Behavioral/Cognitive

Single-Trial Decoding of Visual Attention from Local Field Potentials in the Primate Lateral Prefrontal Cortex Is Frequency-Dependent

Sébastien Tremblay,[1,2] Guillaume Doucet,[1] Florian Pieper,[3] Adam Sachs,[4] and Julio Martinez-Trujillo[1,5]

[1]Cognitive Neurophysiology Laboratory, Department of Physiology, McGill University, Montreal, Quebec, QC H3G 1Y6, Canada, [2]Integrated Program in Neuroscience, Montreal Neurological Institute, McGill University, Montreal, Quebec, QC H3A 2B4, Canada, [3]Institute for Neuro- & Pathophysiology, University Medical Center Hamburg-Eppendorf, 20246 Hamburg, Germany, [4]Division of Neurosurgery, Department of Surgery, Ottawa Hospital Research Institute, University of Ottawa, Ottawa, Ontario, ON K1Y 4E9, Canada, and [5]Robarts Research Institute, Departments of Psychiatry, Physiology and Pharmacology, University of Western Ontario, London, Ontario, ON N6A 5C1, Canada

Local field potentials (LFPs) are fluctuations of extracellular voltage that may reflect the physiological phenomena occurring within a volume of neural tissue. It is known that the allocation of spatial attention modulates the amplitude of LFPs in visual areas of primates. An issue that remains poorly investigated is whether and how attention modulates LFPs in executive brain areas, such as the lateral prefrontal cortex (LPFC), thought to be involved in the origins of attention. We addressed this issue by recording LFPs from multielectrode arrays implanted in the LPFC of two macaques. We found that the allocation of attention can be reliably decoded on a single-trial basis from ensembles of LFPs with frequencies >60 Hz. Using LFP frequencies <60 Hz, we could not decode the allocation of attention, but we could decode the location of a visual stimulus as well as the endpoint of saccades toward that stimulus. The information contained in the high-frequency LFPs was fully redundant with the information contained in the spiking activity of single neurons recorded from the same electrodes. Moreover, the decoding of attention using γ frequency LFPs was less accurate than using spikes, but it was twice more stable across time. Finally, decorrelating the LFP signals from the different electrodes increased decoding performance in the high frequencies by up to ~14%. Our findings suggest that LFPs recorded from chronically implanted multielectrode arrays in the LPFC contain information about sensory, cognitive, and motor components of a task in a frequency-dependent manner.

Key words: **attention; decoding; local field potentials; multielectrode array; prefrontal cortex; primate**

Introduction

Local field potentials (LFPs) are slow-voltage fluctuations in the extracellular space of neural tissue that are thought to mostly represent the aggregate of postsynaptic potentials averaged over a few hundred micrometers volume around a recording electrode (Buzsáki et al., 2012; Einevoll et al., 2013). Certain cognitive operations are known to modulate LFP oscillations. For example, allocating visual attention into the receptive field of neurons in the vicinity of a recording electrode has been associated with increases in local γ power, γ spike-field synchronization, and between-area field-field synchrony in the β and γ bands (Fries et al., 2001, 2008; Womelsdorf et al., 2006; Buschman and Miller, 2007; Gregoriou et al., 2009, 2012; Khayat et al., 2010; Bosman et al., 2012; Vinck et al., 2013). One issue that remains unclear is how much information LFP signals can provide regarding the allocation of visual attention and how this information compares with the one provided by the spiking activity of neurons in a certain brain region.

We have recently reported that the spiking activity of neuronal ensembles in lateral prefrontal cortex (LPFC) area 8A contains sufficient information to predict the allocation of covert attention on a single-trial basis (Tremblay et al., 2015). Supporting this finding, several studies have demonstrated that the LPFC is involved in target selection and attention (Petrides and Pandya, 1999; Everling et al., 2002; Lebedev et al., 2004; Petrides, 2005; Lennert and Martinez-Trujillo, 2011, 2013). Thus, we proposed that area 8A contains a neuronal map of the behavioral relevance of stimuli, and that this map can provide top-down signals to modulate the activity of visual neurons in upstream areas during the allocation of attention.

We further speculated that the neuronal ensemble spiking activity recorded from chronic multielectrode arrays (MEAs) could be instrumental in the implementation of cognitive neural

Received March 17, 2015; revised April 21, 2015; accepted May 5, 2015.

Author contributions: F.P., A.S., and J.M.-T. designed research; F.P., A.S., and J.M.-T. performed research; S.T. and G.D. analyzed data; S.T., G.D., and J.M.-T. wrote the paper.

This work was supported by Canadian Institutes of Health Research, Natural Sciences and Engineering Research Council, EJLB Foundation, and Canada Research Chair program grants to J.M.-T. The Natural Sciences and Engineering Research Council Alexander Graham Bell Canada Graduate Scholarship supported S.T. We thank Mr. Walter Kucharski, Mr. Stephen Nuara, and Mr. Rishi Rajalingham for technical support; and members of the J.M.-T. laboratory for useful comments on previous versions of this manuscript.

The authors declare no competing financial interests.

Correspondence should be addressed to Dr. Julio Martinez-Trujillo, Robarts Research Institute, Departments of Psychiatry, Physiology and Pharmacology, University of Western Ontario, London, Ontario N6A 5C1, Canada. E-mail: julio.martinez@robarts.ca.

DOI:10.1523/JNEUROSCI.1041-15.2015

Copyright © 2015 the authors 0270-6474/15/359038-12$15.00/0

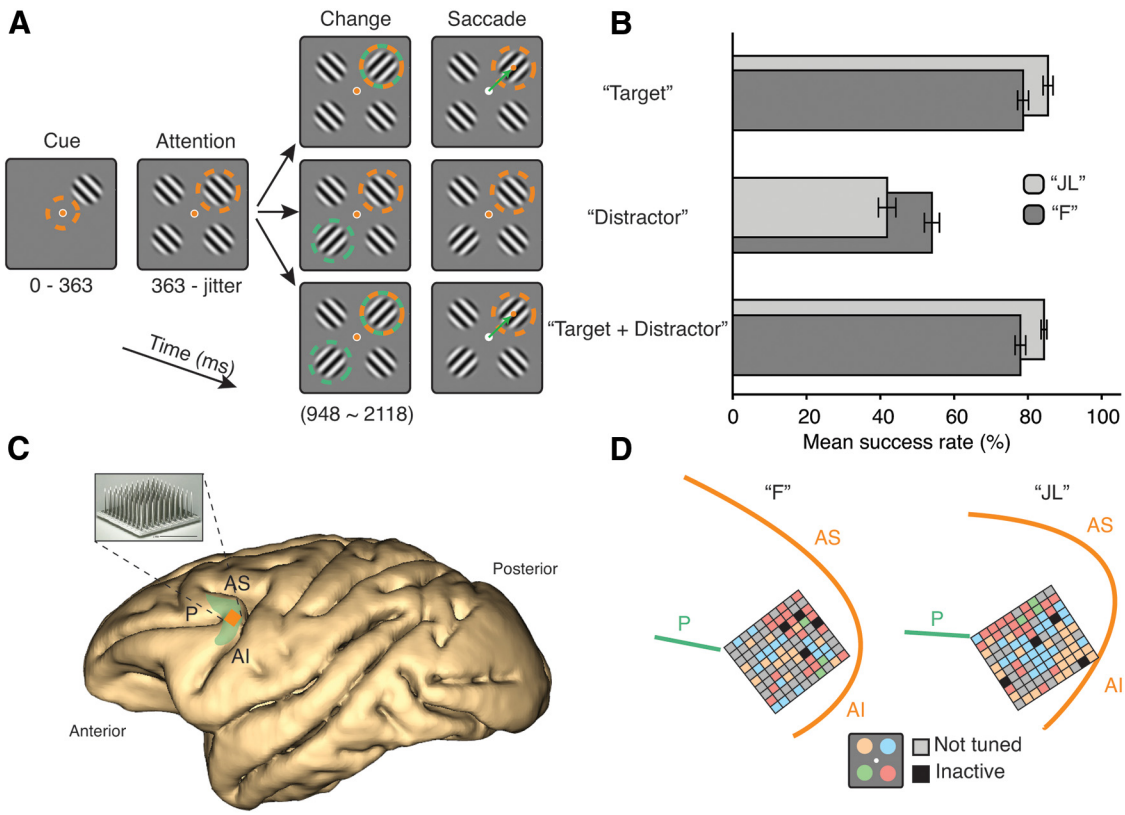

Figure 1. Task performance and MEA recording sites. ***A***, Monkeys were required to saccade to the stimulus that changed orientation if, and only if, it was previously cued. Orange circle represents focus of attention. Green circle represents orientation change. Orange dot indicates gaze position. Green arrow indicates saccade direction. Colored elements were not displayed during task. ***B***, Behavioral performance of Monkey JL and Monkey F on the three trial types. Error bars indicate SEM. ***C***, Location of chronic implant on cortex within the left dorsolateral prefrontal area 8A of a macaque brain. Orange square represents the position of the MEA. Green shaded area represents the location of macaque area 8A. P, Principal sulcus; AS, superior arcuate sulcus; AI, inferior arcuate sulcus. ***D***, Precise location of implants according to intraoperative photography. Each small square represents an electrode site. Color codes for the attentional tuning of multiunit clusters recorded at each electrode site.

prosthetics able to decode the behavioral relevance of objects across the visual field (Andersen et al., 2010). However, one potential problem with the use of spike signals recorded from chronic MEAs is that the isolation of single neurons deteriorates over short time periods (Dickey et al., 2009; Chestek et al., 2011; Perge et al., 2013), leaving LFPs as the only detectable signal. Thus, it is important to determine whether LFP signals can convey similar information as the spikes, and whether this information is sufficient for decoding cognitive signals within a behaviorally relevant timeframe.

