



Published as: *Nat Rev Neurosci.* 2001 August ; 2(8): 539–550.

CORRELATED NEURONAL ACTIVITY AND THE FLOW OF NEURAL INFORMATION

Emilio Salinas^{*,‡} and Terrence J. Sejnowski^{*,§}

^{*} Computational Neurobiology Laboratory, Howard Hughes Medical Institute, The Salk Institute for Biological Studies, 10010 North Torrey Pines Road, La Jolla, California 92037, USA

Abstract

For years we have known that cortical neurons collectively have synchronous or oscillatory patterns of activity, the frequencies and temporal dynamics of which are associated with distinct behavioural states. Although the function of these oscillations has remained obscure, recent experimental and theoretical results indicate that correlated fluctuations might be important for cortical processes, such as attention, that control the flow of information in the brain.

Research in systems neuroscience has traditionally focused on how neurons represent the world and on the mechanisms that endow neurons with their response properties or receptive fields. However, an equally important, but less well understood, aspect of brain function is how neurons communicate. For instance, the presence of a red light in the visual field might be irrelevant if in a theatre, but crucial if about to cross a road. The neural representation of the red light might be, at some level, the same in the two situations, but this information is then redirected and prioritized in totally different ways. Little is known about how, depending on the current behavioural requirements, neural signals are routed or assessed in the nervous system. There is evidence that timing is crucial: a recent study¹ showed that whether intracortical microstimulation influences performance in a sensory discrimination task depends on the time at which the microinjected current is delivered relative to the natural stimulus onset. This indicates that even a simple discrimination paradigm is executed according to an internal schedule, such that the information provided by the sensory neurons is effectively transmitted only during a certain time window. So, the temporal dynamics of neuronal interactions seem to be important for the gating processes that control the information that goes through at a given time.

On the other hand, networks of neurons show highly complex temporal dynamics. It is well known from electroencephalographic studies² that the small voltage signals recorded from the scalp fluctuate at various frequencies, with dominant frequency components shifting according to behaviour. Slow oscillations with a strong 0.75–4-Hz component are associated with certain stages of sleep, whereas oscillations dominated by the 14–40-Hz band are typical of active, awake states^{3,4}. Direct measurement of field potentials from the cortex reveals even higher-frequency components in the 40–200-Hz range⁵. At the single-neuron level, collective oscillations in cortical neurons have been documented for several years^{6–8}. One functional interpretation is that this rhythmic behaviour is, again, related to higher-order sensory representations, but this idea continues to be hotly debated^{9,10}. The problem is not that oscillations are not there, but that linking them to behaviour has been difficult.

Correspondence to T.J.S. terry@salk.edu.

[‡]Present address: Department of Neurobiology and Anatomy, Wake Forest University School of Medicine, Medical Centre Boulevard, Winston-Salem, North Carolina 27157-1010, USA.

[§]Department of Biology, University of California at San Diego, La Jolla, California 92093, USA

Synchrony is another form of temporal relationship between neurons that has been intensely studied. As with ‘oscillations’ and ‘rhythmic activity’, the term synchrony encompasses a spectrum of neuronal behaviours with various spatial and temporal scales¹¹. Here, we will label all of these phenomena as temporally correlated activity, which describes a common feature: when two neurons are correlated, they do not fire independently of each other; when one fires, the other is more or less likely to fire. This is an extremely broad generalization, especially as the underlying mechanisms and potential functions can vary greatly, but here we discuss a broad idea, so it is better to think in general terms.

Recent theoretical and experimental findings have brought the study of temporally correlated neuronal activity into a new perspective. This has emerged by investigating two related questions. First, how is a post-synaptic neuron affected by the presence of correlated activity in its inputs? The response of a neuron depends on the rates at which excitatory and inhibitory input spikes impinge on it, but the temporal pattern of those spikes can also modulate postsynaptic activity. When and how does this happen? This is a biophysical problem. Second, what is the relationship between such temporal patterning and behaviour? It is important to understand how networks generate and react to oscillatory signals, but it is just as important to determine their function. This is a systems-level problem.

Here, we review recent findings on these two issues. First, we briefly discuss some examples of the traditional interpretation in which correlation is viewed as an additional coding dimension for building internal representations. Then we discuss some common correlation patterns and review results which show that neurons can be highly sensitive to their presence. We propose that correlations could be controlled independently of firing rate and that this would serve to regulate the flow of information rather than its meaning. Finally, we discuss several experiments in which changes in correlations have been measured and reported to be independent of changes in mean firing rate. In these cases, correlations covary with expectation, attention, response latency or rivalry — all processes that affect the transit of information but not how sensory stimuli are represented. The idea that correlations can gate the flow of neural information is a recent viewpoint that could give rise to new theoretical and experimental studies.

A short digression on coding strategies

For the sake of comparison, it is instructive to mention a few studies that are representative of the more traditional approach, in which neural correlations are investigated in terms of their potential value for stimulus representation. This list is by no means exhaustive.

The antennal lobe of insects is an interesting preparation because various manipulations are possible. Spikes in this structure are typically synchronized by 20-Hz oscillations^{12–14}. When these neurons are artificially desynchronized¹⁵, the specificity of downstream responses is strongly degraded; selectivity for different odours decreases and responses to new odours arise, even though this loss of information does not occur upstream. Further experiments¹⁶ indicate that the disruption of synchrony has a real impact on behaviour, impairing odour discrimination. Oscillations have also been observed in the mammalian olfactory bulb¹⁷, but whether they serve a similar function is unknown.

Consider two neurons with overlapping receptive fields and, therefore, a considerable degree of synchrony. Analysis of the activity of such visual neurons in the lateral geniculate nucleus has shown¹⁸ that significantly more information about the stimulus (around 20% more) can be extracted from their spike trains if the synchronous spikes are taken into account separately from the non-synchronous ones. In a similar vein, recordings from primary auditory cortex indicate¹⁹ that, when a stimulus is turned on, neurons respond by changing their firing rates and their correlations. In many cases, the firing rate modulations are transient: if the sound is

sustained, they tend to disappear. However, the evoked changes in correlation can be sustained¹⁹. So, the correlation structure can signal the presence of a stimulus in the absence of changes in firing rate. In an example that is usually associated with the BINDING PROBLEM^{7,9,10}, the receptive fields of two visual neurons were stimulated in two conditions²⁰, one in which a single object was presented, and another in which two objects were presented, but in a way that evoked practically the same firing rates as the single stimulus. In this case, the synchrony between pairs of neurons reflected whether one or two stimuli were shown, even when both firing rates did not vary across conditions²⁰.

These examples show that, if neurons are sensitive to correlations, they are able to extract more information from their inputs. So, the neural codes used to represent the physical world could be made more efficient by taking into account the second-order statistics of neural responses. The degree to which this is actually the case is not the main issue here; the key observation is that, from this point of view, correlations are stimulus-dependent, just like sensory-evoked firing rates. The studies discussed below point to a more dynamic picture, in which correlations change rapidly as functions of internal events, and can regulate the flow of neural information, rather than its meaning.

Common patterns of correlated activity

A popular analytical tool used by neuroscientists to study the joint activity of neurons is the cross-correlation histogram or cross-correlogram (REFS^{21–23}; see also REF.²⁴). It is constructed from the spike trains of two neurons, and shows the probability (or some quantity proportional to it) that neuron B fires a spike τ milliseconds before or after a spike from neuron A; τ is called the time shift or time lag. When the two spike trains are independent, the cross-correlogram is flat; if there is any covariation in the spike trains, one or more peaks appear²³. For instance, a peak at zero time shift means that the two neurons tend to fire at the same time more often than expected by chance. Usually, cross-correlograms are corrected so that peaks caused by covariations in mean firing rate, computed over several tens or hundreds of milliseconds, are eliminated^{21–23}.

FIGURE 1 shows artificial, computer-generated spike trains and illustrates two common forms of correlated activity: synchrony and oscillations. Each raster plot shows twenty spike trains generated simultaneously. The artificial neurons fire randomly but always at the same mean rate of 27 spikes s^{-1} . What varies across panels is the temporal relationship between spikes from different neurons. In FIG. 1a–c, the neurons fire with increasing degrees of synchrony, caused by common inputs. The red traces superimposed on the rasters show the overall spike density or instantaneous firing rate, averaged over all neurons; the cross-correlograms are shown below, in blue. Typically, cross-correlograms from experimental data also have single peaks, although they can vary in width from a few to several hundred milliseconds^{11,25–28}. As the peak increases, different neurons tend to fire more often at the same time, causing larger fluctuations in spike density. So, when inputs to a neuron are synchronized, the total synaptic drive generated might be much more variable than when inputs are independent (this also depends on whether the inputs are excitatory or inhibitory, as discussed below).

In FIG. 1d–f, the 20 neurons fired at a rate that varied sinusoidally with a frequency of 25 Hz, with identical phase for all units. Here, by construction, fluctuations in spike density increase with oscillation amplitude. Cross-correlograms show the strength and frequency of the oscillations. In the examples shown, oscillatory activity could also be detected from the spike trains of individual neurons, as indeed it has been^{29,30}. However, a subtle technicality should be mentioned: the spike densities in FIG. 1 involve an ensemble average (an average over neurons), whereas spike densities are often constructed from single-neuron peristimulus time histograms (PSTHs), which involve an average over trials. The PSTH is useful in detecting

synchronous activity that is triggered by external stimuli, because there is an absolute time reference; however, some sort of ensemble average is necessary to analyse oscillations that are generated intrinsically, as these have their own timing. Oscillations observed in actual experimental records are complex, typically involving broad frequency bands^{30–33}, so they might be difficult to detect from single-neuron spike trains^{32,33}.

A significant drawback of cross-correlation methods is that they require large amounts of data to resolve significant deviations from independence. Alternative techniques that overcome this and other limitations are discussed below.

