections of cortical association areas. In: *Cerebral cortex*, ed. A. Peters & E. G. Jones. Plenum Press. [GB]
- Panzeri, S., Biella, G., Rolls, E. T., Skaggs, W. E. & Treves, A. (1996) Speed, noise, information and the graded nature of neuronal responses. *Network* 7:365–70. [AT]
- Pece, A. E. C. (1992) Redundancy reduction of a Gabor representation: A possible computational role for feedback from primary visual cortex to lateral geniculate nucleus. In: *Artificial neural networks 2*, ed. I. Aleksander & J. Taylor. Elsevier. [aWAP]
- Phillips, W. A. (1971) Does familiarity affect transfer from an iconic to a short-term memory? *Perception and Psychophysics* 10:153–57. [aWAP]
- (1997) Theories of cortical computation. In: *Cognitive neuroscience*, ed. M. D. Rugg. University College Press. [aWAP]
- Phillips, W. A. & Goodall, W. C. (1994) Lexical writing can be non-semantic and fluent without practice. *Cognitive Neuropsychology* 12:149–74. [aWAP]
- Phillips, W. A., Hay, I. M., & Smith, L. S. (1993) Lexicality and pronunciation in a simulated neural net. *British Journal of Mathematical and Statistical Psychology* 46:193–205. [aWAP]
- Phillips, W. A., Kay, J., & Smyth, D. (1995a) How local cortical processors that maximise coherent variation could lay foundations for representation proper. In: *Neural computation and psychology*, ed. L. S. Smith & P. J. B. Hancock. Springer-Verlag. [aWAP]
- Phillips, W. A., Kay, J. & Smyth, D. (1995b) The discovery of structure by multi-stream networks of local processors with contextual guidance. *Network* 6:225–46. [aWAP]
- Phillips, W. A. & Singer, W. (1974) Function and interaction of on and off transients in vision: 1. Psychophysics. *Experimental Brain Research* 19:493–506. [aWAP]
- Piaget, J. (1954) *The construction of reality by the child*. Basic Books. [aWAP]
- Picard, N. & Strick, P. L. (1996) Motor areas of the medial wall: A review of their location and functional activation. *Cerebral Cortex* 6:342–53. [MI]
- Pinker, S. (1994) *The language of instinct*. Penguin. [rWAP]
- Place, E. J. S. & Gilmore, G. C. (1980) Perceptual organization in schizophrenia. *Journal of Abnormal Psychology* 89:409–18. [rWAP]
- Plaut, D. C. & McClelland, J. L. (1993) Generalization with componential attractors: Word and nonword reading in an attractor network. In: *Proceedings of the 15th Annual Conference of the Cognitive Science Society*. Erlbaum. [aWAP]
- Poggio, T., Fahle, M. & Edelman, S. (1992) Fast perceptual learning in visual hyperacuity. *Science* 256:1018–20. [aWAP]
- Polat, U. & Norcia, A. M. (1996) Neurophysiological evidence for contrast dependent long-range facilitation and suppression in the human visual cortex. *Vision Research* 36(14):2099–2109. [DS]
- Polat, U. & Sagi, D. (1993) Lateral interactions between spatial channels: Suppression and facilitation revealed by lateral masking experiments. *Vision Research* 33:993–99. [aWAP]

- (1994a) The architecture of perceptual spatial interactions. *Vision Research* 34:73–78. [aWAP]
- (1994b) Spatial interactions in human vision: From near to far via experience-dependent cascades of connections. *Proceedings of the National Academy of Sciences USA* 91:1206–9. [arWAP]
- Poppel, E. (1970) Excitability cycles in central intermittency. *Psychologische Forschung* 34:1–9. [VGH]
- (1978) Time perception. In: *Handbook of sensory physiology. Vol. VIII: Perception*, ed. R. Held, H. W. Leibowitz & H. -L. Teuber. Springer. [VGH]
