LETTER Communicated by William Lytton

Relating Macroscopic Measures of Brain Activity to Fast, Dynamic Neuronal Interactions

D. Chawla
E. D. Lumer
K. J. Friston
Wellcome Department of Cognitive Neurology, Institute of Neurology, Queen Square, London WC1N 3BG, U.K.

In this article we used biologically plausible simulations of coupled neuronal populations to address the relationship between phasic and fast coherent neuronal interactions and macroscopic measures of activity that are integrated over time, such as the BOLD response in functional magnetic resonance imaging. Event-related, dynamic correlations were assessed using joint peristimulus time histograms and, in particular, the mutual information between stimulus-induced transients in two populations. This mutual information can be considered as an index of functional connectivity. Our simulations showed that functional connectivity or dynamic integration between two populations increases with mean background activity and stimulus-related rate modulation. Furthermore, as the background activity increases, the populations become increasingly sensitive to the intensity of the stimulus in terms of a predisposition to transient phase locking. This reflects an interaction between background activity and stimulus intensity in producing dynamic correlations, in that background activity augments stimulus-induced coherence modulation. This is interesting from a computational perspective because background activity establishes a context that may have a profound effect on event-related interactions or functional connectivity between neuronal populations. Finally, total firing rates, which subsume both background activity and stimulus-related rate modulation, were almost linearly related to the expression of dynamic correlations over large ranges of activities. These observations show that under the assumptions implicit in our model, rate-specific metrics based on rate or coherence modulation may be different perspectives on the same underlying dynamics. This suggests that activity (averaged over all peristimulus times), as measured in neuroimaging, may be tightly coupled to the expression of dynamic correlations.

1 Introduction

Previously (Chawla, Lumer, & Friston, 1998), we found, using computer simulations of coupled neuronal populations, that mean activity and syn-

chronization were tightly coupled during relatively steady-state dynamics. This allowed us to make inferences about the degree of phase locking or synchronization among, or within, neuronal populations given macroscopic measures of activity such as those provided by neuroimaging. This article is about the relationship between fast dynamic interactions among neurons, as characterized by multiunit electrode recordings of separable spike trains, and measures of neural activity that are integrated over time (e.g., functional neuroimaging). In particular we address the question, Can anything be inferred about fast coherent or phasic interactions based on averaged macroscopic observations of cortical activity? This question is important because a definitive answer would point to ways in which electrophysiological findings (in basic neuroscience) might inform functional neuroimaging studies that employ a train of stimulus or task events to detect changes in time-integrated activity.

Our basic hypothesis is that fast, dynamic interactions between two neuronal populations are a strong function of their background activity. This hypothesis derives from a series of compelling computational studies (e.g., Boven & Aertsen, 1990; Aertsen & Preissl, 1991; Aertsen, Erb, & Palm, 1994). In other words, the dynamic coupling between two populations, reflected in changes in their coherent activity over a timescale of milliseconds, cannot be separated from the context in which these interactions occur. This context is shaped by the population dynamics expressed over extended periods of time and, in particular, the overall level of activity. This is based on the commonsense observation that the responsiveness of one unit, to the presynaptic input of another distant unit, will depend on postsynaptic depolarization extant at the time the presynaptic input arrives. In a previous modeling study, using relatively steady-state dynamics (i.e., in the absence of induced transients), we showed that the mean firing rate and average phase locking between two populations were tightly coupled in all regions of the model's parameter space. There could therefore be a link between mean activity and the emergence of dynamic correlations over a timescale of milliseconds. Previous modeling work has shown it to be the case that functional and effective connectivity vary strongly with background population activity (Boven & Aertsen, 1990; Aertsen & Preissl, 1991; Aertsen et al., 1994). In this article, we pursue this same question but with a more biologically motivated neuronal model, which uses Hodgkin-Huxley-like dynamics and a more refined analysis of dynamic correlations that is statistically grounded and uses mutual information.

We expected that the emergence of phasic coherent interactions between two populations is both facilitated by and results in high mean population activity, suggesting that high background population activity levels may be a necessary condition for the emergence of faster interactions. In order to examine this, we measured the short-term correlation structures between two simulated time series as characterized by the joint peristimulus time his-

togram (J-PSTH). The advantages of this characterization include a proper assessment of phasic and stochastic interactions over peristimulus time, where these interactions are referred to a stimulus or behavioral event. Using simulations, we show that the expression of dynamic correlations is a strong function of the mean activity (averaged over time) extant in the two populations at the time that these interactions are expressed. Furthermore, we show an interaction between background and evoked firing-rate changes that is mediated by activity-dependent changes in functional coupling.

Although previous work has established that the effective connectivity among neurons is sensitive to mean levels of population activity, the specific issue we wanted to address in this work is how this activity-dependent change in functional coupling would be expressed in terms of integrated firing rates. This is important from the point of view of neuroimaging, where only time-integrated measures of activity are available. These averages include a number of components, including the background activity and stimulus-related rate modulation. The latter component may be a strong function of the effective connectivity within and among neuronal populations and, consequently, the background activity itself. The interaction between background and evoked rate modulation is therefore an important phenomenon when trying to interpret responses observed with functional neuroimaging. For example, consider the cortical responses to a train of stimuli measured when the subject was attending and not attending to these stimuli. Increased time-integrated responses may be due to attentional modulation of background activity, increased stimulus-related rate modulation, or both. Demonstrating an obligatory increase in rate modulation with background activity in neuronal stimulations would greatly simplify the interpretation of imaging results because it would suggest that both mechanisms were being expressed. In order to address the interaction between background activity and stimulus intensity in modulating event-related responses, we varied both while measuring the total integrated activity and dynamic correlations.