When are neurons sensitive to correlated input?

Having described some correlation patterns that are commonly observed in various experimental preparations, we now turn to the first main issue: the sensitivity of a postsynaptic cortical neuron to the presence of such correlations in its inputs. The goal is to spell out, at least to a first approximation, the factors that determine such sensitivity.

Coincidence detection

In theory, neurons might be exquisitely sensitive to certain temporal input patterns. The classical mechanism proposed for this is coincidence detection, which occurs when a neuron is sensitive to the arrival of spikes from two or more inputs within a short time window^{34–37}. There are examples, most notably in the auditory system^{38,39}, in which highly accurate coincidence detection takes place, but the question is whether this mechanism is commonly used throughout the cortex.

In the traditional view, coincidence detection is based on a very short MEMBRANE TIME CONSTANT^{34–37}. However, it can be greatly enhanced by the spatial arrangement of synapses and by nonlinear processes. For example, neighbouring synapses might interact strongly, forming clusters in which responses to simultaneous activation are much stronger than the sum of individual, asynchronous responses^{40,41}. A neuron could operate with many such clusters which, if located on ELECTROTONICALLY DISTANT parts of the dendritic tree, could act independently of each other, increasing storage and computational capacity^{42,43}. Voltage-dependent channels in the dendrites might mediate or boost such nonlinear interactions between synapses^{40–43}. These nonlinearities could, in principle, increase the capacity for coincidence detection to the point of making the neuron selective for a specific temporal sequence of input spikes; they could also serve as a basis for gating or selective amplification mechanisms^{42,44}. However, the degree to which the cortex exploits such nonlinearities is unclear.

An alternative mechanism for enhancing the coincidence detection capabilities of a neuron is to adjust the kinetics of its voltage-sensitive channels to favour large, transient depolarizing events like those that would result from synchronous inputs (FIG. 1a–c). Recent *in vivo* recordings⁴⁵ give some support to this possibility.

In these examples, the question is whether, at given input rates, the temporal alignment of spikes is important for the postsynaptic response. But the coincidence detection problem can also be posed as follows^{46,47}: if a neuron receives a volley of input spikes, what is the likelihood of evoking a response (reliability), and what will its timing be relative to the centre of mass of the input volley (precision)? Theoretical studies indicate that, in this case, the temporal precision of the response spikes is not limited by the membrane time constant, but rather by the rise time of excitatory synaptic events⁴⁷. The end result is that a volley of synchronized action potentials can propagate in a stable way through many layers.

Neurons can be driven by fluctuations

The flip side of coincidence detection is integration. Neurons can also act as integrators: they can sum or average their inputs to generate an action potential^{34,37,48}. Earlier theoretical arguments suggested that neurons acting as integrators would not be sensitive to temporal correlations^{48,49}, or that these would matter only at high firing rates, when refractory effects become important^{50,51}. However, these conclusions were based on models in which parameter space had been explored in a limited way; in particular, the role of inhibition had been underestimated. Recent results⁵² show that neurons can still be highly sensitive to weak correlations in their inputs, even if there is no spatial segregation along the dendritic tree, no variation in spike threshold, and no synaptic interaction beyond the expected temporal summation of postsynaptic currents. A key quantity in this case is the ‘balance’ of the neuron, which refers to the relative strength between inhibitory and excitatory inputs^{52–54}. When the neuron is not balanced, excitation is, on average, stronger than inhibition, such that the net synaptic current is depolarizing and the mean steady-state voltage is above threshold. In this case, the main driving force is the drift towards steady state, and input fluctuations have a small effect on the rate of output spikes^{52,55}. Conversely, when the neuron is balanced, both excitation and inhibition are strong, such that the mean input current is zero or very small, and the mean steady-state voltage remains below threshold. However, the neuron might still fire, because there are large voltage fluctuations that lead to random threshold crossings. Networks of balanced neurons have rich dynamics^{56–59} and can react to external stimuli on effective timescales that are much smaller than the membrane time constant of a single neuron^{57,58}. Crucially, in this mode, any factor that enhances the fluctuations will produce more intense firing^{52,60}.

There is a subtle but important distinction between mechanisms that alter input fluctuations. Higher firing rates should be seen in a balanced neuron if fluctuations increase without affecting the mean synaptic conductances, as when only the correlations are changed⁵². But if stronger fluctuations are accompanied by increases in conductance, as when both excitatory and inhibitory inputs fire more intensely, the firing rate can decrease^{60–62}. In a complex network, these effects can be hard to disentangle.

FIGURE 2 compares the responses of balanced (red traces) and unbalanced (green traces) model neurons. These are driven by excitatory and inhibitory input spike trains, like those illustrated in FIG. 1. The four panels correspond to different correlation patterns in the inputs. Cell responses were obtained using leaky integrate-and-fire models (see BOX 1). In FIG. 2a, all inputs are independent. The voltage traces show a typical difference between balanced and unbalanced modes: although the output rate is approximately the same, the subthreshold voltage of the balanced neuron is noisier and its interspike intervals are more variable^{52,54}. FIG. 2b shows what happens when the excitatory inputs fire partly synchronously (10% of their inputs being shared, as in FIG. 1a). The firing rate of the balanced neuron always increases relative to the response to independent inputs, whereas the rate of the unbalanced neuron might show smaller (although still substantial) increases, or might decrease^{50,51} if the output rate is already high without correlations. Another effect of synchrony is to increase the variability of the output spike trains, for both balanced and unbalanced configurations^{52,63–65}; this can be seen by comparing FIG. 2b and d with FIG. 2a. Correlations between inhibitory inputs can also produce stronger responses (FIG. 2c). When the inhibitory drive oscillates sinusoidally, as in FIG. 1e, the firing rate of the balanced neuron increases to practically double that recorded with no oscillations; by contrast, the firing rate of the unbalanced neuron does not change.

Box 1**The leaky integrate-and-fire model**

In the leaky integrate-and-fire model, driven by conductance changes^{52,54,124,125}, membrane potential V evolves according to $\tau_m(dV/dT) = -(V + E_{leak}) - g_{exc}(V - E_{exc}) - g_{inh}(V - E_{inh})$, but the spike-generating currents are substituted by a simple rule: whenever V exceeds a threshold (-54 mV), a spike is emitted and V is clamped to a reset value (-60 mV) for a refractory period (1.72 ms). After that, V continues to evolve according to the above equation. Every time an excitatory input spike arrives, the excitatory synaptic conductance g_{exc} increases instantaneously by an amount \bar{g}_{exc} ; otherwise, it decreases exponentially towards zero with a certain time constant (5 ms). The inhibitory conductance is modelled in the same way, except that it increases by \bar{g}_{inh} whenever an inhibitory spike arrives. In FIG. 2, for the balanced neuron: $\bar{g}_{exc} = 0.05$ and $\bar{g}_{inh} = 0.708$. For the unbalanced neuron: $\bar{g}_{exc} = 0.0167$ and $\bar{g}_{inh} = 0.0822$. All other parameters were identical for the two conditions: $\tau_m = 20$ ms, $E_{leak} = -74$ mV, $E_{exc} = 0$ mV, $E_{inh} = -63$ mV, where τ_m is the membrane time constant, E_{leak} is the resting membrane potential, and E_{exc} and E_{inh} are, respectively, the reversal potentials of the excitatory and inhibitory synapses. See REF.⁵² for further details.

The balance of a neuron is important in determining its sensitivity to correlations, but there is another key factor⁵². There are three correlation terms: correlations between pairs of excitatory neurons, between pairs of inhibitory neurons, and between excitatory–inhibitory pairs. The last term acts in the opposite direction, decreasing the fluctuations, and the total effect on the postsynaptic neuron is a function of the three terms. In FIG. 2d, all inputs to the model neurons are equally correlated, but the balanced model shows no change in firing rate. So, it is possible to have strong correlations between all inputs, but still not see a change in the firing rate of the postsynaptic neuron relative to the case of independent inputs.

In summary, a balanced neuron is much more sensitive to input correlations than an unbalanced one because correlations affect the fluctuations in synaptic drive, which cause the balanced neuron to fire. However, the postsynaptic response depends on the relative values of the three correlation terms, which might cancel out. The key point here is that the output of the neuron will be determined, not only by the firing rates of its inputs, but also by their correlations. Neurons can, in a statistical sense, be highly sensitive to the temporal patterns of their input spikes.

How correlations affect downstream activity?

From the examples in FIG. 2, it is clear that input correlations can have various effects. For instance, everything else being equal, oscillations can increase the gain of a balanced neuron (FIG. 2c, right) by increasing the output rate by approximately the same factor across a large dynamic range. On the other hand, notice the unbalanced configuration in FIG. 2b: for output rates below 30 spikes⁻¹ or so, correlations might act as a switch, turning on or off the transmission of spikes. FIGURE 2b–d shows simple examples, but other, more complex correlation patterns might have different effects.

In previous studies^{36,66,67}, it was noted that synchronous excitatory inputs can be more effective than independent inputs, but that the range of effects that correlations can give rise to⁵² seems to have been underestimated⁴⁹. Many variants of the following scenario are plausible. Consider three groups of neurons — A, B and M. Group A drives group B, which is sensitive to correlations, and group M provides input to A in a way that changes the correlations in A, but not the mean firing rate. The response of B will change as M is activated,

and the nature of the change will depend on the details of the circuitry⁶⁸. The plots in FIG. 2 indicate possible modulations in gain^{69–71} (FIG. 2band c , balanced), or full on–off switching⁷² for low output rates (FIG. 2b, unbalanced). Regardless of the final effect, a group of neurons M might affect another group A in two, not necessarily exclusive, ways: by changing the firing rates of A or the correlations between local neurons in A. There are two knobs that can be turned, not just one.