We investigated this issue by recording LFPs from area 8A of two monkeys using chronically implanted MEAs while the animals performed a visual attention task. We examined LFP decoding accuracy across multiple frequency bands. We found that the allocation of attention could be decoded from LFPs in the 60–250 Hz range, but not from frequencies <60 Hz. High-frequency LFPs and spikes contained redundant information, similar attentional codes, and reacted similarly to distractors. Finally, decoding was lower using LFPs compared with spikes but was twice as stable over a 30 d time period and substantially improved when decorrelating the simultaneously recorded signals.

Materials and Methods

Two male macaque monkeys (*Macaca fascicularis*) were trained to perform a visual attention task (Fig. 1). The animals were instructed to covertly attend to one of four identical Gabor stimuli, to detect a subtle change in its orientation, and to saccade to it within 400 ms of the change to obtain a juice reward (Fig. 1*A*). The target Gabor appeared 363 ms earlier (cue period) than the other three distractors. After a variable delay period, orientation changes could happen in either the target ("Target" trials) or in the opposite distractor ("Distractor" trials). In the latter case, the monkey had to ignore the changing distractor and maintain fixation on the center dot until the end of the trial. In a third trial type, orientation changes occurred simultaneously in the target and in the opposite distractor ("Target + Distractor" trials), and the monkey had to saccade to the target change and ignore the distractor change. Within a given session, all three trial types were randomly interleaved, so that it was impossible for the animals to: (1) predict the location of the target, (2) know whether or not a saccade would be required, and (3) know whether the change would happen in the target, in a distractor stimulus, or in both.

Both monkeys performed above chance in all trial types: "Target," "Target + Distractor" (both ~80% hit rate, including fixation breaks), and "Distractor" (~60%; Fig. 1*B*). The performance decrease in the "Distractor" condition is explained by the animals failing to inhibit their saccade and responding to the distractor change in a number of trials. Importantly, this was not due to the animals ignoring the cue and saccading to a change in any stimulus. If this were the case, the performance in the "Target + Distractor" trials would have been close to 50% and the performance in the "Distractor" trials would have been close to 0%.

We recorded neural activity during the different trials using a 96-channel MEA (Blackrock Microsystems) chronically implanted in area 8A of the left prefrontal cortex (Fig. 1*C*,*D*). Area 8A is a cytoarchitectonic subdivision of the LPFC located rostral to the frontal eye fields (FEFs) on the prearcuate convexity (Petrides and Pandya, 1999; Petrides, 2005). Spikes were extracted using standard thresholding and spike sorting techniques. LFPs were obtained by low-pass filtering the broadband signal at 250 Hz sampled at 1 kHz. Instantaneous power in 400 ms windows was obtained using complex wavelet convolution for each center fre-

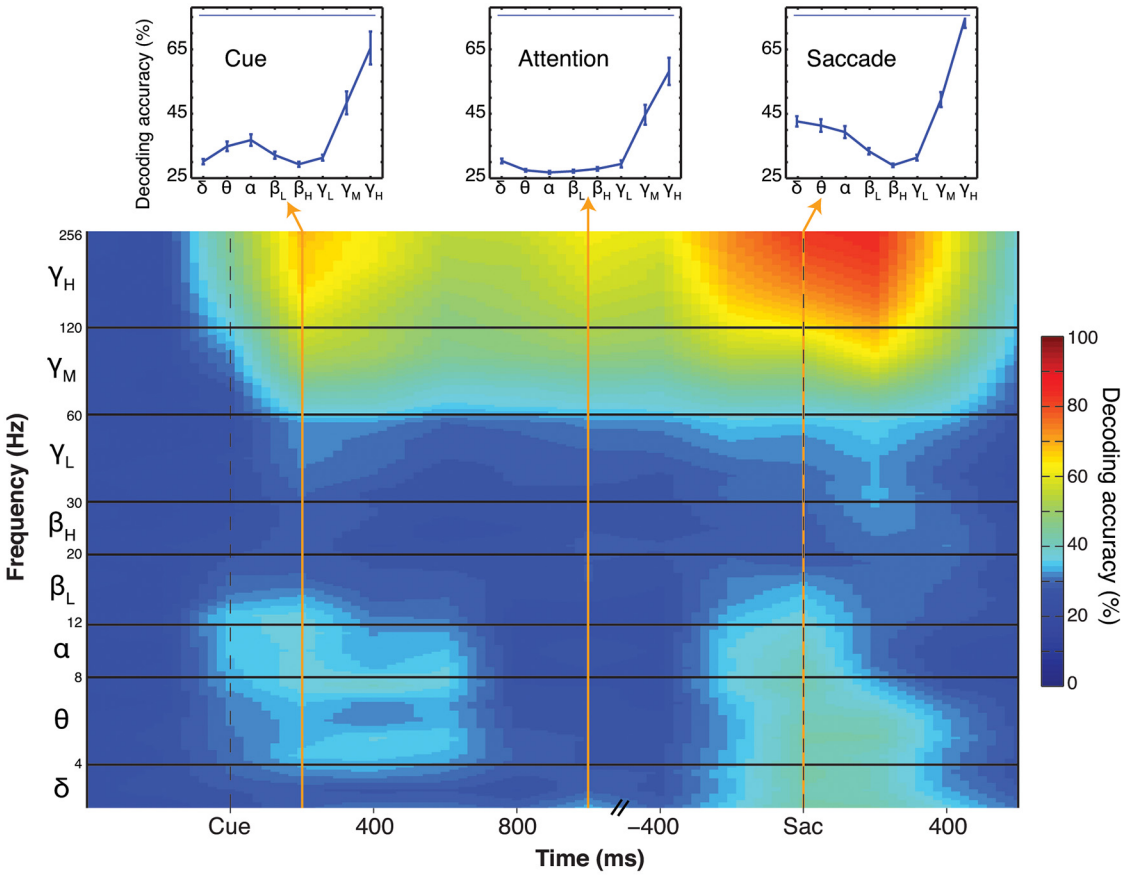

Figure 2. Decoding of LFP. This figure represents the single-trial decoding accuracy as a function of task epoch for each LFP frequency band. Only correct trials from "Target" trials were included. Cutoff frequencies for each band are indicated on the *y*-axis. The results were aligned to cue onset or saccade onset, and merged along the *x*-axis between the 1000 and −400 time points. Insets, Decoding accuracy as a function of LFP band for three representative 400 ms time bins within the cue, attention, and saccade epochs. Thin blue bars on top represent statistically higher than chance decoding ($p < 0.001$). Non-zero decoding accuracies just before cue onset are due to a smoothing artifact caused by the integration window (400 ms). The time axis is truncated because of the time jitter between cue and response event onsets.

quency (Liebe et al., 2012). Frequency bands were defined as follows: δ (0–4 Hz), θ (4–8 Hz), α (8–12 Hz), low-β (β_L: 12–20 Hz), high-β (β_H: 20–30 Hz), low-γ (γ_L: 30–60 Hz), mid-γ (γ_M: 60–120 Hz), and high-γ (γ_H: 120–250 Hz).

The instantaneous power of each frequency band was decoded using a support-vector machine (Chang and Lin, 2011), which has been shown to reliably decode attentional signals from prefrontal cortex neurons (Astrand et al., 2014; Tremblay et al., 2015). The decoder was trained to decode both the spike and LFP signals at various epochs over the course of each trial. Depending on the epoch, the decoded information can be related to: (1) the cue position (cue epoch), (2) the attended location (attention epoch), or (3) the saccade direction (saccade epoch). The decoder was trained iteratively on four-fifths of available trials with the true label information provided. Next, the decoder's single-trial predictions were tested on an independent testing set (i.e., the remaining one-fifth of the trials) for which the true label information was not provided. This cross-validation procedure was repeated 5 times so that every trial was part of the testing set once (K-fold = 5). The chance performance of the decoder was determined by shuffling the trial labels of the training set before the cross-validation procedure.

Results

Decoding visual, attentional, and saccadic information from LFP

To decode trial information, we separated the LFP signal in different frequency bands. Figure 2 shows the decoding accuracy as a function of task epoch and frequency band. In the high-γ band (γ_H: 120–250 Hz), decoding accuracy reached 65% for the cue epoch (cue onset to 400 ms after cue onset), 58% for the attentional epoch (800–1200 ms after cue onset), and 75% for the saccade epoch (−200 to +200 ms after saccade onset). In the mid-γ band (γ_M: 60–120 Hz), cue decoding reached 48%, attention decoding 44%, and saccade decoding 50%. For frequencies <60 Hz, decoding accuracies sharply dropped near chance level (25%) for all task epochs. The decoding of cue and saccade direction recovered slightly for α (α: 8–12 Hz) and theta bands (θ: 4–8 Hz). The δ band (δ: 0–4 Hz) also displayed significant decoding of the saccade epoch but did not yield good predictions of cue position. In general, there was little attentional decoding for any frequency band lower than mid-γ (<60 Hz). Thus, decoding of attention was selective for frequencies >60 Hz.

