

ORIGINAL ARTICLE

Decoding of Attentional State Using High-Frequency Local Field Potential Is As Accurate As Using Spikes

Surya S. Prakash¹, Aritra Das¹, Sidrat Tasawoor Kanth^{1,2},
J. Patrick Mayo³ and Supratim Ray^{1,2}

¹Centre for Neuroscience, Indian Institute of Science, Bangalore 560012, India, ²IISc Mathematics Initiative, Indian Institute of Science, Bangalore 560012, India and ³Department of Ophthalmology, University of Pittsburgh, Pittsburgh, PA 15213, USA

Address correspondence to Supratim Ray, Centre for Neuroscience, Indian Institute of Science, old TIFR building, Bangalore 560012, India.
Email: sray@iisc.ac.in

Abstract

Local field potentials (LFPs) in visual cortex are reliably modulated when the subject's focus of attention is cued into versus out of the receptive field of the recorded sites, similar to modulation of spikes. However, human psychophysics studies have used an additional attention condition, neutral cueing, for decades. The effect of neutral cueing on spikes was examined recently and found to be intermediate between cued and uncued conditions. However, whether LFPs are also precise enough to represent graded states of attention is unknown. We found in rhesus monkeys that LFPs during neutral cueing were also intermediate between cued and uncued conditions. For a single electrode, attention was more discriminable using high frequency (>30 Hz) LFP power than spikes, which is expected because LFP represents a population signal and therefore is expected to be less noisy than spikes. However, previous studies have shown that when multiple electrodes are used, spikes can outperform LFPs. Surprisingly, in our study, spikes did not outperform LFPs when discriminability was computed using multiple electrodes, even though the LFP activity was highly correlated across electrodes compared with spikes. These results constrain the spatial scale over which attention operates and highlight the usefulness of LFPs in studying attention.

Key words: gamma, high-gamma, linear discriminant analysis, population coding, spatial attention

Introduction

Power at different frequency bands like alpha (8–12 Hz) and gamma (30–80 Hz) of the local field potential (LFP) is reliably modulated by selective attention (Fries et al. 2001, 2008; Chalk et al. 2010; Khayat et al. 2010; Khamechian et al. 2019) and can be used to predict behavior (Womelsdorf et al. 2006). These results have led to suggestions that gamma oscillations could play a role in routing information across cortical areas (Fries 2015). However, most of these studies employed a standard attention task in which the animal's attention was either cued inside (attend-in) or outside (attend-out) the receptive field (RF) of the recording sites.

Additional attentional conditions are often used in psychophysical tasks in humans. Here, a spatial attention task

typically involves cueing the subject to deploy attention to a desired location at which a target is anticipated (Posner 1980). The target usually occurs at the cued location (which are called “valid” trials) but can infrequently also occur at an uncued location (“invalid” trials), such that the behavioral effect of spatial attention can be measured by comparing the performance (in terms of detection rates) and reaction times for valid versus invalid trials (Posner 1980; Carrasco 2011). In addition, such studies often use an ambiguous or “neutral” cueing condition where the target change is equally likely at both locations. These studies have shown that behavioral performance during neutral cueing is intermediate between valid and invalid conditions (Posner et al. 1978; Posner 1980; Mangun and Hillyard 1990; Montagna et al. 2009). While

neurophysiological studies of spatial attention have compared neuronal responses for attend-in versus attend-out conditions for decades (Moran and Desimone 1985; Treue and Maunsell 1996; McAdams and Maunsell 1999), the effect of neutral cueing on spiking activity has been studied only recently and found to be intermediate between cued and uncued conditions (Mayo and Maunsell 2016) (but see Denfield et al. 2018). However, whether LFP power in different frequency bands during neutral cueing is also intermediate between cued and uncued conditions is unknown.

While earlier neurophysiological studies focused mainly on responses of individual neurons, simultaneous recordings from multiple neurons have allowed a more comprehensive investigation of the effect of attention on the neuronal population. In particular, attention has been shown to reduce spike-count correlation between pairs of neurons (Cohen and Maunsell 2009; Mitchell et al. 2009), which can affect the way the neural population encodes a stimulus (Averbeck et al. 2006; Cohen and Kohn 2011). The spiking activity of a simultaneously recorded population of neurons has also been used to predict the animals' focus of attention and behavior in a single trial (Cohen and Maunsell 2010, 2011). Since handling high-dimensional data is difficult, many studies have devised methods to obtain a one-dimensional metric from a high-dimensional dataset to predict behavior (Cohen and Maunsell 2010; Cunningham and Yu 2014; Yates et al. 2020). Using a simple method of projecting trials onto a one-dimensional axis constructed using the mean spiking activity of a neuronal population of two different attention conditions (Cohen and Maunsell 2010; Mayo et al. 2015), Mayo and Maunsell (2016) showed that the projections of neutral cueing condition were also intermediate between cued and uncued conditions. However, whether the projections remain similarly graded even when incorporating higher-order statistics like covariance across electrodes was not tested. Further, it is unclear how well different attentional states can be discriminated using LFP power at different frequency bands as compared with spiking activity, and how this discriminability varies with the number of microelectrodes.

LFPs are thought to reflect the summed synaptic potentials of at least a few thousand neurons around the microelectrode tip (Buzsáki et al. 2012). Furthermore, a high-frequency (>80 Hz) component of the LFP reflects the overall spiking activity of the neural population (Ray et al. 2008; Ray and Maunsell 2011). Since spiking activity recorded from a single electrode represents only a noisy estimate of the true underlying attentional state, averaging across a larger population of neurons (as done implicitly in population signals like the LFP) reduces noise and provides a more reliable estimate of the attentional state. However, when the decoding of behavior is performed using multiple electrodes, the trends are often reversed, with multichannel spikes outperforming multichannel LFPs (Hwang and Andersen 2013). This is because firing rates are more variable across electrodes than LFPs, so a decoding algorithm can give more weight to the more informative channels. This strategy may not work as well for LFPs, since these signals represent the average activity of thousands of neurons that may not be informative.

To test these predictions, we recorded spikes and LFPs from two rhesus macaques performing an attention task under different cueing conditions and compared how power in different frequency bands varied across attention conditions. We also compared LFP-LFP and spike-LFP coherence. Finally, we compared different weighting strategies to obtain one-dimensional projections of data using spikes or power in

different frequency bands and compared the discriminability of attention conditions across neural measures as a function of population size. While single-channel LFPs expectedly outperformed single-channel spikes, surprisingly, multichannel spikes could not outperform multichannel LFPs.

Materials and Methods

Electrophysiological Recordings

Data used in this study and the experimental procedures were detailed in a previous report (Mayo and Maunsell 2016) that described the V4 spiking data in detail. All animal procedures were approved by the Institutional Animal Care and Use Committee of Harvard Medical School. Two adult male rhesus monkeys (*Macaca mulatta*) were surgically implanted with a titanium head post and scleral eye coil before training. After they learned the behavioral task, a 6×8 array of microelectrodes (Blackrock Microsystem, 48 electrodes in each) was implanted in visual area V4 in both cerebral hemispheres. The signals were referenced with respect to a single wire that was wrapped around the titanium straps that were used to hold the bone flap in place on top of the arrays. Because popular referencing techniques such as average or bipolar referencing change the power and phase of the LFP signals in nontrivial ways (Shirhatti et al. 2016), the signals were not subsequently re-referenced. The electrodes were 1 mm long and 400 μm apart with impedance typically between 0.2 and 1 $\text{M}\Omega$ at 1 kHz. Area V4 was identified using stereotaxic coordinates and by locating the lunate and superior temporal sulci during surgery. The azimuths and elevations of the centers of the RFs of neurons recorded from the two arrays were around (2° , -7°) and (-7° , -4°) for Monkey A and (6° , -7°) and (-6° , -3°) for Monkey W.

Signals from 96 channels were recorded using a 128-channel Cerebus Neural Signal Processor (Blackrock Microsystems). Signals were filtered between 0.3 Hz (Butterworth filter, 1st order, analog) and 2.5 kHz (Butterworth filter, 4th order, digital), sampled at 2 kHz, and digitized at 16-bit resolution to obtain the LFP signals. Note that the low-pass LFP filter was inadvertently set to a value greater than half the sampling rate and hence could have led to some aliasing. However, because the power of the LFP signal falls rapidly with frequency (log PSDs shown in Figure 1C kept decreasing beyond 200 Hz to about -1.3 at ~ 700 Hz after which they flattened), the aliased power above 1000 Hz was negligible. Signals were separately filtered between 250 Hz (Butterworth filter, 4th order, digital) and 7.5 kHz (Butterworth filter, 3rd order, analog) and subjected to a threshold followed by spike sorting using spike sorter software (Plexon) to extract single and multi-units, which we also refer to as "neurons."

Behavioral Task

Monkeys learned an orientation change detection task in which they held their gaze within a 1.8° square window centered on a fixation spot in the center of the screen throughout the trial. During fixation, two counterphasing Gabor stimuli appeared simultaneously and synchronously in the RFs of one of the units recorded from each hemisphere. At an unsignaled time, chosen from an exponential distribution (mean 3000 ms, range 500–5000 ms) and when the contrasts of the counterphasing Gabor stimuli were both at 0%, the orientation of one of the stimuli changed. Monkeys were rewarded for making a saccade to the location of the changed stimulus between 100 and 550 ms after the orientation change.

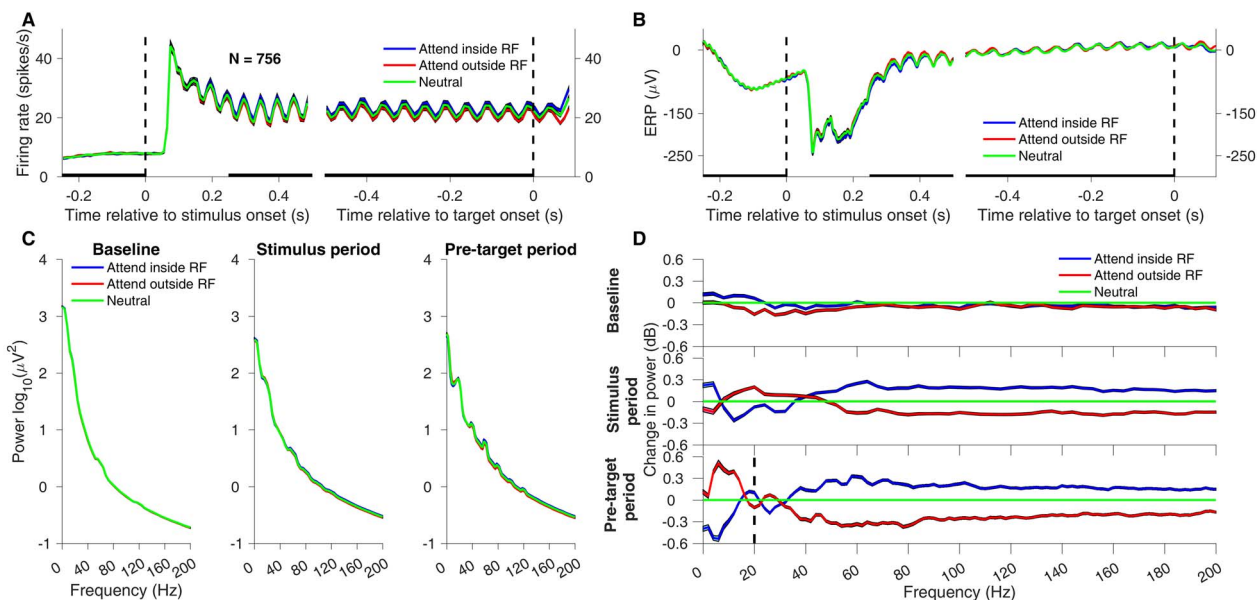


Figure 1. Comparison of firing rate, evoked potential, and PSDs across attend-in, attend-out, and neutral attention conditions. (A) Mean PSTH of all 756 electrodes relative to stimulus onset (left plot) and target onset (right plot) for attend-in (blue), attend-out (red), and neutral (green) attention conditions. The vertical dashed line indicates the stimulus onset time (left) and target onset time (right). The bars on the X-axis indicate the periods chosen for analysis. We only consider the valid and correct responses here (blue: Attend-in Valid Hit, red: Attend-out Valid Hit, green: Target-out Neutral Hit and Target-in Neutral Hit combined). The results for other attention conditions are shown in Figure 3. (B) Mean ERP of the same electrodes and periods as in (A) for the three attention conditions. (C) PSD of the baseline period (–0.25 to 0 s relative to stimulus onset), stimulus period (0.25–0.5 s relative to stimulus onset), and pre-target period (–0.5 to 0 s relative to target onset) for the three attention conditions. (D) Change in power (in dB) of the three attention conditions, obtained by subtracting the green trace from the other traces shown in (C). Green trace represents the change in power of the neutral condition with itself and hence is trivially zero. The black vertical dashed line in the bottom plot indicates the power at 20 Hz SSVEP frequency.