Section 2 describes the synthetic neural model on which our simulations were based. Section 3 describes a characterization of the model dynamics in terms of short-term interactions using J-PSTHs and mutual information. Section 3 establishes a relationship between the expression of short-term interactions (dynamic correlations) and macroscopic descriptors of the population dynamics (mean activity), revealed by varying the strength of the simulated stimulus and the background tonic activity levels. On the basis of these simulations we were able to characterize the specific form for the relationship between fast dynamic interactions and mean activity in two neuronal populations and look at the interaction between background activity and stimulus intensity in mediating changes in these measures.

2 The Neural Model

Individual neurons, both excitatory and inhibitory were modeled as single-compartment units. Spike generation in these units was implemented according to the Hodgkin-Huxley formalism for the activation of sodium and potassium transmembrane channels. (Specific equations governing these channel dynamics can be found in appendix A.) In addition, synaptic channels provided fast excitation and inhibition. These synaptic influences were modeled using exponential functions, with the time constants and reversal potentials for AMPA (excitation) and $GABA_a$ (inhibition) receptor channels taken from the experimental literature (see Lumer, Edelman, & Tononi, 1997a, 1997b) for the use and justification of similar parameters to those used in the present model). Intrinsic (intra-area) connections were 20% inhibitory and 80% excitatory (Beaulieu, Kisvarday, Somogyi, & Cynader, 1992). Extrinsic (inter-area) connections were all excitatory. Transmission delays for individual connections were sampled from a noncentral gaussian distribution. Intra-area delays had a mean of 2 ms and a standard deviation of 1 ms and inter-area delays had a mean and standard deviation of 5 ms and 1 ms, respectively. We modeled two areas that were reciprocally connected. Both consisted of a hundred cells that were 90% intrinsically connected and 5% extrinsically connected. Excitatory NMDA synaptic channels were incorporated in the model (see appendix A), in addition to the excitatory AMPA and inhibitory $GABA_a$ synaptic channels. These NMDA channels were used only in the feedback connections.

Transient dynamics were evoked by providing a burst of noise to population 1. The simulated spike trains from units in both populations were averaged over the population, binned into 4 ms bins, and then smoothed using a gaussian kernel with a half-height full width of 16 ms. These spike trains were then analyzed, using the J-PSTH (Gerstein & Perkel, 1969, 1972; Gerstein, Bedenbaugh, & Aertsen, 1989; Aertsen et al., 1991). The stimulus was provided to population 1 for a duration of 35 ms at intervals of 500 ms. For each analysis, the model was run for a total of 64 seconds of simulated time. This was repeated under different levels of background noise and stimulus intensity, using either AMPA or NMDA feedback receptors.

3 Characterizing Dynamic Correlations with the J-PSTH

3.1 Peristimulus Time Histograms. The display format (see Figure 2) has three components. Plotted along each side of the square matrix (the J-PSTH) are ordinary peristimulus time (PST) histograms, which represent the stimulus time-locked, average-rate modulation. As an index of the total background activity and stimulus-induced rate modulation, we measured the integral under the PSTH of the first population. This served as our macroscopic measure of neural activity that would be measured by, for example,

functional magnetic imaging (fMRI) or positron emission tomography (PET) and represents the first dependent variable in our characterization. The second dependent variable was a measure of the dynamic correlations based on the J-PSTH or crosscorrelation matrix expressed in terms of mutual information (see below). (A detailed explanation of how to read J-PSTHs is given in appendix B.)

3.2 Coincidence Time Histogram and Cross Correlogram. The component of the analysis in the right panel of Figure 2 is the PST coincidence histogram or coincidence time histogram (CTH). This represents the stimulus time-locked average of near-coincident firing (which is simply the leading diagonal of the J-PSTH). This graph thus shows how the level of coherent firing or synchrony (plotted vertically) varies with PST (plotted horizontally). The cross correlogram, the third component, characterizes the degree of coherence averaged over all PSTs at some time lag. Because it is not sensitive to dynamic modulation of coherence, it is not used further in this article. Similarly, we do not use the CTH because, being a metric of coincident firing at near-zero time lags, it is an impoverished metric of dynamic correlations that could be expressed at nonzero time lags.

3.3 A Mutual Information Measure of Dynamic Correlations. We were interested in how stimulus-induced dynamic correlations varied as a function of background noise, stimulus strength, and the interaction between these two factors. As a measure of the dynamic correlations, induced between our two simulated populations, we used the mutual information between the stimulus-induced transients having corrected for mean rate modulation. The calculation and interpretation of mutual information are described in appendix C.