Effects of spike leakage on decoding accuracy

High-frequency LFPs can be contaminated by frequency components of spikes recorded from the same electrode, specifically in the frequencies >60 Hz (Zanos et al., 2011). Because spikes recorded in area 8A contain sufficient information to encode the allocation of attention on a single-trial basis (Tremblay et al., 2015), it is possible that spike leakage could by itself explain the high decoding accuracy observed in the attention period for the mid- and high-γ bands (Fig. 2). To investigate this issue, we removed spike components from the recorded LFPs using the method proposed by Zanos et al. (2011). This method computes the spike-triggered average LFP traces and subtracts the signal-

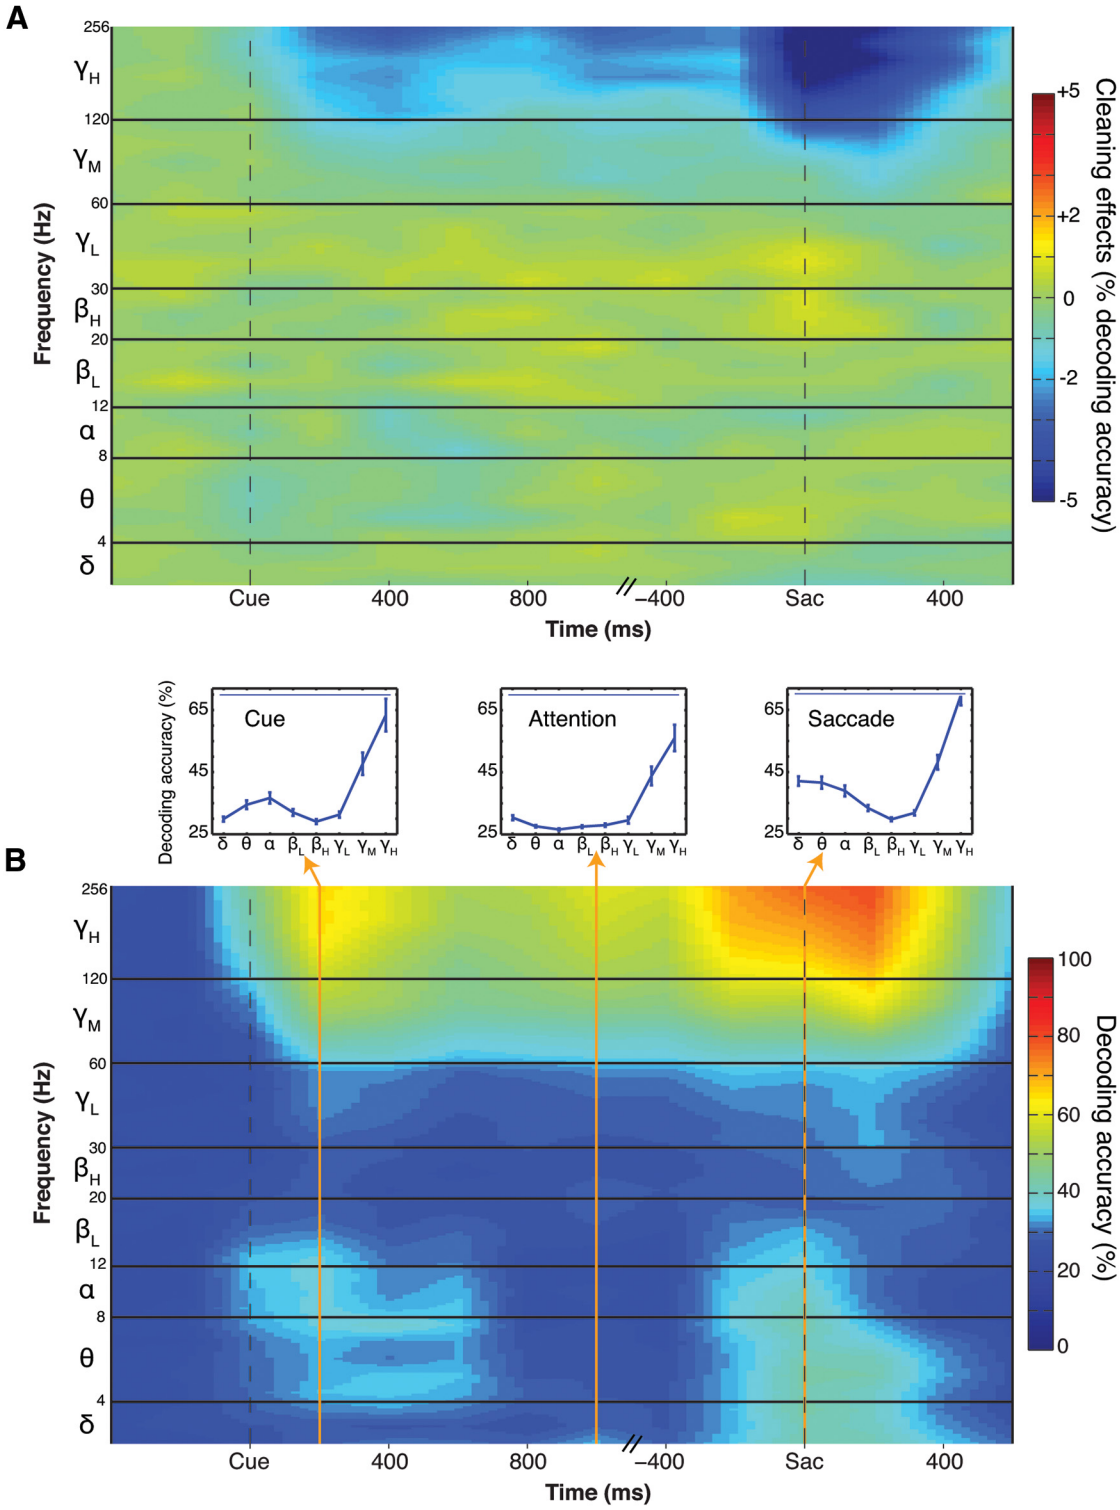

Figure 3. The effects of spike cleaning on decoding from LFP. ***A***, Absolute change in decoding accuracy due to the spike cleaning procedure. Display is as in Figure 2. Color codes for either absolute increase or decrease in decoding accuracy compared with using a noncleaned LFP signal. Most differences are observed in the mid- and high-γ bands. ***B***, Decoding accuracy after spike cleaning. Display is as in Figure 2. Decoding of cue, attention, and saccade is still possible despite spike cleaning of LFP.

to-noise corrected average spike waveform from each point in the LFP signal at which a spike occurred. The despiked LFP signal is thus free from frequency components of the spikes isolated on the corresponding recording electrode. However, spikes originating from neurons isolated on distant electrodes might still contribute, although modestly (Lindén et al., 2011), to the potential fluctuations recorded from one electrode (our electrodes were spaced by 400 μm).

We conducted the decoding analysis on the spike-cleaned LFP dataset and contrasted it with the previous results. Spike cleaning reduced the decoding accuracy in the mid- and high-γ band by a maximum of 5% during the saccade epoch and 2.5% during the

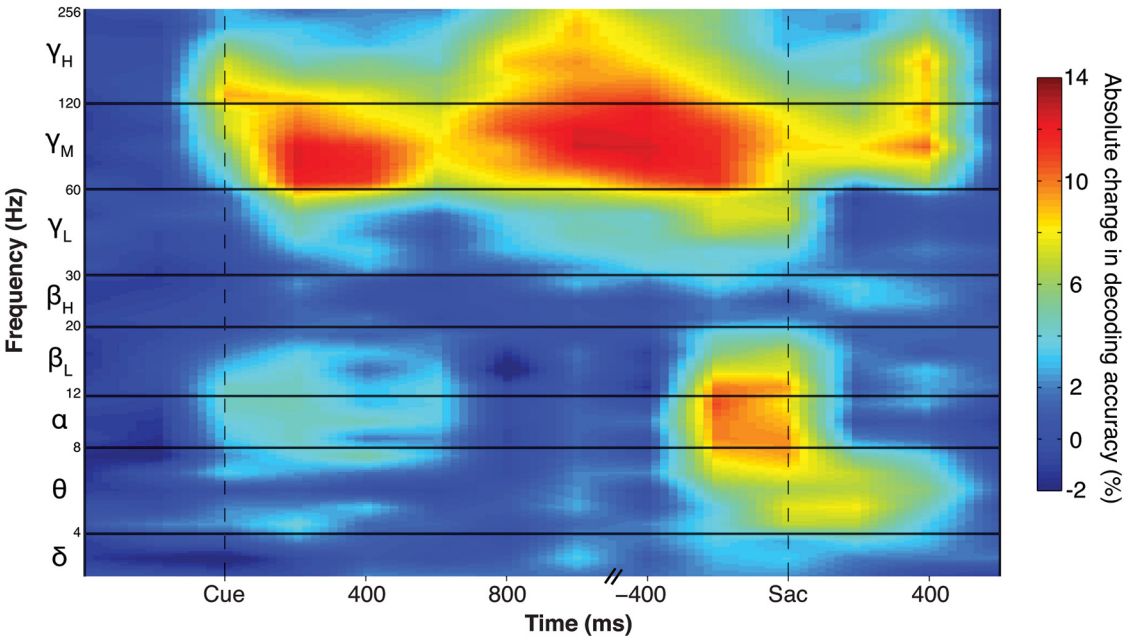

Figure 4. The effects of removing shared trial-to-trial variability on decoding accuracy. Absolute change in decoding accuracy, in percentage points, from removing shared variability (i.e., noise correlations) between electrodes. Hot colors code for increases in decoding accuracy after shuffling trials. The *x*-axis and *y*-axis are as in Figure 2.

cue and attention epochs (Fig. 3*A*). There was no effect of spike cleaning in frequency bands <60 Hz, which is in agreement with spike leakages mainly affecting high-frequency LFPs (Zanos et al., 2011). Remarkably, despite the removal of spike remnants from LFP traces, cue (48%; 63%), attended (44%; 56%), and saccade (48%; 70%) position could still be decoded on a single-trial basis from both mid- and high-γ frequency bands, respectively (Fig. 3*B*).