How neurons generate rhythmic activity

In addition to synchronized inputs arising from sensory codes, oscillations can be generated intrinsically. The phase of a spike relative to these oscillations can also carry information. For example, there are cells in the hippocampus of rats that fire according to the location of the animal in the world: they have a ‘place field’, and the phase of their spikes relative to an underlying oscillation encodes where the animal is located within this field⁷³. The phase of a spike from a pyramidal cell can also encode the number of inhibitory inputs impinging on it, independently of the information carried by the firing rate⁷⁴.

Cortical structures have a wide range of intrinsic mechanisms that could generate synchronous activity⁷⁵. Inhibitory interneurons might be particularly important. They are highly effective at entraining cortical neurons^{76–79}. In addition, recent evidence indicates that fast-spiking interneurons are highly sensitive to single excitatory inputs^{80,81}, and that there is direct electrical coupling between specific classes of cortical inhibitory interneurons⁸². All of these factors make inhibitory cells good candidates for generating synchronized oscillations in the 20–40-Hz range.

Cortical oscillations are regulated by neuromodulatory substances. Neurons in the hippocampus, for example, show several kinds of oscillation, each with different dominant frequencies^{83,84}. Rhythms in the delta (0.5–2-Hz), theta (4–12-Hz) and gamma (30–80-Hz) frequency bands are fairly common, and vary with behavioural conditions⁸³. Interestingly, the neurotransmitter acetylcholine modulates hippocampal circuitry such that the transitions between these three frequency bands depend on its concentration⁸³.

Although the functional roles of these oscillations still need to be worked out in detail, they do correlate with certain stereotyped behaviours^{4,73,85–87}. If different frequency bands correspond to distinct functions, then modulating the level of acetylcholine or other neurotransmitters would cause a switch from one to another. The lesson is that neuromodulators can shift the states of cells, giving rise to oscillatory activity or switching from one oscillatory regime to another. Experiments in invertebrates^{88,89} show that neurons can perform such switching, with highly specific behavioural consequences.

What can correlations tell us?

So far, we can conclude that several mechanisms are available to cortical neurons that allow them to generate and to respond to concerted activity as part of their everyday dynamics. Now we turn to the relationship between correlations and specific functions. First, we consider the advantages and limitations of measuring correlations, and then we discuss a broad range of experiments in which correlations are linked to expectation, attention, sensory latencies and rivalry —processes that regulate the strength but not the content of neural signals.

In itself, the presence of correlations between pairs of neurons is not particularly meaningful: either common inputs or synaptic interactions between them will give rise to some form of correlated activity^{8,22,49,56–59}. Synchrony *per se* could reflect, for example, nothing more than receptive field overlap. Likewise, oscillations can arise naturally from the intrinsic properties of neurons^{90–92} or from the recurrent nature of cortical circuits^{59,77,78,93–95}. This is not to say

that correlations are not important, but simply that their presence alone is not particularly informative.

What is important, however, is that changes in the correlation structure of a neural circuit reflect changes in its functional connectivity^{22,96,97}. We should emphasize this: it is the change from one condition to another, or from one moment to another, that serves as a probe for how the circuit works.

The caveat here is that changes in correlated activity could easily be confounded by simultaneous variations in the average firing rates of the recorded neurons. Therefore, studies that rely on the measurement of correlations face three technical problems: first, two or more neurons need to be recorded simultaneously; second, two conditions should be compared to make a differential measurement; and third, other aspects of the evoked neural activity should change as little as possible across conditions. In particular, when the two conditions involve different stimuli, it is likely that the evoked firing rates from the recorded neurons will change; even the populations that respond within a given area might be different. This is one of the main factors that muddles the interpretation of experiments in which correlations have been measured^{9,10}. To circumvent these obstacles, investigators have studied correlated activity in systems in which, across trials, variations in stimulation conditions are kept to a minimum, and the most significant changes are in the internal state of a subject.

Expectation boosts synchrony in motor cortex

In one paradigm, used to study the relationship between activity in the primary motor cortex (M1) and behaviour⁹⁸, monkeys were trained to perform a simple delayed-response task in which two cues were presented. The first cue indicated a target position and instructed the animal to get ready, whereas the second cue gave the 'go' signal for the requested hand movement. Crucially, the go signal could appear 600, 900, 1,200 or 1,500 ms after the cue, and this delay varied randomly from trial to trial. Neurons recorded in M1 increased their synchrony around the time of the actual sensory stimulus, or when the animal expected the go signal but it did not appear⁹⁸. In the former case, which is more like a sensory-evoked response, synchronization was accompanied by changes in firing rate, although these were separate effects (according to the analysis, the changes in synchrony were not due to observed changes in rate). However, in the latter case, which depends exclusively on the internal state of the monkey, firing rates did not change.

Another interesting observation⁹⁸ was that the patterns of synchronization were diverse: some pairs of neurons synchronized only before the first cue, others before and after the cue, others at only some of the expected go times, and so on. So, the pattern of correlations among neurons can change rapidly as a function of internal state without accompanying variations in mean firing rate.

Attention synchronizes somatosensory activity

Neurons in the secondary somatosensory cortex (S2) respond to tactile stimuli, but are also sensitive to attention^{99–101} and behavioural context¹⁰². Motivated by an earlier theoretical proposal¹⁰³, the correlations between pairs of neurons recorded in S2 were analysed as functions of attention¹⁰⁴. The monkeys used for these experiments were trained to perform two tasks, one visual and one tactile. In the tactile task^{99,104}, a raised pattern was presented to the fingertips, and the monkey had to indicate whether it matched a visual pattern shown on a monitor. In the visual task, identical tactile stimuli were presented, but the animal had to ignore them; the actual task leading to a reward was to detect the dimming of a target spot shown on the monitor. The animals were cued to perform blocks of trials of either task. In the recording sessions from which data were used for the analysis, the monkeys switched tasks at least four

times¹⁰⁴. So, responses to the same tactile stimuli were obtained in two conditions, when the monkeys had to pay attention to them and when they had to ignore them.

Cross-correlations between pairs of neurons were computed in the two conditions¹⁰⁴. The red curves in FIG. 3 are correlograms obtained when the animals performed the tactile task; the blue curves are those obtained when animals performed the visual task. These histograms are essentially the same constructs as those shown in FIGS 1 and 2, but the units on the y axes are coincidences s^{-1} , and the normalization is slightly different. The top two graphs in FIG. 3 are representative of the more common case, in which higher synchrony was observed with attention focused on the tactile stimuli. In some cases, the opposite effect was observed, as shown in the bottom graph; however, overall, directing attention to the fingertips tended to synchronize the S2 neurons responding to stimuli presented there. Interestingly, changes in synchrony were stronger when the somatosensory discrimination task was more difficult¹⁰⁴.

In these experiments, the firing rates of S2 neurons also varied with attention, consistent with previous reports^{99–102}. About 78% of the neurons showed significant changes in firing rate across conditions, whereas only 11% of all neuron pairs showed significant changes in synchrony¹⁰⁴. The analysis eliminated changes in correlation that might arise when pairs of neurons vary their mean firing rates jointly as functions of attention. Besides, these two effects were unrelated because changes in synchrony were not correlated with changes in firing rate. But even if there are two separate effects here, the changes in rate do pose a problem. Consider a downstream neural population driven by S2 and sensitive to input synchrony: will the observed changes in synchrony make a difference when accompanied by large changes in firing rate? This is unknown.

Attention synchronizes visually evoked activity

Attention modulates the firing rates of neurons in many parts of the visual system^{105–111}. However, the strength of this modulation can vary; in particular, the contrast at which stimuli are displayed can be adjusted so that changes in firing rate are minimized¹¹². This was done recently¹¹³ to investigate how correlations between visual neurons in area V4 change with attention, under conditions in which sensory input is constant and mean firing rates vary minimally.

Monkeys were trained to fixate on a central spot and to attend to either of two stimuli presented simultaneously and at the same eccentricity¹¹³ (FIG. 4). One of the stimuli fell inside the receptive field of a neuron, the activity of which was recorded. So, the responses to the same stimulus could be compared in two conditions, with attention inside or outside the neuron's receptive field. At the same time, the local field potential (LFP) was recorded from a nearby electrode. The LFP is the electric field caused by transmembrane currents flowing near the electrode, so it gives some indication of average local activity. This is a key premise in the interpretation of many synchrony results, but it seems reasonable because variations in LFP and membrane potential recorded intracellularly are highly correlated¹¹⁴. The LFP is particularly useful when searching for oscillatory^{31–33} and synchronous¹¹⁵ activity. Examples of LFP traces and spike trains from the two electrodes are shown in FIG. 4a and b.

The correlation that was studied in these experiments¹¹³ was that between the LFP and the recorded neuron's spikes. The key quantity here is the spike-triggered average of the LFP, or STA. The STA is obtained by adding, for each spike recorded, a segment of the LFP centred on the time of the spike; the final sum is then divided by the total number of spikes. The result is the average LFP waveform that is observed around the time of a spike. STAs computed for attention outside and inside the receptive field are shown in FIG. 4c and d, respectively. A rigorous comparison is shown in FIG. 4e, which plots the power spectra of the two STAs. Note, in FIG. 4c and d, that the spike occurs very near to the trough of both the low- and high-

frequency components^{31–33}. This is consistent with the idea that the LFP is related to the average voltage of the cells near the electrode¹¹⁴, but with opposite sign because negative currents depolarize the membrane. The similarity in phase is also evidence for coupled oscillations in the local population of neurons. The two STAs are similar, but they are not identical: the rapid fluctuations are more pronounced when attention is directed inside the receptive field; power in the low-frequency band (0–17 Hz) decreases, whereas power in the high-frequency band (30–70 Hz) increases. Because the STA reflects the correlation between one neuron and the neighbouring population, the interpretation is that, as attention shifts to the receptive fields of a cluster of neurons, these become more synchronized at high frequencies and less so at low frequencies. These changes in synchrony were quantified rigorously using a further measure that was independent of firing rate and of LFP power.