- (1994) Temporal mechanisms in perception. *International Review of Neurobiology* 37:185–202. [VGH]
- Posner, M. I. & Rothbart, M. K. (1994) Constructing neuronal theories of mind. In: *Large-scale neuronal theories of brain function*, ed. C. Koch & J. L. Davis. MIT Press. [aWAP]
- Proctor, R. W. & Dutta, A. (1993) Do the same stimulus-response relations influence choice reactions initially and after practice? *Journal of Experimental Psychology: Learning, Memory, & Cognition* 19:922–30. [MI]
- Protopapas, A. & Bower, J. M. (1994) Sensitivity in the response of piriform cortex pyramidal cells to fluctuations in synaptic timing. In: *Computation in neurons and neural systems*, ed. F. Eeckman. Kluwer Press. [JMB]
- Purves, D., Riddle, D. R. & LaMantia, A.-S. (1992) Iterated patterns of brain circuitry (or how the cortex gets its spots). *Trends in the Neurosciences* 15:362–68. [aWAP]
- Rabinowicz, E. F., Opler, L. A., Owen, D. R. & Knight, R. A. (1996) The dot enumeration perceptual organization task (DEPOT): Evidence for a short-term memory deficit in schizophrenia. *Journal of Abnormal Psychology* 105:336–48. [rWAP]
- Rakic, P. & Singer, W. (1988) *Neurobiology of neocortex*. Wiley. [aWAP]
- Redlich, A. N. (1993) Redundancy reduction as a strategy for unsupervised learning. *Neural Computation* 5:289–304. [aWAP]
- Reeke, G. N. Jr., Finkel, L. H., Sporns, O. & Edelman, G. M. (1990) Synthetic neural modelling: A multilevel approach to the analysis of brain complexity. In: *Signal and sense: Local and global order in perceptual maps*, ed. G. M. Edelman, W. E. Gall & W. M. Cowan. Wiley. [aWAP, GT]
- Reicher, G. M. (1969) Perceptual recognition as a function of meaningfulness of stimulus material. *Journal of Experimental Psychology* 81:274–80. [aWAP]
- Richman, H. B. & Simon, H. A. (1989) Context effects in letter perception: Comparison of two theories. *Psychological Review* 96:417–32. [aWAP]
- Robinson, P. A., Rennie, C. J. & Wright, J. J. (1997a) Propagation and stability of waves of electrical activity in the cerebral cortex. *Physical Review E*. (in press) [JJW]
- Robinson, P. A., Wright, J. J. & Rennie, C. J. (1997b) Synchronous oscillation in the cerebral cortex. *Physical Review E*. (submitted). [JJW]
- Rockel, A. J., Hiorns, R. W. & Powell, T. P. S. (1980) The basic uniformity in structure of the neocortex. *Brain* 103:221–44. [aWAP]
- Rockland, K. & Lund, J. S. (1983) Intrinsic laminar lattice connections in primate visual cortex. *Journal of Comparative Neurology* 216:303–18. [aWAP]
- Roelfsma, P. R., Engel, A. K., König, P. & Singer, W. (1994a) Oscillations and synchrony in the visual cortex: Evidence for their functional relevance. In: *Oscillatory event-related brain dynamics*, ed. C. Pantev. Plenum. [arWAP]
- (1997) Visuomotor integration is associated with zero time-lag synchronization among cortical areas. *Nature* 385:157–61. [rWAP]
- Roelfsma, P. R., König, P., Engel, A. K., Sireteanu, R. & Singer, W. (1994b) Reduced synchronization in the visual cortex of cats with strabismic amblyopia. *European Journal of Neuroscience* 6:1645–55. [GB, arWAP]
- Rolls, E. T., Treves, A. & Tovee, M. J. (1997) The representational capacity of the distributed encoding of information provided by populations of neurons in the primate temporal visual cortex. *Experimental Brain Research* 114:149–62. [AT]