Monkeys were cued to attend to one (valid or invalid cue) or both (neutral cue) of the stimulus locations for a block of 50 trials using four instruction trials in which a white spot (0.45° radius) appeared for 50–100 ms indicating the likely location of the stimulus change in the upcoming cued block or simultaneously at both stimulus locations for the neutral block. In cued blocks, the change occurred at the cued location with 80% probability (valid cue) and 20% at the uncued side (invalid cue). In neutral blocks, the change occurred at each stimulus location with equal probability. Five percent of all trials were catch trials in which neither of the stimuli changed orientation, and monkeys were rewarded for maintaining fixation throughout the trial. Catch trials and instruction trials were excluded from analysis.

Six randomly interleaved orientation change magnitudes were used in each session to characterize the effects of attention on behavioral performance. While all of these orientation changes were used during valid and neutral trials, only the second and third smallest orientation changes were used for invalid trials (see (Mayo and Maunsell 2016), for details). For these orientation changes, the mean behavioral performance was approximately 73% for validly cued trials, approximately 64% for neutral, and approximately 48% for invalid trials. As in the previous report (Mayo and Maunsell 2016), all analyses were restricted to only trials with these orientation changes for proper comparison of valid and neutral conditions with the invalid condition.

Stimuli

Two odd symmetric, Gabor stimuli sinusoidally counterphased at 10 Hz of maximum contrast were used during each recording session. Each had a size (σ range: 0.45°–1.6°) and spatial

frequency (range: 0.3–2.5 cycles/deg) optimized for one of the units from each hemisphere. Stimuli were displayed on a CRT monitor (100 Hz frame rate, 1024 × 768 pixels, 8-bit DACs) over a uniform gray background. The monitor was calibrated to produce linear steps of luminance and positioned at 57 cm from the subject. Such counterphasing stimuli produce a salient neuronal response at twice the stimulus frequency (20 Hz; see Fig. 1A,B). These responses are referred to as steady-state visually evoked responses (SSVEPs). Eye position was sampled at 200 Hz using the scleral eye coil technique (Judge et al. 1980).

Attention Conditions

Attention conditions were defined in two ways. The first formulation was in reference to the RF of the neurons (Figs 3 and 4) and yielded the standard attend-in and attend-out conditions. The second formulation was used for population discriminability analysis that included neurons from both hemispheres (Fig. 5) and was therefore based only on the cued hemifield. Each formulation consisted of 12 conditions, as described below.

First Formulation

During the cued and uncued conditions, the animals' attention was cued either inside or outside the RF. The cue could be either valid or invalid, and the animal could either detect (hit) or not detect (miss) the target. This yielded the following eight conditions: Attend-in Valid Hit, Attend-out Valid Hit, Attend-in Valid Miss, Attend-out Valid Miss, Attend-in Invalid Hit, Attend-out Invalid Hit, Attend-in Invalid Miss, and Attend-out Invalid Miss. During the neutral cueing conditions, trials cannot be designated as attend-in and attend-out but can be separated based

on where the final target appeared. This yielded the following four conditions: Target-in Neutral Hit, Target-out Neutral Hit, Target-in Neutral Miss, and Target-out Neutral Miss. Therefore, each attention condition was denoted by three words. The first word denoted the likely focus of attention/actual target location (two values: Attend-in/Attend-out or Target-in/Target-out), the second word denoted the validity of the cue (three values: Valid, Neutral, or Invalid), and the third word indicated the behavioral outcome (two values: Hit or Miss).

Second Formulation

This formulation was independent of RF location because electrodes from both hemispheres were used. Therefore, instead of using Attend-in or Attend-out as used in the previous formulation, we used Attend-left or Attend-right, where left/right indicated the attended visual hemifield.

Session and Electrode Selection

Data were collected in 26 recording sessions (14 in Monkey A and 12 in Monkey W). One session was discarded because of an insufficient number of trials in one attention condition, yielding 25 sessions (13 and 12 in the two monkeys). Because LFP power was modulated due to attention whether or not the electrode picked up any spiking activity, choosing electrodes with no or unresponsive spiking activity would have reduced the performance of a spike-based decoder. Therefore, for a fair comparison between spiking and LFP data, we only considered electrodes that had neuronal firing rates in the stimulus period that were ≥ 5 spikes/s greater than during the baseline period. This yielded a total of 756 electrodes (average: 30.2 per session; min: 8, max: 46). Only sessions for which at least 15 trials were available for a particular attention condition were used for the analysis of that condition.

Data Analysis

Spike and LFP data were analyzed in three periods: 1) baseline: 250–0 ms before stimulus onset, 2) stimulus: 250–500 ms after the stimulus onset, and 3) pre-target: 500–0 ms before the stimulus changed to a different orientation (target). Peristimulus time histograms (PSTHs) were obtained by binning spike counts in 10 ms windows. Power spectral density (PSD) was calculated using multitaper method (Mitra and Pesaran 1999) with three tapers using the Chronux toolbox in MATLAB (Bokil et al. 2010). Frequency resolution was 4 Hz in baseline and stimulus periods and 2 Hz in the pre-target period. Power in the specific bands was computed by adding the power at each frequency point in specified range for each electrode and then averaging across electrodes of all the sessions. The following bands were used: alpha (8–12 Hz), gamma (42–78 Hz), and high-gamma (122–198 Hz). The power at the harmonics of 20 Hz SSVEP was excluded from all these bands. We also studied the power at the SSVEP frequency (20 Hz). The phase consistency between 1) spikes and LFPs and 2) LFPs and LFPs was computed across electrode pairs using pairwise phase consistency (PPC), which gives an unbiased estimate of the square of the phase consistency (Vinck et al. 2010, 2012).

Population Projection

The goal of the population projection analysis was to quantify how discriminable the population neural activity was when

the monkey paid attention to the left versus the right stimulus location (i.e., the discriminability in the neural data between the Attend-left Valid Hit and Attend-right Valid Hit conditions). These two conditions were chosen because the animals were most likely to be paying attention to the cued locations when they detected a target on the validly cued side. Because the population data were high-dimensional (equal to the number of electrodes in both hemispheres), we wanted to find a one-dimensional projection (a weighted sum) of this high-dimensional data so that the projections were maximally separated for the two conditions. For this purpose, we used a linear discriminant analysis (LDA), which is a technique based on Fisher's Linear Discriminant (Fisher 1936) that maximizes the separation between two classes in a lower-dimensional space. This method works well when the size of the test data (number of trials in the Attend-in and Attend-out Valid Hit conditions) is comparable to the dimensionality (number of electrodes), or if assumptions of common covariance matrix across classes are not met (Li et al. 2006). More complex classifiers like neural networks were not considered here because they typically require a larger training dataset set to prevent overfitting and they are less interpretable.

We used three projection methods of increasing complexity, progressively building to the LDA because this allowed us to study the effect of accounting for mean differences across classes, variances, and covariances separately (see below for details). Furthermore, we wanted to also compare with a technique called "Attention Axis" analysis (Cohen and Maunsell 2010; Mayo et al. 2015) (Fig. 5), which is a variant of the first method. All analyses were performed both with and without 5-fold cross-validation to better illustrate the effect of cross-validation on our results.

Consider two data matrices, D_L and D_R , corresponding to the neural data for the Attend-left Valid Hit and Attend-right Valid Hit conditions. The sizes of these matrices are $N \times T_L$ and $N \times T_R$, where N is the total number of electrodes from both hemispheres, and T_L and T_R are the number of trials for the left and right conditions. Our aim is to find an $N \times 1$ weight vector W , which is used to obtain the projections $P = W^T D$. Although W is obtained using D_L and D_R only, it can be used to project the data (D) of any attention condition. The three methods are as follows:

1) Mean difference: Here $W = (\mu_L - \mu_R)$, where μ_L and μ_R represent the mean of the data matrices D_L and D_R across trials, respectively, and therefore are $N \times 1$ matrices. Here, we assign a weight to each electrode that depends on how different its responses are for the two attention conditions. If the responses are equal for the two attention conditions, the weight for that electrode is zero. The first step of the "Attention Axis" projection is similar to this step. This method does not account for the variance across trials for individual electrodes or the covariances across electrodes.

2) Uncorrelated LDA: Here, we divide each weight obtained in the previous case by the average weighted variance across trials for the two conditions. For the i th electrode, the weight is given by

$$(W_{\text{Uncorrelated LDA}})_i = \left(\frac{(\mu_{Li} - \mu_{Ri})}{((T_L - 1)\sigma_{Li}^2 + (T_R - 1)\sigma_{Ri}^2 / (T_L + T_R - 2))} \right)$$

where $\mu_{Li/Ri}$ is the mean activity for the i th electrode as before, while $\sigma_{Li/Ri}^2$ is the variance in neural activity for left/right condition for the i th electrode, and T_L and T_R are the number of

trials for the left and right conditions as before. In vector form, $W_{\text{UncorrelatedLDA}} = \Sigma^{-1}(\mu_L - \mu_R)$, where Σ is an $N \times N$ diagonal matrix containing the pooled variances of the electrodes along the diagonals.

3) LDA: To also account for covariations in neural data across electrode pairs, we use the full covariance matrix instead of a diagonal matrix:

$$\Sigma = \text{Pooled covariance} = \frac{(T_L - 1)\Sigma_L + (T_R - 1)\Sigma_R}{T_L + T_R - 2}$$

where $\Sigma_{L/R}$ is the $N \times N$ covariance matrix of $D_{L/R}$ and $W_{\text{LDA}} = \Sigma^{-1}(\mu_L - \mu_R)$ as before. We performed LDA using `fitcdiscr` function in MATLAB. This function fits the LDA model to the input data (D_L and D_R) by computing the class means and pooled covariance and estimates the discriminant function (a linear boundary) to classify the data into different classes. We used this function to estimate the weights for each dimension to generate an LDA axis along which the data had maximum separation between the class means and minimum variance within classes.

Statistical Tests

We compared the medians of two groups using nonparametric two-sided Wilcoxon rank sum test, which does not assume a normal distribution of the data. Kruskal-Wallis test was done when comparing medians of more than two groups simultaneously.

Simulations

To study the effects of cross-validation, regularization, and sample size in LDA analysis, we generated two datasets of 50 000 samples, each drawn from two 30-dimensional multivariate normal distributions with different mean vectors and the same covariance matrix such that the true d-prime (ratio of the difference of the means and the pooled variance) of the projection of data onto the optimal LDA axis was between 0.5 and 1.5. We then randomly drew samples of different sizes from these datasets and computed the d-prime of projections onto the LDA axis estimated with or without 5-fold cross-validation for each sample size. Finally, we normalized the estimated d-prime by taking the ratio with the true d-prime. This process was repeated 50 times to compute the standard deviations.

These simulations highlight potential issues with neurophysiological datasets in which the number of dimensions (electrodes, ~30 on average in our case) is not much smaller than the number of data points. In such cases, not cross validating leads to severe overfitting, while cross-validation leads to poor performance because it is difficult to estimate a large covariance matrix and the means of each class with few data. For example, we noticed in our simulations that when the number of data points per class (trials per attention condition) was small (say, ~100 trials), the estimated d-primes normalized relative to the true d-prime were approximately 1.51 ± 0.03 (mean \pm SEM) when not cross-validated, but only approximately 0.6 ± 0.03 when cross-validated. Only when the number of trials was several folds (>10 times) larger than the number of dimensions was the estimated d-prime close to the true d-prime.

If the number of trials is less than the dimensionality, the covariance matrix is no longer invertible and needs to be “regularized” (Friedman 1989; Guo et al. 2007). The `fitcdiscr` function in MATLAB that was used to perform the LDA analysis applies

some regularization by default, but it can be further modified by using the `covshrink` function in MATLAB that can be used to diagonalize and reduce the dimensions of the covariance matrix. The modification of the covariance matrix depends on two regularization parameters, *gamma* and *delta*. The parameter *gamma*, which ranges from 0 to 1, shrinks the covariance matrix toward a diagonal matrix. In our simulations, we used 30 different *gamma* values and chose the one that minimized the classification error. The additional parameter *delta* sets a threshold for choosing the useful predictors (dimensions) and removes the remaining ones, essentially shrinking the covariance matrix. However, varying *delta* did not give any advantage over varying just *gamma* in our data, and varying *gamma* also did not improve the performance over the default regularization performed by `fitcdiscr` (Supplementary Fig. 2). For real data, therefore, we did not do any specific regularization.