The measures of synchrony used in this study are much more sensitive than the traditional correlogram, because the LFP involves averaging over a population¹¹³. Although the changes in synchrony were modest — on average, low-frequency synchronization decreased by 23% and high-frequency synchronization increased by 19% — changes in firing rate were also small; these were enhanced by a median of 16% with attention inside the receptive field. Therefore, under these conditions, the changes in synchrony can be significant in terms of their impact on the responses of downstream neurons. Another interesting observation is that, whereas the attentional modulation of firing rate started about 450 ms after stimulus onset, significant changes in synchrony could be detected very early in the response (50 ms after stimulus onset). These attentional effects were spatially specific; they did not result from a generalized change in arousal.

What exactly is the neural correlate of attention? What happens to the S2 or the V4 circuitry as attention is focused inside or outside their response fields? Unfortunately, we do not yet know. The plots in FIG. 2c, for example, indicate that an increase in oscillatory coherence could lead to higher firing rates. But things are not so simple^{105–112}. The first problem is that attention can lead, not only to increases, but also to reductions in firing. It is not known whether these decreases in rate are correlated with decreases in synchrony at some level. Second, the magnitude of the modulation depends on stimulus configuration, on the neuron's selectivity and tuning properties, and on contrast. Third, our mechanistic understanding of how one circuit can synchronize another, or how an increase in input synchrony might enhance or suppress the activity of a postsynaptic target, is still crude.

However, these results are important for two reasons: first, because they constrain the space of possible interactions between a local circuit and its attentional inputs; and second, because they support the notion that the correlation structure of a neural population may change dynamically, and may be crucial in determining the responses of its downstream targets.

Latencies are correlated by gamma oscillations

The study discussed above¹¹³ shows that synchrony, specifically in the gamma frequency (roughly 30–80 Hz), might enhance the processing of information in some way. But what exactly is the impact of such synchronization? Another recent study¹¹⁶ indicates at least one measurable consequence: the latencies of synchronized neurons responding to a stimulus can shift in unison. In this case, the paradigm was very simple: oriented bars of light were flashed and the responses of two or more neurons in the primary visual cortex (V1) were recorded, along with LFPs. Neurons were activated by the stimuli, and the key quantity examined was the time taken for the neurons to respond — the latency — which was calculated on each trial. Latencies covaried fairly strongly from trial to trial (mean correlation coefficient of 0.34, with a range from 0.18 to 0.55), so pairs of neurons tended to fire early or late together. This tendency depended on the amount of gamma power in the LFPs immediately before the stimulus. When

the LFPs from two electrodes both had a strong gamma component, the latency covariation between the two recorded neurons from the same pair of electrodes was high. Note that the spectral composition of the LFPs was only weakly related to changes in firing rate, so short latencies were probably not due to changes in excitability. This means that, if neurons are synchronized at around 40 Hz, just before a stimulus is presented, they will respond at about the same time¹¹⁶. In other words, although the mean firing rates are mostly insensitive to shifts in oscillation frequencies, the time spread in the evoked spikes from multiple neurons is much smaller when the gamma oscillations are enhanced. This could certainly have an impact on a downstream population driven by these neurons^{46,47,52}. So, the modulation of latency covariations¹¹⁶ is a concrete example of how the synchrony of a local circuit can be used to control the strength of a neural signal.

Rivalry induces changes in synchrony in V1

Another study¹¹⁷ investigated the synchronization of V1 neurons, this time using an interocular rivalry paradigm. In rivalry experiments^{118,119}, different images are shown to the two eyes, but only one image is perceived at any given moment. The percept can flip from one image to the other randomly, but with a characteristic timescale that depends on the experimental setup. The study in question¹¹⁷ was done on awake STRABISMIC CATS, a preparation with two advantages: V1 neurons are dominated by a single eye, so their firing rates essentially depend on what their dominant eye sees regardless of the other one, and it is relatively easy to know which of the two images is perceived (at equal contrasts for the two images, one eye always suppresses the other, and this can be measured by tracking the cat's eye movements in response to conflicting moving stimuli). The two conditions compared were a single image presented to the eye driving the recorded neurons, or the same stimulus shown to the driving eye plus a conflicting image presented to the other eye. The firing rates in these two conditions should be the same, because of the strabismic condition; indeed, the rates did not change significantly across conditions and did not depend on which image was perceived. However, synchrony did change across conditions¹¹⁷. When neurons were driven by the eye providing the percept, synchrony was much stronger in the rivalrous condition than in the monocular one. By contrast, when neurons were driven by the eye for which the image became suppressed, synchrony was much lower in the rivalrous condition than in the monocular one. In other words, when conflicting images were presented, neurons responding to the image being perceived were always more synchronized. Effects were seen in cross-correlograms among spike trains, but were much more robust and evident in the STAs.

In this case, the stimuli were not strictly identical in the two conditions being compared, but the feedforward thalamic inputs driving any single V1 neuron must have been extremely similar. It is tempting to conclude then that the perceptual experience is expressed in V1 as a synchronization pattern, but going that far is unnecessary; what matters is that synchrony is consistently modulated across conditions independently of firing rate, and that the modulation probably originates in the cortex. As in the case of latency covariations, what we understand as the neural code for a stimulus does not change appreciably.

Conclusion

The literature showing that neural circuits have rhythmic and synchronous activity is vast, but this is hardly unexpected given the divergent and convergent interconnections of the cortex^{120,121}, which lead to receptive field overlap, and the intrinsic and emergent oscillatory properties of single and interconnected neurons, respectively. Nevertheless, the dynamics of these forms of correlated activity can still be highly informative about the functional connectivity of cortical circuits. This observation, and the idea that correlations between spikes

can be used for information processing, have been discussed before^{96,97,122}, but the results reviewed here lead to a more specific conclusion.

First, correlation measurements have been compared across conditions using paradigms in which other attributes of the evoked neural activity — in particular, the mean firing rates — vary as little as possible. These experimental studies indicate that large variations in correlations can be observed in the absence of simultaneous variations in mean firing rates. Second, through theoretical work, various factors have been identified that might endow neurons with a high sensitivity to correlations. These two sets of findings support the idea that correlations might be modulated independently of mean firing rates. In our view, this point is of enormous interest and its generality might have been underappreciated; using such independent modulations for object representation is just one possibility, the relevance of which is still unclear^{9,10}. What else could correlations be used for? There could be many other alternatives. For example, according to a recent theoretical proposal¹²³, the transient oscillatory activity evoked by an auditory stimulus could serve as the basis for speech recognition, with the advantage that this mechanism would be invariant to uniform time warp and intensity change in input sounds. However, the experimental studies mentioned above point to a third key piece of evidence: rate-independent modulations in synchrony have been linked to changes in expectation, attention, response latency and rivalry — processes that adjust the flow of information, but have little bearing on stimulus representation. Therefore, we suggest that a more natural role for temporal correlations might be to control the strength of a signal, and hence the downstream circuits that it reaches, rather than the nature of the information that it conveys. Correlations might also regulate synaptic plasticity through spike-timing-dependent mechanisms, enhancing the memory of attended stimuli.

The challenge that lies ahead is to work out how synchrony in a neural circuit can be controlled by other circuits to perform useful operations.

Acknowledgments

Research was supported by the Howard Hughes Medical Institute. We thank P. Steinmetz for providing us with Figure 3, and P. Fries for providing us with Figure 4. We also thank J. Reynolds and P. Tiesinga for helpful comments.

Glossary

BINDING PROBLEM	The problem of binding together representations of the different properties of an object (for example, its colour, form and location)
MEMBRANE TIME CONSTANT	A quantity that depends on the capacitance and resistance of the cell membrane, and which sets a timescale for changes in voltage. A small time constant means that the membrane potential can change rapidly
ELECTROTONICALLY DISTANT	Two points on the dendritic tree are electrotonically distant if the electrical interactions between them are minimal, regardless of the actual physical distance between the points
STRABISMIC CAT	A condition in which the eyes are not straight or properly aligned. The misalignment reflects the failure of the eye muscles to work together. One eye may turn in (crossed eyes), turn out (wall eyes), turn up or turn down. Although some cats are congenitally strabismic, strabismus can also be achieved by cutting the tendon of one of the eye muscles

