

Supporting Information

Siegel et al. 10.1073/pnas.0908193106

SI Materials and Methods

Subjects and Surgery. All experiments were performed in two rhesus monkeys (*Macaca mulatta*), one male and one female, weighing 6.0 and 6.5 kg. Under general anesthesia, each monkey was implanted with a head bolt to immobilize the head and a recording chamber for intracortical recordings. All of the surgeries were performed under aseptic conditions with postoperative antibiotics and analgesia. All of the procedures were carried out in accordance with the guidelines of the National Institutes of Health and the Massachusetts Institute of Technology Animal Care and Use Committee.

Behavioral Task. Monkeys performed a visual two-object short-term memory task that required them to remember a sequence of two natural objects over a brief delay. Each trial was started when the animals acquired fixation of a white dot at the center of the stimulus screen. Fixation was required in a $\pm 1.5^\circ$ window around the fixation dot throughout the trial until the presentation of the test screen (see below). After a blank fixation baseline of 1 s, two target objects were presented centrally for 500 ms each, separated by a blank delay interval of 1 s. After another delay interval of 1 s, three test objects were displayed surrounding the fixation dot, two of which were the target objects previously shown during this trial. The three test objects were positioned at 5° eccentricity with the same mutual distance between them (8.7°).

With presentation of the test screen, the monkeys had to report the remembered target objects by performing two successive saccades to the two target objects in the order of their preceding presentation. The monkeys were required to initiate the first saccade up to 2 s after the presentation of the test screen and the second saccade up to 2 s after finishing the first saccade. Direct saccades were enforced by a maximum saccade duration of 70 ms. Trials were aborted if the monkeys broke fixation before test screen presentation or did not meet the required response timing. After successful completion of the second saccade, all of the objects were switched off, and if the correct object sequence was reported, then the monkeys were rewarded with apple juice. Thus, during the second delay interval, the monkeys had to remember the identities of both target objects along with their order of presentation. The monkeys were proficient at this task with an average correct performance of 63%, well above the chance level of 8% ($P < 10^{-16}$ for all recording sessions, binomial test).

For each trial, the two target objects and the third test object were drawn randomly from a set of four objects. For each recording session, four novel objects never before seen by the animals were chosen at random from a database of images (Corel). Using four objects ensured that the animals could not predict the identity of the third, nonmatching object in the test array at the end of the trial. The objects were small complex objects approximately $2^\circ \times 2^\circ$ in size. Importantly, the experimental design was fully balanced for object identity, object order, and the positions of the objects in the test array. In other words, all of the possible object and response sequences occurred with equal probabilities. The objects were presented on a computer screen positioned directly in front of the animal (viewing distance of 60 cm). Object presentation and behavioral monitoring were controlled using two computers running the CORTEX real-time control system. Eye movements were monitored and recorded using an infrared eye-tracking system (ISCAN).

Recording Techniques. Electrode penetration sites were determined using MRI scans obtained before surgery. The recording chambers were positioned stereotactically over the left lateral prefrontal cortex of each animal such that the principal sulcus and lateral prefrontal cortex were readily accessible. During the course of our experiments, we qualitatively determined that the cells in the ventral halves of the recording chambers were more likely to respond selectively to objects and therefore concentrated our recordings in this region.

Monkeys were seated in primate chairs within sound-attenuating enclosures (Crist Instruments), their heads were immobilized using the implanted head bolt, and a juice spout was placed at their mouths for automated reward delivery. Electrophysiological signals were recorded simultaneously from arrays of eight independently movable dura-puncturing tungsten microelectrodes (FHC Instruments). Electrodes were advanced using custom-made screw-driven mini-microdrives (1) mounted on a plastic grid (Crist Instruments) with spacing of 1 mm between adjacent locations. Neuronal activity was amplified, filtered, and stored using an integrated multichannel recording system (Plexon Neurotechnology Research Systems).

