ISSN 1539-2791

Volume 2 • Number 2 • 2004

Humanalournals.com

# Neuroinformatics

**Editors** 

Giorgio A. Ascoli Erik De Schutter David N. Kennedy

### **Special Issue on Functional Connectivity**

Guest Editors: Ed Bullmore, Lee Harrison, Lucy Lee, Andrea Mechelli, and Karl Friston

Indexed and Abstracted in: Medline/Pubmed/Index Medicus Science Citation Index®

Neuroinformatics Online www.NeuroinformaticsONLINE.com

MANA

*Neuroinformatics* Copyright ©Humana Press Inc. All rights of any nature whatsoever are reserved. ISSN 1539-2791/04/227–238/\$25.00

**Review Article** 

## Inferential Constraint Sets in the Organization of Visual Expectation

Steven L. Bressler

Center for Complex Systems and Brain Sciences, Florida Atlantic University, 777 Glades Road, Boca Raton FL 33431; E-mail: bressler@fau.edu

#### Abstract

Many lines of evidence indicate that considering visual perception as a passive, stimulusdriven, feedforward decoding process is no longer tenable. Visual perception naturally occurs within the context of an integrated array of ongoing cognitive processes involving memory, perception in other modalities, and motor control. In many situations, these processes allow expectations to be formed for likely visual events. This article explores the idea that the formation of visual expectations involves the active organization of visual cortical areas, providing a framework of contextual information within which expected events are interpreted. Retinal inputs are treated as constraints that feed into a complex system of interacting visual cortical areas and thalamic nuclei, which are concurrently imposing constraints on one another. Although the nature of expectational organization in the visual

cortex is not well-understood, a reasonable hypothesis is that expectation involves the mutual constraint of spatiotemporal activity patterns in multiple visual cortical areas. In this scenario, expectation is instantiated by a set of activity patterns in high-level visual cortical areas that impose constraints on one another as well as on low-level areas according to the partial information that is available about expected retinal inputs. One approach to testing this proposal is through the analysis of simultaneously recorded local field potentials (LFPs) from local neuronal assemblies in multiple visual cortical areas. Analysis of LFPs by multivariate autoregressive modeling is showing promise in revealing the organization of expectation in visual cortex.

**Index Entries:** Visual perception; visual cortex; inference; expectation; context; constraint set; coherence; local field potential; Granger causality.

#### Introduction

Indeterminacy, or uncertainty, is a fundamental property of the sensory environment with which the informational systems of the brain must contend. This indeterminacy suggests that perceptual processes in the cerebral cortex are inferential rather than strictly analytical (Pouget et al., 2003; Kersten and Yuille, 2003). The visual system, for example, must detect invariant structure in retinal activity that by itself does not specify what constitutes the invariance (Bressler, 1999). Theoretical considerations imply that the visual system must actively construct the structure of visual space by a process that is inherently inferential in order to operate within an arena of everchanging contexts (Mumford, 1992; Barlow, 2001; Friston, 2002). A common view is that inferential processing in visual perception involves interactions of top-down influences, providing contextual information, with bottom-up influences from the sensorium (Desimone and Duncan, 1995; Rao, 1999; Lee, 2002; Yu and Dayan, 2002).

Young (2000) convincingly argues that the synaptic connectivity of the mammalian visual cortex is especially conducive to inferential processing. For example, geniculo-recipient neurons in cortical area V1 typically receive only around 5% of their excitatory synapses from the lateral geniculate nucleus (LGN), but 30–40% from distant cortical areas and thalamic nuclei. This strongly suggests that processing in V1 depends heavily on constraints originating in multiple other areas and that direct sensory input serves as only one constraint of many. The same is even more likely to be true for high-level visual and non-sensory-specific cortical areas (Galuske et al., 2002).

That top-down constraints play an important role in the processing of visual stimuli is supported by a number of recent studies. Pascual-Leone and Walsh (2001), for example, used transcranial magnetic stimulation to demonstrate a crucial role for top-down influences in the awareness of visual motion. Murray et al. (2002) reported that blood flow in V1 was lowered, and that in higher areas of the lateral occipital complex (LOC) was elevated, when visual scene elements were structured as compared to random, perhaps reflecting greater activation of LOC areas by structured than by random inputs. The imposition of top-down constraints by these LOC areas on V1 neurons may have eliminated irrelevant V1 activity, resulting in lowered V1 metabolic demand and the observed decrease in blood flow.