- Rucci, M., Ttoni, G. & Edelman, G. M. (1997) Registration of neural maps through value-dependent learning: Modeling the alignment of auditory and visual maps in the brain owl's optic tectum. *Journal of Neuroscience* 17:334–52. [GT]
- Ruhau, E. (1995) Time-gestalt and the observer. In: *Conscious experience*, ed. T. Metzinger. Imprint Academic Schöningh. [RAMG]
- Rumelhart, D. E., Hinton, G. E. & Williams, R. J. (1986) Learning internal representations by back-propagating errors. *Nature* 323:533–36. [DF, aWAP]
- Rumelhart, D. E. & McClelland, J. L. (1982) An interactive activation model of context effects in letter perception: 2. The contextual enhancement effect and some tests and extensions of the model. *Psychological Review* 89:60–84. [aWAP]
- (1986) *Parallel distributed processing: Explorations in the microstructure of cognition*, vol. 1. MIT Press. [aWAP]
- Sagi, D. (1995) The psychophysics of texture segmentation. In: *Early vision and beyond*, ed. T. V. Papathomas et al. MIT Press. [rWAP]
- Salin, P. & Bullier, J. (1995) Corticocortical connections in the visual system: Structure and function. *Physiological Reviews* 75:107–54. [rWAP]
- Samuel, A. G. (1981) Phonemic restoration: Insights from a new methodology. *Journal of Experimental Psychology: General* 110:474–94. [aWAP]
- Does lexical information influence the perceptual restoration of phonemes? *Journal of Experimental Psychology: General* 125:28–51. [arWAP]
- Sanes, J. N. & Donoghue, J. P. (1993) Oscillations in local field potentials of the primate motor cortex during voluntary movement. *Proceedings of the National Academy of Sciences (USA)* 90:4470–74. [VGH]
- Sanger, T. D. (1994) Theoretical considerations for the analysis of population coding in motor cortex. *Neural Computation* 6:29–37. [PM]
- Sanguineti, V., Morasso, P. & Frisone, F. (1997) Cortical maps of sensorimotor spaces. In: *Self-organization, computational maps, and motor control*, ed. P. Morasso and V. Sanguineti. Elsevier. [PM]
- Schillen, T. B. & König, P. (1994) Binding by temporal structure in multiple feature domains of an oscillatory neuronal network. *Biological Cybernetics* 70:397–405. [aWAP]
- Schlaggar, B. L. & O'Leary, D. D. M. (1991) Potential of visual cortex to develop an array of functional units unique to somatosensory cortex. *Science* 252:1556–60. [aWAP]
- Schmidhuber, J. & Prelinger, D. (1993) Discovering predictable classifications. *Neural Computation* 5:625–35. [aWAP]
- Schmidt, K. E., Löwel, S., Goebel, R. & Singer, W. (1997) The perceptual grouping criterion of colinearity is reflected by anisotropies of connections in primary visual cortex. *European Journal of Neuroscience* 9:1083–89. [aWAP]
- Schwartz, C. & Bolz, J. (1991) Functional specificity of the long-range horizontal connections in cat visual cortex: A cross-correlation study. *Journal of Neuroscience* 11:2995–3007. [aWAP]
- Schwartz, E. L. (1977) Spatial mapping in the primate sensory projection: Analytic structure and relevance to perception. *Biological Cybernetics* 25:181–94. [GM]
- Sejnowski, T. J., Koch, C. & Churchland, P. S. (1988) Computational neuroscience. *Science* 241:1299–1306. [aWAP]
- Shadlen, M. N. & Newsome, W. T. (1994) Noise, neural codes and cortical organization. *Current Biology* 4:569–79. [rWAP]
- Shallice, T. (1988) *From neuropsychology to mental structure*. Cambridge University Press. [aWAP]
- Shallice, T. (1991) Précis of *From neuropsychology to mental structure*. *Behavioral and Brain Sciences* 14:429–69. [aWAP]