In what follows we examine the way in which the mutual information or functional connectivity changes with integrated firing rate. We did this by manipulating the strength of the stimulus under different levels of background activity. This enabled us not only to assess the effects of changing background activity and stimulus strength on mutual information (and integrated rate) but also to characterize any interaction between these two manipulations. The background noise levels were characterized in terms of the average depolarization produced: −65.9, −63.4, −63.2, and −61.2 mV. These values were calculated after applying a given noise level to both populations, running the simulation for 64 seconds, and computing the mean membrane potential over units and time. The stimulus intensities used were 10, 25, 50, 75, 125, 150, 175, 200, 225, and 250 Hz. Noise level and stimulus intensity represent our two independent variables that were expected to produce changes in the two dependent variables (integrated rate and mutual information).

4 Results

4.1 Relationship Between Dynamic Correlations and Integrated Rate. We found that increases in either background noise or the strength of the stimulus were universally associated with increases in both integrated rate and mutual information. Furthermore, both dependent variables (rate and mutual information) were highly coupled in an almost linear fashion. Figure 1b shows plots of mutual information against integrated rate for three different levels of background noise, demonstrating the coupling among these measures, irrespective of how different levels of either were elicited. Because of this tight relationship, we focused on how the independent manipulations of background noise and stimulus intensity affect mutual information (equivalent effects were observed on integrated rate). Two examples of the mean time course of activity (local field potential, LFP) in population 1 under two levels of background activity are illustrated in Figure 1a.

4.2 Effect of Background Activity and Stimulus Intensity on Dynamic Correlations. Figure 1c shows mutual information as a function of stimulus intensity for the three levels of background activity. As background noise increased, the gradient of the mutual information versus stimulus intensity plot also increased. A formal test for the differences in regression slopes confirmed the significance of this effect (t-statistic = 2.2, residual degrees of freedom = 53, p-value = 0.016 for the average increase from low- to high-background activities over NMDA and AMPA simulations using multiple regression and the appropriate contrast). This is clear evidence of an interaction between tonic background activity and stimulus-induced rate modulation in the genesis of dynamic correlations. In other words, high-background activity increased the sensitivity of evoked dynamic correlations to increased stimulus intensity. This is demonstrated more clearly below. These phenomena were evident irrespective of whether we used NMDA (upper panels) or AMPA-like (lower panels) feedback receptors.

Figure 1: *Facing page.* (a) Mean membrane potential for one stimulus strength at two of the background noise levels. This graph shows the mean membrane potential over three interstimulus intervals (1500 ms). (b) Plots of stimulus-induced mutual information between the two populations against integrated rate (as indexed by the integral under the PSTH) for three different levels of background activity. The stimulus intensity was varied through 10, 25, 50, 75, 125, 150, 175, 200, 225, and 250 Hz. The upper panel shows the results when the model was implemented with NMDA feedback receptors and the lower panel with AMPA feedback receptors. (c) The same as in *b*, but now mutual information is plotted against the stimulus intensities used. In *c*, the regression slopes of mutual information on stimulus intensity are plotted.

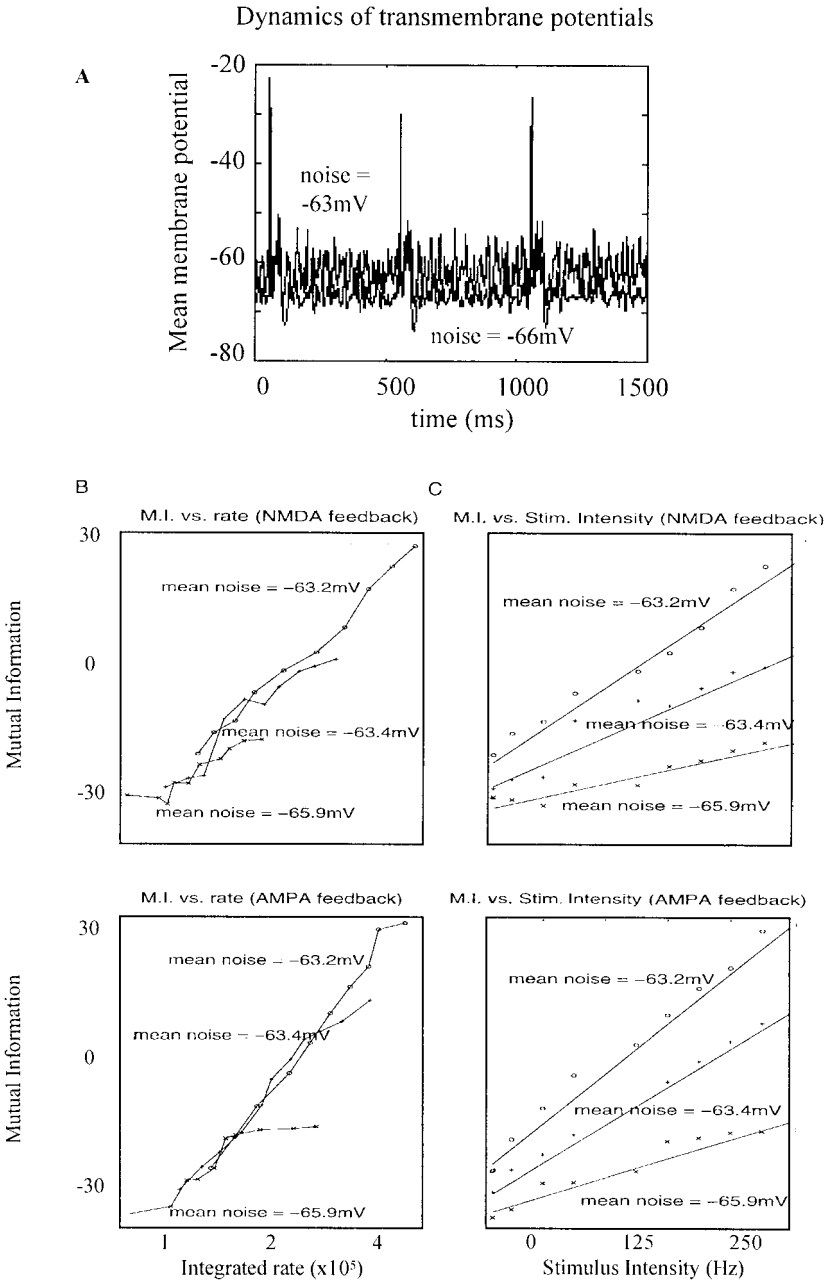
Dynamics of transmembrane potentials
A
Mean membrane potential
-20
-40
-60
-80
noise = -63mV
noise = -66mV
0
500
1000
1500
time (ms)
B
C
M.I. vs. rate (NMDA feedback)
M.I. vs. Stim. Intensity (NMDA feedback)
Mutual Information
30
0
-30
mean noise = −63.2mV
mean noise = −63.4mV
mean noise = −65.9mV
M.I. vs. rate (AMPA feedback)
M.I. vs. Stim. Intensity (AMPA feedback)
1
2
4
Integrated rate (x10⁵)
0
125
250
Stimulus Intensity (Hz)