A critical problem for cognitive neuroscience is thus to understand the nature of inter-areal interactions in the visual cortex. A popular approach to this question has been to consider these interactions from the perspective of Bayesian inference: top-down effects implement contextual priors, whereas bottom-up effects provide observations, with concurrent Bayesian inference occurring at multiple levels in the visual hierarchy (Lee and Mumford, 2003). A related proposal (Bressler, 1996, 2002, 2003) is that top-down, bottom-up, and lateral influences, although different in the specific nature of their effect, all combine to constrain the local processing in individual visual cortical areas. The joint satisfaction of consistent constraints, according to coherence and incoherence relations (Thagard and Verbeurgt, 1998), leads to a unified consensual state in multiple interconnected areas.

Perceptual inference involves the expectation of impending sensory events as well as the interpretation of current events. Inferential topdown constraints from high-level to low-level visual cortical areas should be active not only in processing visual stimuli, but also in visual processing prior to stimulus presentation, when there is an expectation of stimulus properties. Kastner et al. (1999) showed elevated blood flow in area V1 and higher visual cortical areas with covert visual attention to spatial location prior to visual stimulus presentation. Elevated blood flow in higher areas may have been owing to activations related to expectation of the impending visual stimuli, and that in V1 by the top-down imposition of constraining influences from those higher areas.

The highly distributed and hierarchical anatomical organization of the visual cortex implies that the expectation of visual stimulus events should involve multiple high-level areas mutually constraining one another while also constraining activity in low-level areas. The specific visual cortical areas involved in supporting the set of active constraints in any given context will depend on the particular visual features about which information is available. For example, a constraint set would involve areas that express shape information if the situation allowed inference about impending shape features, whereas it would involve motion-related areas if only motion features could be inferred. The rest of this article is concerned with the issue of how inferential constraint sets may operate in visual expectation.

#### Structural and Functional Bases of Inferential Constraint

The anatomical structure of the visual cortex involves more than twenty-five areas with multiple bottom-up, top-down, and lateral connections between them (Felleman and Van Essen, 1991). It has been shown that visual cortical areas are organized within this structure in strongly interconnected clusters that are critical for their function (Hilgetag et al., 2000). For each area, there is a unique set of other areas, called a connection set, to which that area is preferentially connected (Bressler, 2002). The connection set of a visual cortical area represents a structural complex of connected areas at lower, higher, and equivalent hierarchical levels, with which that area interacts before and during visual stimulus processing. Statistical analysis has shown that each area's connection set is unique, suggesting a unique functional role (Passingham et al., 2002).

This section focuses on the idea that cortical processes underlying visual expectation involve the imposition of constraints between different visual cortical areas. Partial information about expected visual stimuli will be distributed across high-level visual cortical areas at different hierarchical levels, with different areas expressing different aspects of that information. An area that is specialized for processing a particular visual stimulus feature will express information relating to the impending stimulus with a degree of certainty allowed by the information available about that feature. (A putative mechanism by which certainty could be represented in a cortical area is presented later.) The expression of specialized information in one area will occur not only locally within this area but will also be transmitted to the other areas to which it has axonal projection paths.

The transmission from one (transmitter) area to another (target) area is envisioned as causing the constraint of ongoing activity in the target (Bressler, 1987a,b). Long-range projection neurons, although excitatory, may synapse on both inhibitory and excitatory target neurons. Thus, the effect of transmission from the transmitter area to its target may be a mix of feedforward excitation and inhibition, and projection neurons may both enhance and suppress ongoing activity in the target area. Projection neurons may exert effects of varying strength on the target area, depending on the number of axon branches, the strengths of their synapses on target neurons, and the relative numbers of target excitatory and inhibitory cells. Target neurons may receive influences of varying strength, depending on the number, strength, and timing of their synaptic activations. The number of active synapses at any time may in turn depend on the degree of convergence of input fibers from different projection neurons, as well as the degree of synchronicity in the pulse activity on those input fibers. The result is that, in the projection of activity from a transmitting cortical area to its target, there is likely to be spatial variation in the strength of influence in the target area. This may result in a complex spatiotemporal interaction between an imposed pattern of excitation/inhibition and ongoing activity in the target area. In general, the imposed pattern may be considered to constrain the ongoing target activity pattern.