References

1. Seidemann E, Zohary U, Newsome WT. Temporal gating of neural signals during performance of a visual discrimination task. *Nature* 1998;394:72–75. A rare inquiry into how neural signals are gated. Microstimulation pulses were applied in the middle temporal visual cortex (MT) during a visual motion discrimination task. Their effect depended critically on the timing of the pulses relative to the time of stimulus presentation. [PubMed: 9665129]
2. Barlow, JS. *The Electroencephalogram: its Patterns and Origins*. MIT Press; Cambridge, Massachusetts: 1993.
3. Borbély, AA.; Hayaishi, O.; Sejnowski, TJ.; Altman, JS. *Human Frontier Science Program*. Strasbourg: 2000. *The Regulation of Sleep*.
4. Destexhe A, Sejnowski TJ. Why do we sleep? *Brain Res* 2000;886:208–223. [PubMed: 11119697]
5. Chrobak JJ, Buzsáki G. High-frequency oscillations in the output networks of the hippocampal–entorhinal axis of the freely-behaving rat. *J Neurosci* 1996;16:3056–3066. [PubMed: 8622135]
6. Engel AK, König P, Schillen TB. Why does the cortex oscillate? *Curr Biol* 1992;2:332–334. [PubMed: 15335950]
7. Singer W, Gray CM. Visual feature integration and the temporal correlation hypothesis. *Annu Rev Neurosci* 1995;18:555–586. [PubMed: 7605074]
8. Usrey WM, Reid RC. Synchronous activity in the nervous system. *Annu Rev Neurosci* 1999;61:435–456.
9. Shadlen MN, Movshon JA. Synchrony unbound: a critical evaluation of the temporal binding hypothesis. *Neuron* 1999;24:67–77. [PubMed: 10677027]
10. Gray CM. The temporal correlation hypothesis of visual feature integration: still alive and well. *Neuron* 1999;24:31–47. [PubMed: 10677025]
11. Aertsen A, Arndt M. Response synchronization in the visual cortex. *Curr Opin Neurobiol* 1993;3:586–594. [PubMed: 8219726]
12. Wehr M, Laurent G. Relationship between afferent and central temporal patterns in the locust olfactory system. *J Neurosci* 1999;19:381–390. [PubMed: 9870967]
13. Laurent G, et al. Odor encoding as an active, dynamical process: experiments, computation, and theory. *Annu Rev Neurosci* 2001;24:263–297. [PubMed: 11283312]
14. Bazhenov M, et al. Model of transient oscillatory synchronization in the locust antennal lobe. *Neuron* 2001;30:553–567. [PubMed: 11395014]
15. MacLeod K, Bäcker A, Laurent G. Who reads temporal information contained across synchronized and oscillatory spike trains? *Nature* 1998;395:693–698. This and the next paper are complementary studies on the functional role of synchrony in the olfactory system of insects. When neurons in the antennal lobe are artificially desynchronized, the responses of downstream neurons are markedly distorted and the animals' ability to discriminate odours is impaired. Similar experiments investigating the impact of oscillations on neural circuits and on behaviour should, at some point, be possible in mammals. [PubMed: 9790189]
16. Stopfer M, Bhagavan S, Smith BH, Laurent G. Impaired odour discrimination on desynchronization of odour-encoding neural assemblies. *Nature* 1997;390:70–74. [PubMed: 9363891]
17. Kashiwadani H, Sasaki YF, Uchida N, Mori K. Synchronized oscillatory discharges of mitral/tufted cells with different molecular receptive ranges in the rabbit olfactory bulb. *J Neurophysiol* 1999;82:1786–1792. [PubMed: 10515968]
18. Dan Y, Alonso JM, Usrey WM, Reid RC. Coding of visual information by precisely correlated spikes in the lateral geniculate nucleus. *Nature Neurosci* 1998;1:501–507. [PubMed: 10196548]
19. DeCharms RC, Merzenich MM. Primary cortical representation of sounds by the coordination of action potential timing. *Nature* 1995;381:610–613. [PubMed: 8637597]
20. Kreiter AK, Singer W. Stimulus-dependent synchronization of neuronal responses in the visual cortex of the awake macaque monkey. *J Neurosci* 1996;16:2381–2396. [PubMed: 8601818]
21. Perkel DH, Gerstein GL, Moore GP. Neuronal spike trains and stochastic point processes. II. Simultaneous spike trains. *Biophys J* 1967;7:419–440. [PubMed: 4292792]
22. Aertsen AMHJ, Gerstein GL, Habib MK, Palm G. Dynamics of neuronal firing correlation: modulation of 'effective connectivity'. *J Neurophysiol* 1989;61:900–917. [PubMed: 2723733]

23. Brody CD. Correlations without synchrony. *Neural Comput* 1999;11:1537–1551. Points out several ways in which peaks might arise in cross-correlation histograms. An important reference for anyone using this method. [PubMed: 10490937]
24. Vaadia E, et al. Dynamics of neuronal interactions in monkey cortex in relation to behavioural events. *Nature* 1995;373:515–518. [PubMed: 7845462]
25. Ts'o DY, Gilbert CD, Wiesel TN. Relationships between horizontal interactions and functional architecture in cat striate cortex as revealed by cross-correlation analysis. *J Neurosci* 1986;6:1160–1170. [PubMed: 3701413]
26. Gochin PM, Miller EK, Gross CG, Gerstein GL. Functional interactions among neurons in inferior temporal cortex of the awake macaque. *Exp Brain Res* 1991;84:505–516. [PubMed: 1864322]
27. Nelson JI, Salin PA, Munk MHJ, Arzi M, Bullier J. Spatial and temporal coherence in cortico–cortical connections: a cross-correlation study in areas 17 and 18 in the cat. *Vis Neurosci* 1992;9:21–37. [PubMed: 1321657]
28. Brosch M, Schreiner CE. Correlations between neural discharges are related to receptive field properties in cat primary auditory cortex. *Eur J Neurosci* 1999;11:3517–3530. [PubMed: 10564360]
29. Livingstone MS. Oscillatory firing and interneuronal correlations in squirrel monkey striate cortex. *J Neurophysiol* 1996;75:2467–2485. [PubMed: 8793757]
30. Gray CM, Viana Di Prisco G. Stimulus-dependent neuronal oscillations and local synchronization in striate cortex of the alert cat. *J Neurosci* 1997;17:3239–3253. [PubMed: 9096157]
31. Murthy VN, Fetz EE. Oscillatory activity in sensorimotor cortex of awake monkeys: synchronization of local field potentials and relation to behavior. *J Neurophysiol* 1996;76:3949–3967. [PubMed: 8985892]
32. Murthy VN, Fetz EE. Synchronization of neurons during local field potential oscillations in sensorimotor cortex of awake monkeys. *J Neurophysiol* 1996;76:3968–3982. [PubMed: 8985893]
33. Donoghue JP, Sanes JN, Hatsopoulos NG, Gaàl G. Neural discharge and local field potential oscillations in primate motor cortex during voluntary movements. *J Neurophysiol* 1998;79:159–173. [PubMed: 9425187]
34. Abeles M. Role of the cortical neuron: integrator or coincidence detector? *Isr J Med Sci* 1982;18:83–92. [PubMed: 6279540]
35. Softky WR. Sub-millisecond coincidence detection in active dendritic trees. *Neuroscience* 1993;58:13–41. [PubMed: 8159289]
36. Softky WR, Koch C. The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *J Neurosci* 1993;13:334–350. [PubMed: 8423479]
37. König P, Engel AK, Singer W. Integrator or coincidence detector? The role of the cortical neuron revisited. *Trends Neurosci* 1996;19:130–137. [PubMed: 8658595]
38. Carr CE, Konishi M. A circuit for detection of interaural time differences in the brain stem of the barn owl. *J Neurosci* 1990;10:3227–3246. [PubMed: 2213141]
39. Agmon-Snir H, Carr CE, Rinzel J. The role of dendrites in auditory coincidence detection. *Science* 1998;281:268–272.
40. Mel BW. Synaptic integration in an excitable dendritic tree. *J Neurophysiol* 1993;70:1086–1101. [PubMed: 8229160]
41. Margulis M, Tang CM. Temporal integration can readily switch between sublinear and supralinear summation. *J Neurophysiol* 1998;79:2809–2813. [PubMed: 9582247]
42. Mel, BW. Dendrites. Stuart, G.; Spruston, N.; Häusser, M., editors. Oxford Univ. Press; Oxford: 1999. p. 271–289.
43. Poirazi P, Mel BW. Impact of active dendrites and structural plasticity on the memory capacity of neural tissue. *Neuron* 2001;29:779–796. [PubMed: 11301036]
44. Archie KA, Mel BW. Structural plasticity, dendritic subunits, and the development of nonlinear visual receptive field properties. *Soc Neurosci Abstr* 2000;26:1362.
45. Azouz R, Gray CM. Dynamic spike threshold reveals a mechanism for synaptic coincidence detection in cortical neurons *in vivo*. *Proc Natl Acad Sci USA* 2000;97:8110–8115. [PubMed: 10859358]
46. Burkitt AN, Clark GM. Analysis of integrate-and-fire neurons: synchronization of synaptic input and spike output. *Neural Comput* 1999;11:871–901. [PubMed: 10226187]