To minimize any sampling bias of neuronal activity, we did not prescreen activity for visual responsiveness or object selectivity. Rather, we advanced each electrode until the activity of one or more neurons was isolated well and then began data collection. From each electrode, we simultaneously recorded spiking activity and the local field potential (LFP). Both signals were referenced to ground (rather than to one of the electrodes). This eliminated the possibility of artifactual synchrony due to neural signals measured by the reference itself. The signal from each electrode was divided into spiking activity and the LFP. The spike signal (passband 250 Hz to 8 kHz) was threshold-triggered to separate neuronal spikes from background noise, and individual spike waveforms were stored at 40 kHz. The LFP was recorded continuously with a sampling rate of 1 kHz. We offline-corrected the LFP signals for any phase shifts induced by the microelectrode amplifier circuit and filtering as described in ref. 2. In brief, we empirically measured the frequency-dependent transfer function of our entire recording system (including electrodes, amplifiers, and filters) from 0.5 to 450 Hz and corrected for signal distortions by division with the transfer function in the frequency domain. This procedure also effectively corrected for any amplitude attenuation of the LFP signal due to the recording system's band-pass across the frequency range from 1 to 250 Hz (the signal was well above the noise floor across this range). After correcting for phase and amplitude distortions, we band-pass-filtered the LFP with a sixth-order zero-phase forward–reverse Butterworth filter (1–250 Hz).

In particular, during the delay interval of interest, average firing rates of individual neurons were typically well below 2 Hz. This low number of spikes compromises the sensitivity to measure neural synchrony and to quantify phase-dependent information (which effectively further reduced the number of spikes by binning; see below). Thus, for each electrode, we pooled all of the recorded spikes into multi-unit activity and analyzed activity from those electrodes that showed stable spiking activity for a minimum of 200 trials with a minimum average firing rate of 2 Hz during the second delay interval. This included data from a total of 140 recording sites, with a median of 498 correctly performed trials each.

Average Spike-Rate Analysis. All of the data analyses were performed with custom software written in MATLAB (Mathworks) and C.

We analyzed the time course of the average firing rates (Fig. 1*B*) and their information about the identities of the two objects (Fig. 3*A*) using a sliding-window approach (window size of 100 ms and step size of 25 ms). For each window, we computed the average number of spikes per second and quantified the information about both objects as the percentage variance in the firing rate across trials explained by the identity of each object (one-way ANOVA). The two target objects were always different in identity, which yields an incomplete design matrix. To compute information about the first object, we thus averaged the explained variance across four independent ANOVAs for each identity of the second object and vice versa for information about the second object. For finite sample sizes, the ANOVA-explained variance is biased positively. We estimated and subtracted this bias by computing the average explained variance after randomly permuting object identities across trials (10^3 permutations). The distributions of the explained variance for the permuted data also were used to estimate the probability (P value) of measuring, under the null hypothesis, an explained variance equal to or larger than the nonpermuted explained variance (permutation test).

Spectral Analyses. All of the spectral analyses were based on the same time-frequency transformation of the LFP. We derived the phase and amplitude of the LFP as a function of frequency and time by convolution of the LFP with frequency-dependent kernels k (Hanning-windowed complex sinusoids)

$$k(f,t) = A(1 - \cos(2\pi t f/q))e^{2i\pi t f} \quad [1]$$

where f is the frequency of interest, q is the kernel length in cycles, t is time ranging from 0 to q/f , and A is a constant normalizing k to unit power. We used kernels k for 39 logarithmically scaled frequencies f ranging from 2 to 161 Hz of length $q = 2$. The convolution is equivalent to a sliding-window Fourier transform using a Hanning-tapered window of two times the frequency of interest period. This time-frequency transformation yields an adaptive spectro-temporal resolution with a temporal bandwidth of approximately $1/f$ (FWHM of the Hanning window) and a spectral bandwidth of ≈ 1 octave. The time-varying power of the LFP $P(f,t)$ at frequency f is the squared norm of the convolution of the LFP signal $S(t)$ with the kernel $k(f,t)$