- Shastri, L. & Ajjanagadde, V. (1993) From simple associations to systematic reasoning: A connectionist representation of rules, variables and dynamic bindings using temporal synchrony. *Behavioral and Brain Sciences* 16:417–94. [aWAP]
- Shepherd, G. M., ed. (1990) *The synaptic organization of the brain*. Oxford University Press. [aWAP]
- Shepherd, G. M. & Koch, C. (1990) Introduction to synaptic circuits. In: *The synaptic organization of the brain*, ed. G. M. Shepherd. Oxford University Press. [aWAP]
- Silberstein, R. B. (1995) Neuromodulation of neocortical dynamics. In: *Neocortical dynamics and human EEG rhythms*, ed. P. L. Nunez. Oxford University Press. [PLN]
- Sillito, A. M., Jones, H. E., Gerstein, G. L. & West, D. C. (1994) Feature-linked synchronization of thalamic relay cell firing induced by feedback from the visual cortex. *Nature* 369:479–82. [GB, SG, aWAP]
- Silverstein, S. M., Bakshi, S., Chapman, R. M. & Nowlis, G. (1996c) Reduced plasticity of the visual information processing system in schizophrenia. Evidence from the processing of visual patterns (submitted). [SMS]
- Silverstein, S. M., Knight, R. A., Schwarzkopf, S. B., West, L. L., Osborn, L. M. & Kamin, D. (1996a) Configural and context effects in perceptual organization in schizophrenia. *Journal of Abnormal Psychology* 105:410–20. [SMS]
- Silverstein, S. M., Matteson, S. & Knight, R. A. (1996b) Reduced top-down influence in auditory perceptual organization in schizophrenia. *Journal of Abnormal Psychology* 105:663–67. [SMS]
- Silverstein, S. M., Raulin, M. R., Pristach, E. A. & Pomerantz, J. R. (1992) Perceptual organization and schizotypy. *Journal of Abnormal Psychology* 101:265–70. [rWAP]
- Silverstein, S. M., Schenkel, L. S., Light, G. A. & Schwarzkopf, S. B. (1996d) Clinical features associated with perceptual organization dysfunction in schizophrenia (submitted). [SMS]
- Simpson, J. I., Wylie, D. R. & De Zeeuw, C. I. (1996) On climbing fiber signals and their consequence(s). *Behavioral and Brain Sciences* 19:380–94. [GM]
- Singer, W. (1987) Activity-dependent self-organization of synaptic connections as a substrate of learning. In: *The neural and molecular bases of learning*, ed. J.-P. Changeux & M. Konishi. Wiley. [aWAP]

- (1990) Search for coherence: A basic principle of cortical self-organization. *Concepts in Neuroscience* 1:1–26. [VGH, aWAP]
- (1993) Synchronization of cortical activity and its putative role in information processing and learning. *Annual Review of Physiology* 55:349–74. [VGH, PLN, aWAP]
- (1994a) Coherence as an organizing principle of cortical functions. *International Review of Neurobiology* 37:153–83. [aWAP]
- (1994b) Time as a coding space in neocortical processing: A hypothesis. In: *Temporal coding in the brain*, ed. G. Buzsáki, R. Llinás, W. Singer, A. Berthoz & Y. Christen. Springer. [VGH]
- (1995) Development and plasticity of cortical processing architectures. *Science* 270:758–64. [aWAP]
- Singer, W. & Artola, A. (1994) Plasticity of the mature cortex. In: *Cellular and molecular mechanisms underlying higher neural functions*, ed. A. I. Selverston & P. Ascher. Wiley. [aWAP]
- Singer, W. & Gray, C. M. (1995) Visual feature integration and the temporal correlation hypothesis. *Annual Review of Neuroscience* 18:555–86. [aWAP, JJW]
- Singer, W. & Phillips, W. A. (1974) Function and interaction of on and off transients in vision: 2. Neurophysiology. *Experimental Brain Research* 19:507–21. [aWAP]