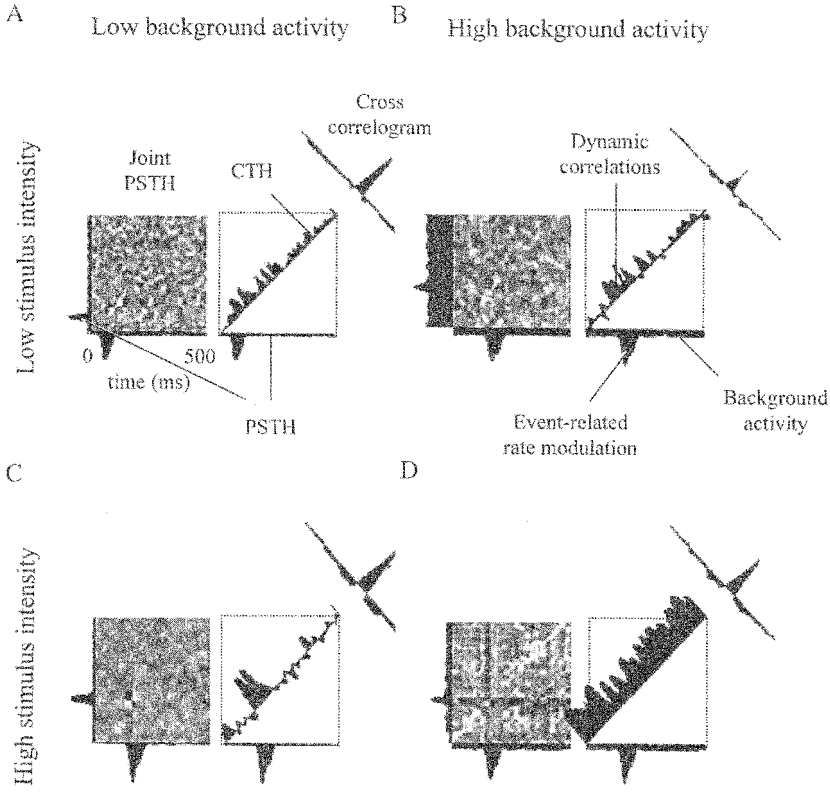

Figure 2: (a) J-PSTH for the simulated populations at the lowest background noise level (see Figure 1) and similarly with a very weak stimulus applied every 500 ms for 35 ms. (b) J-PSTH at the highest background noise level (the same as in Figure 3) and with the low intensity stimulus as in Figure 2a. (c) J-PSTH for the neuronal populations at the low background activity as in Figure 2a and with a high-intensity stimulus. (d) J-PSTH at a high background activity level as in Figure 2b and with the high-intensity stimulus as in Figure 2c.

4.3 Examples of These Effects Demonstrated with J-PSTHs. Figure 2 presents J-PSTHs between the two populations at two different levels of background activity and with two different stimulus intensities. It can be seen that when the stimulus was very weak and the background noise was low, the presence of the stimulus had almost no effect on the synchronous interactions between the populations because it was not strong enough to enable the populations to entrain each other to any extent (see CTH of Figure 2a). However, when the background noise was increased, this same stimulus had a definite effect on the dynamic correlations, as can be seen

