HHS Public Access

Author manuscript

Nat Neurosci. Author manuscript; available in PMC 2015 November 11.

Published in final edited form as:

Nat Neurosci. 2015 May ; 18(5): 736–743. doi:10.1038/nn.3979.

Global network influences on local functional connectivity

Adam C. Snyder[1,2], **Michael J. Morais**[1,3], **Cory M. Willis**[1], and **Matthew A. Smith**[1,2,3,4,#]

[1]Dept. of Ophthalmology, Univ. of Pittsburgh, Pittsburgh, PA, USA

[2]Center for the Neural Basis of Cognition, Univ. of Pittsburgh, Pittsburgh, PA, USA

[3]Dept. of Bioengineering, Univ. of Pittsburgh, Pittsburgh, PA, USA

[4]Fox Center for Vision Restoration, Univ. of Pittsburgh, Pittsburgh, PA, USA

Abstract

A central neuroscientific pursuit is understanding neuronal interactions that support computations underlying cognition and behavior. Although neurons interact across disparate scales – from cortical columns to whole-brain networks – research has been restricted to one scale at a time. We measured local interactions through multi-neuronal recordings while accessing global networks using scalp EEG in rhesus macaques. We measured spike count correlation, an index of functional connectivity with computational relevance, and EEG oscillations, which have been linked to various cognitive functions. We found a surprising non-monotonic relationship between EEG oscillation amplitude and spike count correlation, contrary to the intuitive expectation of a direct relationship. With a widely-used network model we replicated these findings by incorporating a private signal targeting inhibitory neurons, a common mechanism proposed for gain modulation. Finally, we report that spike count correlation explains nonlinearities in the relationship between EEG oscillations and response time in a spatial selective attention task.

Introduction

Action potentials, or spikes, are widely held to be the computational currency of the brain. Decades of research have identified numerous ways in which the activity of individual neurons is related to stimuli in the outside world and to our perception of those stimuli. Cognitive and perceptual processes, however, are not the product of any individual neuron's activity, but instead are network-level phenomena in which groups of neurons act in concert. These phenomena can be studied by local recordings of many neurons simultaneously, via a multielectrode array or imaging of a voltage-sensitive dye, or through more global measures of neuronal activity, such as functional magnetic resonance imaging (fMRI) or

#Address correspondence to: Department of Ophthalmology University of Pittsburgh Eye and Ear Institute 203 Lothrop St., 9th Fl. Pittsburgh, PA, 15213 Tel: (412) 647-2313 smithma@pitt.edu.

Conflict of interest: "The authors declare no competing financial interests."

Author contributions: ACS and MAS designed the experiments, analyzed the data, and wrote the manuscript. ACS and CMW conducted the experiments. ACS and MJM conducted the network model simulations. MAS supervised the project.

electroencephalography (EEG). A complete understanding of the neural basis of perception and behavior requires a bridge across these levels of analysis.

Investigations of small populations of neurons have focused on pairwise interactions such as correlated variability in firing rates from trial to trial, termed spike count correlation (or "noise" correlation; r_{sc}), a measure of functional connectivity with known implications for coding[1]. Several recent investigations of spike count correlation have found that it is highly structured[2-4] and modulated by cognitive and perceptual context[5-9]. However, identification of the signals that generate these dynamics has proved elusive, instead relying on speculation about the large-scale networks involved. Previous attempts to measure the interdependence of activity between brain areas have made tantalizing suggestions that spiking activity can be related to oscillations supported by large-scale networks[10,11], but generally speaking such networks have proved inaccessible to micro-scale methods.

The most widely used methods for measuring large-scale network activity are fMRI and EEG. With these methods, an explanatory gap persists as to how the large-scale signals are related to the spiking activity of small populations of neurons. A fair amount of investigation has been directed at linking spiking activity to the fMRI blood oxygenation level dependent (BOLD) response[12,13], but far less research has aimed to relate spiking activity and EEG[14]. The EEG is thought to reflect the post-synaptic potentials in the apical dendrites of pyramidal cells due to their mutual alignment that allows summation of electric fields[15]. The strength of the signal is related both to the magnitude of the post-synaptic activity, as well as its coherence: post-synaptic currents with low spatio-temporal coherence tend to destructively interfere at the level of the scalp[15,16]. The common postsynaptic activity that drives variability in the EEG signal likely also generates spike count correlation across neurons.

We sought to test whether EEG oscillations index the coordination of the spiking activity of the underlying neuronal population, using simultaneous recordings of evoked and spontaneous activity at the scalp and in the cortex of behaving macaque monkeys. We found that oscillations at the level of the EEG do, in fact, relate to spike count correlation, but they do so in a non-monotonic fashion. However, we found that a variation of a widely used simple network model incorporating excitatory and inhibitory subpopulations[17,18] can account for this surprising relationship. Finally, we report that knowledge of the non-monotonic relationship between EEG oscillations and spike count correlation can explain the connection between EEG oscillations and performance on a spatial selective attention task.

Results

We simultaneously recorded EEG from the scalp along with spiking activity from a "Utah" microelectrode array implanted in area V4 of two macaque monkeys (Fig. 1a) performing a fixation task. We started by isolating a snippet of EEG around the time of each recorded spike. We then Fourier-transformed those EEG snippets to determine the phase and amplitude of the global oscillations relative to the local spiking activity. In order to measure spike count correlation among neurons, we created surrogate spike trains (subsets of the full

recording) in which the amplitude and phase of the EEG were within a specified range. This allowed us to ascertain how spike count correlation related to particular properties in the EEG. A simplified understanding of our approach can be visualized by thresholding the envelope of the continuously filtered EEG signal and apportioning spikes accordingly (Fig. 1b). On the full data set, our actual method employed the spike-triggered Fourier coefficients to create surrogate spike trains (see Methods).

Relationship between EEG and spike count correlation

Surrogate data from a representative session illustrate one of the more robust and surprising effects we found (Fig. 2). We had predicted that spike count correlation would be directly related to oscillation amplitude, based on the assumption that greater amplitude oscillations are a result of more coherent input to the underlying brain area. In contrast to this prediction, we found that spiking activity during the intermediate amplitudes of alpha oscillations not only had relatively less spike count correlation than spiking activity during higher-amplitude oscillations (which we predicted) but also had less spike count correlation than spiking activity during lower-amplitude oscillations. In other words, the magnitude of spike count correlation followed a U-shaped relationship with EEG oscillation amplitude. While the overall magnitude of spike count correlation varied somewhat from session to session, the U-shaped relationship between oscillation amplitude and spike count correlation was highly reliable (Fig. 3a). This relationship was statistically significant ($p = 0.01$), as determined by a permutation test consisting of randomly shuffled data subjected to the identical analysis procedure (Fig. 3b; see Methods).

In addition to the U-shaped relationship between EEG oscillation amplitude and spike count correlation, we also found a sinusoidal relationship between EEG oscillation phase and spike count correlation in the ‘spontaneous’ activity condition (Fig. 3a). This effect was seen for all frequencies. Changes in correlation state with the phase of ongoing oscillations could potentially account for previous observations in both the sensory[19,20] and motor[21] domains that the effectiveness of information transmission can be modulated by oscillation phase. However, in parallel to the relationship between EEG phase and spike count correlation that we observed, we also found a sinusoidal relationship between EEG oscillation phase and the firing rate of the population (Fig. 4), as has been previously reported for the gamma band[14,22]. In our case, the phase with the lowest average firing rate was also the phase with the lowest spike count correlation. Varying rate complicates the interpretation of spike count correlation measurements, because previous modeling work has shown that the input correlation of a pair of neurons is underestimated by spike count correlation when firing rates are low[23,24]. This was not a potential issue for the analysis of oscillation amplitudes, because our approach of binning spikes based on amplitude deciles ensured equal numbers of spikes (and therefore equal average rates) across bins. To determine whether changes in firing rate with EEG phase could account for the effects on spike count correlation, we repeated the analysis with phase bins chosen such that the spike counts in each bin were equal[22]. We did not find any relationship between EEG phase and spike count correlation with this equal-count binning (Fig. 3a), which suggests that nearly all of the modulation of spike count correlation that we observed with respect to EEG phase could be accounted for by variations in firing rate.

Our initial analysis focused on alpha oscillations because of their historical prominence in the EEG literature[25,26], but we found that the U-shaped relationship between EEG oscillation amplitude and spike count correlation was evident across frequencies. Each of the frequency bands of interest revealed a U-shaped relationship between amplitude and correlation, but the strength of this relationship was directly related to frequency: low-frequency oscillations showed relatively less modulation of spike count correlation with amplitude and high-frequency oscillations showed relatively greater modulation of correlation with amplitude (Fig. 5a). The magnitude of the relationship between EEG oscillation amplitude and spike count correlation also displayed spatial specificity, with the strongest relationships seen for EEG measured at the electrode nearest the array and at the electrode diametrically across skull (Fig. 5b). This spatial specificity was most evident for the frequencies with the greatest effect sizes. The antipodal topographical pattern we observed is generally consistent with both the positive pole and negative pole of a single dipole source located near the array reflecting the amplitude-correlation relationship.