Results

Effect of Neutral Cueing on LFP

Neural recordings were made using two 6×8 (48 channels) arrays implanted bilaterally in visual area V4 of both hemispheres in two rhesus monkeys (*M. mulatta*). Figure 1A shows the average PSTH of 756 neurons from 25 recording sessions in two monkeys, aligned to either the stimulus onset (left vertical dashed line) or target onset (right vertical dashed line). Here, attend-in and attend-out conditions are for validly cued stimulus changes that the monkeys correctly detected (Attend-in Valid Hit and Attend-out Valid Hit; see Materials and Methods for details), while the neutral condition includes all neutrally cued trials in which the monkeys detected the target, irrespective of target location (Target-in Neutral Hit + Target-out Neutral Hit; see Materials and Methods for details). There was a transient response after stimulus onset, followed by a 20 Hz SSVEP in the PSTH due to the 10-Hz counterphase stimulus. Throughout each trial, attention state had a small effect on spiking, with activity during the neutral condition (green trace) below attend-in (blue) and above attend-out (red) conditions, as reported previously (Mayo and Maunsell 2016).

Figure 1B shows the event-related potential (ERP), obtained by averaging the raw LFP traces across trials without any baseline correction, of the same three attention conditions. Like the PSTH, the ERP had a stimulus onset-related transient and a prominent 20 Hz SSVEP. Figure 1C shows the PSD of the attention conditions for baseline, stimulus, and pre-target periods (indicated as thick horizontal black lines in Fig. 1A). In all conditions, the PSDs showed a typical “1/f” power-law decay. The PSDs for the baseline condition were largely overlapping. In the stimulus onset and pre-target conditions, there was a small separation in the three traces, with the neutral (green) trace almost always in between the attend-in and attend-out conditions. These PSDs also showed prominent SSVEP bumps at 20 Hz and its harmonics. These SSVEP peaks were best observed in the pre-target period because of better frequency resolution (2 Hz, compared to 4 Hz for the other two periods). These plots show that the effect of attention on PSDs is weak at an absolute scale.

Figure 1D shows the change in power in the attend-in and attend-out conditions relative to the neutral attention condition for the three periods, obtained by subtracting the green trace from the other traces. Because PSDs are shown on a log scale in Figure 1C, subtraction on a log scale is simply the log of the

ratio of power change across conditions (after subtraction, the log ratio is multiplied by 10 to get units of decibels). Previous studies have shown that attention suppresses power at low frequencies and increases power at higher frequencies (Fries et al. 2001, 2008). Consistent with this, during the pre-target period (Fig. 1D, bottom row, where the sampled brain activity presumably reflected the animals' state of attention most accurately before the impending, successfully detected, target change), we observed that the power in lower frequencies was suppressed for attend-in condition relative to attend-out and the effect reversed at higher frequencies (>30 Hz), with the neutral condition almost always between the two. We also observed an increase in power at 20 Hz SSVEP (indicated by a black vertical dashed line). The increase in power at high frequencies was broadband and extended up to 200 Hz and beyond, with a slight peak in the gamma range (30–80 Hz) in the pre-target condition. An increase of 0.5 dB corresponds to a change in $10^{0.05} = 1.12$, or 12% increase in power, comparable to the size observed in firing rates in previous spatial attention studies (Boudreau et al. 2006; Mayo and Maunsell 2016; Snyder, Yu, et al. 2018b). These changes were highly reliable across electrodes and sessions, yielding very small standard error of mean (shaded region around the traces). As shown below, this allowed better discrimination between different attention conditions using LFP power compared with spiking activity.

Previous studies have suggested that phase synchronization between LFPs from two electrodes as well as between spikes and LFPs increases in the gamma range with spatial attention in area V4 (Fries et al. 2008; Fries 2015). We therefore measured phase consistency across electrode pairs using PPC, an unbiased estimator of the square of the phase coherence. Figure 2A shows the field–field (LFP–LFP) PPC across all pairs of selected electrodes within the hemisphere. As before, the neutral condition was between the attend-in and attend-out at almost all frequencies. Peaks were observed at 20 Hz and its harmonics because the counterphasing stimulus forced phase consistency across electrodes. Figure 2B shows the change in PPC from the neutral condition, obtained by subtracting the green trace from each of the traces in Figure 2A. In contrast to the change in power (Fig. 1D), the change in PPC was prominent mainly in the gamma range (30–80 Hz), with negligible differences at higher frequencies (>120 Hz). This could be because phase synchronization measures are less sensitive at high frequencies because even tiny absolute shifts in the time domain correspond to large shifts in phase, leading to poorer phase consistency in the presence of small time shifts (Ray 2015). The spike–field (spike–LFP) PPC (Fig. 2C) showed strong peaks at 20 Hz and harmonics, indicating that spikes were tightly locked to the SSVEP. The change in PPC relative to the neutral condition (Fig. 2D) showed similar trends to that in field–field PPC, but the magnitude of the effect was far smaller. There was also a reduction in phase coherence with attention at very low frequencies (<5 Hz) in both field–field and spike–field coherence, possibly due to the reduction in power at those frequencies with attention.

We also studied the phase relationship between the LFP signals across hemispheres and observed peaks at 20 Hz SSVEP and its harmonics, as in the within-hemisphere case, in both field–field and spike–field PPC, although the magnitude was reduced (Supplementary Fig. 1A,C). In this case, the electrodes of each across-hemisphere pair have RFs in opposite visual hemifields. So, there were no “attend-in” and “attend-out” conditions; we instead took the average of attend-left and attend-right conditions and compared them with the neutral condition. The

difference between attend-left/attend-right and neutral conditions showed inconsistent trends (Supplementary Fig. 1B,D) with small peaks at 20 Hz SSVEP and harmonics.

Comparison of Different Neural Measures across Attention Conditions

Our main goal was to compare how well the animals' behavior could be predicted using spiking activity versus LFP power. To see the consistency of effects, we performed this comparison over 12 different attention conditions (see Methods for details). We first compared how firing rates and LFP power at different frequencies varied across different attention conditions. Rather than using LFP power at individual frequency bins, we computed LFP power in three popular frequency bands—alpha (8–12 Hz), gamma (42–78 Hz), and high-gamma (122–198 Hz). Averaging of power within a frequency range is needed because single-trial estimates of power at individual frequency bins are noisy (Jarvis and Mitra 2001; Srinath and Ray 2014; Chandran et al. 2018). However, because the noise in spectral estimation is largely independent across frequencies, averaging power within a large frequency range provides reliable estimates of the actual power on single trials. The lower limit of high-gamma range was set at a higher value than usual (122 Hz instead of 82 Hz) because we wanted a frequency range in which field–field coherence was close to zero and therefore unlikely to simply reflect the dynamics of the traditional gamma band. However, choosing the high-gamma range between 82 and 198 Hz yielded similar results.

In addition, we also considered the SSVEP power at 20 Hz. These results should be interpreted with caution; since while we observed an increase in power at 20 Hz with attention, there was a general reduction in power around this frequency (bottom plot in Fig. 1D), suggesting that the overall power at this frequency depended on two antagonistic mechanisms. Further, as described above, single-trial estimates of power are noisier when not averaged over several frequency bins.

Figure 3A shows the average firing rates (top row), as well as raw power in the high-gamma (second row), gamma (third row), alpha ranges (fourth row), and SSVEP (bottom row) for 12 different attention conditions. The conditions are grouped based on cue validity (left column: valid, middle column: neutral, and right column: invalid), and within each group, conditions are arranged such that the firing rate was expected to increase going from left to right, as discussed below.

For the valid conditions (Fig. 3A; left column), firing rates were lower for Attend-out Valid Hit (first data point) than Attend-in Valid Hit (fourth data point), as shown previously (Fig. 1). Similar trends were observed in high-gamma, gamma power, and SSVEP power, while alpha power showed the opposite result. For the same attentional conditions, the monkeys were expected to miss the target if their attention was deployed on the “wrong” side, and therefore, the corresponding miss conditions (second and third data points) were expected to lie between the two hit conditions. This was the case for firing rate, high-gamma power, gamma power, and alpha power, barring a few exceptions (Attend-in Valid Miss condition for firing rates and high-gamma power, and Attend-out Valid Miss condition for alpha power). Interestingly, the SSVEP results were different, in which power was lower for both miss conditions compared with the hit conditions. Potential reasons behind these differences are discussed later.

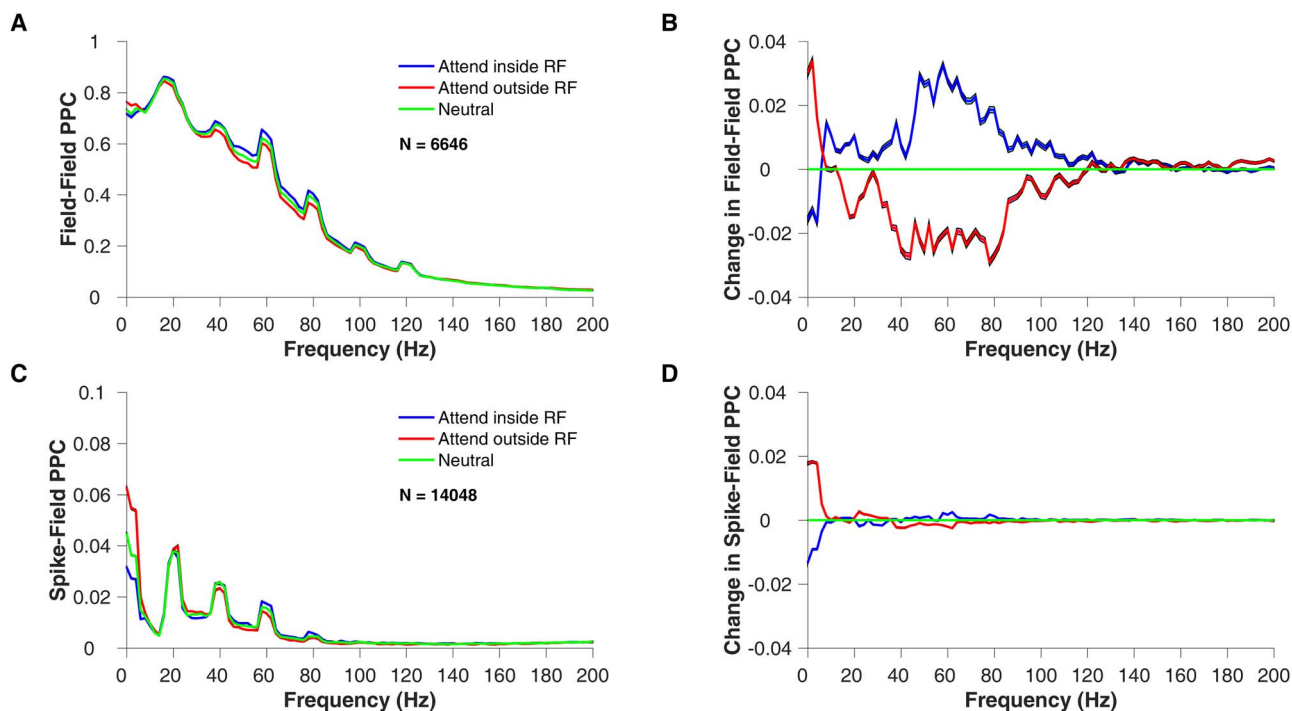


Figure 2. Comparison of intrahemispheric field-field (LFP-LFP) and spike-field (Spike-LFP) phase relationship across attend-in, attend-out, and neutral attention conditions. (A) Mean field-field PPC, averaged across 6646 electrode pairs of attend-in (blue trace), attend-out (red), and neutral conditions (green). (B) Mean change in field-field PPC of the three attention conditions, relative to the neutral condition. As before, the neutral condition is trivially zero. (C) Mean spike-field PPC averaged across 14048 electrode pairs (including spike and LFP from the same electrode) for the three attention conditions. (D) Mean change in spike-field PPC for the three attention conditions, relative to the neutral condition.