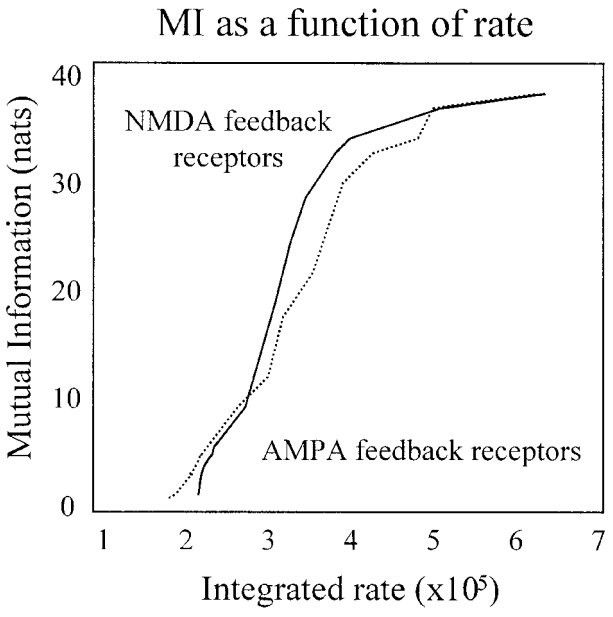

Figure 3: Mutual information plotted against integrated rate for the highest noise level, as depicted in Figures 2b and 2d, under both AMPA and NMDA feedback receptors. The stimulus intensity is varied through the same values as in Figure 1.

in Figure 2b (the background activity level in Figure 2b is higher than any of the background levels in Figure 1). At the low background activity level as in Figure 2a, when the stimulus was very strong, extremely significant dynamic correlations occurred (see Figure 2c), in contrast to when the stimulus was weak and induced minimal dynamic correlations (see Figure 2a). At the high-background noise level as in Figure 2b, when the stimulus was very strong as in Figure 2c, the synchronization induced never died away, and high levels of synchrony were maintained (see Figure 2d).

Figure 3 shows mutual information as a function of integrated rate at the very high background activity level as in Figures 2b and 2d. This shows that at such high background levels, the plot of mutual information versus integrated rate eventually levels off, at which point increasing the stimulus intensity will no longer facilitate an increase in the mutual information between the two populations. This demonstrates nicely the saturation phenomenon as seen in Figure 2d. Figure 2d showed that with a very high background noise, the stimulus intensity may reach a level at which the synchronization induced never dies away. Increasing the stimulus intensity further then has little or no effect.

5 Discussion

Several lines of evidence support our findings that a systematic relationship between fast dynamic interactions (measured here in terms of mutual information or functional connectivity) and macroscopic measures exist. Aertsen and Preissl (1991) investigated the behavior of artificial networks, analytically and using simulations. They concluded that short-term effective connectivity varies strongly with, or is modulated by, pool activity. Pool activity is defined as the product of the number of neurons and their mean firing rate. Effective connectivity is defined as the influence one neural system exerts over another at a synaptic level. The mechanism is simple: the efficacy of subthreshold excitatory postsynaptic potentials (EPSPs) in establishing dynamic interactions is a function of postsynaptic depolarization, which in turn depends on the tonic background activity. This idea can be elucidated in the following way. If the network activity is very low, the inputs to a single neuron (say, neuron j) will cause only a very subthreshold EPSP in that neuron. If some presynaptic neuron (say neuron i) fires so that it provides input to neuron j, this input will be insufficient to cause neuron j to fire. However, if the pool activity is high enough to maintain a slightly subthreshold transmembrane potential in neuron j, then an input from neuron i to neuron j is more likely to push the membrane potential of neuron j over the reversal threshold and elicit an action potential.

In previous work, we demonstrated that sustained synchrony shows a monotonic relationship with mean activity (Chawla et al., 1998). As mean activity in the network increases, the mean instantaneous membrane time constants decrease, giving rise to a higher level of synchrony. The decrease in time constants is a natural consequence of conjointly increasing membrane conductances through excitatory and inhibitory channels at high levels of activity. Hence, as activity level increases, smaller membrane time constants increase the synchronous gain in the network, that is, individual neurons became more sensitive to temporal coincidences in their synaptic inputs, responding with a higher firing rate to synchronous rather than asynchronous inputs. Therefore, in the event-related context, as background noise increases, the network should become more prone to stimulus-induced synchronous transients. This is reflected in both the time that the poststimulus synchronization endures and a progressive increase in the mutual information versus stimulus intensity regression slope. The latter effect constitutes an interaction and can be viewed as a stimulus-dependent effect that is context sensitive. In this instance, the context is set by the tonic background of activity and could be mediated through a progressive diminution of the effective membrane time constants.

Our findings provide the basis for two fairly important conclusions. The first is that increasing the tonic or background activity can potentiate the transient or dynamic correlations induced by a salient stimulus or behavioral event. This simple phenomenon might provide a useful mechanism

in the brain for exerting control over functional integration between neuronal populations in a context-dependent fashion. For example, attentional modulation of background activity in distinct sensory neuronal populations could be used selectively to enhance neuronal interactions in a topographically constrained way (Frith & Friston, 1997). The fact that a higher mutual information between our simulated neuronal populations was induced under conditions of higher noise is presumably very similar to stochastic resonance (Wiesenfeld & Moss, 1995), wherein small amounts of stochastic noise facilitate nonlinear transformations, in this instance effected by interacting neuronal populations. This phenomenon is interesting because of its almost counterintuitive nature. Indeed, it might be thought that increased background noise may lead to greater difficulty in distinguishing a transient signal from noise. However, this is not necessarily the case. Mainen and Sejnowski (1995) have shown that noisier neuronal input can increase the precision of the cell spikes, thus increasing the sensitivity of the cells to their inputs. Their data suggests "a low intrinsic noise level in spike generation, which could allow cortical neurons to accurately transform synaptic input into spike sequences, supporting a possible role for spike timing in the processing of cortical information by the neocortex" (Mainen & Sejnowski, 1995).