All of the above analyses were performed based on spontaneous spiking and EEG activity recorded during fixation. When a stimulus was presented in the receptive field of the neurons recorded on the array, we saw a similar U-shaped relationship between EEG oscillation amplitude and spike count correlation to what we found for the spontaneous condition (Fig. 6a). However, we did not observe a relationship between correlation and EEG phase in the evoked condition, indicating that the visual stimulus abolished the phase-locking between the spikes and the EEG. Analysis of the firing rate of our pairs (Fig. 6b) revealed that we were able to control for responsiveness in both phase and amplitude bins (note the much higher firing rates in the evoked condition compared with spontaneous data in Fig. 4).

Although the subjects were required to maintain fixation, small (< 1°) eye movements were possible, and prior reports have suggested that microsaccades are related to EEG oscillations[27,28]. In addition, variability in gaze position between trials could lead to spurious measurements of “noise” correlation due to similarities in tuning (i.e., “signal” correlation), since the image in the receptive field would shift slightly from trial to trial (although this caveat would not apply to our spontaneous data). Thus, we tested whether small eye movements could confound our results by dividing each session into two sets of trials using a median split based on the variance of the gaze position over the trial. We found that spike count correlation was nearly identical in the more variable and more stable gaze position trials (all p’s > 0.17, median: $p = 0.77$), which indicated that the relationship between EEG oscillations and spike count correlation reported here was not due to eye movements.

Relationship between local field potentials and spike count correlation

We wondered whether the effects we observed were specific to the EEG, because there is a demonstrated relationship between the EEG and local field potentials (LFPs)[16]. Although there is unquestionable value in relating EEG (the predominant electrophysiological method in humans) to spiking activity, an important question is whether the same mechanisms operate at the scale of hundreds of micrometers (LFP) as at the scale of centimeters (EEG).

We therefore performed the same analysis on the stimulus-evoked condition using the LFP from a randomly chosen electrode in the array (3 of 18 datasets were excluded because of errors saving the LFP data). We found that the LFP-r_{sc} relationship closely mirrored the EEG-r_{sc} relationship both qualitatively and quantitatively. To assess this we measured the Pearson's correlation between the data points comprising the two sets of curves, and we found significant positive correlations at each frequency (all r's > 0.72, all p's < 0.001, one-sample t-test). The magnitude of the modulation (max – min) of r_{sc} did not differ significantly whether based on oscillations in the EEG or the LFP ($p = 0.24$ across all frequencies, paired-samples t-test). These results are consistent with the EEG and LFP both reflecting the same mechanisms that modulate spiking activity.

Potential role of cross-frequency interactions in EEG

We were surprised to find a U-shaped relationship of spike count correlation to oscillation amplitude. Our guiding prediction had rather been that spike count correlation would increase monotonically with oscillation amplitude, under the reasoning that greater-amplitude oscillations would be generated by more coherent postsynaptic currents in the underlying brain area, and that more coherent postsynaptic currents would lead to correlated spike output. This reasoning can account for the pattern of results we observed at high amplitudes, where amplitude and spike count correlation were directly related, but it cannot explain the pattern of results we observed at low amplitudes, where amplitude and spike count correlation were inversely related.

Since cross-frequency EEG interactions have been linked to multiunit activity[14], we wondered whether dependencies between the frequency bands might lead to the non-monotonic effect we observed. If the amplitudes of some of our frequency bands tended to be inversely related to each other, then this could potentially account for the non-monotonicity while still remaining consistent with the essentially monotonic mechanism that guided our initial predictions. In other words, two oscillations that consistently trade-off in amplitude from trial to trial could be driving changes in spike count correlation, giving merely the appearance of a non-monotonic mechanism. This was not the case, however, as we calculated the Pearson correlation of amplitude between all pairings of the six frequencies of interest and found that no pair of frequencies was inversely related. We did find significant positive correlations between some frequencies, however. In particular, lower frequencies (delta, theta and alpha) were correlated with each other in the range of r = 0.02 to 0.08, and higher frequencies (alpha, beta, low gamma and high gamma) were correlated with each other in the range of r = 0.08 to 0.26. The highest frequency pairs had the greatest correlation, which is consistent with the diminishing frequency resolution achieved as frequency increases. In general, it is not surprising that nearby frequencies were correlated, simply due to their proximity. The only significant correlation we observed between distant frequency bands was between delta and low gamma (r = 0.05, $p = 0.002$), which is noteworthy because interactions between these two frequency bands in particular has recently been linked to multiunit activity[14]. To summarize, the main finding that none of the pairs of frequencies was inversely correlated rules out amplitude trade-offs between frequency bands as a potential cause of the non-monotonic amplitude-correlation relationship.

Neural network modelling

We next asked whether a network of interconnected neurons could produce the non-monotonic relationship by virtue of the complexity of the dynamics that a neural network can generate. We modeled a balanced network of excitatory and inhibitory neurons driven by an external population of random, doubly-poisson neurons and an additional applied current[17]. The parameters for the model were set based on the previous literature[17] and we did not modify them; rather, we only modified the input currents that we applied to this "off-the-shelf" model. We used band-filtered (alpha band, 8 – 12 Hz) white noise signals for our applied current to model a coherent oscillatory input to the network. We modeled the resulting spike train for each neuron in the network, and we modeled the corresponding EEG as the summed post-synaptic potential of the excitatory neurons in the network. We divided our set of modeled trials into decile bins by EEG alpha amplitude, and calculated the average spike count correlation within each bin.

When we applied a single oscillatory current ranging from 0 – 0.3 nA in amplitude identically to all the neurons in the network, we found that correlation between model excitatory neurons was on the order of 1/N ($r_{sc}^{EE} \approx 0.001$; Fig. 7a), consistent with prior calculations[17] demonstrating the ability of a balanced network to cancel input correlations. When much larger-amplitude oscillation amplitudes were included (> 0.3 nA), correlation increased monotonically with alpha amplitude as the spiking activity became unnaturally coherent with the driving oscillation (not shown). Because the spiking behavior induced by these large-amplitude input oscillations was physiologically implausible, we restricted our subsequent analyses to input oscillations below that range. To summarize, the simple balanced neural network with a single input shared by the entire network cannot reproduce the non-monotonic behavior we observed in our experiment.

We next tested whether applying separate input oscillations to the excitatory and inhibitory subpopulations of our network might produce a non-monotonic relationship between EEG amplitude and correlation. We were guided in this hypothesis by the observation that feedback signals tend to target these two cell types differentially. For example, in V4 attention-dependent modulation of the activity of putative inhibitory neurons is stronger than that for putative excitatory neurons (using waveform shape to separate neurons into subclasses)[29]. When we applied a separate input oscillation to the inhibitory subpopulation (0 – 0.2 nA), in addition to the global input oscillation that was shared by all the neurons, this resulted in a U-shaped relationship between EEG alpha amplitude and spike count correlation that matched our experimental data (Fig. 7a). This effect was robust over the entire range of input amplitudes of the model we tested that produced physiologically plausible spike trains (Fig. 7b). This indicates that a balanced neural network is sufficient to generate the non-monotonic relationship we found if the inhibitory subpopulation receives a separate input oscillation in addition to that shared with the population at large.

It is notable that our neural network model is spatially homogenous – that is, neurons are connected to each other randomly without regard for their relative spatial positions. Functional connections among real neurons are known to decay with distance and tuning similarity[3,4] and in our own data, we see that the relationship between EEG and spike count correlation is strongest for scalp locations closest to the multielectrode array. Extensions of

this model will need to incorporate these distance effects as well as tissue conduction in order to more fully model the EEG. An additional important question is how critical it was that we applied band-limited input to our neural network model, because such input might not be ecologically typical (although intrinsic filtering properties of dendritic membranes may make individual synapses particularly sensitive to limited frequency bands[30,31]). Two observations suggest that band-limited input is not critical. First, although our input oscillation was band-limited to the alpha range, the resulting modeled EEG had a broad-band power spectrum by virtue of the internal network dynamics (Supplemental Fig. 1). Second, we also observed a U-shaped relationship between correlation and EEG amplitude at all other resolvable frequency bands in the model (delta was below our resolution) that were qualitatively similar to what we observed *in vivo*. Because both the field potentials and the effects on correlation generated by our neural network model were not limited to the input frequency band, this neural network model is consistent with the essentially broadband field potentials that have been previously linked to spiking activity *in vivo*[32,33].

Spatial attention task

In a variety of contexts, psychophysical performance has been found to follow a U-shaped relationship with EEG oscillation power[34-37], much like the relationship we found between the EEG and spike count correlation. To date, interpretations of this observation have been speculative, typically invoking the notion of stochastic resonance. Our results, however, hint at an alternative explanation – that changes in spike count correlation could serve as the computational basis for this effect, as decorrelation has been previously linked to improved performance in an attention task[5], potentially due to improved signal averaging during the decorrelated state[1]. We therefore sought to test this hypothesis by performing an additional experiment.