For the neutral case (Fig. 3A; middle column), the attention conditions were divided based on the trial outcome and target location rather than cue location since monkeys were cued to attend to both locations. Like the changes in firing rate, high-gamma power and gamma power were lower when monkeys detected the target that appeared outside the RF than when it appeared inside the RF, while alpha power showed the opposite trend. As before, barring a few exceptions (Target-in Neutral Miss condition for high-gamma, gamma, and alpha power), missed conditions lay between the hit conditions for the four neural measures except SSVEP, for which the missed conditions had lower power than the two hit conditions. Importantly, almost all the neutral attention conditions lay between validly cued attend-out and attend-in conditions (marked by solid horizontal lines for clarity), for all the measures excluding SSVEP.

In the invalid case (Fig. 3A; right column), the target always appeared at the uncued location. Firing rate, high-gamma, and gamma power were lowest when attention was directed outside the RF and the monkey missed the target, as expected (first data point). Likewise, these measures were highest when the monkey attended inside the RF but missed an invalidly cued stimulus change outside the RF, likely because attention remained focused on the unchanged stimulus within the RF (fourth data point). These two miss conditions were comparable to the Attend-out Valid Hit and Attend-in Valid Hit conditions, respectively, regarding the focus of attention. As before, alpha power showed the opposite effect. The hit trials, for which the monkeys were likely attending to the uncued side where the target actually appeared, lay between the missed trials for the invalid condition for all the neural measures except SSVEP,

which showed higher power in hit trials compared with the miss trials. Overall, unlike the other four measures, SSVEP power was higher for hit versus miss trials, irrespective of attentional location.

These results show that both firing rate and LFP power in different frequency bands vary systematically with the attentional state and therefore could potentially be used to predict the animals' behavior. To directly compare the discriminability of different neural measures, we z-scored all measures separately for each electrode and attention condition, by subtracting the mean and dividing by the standard deviation across trials of the Attend-out Valid Hit condition. The z-scored values therefore indicated how far a particular attention condition was from the Attend-out Valid Hit condition, in units of the standard deviation of Attend-out Valid Hit condition across trials. This condition was chosen for simplicity: The discriminability remained similar if we normalized all measures relative to Attend-in Valid Hit condition or used the average standard deviation of the Attend-out and Attend-in Valid Hit conditions.

Figure 3B shows the mean z-scored values (averaged across electrodes) for different attention conditions for all five neural measures, as indicated by line colors. The difference in z-scored activity between Attend-out Valid Hit (which was zero by definition) and Attend-in Valid Hit was greatest for gamma power ($z = 0.34 \pm 0.0091$; mean ± 1 SEM), followed by high-gamma (0.31 ± 0.0094), firing rate (0.2 ± 0.0083), alpha (-0.1 ± 0.0051), and finally SSVEP (0.09 ± 0.009) in terms of magnitude. The difference in z-scores between gamma and high-gamma was not significant ($P = 0.24$; Wilcoxon rank sum test), but the difference was highly significant between high-gamma

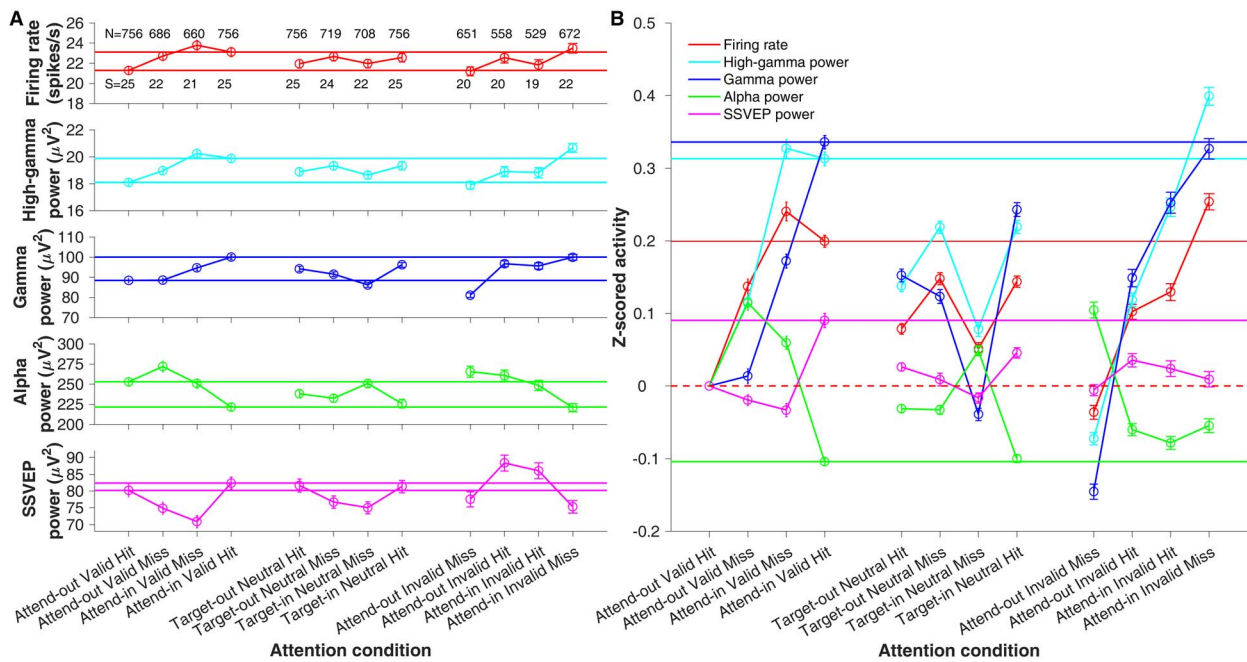


Figure 3. Comparison of firing rate, high-gamma power, gamma power, alpha power, and SSVEP power across 12 attention conditions. (A) First row: mean firing rate during the pre-target period for the Hit and Miss trials where attention was validly cued out and into the RF (left column), for Hit and Miss trials of neutral condition where the target appeared out and in the RF (middle column), and for Hit and Miss trials where attention was invalidly cued out and into the RF (right column). The numbers above the error bars indicate the number of electrodes averaged across for the respective conditions, while the numbers below the error bars indicate the number of sessions. We only considered sessions that had more than 15 trials for an attention condition. Second, third, fourth, and fifth rows: mean high-gamma, gamma, alpha, and SSVEP power, respectively. The same number of electrodes and sessions were used for these measures as well. Error bars indicate the SEM. The validly cued attend-out and attend-in conditions are indicated by horizontal lines for better comparison with the other conditions. (B) Same data as (A), but after z-scoring the data of each electrode and attention condition by subtracting the mean and dividing by the standard deviation (computed across trials) of the attend-out valid hit condition. The colored horizontal lines indicate the validly cued attend-in values for each neural measure. The red horizontal dashed line indicates the validly cued attend-out condition, which is trivially zero.

and firing rate ($P = 2.3 \times 10^{-21}$; Wilcoxon rank sum test) as well as gamma and firing rate ($P = 3.84 \times 10^{-27}$; Wilcoxon rank sum test). On the other hand, magnitudes of the z-scored values for alpha power ($P = 2.68 \times 10^{-10}$; Wilcoxon rank sum test) and SSVEP power ($P = 6.83 \times 10^{-25}$; Wilcoxon rank sum test) were significantly less than firing rates. Therefore, if a single electrode is available for analysis, gamma/high-gamma power allows superior discrimination of attend-in and attend-out conditions compared with firing rates.

Previous studies have shown that when firing rates are recorded from multiple electrodes simultaneously, a suitable linear combination of these signals can substantially improve the discriminability across attentional conditions, and this discriminability further depended on the correlation between the total number of spikes recorded from two electrodes (“spike-count” correlation; Cohen and Maunsell 2009). Analogous correlations between LFP power across electrode pairs have not been compared, to our knowledge, with spike-count correlations (but see Snyder et al. 2015 for related work). We found that the absolute Pearson correlation between power recorded from two electrodes within the same hemisphere was highest for alpha power (Attend-out Valid Hit condition: 0.84 ± 0.0015), progressively decreasing for SSVEP, gamma, and high-gamma power (0.76 ± 0.0019 , 0.69 ± 0.0018 , and 0.46 ± 0.0019), and lowest for spike-count correlations (0.22 ± 0.002). However, the relative change in correlation from the Attend-out Valid Hit condition, as shown in Figure 4A, showed similar trends across the

neural measures, except SSVEP. For valid cues (left column), spike-count correlations decreased from attend-in to attend-out conditions, consistent with previous studies (Cohen and Maunsell 2009; Mitchell et al. 2009). Interestingly, we found that correlation of high-gamma, gamma, and alpha power was also reduced in Attend-in Valid Hit condition compared with Attend-out Valid Hit. However, the correlation of SSVEP power increased. As before, the “missed” cases showed intermediate correlations between the validly cued hit conditions in most cases.

For the neutral conditions (middle column), previous study (Mayo and Maunsell 2016) showed that spike-count correlations were intermediate between attend-in and attend-out. However, in their data, the correlations were computed for all correct trials, irrespective of the target position (in other words, Target-out Neutral Hit and Target-in Neutral Hit conditions were pooled, as done in Figs 1 and 2 here as well). However, when we computed the correlation separately for these two Hit conditions, we found that spike-count correlations for Target-out Neutral Hit (0.218 ± 0.002) were comparable to Attend-out Valid Hit (0.22 ± 0.002), although the difference was significant ($P = 0.007$; Wilcoxon rank sum test). Similarly, Target-in Neutral Hit (0.189 ± 0.002) was similar to Attend-in Valid Hit (0.19 ± 0.002 , $P = 0.64$; Wilcoxon rank sum test). The Miss conditions were again intermediate between the two Hit conditions in most cases. Finally, the invalid conditions (right column) showed predictable trends for firing rates, but variable results for the LFP power, especially correlations in gamma power that increased

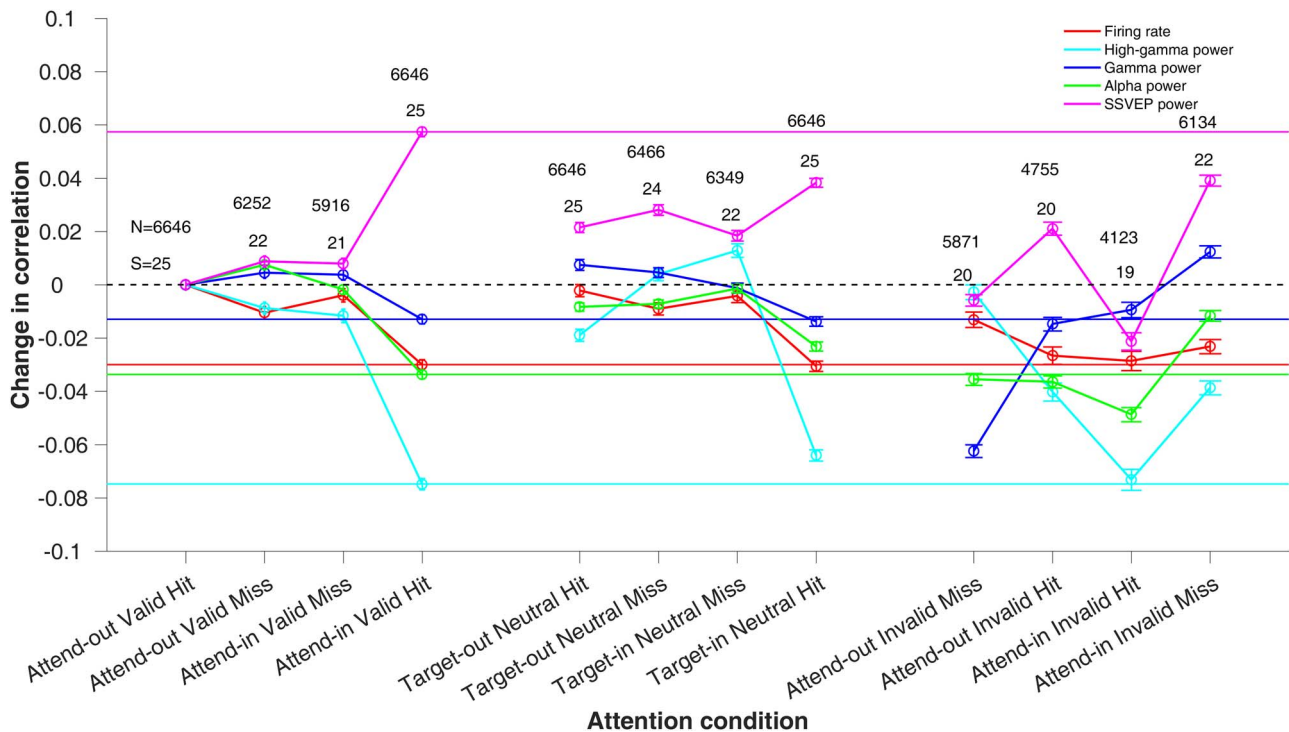


Figure 4. Intrahemispheric change in Pearson correlation of firing rate, high-gamma, gamma, alpha, and SSVEP power. Mean change in correlations between intrahemispheric electrode pairs for the five neural measures relative to attend-out valid hit condition (same format as Fig. 3B). Numbers above the error bars of the blue line indicate the total number of electrode pairs (N) and sessions (S) considered for each condition. As in Figure 3B, the colored horizontal lines indicate the validly cued attend-in values for each neural measure. The black horizontal dashed line indicates the validly cued attend-out condition, which is trivially zero.

between Attend-out Invalid Miss versus Attend-in Invalid Miss. Potential reasons for this surprising effect are addressed in the Discussion. Interhemispheric correlations were relatively weak (absolute correlation values for the valid hit condition were 0.273 ± 0.0018 , 0.235 ± 0.0015 , 0.167 ± 0.0014 , 0.104 ± 0.0015 , and 0.044 ± 0.0012 for alpha power, SSVEP power, gamma power, high-gamma power, and firing rate, respectively) and difficult to interpret (data not shown).