47. Diesmann M, Gewaltig M-O, Aertsen A. Stable propagation of synchronous spiking in cortical neural networks. *Nature* 1999;402:529–533. A different perspective on coincidence detection and synchrony. A network of neurons is activated by a volley of input spikes that occurs at a certain time and with a given temporal width. The response is another volley of spikes, the timing and width of which depend on network parameters. Analytical and simulation results show that such volleys can propagate stably and with minimal time jitter through multiple layers. [PubMed: 10591212]
48. Shadlen MN, Newsome WT. Noise, neural codes and cortical organization. *Curr Opin Neurobiol* 1994;4:569–579. [PubMed: 7812147]
49. Shadlen MN, Newsome WT. The variable discharge of cortical neurons: implications for connectivity, computation and information coding. *J Neurosci* 1998;18:3870–3896. [PubMed: 9570816]
50. Bernander Ö, Koch C, Usher M. The effects of synchronized inputs at the single neuron level. *Neural Comput* 1994;6:622–641.
51. Murthy VN, Fetz EE. Effects of input synchrony on the firing rate of a three-conductance cortical neuron model. *Neural Comput* 1994;6:1111–1126.
52. Salinas E, Sejnowski TJ. Impact of correlated synaptic input on output firing rate and variability in simple neuronal models. *J Neurosci* 2000;20:6193–6209. Using theoretical models and computer simulations, the authors study the possible effects of input correlations on a postsynaptic neuron, as well as the properties that make the postsynaptic neuron more or less sensitive to them. A key condition for high sensitivity is the balance between excitation and inhibition; a balanced neuron might be strongly affected by small changes in input synchrony or by oscillations in input firing rates. [PubMed: 10934269]
53. Bell, AJ.; Mainen, ZF.; Tsodyks, M.; Sejnowski, TJ. Technical Report INC-9502. Institute for Neural Computation, Univ. California; San Diego, California 92093-0523: 1995. ‘Balancing’ of conductances may explain irregular cortical spiking.
54. Troyer TW, Miller KD. Physiological gain leads to high ISI variability in a simple model of a cortical regular spiking cell. *Neural Comput* 1997;9:971–983. [PubMed: 9188190]
55. Feng J, Brown D. Impact of correlated inputs on the output of the integrate-and-fire model. *Neural Comput* 2000;12:671–692. [PubMed: 10769326]
56. Tsodyks MV, Sejnowski TJ. Rapid state switching in balanced cortical network models. *Network* 1995;6:111–124.
57. Van Vreeswijk C, Sompolinsky H. Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science* 1996;274:1724–1726. [PubMed: 8939866]
58. Van Vreeswijk C, Sompolinsky H. Chaotic balanced state in a model of cortical circuits. *Neural Comput* 1998;10:1321–1371. [PubMed: 9698348]
59. Brunel N, Hakim V. Fast global oscillations in networks of integrate-and-fire neurons with low firing rates. *Neural Comput* 1999;11:1621–1671. [PubMed: 10490941]
60. Doiron B, Longtin A, Berman N, Maler L. Subtractive and divisive inhibition: effect of voltage-dependent inhibitory conductances and noise. *Neural Comput* 2001;13:227–248. [PubMed: 11177434]
61. Tiesinga PHE, José JV, Sejnowski TJ. Comparison of current-driven and conductance-driven neocortical model neuron with Hodgkin–Huxley voltage-gated channels. *Phys Rev E* 2000;62:8413–8419.
62. Chance FS, Abbott LF. Multiplicative gain modulation through balanced synaptic input. *Soc Neurosci Abstr* 2000;26:1064.
63. Stevens CF, Zador AM. Input synchrony and the irregular firing of cortical neurons. *Nature Neurosci* 1998;1:210–217. [PubMed: 10195145]
64. Svirskis G, Rinzel J. Influence of temporal correlation of synaptic input on the rate and variability of firing in neurons. *Biophys J* 2000;79:629–637. [PubMed: 10919997]
65. Salinas E, Sejnowski TJ. Exact solutions for the non-leaky integrate-and-fire model neuron driven by correlated stochastic inputs. *Soc Neurosci Abstr* :27. in the press.
66. Ritz, R.; Sejnowski, TJ. *Artificial Neural Networks—ICANN ‘97*. Gerstner, W.; Germond, A.; Hasler, M.; Nicoud, J-D., editors. Springer; Lausanne, Switzerland: 1997. p. 79-84.
67. Usrey WM, Reppas JB, Reid RC. Paired-spike interactions and synaptic efficacy of retinal inputs to the thalamus. *Nature* 1998;395:384–387. [PubMed: 9759728]

68. José JV, Tiesinga PHE, Fellous J-M, Salinas E, Sejnowski TJ. Synchronization as a mechanism for attentional modulation. *Soc Neurosci Abstr* :27. in the press.
69. Salinas E, Their P. Gain modulation - a major computational principle of the central nervous system. *Neuron* 2000;27:15–21. [PubMed: 10939327]
70. Salinas, E.; Abbott, LF. *Progress in Brain Research*. Nicolelis, M., editor. Vol. 130. Elsevier; Amsterdam: in the press
71. Salinas E, Sejnowski TJ. Gain modulation in the central nervous system: where behavior, neurophysiology and computation meet. *The Neuroscientist*. in the press.
72. Zhang M, Barash S. Neuronal switching of sensorimotor transformations for antisaccades. *Nature* 2000;408:971–975. A study in awake monkeys showing that visually triggered sensory responses in the lateral intraparietal area (LIP) can be switched on or off depending on the contingencies of the task. As in reference 1, the emphasis is on how neurons communicate, rather than on how they represent the sensory world. [PubMed: 11140683]
73. O’Keefe J, Recce ML. Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus* 1993;3:317–330. [PubMed: 8353611]
74. Tiesinga PHE, Fellous JM, José JV, Sejnowski TJ. Optimal information transfer in synchronized neocortical neurons. *Neurocomputing* 2001;38:397–402.
75. Ritz R, Sejnowski TJ. Synchronous oscillatory activity in sensory systems: new vistas on mechanisms. *Curr Opin Neurobiol* 1997;7:536–546. [PubMed: 9287205]
76. Lytton WW, Sejnowski TJ. Simulations of cortical pyramidal neurons synchronized by inhibitory interneurons. *J Neurophysiol* 1991;66:1059–1079. [PubMed: 1661324]
77. Cobb SR, Buhl EH, Halasy K, Paulsen O, Somogyi P. Synchronization of neural activity in hippocampus by individual GABAergic interneurons. *Nature* 1995;378:75–78. [PubMed: 7477292]
78. Bush P, Sejnowski TJ. Inhibition synchronizes sparsely connected cortical neurons within and between columns in realistic network models. *J Comput Neurosci* 1996;3:91–110. [PubMed: 8840227]
79. Jefferys JGR, Traub RD, Whittington MA. Neuronal networks for induced ‘40 Hz’ rhythms. *Trends Neurosci* 1996;19:202–208. [PubMed: 8723208]
80. Fricker D, Miles R. EPSP amplification and the precision of spike timing in hippocampal neurons. *Neuron* 2000;28:559–569. [PubMed: 11144364]
81. Galarreta M, Hestrin S. Spike transmission and synchrony detection in networks of GABAergic interneurons. *Science* 2001;292:2295–2299. [PubMed: 11423653]
82. Galarreta M, Hestrin S. Electrical synapses between GABA-releasing interneurons. *Nature Rev Neurosci* 2001;2:425–433. [PubMed: 11389476]
83. Fellous J-M, Sejnowski TJ. Cholinergic induction of spontaneous oscillations in the hippocampal slice in the slow (0.5–2 Hz), theta (5–12 Hz) and gamma (35–70 Hz) bands. *Hippocampus* 2000;10:187–197. Three kinds of rhythmic activity are observed in a hippocampal slice preparation, and a single neuromodulator can shift the dynamics from one mode to another. A model for this concentration-dependent switching is developed in the reference below. [PubMed: 10791841]
84. Tiesinga PHE, Fellous JM, José JV, Sejnowski TJ. Computational model of carbachol-induced delta, theta and gamma oscillations in the hippocampus. *Hippocampus* 2001;11:251–274. [PubMed: 11769308]
85. Bland BH. The physiology and pharmacology of hippocampal formation theta rhythms. *Prog Neurobiol* 1986;26:1–54. [PubMed: 2870537]
86. Lisman JE. Relating hippocampal circuitry to function: recall of memory sequences by reciprocal dentate–CA3 interactions. *Neuron* 1999;22:233–242. [PubMed: 10069330]
87. Siapas AG, Wilson MA. Coordinated interactions between hippocampal ripples and cortical spindles during slow-wave sleep. *Neuron* 1998;21:1123–1128. [PubMed: 9856467]
88. Hooper SL, Moulins M. Switching of a neuron from one network to another by sensory-induced changes in membrane properties. *Science* 1989;244:1587–1589. [PubMed: 2740903]
89. Weimann JM, Marder E. Switching neurons are integral members of multiple oscillatory networks. *Curr Biol* 1994;4:896–902. [PubMed: 7850423]

90. Llinás RR. The intrinsic electrophysiological properties of mammalian neurons: insights into central nervous system function. *Science* 1988;242:1654–1664. [PubMed: 3059497]
91. Gray CM, McCormick DA. Chattering cells: superficial pyramidal neurons contributing to the generation of synchronous oscillations in the visual cortex. *Science* 1996;274:109–113. [PubMed: 8810245]
92. Lüti A, McCormick DA. H-current: properties of a neuronal and network pacemaker. *Neuron* 1998;21:9–12. [PubMed: 9697847]
93. Wilson M, Bower JM. Cortical oscillations and network interactions in a computer simulation of piriform cortex. *J Neurophysiol* 1992;67:981–995. [PubMed: 1316954]
94. Fuentes U, Ritz R, Gerstner W, Van Hemmen JL. Vertical signal flow and oscillations in a three-layer model of the cortex. *J Comput Neurosci* 1996;3:125–136. [PubMed: 8840229]
95. Timofeev I, Grenier F, Bazhenov M, Sejnowski TJ, Steriade M. Origin of slow cortical oscillations in deafferented cortical slabs. *Cereb Cortex* 2000;10:1185–1199. [PubMed: 11073868]
96. Von der Malsburg, C. *Models of Neural Networks II*. Domany, E.; Van Hemmen, JL.; Schulten, K., editors. Springer; Berlin: 1994. p. 95–119.
97. Aertsen A, Erb M, Palm G. Dynamics of functional coupling in the cerebral cortex: an attempt at a model-based interpretation. *Physica D* 1994;75:103–128.
98. Riehle A, Grün S, Diesmann M, Aertsen A. Spike synchronization and rate modulation differentially involved in motor cortical function. *Science* 1997;278:1950–1953. This study exploits a simple, yet creative, behavioural model to study synchronization in the primary motor cortex. Neurons in this area become transiently synchronized when a stimulus appears, or when it is expected to appear but it does not. In the former case, mean firing rates typically change with (but independently of) synchrony, but in the latter case they typically do not. [PubMed: 9395398]
99. Hsiao SS, Johnson KO, O’Shaughnessy DM. Effects of selective attention of spatial form processing in monkey primary and secondary somatosensory cortex. *J Neurophysiol* 1993;70:444–447. [PubMed: 8360721]
100. Burton H, Sinclair RJ, Hong SY, Pruett JR, Whang KC. Tactile-spatial and cross-modal attention effects in the second somatosensory and 7b cortical areas of rhesus monkeys. *Somatosens Mot Res* 1997;14:237–267. [PubMed: 9443366]
101. Johansen-Berg H, Lloyd DM. The physiology and psychology of attention to touch. *Front Biosci* 2000;5:D894–904. [PubMed: 11056079]
102. Salinas E, Hernández H, Zainos A, Romo R. Periodicity and firing rate as candidate neural codes for the frequency of vibrotactile stimuli. *J Neurosci* 2000;20:5503–5515. [PubMed: 10884334]
103. Niebur E, Koch C. A model for the neuronal implementation of selective visual attention based on temporal correlation among neurons. *J Comput Neurosci* 1994;1:141–158. [PubMed: 8792229]
104. Steinmetz PN, et al. Attention modulates synchronized neuronal firing in primate somatosensory cortex. *Nature* 2000;404:187–190. A study in which tactile stimuli were delivered and neurons in the secondary somatosensory cortex responded to them. When attention is focused on the tactile stimuli, the neurons respond more intensely and become more synchronized than when attention is directed towards a visual display. So, attention might regulate, through changes in synchrony, the strength of the somatosensory response. [PubMed: 10724171]
105. Moran J, Desimone R. Selective attention gates visual processing in the extrastriate cortex. *Science* 1985;229:782–784. [PubMed: 4023713]
106. Motter BC. Focal attention produces spatially selective processing in visual cortical areas V1, V2, and V4 in the presence of competing stimuli. *J Neurophysiol* 1993;70:909–919. [PubMed: 8229178]
107. Connor CE, Preddie DC, Gallant JL, Van Essen DC. Spatial attention effects in macaque area V4. *J Neurosci* 1997;17:3201–3214. [PubMed: 9096154]
108. McAdams CJ, Maunsell JHR. Effects of attention on orientation tuning functions of single neurons in macaque cortical area V4. *J Neurosci* 1999;19:431–441. [PubMed: 9870971]
109. Reynolds J, Desimone R. Competitive mechanisms subserve attention in macaque areas V2 and V4. *J Neurosci* 1999;19:1736–1753. [PubMed: 10024360]
110. Treue S, Martínez-Trujillo JC. Feature-based attention influences motion processing gain in macaque visual cortex. *Nature* 1999;399:575–579. [PubMed: 10376597]