Two animals were trained on a selective attention task in which they maintained central fixation while two drifting gratings were presented; one in the receptive field area, and one in the mirror symmetric position in the opposite hemifield. The task was to detect a change in the drift speed of one of the two gratings and to make an eye movement to the grating, if any, that changed, or to maintain fixation until the end of the trial (1.2 s) otherwise. On most trials a briefly flashed visual cue indicated the stimulus location that was most likely to contain the target, and analysis of behavioral data indicated that this cue resulted in a robust enhancement of performance at the cued location.

For the analysis of this experiment, we focused on the amplitude in the alpha frequency band of the EEG, which has been most consistently linked to selective attention[26]. Our goal was to determine whether the underlying spike count correlation mediated a U-shaped relationship between EEG alpha amplitude and response time. We divided trials into six bins based on alpha amplitude, and calculated spike count correlation over the trials in each bin. Following the procedure outlined by Baron and Kenny[38], we used a regression analysis to determine the degree to which alpha band amplitude could explain variance in behavioral response time to targets contralateral to the array, and the degree to which this relationship was mediated by spike count correlation. We found that response time indeed followed a U-shaped relationship with alpha amplitude measured at the parieto-occipital electrode over the

array, with the fastest response times observed at intermediate alpha amplitudes (Fig. 8). Inclusion of a linear term for correlation improved the variance-accounted-for by the regression equation (by 118% for Monkey B, $p = 0.011$; 240% for Monkey R, $p = 0.042$). Most critically for our predictions, we found that inclusion of spike count correlation in the regression equation led to a decrease of the quadratic term for alpha amplitude (by 15% for Monkey B, $p = 0.047$; 16% for Monkey R, $p = 0.030$). In other words, the U-shaped relationship between response time and alpha amplitude was substantially mediated by the degree of spike count correlation in the underlying area. To illustrate the effect of this mediation, we used the resultant regression equation to model the predicted response times for each alpha amplitude while adjusting for the effect of spike count correlation, which resulted in a marked flattening of the relationship (Fig. 8). This relationship with behavior was specific to the alpha band (for all other frequency bands for both subjects: all p's > 0.186).

Discussion

We found that functional connectivity among spiking neurons (measured by spike count correlation) followed a U-shaped relationship with EEG spectral power measured at the scalp. This surprising result is counter to the basic intuition that increasing EEG power would reflect increasingly coherent input to the underlying neurons, and in turn lead to an increase in spike count correlation. The relationship we found was topographically specific, with the EEG electrode sites nearest the multielectrode array showing the strongest effects, but with additional contributions from the diametrically opposing sites, consistent with the two poles of a dipole source. We also found the general pattern to be robust across frequencies, although higher frequencies had stronger amplitude-correlation relationships. Our results provide a powerful insight into the relationship between local neuronal populations and global brain activity, and advance our understanding of how the functional connectivity of small scale networks is modulated by the context of global brain states, which is a fundamental aspect of brain function.

To further understand our results, we employed a popular computational model using a balanced network of excitatory and inhibitory neurons[17,18]. In its basic form, this model effectively cancels correlations in common input and produces an asynchronous network of neurons with near-zero mean correlations. We first considered the simplest modification to this model, a global excitatory oscillatory input applied to all of the neurons, which failed to produce the behavior we observed *in vivo* and instead produced a monotonic increase in correlations with excitatory oscillations of sufficient magnitude. We then considered a second addition to the model, in which we delivered an independent oscillation exclusively to the inhibitory subpopulation of neurons. Although the spatially homogenous organization of dendrites in inhibitory interneurons[39] would largely hide such an oscillation at the level of the EEG, it produced a non-monotonic relationship between EEG power and spike count correlation that was an excellent match to the physiological data. That a relatively simple and well-studied model reproduces our *in vivo* findings while implicating a specific subpopulation of neurons provides a powerful link to phenomena proposed to depend on selective activation of inhibition, such as attention[29], sensory binding[40] and memory formation[40].

For example, several recent studies have reported an inverted-'U' effect of EEG oscillations during attention-demanding detection tasks, with performance varying non-monotonically with EEG power[34-37]. Without concurrent population recordings, such studies are left to speculate high-level functional explanations of behavior, invoking phenomena such as stochastic resonance. Our data confirm that the inverted 'U' shape relating EEG oscillations to performance during attention in these studies can be explained by the decorrelation in spiking observed at intermediate oscillation amplitudes, which reflects a change in processing state of the neuronal population[1]. Moreover, our model posits a key role for inputs to inhibitory neurons underlying this non-monotonic effect. Indeed, modulation of inhibitory subpopulations by long-range feedback signals appears to be a key mechanism to adjust gain during selective attention[29] and may be mediated by cholinergic receptors on inhibitory neurons[41,42]. Our observations thus connect studies of attention in humans, in which changes in EEG alpha power are the critical observation[26], and in monkeys, where a reduction in correlated variability with attention has been reported[5,6]. This link provides explanatory power for interpreting electrical potentials at the scalp in terms of the activity of groups of local neurons, and in linking studies of spatial attention and its mechanisms across species and electrophysiological methods.

While knowledge of the relationship between spiking activity and the EEG has clear benefits for future EEG research, the inferential power actually moves in both directions. As advances in multielectrode array technology have spurred a shift in focus from the first-order spiking statistics available from single neurons to higher-order spiking statistics available in population recordings, a key question has emerged: What are the signals that give rise to correlated spiking variability, its structure[3,4], and its dynamics[5-8,43]? One hypothesis is that weakly correlated noise in the sensory input is inherited in a feedforward fashion across the cortex[44]. Recent observations that the degree of correlated variability can be modulated in an exquisitely complex and goal-directed fashion[8] suggest an alternative hypothesis, however: that correlated variability is due at least in part to feedback input from the large-scale networks that support cognitive processes involving coordination between brain areas. Researchers have attempted to account for the influence of the broader network context using intracranial local field potentials[10,45], which reflect the summed activity of all the neurons within a sub-millimeter volume[46] of cortex. Another approach has leveraged local recordings of neuronal populations to infer the properties of modulatory signals driving correlated variability[47]. Both of these methods are a far cry from the truly global scale networks probed by EEG recordings. For researchers studying complex cognitive processes at the level of small populations of neurons, the ability to relate their findings to studies of whole-brain EEG signals is a vital advancement.

Our study uncovered a simple and yet non-intuitive relationship between correlated variability measured at the level of a small population of neurons and the amplitude of EEG oscillations measured at the scalp. Both correlated variability and EEG oscillations have been demonstrated to have important functional consequences across a variety of brain processes, and here we establish a definitive link between these measurements. Our modification of a simple and widely employed model suggests that selective feedback to inhibitory subpopulations, which are largely invisible to EEG, is a critical mechanism for coordination of local neuronal populations by global brain networks. These findings

highlight the critical role of observations across scale in revealing the mechanisms by which neuronal networks give rise to complex perceptual and cognitive experiences.

Methods

Experimental procedures were approved by the Institutional Animal Care and Use Committee of the University of Pittsburgh.

Subjects

We implanted one, 100-electrode "Utah" array (Blackrock Microsystems) in V4 in each of three adult male rhesus macaques (*Macaca mulatta*). We implanted the arrays in the right hemisphere for Monkeys W and B, and in the left hemisphere for Monkey R. The basic surgical procedures have been described previously[4], and were conducted in aseptic conditions under isoflurane anesthesia. In addition to the microelectrode arrays, the animals were implanted with a titanium head post to immobilize the head during experiments. We recorded neurons with receptive fields centered 3.16°, 5.66° or 10.23° from the fovea in the lower visual field of each animal.

Behavioral tasks

We presented visual stimuli with custom software written in MATLAB (The MathWorks) using the Psychophysics Toolbox extensions[48-50]. Monkeys W and B performed the following tasks. We trained the subjects to maintain fixation on a 0.6° blue dot at the center of a flat-screen cathode ray tube monitor positioned 36 cm from their eyes. The background of the display was 50% gray. We measured the monitor luminance gamma functions using a photometer and linearized the relationship between input voltage and output luminance using lookup tables. In the 'spontaneous' task, subjects were trained to maintain fixation on the central dot for 2 s, at which time the fixation point would be moved 11.6° in a random direction and the animal received a liquid reinforcement for making a saccade to the new location. No other stimuli were presented for the spontaneous task. The 'evoked' task was identical, except we presented a drifting sinusoidal grating (100% contrast, maximum luminance of 145 cd/m^2) in the aggregate receptive field (RF) area of the neurons recorded on the microelectrode arrays during the 2 s fixation interval. Both subjects completed the spontaneous task (Monkey B: 14 sessions; Monkey W: 10 sessions) prior to performing the evoked task (Monkey B: 10 sessions; Monkey W: 8 sessions).

Microelectrode array recordings

Signals from the microelectrode arrays were band-pass filtered (0.3 – 7500 Hz), digitized at 30 kHz and amplified by a Grapevine system (Ripple). Signals crossing a threshold (periodically adjusted using a multiple of the root-mean-squared [RMS] noise for each channel) were stored for offline analysis. These waveform segments were sorted using an automated clustering algorithm[51] followed by manual refinement using custom MATLAB software[52] (available at http://www.smithlab.net/spikesort.html), taking into account the waveform shapes and interspike interval distributions. After sorting, we calculated the signal-to-noise (SNR) ratio of each candidate unit as the ratio of the average waveform

amplitude to the standard deviation of the waveform noise[52]. Candidates with an SNR below 2.5 were discarded.