Shared trial-to-trial variability, also termed noise correlations, is known to impact population coding (Averbeck et al., 2006; Cohen and Kohn, 2011) and to play an important role in the neural network activity underlying visual attention (Cohen and Maunsell, 2009; Mitchell et al., 2009; Ruff and Cohen, 2014). We previously found that decorrelating the spiking activity of single neurons by shuffling trial order led to improvements in the decoding accuracy of visual attention across the visual field (Tremblay et al., 2015). We investigated whether decorrelating the LFP signal would lead to similar beneficial effects.

We applied the same shuffling procedure to the LFP signal to destroy the temporal order of the trials and thus remove naturally occurring shared trial-to-trial variability. The procedure implies shuffling the trial order for each neuron independently, so that the responses from the different neurons in the ensemble would not be extracted from the same trial. This shuffling is performed only within subset of trials belonging to the same stimulus condition, so that the tuning of the neurons' responses is left untouched (Leavitt et al., 2013). Figure 4 displays the change in decoding accuracy for the different frequency bands relative to the unshuffled recorded data. Surprisingly, decorrelating the signals led to frequency- and epoch-specific improvements in decoding accuracy of up to 14% in the mid-γ band during the attention epoch, and of up to 10% in the α band during the saccade epoch. These improvements are much larger than what we have previously reported when destroying correlated spiking activity in the same area.

Redundancy of information between spikes and LFP

We found that both spikes and high-frequency LFPs encode information about the cue position, allocation of attention, and saccade endpoint. One question arising from this result is whether these two signals overlap in terms of their information content. To investigate this issue, we first contrasted the absolute decoding accuracy of spikes with the decoding accuracy of LFPs in each frequency band (Fig. 5*A*). For all LFP bands, spikes provided more accurate decoding of the cue position, allocation of attention, and saccade endpoint. This difference averaged 60% for frequency bands <60 Hz, and 25% for the mid- and high-γ bands (>60 Hz). Thus, in all cases, the spiking signal contained more information than the LFPs. We also compared the error patterns of the decoder when using the spike or the LFP signal to determine whether or not they would differ. In Figure 5*B*, we can see the confusion matrices for both signals and for all three epochs. We found that the patterns of error were qualitatively very similar between the two signals.

Next, we explored how the information in the LFPs from each recording site relates to the information carried by the spiking activity of single neurons recorded from the same site. To do so, we correlated the decoding accuracy of the LFPs recorded on every single electrode against the decoding accuracy of the spiking activity of units recorded on the corresponding electrode. We hypothesized that, if spikes from units around an electrode carry a large amount of information about the task, the LFP recorded from the same electrode should also be very informative. Figure 6 illustrates the correlation coefficient for each LFP frequency band and for the three task epochs (cue, attention, and saccade). We found particularly strong unit/LFP correlations for the mid- and high-γ frequency bands ($r > 0.50$, $p < 0.001$) across all epochs. Correlations of up to $r = 0.70$ indicate a strong relationship between the information content of the spiking signal and that of the high-frequency LFP recorded from the same electrode. This correlation was much weaker for frequency bands <60 Hz; that

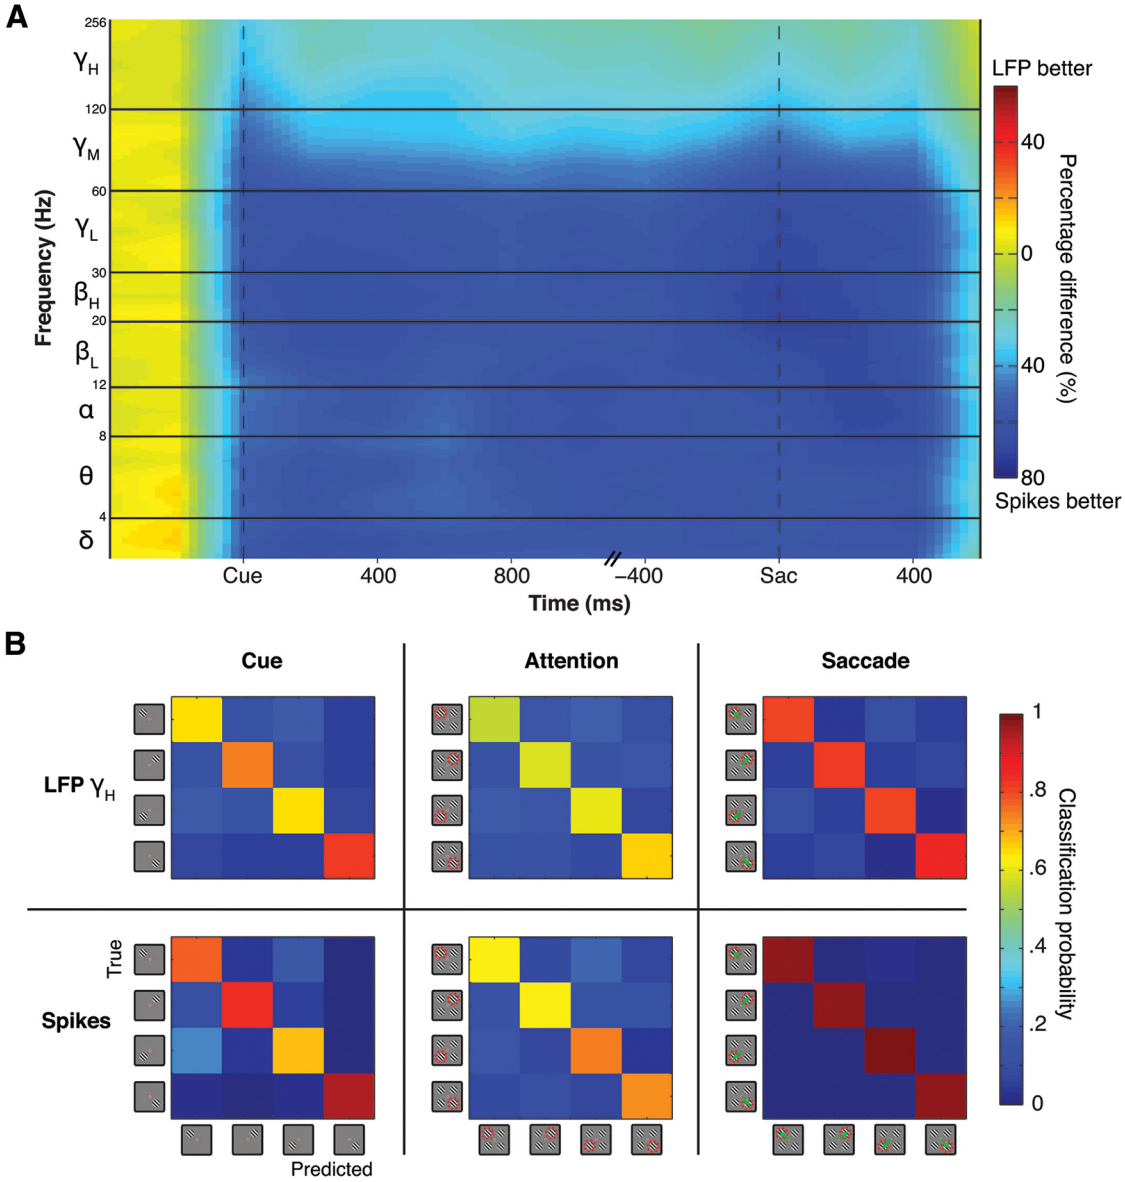

Figure 5. Spikes versus LFP decoding. ***A***, Contrast between the decoding accuracy of spikes and the decoding accuracy of all LFP frequencies. Display is as in Figure 2. Color codes for the relative percentage difference between the LFP and spike decoding accuracy for each epoch. Percentages are relative to spike decoding accuracy: (LFP − spikes)/spikes. For example, deep blue represents that the spike signal is 1.8 times more accurate than the LFP signal. ***B***, Confusion matrices of error patterns for spikes and for the best LFP band (γ_H), separated by epoch. Color codes for the probability that a given label will be predicted (columns) by the decoder given a true label (rows). Diagonal terms are correct classifications. Off-diagonal terms are misclassifications.

is, the maximum correlation coefficient obtained <60 Hz was $r = 0.40$ for the low-γ band (30–60 Hz).

We further compared the information content of the ensemble spiking activity with the information content of the ensemble LFPs to assess whether these two signals contain redundant, or complementary, information about the task. Decoding techniques are well suited for such a comparison between different signals because both are projected onto a common scale (i.e., decoding accuracy) (Quian Quiroga and Panzeri, 2009). In this analysis, we first entered the information contained in the spiking signal and obtained a baseline decoding accuracy. In a second step, we added the information contained in specific LFP bands and repeated the decoding analysis. Because some LFP bands contain information about the task (see Fig. 3*B*), this analysis could result in two possible outcomes. One outcome is that the decoding accuracy improves with the added information from LFPs. In this case, it would mean that some of the information in the LFPs is complementary to that of spikes. Alternatively, the decoding accuracy might not improve, which would imply that the information contained in the LFPs is redundant with that of the spikes.