111. Kastner S, Ungerleider L. Mechanisms of visual attention in human cortex. *Annu Rev Neurosci* 2000;23:315–341. [PubMed: 10845067]
112. Reynolds JH, Pasternak T, Desimone R. Attention increases sensitivity of V4 neurons. *Neuron* 2000;2:703–714. [PubMed: 10896165]
113. Fries P, Reynolds JH, Rorie AE, Desimone R. Modulation of oscillatory neuronal synchronization by selective visual attention. *Science* 2001;291:1560–1563. The responses of visual neurons were compared when attention was directed inside or outside their receptive fields, for the same stimulus. Extreme care was taken to minimize changes in mean firing rate and to measure synchrony in an unbiased way. When attention shifts to the recorded neuron's receptive field, the unit and its neighbours become more synchronized with respect to rapid (50-Hz) fluctuations, but less so with respect to slow (10-Hz) fluctuations. Attention seems to cause a complex yet stereotyped change in the dynamics of the local circuit of visual neurons. [PubMed: 11222864]
114. Frost JD. Jr An averaging technique for detection of EEG–intracellular potential relationships. *Electroencephalogr Clin Neurophysiol* 1967;23:179–181. [PubMed: 4166700]
115. Goto Y, O'Donnell P. Network synchrony in the nucleus accumbens *in vivo*. *J Neurosci* 2001;21:4498–4504. [PubMed: 11404437]
116. Fries P, Neuenschwander S, Engel AK, Goebel R, Singer W. Rapid feature selective neuronal synchronization through correlated latency shifting. *Nature Neurosci* 2001;4:194–200. Latencies in the responses evoked by visual stimuli were measured simultaneously for pairs of neurons. These latencies covaried across trials, with stronger covariations observed for pairs that were more synchronized in the band around 50 Hz. Covariations in latency were independent of covariations in firing rate, and were not caused by common input. A functional role for oscillations in the 50-Hz range is suggested: to temporally align the responses of the synchronized neural population to a forthcoming stimulus. [PubMed: 11175881]
117. Fries P, Roelfsema PR, Engel AK, König P, Singer W. Synchronization of oscillatory responses in visual cortex correlates with perception in interocular rivalry. *Proc Natl Acad Sci USA* 1997;94:12699–12704. A study of V1 responses in an experimental set-up in which firing rates did not vary, but perceptual experience did. Robust changes in synchrony were observed from one perceptual condition to another. Even if the nature of the perceptual process is questioned, it is remarkable that synchrony in V1 can be so strongly modulated by changes in internal state. [PubMed: 9356513]
118. Logothetis NK, Schall JD. Neuronal correlates of subjective visual perception. *Science* 1989;245:761–763. [PubMed: 2772635]
119. Leopold DA, Logothetis NK. Activity changes in early visual cortex reflect monkeys' percepts during binocular rivalry. *Nature* 1996;379:549–553. [PubMed: 8596635]
120. Braitenberg, V.; Schüz, A. *Cortex: Statistics and Geometry of Neuronal Connectivity*. Springer; Berlin: 1997.
121. White, EL. *Cortical Circuits*. Birkhäuser; Boston: 1989.
122. Sejnowski, TJ. *Parallel Models of Associative Memory*. Hinton, GE.; Anderson, JA., editors. Lawrence Erlbaum Associates; Hillsdale, New Jersey: 1981. p. 189-212.
123. Hopfield JJ, Brody CD. What is a moment? Transient synchrony as a collective mechanism for spatiotemporal integration. *Proc Natl Acad Sci USA* 2001;98:1282–1287. A model for speech recognition in which a set of sensory units responds, a downstream population becomes activated and synchronized, and a third population further downstream responds selectively to the evoked synchrony patterns. The model shows how oscillations generated centrally could confer a functional advantage to a neural circuit. [PubMed: 11158631]
124. Tuckwell, HC. *Introduction to Theoretical Neurobiology*. Vol. 1 & 2. Cambridge Univ. Press; New York: 1988.
125. Koch, C. *Biophysics of Computation*. Oxford Univ. Press; New York: 1999.

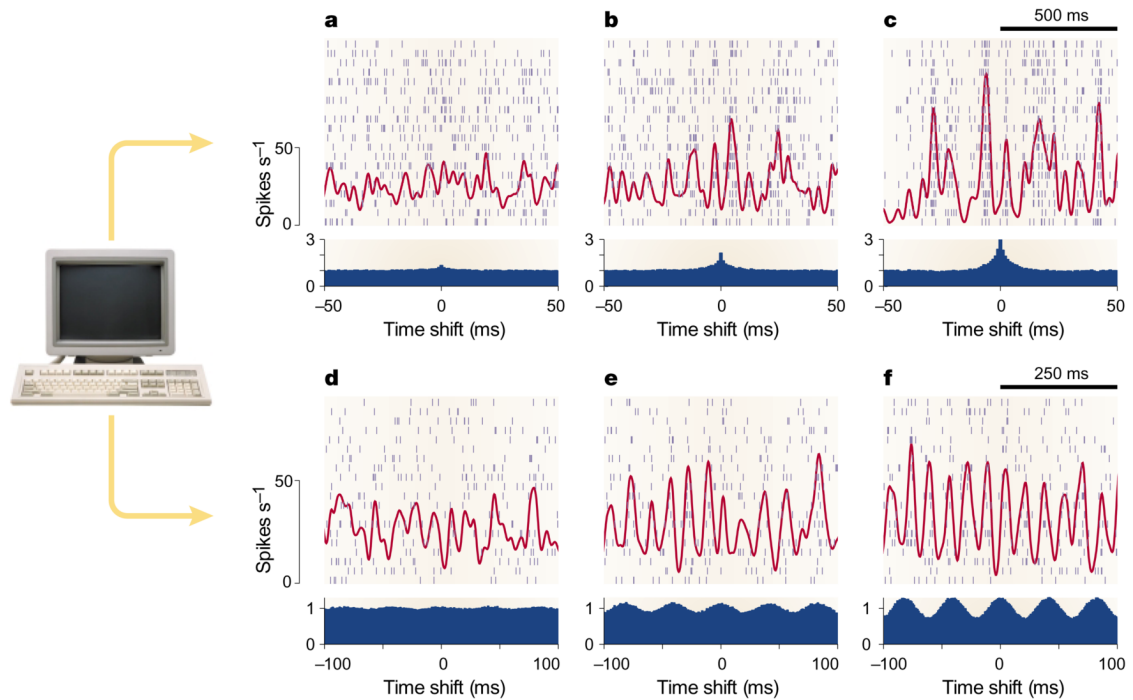


Figure 1. Synthetic computer-generated spike trains with various correlation patterns

Each panel includes a raster plot with 20 simulated spike trains generated simultaneously; each row corresponds to one artificial neuron and each small vertical line to a spike. All neurons were set to fire at a mean rate of 27 spikes s^{-1} and with a CV_{ISI} near 1, as for a Poisson process (the CV_{ISI} is equal to the standard deviation of the interspike intervals divided by their mean). Red traces show instantaneous firing rate or spike density, obtained by smoothing the spike traces with a Gaussian function ($\sigma = 10 \text{ ms}$ for top row; $\sigma = 5 \text{ ms}$ for bottom row) and averaging across neurons. Blue histograms show the average cross-correlation between all possible distinct pairs of units. Cross-correlograms were computed from 20-s segments of simulated data, which included the short segments shown. The y axes are proportional to the probability that two spikes from two different neurons are separated in time by the amount indicated in the x axis. The normalization is such that the probability expected by chance, assuming independence, is set to 1. **a–c**] Each neuron was driven by 1,000 random inputs⁵² and, on average, individual pairs of neurons shared 10% (**a**), 25% (**b**) or 50% (**c**) of those inputs. As the fraction of shared inputs rises, neurons tend to fire closer together in time, which produces larger fluctuations in the average spike density. **d–f**] Here, the neurons fired through independent Poisson processes, but the underlying firing rate was equal to $27(1 + A \sin(2\pi 25t))$, where t is the time in seconds, and was identical for all units. So, the mean rate was still 27 spikes s^{-1} , but it oscillated with a frequency of 25 Hz. The amplitude of the oscillations was $A = 0.25$ (**d**), $A = 0.50$ (**e**) or $A = 0.75$ (**f**). See REF.⁵² for further details.