$$P(f,t) = |k(f,t) \times S(t)|^2 \quad [2]$$

The time-varying phase of the LFP $\varphi(f,t)$ at frequency f is the complex argument of the convolution of the LFP signal $S(t)$ with the kernel $k(f,t)$

$$\varphi(f,t) = \arg(k(f,t) \times S(t)) \quad [3]$$

The LFP power decays with frequency ($1/f^n$ power-law decay), which introduces a spectral asymmetry and leakage in estimating the LFP phase. To minimize this, we differentiated the LFP before computing its instantaneous phase (3, 4). Differentiating the signal scales its amplitude by f and rotates its phase by 90° , which results in a simple prewhitening correcting for a $1/f^2$ power decay. All of the estimated phases were corrected for the 90° rotation caused by the differentiation.

LFP Power. To enhance readability of the LFP power at high frequencies, which are masked by the power decay, we normalized the power by $1/f$ (see also Fig. S1). We applied the same temporal binning as for the synchrony analysis (see below) with a sliding window with a length of 250 ms and a step size of 25 ms.

This allowed for a direct comparison of power and synchrony effects. To investigate which LFP activity reflects signal components phase-locked to trial events (event-related potentials), we repeated the power analysis after subtracting from each trial the LFP signal averaged across all of the trials (“evoked potential”; Fig. 1*D Left*). The subtraction had no effect on the prominent activity at ≈ 32 Hz, which thus was not phase-locked to trial events (compare Fig. 1*C* and *D Right*). In contrast, the transient low-frequency activity at ≈ 2 –4 Hz around object presentation was largely removed and thus attributable to the prominent low-frequency fluctuations of the evoked potential.

Spike-LFP Synchrony. We measured phase synchrony between spikes and LFPs recorded simultaneously from neighboring electrodes (distance of 1 mm, $n = 140$ pairs). This precluded measuring artificial synchrony caused by spillover from the spike signal to the LFP if recorded from same electrode. We derived the time-varying LFP phase from 2 to 161 Hz according to Eq. 3 and measured for each spike its instantaneous LFP phase. We characterized the temporal evolution of spike-LFP synchronization by a sliding-window analysis (window length of 250 ms and step size of 25 ms). For n spikes in each temporal window t with phase $\varphi(f,t)$, the strength of the phase synchronization is quantified by the phase-locking value PLV(f,t) (length of the vector average, representing all of the spikes as unit vectors)

$$PLV(f,t) = \left| \frac{1}{n} \sum_{k=1}^n e^{i\varphi_k(f,t)} \right| \quad [4]$$

The preferred phase of spiking $\varphi_{pref}(f,t)$ is the argument (phase) of the vector average

$$\varphi_{pref}(f,t) = \arg\left(\sum_{k=1}^n e^{i\varphi_k(f,t)} \right) \quad [5]$$

To assess the statistical significance of synchronization, we used a Rayleigh test against the null hypothesis of circular uniformity (5).

Importantly, with a decreasing number of spikes n , the PLV is biased positively toward 1, and the sensitivity of the Rayleigh test is reduced. Thus, before quantifying phase synchrony, we stratified the number of spikes across all of the temporal windows. For each window, we randomly removed the number of spikes in excess of the number of spikes in the window of the minimum firing rate. Thus, changes in the firing rate across the trial could not bias the measured phase synchrony.

We displayed the spectro-temporal evolution of spike-LFP synchrony as the percentage of electrode pairs showing a significant synchronization ($P < 0.01$; Fig. 2*A*). The average PLVs across all of the pairs are displayed in Fig. S2. In analogy to the LFP power analysis, we assessed which synchrony is due to phase-locking to stimulus presentation by removing from each trial the average LFP across trials (Fig. S2).

