

Temporal structure in neuronal activity during working memory in macaque parietal cortex

Bijan Pesaran¹, John S. Pezaris^{2,3}, Maneesh Sahani^{2,4}, Partha P. Mitra⁵ and Richard A. Andersen^{2,6}

¹ Division of Physics, Mathematics & Astronomy, California Institute of Technology, Pasadena, California 91125, USA

² Computation and Neural Systems Program, California Institute of Technology, Pasadena, California 91125, USA

³ Present address: Department of Neurobiology, Harvard Medical School, Boston, Massachusetts 02115, USA

⁴ Present address: Gatsby Computational Neuroscience Unit, Alexandra House, 17 Queen Square, London, WC1N 3AR, UK

⁵ Bell Labs, Lucent Technologies, Murray Hill, New Jersey 07974, USA

⁶ Division of Biology, Mail Code 216-76, California Institute of Technology, Pasadena, California 91125, USA

Correspondence should be addressed to R.A.A. (andersen@vis.caltech.edu)

Published online: 22 July 2002, doi:10.1038/nn890

Many cortical structures have elevated firing rates during working memory, but it is not known how the activity is maintained. To investigate whether reverberating activity is important, we studied the temporal structure of local field potential (LFP) activity and spiking from area LIP in two awake macaques during a memory-saccade task. Using spectral analysis, we found spatially tuned elevated power in the gamma band (25–90 Hz) in LFP and spiking activity during the memory period. Spiking and LFP activity were also coherent in the gamma band but not at lower frequencies. Finally, we decoded LFP activity on a single-trial basis and found that LFP activity in parietal cortex discriminated between preferred and anti-preferred direction with approximately the same accuracy as the spike rate and predicted the time of a planned movement with better accuracy than the spike rate. This finding could accelerate the development of a cortical neural prosthesis.

The neural basis of working memory is typically studied in non-human primates by recording activity during delayed response tasks¹. Cue-selective elevations in mean firing rates are found during the delay period in many brain areas using different versions of these tasks^{2–5}. Parietal cortex is important for spatial cognition⁶; spatially tuned increases in firing rate during working memory were first reported in parietal cortex in the lateral intraparietal area (area LIP)³. In analogy to receptive fields in the sensory system, such activity can be said to form memory fields⁷ and is thought to reflect the plan to make a movement^{8,9}. Converging evidence indicates memory fields may be important for understanding the neural basis of working memory¹⁰, and this activity is the subject of current theoretical work^{11,12}.

Temporally correlated neuronal activity in the form of reverberations is proposed to be important to cognitive processing^{13,14}. However, temporal structure is not captured by measures of mean activity, such as the firing rate, so measures of the variance, such as the spectrum, are required. We characterize temporal structure by peaks in the spectrum corresponding to concentrated power at certain frequencies and times. If reverberant activity is present in neuronal activity, temporal structure localized in frequency may also predict planned saccades. This would be evidence for memory fields of temporal structure, or 'dynamic memory fields', defined by

tuned changes in the spectrum of neuronal activity. The temporal structure of neuronal activity has been related to perception^{15–18}, attention^{19,20} and action^{21–26}. Temporal structure during object working memory occurs in spike activity in monkeys²⁷ and in EEG activity in humans²⁸. Thus, the presence of tuned temporal structure in neuronal activity could have implications for neural information processing.

Here we present spectral analysis of spiking and local field potential (LFP) activity recorded in area LIP from two macaques during a memory-saccade task. Spike activity has been examined in area LIP during this task^{3,29}, but LFP activity has not. LFP activity is generated by extracellular currents due to synchronized activity in a local population of cells³⁰. Temporal structure is present in LFP activity in cats¹⁷ and monkeys^{15,18}. We used multi-taper spectral analysis techniques, which were important for obtaining the present results, and found significant structure in the spectrum of LFP activity, spiking and the coherency between them. The results are evidence for dynamic memory fields in neuronal activity that are specific to both the direction and time of a planned movement. We suggest the movement can be decomposed into a plan state and an execution state and show that both the direction and the state can be decoded from LFP activity on a single-trial basis. These results may accelerate the development of a cortical neural prosthesis, as the LFP is easier to record than action potentials.