Consider next that, during expectation of an impending visual stimulus, multiple inputs arrive in a target visual cortical area over converging projection pathways from the multiple areas of its connection set. This multiplicity of converging inputs would appear to dictate that the ongoing activity in that target area be constrained by a combination of inputs, and not by any single input. These inputs, particularly in the case of top-down inputs to low-level areas, may mainly play a modulatory role, with synapses terminating primarily in the supragranular layers where NMDA receptors, having slow dynamics, are located (Friston, 2002). These modulatory constraints could limit the range of variability of the membrane potential in the postsynaptic target neurons, thus limiting the range of response of these neurons when the stimulus is presented.

We now examine the relation of high-level visual cortical areas to their connection sets during visual expectation. Allowing that the cortex has a history of exposure to stimuli of a particular class, areas that express stimulusrelated information will have previously been activated by stimuli of that class. It is usually assumed that cell assemblies are formed in those areas by local Hebbian learning. Then, when partial information is available about the impending stimulus, local Hebbian assemblies in those areas will manifest activity patterns that are expressions of their local associative memories, and those activity patterns will critically depend on inputs from their connection sets (Miller, 1996). In other words, visual cortical areas express information about an impending visual stimulus by manifesting activity patterns that reflect their stored associative memories under constraint by inputs from their connection sets.

Figure 1 illustrates the situation in which a given area (A) receives inputs from the areas of its connection set in expectation of an impending stimulus. Some of these areas (light shading) express information about the impending stim-

ulus since they have previously been co-activated by that class of stimulus, and some (dark shading) do not. Through Hebbian learning, which is thought to involve the strengthening of synaptic connections between neurons in different areas as well as within single areas (Pulvermuller, 1999; Fuster, 2000), the synaptic connections that area A receives from these co-active connection set areas have previously been strengthened. Area A thereby receives inputs from the areas of its connection set that express stimulus-related information, and these inputs exert constraints which reinforce a consistent stimulus-related activity pattern in area A.

Furthermore, projection pathways between areas in the visual system are most often bidirectional (Felleman and Van Essen, 1991). Allowing that two areas may be both transmitter and target concurrently, the imposed patterns from each may concurrently constrain the ongoing patterns of the other (Sommer and Wennekers, 2003). Initially after the introduction of information about an impending stimulus, the activity patterns of an area (e.g., A in Fig. 1) and its connection set areas may be inconsistent. However, as the area and its connection set areas interact over projection pathways that are largely bi-directional, they impose constraints on one another. It has been postulated (Bressler, 2002, 2003) that under these conditions, cortical interactions lead to some areas (light shading in Fig. 1) relaxing into consistent patterns that satisfy coherence relations, where the term "coherence" here is used in the sense of informational consistency (Thagard and Verbeurgt, 1998). The lack of pattern consistency between other areas (dark shading in Fig. 1) may satisfy incoherence relations. The result is proposed to be a distributed set of visual areas expressing consistent activity patterns, representing in its entirety the visual cortical organization of expectation for the anticipated stimulus (Engel et al., 2001).

The prior discussion has been intentionally vague about the form in which activity patterns



Fig. 1. Schematic representation of a cortical area (A) and its connection set. The cortex is idealized as a two-dimensional sheet lacking laminar organization. Area A interacts with the areas of its connection set in expectation of an impending stimulus. For simplicity, possible connections among the connection set areas are not represented. Some areas (light shading) express stimulus-related activity patterns through mutual imposition of consistent constraints (solid arrows) satisfying coherence relations. Other areas (dark shading), although anatomically connected (dashed arrows) to area A, do not express consistent patterns, and thus satisfy incoherence relations. (Patterned after Houk and Wise, 1995.)

are manifested in the visual cortex, and the way that they may interact. This is because relatively little evidence exists on these critical questions. Those studies that have addressed these questions suggest that cortical areas sustain time-varying, spatially coherent, oscillatory activity patterns that carry information in their spatially modulated amplitude patterning (Arieli et al., 1995; Bullock et al., 1995; Miltner et al., 1999; Meador et al., 2002; Cossart et al., 2003). Here, unlike previously, the term "coherence" refers to frequency-specific phase synchronization of neural activity across the spatial extent of a cortical area. Coherence has previously been proposed as an ongoing principle for cortical operations (Singer, 1994), and has been suggested to play a role in predictive processing (Engel et al., 2001).