The average preferred phase of spiking at ≈ 3 and 32 Hz (Fig. 2*B* and *D*) was measured by averaging the preferred phase across the entire trial. We fitted the histogram of the preferred phases across all electrode pairs with a von Mises distribution (5) and estimated the SEM of the overall preferred phase by bootstrap with 10^3 repetitions. We tested for a significant phase concentration of preferred phases using a Rayleigh test. To estimate the strength of spike-rate modulation by LFP phase (Fig. 2*C* and *E*), for each electrode pair we computed the firing rate in 12 equally spaced phase bins relative to each pair’s preferred phase of spiking. We normalized firing rates by the average rate across all bins and fitted the data with a von Mises distribution. We estimated the SEM of the modulation strength by bootstrap with 10^3 repetitions.

Phase-Dependent Coding. To test for a relationship between LFP phase and the encoding of object information, we focused on the second delay interval and those recording sites for which average firing rates conveyed information about both object identities during that interval (ANOVA, permutation test, $P < 0.01$, $n = 103$ of 140 or 74%). For each frequency f (3–161 Hz), we binned all spikes by their phase into 12 equally spaced bins b and for each bin measured the information $I(f, b, s)$ that spikes at this phase conveyed about the identity of the first or second object s (one-way ANOVA explained variance). As described for the average firing rates above, for each frequency phase bin we subtracted the bias in explained variance estimated by permuting object identities across trials (10^3 permutations). To account for possible nonstationarities between the early and the late part of the second delay interval (Figs. 1B and 3A), we conducted this analysis independently for the first and second 500 ms of that interval and then averaged information across these. We smoothed $I(f, b, s)$ with a two-dimensional Hanning kernel (0.5 octave \times 90°, FWHM) and normalized, for each frequency, the information (explained variance) by the average across all of the phases (Fig. 3B).

We quantified the phase-dependent information $PDI(f, s)$ about stimulus s at frequency f (Fig. 3C) as the peak-to-mean modulation of a cosine fit to the average Information I as a function of phase, normalized by the average information across all of the phases

$$PDI(f, s) = 4 \frac{\left| \sum_{k=1}^K \sum_{l=1}^L (I_k(f, b_l, s) e^{i\lambda_l}) \right|}{\sum_{k=1}^K \sum_{l=1}^L (I_k(f, b_l, s))} \quad [6]$$

where $K = 103$ is the number of recording sites and $L = 12$ is the number of phase bins b with center phases λ . In other words, the PDI measures the peak increase and decrease in information at the best and worst encoding phases in percentage of the average information across all of the phases. To test for a PDI larger than 0, while correcting for multiple comparisons across frequency, we used a cluster-based permutation test: For each site, frequency and both objects, we randomly permuted $I(f, b, s)$ across phase bins (10^4 permutations, constant across frequency), which yields a distribution of PDI spectra under the null hypothesis of equal information across phases. For the measured and permuted PDI spectra, we defined clusters as continuous frequency bands with PDIs exceeding the 95th percentile of the permuted PDIs and computed each cluster's average PDI. The P value of each measured cluster equals the probability of clusters with an equal or larger average PDI in the distribution of permuted PDI spectra. We applied the same approach without clustering and smoothing across the frequency dimension to investigate the phase-dependent coding at 3 and 32 Hz (Fig. 4).

The optimally encoding phase $\varphi_{opt}(f, s)$ is the peak phase of a cosine fit to the information–phase relation (Fig. 4)

$$\varphi_{opt}(f, s) = \arg \left(\sum_{k=1}^K \sum_{l=1}^L (I_k(f, b_l, s) e^{i\lambda_l}) \right) \quad [7]$$

where $K = 103$ is the number of recording sites and $L = 12$ is the number of phase bins b with center phases λ . We tested for a significant difference between the optimally encoding phases for objects 1 and 2 using a permutation test: We tested the measured difference in φ_{opt} against its distribution under the null hypothesis of no phase difference generated by randomly permuting $I(f, b, s)$ between objects (constant across phases) for all recording sites (10^4 permutations). The optimally encoding phases and the significance of a phase difference for all frequencies from 2 to 161 Hz with a significant phase-dependent coding are dis-

played in Fig. S3. The SEMs of PDI and φ_{opt} were estimated by bootstrap across recording sites (10^3 repetitions).