- Smyth, D. (1994) Simulations of networks of local processors with contextual guidance. M.Sc. thesis, Centre for Cognitive and Computational Neuroscience, University of Stirling, Scotland, UK. [aWAP]
- Smyth, D. & Der, R. (1995) Learning to bind. Poster presented at the Symposium on Phenomena and Architectures of Cognitive Dynamics, Universität Leipzig, June. [aWAP]
- Smyth, D., Kay, J. & Phillips, W. A. (1994) Discovery of high-order functions in multi-stream, multi-stage nets without external supervision. Paper presented to the Neural Computation and Psychology Workshop, Centre for Cognitive and Computational Neuroscience, University of Stirling, Scotland, UK. [aWAP]
- Smyth, D., Phillips, W. A. & Kay, J. (1996) Measures for investigating the contextual modulation of information transmission. *Network* 7:307–16. [aWAP]
- Softky, W. R. (1995) Commentary: Simple codes versus efficient codes. *Current Biology* 5(2):239–47. [rWAP]
- Somers, D. & Kopell, N. (1993) Rapid synchronization through fast threshold modulation. *Biological Cybernetics* 68:393–407. [rWAP, DLW]
- Soutif, M. (1970) *Vibrations, propagation, diffusion*. Dunod. [VGH]
- Spinelli, D. N., Starr, A. & Barrett, T. W. (1968) Auditory specificity in unit recordings from cat's visual cortex. *Experimental Neurology* 22:75–84. [aWAP]
- Sporns, O., Gally, J. A., Reeke, G. N., Jr. & Edelman, G. M. (1989) Reentrant signaling among simulated neuronal groups leads to coherency in their oscillatory activity. *Proceedings of the National Academy of Sciences USA* 86:7265–69. [aWAP]
- Sporns, O., Tononi, G. & Edelman, G. M. (1991) Modeling perceptual grouping and figure-ground segregation by means of active re-entrant connections. *Proceedings of the National Academy of Sciences USA* 88:129–33. [aWAP, GT, DLW]
- Squire, L. R. (1992) Memory and the hippocampus: A synthesis from findings with rats, monkeys, and humans. *Psychological Review* 99:195–231. [aWAP]
- Stemmler, M., Usher, M. & Niebur, E. (1995) Lateral interactions in primary visual cortex: A model bridging physiology and psychophysics. *Science* 269:1877–80. [DS]
- Steriade, M., Curro-Dossi, R. & Contreras, D. (1993) Electrophysiological properties of intralaminar thalamocortical cells discharging rhythmic approx. 40 Hz spike bursts at approx. 1000 Hz during waking and rapid eye movement sleep. *Neuroscience* 56:1–9. [JJW]
- Stone, J. V. (1996a) A canonical microfunction for learning perceptual invariances. *Perception* 25(2):207–20. [JVS]
- (1996b) Learning perceptually salient visual parameters through spatio-temporal smoothness constraints. *Neural Computation*. [aWAP, JVS]
- Stone, J. & Bray, A. (1995) A learning rule for extracting spatio-temporal invariances. *Network* 6:429–36. [aWAP]
- Stryker, M. P. (1988) Group report. Principles of cortical self-organization. In: *Neurobiology of neocortex*, ed. P. Rakic & W. Singer. Wiley. [aWAP]
- Stuart, G. J. & Sakmann, B. (1994) Action propagation of somatic action potentials into neocortical pyramidal cell dendrites. *Nature* 367:69–72. [rWAP]
- Sur, M., Garraghty, P. & Roe, A. (1988) Experimentally induced visual projections into auditory thalamus and cortex. *Science* 242:1437–41. [aWAP]
- Swindale, N. V. (1990) Is the cerebral cortex modular? *Trends in the Neurosciences* 12:487–92. [aWAP]
- Szenatgoathai, J. (1978) The neuron network of the cerebral cortex: A functional approach. *Proceedings of the Royal Society London (B)* 201:219–48. [JVS]