What part might spatial coherence play in visual expectation? I have discussed the idea that the features of an expected visual stimulus are

represented by spatial activity patterns in distributed visual cortical areas. It is generally believed that expectation may have associated with it different degrees of certainty (Yu and Dayan, 2002). Therefore, some features of the expected stimulus may have a higher degree of certainty than others. For example, information that an image of a book will appear on a currently blank display screen will likely entail more certainty about the expected shape than about the expected color. It thus appears reasonable to assume that the various representations of expected features in different visual cortical areas will have differing levels of certainty. I propose spatial coherence of activity patterns within a cortical area as the basis for its local certainty. Consider the findings of Freeman and colleagues showing that spontaneous LFP activity in sensory cortical areas is often organized as oscillatory "wave packets" (Freeman and Rogers, 2002). Spatial coherence

may be viewed as a quantity that measures the degree of spatial phase dispersion of LFP oscillations within a cortical area. Expected stimulus features about which a high level of certainty exists are represented, then, by wave packets with a high degree of spatial coherence, and those with a low level of certainty have a correspondingly low spatial coherence. An advantage of this putative mechanism would be that when a high-level transmitter area constrained the activity of a target area, as described previously, the level of spatial coherence in the transmitter would directly dictate the strength of its constraint on target neurons. The greater the degree of spatial coherence in the transmitter area, the greater would be the number of transmitter neurons providing in-phase inputs over converging projection fibers onto the dendrites of target neurons, and the stronger would be the driving effect in the target area.

The result of this would be that high-level areas which manifested wave packets having the highest spatial coherence would have the greatest influence in organizing states of expectation in the visual cortex, and those with progressively lower spatial coherence would have progressively less influence. This form of dynamic organization in the visual cortex would presumably be advantageous for stimulus recognition: top-down constraints on V1 and other low-level visual cortical areas would be dominated by expectations having the greatest certainty. For the numerous situations in which those expectations were valid, these constraints could enhance the speed and efficiency of early stimulus processing, and, when expectations were invalid, could play a role in signaling the occurrence of an unexpected stimulus.

#### Experimental Evidence for Visual Constraint Sets

Experimental investigation of the operation of constraint sets in the visual cortex ideally

would require simultaneous high-density spatiotemporal recording of neuroelectric activity from two or more visual cortical areas in subjects expecting visual stimuli. However, although the spatiotemporal patterning of spontaneous activity in visual cortical areas has been investigated (Freeman and van Dijk, 1987; Arieli et al., 1995; Barrie et al., 1996; Chiu and Weliky, 2001; Warren et al., 2001; Leopold and Logothetis, 2003), there have been no detailed studies, to my knowledge, relating simultaneous spatial activity patterns in different visual cortical areas in a well-controlled visual perceptual task. A more limited approach to the problem is to measure the phase synchronization of simultaneously recorded LFPs from different visual cortical areas in relation to visual stimulus expectation and processing. This approach is based on the idea that LFP oscillations from different cortical areas become phase synchronized when, during the course of visual task processing, those areas reach consensual informational states (Bressler and Kelso, 2001).

Recent work in my laboratory has focused on the study of simultaneously recorded LFPs from different areas in the visual cortex of macaque monkeys as they waited for the appearance of a visual stimulus in a visual pattern discrimination task (Bressler et al., 1993). The task required discrimination between two types of four-dot images in a stimulus set consisting of four such images. Since the monkeys performed the task thousands of times prior to the sessions that were analyzed, they had a high degree of certainty about the nature of the impending stimulus during the wait period. We have observed that, in this period, LFP oscillations from specific sets of sites located in different visual cortical areas are phase-synchronized in the beta-frequency range (15–30 Hz). These results were obtained by spectral analysis of LFP time series using the technique of multivariate autoregressive (MVAR) modeling (Bressler et al., 1999; Ding et al., 2000).

MVAR modeling has also been successfully applied to human electroencephalographic (Franaszczuk et al., 1985; Moller et al., 2001) and functional magnetic resonance imaging (Harrison et al., 2003) data.