Receptive field mapping and tuning curves

Prior to beginning the experiments, we mapped the RF areas of the units recorded on our arrays by presenting small (~1°) sinusoidal gratings at four orientations positioned one at a time on the vertices of a lattice in the likely RF area per the anatomical location of the implant. After inspecting the responses of the units to these small probe stimuli, we picked a stimulus size and position to roughly cover the aggregate RF area. For Monkey B this was 5.87° diameter centered 4.00° below and 4.00° to the left of fixation, for Monkey W this was 4.70° diameter centered 1.18° below and 2.94° to the left of fixation, and for Monkey R this was 5.87° diameter centered 8.94° below and 4.99° to the right of fixation. Throughout this report we use the term 'RF' to refer to this aggregate area. We next measured tuning curves for the recorded units by presenting sinusoidal gratings to the RF area in four orientations and at a variety of spatial and temporal frequencies. For each subject we chose a compromise temporal and spatial frequency that evoked the maximum response from the array as a whole.

EEG recordings

We recorded EEG from 8 Ag/AgCl electrodes (Grass Technologies) adhered to the scalp with conductive paste (see Fig. 1a for positioning). Signals were referenced online to the head post, digitized at 1 kHz and amplified by a Grapevine system (Ripple) and low-pass filtered online at 250 Hz. Data were first divided into trial epochs time-locked to the onset of the grating stimuli (−0.5 to 3 s); analyses focused on subsets of this interval. Trials containing excessive muscle activity or transient artifacts were rejected using an automatic threshold criterion (± 300 μV) and visual inspection. After artifact rejection, data were rereferenced to the average activity across all electrodes.

Eye tracking

We tracked the gaze of the subjects using an infrared eye tracking system (EyeLink 1000; SR Research). Gaze was monitored online by the experimental control software to ensure that subjects maintained fixation within 1.17° of the central fixation point during each trial.

Data analysis

To examine the relationship between spectral properties of the EEG and the spike count correlation of the neurons recorded on our microelectrode arrays, we constructed surrogate spike trains characterized by particular phase and amplitude properties (similar to the procedure of Womelsdorf et al.[22]). Specifically, we focused on the six frequency bands analyzed by Musall et al.[16], who compared spiking activity to a single EEG ring electrode: 2 – 4 Hz (delta), 4 – 8 Hz (theta), 8 – 12 Hz (alpha), 12 – 30 Hz (beta), 30 – 60 Hz (low gamma) and 60 – 100 Hz (high gamma). Average power spectra for Monkeys B and W, measured over the 2 s fixation periods, are illustrated in Supplementary Figure 2. For each spike, we calculated the fast Fourier transform of the 500 ms Hanning-windowed segment of each channel of EEG data centered at the time of the spike. Note that spikes occurring in the

first or final 250 ms of a trial were therefore excluded because the window would have overlapped the trial boundary. These spike-triggered spectra were used to compute the average amplitude across each frequency band around the time of each spike, as well as the phase of the center frequency for each frequency band around the time of each spike. We next created virtual spike trains for each combination of 10 amplitude ranges and 12 phase ranges measured at each EEG channel and for each frequency band. We divided amplitude measurements into 10 deciles that each contained an equal number of spikes. We divided phase into 12 bins, each subtending 60 degrees of phase, with 50% overlap between adjacent bins. In a separate control analysis, we used six, non-overlapping phase bins spaced to equate as closely as possible the number of spikes in each bin for each session. Each virtual spike train spanned 1.5 s of real time, but consisted only of spikes that occurred when the amplitude and phase of a given frequency band measured at a given electrode was within the specified ranges. We summed the spikes of these surrogate spike trains to produce a single spike count for each 2 s trial. We normalized the spike counts of each neuron for each grating orientation by z-scoring. We then computed the spike count correlation (r_{sc}) as the Pearson product moment correlation coefficient of normalized spike counts for all pairs of neurons recorded simultaneously from separate microelectrodes. Correlation was measured separately for each recording session.

Statistical analysis

We used curve-fitting combined with a non-parametric resampling procedure to quantify our statistical confidence in the relationships between spectral properties of the EEG and spike count correlation. We first fit curves (described below) to our observed data in the ordinary least-squares sense, and quantified the strength of the relationship according to the fitted function. We then randomly shuffled the pairings between spikes and EEG data windows, so that the relationship between spikes and EEG was random, which represents the null hypothesis. We repeated the analysis making surrogate spike trains for specific EEG spectral properties, this time using the randomly-shuffled data. We then fit curves to the results from the shuffled-data analysis. The data shuffling and reanalysis were iterated a total of 1000 times to create a distribution of resultant curve fits under the null hypothesis. The proportion of this distribution that exceeded the observed value was the p-value for the claim that the relationship we observed was stronger than would be expected by chance. We considered $p < 0.05$ to be significant. All statistics were performed across sessions, treating the data from each day as a unit of observation. We viewed this to be the most conservative approach, when compared with the alternative of adding together the pairs across days and treating each pair as an observation.

To characterize the U-shaped relationship between EEG oscillation amplitude and spike count correlation, the function we fit to the data was a ratio of two quadratics. This is the simplest function that can be non-monotonic and can also have a bounded range (since correlation must be bounded at least between −1 and 1). We collapsed across phase for this fit, and considered only amplitude. We measured the coefficient of determination, adjusted for the number of free parameters (i.e., $\overline{R}^2$), for both the rational function fit, as well as for a simple linear fit. We quantified the strength of the non-monotonic relationship as the change in adjusted coefficient of determination from the linear fit to the rational fit ($\Delta\overline{R}^2$).

To test whether the presence of small eye movements could have confounded our results, we performed a median split of trials for each experimental session based on gaze position variance, and repeated the analysis described in the "Data analysis" section above. Since gaze position is two-dimensional, we derived a scalar measure of gaze position variance using the generalized variance:

$$\Sigma = \sqrt{det\,(X^{\mathrm{T}}X)}$$

where X is a $n \times 2$ matrix of gaze position coordinates, and n is the number of samples in the 2 s analysis window. After we repeated our analysis on both halves of trials for each session, we tested for differences between trials with relatively stable gaze positions and trials with relatively variable gaze positions using independent samples t-tests at each amplitude of each frequency for both conditions (a total of 120 tests). We did not adjust the reported p-values for the familywise error rate.

To test for amplitude covariance between the frequency bands of interest, we measured the amplitude of each of the six frequency bands of interest by windowing each 2 s trial with a Hann function and then applying the fast Fourier transform (FFT). We calculated the Pearson product-moment correlation coefficient of amplitude over trials for each pair of frequencies for each recording session (spontaneous and evoked conditions). We applied Fisher's r-to-z transformation to the correlation values, and tested the distributions of transformed values against the null hypothesis of zero correlation with a two-tailed Student's t-test at $\alpha = 0.05$, Bonferroni-corrected for the 15 pairs of frequency bands.

Neural network modeling

We simulated spiking neuron populations using a conductance-based leaky integrate-and-fire network with band-pass (alpha band, 8 – 12 Hz) noise oscillatory input, adapted from that presented by Renart and colleagues[17]. Briefly, the network was comprised of three populations— one excitatory ('E'; $N_E = 1000$ neurons), one inhibitory ('I'; $N_I = 250$ neurons), and one input ('X'; $N_X = 1000$ neurons) population— and was iterated with time step $\Delta t = 0.05$ ms in a Runge-Kutta-2 integration scheme using MATLAB. The membrane voltage of the i^{th} neuron in the α population, V_i^α, was governed by

$$C_m \frac{dV_i^\alpha}{dt} = -g_L\,(V_i^\alpha - V_L) + I_i^{\alpha E} + I_i^{\alpha I} + I_i^{\alpha X} + I_{app}^\alpha$$

where $C_m = 0.25$ nF was the membrane conductance, $g_L = 16.7$ nS was the leak conductance (such that the membrane time constant $C_m/g_L = 15$ ms), $V_L = -70$ mV was the resting potential and $\theta = -50$ mV was the spiking threshold. Following a spike, the voltage V_i^α was reset to $V_R = -60$ mV, after which a refractory period of 2 or 1 ms was assumed for the E and I populations, respectively. All parameters were consistent with the original model[17].