- Tass, P. & Haken, H. (1996) Synchronized oscillations in the visual cortex - a synergetic model. *Biological Cybernetics* 74:31–39. [JJW]
- Terman, D. & Wang, D. L. (1995) Global competition and local cooperation in a network of neural oscillators. *Physica D* 81:148–76. [DLW]
- Theunissen, F. & Miller, J. P. (1995) Temporal encoding in nervous systems: A rigorous definition. *Journal of Computational Neuroscience* 2:149–62. [VGH]
- Thorpe, S. J., O'Regan, K. & Pouget, A. (1989) Humans fail on XOR pattern classification problems. In: *Neural networks from models to applications*, ed. L. Personnaz & G. Dreyfus. IDSET. [aWAP]
- Tiitinen, H., Sinkkonen, J., Reinikainen, K., Alho, K., Lavikainen, J. & Näätänen, R. (1993) Selective attention enhances the auditory 40-Hz transient response in humans. *Nature* 364:59–60. [aWAP]
- Tononi, G., Sporns, O. & Edelman, G. M. (1992a) The problem of neuronal integration: Induced rhythms and short-term correlations. In: *Induced rhythms in the brain*, ed. E. Basar & T. H. Bullock. Birkhauser. [aWAP]
- Tononi, G., Sporns, O. & Edelman, G. M. (1992b) Reentry and the problem of integrating multiple cortical areas: Simulation of dynamic integration in the visual system. *Cerebral Cortex* 2:310–35. [aWAP, GT]
- Tononi, G., Sporns, O. & Edelman, G. M. (1994) A measure for brain complexity: Relating functional segregation and integration in the nervous system. *Proceedings of the National Academy of Sciences USA* 91:5033–37. [aWAP, GT]
- Tononi, G., Sporns, O. & Edelman, G. M. (1996) A measure for selective matching of signals by the brain. *Proceedings of the National Academy of Sciences USA* 93:3422–27. [aWAP, GT]
- Tooby, J. & Cosmides, L. (1995) Mapping the evolved functional organization of mind and brain. In: *The cognitive neurosciences*, ed. M. S. Gazzaniga. MIT Press. [aWAP]
- Tovee, M. J. & Rolls, E. T. (1992a) Oscillatory activity is not evident in the primate temporal visual cortex with static stimuli. *NeuroReport* 3:369–72. [VGH]
- (1992b) The functional nature of neuronal oscillations. *Trends in the Neurosciences* 15:387. [aWAP]
- Toyama, K., Kimura, M. & Tanaka, K. (1981) Organization of cat visual cortex as investigated by cross-correlation techniques. *Journal of Neurophysiology* 46:202–14. [aWAP]
- Traub, R. D., Whittington, M. A., Stanford, I. M. & Jeffreys, J. G. R. (1996) A mechanism for generation of long-range synchronous fast oscillations in the cortex. *Nature* 383:621–24. [rWAP, JJW]
- Treves, A. & Panzeri, S. (1995) The upward bias in measures of information derived from limited data samples. *Neural Computation* 7:399–407. [AT]
- Treves, A., Skaggs, W. E. & Barnes, C. A. (1996) How much of the hippocampus can be explained by functional constraints? *Hippocampus* 6:666–74. [AT]
- Troscianko, T. (1994) Contribution of colour to the motion aftereffect and motion perception. *Perception* 23:1221–31. [aWAP]
- Troscianko, T., Davidoff, J., Humphreys, G., Landis, T., Fahle, M., Greenlee, M., Brugger, P. & Phillips, W. A. (1996) Human colour discrimination based on a non-parvocellular pathway. *Current Biology* 6:200–10. [aWAP]
- Troscianko, T., Landis, T. & Phillips, W. A. (1993) Chromatic discrimination in cerebral achromatopsia: Additional evidence favouring magno-based perception, and a neural-net model. *Perception* 22(suppl):8–9. [aWAP]
- Troscianko, T., Prince, C., Fahle, M. & Regard, M. (1995) The uses of high-temporal-frequency chromatic information in visual perception. *Perception* 24(suppl):59. [aWAP]