What is particularly relevant is that betafrequency phase synchronization appears to support top-down inter-areal constraints in the wait period of the visual discrimination task. To study inter-areal constraints, we measured causal influences between cortical areas using Granger causality analysis. Granger (1969) defined causality in terms of a bivariate autoregressive model: one stochastic process is causal to a second if the predictability of the second's autoregressive model at a given time point is improved by including measurements from the immediate past of the first. In our application, stochastic processes are considered to generate LFP time series at different cortical sites. Two important asymmetries in this definition should be noted. First, the Granger causality measure is asymmetric in time: past time values predict the present, and not vice versa. Second, it is asymmetric in space: the stochastic process at one cortical site may causally influence that at a second site without that second site causally influencing the first. Following the lead of previous authors (Bernasconi and König, 1999; Baccala and Sameshima, 2001; Hesse et al., 2003), we implemented Geweke's (1982) decomposition of Granger causality in the frequency domain.

The interpretation of causal influence as measured by Granger causality is very much in keeping with the concept of inter-areal constraint: one area exerts an effect that limits or confines the ongoing activity in another area. However, a drawback of using Granger causality to measure inter-areal constraint is that it does not indicate whether the effect is exerted by direct or indirect pathways. Nonetheless, Granger causality analysis can be used to distinguish among four general types of relation between any pair of sites (e.g., *x* and *y*): (1) *x* is

causal to y, but y is not causal to x; (2) y is causal to x, but x is not causal to y; (3) x and y are both causal to each other; and (4) neither x nor y is causal to the other.

Initial Granger causality analysis of LFPs from the visual cortex of one monkey has provided provisional evidence for top-down interareal constraints in visual expectation. LFPs from three sites in striate cortex (V1), two in prestriate cortex, and one in inferior temporal cortex were included in the analysis, thus yielding 30 Granger causality spectra (15 pairwise combinations of these 6 sites x 2 directions per pair). Figure 2 presents four representative sets of Granger causality spectra, each from a different site pair. The two Granger causality spectra in each set represent the two directions of causal influence between the two sites of the pair. Each spectrum is displayed as a function of frequency from 5 to 50 Hz, and represents the mean Granger causality from recording sessions on four different days (with error bars representing the inter-session variability).

The two upper panels in Fig. 2 present Granger causality spectra from site pairs showing significant top-down influence in the beta frequency range. The upper left panel illustrates the relation between a striate (Striate3) and a prestriate (Prestriate2) site. It indicates that a top-down causal influence exists from the prestriate to the striate site (with a peak at 17 Hz) that is significantly larger than the bottom-up (striate-to-prestriate) influence. The upper right panel shows a top-down influence from the inferior temporal (InfTemporal) site to this same prestriate (Prestriate2) site (also with a peak at 17 Hz) that is also significantly greater than its bottom-up counterpart.

For comparison, the two lower panels in Fig. 2 present examples of site pairs having relatively little Granger causality in either direction. The lower-left panel shows that site Prestriate2 exerts a top-down influence on another striate site (Striate2) with a 17-Hz peak, but with an exceedingly low magnitude. Also

0.1 0.1 From Prestriate2 to InfTemporal ...... From Striate3 to Prestriate2 ..... Prestriate2 to Striate3 From InfTemporal to Prestriate2 From 0.08 0.08 0.06 0.06 0.04 0.04 0.02 0.02 n 0 15 20 25 30 35 <u>4</u>0 45 50 25 30 35 40 45 50 0.1 0.1 From Striate2 to Prestriate2 ...... From Prestriate1 to InfTemporal ...... From Prestriate2 to Striate2 From InfTemporal to Prestriate1 0.08 0.08 0.06 0.06 S 0.04 0.04 0.02 0.02 0 0 25 30 35 40 45 20 25 30 35 40 45 50 5 15 20 -50 5 10 15 FREQUENCY (HZ)

Fig. 2. Four representative sets of Granger causality spectra are displayed, each set from a different pair of cortical sites in one monkey. The two Granger causality spectra in each set represent the causal influences exerted in both directions between cell assemblies at the two sites. Each Granger causality spectrum is displayed as a function of frequency from 5 to 50 Hz. At each frequency, the mean Granger causality over recording sessions from four different days is displayed, with the inter-session standard deviation shown by the error bars.

for comparison, the lower right panel shows the causal influences between the InfTemporal site and another prestriate site (Prestriate1). Here, the bottom-up influence is slightly stronger than the top-down, but the influences in both directions have low magnitude and are not significantly different.