The second conclusion has a more practical importance and pertains to the interpretation of neuroimaging studies in which only macroscopic observations of rate modulation are generally allowed. By simply demonstrating a systematic and consistent relationship between rate modulation (integrated rate over PST) and coherence modulation (the mutual information associated with dynamic correlations), one can be comfortable with the probability that functional neuroimaging is not totally insensitive to event-related fast coherent interactions of the sort mediated by dynamic correlations. This is important given the possible role that dynamic correlations may play in sensorimotor and cognitive operations (Vaadia et al., 1995). This conclusion leads to the more general point that specific metrics based on rate or coherence modulation may be different perspectives on the same underlying dynamics. In this view, synchronized, mutually entrained signals enhance overall firing levels and can be thought of as mediating an increase in the effective connectivity between the two areas. Equivalently, high levels of discharge rates increase the effective connectivity between two populations and augment the fast, synchronous exchange of signals. In this sense there is an almost circular causality in the relationship between rate and transient synchronization. Although we manipulated mean background activity in our simulations, much of the variability in integrated rates over PST can be accounted for by the dynamic correlations induced by the stimulus. In other words, a high mean level of activity facilitates transient coherent interactions, above and beyond those predicted by dynamic rate modulation itself. These dynamic correlations in turn cause a mutual entrainment of the interacting populations and augment activity levels for

a period of time. At very high levels of activity, any stimulus-evoked transient might ignite the system, leading to high levels of synchrony and mean activity that are self-maintaining (see Figures 2 and 3).

The statement that the distinction between temporal and rate coding is simply a matter of perspective is a strong one that is made with some expectation of being refuted. Although it is clearly possible that the information conveyed by the precise timing of spikes is very different from that conveyed by discharge rates, from the point of view of population dynamics, it may be the case that changes in spike timing cannot be divorced from changes in firing rate given the neuronal infrastructure employed by the brain. The point being made here is that due to the intimate relationship between the temporal patterning of presynaptic events (in terms of either phase locking as discussed in Chawla et al., 1998, or in terms of dynamic correlations, as considered in this article), and postsynaptic discharge probabilities, an increase in synchronized input inevitably will result in higher population discharge rates. The mechanisms that underlie this relationship may involve increased membrane conductances, decreased effective membrane time constants, and an increase in synchronous gain mediated by impoverished temporal integration. Put simply, under the constraints imposed by the emergent nonlinear dynamics of neuronal circuits, one cannot change the fine temporal structure of discharge patterns without changing population activity (this point is made by the results in Figure 1, which show a monotonic relationship between integrated firing rate and mutual information). If changes in one metric of neuronal dynamics, such as spike timing, are universally associated with changes in another metric, such as population activity, then the two metrics are mutually redundant and reflect different measures of the same underlying dynamics. It should be noted that these observations pertain only to population codes.

In summary, we have shown that background activity levels in simulated neuronal populations facilitate and are facilitated by the expression of stimulus-induced dynamic correlations. These findings have implications for the context-dependent aspects of stimulus-related neuronal interactions and also inform the interpretation of neuroimaging measures of neurophysiology.

Appendix A: Modeling Neuronal Dynamics

The neuronal dynamics were produced by simulations based on the equations from the Yamada, Koch, and Adams (1989) single neuron model, using the Hodgkin and Huxley formalism:

$$\begin{aligned} dV/dt = & -1/C_M\{(g_{Na}m^2h(V - V_{Na}) + g_K n^2 y(V - V_K) + g_l(V - V_l) \\ & + g_{AMPA}(V - V_{AMPA}) + g_{GABA}(V - V_{GABA})\}, \end{aligned}$$

$$dm/dt = \alpha_m(1 - m) - \beta_m m$$

Table 1: Table of the Parameter Values of the Neuronal Model.

Receptor/Channel	$g_{\text{peak(mS)}}$	τ (ms)	V_j (mV)
AMPA	0.05	3	0
$GABA_a$	0.175	7	−70
NMDA	0.01	100	0
Na^+	200		50
K^+	170		−90
Leak	1		−60

$$\mathrm{d}n/\mathrm{dt} = \alpha_{\mathrm{n}}(1-n) - \beta_n n$$

$$\mathrm{dg}_{\mathrm{AMPA}}/\mathrm{dt} = -g_{\mathrm{AMPA}}/\tau_{\mathrm{AMPA}}$$

$$\mathrm{d}h/\mathrm{dt} = \alpha_h(1-h) - \beta_h h$$

$$\mathrm{d}y/\mathrm{dt} = \alpha_y(1-y) - \beta_y y$$

$$\mathrm{dg}_{\mathrm{GABA}}/\mathrm{dt} = -g_{\mathrm{GABA}}/\tau_{\mathrm{GABA}}$$