$I_i^{\alpha\beta}$, the synaptic current from population β (E, I, or X) onto population α (E or I), was governed by

$$I_i^{\alpha\beta}(t) = -\left|\sum_{j=1}^{N_\beta} p_{ij}^{\alpha\beta} g_{ij}^{\alpha\beta} s_{ij}^{\alpha\beta}(t)\right| \left(V_i^\alpha - V_R^\beta\right)$$

where V_R^β is the reversal current of the efferent population ($V_R^E = V_R^X = 0$ mV and V_R^I mV). $p_{ij}^{\alpha\beta}$ was the probability that cell *i* in population α was connected to cell *j* in population β, defined for all populations as a binary random variable with a mean of 0.2 (i.e. densely but randomly connected). $g_{ij}^{\alpha\beta}$ was the synaptic conductance from cell *i* in population α to cell *j* in population β, defined as a Gaussian random variable with mean $g^{\alpha\beta}$ and standard deviation $0.5g^{\alpha\beta}$. We used g^{EE} = 2.4 nS, g^{EI} = 40 nS, g^{IE} = 4.8 nS, g^{II} = 40 nS, and g^{EX} = g^{IX} = 5.4 nS. $s_{ij}^{\alpha\beta}$ was a synaptic gating term, which integrated post-synaptic potentials from the β population at the *ij* synapse according to the following double-exponential ODE (with intermediate variable $x_{ij}^{\alpha\beta}$):

$$\begin{aligned} \tau_d \frac{ds_{ij}^{\alpha\beta}}{dt} &= -s_{ij}^{\alpha\beta} + x_{ij}^{\alpha\beta} \\ \tau_r \frac{dx_{ij}^{\alpha\beta}}{dt} &= -x_{ij}^{\alpha\beta} + \sum_{t_j^\beta} \delta\left(t - t_j^\beta - d_{ij}^{\alpha\beta}\right) \end{aligned}$$

where t_j^β were spike times of neuron *j* and τ_d = 5 ms and τ_r = 1 ms were the decay and rise time constants, respectively. $d_{ij}^{\alpha\beta}$ was the conductance delay between pairs of neurons, defined as a uniform random variable with resolution Δt = 0.05 ms with $d_{ij}^{\alpha E}, d_{ij}^{\alpha X} \in [0.5, 1.5]\, ms$ and $d_{ij}^{\alpha I} \in [0.1, 0.9]\, ms$. The spike times of neurons of the input population (X) were independently Poisson-distributed with mean 10 Hz.

I_{app}^α was an applied oscillatory current to the α population, constructed as white noise that was bandpass-filtered to the alpha band, 8 – 12 Hz. We used two oscillations: one global oscillation delivered to both E and I populations and a private oscillation delivered only to the I population. To represent varied input oscillation amplitudes, the RMS amplitude of the global and private oscillations were scaled independently; 10 different global oscillations ranged from 0 – 0.60 nA, 10 different private oscillations ranged from 0 – 0.20 nA, and all pairs of amplitudes were simulated. This range was chosen to cover the full range of possible network behaviors under the oscillations. We found that global oscillation amplitudes greater than 0.30 nA yielded ecologically implausible bursting activity coherent with the input oscillation, and therefore focused only on oscillations below that value.

We simulated 48 different networks, each for 6 – 12 s for each individual pairing of E and I amplitudes. Every segment of data was decomposed into 200 ms “trials,” giving us frequency resolution at and above 5 Hz. By this method, we accrued at least 1920 trials for each E and I amplitude pair. We modeled the EEG as the sum of postsynaptic potentials in the excitatory neurons, reasoning that the EEG signal is primarily generated by currents in the apical dendrites of pyramidal cells[15]. We sorted these trials into deciles of the alpha-

band amplitude of the modeled EEG (Figure 7), and pseudo-randomly subsampled the population in each decile bin so that the resulting average pairwise geometric mean firing rate was as nearly constant as possible across bins. Finally, we calculated r_{sc} between pairs of excitatory neurons, which are the predominant cell type recorded on our multielectrode arrays.

Attention experiment

Monkeys B and R participated in the attention experiment. The animals were trained to fixate a central yellow dot. After fixating for a randomly chosen duration of 300 or 700 ms, a peripheral visual cue was presented for 120 ms on 89% of trials (no-cue trials were identical in all other respects). For Monkey B, the cue was a 0.6° dim gray dot (87 cd/m^2; 9% contrast) centered on one of the two locations of subsequently presented grating stimuli. For Monkey R, the cue was a yellow annulus (isoluminant with the display background, 6.45° inner diameter, 7.03° outer diameter, masked to prevent crossing the vertical meridian) that encircled one of the two stimulus locations. The cue onset was followed by another randomly-chosen duration of 300 or 700 ms, after which the two drifting sinusoidal gratings were presented. One grating was presented at 0° orientation and the other was presented at 90° orientation. Orientation was counterbalanced between stimulus locations across trials. All other properties of the gratings were the same as described for the ‘evoked’ task above. One grating was presented in the receptive field area of the neurons recorded on the array, as described above, and the other grating was presented in the mirror symmetric location in the opposite hemifield. The animal’s task was to detect a speed change (acceleration or deceleration) of one of the two gratings and to make a saccade to the stimulus that changed within 800 ms of the change onset. Correct responses were reinforced with juice or water. The speed change was governed by a triangular ramp function (i.e., a gradual increase followed by a return to baseline, or vice versa). We titrated the maximum magnitude of the speed change in each direction prior to the experiment using a staircasing procedure to set an overall correct detection rate for each type of speed change between 70% and 80% correct. On 40% of trials, neither grating changed speed, and the animal was rewarded for maintaining fixation for the full stimulus duration, 1.2 s. The speed change, if any, could begin between 250 ms and 700 ms after grating onset, uniformly distributed. If a cue and speed change both occurred, the cue validly indicated the correct location of the target with 80% probability. The attention task design is shown in Supplementary Figure 3.

While the animals performed the task, we recorded EEG data and spiking data as described above. Initial preprocessing (e.g., EEG artifact rejection, spike identification and sorting) also proceeded as described above. We excluded from this analysis recording sessions with less than 72 correct responses to targets contralateral to the array (15 sessions in Monkey B and 4 sessions in Monkey R), because fewer responses produced highly variable estimates of r_{sc}. We also excluded sessions for which an animal was both slower and less accurate for targets preceded by valid cues (zero sessions in Monkey B and 4 sessions in Monkey R), which indicated poor attention performance. This resulted in a total of 8 recording sessions in the analysis for Monkey B, and a total of 34 recording sessions in the analysis for Monkey R. We also excluded trials with targets that occurred ipsilateral to the array, uncued trials, and trials with targets onsetting less than 500 ms after the grating onset or for which

the animal prematurely left the fixation window, so that we always had a 500 ms interval of constant stimulation over which to compute alpha-band amplitude and r_{sc} (n = 91.25 ± 7.18 trials per dataset in the analysis for Monkey B; n = 179.56 ± 8.30 trials per dataset for Monkey R). We calculated alpha amplitude at the parieto-occipital EEG electrode over the array by multiplying the 500 ms of EEG data following grating onset by a Hanning window, taking the Fourier transform of the windowed data, and averaging the amplitude values over 8 – 12 Hz. We log transformed the alpha amplitudes to approximate a normal distribution. Because we were particularly interested in the effect of within-session variation in alpha amplitude, we calculated normalized amplitude by dividing each dataset's amplitude values by the mean amplitude for that session. We likewise normalized response times (RTs) and r_{sc} by dividing by the average for the session. We next binned each session's trials into six quantiles based on alpha amplitude, and calculated mean RT, mean alpha amplitude and r_{sc} within each quantile. Values exceeding 3 standard deviations from the mean were excluded as outliers (2 data points excluded). We calculated r_{sc} by counting each unit's spikes in the 500 ms interval following grating onset, standardizing each unit's spike counts over trials separately for each grating orientation with a z-transform, and then calculating Pearson product-moment correlation over trials for each pair of units. We next used a regression analysis to quantify the relationship between alpha amplitude and RT, and the degree to which any such effect was mediated by r_{sc}. Although the regression of RT on alpha amplitude could have in principle been performed using single trials, we used the means of the amplitude bins for this analysis so that the result could be meaningfully compared to a model that included r_{sc}, which is only defined for sets of multiple trials. In line with our driving hypothesis, we first performed a second-order polynomial regression of RT on alpha amplitude in the ordinary least squares sense using the following equation:

$$RT = \beta_2^{\alpha} x^2 + \beta_1^{\alpha} x + \beta_0$$

Where *x* represents mean normalized alpha amplitude. Next, we added a linear term for r_{sc}, again in line with our driving hypothesis:

$$RT = \beta_2^{\alpha} x^2 + \beta_1^{\alpha} x + \beta_1^{rsc} y + \beta_0$$

Where *y* represents r_{sc}. Since we were testing the specific model where RT follows a U-shaped relationship with alpha power and a direct linear relationship with r_{sc}, we constrained β_2^{α} and β_1^{rsc} to each be non-negative.

We calculated the coefficient of determination (R^2) for each model, and statistically tested the improvement in the proportion of variance-accounted-for by inclusion of the r_{sc} term against the null hypothesis that r_{sc} adds no information by using a bootstrap resampling procedure. For the bootstrap, we created 1000 surrogate samplings of r_{sc} values by drawing at random an equal number with replacement from the original sample. For each bootstrap sample we repeated the bivariate regression described above and calculated R^2 to derive a distribution of R^2 values reflecting the null hypothesis. The proportion of the bootstrapped R^2 values exceeding the observed R^2 value is the *p*-value for the statistical test.