Although preliminary, our Granger causality analysis thus far suggests some tentative conclusions that may relate to inter-areal constraint in visual expectation. First, considering all pairwise combinations of striate, prestriate, and inferior temporal sites, the preponderance of significant causal influence during the wait period of the visual discrimination task was top-down. This finding suggests that expectation may largely involve the constraint of low-level areas by high-level areas expressing stimulus-specific information. Second, a pronounced path of top-down influence was observed from the inferior temporal site to a single prestriate site (Prestriate2), and from there to a single striate site (Striate3). A possible interpretation of this finding is that this striate site contained neurons having receptive fields that covered the most likely location of one of the dots in the stimulus images. Given the high level of certainty about the stimulus set in this task, the observed path of top-down influence may correspond to the constraint of those striate cortical neurons by specific prestriate and inferior temporal neurons that represented the dot features.

A third conclusion suggested by these results is that oscillations in the beta frequency range may represent a general cortical mechanism that mediates top-down causal effects. Bekisz and Wróbel (2003) have presented evidence suggesting that beta-frequency oscillations in descending cortico-geniculate pathways selectively lower the threshold of LGN neurons, thereby increasing the gain of lateral geniculate inputs to primary visual cortex. It is reasonable that similar beta-frequency-dependent facilitatory mechanisms could be used by descending cortico-cortical pathways during expectation to enhance the sensitivity of striate cortical cells even before geniculo-striate inputs arrive. This conclusion is consistent with the view of Phillips and colleagues that top-down contextual modulation can strongly affect the processing of neurons in low-level visual cortical areas by changing their sensitivity to the primary inputs they receive from their receptive fields (Phillips and Singer, 1997; Phillips, 2001).

Clearly, there is a pressing need for more extensive studies of inferential processing in visual expectation. Technologies are currently being developed to allow the simultaneous recording of LFP and unit activity from far greater numbers of electrodes than previously possible. These methods are allowing the measurement of activity patterns with greater spatial resolution from greater numbers of cortical (and subcortical) areas in alert, behaving monkeys. The MVAR methodology, including Granger causality analysis, is potentially scalable to accommodate the larger numbers of simultaneous recordings envisioned. When used in conjunction with well-controlled behavioral paradigms, these hardware and software developments should make it possible to test detailed hypotheses about the nature of inferential constraint sets in the visual and other sensory systems. It is to be hoped that the capacity for such detailed hypothesis testing will, in turn, lead to deeper understanding of the basic functionality of these systems.

#### Acknowledgments

This work has been supported by NSF grant IBN0090717, and NIMH grants MH64204 and MH42900.

#### References

- Arieli, A., Shoham, D., Hildesheim, R., and Grindvald, A. (1995) Coherent spatiotemporal patterns of ongoing activity revealed by real-time optical imaging coupled with single-unit recording in the cat visual cortex. J. Neurophysiol. 73, 2072–2093.
- Baccala, L. A. and Sameshima, K. (2001) Partial directed coherence: a new concept in neural structure determination. Biol. Cybern. 84, 463–474.
- Barrie, J. M., Freeman, W. J., and Lenhart, M. D. (1996) Spatiotemporal analysis of prepyriform, visual, auditory, and somesthetic surface EEGs in trained rabbits. J. Neurophysiol. 76, 520–539.
- Barlow, H. (2001) Redundancy reduction revisited. Network 12, 241–253.
- Bekisz, M. and Wróbel, A. (2003) Attentiondependent coupling between beta activities recorded in the cat's thalamic and cortical representations of the central visual field. Eur. J. Neurosci. 17, 421–426.
- Bernasconi, C. and König, P. (1999) On the directionality of cortical interactions studied by structural analysis of electrophysiological recordings. Biol. Cybern. 81,199–210.
- Bressler, S. L. (1987a) Functional relation of olfactory bulb and cortex: I. Spatial variation of bulbocortical interdependence. Brain Res. 409, 285–293.
- Bressler, S. L. (1987b) Functional relation of olfactory bulb and cortex: II. Model for driving of cortex by bulb. Brain Res. 409, 294–301.
- Bressler, S. L. (1996) Interareal synchronization in the visual cortex. Behav. Brain Res. 76, 37–49.
- Bressler, S. L. (1999) The dynamic manifestation of cognitive structures in the cerebral cortex, in Understanding Representation in the Cognitive Sciences, Riegler, A., Peschl, M., and von Stein, A., eds., Kluwer Academic, New York, pp. 121–126.
- Bressler, S. L. (2002) Understanding cognition through large-scale cortical networks. Curr. Dir. Psychol. Sci. 11, 58–61.
- Bressler, S. L. (2003) Cortical coordination dynamics and the disorganization syndrome in schizophrenia. Neuropsychopharm. 28, S35–S39.
- Bressler, S. L. and Kelso, J. A. (2001) Cortical coordination dynamics and cognition. Trends Cogn. Sci. 5, 26–36.
- Bressler, S. L., Coppola, R., and Nakamura, R. (1993) Episodic multiregional cortical coherence at multiple frequencies during visual task performance. Nature 366, 153–156.