- Ts'o, D., Gilbert, C. & Weisel, T. N. (1986) Relationship between horizontal interactions and functional architecture in cat striate cortex as revealed by cross-correlation analysis. *Journal of Neuroscience* 6:1160–70. [aWAP]
- Ullman, S. (1994) Sequence seeking and counterstreams: A model for bidirectional information flow in the cortex. In: *Large-scale neuronal theories of brain function*, ed. C. Koch & J. L. Davis. MIT Press. [aWAP]
- Usher, M. & Niebur, E. (1996) Modeling the temporal dynamics of IT neurons in visual search: A mechanism for top-down selective attention. *Journal of Cognitive Neuroscience* 8:311–27. [rWAP]
- Van Essen, D. C., Anderson, C. H. & Olshausen, B. A. (1994) Dynamic routing strategies in sensory, motor and cognitive processing. In: *Large-scale neuronal theories of brain function*, ed. C. Koch & J. L. Davis. MIT Press. [aWAP]
- Van Essen, D. C. & Maunsell, J. H. R. (1983) Hierarchical organization and functional streams in the visual cortex. *TINS* 6:370–75. [PK]
- van Rotterdam, A., Lopes da Silva, F. H., van der Ende, J., Viergever, M. A. & Hermans, A. J. (1982) A model of the spatial-temporal characteristics of the alpha rhythm. *Bulletin of Mathematical Biology* 44:283–305. [PLN]
- Volgushev, M., Chistiakova, M. & Singer, W. (submitted) Modification of discharge patterns of neocortical neurons by induced oscillations of the membrane potential. [rWAP]

- Volgushev, M., Voronin, L. L., Chistiakova, M. & Singer, W. (1997) Relations between long-term synaptic modifications and paired-pulse interactions in the rat neocortex. *European Journal of Neuroscience* (in press). [rWAP]
- von der Malsburg, C. (1981) *The correlation theory of brain function (Internal report 81-2)*. Department of Neurobiology, Max-Planck-Institute for Biophysical Chemistry, Gottingen, Germany. [aWAP, DLW]
- (1988) Pattern recognition by labeled graph matching. *Neural Networks* 1:141–48. [aWAP]
- von der Malsburg, C. & Schneider, W. (1986) A neural cocktail-party processor. *Biological Cybernetics* 54:29–40. [aWAP]
- von der Malsburg, C. & Singer, W. (1988) Principles of cortical network organization. In: *Neurobiology of neocortex*, ed. P. Rakic & W. Singer. Wiley. [rWAP]
- Von Stein, A., Chiang, C. & König, P. (1996) Expectancy-driven synchronization between primary visual cortex and parietal cortex in cats. *Society for Neuroscience Abstracts* 22:255.3. [PK]
- Wang, D., Buhmann, J. & von der Malsburg, C. (1990) Pattern segmentation in associative memory. *Neural Computation* 2:94–106. [aWAP]
- Wang, D. L. (1993a) Modeling global synchrony in the visual cortex by locally coupled neural oscillators. *Proceedings of the Fifteenth Annual Conference of the Cognitive Science Society* 1058–63. [DLW]
- (1993b) Pattern recognition: Neural networks in perspective. *IEEE Expert* 8:52–60, August. [DLW]
- (1995) Emergent synchrony in locally coupled neural oscillators. *IEEE Transactions on Neural Networks* 6:941–48. [DLW]
- Wang, D. L. & Terman, D. (1995) Locally excitatory globally inhibitory oscillator networks. *IEEE Transactions on Neural Networks* 6:283–86. [DLW]
- (1997) Image segmentation based on oscillatory correlation. *Neural Computation* 9:805–36. [DLW]
- Warren, R. M. (1970) Perceptual restoration of missing speech sounds. *Science* 167:392–93. [aWAP]
- Weisstein, N. & Harris, C. S. (1974) Visual detection of line segments: An object-superiority effect. *Science* 186:752–55. [aWAP]