Most critically for our predictions, we tested whether the inclusion of a linear term for r_{sc} led to a significant decrease in the magnitude of the quadratic term for alpha amplitude (i.e., β_2^{α}). If so, then this would indicate that any specific U-shaped relationship between RT and alpha amplitude was mediated by the linear relationship between RT and r_{sc}. The statistical significance for this test was assessed using a bootstrap resampling procedure as above. One advantage of a bootstrap procedure for this statistical test is that we compared the improvement in the model by including an additional term (i.e., r_{sc}) for the observed data to the same improvement by including an additional term for the resampled data, which reflect the null hypothesis. Therefore, significant differences between the observed and resampled data could not be simply due to having a greater number of model terms. We also performed this analysis for the other five frequency bands of interest besides the alpha band to test for frequency specificity.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgements

ACS was supported by an NIH fellowship (F32EY023456). MJM was supported by an NIH undergraduate research training fellowship through the University of Pittsburgh's Program in Neural Computation (R90DA023426). MAS was supported by NIH grants R00EY018894, R01EY022928 and P30EY008098, a career development grant and an unrestricted award from Research to Prevent Blindness, and the Eye and Ear Foundation of Pittsburgh. We are grateful to Dr. Brent Doiron for helpful advice and discussion, and to Ms. Stephanie Cortes and Ms. Samantha Nelson for help with data collection.

References

1. Averbeck BB, Latham PE, Pouget A. Neural correlations, population coding and computation. Nature reviews. Neuroscience. 2006; 7:358–366. [PubMed: 16760916]
2. Kohn A, Smith MA. Stimulus dependence of neuronal correlation in primary visual cortex of the macaque. J Neurosci. 2005; 25:3661–3673. [PubMed: 15814797]
3. Smith MA, Kohn A. Spatial and temporal scales of neuronal correlation in primary visual cortex. J Neurosci. 2008; 28:12591–12603. [PubMed: 19036953]
4. Smith MA, Sommer MA. Spatial and temporal scales of neuronal correlation in visual area V4. J Neurosci. 2013; 33:5422–5432. [PubMed: 23516307]
5. Cohen MR, Maunsell JH. Attention improves performance primarily by reducing interneuronal correlations. Nat Neurosci. 2009; 12:1594–1600. [PubMed: 19915566]
6. Mitchell JF, Sundberg KA, Reynolds JH. Spatial attention decorrelates intrinsic activity fluctuations in macaque area V4. Neuron. 2009; 63:879–888. [PubMed: 19778515]
7. Gu Y, et al. Perceptual learning reduces interneuronal correlations in macaque visual cortex. Neuron. 2011; 71:750–761. [PubMed: 21867889]
8. Jeanne JM, Sharpee TO, Gentner TQ. Associative learning enhances population coding by inverting interneuronal correlation patterns. Neuron. 2013; 78:352–363. [PubMed: 23622067]
9. Snyder AC, Morais MJ, Kohn A, Smith MA. Correlations in V1 are reduced by stimulation outside the receptive field. J Neurosci. 2014; 34:11222–11227. [PubMed: 25143603]
10. Gregoriou GG, Gotts SJ, Zhou H, Desimone R. High-frequency, long-range coupling between prefrontal and visual cortex during attention. Science. 2009; 324:1207–1210. [PubMed: 19478185]
11. Pesaran B, Nelson MJ, Andersen RA. Free choice activates a decision circuit between frontal and parietal cortex. Nature. 2008; 453:406–409. [PubMed: 18418380]

12. Logothetis NK, Pauls J, Augath M, Trinath T, Oeltermann A. Neurophysiological investigation of the basis of the fMRI signal. Nature. 2001; 412:150–157. [PubMed: 11449264]
13. Shmuel A, Augath M, Oeltermann A, Logothetis NK. Negative functional MRI response correlates with decreases in neuronal activity in monkey visual area V1. Nat Neurosci. 2006; 9:569–577. [PubMed: 16547508]
14. Whittingstall K, Logothetis NK. Frequency-band coupling in surface EEG reflects spiking activity in monkey visual cortex. Neuron. 2009; 64:281–289. [PubMed: 19874794]
15. Nunez, PL.; Srinivasan, R. Electric fields of the brain : the neurophysics of EEG. Oxford University Press; Oxford ; New York: 2006.
16. Musall S, von Pfostl V, Rauch A, Logothetis NK, Whittingstall K. Effects of neural synchrony on surface EEG. Cereb Cortex. 2014; 24:1045–1053. [PubMed: 23236202]
17. Renart A, et al. The asynchronous state in cortical circuits. Science. 2010; 327:587–590. [PubMed: 20110507]
18. van Vreeswijk C, Sompolinsky H. Chaos in neuronal networks with balanced excitatory and inhibitory activity. Science. 1996; 274:1724–1726. [PubMed: 8939866]
19. Mathewson KE, Gratton G, Fabiani M, Beck DM, Ro T. To see or not to see: prestimulus alpha phase predicts visual awareness. J Neurosci. 2009; 29:2725–2732. [PubMed: 19261866]
20. Lakatos P, Karmos G, Mehta AD, Ulbert I, Schroeder CE. Entrainment of neuronal oscillations as a mechanism of attentional selection. Science. 2008; 320:110–113. [PubMed: 18388295]
21. van Elswijk G, et al. Corticospinal beta-band synchronization entails rhythmic gain modulation. J Neurosci. 2010; 30:4481–4488. [PubMed: 20335484]
22. Womelsdorf T, et al. Orientation selectivity and noise correlation in awake monkey area V1 are modulated by the gamma cycle. Proceedings of the National Academy of Sciences of the United States of America. 2012; 109:4302–4307. [PubMed: 22371570]
23. de la Rocha J, Doiron B, Shea-Brown E, Josić K, Reyes A. Correlation between neural spike trains increases with firing rate. Nature. 2007; 448:802–806. [PubMed: 17700699]
24. Cohen MR, Kohn A. Measuring and interpreting neuronal correlations. Nat Neurosci. 2011; 14:811–820. [PubMed: 21709677]
25. Palva S, Palva JM. New vistas for alpha-frequency band oscillations. Trends in neurosciences. 2007; 30:150–158. [PubMed: 17307258]
26. Foxe JJ, Snyder AC. The Role of Alpha-Band Brain Oscillations as a Sensory Suppression Mechanism during Selective Attention. Frontiers in psychology. 2011; 2:154. [PubMed: 21779269]
27. Yuval-Greenberg S, Tomer O, Keren AS, Nelken I, Deouell LY. Transient induced gamma-band response in EEG as a manifestation of miniature saccades. Neuron. 2008; 58:429–441. [PubMed: 18466752]
28. Bosman CA, Womelsdorf T, Desimone R, Fries P. A microsaccadic rhythm modulates gamma-band synchronization and behavior. J Neurosci. 2009; 29:9471–9480. [PubMed: 19641110]
29. Mitchell JF, Sundberg KA, Reynolds JH. Differential attention-dependent response modulation across cell classes in macaque visual area V4. Neuron. 2007; 55:131–141. [PubMed: 17610822]
30. Narayanan R, Johnston D. Long-term potentiation in rat hippocampal neurons is accompanied by spatially widespread changes in intrinsic oscillatory dynamics and excitability. Neuron. 2007; 56:1061–1075. [PubMed: 18093527]
31. Linden H, Pettersen KH, Einevoll GT. Intrinsic dendritic filtering gives low-pass power spectra of local field potentials. Journal of computational neuroscience. 2010; 29:423–444. [PubMed: 20502952]
32. Manning JR, Jacobs J, Fried I, Kahana MJ. Broadband shifts in local field potential power spectra are correlated with single-neuron spiking in humans. J Neurosci. 2009; 29:13613–13620. [PubMed: 19864573]
33. Miller KJ, et al. Human motor cortical activity is selectively phase-entrained on underlying rhythms. PLoS computational biology. 2012; 8:e1002655. [PubMed: 22969416]
34. Ai L, Ro T. The phase of prestimulus alpha oscillations affects tactile perception. J Neurophysiol. 2014; 111:1300–1307. [PubMed: 24381028]