- Bressler, S. L., Ding, M., and Yang, W. (1999) Investigation of cooperative cortical dynamics by multivariate autoregressive modeling of eventrelated local field potentials. Neurocomput. 26-27, 625–631.
- Bullock, T. H., McClune, M. C., Achimowicz, J. Z., Iragui-Madoz, V. J., Duckrow, R. B., and Spencer, S. S. (1995) Temporal fluctuations in coherence of brain waves. Proc. Natl. Acad. Sci. 92, 11568–11572.
- Chiu, C. and Weliky, M. (2001) Spontaneous activity in developing ferret visual cortex in vivo. J. Neurosci. 21, 8906–8914.
- Cossart, R., Aronov, D., and Yuste, R. (2003) Attractor dynamics of network UP states in the neocortex. Nature 423, 283–288.
- Desimone, R. and Duncan, J. (1995) Neural mechanisms of selective visual attention. Ann. Rev. Neurosci. 18, 193–222.
- Ding, M., Bressler, S. L., Yang, W., and Liang, H. (2000) Short-window spectral analysis of cortical eventrelated potentials by adaptive multivariate autoregressive modeling: data preprocessing, model validation, and variability assessment. Biol. Cybern. 83, 35–45.
- Engel, A. K., Fries, P., and Singer, W. (2001) Dynamic predictions: oscillations and synchrony in top-down processing. Nat. Rev. Neurosci. 2, 704–716.
- Felleman, D. J. and Van Essen, D. C. (1991) Distributed hierarchical processing in the primate cerebral cortex. Cereb. Cortex 1, 1–47.
- Franaszczuk, P. J., Blinowska, K. J., and Kowalczyk, M. (1985) The application of parametric multichannel spectral estimates in the study of electrical brain activity. Biol. Cybern. 51, 239–247.
- Freeman, W. J. and Rogers, L. J. (2002) Fine temporal resolution of analytic phase reveals episodic synchronization by state transitions in gamma EEGs. J. Neurophysiol. 87, 937–945.
- Freeman, W. J. and van Dijk, B. W. (1987) Spatial patterns of visual cortical fast EEG during conditioned reflex in a rhesus monkey. Brain Res. 422, 267–276.
- Friston, K. (2002) Functional integration and inference in the brain. Prog. Neurobiol. 68, 113–143.
- Fuster, J. M. (2000) Cortical dynamics of memory. Int. J. Psychophysiol. 35, 155–164.
- Galuske, R. A., Schmidt, K. E., Goebel, R., Lomber, S. G., and Payne, B. R. (2002) The role of feedback in shaping neural representations in cat visual cortex. Proc. Natl. Acad. Sci. 99, 17083–17088.