Discriminability between Attention Conditions Using Population Activity

Although the analysis shown in Figure 3 suggests that high-frequency LFP power is more discriminable than firing rates, this analysis only reflects the average discriminability computed from individual electrodes. Previous studies have shown that while single-channel LFPs outperform single-channel spikes, decoding based on multichannel spikes is often superior to multichannel LFPs (Hwang and Andersen 2013) (see Discussion for details). This is because spiking activity is more dissimilar across electrodes as compared with LFP power (spikes have lower absolute correlations, as seen in our dataset as well), and consequently, discriminability between different attentional states is more variable across electrodes for firing rates than LFP power. Therefore, a weighting strategy that gave more weight to the more discriminable electrodes could improve the overall discriminability of the pooled signal (Jazayeri and Movshon 2006; Graf et al. 2011), and improve the overall population discriminability of firing rates compared with LFP power. This can also be explained from an information theoretic viewpoint since higher variability in firing rates compared with LFP power

implies higher system entropy and consequently higher information capacity. The optimum weights for a linear decoder can be obtained using LDA. We used three pooling strategies (see Materials and Methods) to illustrate three factors (mean, variance, and covariance) that are used in LDA analysis. We show these results first without any cross-validation since the effect size is larger in that case (Fig. 5A). The results after 5-fold cross-validation are shown in Figure 5B.

In the first pooling strategy (Fig. 5A), we assigned a weight to each electrode that was equal to the difference in neural activity between the Attend-left Valid Hit and Attend-right Valid Hit conditions. The weighted average of all electrodes was then taken to obtain the projections for individual trials. The weights were defined such that attention to the left side led to larger projections than attention to the right side. The “Attention Axis” approach, which involves taking an inner product of the data vector with a vector created by connecting the mean of the data points for the Attend-left Valid Hit and Attend-right Valid Hit conditions, essentially performs the same operation (Cohen and Maunsell 2010). To facilitate unbiased comparisons in discriminability across our five neural measures, we normalized all the projections by subtracting the mean value of the projections for the Attend-right Valid Hit condition and dividing by the standard deviation of the same condition (essentially z-scoring, as done in Fig. 3B as well). Therefore, the mean projections for the Attend-right Valid Hit case were zero by definition, while these values were positive for Attend-left Valid Hit conditions, with larger values corresponding to better population discriminability.

Figure 5A (first row) shows the projections for all 12 attention conditions. Note that the computation of the weights was done using only 2 of these 12 conditions, the first and fourth data

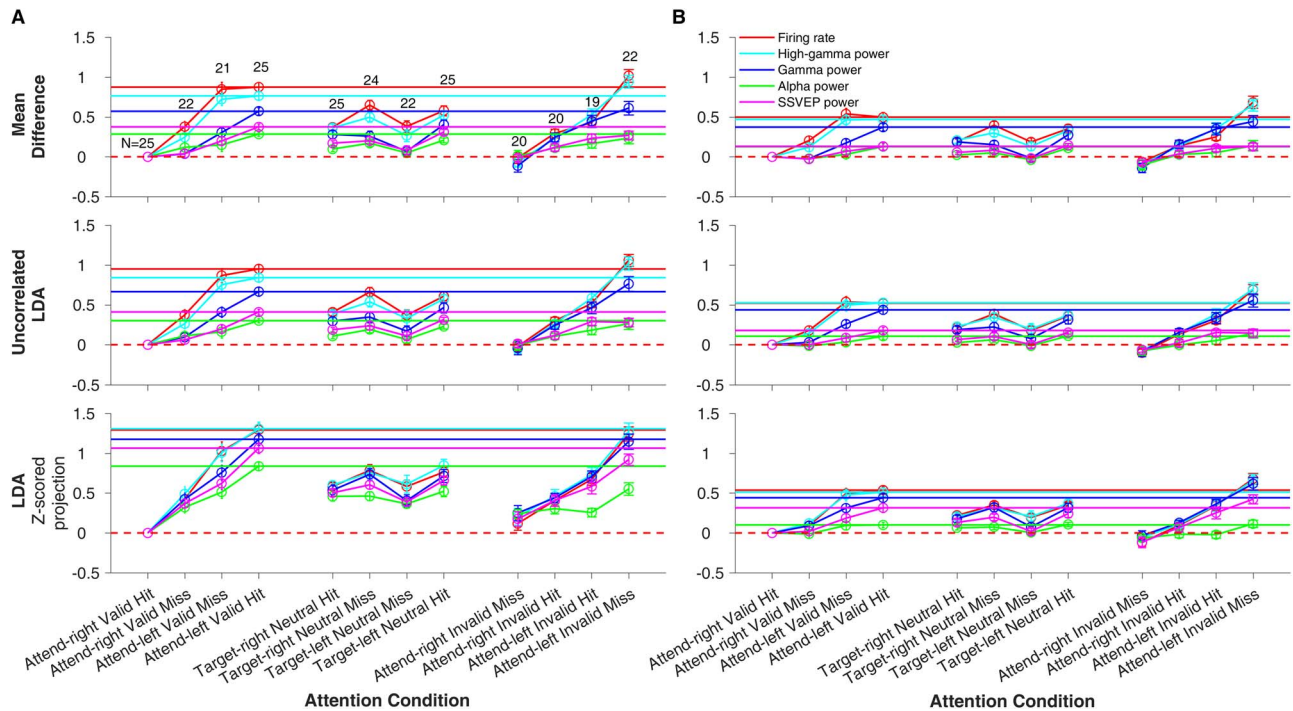


Figure 5. Comparison of projections onto mean difference, uncorrelated LDA, and LDA weight vectors for different attention conditions. (A) Mean z-scored projections of firing rate (red trace), high-gamma power (cyan), gamma power (blue), alpha power (green), and SSVEP power (magenta) of trials of each attention conditions using non-cross-validated weight vectors computed using mean difference (top row), uncorrelated LDA (middle row), and LDA (bottom row) methods (see text for details). The horizontal lines indicate the mean z-scored projections of the attend-left valid condition of the five neural measures (represented by their respective colors). The horizontal red dashed line indicates the z-scored projection of attend-right valid hit condition, which is trivially zero since z-scoring is done relative to that condition (see Materials and Methods). The numbers on top of the data points in the top row indicate the number of sessions as before. Error bars indicate the SEM. (B) Same as that of (A) but with 5-fold cross-validation.

points in the left column, and subsequent z-scoring was done relative to the first data point. For the Attend-left Valid Hit condition, firing rates showed the largest value (0.87 ± 0.06), followed by high-gamma (0.76 ± 0.05); but the difference between the two was not statistically significant, $P = 0.16$, Wilcoxon rank sum test), gamma (0.57 ± 0.05 , significantly smaller than firing rates, $P = 1.67 \times 10^{-4}$ and high-gamma, $P = 0.003$, Wilcoxon rank sum test), SSVEP (0.38 ± 0.04 ; significantly smaller than firing rate: $P = 8.29 \times 10^{-7}$, high-gamma: $P = 1.16 \times 10^{-5}$, and gamma: $P = 0.005$; Wilcoxon rank sum test), and alpha power (0.28 ± 0.03 , significantly smaller than firing rates: $P = 3.21 \times 10^{-8}$, high-gamma: $P = 1.06 \times 10^{-7}$, gamma: $P = 8.1 \times 10^{-6}$; not significantly different from SSVEP power: $P = 0.11$; Wilcoxon rank sum test). The Valid Miss conditions (second and third data points in the left column) were intermediate between the two Valid Hit conditions in all cases (including SSVEPs), as were the projections for all the Neutral conditions (middle column). Invalid conditions (right column) showed similar results as the Valid (with the Miss and Hit labels flipped, as in Figs 3B and 4).

The results were similar when we scaled individual weights by their variances (second row). However, when we also accounted for the covariances in the computation of weights (as done in LDA), we observed a large increase in the projection values for gamma, SSVEP, and alpha power for which the absolute correlation was high (third row). The projections of the Attend-left Valid Hit for firing rate (1.29 ± 0.07) and high-gamma (1.3 ± 0.08) were not significantly different ($P = 1$; Wilcoxon rank sum test), as before. Surprisingly, the projections for gamma power (1.18 ± 0.08) were now statistically indistinguishable to

high-gamma ($P = 0.26$; Wilcoxon rank sum test) and firing rates ($P = 0.3$; Wilcoxon rank sum test). The projections for SSVEP (1.06 ± 0.06) were comparable to that of gamma power ($P = 0.32$; Wilcoxon rank sum test) but were significantly smaller than firing rate ($P = 0.03$; Wilcoxon rank sum test) and high-gamma power ($P = 0.04$; Wilcoxon rank sum test). Comparable changes in discriminability across neural measures were also observed for the Invalid conditions (right column), while the Neutral conditions had intermediate projections (middle column).

The results in Figure 5A were obtained without any cross-validation, which tends to inflate the discriminability (Subramanian and Simon 2013) because the training and testing data are the same, which leads to overfitting. Figure 5B shows the same results using 5-fold cross-validation in which data were divided into five equal parts, out of which the weight vector was obtained using four parts (training data) and the projections were computed for the trials in the remaining part (testing data). This was repeated over by considering different testing set exhaustively among those five parts and then averaging across them. The results were similar to Figure 5A, although now there was not much difference between the three methods. For example, the mean projection values for Attend-left Valid Hit condition for firing rates were 0.48 ± 0.06 , 0.57 ± 0.06 and 0.55 ± 0.05 ($P = 0.45$, Kruskal-Wallis test) for mean difference, uncorrelated LDA, and LDA, respectively. We show using simulations that this reduction in the LDA projection values was because the number of trials was not substantially larger than the number of dimensions (electrodes), especially after cross-validation for which the number of training trials

was 80% of the total (Supplementary Fig. 2). Despite this problem that led to lower projection values, the overall trends remained similar, with firing rate projections for the Attend-left Valid Hit (0.55 ± 0.05) not significantly different from high-gamma (0.52 ± 0.06 ; $P = 0.79$; Wilcoxon rank sum test) and gamma (0.44 ± 0.06 ; $P = 0.42$; Wilcoxon rank sum test) but were higher compared with SSVEP (0.32 ± 0.06 ; $P = 0.003$; Wilcoxon rank sum test). The SSVEP projection remained comparable to that of gamma ($P = 0.09$; Wilcoxon rank sum test). Perhaps because alpha power showed the weakest effects without cross-validation, it showed severely limited discriminability after cross-validation. This suggests that using high-frequency LFP is equally useful in discriminating between attention conditions as firing rates, even after using population data from several microelectrodes. As before, the invalid condition (right column) yielded similar results as valid (left column), while neutral condition (middle column) yielded intermediate values.