- Weliky, M., Kandler, K., Fitzpatrick, D. & Katz, L. C. (1995) Patterns of excitation and inhibition evoked by horizontal connections in visual cortex share a common relationship to orientation columns. *Neuron* 15:541–52. [DS]
- Wennekers, T. & Palm, G. (1997) On the relation between neural modelling and experimental neuroscience. *Theory in Biosciences* (in press). [GP]
- Wever, R. (1965) Pendulum vs. relaxation oscillation. In: *Circadian clocks. Proceedings of the Feldafing Summer School*, ed. J. Aschoff. North Holland. [VGH]
- White, E. L. (1989) *Cortical circuits: Synaptic organization of the cerebral cortex: Structure, function and theory*. Birkhauser. [aWAP]
- Wilson, H. R. & Cowan, J. D. (1973) A mathematical theory of the functional dynamics of cortical and thalamic nervous tissue. *Kybernetik* 13:55–80. [PLN]
- Wilson, M. A. & Bower, J. M. (1988) A computer simulation of olfactory cortex with functional implications for storage and retrieval of olfactory information. In: *Neural information processing systems*, ed. D. Anderson. AIP Press. [JMB]
- (1991) A computer simulation of oscillatory behavior in primary visual cerebral cortex. *Neural Computation* 3:498–509. [JMB]
- (1992) Cortical oscillations and temporal interactions in a computer simulation of piriform cortex. *Journal of Neurophysiology* 67:981–95. [JMB]
- Wimmer, H. & Perner, J. (1983) Beliefs about beliefs: Representing and constraining function of wrong beliefs in children's understanding of deception. *Cognition* 13:103–28. [aWAP]
- Wright, J. J. (1997) EEG simulation: Variation of spectral envelope, pulse synchrony and approx. 40 Hz oscillation. *Biological Cybernetics* (in press). [JJW]
- Wright, J. J. & Liley, D. T. J. (1996) Dynamics of the brain at global and microscopic scales: Neural networks and the EEG. *Behavioral and Brain Sciences* 19:285–295. [PLN, JJW]
- Wurtz, R. P. (1995) *Multilayer dynamic link networks for establishing image point correspondences and visual object recognition*, volume 41 of *Reihe Physik*. Verlag Harri Deutsch. [RPW]
- (1997) Object recognition robust under translations, deformations and changes in background. *IEEE Transactions on Pattern Recognition and Machine Intelligence* (in press). [RPW]
- Yamaguchi, Y. & Shimizu, H. (1994) Pattern-recognition with figure-ground separation by generation of coherent oscillations. *Neural Networks* 3:49–63. [aWAP]
- Yan, J. & Suga, N. (1996) Corticofugal modulation of time-domain processing of biosonar information in bats. *Science* 273:1100–03. [rWAP]
- Young, M. P., Tanaka, K. & Yamane, S. (1992) On oscillating neuronal responses in the visual cortex of the monkey. *Journal of Neurophysiology* 67:1464–74. [VGH, aWAP]
- Zaitchik, D. (1990) When representations conflict with reality. *Cognition* 35:41–68. [aWAP]
- Zemel, R. S. & Hinton, G. E. (1991) Discovering viewpoint-invariant relationships that characterize objects. In: *Advances in neural information processing systems* 3:299–305. Morgan Kaufmann. [aWAP]
- Zhadin, M. N. (1984) Rhythmic processes in cerebral cortex. *Journal of Theoretical Biology* 108:565–95. [PLN]
- Zipsper, D. & Andersen, R. A. (1988) A back-propagation programmed network that simulates response properties of a subset of posterior parietal neurons. *Nature* 331:679–84. [DF]
- Zipsper, K., Lamme, V. A. F. & Schiller, P. H. (1996) Contextual modulation in primary visual cortex. *The Journal of Neuroscience* 15:7376–89. [rWAP]