We hypothesized that the observed phase difference between optimal information for both objects may be related to a difference in the frequency of the phase-dependent coding. We thus tested for such a frequency difference using a permutation test: We computed the center of gravity of the PDI spectra for both objects from 8 to 128 Hz (Fig. 3C) and compared their difference against the distribution of differences obtained by randomly permuting the PDI spectra between objects (constant across frequency) for all recording sites (10^4 permutations). There was no significant difference ($P = 0.8$).

To investigate the relationship between phase-dependent coding at 3 and 32 Hz, we measured the information that spikes conveyed about each object's identity as a function of both frequencies' phases (Fig. S4). Again, for both frequencies, we binned information by 12 equally spaced phase bins. We smoothed the information with a two-dimensional Hanning kernel (0.5 octave \times 0.5 octave, FWHM) and fitted it with a two-dimensional cosine (additive modulations for both frequencies). Information was modulated significantly across both frequencies ($P < 10^{-4}$, permutation test). We then tested for an interaction between both frequencies using two approaches. First, we applied a nonparametric two-way ANOVA, and second, we directly tested for a cosine modulation of the 32-Hz phase-dependent information by the 3-Hz phase (permutation tests with 10^4 repetitions). Both tests did not yield a significant interaction between the 3- and 32-Hz phase-dependent coding ($P > 0.4$).

Control Analyses. We performed several control analyses to assess the robustness of the order effect on the optimally encoding phase (Fig. S5). First, we tested if the effect was specific to the particular binning used to separate spikes of different phases (12 bins). We thus repeated the analysis based on eight phase bins (Fig. S5B). We obtained qualitatively the same result.

Second, we tested if the order effect on optimally encoding phases was accountable to our removal of LFP activity phase-locked to trial events. We thus repeated the analysis without removing the phase-locked components from the LFP (Fig. S5C). This yielded qualitatively the same order effect.

Third, we tested if the order effect may be accountable to the specific measure of information used to quantify the dependency between phase-dependent spikes and object identity (explained variance). We thus repeated the analysis based on mutual information (6) (Fig. S5D). As described for the explained variance, we estimated and subtracted the bias in mutual information for finite sample size by permuting object identities across trials (10^3 permutations). The mutual-information-based analysis yielded qualitatively the same result.

Anatomical Specificity. To assess the anatomical (spatial) specificity of the demonstrated effects, we repeated the analyses of spike-field synchrony, phase-dependent coding, and phase–amplitude coupling (see below) for electrode pairs separated by 3 mm. Fig. S6 compares the results obtained for this larger electrode distance with the original results (distance of 1 mm). The larger electrode separation strongly reduced spike-field synchrony, in particular for frequencies at ≈ 32 Hz (compare Fig. S6A and B), and abolished the phase-dependent coding of object information (both objects and bands $P > 0.05$, permutation test; Fig. S6D and F). In contrast, the 3-to-32-Hz phase–amplitude coupling was largely sustained across the distance of 3 mm (Fig. S6H).

Error-Trial Analyses. The monkeys' good performance (average of 63% correct) meant that there were many fewer error trials than correct trials. To ensure equal statistical power and sample sizes