We conducted this analysis in two different ways: adding LFP information on top of spikes, as explained above, or adding spiking information on top of LFPs. Figure 7 illustrates the decoding improvements for each added signal (blue represents LFP; red represents spikes), as a function of frequency band and task epochs. We found that, across all bands, adding the spiking information on top of the LFPs significantly improved the decoding accuracy, meaning that spikes contain novel information not contained in the LFPs. This improvement mainly reflects the higher absolute decoding accuracy of the spikes compared with the LFPs, as illustrated in Figure 5. On the other hand, adding LFP information on top of spike information did not improve the decoding accuracy, regardless of frequency band. Indeed, in some

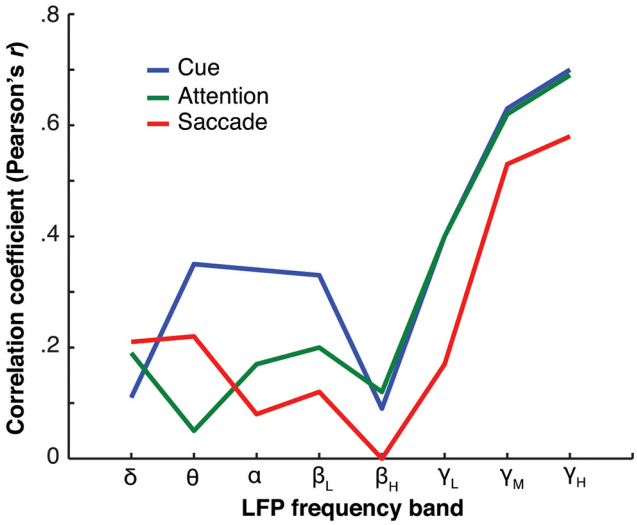

Figure 6. Spike–LFP correlations for each frequency band. The decoding accuracy of single LFP channels is correlated with the decoding accuracy of the single units recorded on the same channel. The Pearson's correlation coefficient is plotted on the *y*-axis. The *x*-axis represents all eight frequency bands. Each colored line indicates a different task epoch (cue, attention, or saccade).

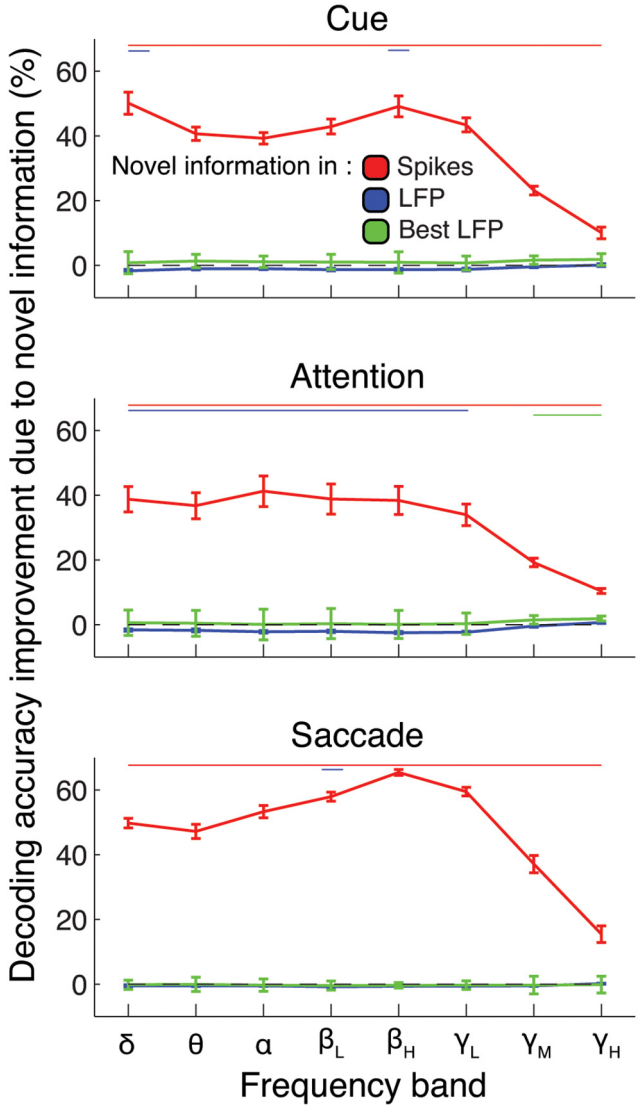

Figure 7. Redundancy of information in spikes and LFPs. These plots represent the decoding accuracy improvements from adding the information from one signal to the information contained in the other signal (e.g., adding the LFPs to the spikes vs the spikes alone; blue curve). Improvements mean that nonredundant information is contained in the added signal. Results are displayed for all frequency bands and for the three task epochs. Thin lines on top indicate statistically significant difference for the corresponding colored curve ($p < 0.001$). Error bars indicate SEM.

frequency bands, the decoding accuracy became worse when adding LFPs to the spikes. This is likely due to the vulnerability of the decoder to large amounts of noisy predictors. To account for this limitation, we ran a control analysis, including only the 10 best LFP channels on top of the spike signal (Fig. 7, green line). We found that the results mirrored those including all LFP channels, apart from a small benefit for mid- and high-γ signals during the attention epoch (<2%). This means that the task-related information contained in the power of the LFPs is redundant with that of the spiking activity.

Another way to investigate whether LFPs and spikes contain redundant information is by examining whether the LFPs and spiking codes are similar to one another. To assess this question, we trained the decoder on the normalized spiking signal and made single-trial predictions about task-related events using the normalized LFP signal. In Figure 8, this cross-signal decoding is contrasted with a decoder trained and tested using the same LFP signal. We found that, for high frequencies (i.e., low, mid, and high γ), it was possible to decode the cue position, allocation of attention, and saccade direction from LFP signals using a code obtained from spikes. The decoding accuracy for these bands was comparable with the one obtained by training and testing on the same LFP signals, meaning that the spiking code and the high-frequency LFP codes are very similar from a decoding perspective. We also found that above-chance decoding was possible in the lower-frequency bands, although limited to the cue and saccade epochs and with a performance much lower than the one obtained by training on the LFP signal. Thus, the similarity between LFP and spike codes varies as a function of the LFP frequency, higher frequencies being more similar to spikes than lower ones. Moreover, the fact that high-frequency LFPs can be well decoded when training with spikes further corroborates that these two signals contain similar information.

Effects of distractor interference on LFP

We have previously shown that a salient distractor transiently interferes with the ensemble spiking code of neurons in area 8A (Tremblay et al., 2015). When the distractor was successfully suppressed at a behavioral level, the code returned to a predistractor state a few hundred milliseconds after distractor offset. Here, we investigated whether the LFP ensemble code is as resilient to a salient visual distractor as the spiking code. To address this question, we obtained the baseline, predistractor, LFP ensemble code by training the decoder during the attention epoch of the task. This baseline attention code was then used to predict the allocation of attention before, during, and after the onset of the salient visual distractor. Any interference with the ensemble code caused by the distractor would be reflected as a decrease in the decoding accuracy. This procedure was applied separately for the mid- and high-γ bands, from which attention could be decoded (see Fig. 3*B*).

Figure 9*A* depicts the accuracy of decoders trained during the baseline predistractor epoch, and tested before, during and after distractor onset (Fig. 9*A*, left, black dashed rectangles). The line

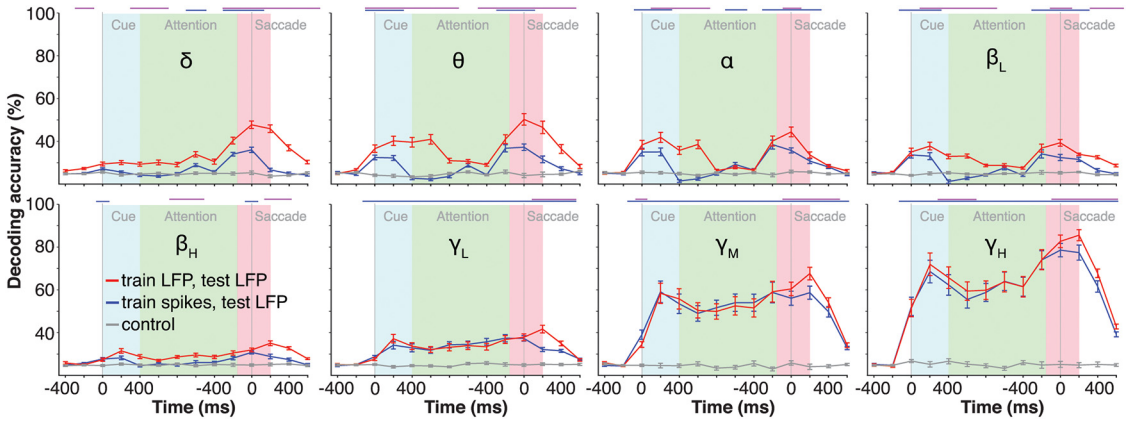

Figure 8. Similarity of ensemble codes between spikes and LFP. This figure illustrates the decoding accuracy as a function of time and frequency band for a decoder trained on spikes and tested on LFP (blue curve), and a decoder trained on LFPs and tested on LFPs (red curve). Decoding of cue, attention, and saccade direction from high-frequency LFP was possible when using a spiking code. Gray curves indicate chance decoding accuracy obtained from shuffling trial labels in the training set. Thin blue line on top indicates better decoding than chance when training on spikes ($p < 0.001$). Thin red line on top indicates better decoding when training on LFP compared with when training on spikes ($p < 0.001$). Error bars indicate SEM.