Performance of Spikes versus LFP in Decoding Attentional State

Finally, we studied how discriminability varied with the number of available electrodes. Figure 6 shows the z-score values of the Attend-left valid hit condition with cross-validation as a function of the number of electrodes, obtained by first selecting a subset of electrodes without replacement and performing the full LDA analysis (third pooling strategy) as before on this subset. When the number of electrodes was small, high-gamma and gamma power performed significantly better than the firing rate, SSVEP, and alpha power (similar trends as in Fig. 3B). As the number of electrodes increased to approximately 12, the discriminability of firing rate became comparable to gamma and high-gamma power. The discriminability reached an asymptote after approximately 15 electrodes, but this could be due to poorer performance of LDA with increasing dimensionality (Supplementary Fig. 2). Discriminability also increased for SSVEP power and alpha power for up to a few electrodes before reaching an asymptote but remained much lower than the other three measures. Overall, these results show that high-frequency LFP power is more useful than spiking activity when few electrodes (roughly 10–12 electrodes) are available and equally useful as firing rate with higher number of electrodes. The results were qualitatively similar without cross-validation, although the absolute values were larger (Supplementary Fig. 3).

Discussion

We show that nearly every measure of neural activity, including PSD, phase consistency, noise correlations, and discriminability, during neutral cueing is intermediate to those in the cued and uncued conditions, consistent with the effects observed earlier in spiking activity and behavior (Mayo and Maunsell 2016). For single-electrode analysis, the discriminability between attention conditions was higher for both gamma (42–78 Hz) and high-gamma (122–198 Hz) power compared with spiking. When multiple electrodes were used, population projections obtained using weights proportional to the difference in mean activity between correctly detected attend-left and attend-right conditions (as done in the “Attention Axis” analysis) were equally discriminable with spikes and high-gamma but poorer with gamma and alpha. However, once the full covariance matrix was incorporated using LDA analysis, the discriminability of high-frequency LFP power (both gamma and high-gamma) was comparable to

spiking activity. These results were largely consistent across a wide range of attention conditions, including valid, invalid, and neutral cues.

Relation to Previous Studies

LFPs have been previously used to decode stimulus information (Belitski et al. 2008; Kanth and Ray 2020), saccade location (Pesaran et al. 2002), allocation of attention (Tremblay et al. 2015), motor movements (Bansal et al. 2012), and in brain-machine interfacing (BMI) applications (Hwang and Andersen 2013; Andersen et al. 2014). Consistent with our findings, many of these studies showed high decodability using high-frequency (>40 Hz) LFP. Similarly, BMI studies have shown that while single-channel LFPs offer better decoding of reach targets than single-channel spikes, decoders using multichannel LFPs performed worse than multichannel spikes because of larger noise correlations in LFPs than spikes, with comparable spike and LFP performance with approximately eight electrodes (Hwang and Andersen 2013).

A recent study analyzed neuronal responses in macaque primary visual cortex (V1) during attend-in, attend-out, and attend-both (neutral) conditions to test whether spike-count correlations show a graded response as a function of the focus of attention (Denfield et al. 2018). They showed that the effect of attention on spike-count correlations depended on the duration over which analysis is performed. For short durations (100–200 ms), attention decreased spike-count correlations and neutral cueing produced intermediate correlations (see their Fig. 5C). However, for longer analysis intervals, correlations for attend-in and attend-out were comparable, and attend-both produced the largest correlations. They suggested that the increase in correlation could be due to a switch in the focus of attention in the attend-both condition when observed over long time scales, which has also been reported in other studies (Landau and Fries 2012; Landau 2018; Fiebelkorn and Kastner 2019). However, here we performed all analyses over 500 ms (mainly to improve the frequency resolution for spectral analysis of the LFP) and found that neutral cueing produced intermediate correlations when averaged over the two target locations (although, when separated based on the target location, correlations for target-in and target-out neutral hits were comparable to attend-out and attention-in valid hit conditions; Fig. 4). We also found no evidence of attention switching, which could have led to negative spike-count correlations for interhemispheric pairs of electrodes. Instead, we found interhemispheric spike-count correlations to be small (0.044 ± 0.0012), consistent with previous reports where this analysis was done over short (200 ms) periods (Cohen and Maunsell 2010; Mayo and Maunsell 2016). Likewise, correlations in the LFP power at different frequencies were also not negative. The results remained similar when we used a 250-ms analysis window. We note that non-negative correlations do not necessarily rule out attentional switching, because it is possible that there is some other source of covariation with greater magnitude than attention shifts, such as a global “slow drift” of neural activity (Cowley et al. 2020). Denfield et al. (2018) could not quantify the attentional switch using interhemispheric spike-count correlations because they did not record from both hemispheres simultaneously. It is possible that attentional switching is difficult to capture using interhemispheric spike-count correlations and may need more sensitive methods (Landau 2018; Fiebelkorn and Kastner 2019). The discrepancy between studies could also be due to differences in experimental design (e.g.,

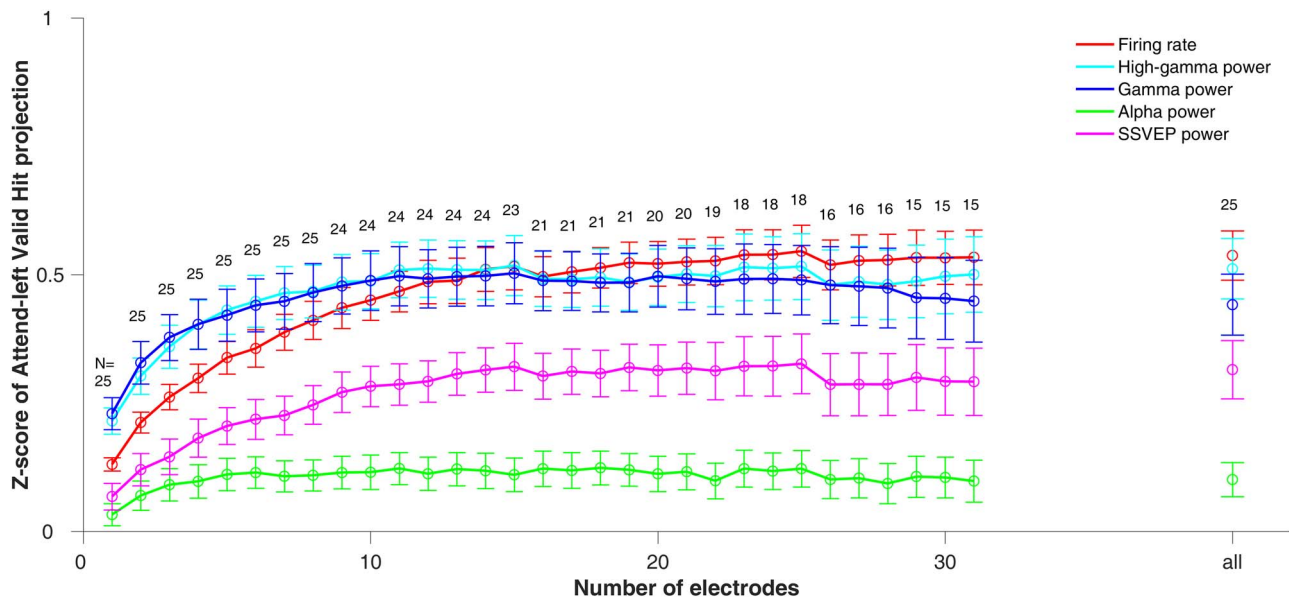


Figure 6. Discriminability as a function of the number of electrodes. Mean z-score of 5-fold cross-validated LDA projection of firing rate (red), high-gamma (cyan), gamma (blue), alpha power (green), and SSVEP power (magenta) of attend-left valid hit condition relative to attend-right valid hit condition as a function of the number of electrodes (dimensions) used to calculate the projections. This analysis is done for electrode numbers for which at least 15 sessions were available, as indicated by the number on top of the error bar. The data point at the right end represents the projection calculated using all the electrodes of a session, which is also plotted in Figure 5.

Denfield et al. (2018) had 100% validity and no catch trials in the non-neutral condition, in order to minimize attention switching) that could have encouraged different attention strategies on the part of their animals.

Effect of Attention on Gamma Oscillations

Previous studies have suggested a potentially important role of the phase synchronization of gamma oscillations in communication (Fries et al. 2001; Fries 2015). Even when gamma oscillations are not strong enough to be visible as a distinct bump in the PSD, measures such as field-field and spike-field coherence often show such a bump, suggesting a potentially important role of phase synchronization (Fries et al. 2008). We found similar results in our data; while power increased over a wide range of frequencies between approximately 40 and 200 Hz and beyond with attention, the increase in phase synchronization was localized to a narrower range with a clear peak between 40 and 80 Hz. However, this does not necessarily imply the presence of narrow-band phase synchronization, because measures of phase synchronization could themselves be less sensitive at higher frequencies. For example, small offsets in time correspond to larger fluctuations in phase at higher frequencies (Ray 2015). Indeed, in spite of negligible coherence, the discriminability of high-gamma (>122 Hz) was comparable (using LDA) or higher (other methods) than gamma (42–78 Hz) that showed higher phase synchronization. Further, spikes were not locked to the phase of gamma rhythm but instead to the stimulus that was counterphasing at 10 Hz. Therefore, even though we observed increased field-field synchronization with attention in the gamma range, whether this synchronization has a functional role is unclear (Ray and Maunsell 2015).

While correlations in power across electrodes showed similar trends as spiking activity during valid and neutral cueing, the trends were different during invalid cueing, especially for

gamma oscillations (Fig. 4). We interpret these results with caution since the number of invalid trials was smaller than valid and neutral trials. Nonetheless, the striking difference in the correlation trends for gamma versus other measures could provide clues about switching of attention in invalid trials. Attention has been divided into several “subcomponents” such as orienting, target detection, and vigilance, which may be carried out by distinct brain areas (for a review, see Posner and Petersen 1990). Even orientation of attention involves disengagement from the present location, shifting, and finally redeployment to another location, which may involve the interaction of many brain areas (Posner and Petersen 1990). Importantly, to detect an invalid target, a potentially large shift of attention may be required (see Fig. 3 of Mayo and Maunsell 2016). Since gamma oscillations have been linked to a variety of neural mechanisms such as normalization (Ray et al. 2013) and gain control (Ni et al. 2016) that could be involved in attentional switching, the modulation of gamma could be different from other measures in the invalid case. More detailed experiments in which various subcomponents of attention are manipulated independently are needed to test these predictions.

Correlations were also modulated differently for firing rates versus LFP power when comparing hits versus misses in some cases (Fig. 4). These differences could be related to the variety of different reasons behind a miss, including focused attention at the incorrect location, a unique momentary distraction, or a reduction in global attention at both locations. Power at alpha and gamma frequency bands, which could be influenced by global synchronization (Buzsaki 2006) and could be modulated differently for these conditions.

Effect of Attention on Alpha Power

While EEG studies in humans have shown a marked effect of attention on alpha power (Klimesch et al. 1998; Thut et al. 2006),

we did not observe a comparable effect in our study. There are several possible reasons for this. First, EEG signals reflect a more global population than LFP, and the changes in alpha power could be more prominent in spatially broader areas that could be captured in the EEG but not the LFP (Snyder, Issar, et al. 2018a). Second, we used very short analysis windows (250 or 500 ms duration, which correspond to only 2 or 3 data points within the alpha range) that may not have been long enough to capture alpha power reliably. Finally, we used a 10-Hz counterphasing stimulus. Although the most salient change was observed at twice the fundamental frequency (20 Hz), it is possible that a small fundamental at 10 Hz was also present, which interfered with the endogenous alpha rhythm.

Effect of Attention on SSVEP Power

The SSVEP power (Fig. 3A, bottom row) showed a stronger variation with behavioral outcome (hits versus misses; compare first with second and third with fourth data point of all the columns) than focus of attention (compare first with fourth data point of all the columns). As discussed earlier, overall LFP power was modulated by two antagonistic effects: a suppression in the low frequency range and an increase near the SSVEP frequency. It has been shown previously in EEG that attention modulates SSVEP power differently depending on its frequency (Gulbinaite et al. 2019). For example, attention reduces SSVEP power when its frequency is in the alpha range but increases when the SSVEP frequency is in the gamma range. In general, there are some important differences between a standard SSVEP paradigm used in EEG studies and our design (which was not specifically tailored to study SSVEPs) that makes a direct comparison of our SSVEP results with previous results difficult. In EEG studies, different stimuli are tagged at different frequencies, but in our case both stimuli were counterphasing at 10 Hz. Further, SSVEP responses are typically analyzed over long durations to obtain sufficient spectral resolution (typically at least a second or more), but we only analyzed segments of 250–500 ms. A more comparable SSVEP experimental design with multiple SSVEP frequencies that span a wider range is needed to test whether SSVEPs are able to capture some aspect of attentional mechanisms that are not captured by power in standard frequency bands.