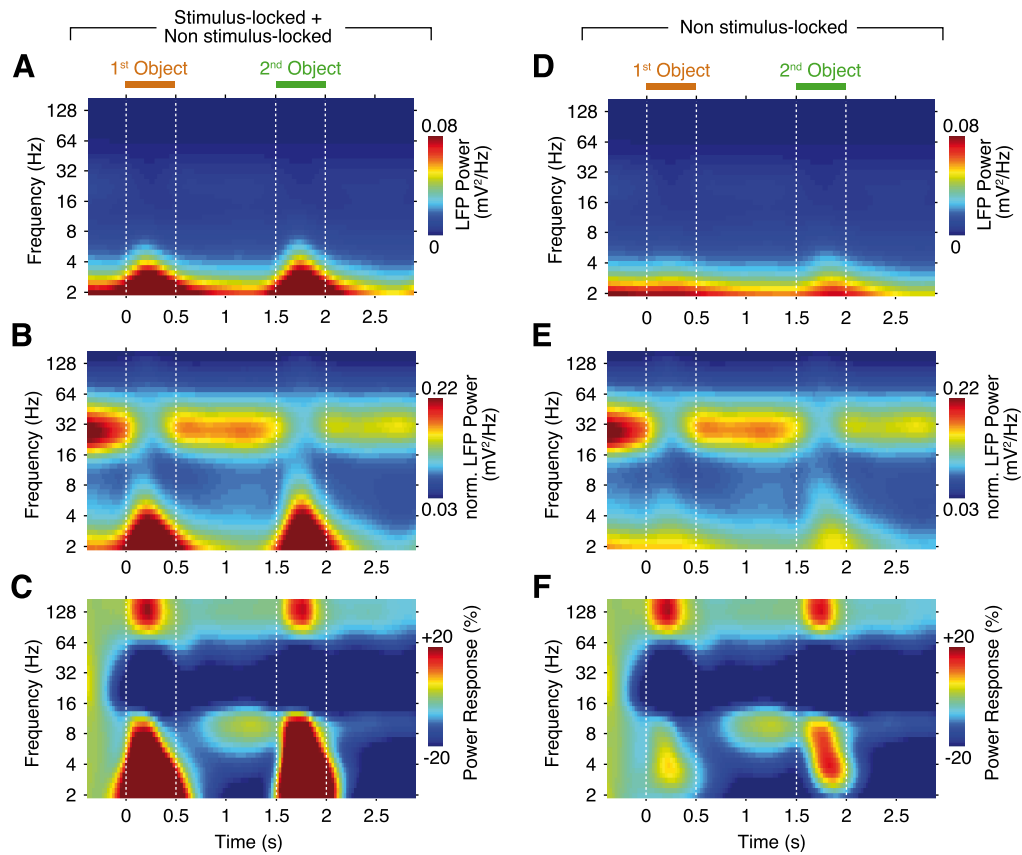


Fig. S1. (A) Time-frequency representation of the average non-normalized local field potential (LFP) power ($n = 140$). (B) Time-frequency representation of the average LFP power normalized by $1/\text{frequency}$ as displayed in Fig. 1C ($n = 140$). (C) Time-frequency representation of the average percentage change in LFP power relative to the prestimulus fixation baseline ($n = 140$). This representation of the data well displays the increase in signal power at frequencies >60 Hz. (D–F) Same time-frequency representations of the LFP power after subtracting out all of the signal components phase-locked to stimulus presentation. This primarily removes the low-frequency responses around the time of stimulus presentation corresponding to the “evoked fields.”

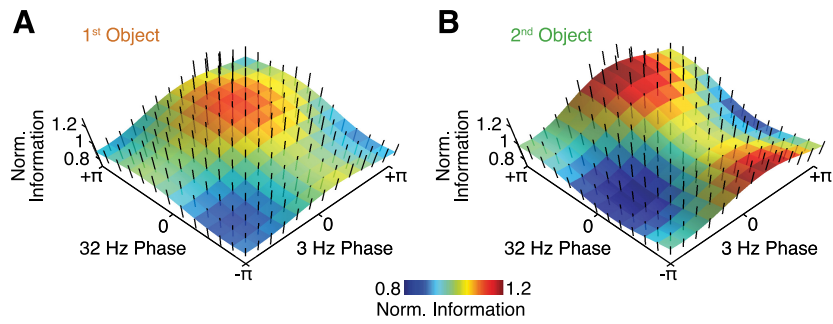


Fig. S4. Phase-dependent coding for the combined 3- and 32-Hz phases. (A) Normalized information in spikes about the first object's identity during the second delay interval as a function of the simultaneous 3- and 32-Hz LFP phases (distance of 1 mm, $n = 103$ pairs). The black lines display the information plus or minus one bootstrap SEM. The semitransparent surface displays a two-dimensional cosine fit. (B) Same analysis for information about the identity of the second object.