V represents the membrane potential of the neuron, C_M represents the membrane capacitance ($1\mu F$), and g_{Na}, g_{K}, and g_{l} represent the maximum Na+ channel, K+ channel, and leakage conductances, respectively. V_{Na} represents the Na+ equilibrium potential, and similarly for V_{K} and V_{l}. m, h, n, and y are the fraction of Na+ and K+ channel gates that are open. g_{AMPA} and g_{GABA} are the conductances of the excitatory (AMPA) and inhibitory ($GABA_a$) synaptic channels, respectively. The peak conductances are constant for each receptor type and are given in Table 1. τ represents the excitatory and inhibitory decay time constants. α_{n}, β_{n}, α_{m}, β_{m}, α_{h}, β_{h}, α_{y}, β_{y} are nonnegative functions of V that model voltage-dependent rates of channel configuration transitions. The implementation of NMDA channels was based on Traub, Wong, Miles, and Michelson (1991):

$$I_{\mathrm{NMDA}} = g_{\mathrm{NMDA}}(\mathrm{t})M(\mathrm{V} - V_{\mathrm{NMDA}})$$

$$\mathrm{dg}_{\mathrm{NMDA}}/\mathrm{dt} = -g_{\mathrm{NMDA}}/\tau_2$$

$$M = 1/(1 + (Mg^{2+}/3)(\exp[-0.07(V-\xi)])$$

I_{NMDA} is the current that enters linearly into the equation for $\mathrm{d}V/\mathrm{dt}$, above. g_{NMDA} is a ligand-gated virtual conductance. M is a modulatory term that mimicks the voltage-dependent affinity of the Mg^{2+} channel pore. ξ is −10 mV and Mg^{2+} is the external concentration of Mg^{2+} often used in hippocampal slice experiments (2 mM). These and other parameters are given in Table 1.

Appendix B: How to Read J-PSTHs

The J-PSTH is a raster plot of the two PSTHs plotted against each other, and as such, coincident firings are shown along the leading diagonal of the J-PSTH. Time-lagged synchronized firings are shown as diagonal bands that are shifted relative to this diagonal. The displacement of the band is therefore a measure of the latency of phase locking. The width and structure of the band depend on the details of the coherent interactions. Direct synaptic connections from population 1 to population 2 will produce a 45 degree band of differing density lying below the principal diagonal at a distance proportional to the latency of the interaction. Connections from population 2 to population 1 will produce a similar band lying above the principal diagonal. If the interactions between the two populations are affected by the stimulus, then the diagonal band will show changes in density along its length, evidencing dynamic correlations or coherence modulation as a function of PST. Correlations due purely to rate modulation evoked by the stimulus are removed by subtracting the cross-product matrix of the individual PSTHs from the raw J-PSTH and then dividing the resulting difference matrix (bin by bin) by the cross-product matrix of the standard deviations of the PSTHs. This is mathematically the same as computing the cross-correlation matrix between the binned activities from both populations over stimulus epochs.

Appendix C: Explanation of Mutual Information

Our measure of mutual information is equivalent to testing the null hypothesis that all the elements of the J-PSTH are jointly zero. The mutual information can be thought of as the generalization of a correlation for multivariate data (in this case, the activity over different PSTs). As such, it serves as a measure of functional connectivity. The critical idea in this instance is that by predicating our measure of functional connectivity on the J-PSTH, we properly include dynamic correlations that would otherwise be missed if we simply looked at average correlations as implicit in the cross correlogram. The importance of using the J-PSTH in the context of this article is discussed fully in Aertsen et al. (1994). A simple example makes the distinction between these two approaches clear: Imagine a stimulus-induced strong, positive correlation for 100 milliseconds followed, systematically, by equally strong negative correlations for the subsequent 100 milliseconds. The average correlation over all PSTs, as measured by the cross correlogram, would be zero. However, the mutual information mediated by the dynamic correlations in the J-PSTH would be extremely significant.

The reason that the mutual information is an implicit test of the null hypothesis that the dynamic correlations are jointly zero follows from the fact that the maximum likelihood statistic for the latter test is Wilks' lambda. Under gaussian assumptions, the log of this statistic is proportional to the mutual information between the stimulus-induced transients that produce

the dynamic correlations. In this article, we restricted ourselves to using the mutual information and refer interested readers to Chatfield and Collins (1982) for a discussion of Wilks' lambda in the context of multivariate analysis of covariance (ManCova).

C.1 Mutual Information and Wilks' Lambda. The J-PSTH is given by X^TY (^{T} denotes transposition) where X is a mean corrected and normalized data matrix from unit or population 1 with a row for every stimulus epoch and a column for every PST bin. Y is the corresponding matrix from the second population. The mutual information between X and Y is given by:

$$I(X, Y) = H(X) + H(Y) - H(X \cap Y),$$

where $H(X)$ is the entropy of X, $H(Y)$ is the entropy of Y, and $H(X \cap Y)$ is the entropy of X and Y considered jointly. Under gaussian assumptions:

$$H(X) = \ln(2\pi e^n |X^TX|)/2,$$

where $|X^TX|$ is the determinant of X^TX (i.e., the autocovariance matrix of X) and n is the number of columns.