Spatial Extent of Selective Attention

A surprising finding of this study is that discriminability of high-gamma remains comparable to spiking activity even when many electrodes are used (Fig. 6), unlike other studies in which multichannel spikes outperform multichannel LFPs (Hwang and Andersen 2013; Tremblay et al. 2015). This is reminiscent of the finding that electrocorticogram, which has a much larger spatial spread than LFPs (Dubey and Ray 2019), nonetheless show better decodability of natural images than LFPs or spikes (Kanth and Ray 2020). The optimal level of neural integration for decodability of a particular stimulus or behavioral state depends on the spatial extent over which neural assemblies remain coherent under that stimulus/state. For example, if spatial attention is mediated through circuits involved in normalization, as suggested recently (Reynolds and Heeger 2009; Verhoef and Maunsell 2017), the spatial extent of normalization could determine the spatial resolution of attention. Since such divisive normalization mechanisms incorporate an inhibitory network that is involved in generating gamma oscillations (Carandini and

Heeger 2011; Ray et al. 2013), the spatial resolution of attention could be related to the spatial spread over which gamma oscillations are coherent. We and others have previously shown that the spatial spread of LFP is local (radius of ~0.5 mm; Xing et al. 2009; Dubey and Ray 2016), with a slightly larger spread in the high-gamma range (Dubey and Ray 2016, 2019). If this spread matches the spatial extent of selective attention (which could be dependent on the constraints of the task), decodability is likely to be best using high-gamma power. Our results therefore provide an indirect measure of the spatial scale over which attention might be operating.

Supplementary Material

Supplementary material can be found at *Cerebral Cortex* online.

Author Contributions

J.P.M. recorded the data; S.S.P., J.P.M. and S.R. conceived the idea of research; S.S.P., A.D., S.T.K. and S.R. analysed the data; S.S.P. and S.R. wrote the initial draft of the manuscript; J.P.M., A.D., S.T.K. reviewed the draft.

Funding

Department of Biotechnology/Wellcome Trust India Alliance (Senior fellowship IA/S/18/2/504003 to S.R.); Department of Biotechnology-Indian Institute of Science Partnership Programme (to S.R.); National Institutes of Health Fellowship (F32EY02259 to J.P.M.). S.S.P., A.D., and S.T.K. were each supported by individual research fellowships in Indian Institute of Science, funded by the Ministry of Human Resource Development, Government of India.

Notes

We thank Drs Marlene Cohen, John Maunsell, and Chris Henry for helpful comments. *Conflict of Interest:* The authors declare no competing financial interests.

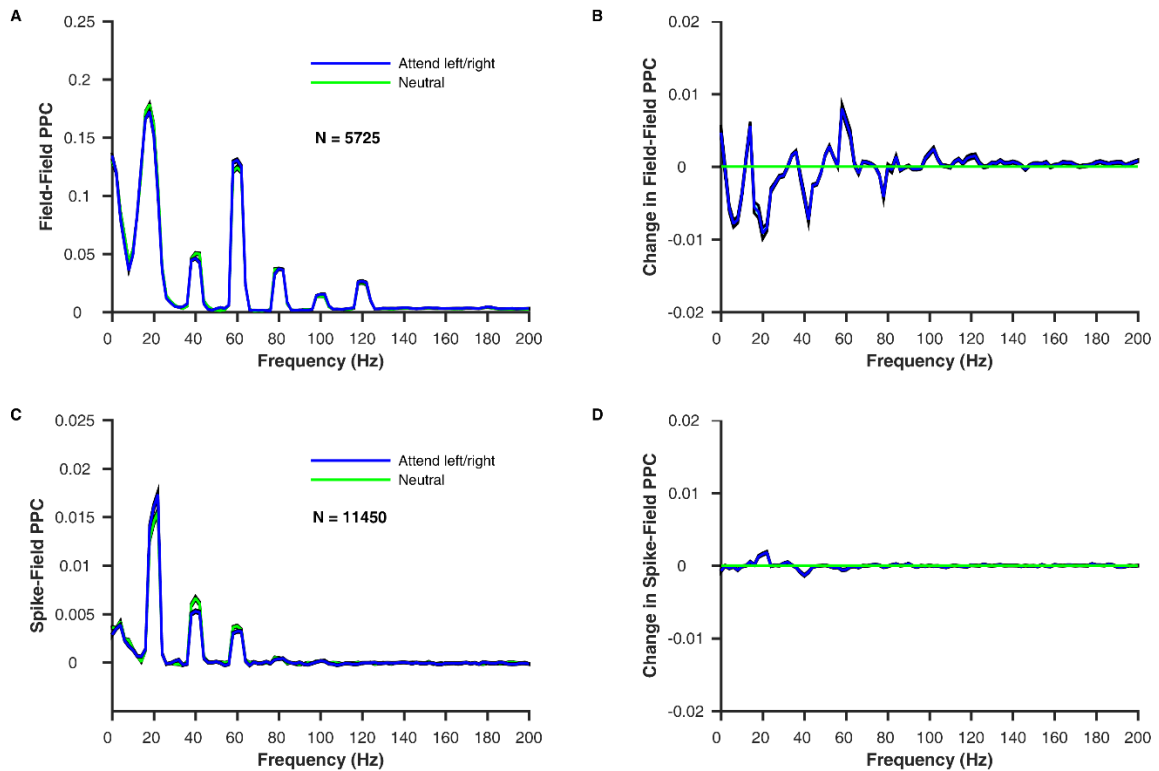
References

- Andersen RA, Kellis S, Klaes C, Aflalo T. 2014. Toward more versatile and intuitive cortical brain-machine interfaces. *Curr Biol*. 24:R885–R897.
- Averbeck BB, Latham PE, Pouget A. 2006. Neural correlations, population coding and computation. *Nat Rev Neurosci*. 7:358–366.
- Bansal AK, Truccolo W, Vargas-Irwin CE, Donoghue JP. 2012. Decoding 3D reach and grasp from hybrid signals in motor and premotor cortices: spikes, multiunit activity, and local field potentials. *J Neurophysiol*. 107:1337–1355.
- Belitski A, Gretton A, Magri C, Murayama Y, Montemurro MA, Logothetis NK, Panzeri S. 2008. Low-frequency local field potentials and spikes in primary visual cortex convey independent visual information. *J Neurosci*. 28:5696–5709.
- Bokil H, Andrews P, Kulkarni JE, Mehta S, Mitra PP. 2010. Chronux: a platform for analyzing neural signals. *J Neurosci Methods*. 192:146–151.
- Boudreau CE, Williford TH, Maunsell JH. 2006. Effects of task difficulty and target likelihood in area V4 of macaque monkeys. *J Neurophysiol*. 96:2377–2387.

- Buzsáki G. 2006. *Rhythms of the brain*. New York, USA: Oxford University Press.
- Buzsáki G, Anastassiou CA, Koch C. 2012. The origin of extracellular fields and currents — EEG, ECoG, LFP and spikes. *Nat Rev Neurosci*. 13:407–420.
- Carandini M, Heeger DJ. 2011. Normalization as a canonical neural computation. *Nat Rev Neurosci*. 13:51–62.
- Carrasco M. 2011. Visual attention: the past 25 years. *Vision Research, Vision Research 50th Anniversary Issue: Part 2*. 51:1484–1525.
- Chalk M, Herrero JL, Gieselmann MA, Delicato LS, Gotthardt S, Thiele A. 2010. Attention reduces stimulus-driven gamma frequency oscillations and spike field coherence in V1. *Neuron*. 66:114–125.
- Chandran KSS, Seelamantula CS, Ray S. 2018. Duration analysis using matching pursuit algorithm reveals longer bouts of gamma rhythm. *J Neurophysiol*. 119:808–821.
- Cohen MR, Kohn A. 2011. Measuring and interpreting neuronal correlations. *Nat Neurosci*. 14:811–819.
- Cohen MR, Maunsell JH. 2009. Attention improves performance primarily by reducing interneuronal correlations. *Nat Neurosci*. 12:1594–1600.
- Cohen MR, Maunsell JHR. 2010. A neuronal population measure of attention predicts Behavioral performance on individual trials. *J Neurosci*. 30:15241–15253.
- Cohen MR, Maunsell JHR. 2011. Using neuronal populations to study the mechanisms underlying spatial and feature attention. *Neuron*. 70:1192–1204.
- Cowley BR, Snyder AC, Acar K, Williamson RC, Yu BM, Smith MA. 2020. Slow drift of neural activity as a signature of impulsivity in macaque visual and prefrontal cortex. *Neuron*. 108:551–567.e8.
- Cunningham JP, Yu BM. 2014. Dimensionality reduction for large-scale neural recordings. *Nat Neurosci*. 17:1500–1509.
- Denfield GH, Ecker AS, Shinn TJ, Bethge M, Tolias AS. 2018. Attentional fluctuations induce shared variability in macaque primary visual cortex. *Nat Commun*. 9:2654.
- Dubey A, Ray S. 2016. Spatial spread of local field potential is band-pass in the primary visual cortex. *J Neurophysiol*. 116:1986–1999.
- Dubey A, Ray S. 2019. Cortical electrocorticogram (ECoG) is a local signal. *J Neurosci*. 39:4299–4311.
- Fiebelkorn IC, Kastner S. 2019. A rhythmic theory of attention. *Trends Cogn Sci*. 23:87–101.
- Fisher RA. 1936. The use of multiple measurements in taxonomic problems. *Ann Eugen*. 7:179–188.
- Friedman JH. 1989. Regularized discriminant analysis. *J Am Stat Assoc*. 84:165–175.
- Fries P. 2015. Rhythms for cognition: communication through coherence. *Neuron*. 88:220–235.
- Fries P, Reynolds JH, Rorie AE, Desimone R. 2001. Modulation of oscillatory neuronal synchronization by selective visual attention. *Science*. 291:1560–1563.
- Fries P, Womelsdorf T, Oostenveld R, Desimone R. 2008. The effects of visual stimulation and selective visual attention on rhythmic neuronal synchronization in macaque area V4. *J Neurosci*. 28:4823–4835.
- Graf ABA, Kohn A, Jazayeri M, Movshon JA. 2011. Decoding the activity of neuronal populations in macaque primary visual cortex. *Nature Neuroscience*. 14:239–247.
- Gulbinaite R, Roozendaal DHM, VanRullen R. 2019. Attention differentially modulates the amplitude of resonance frequencies in the visual cortex. *Neuroimage*. 203:116146.
- Guo Y, Hastie T, Tibshirani R. 2007. Regularized linear discriminant analysis and its application in microarrays. *Biostatistics*. 8:86–100.
- Hwang EJ, Andersen RA. 2013. The utility of multichannel local field potentials for brain-machine interfaces. *J Neural Eng*. 10:046005.
- Jarvis MR, Mitra PP. 2001. Sampling properties of the spectrum and coherency of sequences of action potentials. *Neural Comput*. 13:717–749.
- Jazayeri M, Movshon JA. 2006. Optimal representation of sensory information by neural populations. *Nat Neurosci*. 9:690–696.
- Judge SJ, Richmond BJ, Chu FC. 1980. Implantation of magnetic search coils for measurement of eye position: an improved method. *Vision Res*. 20:535–538.
- Kanth ST, Ray S. 2020. Electrocorticogram (ECoG) is highly informative in primate visual cortex. *J Neurosci*. 40:2430–2444.
- Khamechian MB, Kozyrev V, Treue S, Esghaei M, Daliri MR. 2019. Routing information flow by separate neural synchrony frequencies allows for “functionally labeled lines” in higher primate cortex. *PNAS*. 116:12506–12515.
- Khayat PS, Niebergall R, Martinez-Trujillo JC. 2010. Frequency-dependent attentional modulation of local field potential signals in macaque area MT. *J Neurosci*. 30:7037–7048.
- Klimesch W, Doppelmayr M, Russegger H, Pachinger T, Schwaiger J. 1998. Induced alpha band power changes in the human EEG and attention. *Neurosci Lett*. 244:73–76.
- Landau AN. 2018. Neuroscience: a mechanism for rhythmic sampling in vision. *Curr Biol*. 28:R830–R832.
- Landau AN, Fries P. 2012. Attention samples stimuli rhythmically. *Curr Biol*. 22:1000–1004.
- Li T, Zhu S, Ogihara M. 2006. Using discriminant analysis for multi-class classification: an experimental investigation. *Knowl Inf Syst*. 10:453–472.
- Mangun GR, Hillyard SA. 1990. Allocation of visual attention to spatial locations: Tradeoff functions for event-related brain potentials and detection performance. *Percept Psychophys*. 47:532–550.
- Mayo JP, Cohen MR, Maunsell JHR. 2015. A refined neuronal population measure of visual attention. *Plos One*. 10:e0136570.
- Mayo JP, Maunsell JHR. 2016. Graded neuronal modulations related to visual spatial attention. *J Neurosci*. 36:5353–5361.
- McAdams CJ, Maunsell JH. 1999. Effects of attention on orientation-tuning functions of single neurons in macaque cortical area V4. *J Neurosci*. 19:431–441.
- Mitchell JF, Sundberg KA, Reynolds JH. 2009. Spatial attention decorrelates intrinsic activity fluctuations in macaque area V4. *Neuron*. 63:879–888.
- Mitra PP, Pesaran B. 1999. Analysis of dynamic brain imaging data. *Biophys J*. 76:691–708.
- Montagna B, Pestilli F, Carrasco M. 2009. Attention trades off spatial acuity. *Vision Res*. 49:735–745.
- Moran J, Desimone R. 1985. Selective attention gates visual processing in the extrastriate cortex. *Science*. 229:782–784.
- Ni J, Wunderle T, Lewis CM, Desimone R, Diester I, Fries P. 2016. Gamma-rhythmic gain modulation. *Neuron*. 92:240–251.
- Pesaran B, Pezaris JS, Sahani M, Mitra PP, Andersen RA. 2002. Temporal structure in neuronal activity during working memory in macaque parietal cortex. *Nat Neurosci*. 5:805–811.
- Posner MI. 1980. Orienting of attention. *Q J Exp Psychol*. 32:3–25.
- Posner MI, Nissen MJ, Ogden WC. 1978. Attended and unattended processing modes: the role of set for spatial location. *Modes of perceiving and processing information*. Hillsdale, New Jersey: Lawrence Erlbaum Associates, Inc., pp. 137–158.