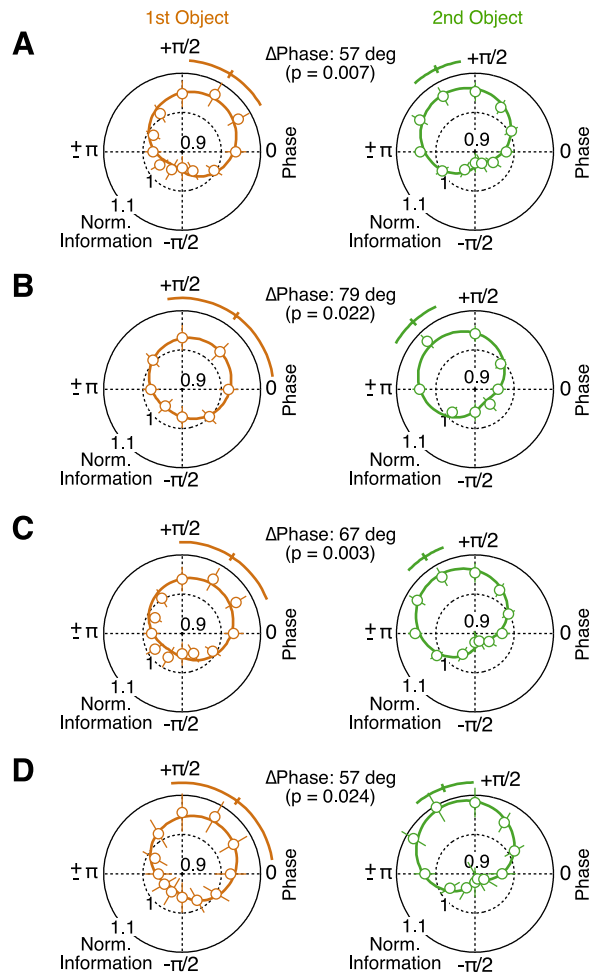


Fig. S5. Robustness of the order effect on the optimally encoding 32-Hz phases across several control analyses ($n = 103$). (A) The original result as displayed in Fig. 4A. (B) Optimally encoding phases based on eight phase bins. (C) Phase-dependent coding without subtraction of LFP components phase-locked to object presentation. (D) Phase-dependent coding based on mutual information.