Let:

$$A = [XY]^T[XY] = \begin{matrix} \lceil X^TX & X^TY \rceil \\ \lfloor Y^TX & Y^TY \rfloor. \end{matrix}$$

then,

$$|A| = |X^TX||Y^TY - Y^TX(X^TX)^{-1}X^TY|,$$

giving,

$$I(X, Y) = \ln(|Y^TY|/|Y^TY - Y^TX(X^TX)^{-1}X^TY|).$$

An alternative (statistical) perspective, mathematically equivalent to testing for dynamic correlations, is to test the null hypothesis that $X^TY = 0$. This can be effected in the context of multivariate analysis using Wilks' maximum liklihood ratio or Wilks' lambda, $\lambda = |R|/|R_0|$, where R and R_0 are the residual sum of squares and products under the alternate and null hypotheses, respectively. In this instance, we treat the test for statistical dependencies between X and Y as a multiple regression problem under the general linear model,

$$Y = X\beta + \varepsilon,$$

where β are the regression coefficients and ε are errors with a multinormal distribution. The parameter estimates are given by:

$$\beta = (X^TX)^{-1}X^TY$$
$$Y^* = X\beta, \qquad \text{where } Y^* \text{ is the fitted data.}$$

Under the null hypothesis, $\beta = 0$ and therefore,

$$R_0 = Y^T Y.$$

Similarly, under the alternate hypothesis:

$$\begin{aligned} R &= [Y - Y^*]^T [Y - Y^*] \\ &= Y^T Y - Y^T Y^*; \end{aligned}$$

therefore:

$$\lambda = |R|/|R_0| = |Y^T Y - Y^T X (X^T X)^{-1} X^T Y|/|Y^T Y| \qquad \text{and}$$

$$I(X, Y) = -\ln(\lambda)$$

The negative log of Wilks' lambda is the mutual information.

Under the null hypothesis of no dynamic correlations, $\ln(\lambda)$ has approximately a chi-squared distribution where $-(r - 1/2)\ln(\lambda) \sim \chi^2(n^2)$ where r are the residual degrees of freedom (number of epochs minus time bins) (Chatfield & Collins, 1982). In practice the number of time bins may exceed the number of epochs. In this instance, one generally reduces the dimensionality of the data X (and Y) using singular value decomposition.

Acknowledgments

This work was supported by the Wellcome Trust.

References

Aertsen, A., Erb, M., & Palm, G. (1994). Dynamics of functional coupling in the cerebral cortex: An attempt at a model-based interpretation. *Physica D, 75*, 103–128.

Aertsen, A., & Preissl, H. (1991). Dynamics of activity and connectivity in physiological neuronal networks. In xx Schuster (Ed.) *Non linear dynamics and neuronal networks* (pp. 281–302). New York: HG VCH Publishers.

Beaulieu, C., Kisvarday, Z., Somogyi, P., & Cynader, M. (1992). Quantitative distribution of GABA-immunopositive and -immunonegative neurons and synapses in the monkey striate cortex (area 17). *Cerebral Cortex, 2*, 295–309.

Boven, K. H., & Aertsen, A. M. H. J. (1990). Dynamics of activity in neuronal networks give rise to fast modulations of functional connectivity. In *Parallel processing and neural systems and computers* (pp. 53–56). Amsterdam: Elsevier.

Chatfield, C., & Collins, A. J. (1982). *Introduction to multivariate analysis*. London: Chapman and Hall.

Chawla, D., Lumer, E. D., & Friston, K. J. (1998). The relationship between synchronization and neuronal populations and their mean activity levels. *Neural Computation, 11*, 1389–1411.

Frith, C. D., & Friston, K. J. (1996). The role of the thalamus in "top down" modulation off attention to sound. *NeuroImage, 4*, 210–215.
Gerstein, G. L., Bedenbaugh, P. & Aertsen, A. M. H. J. (1989). Neuronal assemblies. *IEEE Trans. on Biomed. Engineering, 36*, 4–14.
Gerstein, G. L., & Perkel, D. H. (1969). Simultaneously recorded trains of action potentials: Analysis and functional interpretation. *Science, 164*, 828–830.
Gerstein, G. L., & Perkel, H. (1972). Mutual temporal relationships among neuronal spike trains. *Biophysical Journal, 12*, 453–473.
Lumer, E. D., Edelman, G. M., & Tononi, G. (1997a). Neural dynamics in a model of the thalamocortical system I. Layers, loops and the emergence of fast synchronous rhythms. *Cerebral Cortex, 7*, 207–227.
Lumer, E. D., Edelman, G. M., & Tononi, G. (1997b). Neural dynamics in a model of the thalamocortical system II. The role of neural synchrony tested through perturbations of spike timing. *Cerebral Cortex, 7*, 228–236.
Mainen, Z. F., & Sejnowski, T. J. (1995). Reliability of spike timing in neocortical neurons. *Science, 258*, 1503–1506.
Traub, R. D., Wong, R. K., Miles, R., & Michelson, H. (1991). A model of a CA3 hippocampal pyramidal neuron incorporating voltage-clamp data on intrinsic conductances. *J. Neurophysiol., 66*, 635–650.
Vaadia, E., Haalman, I., Abeles, M., Bergman, H., Prut, Y., Slovin, H., & Aertsen, A. (1995). Dynamics of neuronal interactions in monkey cortex in relation to behavioural events. *Nature, 373*, 515–518.
Wiesenfeld, K., & Moss, F. (1995). Stochastic resonance and the benefits of noise: From ice ages to crayfish and SQUIDs. *Nature, 373*, 33–36.
Yamada, W. M., Koch, C., & Adams, P. R. (1989). Multiple channels and calcium dynamics. In C. Koch & I. Segev (Eds.), *Methods in neuronal modeling* (pp. 97–134). Cambridge, MA: The MIT Press.

Received August 3, 1998; accepted February 9, 2000.