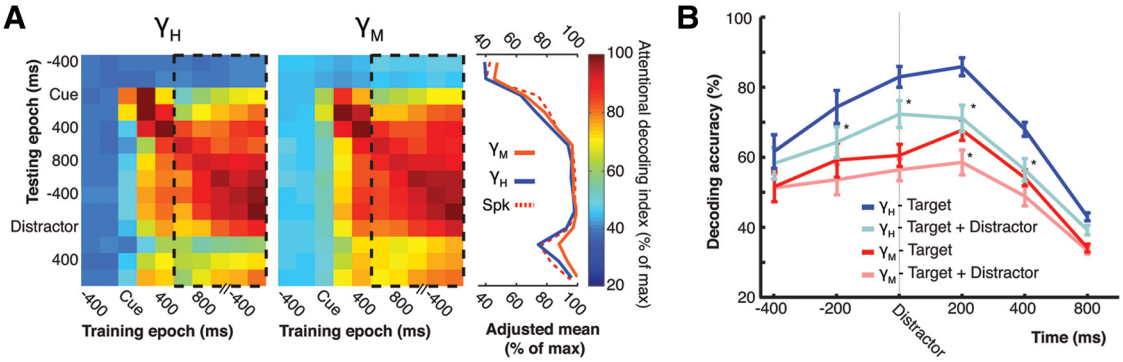

Figure 9. Distractor interference. ***A***, Interference caused by a salient distractor in correct "Distractor" trials. Left, Decoders were trained on every possible pair of time points, generating training by testing epoch matrices for each frequency band. Results are expressed as percentage of maximum decoding accuracy during the attention epoch. Dashed rectangles represent decoders trained during the attention epoch. Lines on the right were computed by averaging each row inside the black dashed rectangles. These lines are adjusted to take into account the normal accuracy decay due to the time interval between training and testing time points. The −400 ms time point on the *x*-axis is aligned to the saccade onset. ***B***, Effect of a salient distractor when contiguous with a change at the target location ("Target + Distractor" trials). Results for frequencies >60 are contrasted to results when no synchronous distractor is present ("Target" trials). *$p < 0.001$, statistically significant differences between curves belonging to the same frequency band. Error bars indicate SEM.

plot on the right displays the averaged results. We observed that the high-γ LFP band (blue line) shows a transient distractor-related decrease in decoding accuracy of ~25%. This interference is very similar to the one previously observed in the spiking code (Fig. 9*A*, red dashed line). In contrast, the interference of the mid-γ was only 10%. This may be explained by the fact that this frequency band initially contains less information about the allocation of attention, leaving less room for interference. In absolute values, the distractor-related decrease in decoding accuracy was of 3% and 10% in the mid- and high-γ bands, respectively. Overall, both LFP band codes were resilient to distracting information, as expressed by their rebound to predistractor levels. Moreover, the relative amplitude of the interference observed in the spiking and high-γ codes was almost identical.

We further explored the effects of a salient distractor when simultaneously presented with the target stimuli ("Target + Distractor" trials). At a behavioral level, this distractor did not severely impair task performance (mean difference between "Target" and "Target + Distractor" trial types = 0.9%, CI = −0.8% to 2.6%, paired *t* test, $p = 0.3$) as opposed to when the distractor was presented alone (i.e., "Distractor" trials, see Fig. 1*B*). This behavioral robustness was previously found to be mirrored in the spiking code of neuronal ensembles in area 8A (Tremblay et al., 2015). We wondered whether the high-frequency LFP signal would show a similar effect. We addressed this question by training the decoder on "Target" trials and testing it on "Target + Distractor" trials. If the distractor interferes with the attention code of LFP ensemble activity, we should observe a decrease in the decoding accuracy relative to testing on "Target" trials.

The simultaneous distractor decreased decoding accuracy in the high-γ band, albeit not comparable with when the distractor was presented alone (Fig. 9*B*; maximum drop = 15% vs 25%). Similarly, the mid-γ band showed a slight decrease in decoding accuracy, smaller than the one observed in the high-γ band (5% vs 15%). Again, the smaller interference in the mid compared with the high-γ band could be explained by lower baseline decoding accuracy in the mid-γ band. Overall, these results mirror the behavioral observations: a distractor presented simultaneously with a target elicits a small and transient interference in both the spikes and the high-frequency LFP signal.

Decoding stability over time

The decoding stability over time is an important factor that could affect the reliability of a cognitive neural prosthetic using LFP signals from prefrontal cortex. As such, we investigated the sta-

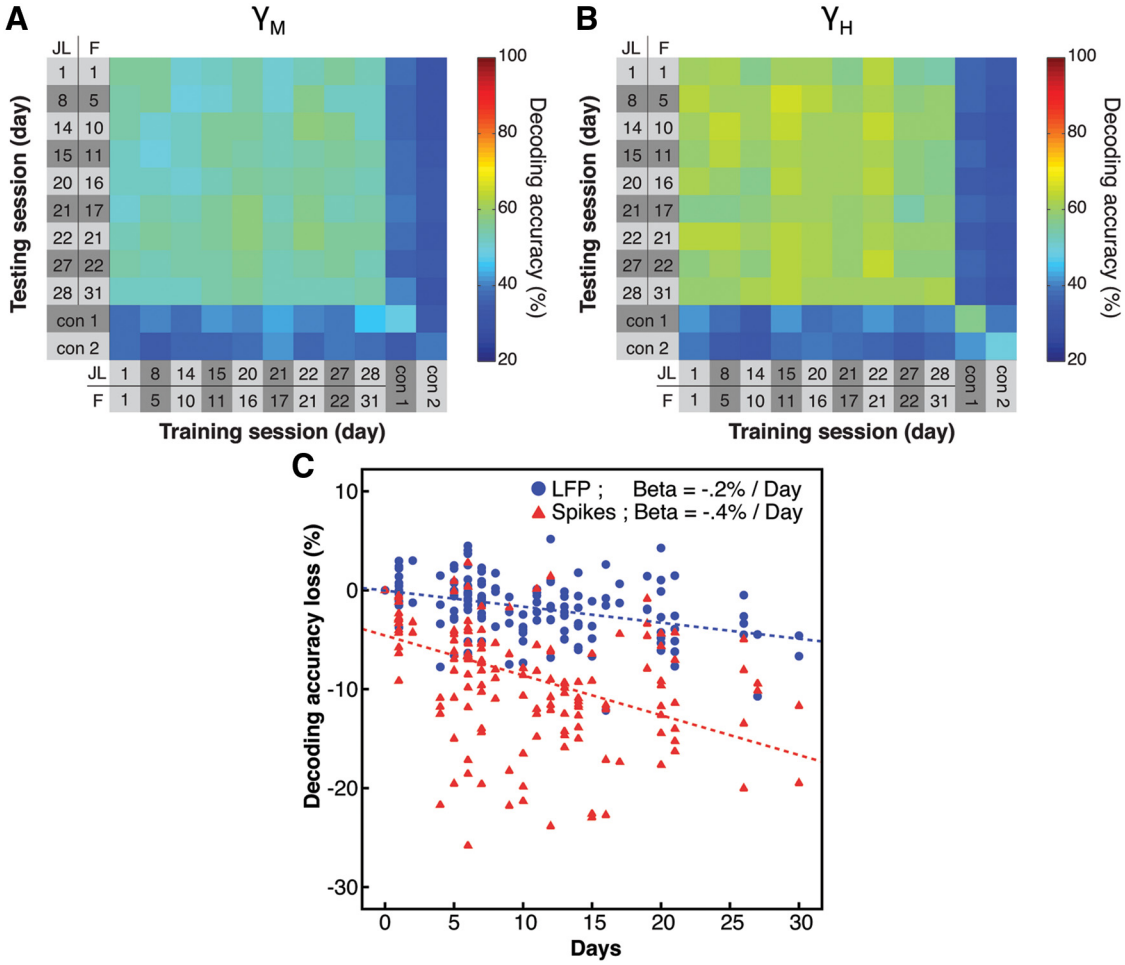

Figure 10. Decoding stability across time. ***A***, Stability of decoding across recording sessions for the mid-γ band LFP signal. The decoder was trained and tested on every pair of sessions within the attention epoch. Color represents decoding accuracy, pooled across monkeys. "con 1" and "con 2" represent control sessions that were recorded from a different set of nearby channels. ***B***, Same as in ***A*** for the high-γ frequency band. ***C***, Decay of the generalizability as a function of days. The *x*-axis represents the number of days separating the training session from the testing session. The *y*-axis represents the absolute decoding accuracy loss compared with training and testing on the same day. Blue circles and line represent the decay of high-γ LFP signals. Red triangles and line represent the decay of the spiking signal.

bility of the LFP ensemble code over a prolonged period of time. We reasoned that, if the code is stable over time, it should be possible, for example, to use a code obtained on the first day of recording to decode activity recorded a month later. We used the mid- and high-γ band activity during the attention epoch to train the decoder on every session, and to test it on every other session during the same epoch. This created a "training session" by "testing session" matrix of decoding accuracies (Fig. 10*A*).

We found that the decoding accuracy was very stable over time. Ensemble codes obtained on one session could successfully decode the allocation of attention in another session that occurred days or weeks later. This demonstrates that the LFP coding of attention is generalizable over time, which is in agreement with previous observations in the spiking data (Tremblay et al., 2015). Importantly, this generalization was limited to the local ensemble; the same code could not be used to decode LFP signals captured by nearby electrodes in the same brain region (Fig. 10*A*, con 1 and con 2 rows and columns).

Finally, we contrasted the stability of the high-γ band LFP signal with the stability of spike decoding. We computed the decoding accuracy loss as a function of days for both signals (Fig. 10*C*) and found that both signals decayed as a function of time, although at different rates (general linear model test on β coefficients; $t = 3.6$, $p < 0.001$). The stability of the LFP signal decayed at a rate of 0.2% per day. The decay of the spiking signal was twice as fast, at a rate of 0.4% per day. These results indicate that the coding of attention is more stable across days using the LFP signal compared with using the spiking signal.