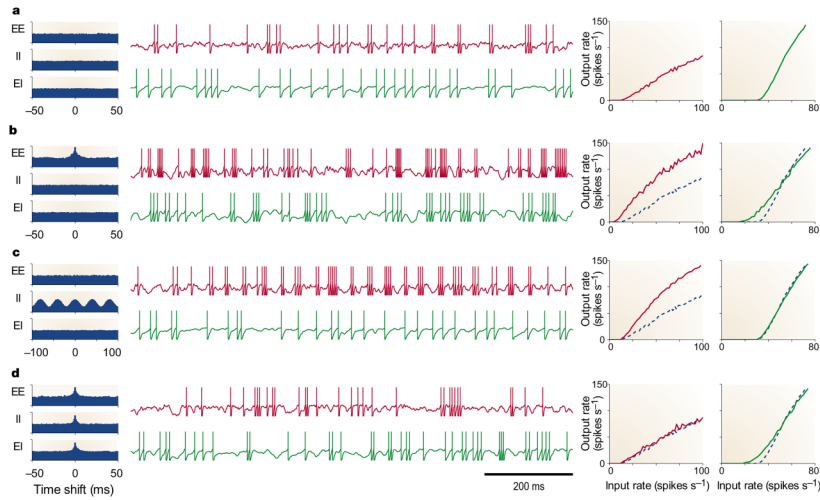


Figure 2. Responses of two model neurons to four input correlation patterns

Histograms on the left show average cross-correlations, like those in FIG. 1, between pairs of excitatory inputs (EE), between pairs of inhibitory inputs (II), and between excitatory–inhibitory pairs (EI). The y axes in the correlograms extend from 0.7 to 1.4. Red and green traces correspond to responses of balanced and unbalanced neurons, respectively, always driven by 160 excitatory and 40 inhibitory inputs. The rate of inhibitory inputs was always 1.7 times the excitatory rate. In the middle traces, all excitatory inputs fired at 42 spikes s^{-1} . In the plots on the right, the mean firing rate of the excitatory inputs varies along the x axes, and the y axes correspond to the output firing rates of the two postsynaptic model neurons. All responses were obtained using leaky integrate-and-fire models (see BOX 1). **a**| All input spike trains were independent. In the middle traces, both postsynaptic neurons are shown to fire at about 30 spikes s^{-1} . **b**| Excitatory inputs were synchronous, with 10% shared inputs, as in FIG. 1a. Balanced and unbalanced neurons fired at 67 and 45 spikes s^{-1} , respectively. **c**| Inhibitory inputs oscillated with an amplitude equal to 50% of the mean rate, as in FIG. 1e. Balanced and unbalanced neurons fired at 59 and 30 spikes s^{-1} , respectively. **d**| All inputs were synchronous, with 10% shared inputs. Balanced and unbalanced neurons fired at 31 and 41 spikes s^{-1} , respectively. For comparison, broken lines in the input–output rate plots (**b–d**) are the curves obtained with independent inputs (**a**). The balanced neuron is much more sensitive to correlations than the unbalanced one.

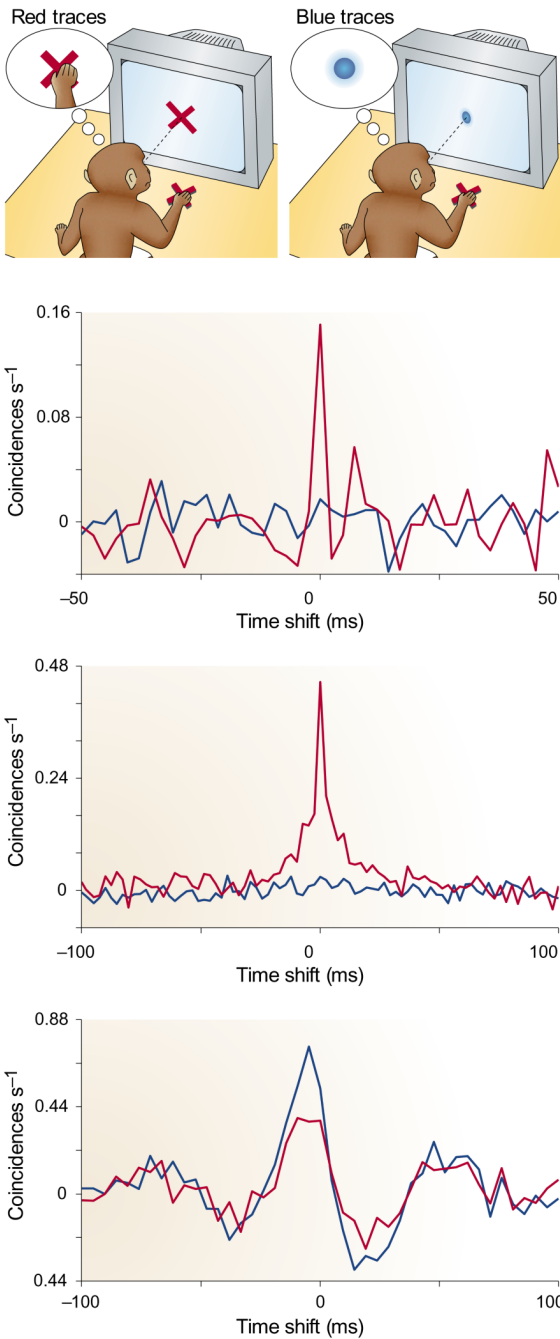


Figure 3. Cross-correlation histograms, with and without attention, from pairs of neurons recorded in the secondary somatosensory cortex of awake monkeys

The y axis indicates the rate of spike coincidences (defined as two spikes within 2.4 ms of each other) when the spike trains from the two neurons are shifted in time by the amount shown on the x axis. These correlograms have been normalized so that a zero rate corresponds to independent spike trains. The three panels correspond to three different pairs. Red traces were calculated from trials in which the monkey paid attention to a tactile stimulus (the cross on the table); blue traces were calculated from trials in which the same tactile stimulus was presented, but the monkey had to pay attention to a visual stimulus on the screen. In the top two examples, more synchrony was observed when attention was focused on the tactile stimuli; this was the

more prevalent effect. An example of lower synchrony with attention on the tactile stimulus — the less frequent effect - is shown in the lower plot. Data modified from REF. ¹⁰⁴ and kindly provided by P. Steinmetz.

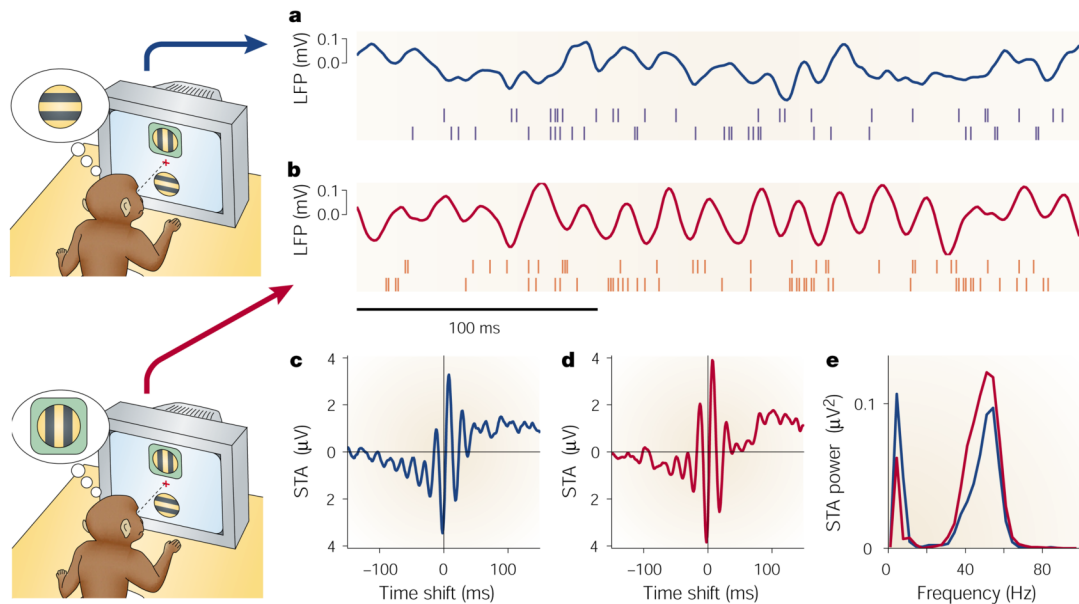


Figure 4. Attention induces changes in synchrony in the visual cortex

Data shown are from experiments in which two visual stimuli were presented, one inside and one outside the receptive field of a neuron in area V4. In the schematics, the green box represents the receptive field: this was not presented on the screen in the trials. Red traces correspond to attention directed inside the receptive field of the recorded neuron; blue traces correspond to attention directed outside. Stimuli were the same in the two conditions. **a** and **b**| The continuous traces show the stimulus-driven local field potentials (LFPs). The spikes below were recorded simultaneously from different electrodes. **c** and **d**| Spike-triggered averages (STAs) computed during the stimulus presentation period. The STA corresponds to the average LFP waveform that is seen at the time of a spike. The y axes indicate the mean LFP; the x axes indicate time relative to the occurrence of a spike. **e**| Power spectra of the two STAs shown in **c** and **d**. When attention is focused inside the receptive field, the recorded neuron tends to fire more in phase with the frequency components around 50 Hz, and less so with respect to the frequencies around 10 Hz. Data modified from REF. ¹¹³ and kindly provided by P. Fries.