- Geweke, J. (1982) Measurement of linear dependence and feedback between multiple time series. J. Am. Stat. Assoc. 77, 304–324.
- Granger, C. W. J. (1969) Investigating causal relations by econometric models and cross-spectral methods. Econometrica 37, 424–438.
- Harrison, L., Penny, W. D., and Friston, K. (2003) Multivariate autoregressive modeling of fMRI time series. Neuroimage 19, 1477–1491.
- Hesse, W., Moller, E., Arnold, M., and Schack, B. (2003) The use of time-variant EEG Granger causality for inspecting directed interdependencies of neural assemblies. J. Neurosci. Methods 124, 27–44.
- Hilgetag, C. C., Burns, G. A., O'Neill, M. A., Scannell, J. W., and Young, M. P. (2000) Anatomical connectivity defines the organization of clusters of cortical areas in the macaque monkey and the cat. Philos. Trans. R. Soc. Lond. B Biol. Sci. 355, 91–110.
- Houk, J. C. and Wise, S. P. (1995) Distributed modular architectures linking basal ganglia, cerebellum, and cerebral cortex: their role in planning and controlling action. Cereb. Cortex 5, 95–110.
- Kastner, S., Pinsk, M. A., De Weerd, P., Desimone, R., and Ungerleider, L. G. (1999) Increased activity in human visual cortex during directed attention in the absence of visual stimulation. Neuron 22, 751–761.
- Kersten, D. and Yuille, A. (2003) Bayesian models of object perception. Curr. Opin. Neurobiol. 13, 150–158.
- Lee, T. S. (2002) Top-down influence in early visual processing: a Bayesian perspective. Physiol. Behav. 77, 645–650.
- Lee, T. S. and Mumford, D. (2003) Hierarchical Bayesian inference in the visual cortex. J. Opt. Soc. Am. A Opt. Image Sci. Vis. 20, 1434–1448.
- Leopold, D. A. and Logothetis, N. K. (2003) Spatial patterns of spontaneous local field activity in the monkey visual cortex. Rev. Neurosci. 14, 195–205.
- Meador, K. J., Ray, P. G., Echauz, J. R., Loring, D. W., and Vachtsevanos, G. J. (2002) Gamma coherence and conscious perception. Neurology 59, 847–854.
- Miller, R. (1996) Neural assemblies and laminar interactions in the cerebral cortex. Biol. Cybern. 75, 253–261.
- Miltner, W. H., Braun, C., Arnold, M., Witte, H., and Taub, E. (1999) Coherence of gamma-band EEG

activity as a basis for associative learning. Nature 397, 434–436.

- Moller, E., Schack, B., Arnold, M., and Witte, H. (2001) Instantaneous multivariate EEG coherence analysis by means of adaptive high-dimensional autoregressive models. J. Neurosci. Methods 105, 143–158.
- Mumford, D. (1992) On the computational architecture of the neocortex. II. The role of corticocortical loops. Biol. Cybern. 66, 241–251.
- Murray, S. O., Kersten, D., Olshausen, B. A., Schrater, P., and Woods, D. L. (2002) Shape perception reduces activity in human primary visual cortex. Proc. Natl. Acad. Sci. 99, 15164–15169.
- Pascual-Leone, A. and Walsh, V. (2001) Fast backprojections from the motion to the primary visual area necessary for visual awareness. Science 292, 510–512.
- Passingham, R. E., Stephan, K. E., and Kotter, R. (2002) The anatomical basis of functional localization in the cortex. Nat. Rev. Neurosci. 3, 606–616.
- Phillips, W. A. (2001) Contextual modulation and dynamic grouping in perception. Trends Cogn. Sci. 5, 95–97.
- Phillips, W. A. and Singer, W. (1997) In search of common foundations for cortical computation. Behav. Brain Sci. 20, 657-683; discussion 683–722.

- Pouget, A., Dayan, P., and Zemel, R. S. (2003) Inference and computation with population codes. Ann. Rev. Neurosci. 26, 381–410.
- Pulvermuller, F. (1999) Words in the brain's language. Behav. Brain Sci. 22, 253-279; discussion 280–336.
- Rao, R. P. (1999) An optimal estimation approach to visual perception and learning. Vision Res. 39, 1963–1989.
- Singer, W. (1994) Coherence as an organizing principle of cortical functions. Int. Rev. Neurobiol. 37, 153–183.
- Sommer, F. T. and Wennekers, T. (2003) Models of distributed associative memory networks in the brain. Theory Biosci. 122, 55–69.
- Thagard, P. and Verbeurgt, K. (1998) Coherence as constraint satisfaction. Cognit. Sci. 22, 1–24.
- Warren, D. J., Fernandez, E., and Normann, R. A. (2001) High-resolution two-dimensional spatial mapping of cat striate cortex using a 100-microelectrode array. Neuroscience 105, 19–31.
- Young, M. P. (2000) The architecture of visual cortex and inferential processes in vision. Spat. Vis. 13, 137–146.
- Yu, A. J. and Dayan, P. (2002) Acetylcholine in cortical inference. Neural Netw. 15, 719–730.