Discussion

Decoding of attention from γ frequencies

Attention-dependent modulation of γ frequency LFPs has been reported across several brain areas (Fries et al., 2001; Chalk et al., 2010; Khayat et al., 2010; Ray and Maunsell, 2011). In the LPFC, most studies reported modulation of γ LFPs in the FEFs and synchronization with activity in visual area V4 (Buschman and Miller, 2007; Monosov et al., 2008; Gregoriou et al., 2009, 2012). In general, these studies suggest that the γ signal contains information about the allocation of attention.

In agreement with these studies, we showed that the allocation of attention across the four visual field quadrants could be decoded from LFP frequencies >60 Hz recorded from the LPFC. Our results complement a previous report of single-trial decoding of attention using the simultaneous spiking activity of neurons recorded from the same area (Tremblay et al., 2015). Together, these results support the hypotheses that area 8A con-

tains a neuronal map of the behavioral relevance of stimuli and is therefore a potential source of attentional signals that influence processing in the rest of the brain (Desimone and Duncan, 1995; Miller and Cohen, 2001).

One question arising from our results concerns the source of the attentional information contained in the mid- and high-γ frequencies. This information may be contained in the input signals arriving into area 8A and therefore be inherited from areas upstream, such as the lateral intraparietal area or the FEF (Monosov et al., 2008; Squire et al., 2013). On the other hand, it might simply reflect local spiking activity within area 8A. An insight into this question may be provided by a previous study in primary visual cortex that dissociated LFP activity in the low-γ range (30–80 Hz) from activity in the mid- and high-γ range (>80 Hz) (Ray and Maunsell, 2011). The authors concluded that LFP signals above a threshold of 50–80 Hz mainly reflect local spiking activity. Numerous supporting observations have been reported in the primary visual cortex (V1) (Belitski et al., 2008; Rasch et al., 2008), secondary somatosensory cortex (S2) (Ray et al., 2008), and posterior parietal regions (Pesaran et al., 2002; Hwang and Andersen, 2011). All these studies support a close relationship between the high-γ activity and spiking activity of cortical neurons. Here we showed that this relationship also holds true within the caudal prefrontal cortex.

Supporting this proposal, we found a strong correlation between the decoding accuracy using LFPs >60 Hz and using neuronal activity recorded from the same electrode (Fig. 6). This finding is also congruent with observations of tuning similarity between single-unit and γ frequency LFPs recorded in areas, such as middle temporal (Liu and Newsome, 2006), V1 (Xing et al., 2009), and the inferior temporal cortex (Kreiman et al., 2006). Moreover, we found that a decoder trained using spike rates can reliably decode the allocation of attention using the information contained in the γ LFPs. This decoding generalizability may be attributed to at least two factors: (1) the similarity between the tuning of multiunit clusters and γ LFPs in single electrodes; and (2) the similarity in the spatial distribution of tuned signals in the cortical volume covered by the MEA (Einevoll et al., 2013).

Finally, we compared the information content of spikes and LFPs ensemble activity. The information about the allocation of attention contained in the γ LFPs was found to be fully redundant with the one contained in the spiking activity of single units and multiunits. This result applied to all frequency bands, although it can only be interpreted correctly in bands that contain attentional information in the first place (i.e., mid- and high-γ bands). Similar observations have been previously reported in the primate motor and premotor cortices (Bansal et al., 2012). These results further support the idea that multiunit and high-γ activity are of similar nature (Ray et al., 2008; Ray and Maunsell, 2011; Einevoll et al., 2013) and start to raise questions about the functional role of these high-frequency oscillations in cortical processing (Ray and Maunsell, 2014).

Decoding from low-frequency LFPs

Evidence for the attention-dependent modulation of non-γ LFPs in LPFC is thinner. Some modulations are observed in the β band, although they seem to primarily affect movement cells in the FEF (Gregoriou et al., 2012), which may not be implicated in the attentional modulation of visual processing (Thompson et al., 2005). Along those lines, we found that the low-frequency LFPs did not contain sufficient information to accurately predict the allocation of attention. This finding apparently contradicts previous reports of attention influencing low-frequency LFPs in area middle temporal (Khayat et al., 2010; Esghaei and Daliri, 2014). This apparent discrepancy may be due to several reasons, including differences in the microstructure of visual and prefrontal areas. For example, early visual areas do not have the same neuronal density, the same connectivity with the thalamus, or the same sizes of neurons as the prefrontal cortex (Barbas and Pandya, 1989). Moreover, retinotopic maps in middle temporal are well defined and are limited to the contralateral visual field (Born and Bradley, 2005). This is very distinct from LPFC, where receptive fields can span the entire visual field and retinotopy is not present (Rainer et al., 1998; Lennert and Martinez-Trujillo, 2013; Tremblay et al., 2015).

Although we could not decode the allocation of attention from the power of low-frequency LFPs, we could decode cue position and saccade endpoint. It may be that these frequencies encode transient sensory and motor events more strongly than sustained attention. Favoring this hypothesis, LFPs in visual area V4 best reflect stimulus tuning during time intervals close to the stimulus onset (Mineault et al., 2013). This issue, however, needs further investigation.

Implications for neural prosthetics

Neural prosthetics are promising therapeutic applications that seek to empower patients by offering them control over artificial limbs or objects in their environment (e.g., computer cursor) (Donoghue, 2008). High degree-of-freedom control signals for these prosthetics can be extracted directly from neural spiking activity using chronic intracranial implants over motor and premotor cortical areas (Hochberg et al., 2006). However, one problem with chronic MEA implants used in brain–machine interface applications is that the isolation of single units deteriorates over time (Dickey et al., 2009; Chestek et al., 2011; Perge et al., 2013). This is in part due to factors such as movement of the array within the cortical tissue, as well as gliosis around the electrodes due to the immune response of the brain to a foreign object (Biran et al., 2005).

Our results show that, even in the absence of spikes, one can decode the allocation of attention using high-frequency LFP signals. However, the decoding accuracy was lower than when using spikes. This may be caused by the LFPs' larger integration volume and noisier nature (Buzsáki et al., 2012). On the other hand, we found that LFPs were twice as stable as spikes across a 30 d interval. In a case where the decoder cannot be retrained on a regular basis, the decay of the decoding accuracy over days for spikes is such that LFPs would surpass spikes decoding after a few weeks of recording. Thus, the greater stability of the LFP signal in the LPFC might come as an important advantage over spikes for long-term cognitive brain–machine interface applications.

We also found that destroying correlated activity between the electrodes led to significant improvements in decoding accuracy, as was previously found with spikes (Tremblay et al., 2015). This improvement was frequency-specific, with the largest increase in the mid-γ band (60–120 Hz) and during the attention epoch. Our results may appear at odds with observations of a previous study reporting that decorrelating LFP signals did not lead to improvements in decoding movement intentions (Markowitz et al., 2011). However, we did not observe the greatest improvement in accuracy during the movement epoch of our task, and the small improvement we observed during this epoch might be best explained by the visual transient happening just before the saccade. Thus, it may be that destroying correlations mainly improves the visual-attentional task component rather than the motor/saccade component. This issue, however, requires further

testing. Importantly, our results could be instrumental in improving decoding accuracy of LFP-based neural prosthetics by decorrelating neural control signals before decoding.

Our study demonstrates that using LFP signals from a single implant in area 8A, it is possible to decode the allocation of attention across both visual hemifields. This is an advantage compared with other cortical areas that contain only a unilateral representation of the contralateral hemifield. Our results also demonstrate that the locus of attention can be decoded in advance of a saccade to the attended location. This may be instrumental to improve the performance of a neural motor prosthetics by restricting future movement goals to information that is behaviorally relevant, potentially shortening reaction times and improving movement accuracy. Moreover, area 8A is easily accessible for a surface MEAs (Maynard et al., 1997), which constitutes an important advantage compared with areas, such as FEF, which is located within a sulcus. Our study provides a proof of principle that LFP signals from this area can be used to decode the allocation of attention as well as movement intentions and therefore improve the performance of existing motor prosthetics, or promote the development of novel cognitive neural prosthetics (Donoghue, 2008; Andersen et al., 2010) to assist people with disabilities.

References

Andersen RA, Hwang EJ, Mulliken GH (2010) Cognitive neural prosthetics. Annu Rev Psychol 61:169–190. CrossRef Medline

Astrand E, Enel P, Ibos G, Dominey PF, Baraduc P, Ben Hamed S (2014) Comparison of classifiers for decoding sensory and cognitive information from prefrontal neuronal populations. PLoS One 9:e86314. CrossRef Medline

Averbeck BB, Latham PE, Pouget A (2006) Neural correlations, population coding and computation. Nat Rev Neurosci 7:358–366. CrossRef Medline

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. CrossRef Medline

Barbas H, Pandya DN (1989) Architecture and intrinsic connections of the prefrontal cortex in the rhesus monkey. J Comp Neurol 286:353–375. CrossRef Medline

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. CrossRef Medline

Biran R, Martin DC, Tresco PA (2005) Neuronal cell loss accompanies the brain tissue response to chronically implanted silicon microelectrode arrays. Exp Neurol 195:115–126. CrossRef Medline

Born RT, Bradley DC (2005) Structure and function of visual area MT. Annu Rev Neurosci 28:157–189. CrossRef Medline

Bosman CA, Schoffelen JM, Brunet N, Oostenveld R, Bastos AM, Womelsdorf T, Rubehn B, Stieglitz T, De Weerd P, Fries P (2012) Attentional stimulus selection through selective synchronization between monkey visual areas. Neuron 75:875–888. CrossRef Medline

Buschman TJ, Miller EK (2007) Top-down versus bottom-up control of attention in the prefrontal and posterior parietal cortices. Science 315: 1860–1862. CrossRef Medline

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. CrossRef Medline

Chalk M, Herrero JL, Gieselmann MA, Delicato LS, Gotthardt S, Thiele A (2010) Attention reduces stimulus-driven γ frequency oscillations and spike field coherence in V1. Neuron 66:114–125. CrossRef Medline

Chang CC, Lin CJ (2011) LIBSVM: a library for support vector machines. ACM Transactions on Intelligent Systems and Technology 2.