- Posner MI, Petersen SE. 1990. The attention system of the human brain. *Annu Rev Neurosci.* 13:25–42.
- Ray S. 2015. Challenges in the quantification and interpretation of spike-LFP relationships. *Curr Opin Neurobiol*, SI: Brain rhythms and dynamic coordination. 31:111–118.
- Ray S, Crone NE, Niebur E, Franaszczuk PJ, Hsiao SS. 2008. Neural correlates of high-gamma oscillations (60–200 Hz) in macaque local field potentials and their potential implications in electrocorticography. *J Neurosci.* 28:11526–11536.
- Ray S, Maunsell JHR. 2011. Different origins of gamma rhythm and high-gamma activity in macaque visual cortex. *PLoS Biol.* 9:e1000610.
- Ray S, Maunsell JHR. 2015. Do gamma oscillations play a role in cerebral cortex? *Trends Cogn Sci (Regul Ed)*. 19:78–85.
- Ray S, Ni AM, Maunsell JHR. 2013. Strength of gamma rhythm depends on normalization. *PLoS Biol.* 11:e1001477.
- Reynolds JH, Heeger DJ. 2009. The normalization model of attention. *Neuron.* 61:168–185.
- Shirhatti V, Borthakur A, Ray S. 2016. Effect of reference scheme on power and phase of the local field potential. *Neural Comput.* 28:882–913.
- Snyder AC, Issar D, Smith MA. 2018a. What does scalp electroencephalogram coherence tell us about long-range cortical networks? *Eur J Neurosci.* 48:2466–2481.
- Snyder AC, Morais MJ, Willis CM, Smith MA. 2015. Global network influences on local functional connectivity. *Nat Neurosci.* 18:736–743.
- Snyder AC, Yu BM, Smith MA. 2018b. Distinct population codes for attention in the absence and presence of visual stimulation. *Nat Commun.* 9:4382.
- Srinath R, Ray S. 2014. Effect of amplitude correlations on coherence in the local field potential. *J Neurophysiol.* 112:741–751.
- Subramanian J, Simon R. 2013. Overfitting in prediction models – is it a problem only in high dimensions? *Contemp Clin Trials.* 36:636–641.
- Thut G, Nietzel A, Brandt SA, Pascual-Leone A. 2006. α -Band electroencephalographic activity over occipital cortex indexes visuospatial attention bias and predicts visual target detection. *J Neurosci.* 26:9494–9502.
- Tremblay S, Doucet G, Pieper F, Sachs A, Martinez-Trujillo J. 2015. Single-trial decoding of visual attention from local field potentials in the primate lateral prefrontal cortex is frequency-dependent. *J Neurosci.* 35:9038–9049.
- Treue S, Maunsell JH. 1996. Attentional modulation of visual motion processing in cortical areas MT and MST. *Nature.* 382:539–541.
- Verhoef B-E, Maunsell JHR. 2017. Attention-related changes in correlated neuronal activity arise from normalization mechanisms. *Nat Neurosci.* 20:969–977.
- Vinck M, Battaglia FP, Womelsdorf T, Pennartz C. 2012. Improved measures of phase-coupling between spikes and the local field potential. *J Comput Neurosci.* 33:53–75.
- Vinck M, van Wingerden M, Womelsdorf T, Fries P, Pennartz CMA. 2010. The pairwise phase consistency: a bias-free measure of rhythmic neuronal synchronization. *Neuroimage.* 51:112–122.
- Womelsdorf T, Fries P, Mitra PP, Desimone R. 2006. Gamma-band synchronization in visual cortex predicts speed of change detection. *Nature.* 439:733–736.
- Xing D, Yeh CI, Shapley RM. 2009. Spatial spread of the local field potential and its laminar variation in visual cortex. *J Neurosci.* 29:11540–11549.
- Yates JL, Katz LN, Levi AJ, Pillow JW, Huk AC. 2020. A simple linear readout of MT supports motion direction-discrimination performance. *J Neurophysiol.* 123:682–694.

Supplementary figures



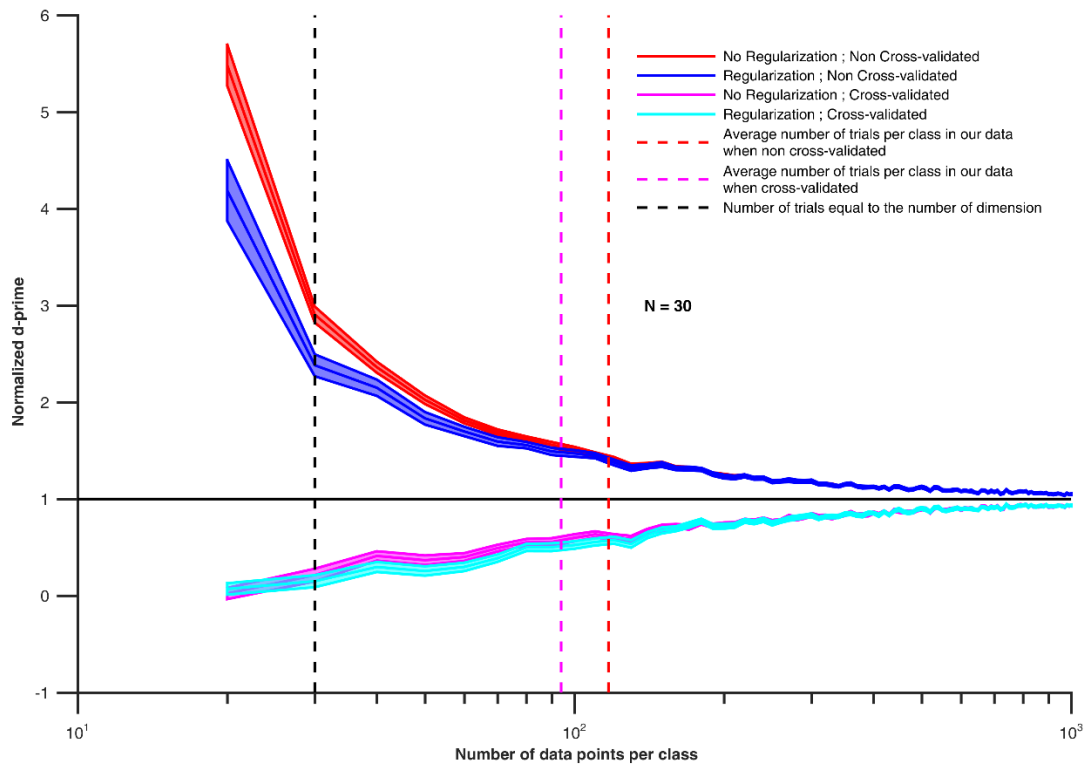
Supplementary figure 1: Inter-hemispheric electrode pairs phase relationship

(A) Mean field-field pairwise phase consistency (PPC) averaged across 5725 electrode pairs of 25 sessions of the attend-right/left (average of attend-left valid hit and attend-right valid hit condition; blue trace) and neutral (target-right neutral hit and target-in neutral hit combined; green trace) condition.

(B) Mean change in field-field PPC of the attend-left/right condition (blue) relative to the neutral condition.

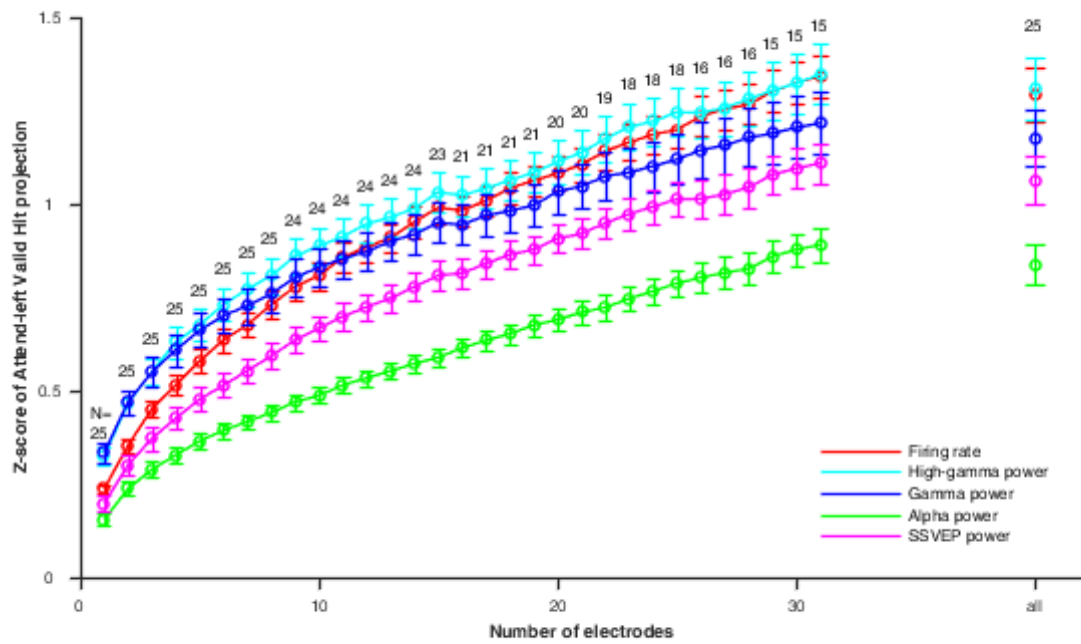
(C) Mean spike-field pairwise phase consistency (PPC) averaged across 11450 electrode pairs (including spike and LFP of the same electrode) for the same attention conditions as above.

(D) Mean change in spike-field PPC relative to the neutral condition.



Supplementary figure 2: Effect of the number of data points per class/category on the estimation of d-prime between two populations using LDA and explicit regularization of LDA, with no cross-validation and 5-fold cross-validation.

Estimated mean d-prime normalized relative to true d-prime between the projections of the populations as a function of the number of data points chosen per class to estimate the linear discriminant analysis (LDA) axis. The cases with no regularization with no cross-validation (red trace) and 5-fold cross-validation (magenta), regularized LDA with no cross-validation (blue) and 5-fold cross-validation (cyan) are compared. $N (=30)$ is the number of dimensions used, which is close to the average number of electrodes per session in our dataset. The black vertical dashed line indicates the case when the data points are equal to the number of dimensions. The magenta and red vertical dashed lines indicate the average number of data points in our dataset with and without cross-validation, respectively.



Supplementary figure 3: Discriminability as a function of the number of electrodes without cross-validation

Same as Figure 6 but without any cross-validation. Note that unlike the cross-validated case, values do not asymptote here, likely due to an increase in the degree of over-fitting for larger dimensionality.