35. Lange J, Halacz J, van Dijk H, Kahlbrock N, Schnitzler A. Fluctuations of prestimulus oscillatory power predict subjective perception of tactile simultaneity. Cereb Cortex. 2012; 22:2564–2574. [PubMed: 22114082]
36. Linkenkaer-Hansen K, Nikulin VV, Palva S, Ilmoniemi RJ, Palva JM. Prestimulus oscillations enhance psychophysical performance in humans. J Neurosci. 2004; 24:10186–10190. [PubMed: 15537890]
37. Zhang Y, Ding M. Detection of a weak somatosensory stimulus: role of the prestimulus mu rhythm and its top-down modulation. Journal of cognitive neuroscience. 2010; 22:307–322. [PubMed: 19400673]
38. Baron RM, Kenny DA. The moderator-mediator variable distinction in social psychological research: conceptual, strategic, and statistical considerations. Journal of personality and social psychology. 1986; 51:1173–1182. [PubMed: 3806354]
39. Houser CR, Hendry SH, Jones EG, Vaughn JE. Morphological diversity of immunocytochemically identified GABA neurons in the monkey sensory-motor cortex. Journal of neurocytology. 1983; 12:617–638. [PubMed: 6352867]
40. Buzsaki G, Chrobak JJ. Temporal structure in spatially organized neuronal ensembles: a role for interneuronal networks. Curr Opin Neurobiol. 1995; 5:504–510. [PubMed: 7488853]
41. Disney AA, Aoki C, Hawken MJ. Gain modulation by nicotine in macaque v1. Neuron. 2007; 56:701–713. [PubMed: 18031686]
42. Disney AA, Aoki C, Hawken MJ. Cholinergic suppression of visual responses in primate V1 is mediated by GABAergic inhibition. J Neurophysiol. 2012; 108:1907–1923. [PubMed: 22786955]
43. Snyder AC, Morais MJ, Kohn A, Smith MA. Correlations in V1 are reduced by stimulation outside the receptive field. J Neurosci. (In Press).
44. Zohary E, Shadlen MN, Newsome WT. Correlated neuronal discharge rate and its implications for psychophysical performance. Nature. 1994; 370:140–143. [PubMed: 8022482]
45. Kelly RC, Smith MA, Kass RE, Lee TS. Local field potentials indicate network state and account for neuronal response variability. Journal of computational neuroscience. 2010; 29:567–579. [PubMed: 20094906]
46. Katzner S, et al. Local origin of field potentials in visual cortex. Neuron. 2009; 61:35–41. [PubMed: 19146811]
47. Goris RL, Movshon JA, Simoncelli EP. Partitioning neuronal variability. Nat Neurosci. 2014; 17:858–865. [PubMed: 24777419]

References for online methods

48. Pelli DG. The VideoToolbox software for visual psychophysics: transforming numbers into movies. Spat Vis. 1997; 10:437–442. [PubMed: 9176953]
49. Kleiner M, Brainard DH, Pelli D. What's new in Psychtoolbox-3? Perception. 2007; 36
50. Brainard DH. The Psychophysics Toolbox. Spat Vis. 1997; 10:433–436. [PubMed: 9176952]
51. Shoham S, Fellows MR, Normann RA. Robust, automatic spike sorting using mixtures of multivariate t-distributions. Journal of neuroscience methods. 2003; 127:111–122. [PubMed: 12906941]
52. Kelly RC, et al. Comparison of recordings from microelectrode arrays and single electrodes in the visual cortex. J Neurosci. 2007; 27:261–264. [PubMed: 17215384]

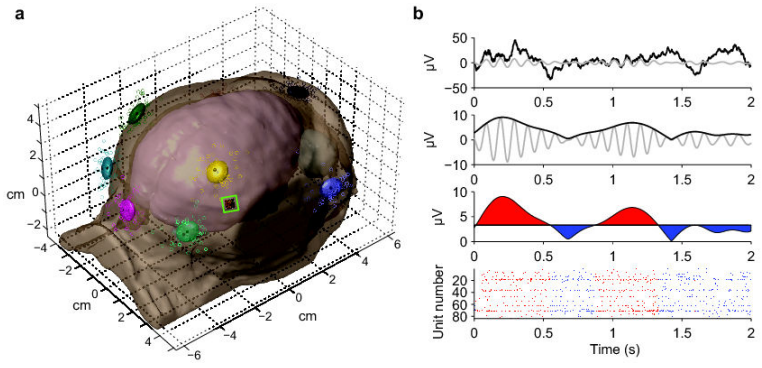

Figure 1).

Recording set-up and analysis strategy. A) Model of a representative subject's head derived from an MRI with co-registered electrode positions (colored scatter points). Each color represents a different position, and each marker of a given color represents a separate recording session. The corresponding ellipsoids each subtend ±1 standard deviation of the position (approximately 5 mm) along each of the three principle axes of the position distribution. The points by the left lateral canthus are occluded in this view. Unless otherwise specified, EEG data from the right occipital electrode (teal) are depicted in subsequent figures. The microelectrode array is shown to scale positioned in right V4 outlined by the green square. B) Conceptual illustration of the procedure for creating surrogate spike trains. From top to bottom: First, the raw EEG signal (black) is filtered at the alpha-band (gray); second, the envelope (black) is taken for the filtered signal (gray); third, periods of high (red) and low (blue) amplitude are identified from the envelope (illustrated here using a median split for clarity, in actuality deciles were used); fourth, spikes are divided into surrogate spike trains consisting wholly of spikes occurring during periods of high alpha amplitude (red raster points) or low alpha amplitude (blue raster points). The continuous filter used here is for conceptual illustration; in actuality the amplitude (and phase) of the EEG was measured using an FFT of a snippet of data centered on each spike (see Methods).

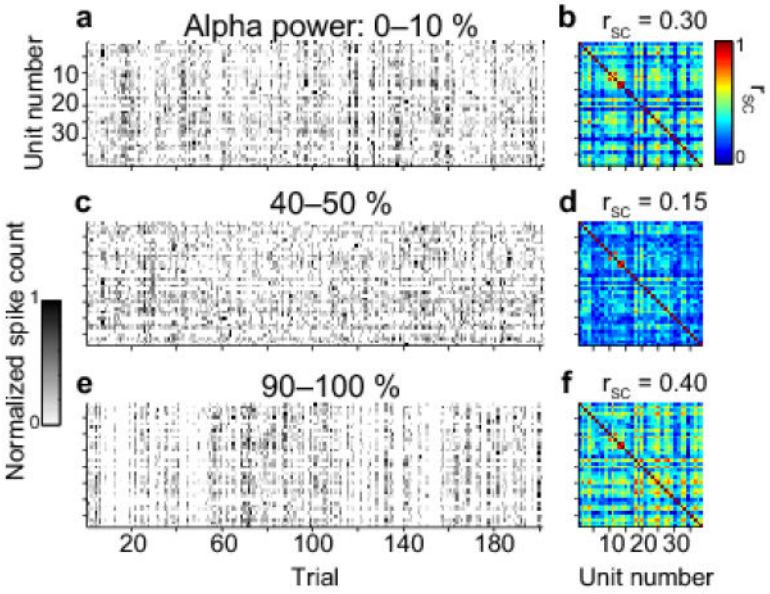

Figure 2).

Single-session example of surrogate datasets. A, C, E) Normalized spike counts of surrogate spike trains of 43 single units from 200 trials, as a proportion of the maximal spike count for each unit. Spikes were split into surrogate spike trains based on which decile of alpha power (relative to the entire session) was observed in a 500 ms window of data around the time of each spike. In the lowest (A) and highest (E) amplitude deciles, common periods of high and low spike counts are evident as dark and light vertical bands, respectively, indicating correlated variability in firing rate. In the intermediate decile (C), relative spike counts are more uniformly distributed across trials, indicating relatively less correlated variability. B, D, F) Matrices of all pair-wise spike count correlations (r_{sc}) for the data depicted in panels (A), (C) and (E). Note that the lowest (B) and highest (F) amplitude deciles have greater overall correlation than the intermediate decile (D).

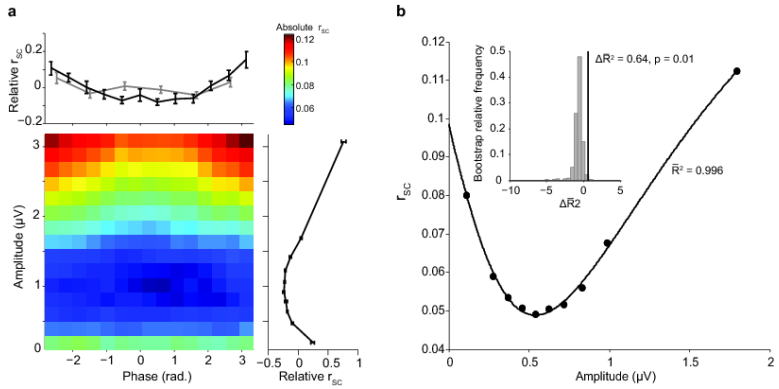

Figure 3).

Correlation results for spontaneous task across sessions (N = 24. Monkey B: 14 sessions; Monkey W: 10 sessions). The results illustrated are for alpha-band EEG measured at the right occipital electrode, but similar results were seen across frequencies and at other electrodes (see Figure 5). A) The false-color plot shows spike count correlation (r_{sc}) as a function of the phase and amplitude of the EEG, interpolated to fill the phase/amplitude space. At top, the change in r_{sc} relative to the session average is plotted as a function of phase measured when alpha amplitudes were in the top 10%. The black line shows phase measured in 12 equally-sized bins. The grey line shows phase measured in 6 bins spaced to equate firing rates. Note that the phase-correlation relationship is greatly reduced when the phase-rate relationship is controlled. At right, the relative change in r_{sc} is plotted as a function of amplitude (midpoint of each decile bin), averaged across phase. Error bars = ± 1 SEM. B) Curve fitting result for amplitude-correlation relationship. A rational function closely approximates the observed data $\left(\overline{R}^2=0.996\right)$, and was better than a simple linear fit $\left(\Delta\overline{\mathrm{R}}^2=0.64\right)$. The inset histogram shows the distribution of rational function fits to bootstrap permutations of the data, relative to linear fits. The observed $\Delta\overline{\mathrm{R}}^2$ is indicated by the vertical dark line and is greater than 99% of the bootstrap values, indicating that the relationship is significantly better explained by a rational function than would be expected by chance.