Chestek CA, Gilja V, Nuyujukian P, Foster JD, Fan JM, Kaufman MT, Churchland MM, Rivera-Alvidrez Z, Cunningham JP, Ryu SI, Shenoy KV (2011) Long-term stability of neural prosthetic control signals from silicon cortical arrays in rhesus macaque motor cortex. J Neural Eng 8:045005. CrossRef Medline

Cohen MR, Kohn A (2011) Measuring and interpreting neuronal correlations. Nat Neurosci 14:811–819. CrossRef Medline

Cohen MR, Maunsell JH (2009) Attention improves performance primarily by reducing interneuronal correlations. Nat Neurosci 12:1594–1600. CrossRef Medline

Desimone R, Duncan J (1995) Neural mechanisms of selective visual attention. Annu Rev Neurosci 18:193–222. CrossRef Medline

Dickey AS, Suminski A, Amit Y, Hatsopoulos NG (2009) Single-unit stability using chronically implanted multielectrode arrays. J Neurophysiol 102:1331–1339. CrossRef Medline

Donoghue JP (2008) Bridging the brain to the world: a perspective on neural interface systems. Neuron 60:511–521. CrossRef Medline

Einevoll GT, Kayser C, Logothetis NK, Panzeri S (2013) Modelling and analysis of local field potentials for studying the function of cortical circuits. Nat Rev Neurosci 14:770–785. CrossRef Medline

Esghaei M, Daliri MR (2014) Decoding of visual attention from LFP signals of macaque MT. PLoS One 9:e100381. CrossRef Medline

Everling S, Tinsley CJ, Gaffan D, Duncan J (2002) Filtering of neural signals by focused attention in the monkey prefrontal cortex. Nat Neurosci 5:671–676. CrossRef Medline

Fries P, Reynolds JH, Rorie AE, Desimone R (2001) Modulation of oscillatory neuronal synchronization by selective visual attention. Science 291: 1560–1563. CrossRef Medline

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. CrossRef Medline

Gregoriou GG, Gotts SJ, Zhou H, Desimone R (2009) High-frequency, long-range coupling between prefrontal and visual cortex during attention. Science 324:1207–1210. CrossRef Medline

Gregoriou GG, Gotts SJ, Desimone R (2012) Cell-type-specific synchronization of neural activity in FEF with V4 during attention. Neuron 73:581–594. CrossRef Medline

Hochberg LR, Serruya MD, Friehs GM, Mukand JA, Saleh M, Caplan AH, Branner A, Chen D, Penn RD, Donoghue JP (2006) Neuronal ensemble control of prosthetic devices by a human with tetraplegia. Nature 442: 164–171. CrossRef Medline

Hwang EJ, Andersen RA (2011) Effects of visual stimulation on LFPs, spikes, and LFP-spike relations in PRR. J Neurophysiol 105:1850–1860. CrossRef Medline

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. CrossRef Medline

Kreiman G, Hung CP, Kraskov A, Quiroga RQ, Poggio T, DiCarlo JJ (2006) Object selectivity of local field potentials and spikes in the macaque inferior temporal cortex. Neuron 49:433–445. CrossRef Medline

Leavitt ML, Pieper F, Sachs A, Joober R, Martinez-Trujillo JC (2013) Structure of spike count correlations reveals functional interactions between neurons in dorsolateral prefrontal cortex area 8a of behaving primates. PLoS One 8:e61503. CrossRef Medline

Lebedev MA, Messinger A, Kralik JD, Wise SP (2004) Representation of attended versus remembered locations in prefrontal cortex. PLoS Biol 2:e365. CrossRef Medline

Lennert T, Martinez-Trujillo J (2011) Strength of response suppression to distracter stimuli determines attentional-filtering performance in primate prefrontal neurons. Neuron 70:141–152. CrossRef Medline

Lennert T, Martinez-Trujillo JC (2013) Prefrontal neurons of opposite spatial preference display distinct target selection dynamics. J Neurosci 33: 9520–9529. CrossRef Medline

Liebe S, Hoerzer GM, Logothetis NK, Rainer G (2012) Theta coupling between V4 and prefrontal cortex predicts visual short-term memory performance. Nat Neurosci 15:456–462. CrossRef Medline

Lindén H, Tetzlaff T, Potjans TC, Pettersen KH, Grün S, Diesmann M, Einevoll GT (2011) Modeling the spatial reach of the LFP. Neuron 72: 859–872. CrossRef Medline

Liu J, Newsome WT (2006) Local field potential in cortical area MT: stimulus tuning and behavioral correlations. J Neurosci 26:7779–7790. CrossRef Medline

Markowitz DA, Wong YT, Gray CM, Pesaran B (2011) Optimizing the de-

coding of movement goals from local field potentials in macaque cortex. J Neurosci 31:18412–18422. CrossRef Medline

Maynard EM, Nordhausen CT, Normann RA (1997) The Utah Intracortical Electrode Array: a recording structure for potential brain-computer interfaces. Electroencephalogr Clin Neurophysiol 102:228–239. CrossRef Medline

Miller EK, Cohen JD (2001) An integrative theory of prefrontal cortex function. Annu Rev Neurosci 24:167–202. CrossRef Medline

Mineault PJ, Zanos TP, Pack CC (2013) Local field potentials reflect multiple spatial scales in V4. Front Comput Neurosci 7:21. CrossRef Medline

Mitchell JF, Sundberg KA, Reynolds JH (2009) Spatial attention decorrelates intrinsic activity fluctuations in macaque area V4. Neuron 63:879–888. CrossRef Medline

Monosov IE, Trageser JC, Thompson KG (2008) Measurements of simultaneously recorded spiking activity and local field potentials suggest that spatial selection emerges in the frontal eye field. Neuron 57:614–625. CrossRef Medline

Perge JA, Homer ML, Malik WQ, Cash S, Eskandar E, Friehs G, Donoghue JP, Hochberg LR (2013) Intra-day signal instabilities affect decoding performance in an intracortical neural interface system. J Neural Eng 10: 036004. CrossRef Medline

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. CrossRef Medline

Petrides M (2005) Lateral prefrontal cortex: architectonic and functional organization. Philos Trans R Soc Lond B Biol Sci 360:781–795. CrossRef Medline

Petrides M, Pandya DN (1999) Dorsolateral prefrontal cortex: comparative cytoarchitectonic analysis in the human and the macaque brain and corticocortical connection patterns. Eur J Neurosci 11:1011–1036. CrossRef Medline

Quian Quiroga R, Panzeri S (2009) Extracting information from neuronal populations: information theory and decoding approaches. Nat Rev Neurosci 10:173–185. CrossRef Medline

Rainer G, Asaad WF, Miller EK (1998) Memory fields of neurons in the primate prefrontal cortex. Proc Natl Acad Sci U S A 95:15008–15013. CrossRef Medline

Rasch MJ, Gretton A, Murayama Y, Maass W, Logothetis NK (2008) Inferring spike trains from local field potentials. J Neurophysiol 99:1461–1476. CrossRef Medline

Ray S, Maunsell JH (2011) Different origins of γ rhythm and high-γ activity in macaque visual cortex. PLoS Biol 9:e1000610. CrossRef Medline

Ray S, Maunsell JHR (2014) Do γ oscillations play a role in cerebral cortex? Trends Cogn Sci (Regul Ed) 1–8.

Ray S, Crone NE, Niebur E, Franaszczuk PJ, Hsiao SS (2008) Neural correlates of high-γ oscillations (60–200 Hz) in macaque local field potentials and their potential implications in electrocorticography. J Neurosci 28: 11526–11536. CrossRef Medline

Ruff DA, Cohen MR (2014) Attention can either increase or decrease spike count correlations in visual cortex. Nat Neurosci 17:1591–1597. CrossRef Medline

Squire RF, Noudoost B, Schafer RJ, Moore T (2013) Prefrontal contributions to visual selective attention. Annu Rev Neurosci 36:451–466. CrossRef Medline

Thompson KG, Biscoe KL, Sato TR (2005) Neuronal basis of covert spatial attention in the frontal eye field. J Neurosci 25:9479–9487. CrossRef Medline

Tremblay S, Pieper F, Sachs A, Martinez-Trujillo J (2015) Attentional filtering of visual information by neuronal ensembles in the primate lateral prefrontal cortex. Neuron 85:202–215. CrossRef Medline

Vinck M, Womelsdorf T, Buffalo EA, Desimone R, Fries P (2013) Attentional modulation of cell-class-specific γ-band synchronization in awake monkey area v4. Neuron 80:1077–1089. CrossRef Medline

Womelsdorf T, Fries P, Mitra PP, Desimone R (2006) γ-band synchronization in visual cortex predicts speed of change detection. Nature 439:733–736. CrossRef Medline

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. CrossRef Medline

Zanos TP, Mineault PJ, Pack CC (2011) Removal of spurious correlations between spikes and local field potentials. J Neurophysiol 105:474–486. CrossRef Medline