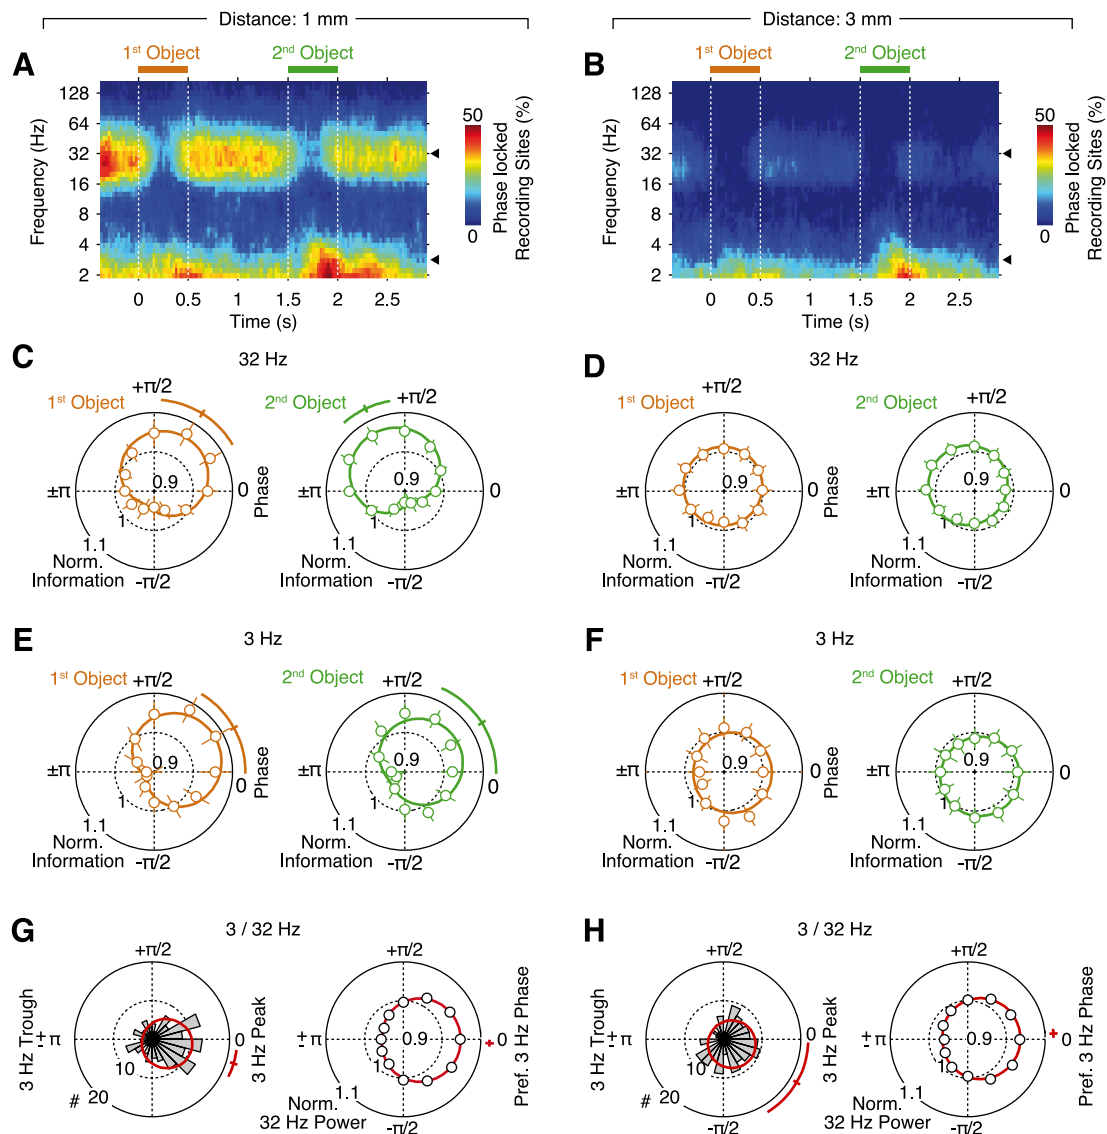


Fig. S6. Spatial specificity of spike–LFP synchrony, phase-dependent coding, and phase–amplitude coupling. (A) Time–frequency representation of the percentage of spike–LFP pairs that showed significant ($P < 0.01$) phase synchrony ($n = 140$ pairs) after subtracting LFP components phase-locked to stimulus presentation (distance of 1 mm; same data as Fig. 2A). (B) Same synchrony analysis as A for an electrode distance of 3 mm. (C, D) Normalized information about the first and second presented objects as a function of the 32-Hz LFP phase for electrode distances of 1 and 3 mm, respectively. Solid traces display a cosine fit, the average optimally encoding phase, and its bootstrap SEM (distance of 1 mm, $n = 103$ pairs; C shows the same data as Fig. 4A). (E, F) Same analysis of phase-dependent coding for 3 Hz (E shows the same data as Fig. 4B). (G, H) The 3-to-32-Hz phase–amplitude coupling for electrode distances of 1 and 3 mm. (Left) Distribution of 3-Hz phases of peak 32-Hz activity across all of the pairs. (Right) Strength of the 32-Hz power modulation by the 3-Hz phase relative to the peak phase ($n = 140$ pairs; G shows the same data as Fig. 5).