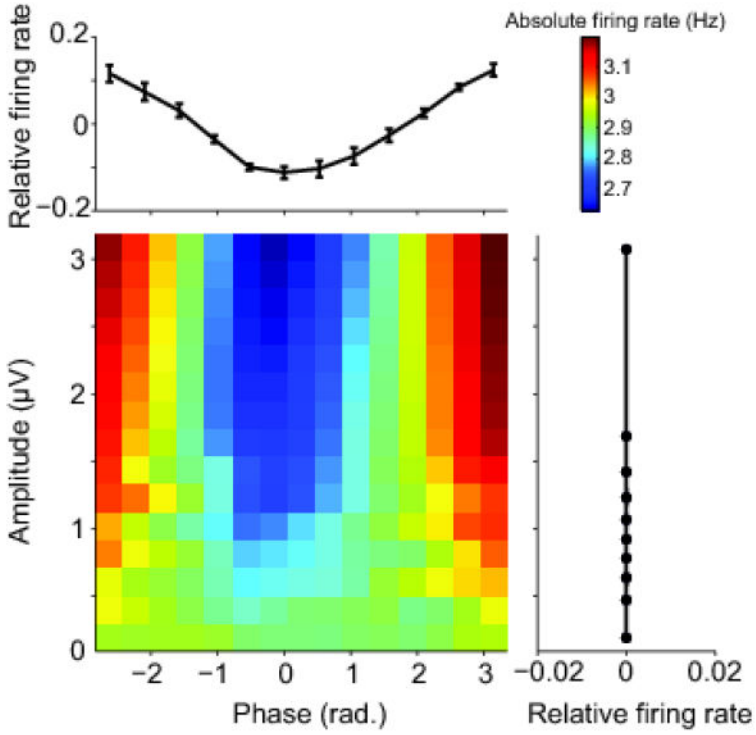

Figure 4).

Firing rate results for spontaneous task across all sessions and subjects. The results illustrated are for alpha-band EEG measured at the right occipital electrode, but similar results were seen across frequencies and at other electrodes. A) The false-color plot shows firing rate as a function of the phase and amplitude of the EEG, interpolated to fill the phase/amplitude space. At top, the change in firing rate relative to the session average is plotted as a function of phase measured when alpha amplitudes were in the top 10%. At right, the relative change in r_{sc} is plotted as a function of amplitude (midpoint of each decile bin), averaged across phase. Firing rate was controlled for amplitude by design. Error bars = ±1 SEM. For the marginal amplitude plot, error bars are smaller than the line thickness.

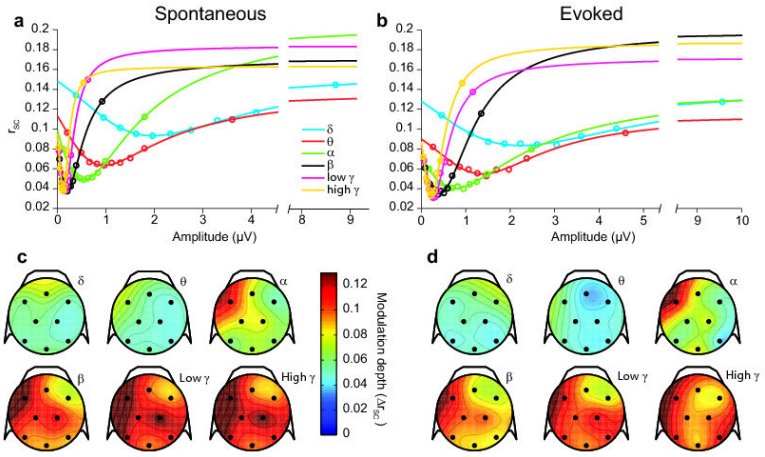

Figure 5).

Amplitude-correlation relationship across all six frequency bands of interest. Panels A and C depict the spontaneous condition, panels B and E depict the stimulus-evoked condition. A,B) Spike count correlation (r_{sc}) as a function of amplitude (midpoint of decile bin) for each frequency. Higher frequencies had greater modulation of r_{sc} as a function of amplitude. Lines depict rational function fits as in Figure 3b. C,D) Spline-interpolated topographical maps of the strength of the amplitude-correlation relationship (Δr_{sc} = max r_{sc} - min r_{sc}) for each frequency. The relationship was qualitatively similar at all electrodes, but was strongest at electrodes nearest the array and at those diametrically opposite the array, and was weakest at orthogonal locations. This antipodal topography was most evident at higher frequencies where the amplitude-correlation relationship was strongest overall.

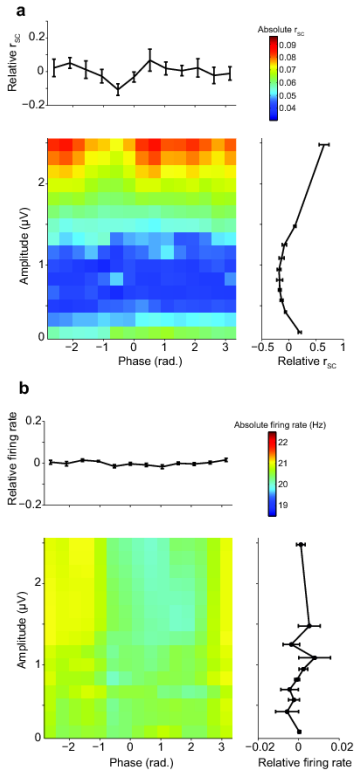

Figure 6).

Results for evoked task. The results illustrated are for alpha-band EEG measured at the right occipital electrode, but similar results were seen across frequencies and at other electrodes. A) Correlation results, plotted as in Figure 3a. B) Firing rate results, plotted as in Figure 5. Note that phase effects are greatly reduced for the evoked task compared to the spontaneous task. Error bars = ± 1 SEM.

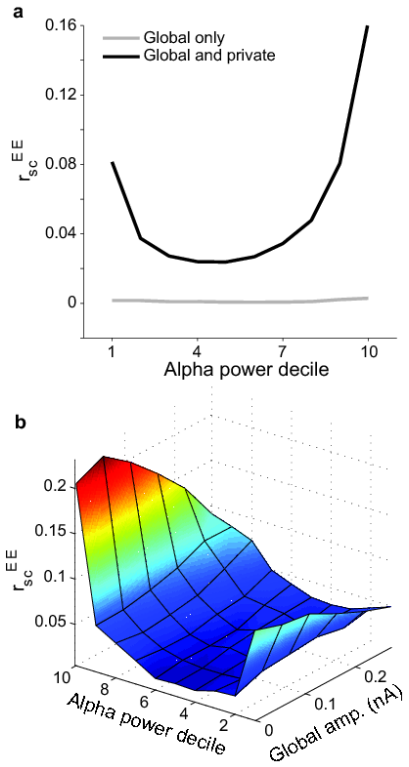

Figure 7).

Neural network modeling results. A) Spike count correlation between excitatory neurons (r_{sc}^{EE}) was on the order of 1/N if global input oscillations only were applied to the model (gray), and did not vary substantially with modeled EEG alpha power. If an independent input oscillation was applied to the inhibitory population, the resulting r_{sc}^{EE} spanned a range similar to our observed results and followed a U-shaped profile as a function of modeled EEG alpha power (black). B) Amplitude-correlation relationship as a function of global oscillation amplitude. The U-shaped amplitude-correlation relationship was robust when considering trials derived using a single global input oscillation (and a mixture of private inhibitory input oscillations), across a broad range of global input oscillation magnitudes.

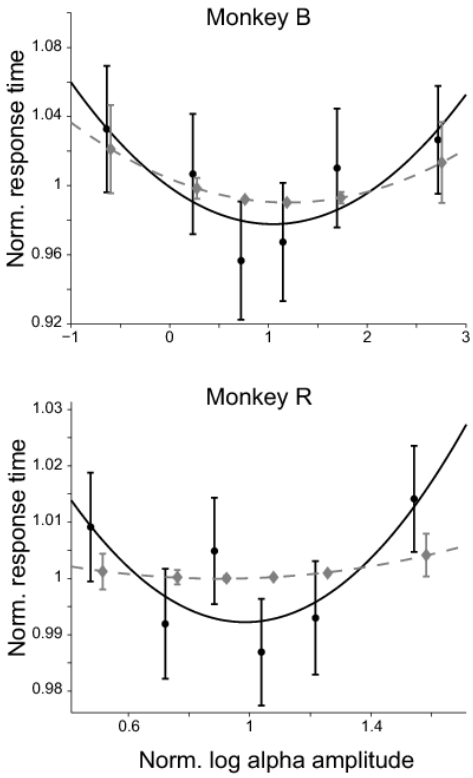

Figure 8).

Psychophysical performance vs. alpha amplitude. Response time followed a U-shaped relationship with alpha amplitude (black). The curve depicts the best-fitting second-order polynomial. After adjusting for a linear relationship between response time and spiking correlation, the relationship between response time and alpha amplitude was more uniform (gray), indicating that spiking correlation mediated the relationship (Monkey B, $p = 0.047$; Monkey R, $p = 0.030$). Error bars: ± 1 SEM.