Fig. 1. The memory-saccade task. **(a)** The monkey performs a saccade to a memorized location in one of eight directions. **(b)** Sample trace of extracellular potential for a trial during a saccade to the preferred direction. The polarity of the potential is reversed. The data is viewed on an expanded time base during the memory period from 0.6–1.1 s below. **(c)** Sample trace for a saccade to the anti-preferred direction.

RESULTS

The database for this study contained activity from 16 cells recorded at 16 sites in one monkey and 24 cells recorded at 18 sites in another monkey. Recordings were made with tetrodes while animals performed a memory-saccade task (Fig. 1a). We found 28 of 40 neurons (70%) had significant memory period activity ($P < 0.05$; ANOVA). Single-unit and LFP activity were extracted from recordings of broadband (1 Hz–10 kHz) extracellular potentials (Methods). In addition to spiking activity, large-amplitude LFP fluctuations were visible in the raw data (see the expanded time base) during saccades to the preferred direction for that site (Fig. 1b). These fluctuations were not visible during saccades in the opposite direction (Fig. 1c).

Temporal structure in LFP activity

We estimated spectrograms to determine how LFP activity was related to the task. In the preferred direction, an increase in gamma-band (25–90 Hz) power was sustained during the memory period through the saccade for a single site and across the population (Fig. 2a and c). In contrast, gamma-band power did not change before saccades to the anti-preferred direction (Fig. 2b and d). LFP gamma-band power was significantly greater ($P < 0.05$; *t*-test) in the preferred direction during the memory period than baseline in 27 of the 34 recorded sites (82%) and significantly greater ($P < 0.05$; *t*-test) in the preferred direction than in the anti-preferred direction at 28 of the 34 sites (85%).

The increase in gamma-band LFP power during the memory period is evidence for dynamic memory fields that may be similar to those seen in the spike rate. We estimated tuning curves pooled across the population of recordings from both monkeys for the mean firing rate (Fig. 3a), gamma-band LFP power (Fig. 3b, thick) and low-frequency (0–10 Hz) LFP (Fig. 3b, thin). Each curve was aligned to the direction that elicited the maxi-

mum firing rate for the single unit at that site before averaging. The gamma band activity showed spatial tuning similar to the tuning seen for the rate of firing. LFP activity at low frequency did not show significant spatial tuning, although activity in this band was larger in amplitude. This result shows tuning of the LFP during memory is not simply due to an overall increase in power of the LFP signal.

Although activity at lower frequencies (< 25 Hz) did not show simple spatial tuning, power in the beta frequency band (~20 Hz) of sensorimotor cortex shows complex dynamics during movement preparation and execution²³. We examined LFP power in the gamma and beta frequency bands during saccades to the preferred direction. Elevated gamma-band power was present during the memory period, as reported above. We also found that activity in the beta band was modulated, increasing toward the end of the memory period before dropping around the saccade. Suppression of beta band activity in the perisaccadic period compared with the memory period was significant ($P < 0.05$; *t*-test) at 34 of 34 recording sites (100%).

Temporal structure in spike activity

Increases in the mean firing rate of cells during working memory tasks have been cited as evidence for memory fields that predict saccades to a remembered location in space^{3,7}. We have shown that gamma-band LFP activity in area LIP also contains similar memory fields. Now we show dynamic memory fields for spiking activity. Dynamic memory fields in spiking are defined by changes in the shape of the spike spectrum during memory compared to the baseline. This is in contrast to memory fields described by the mean firing rate, which are defined by changes in the level of the spectrum, not the shape.

We performed spectral analysis of spike activity from a typical cell with memory activity. During the baseline (Fig. 4a), the level of

Fig. 2. LFP spectrograms averaged across trials during saccades to either the preferred or anti-preferred direction. **(a)** Two-dimensional plot of the spectrogram in the preferred direction at a single site. Time is on the x-axis; frequency is on the y-axis. Power is color-coded on a log scale. Square, time of the cue. Triangle, mean time of the saccade. **(b)** Anti-preferred direction at a single site. **(c)** Preferred direction averaged across a population of recordings from both monkeys. **(d)** Population average in the anti-preferred direction.

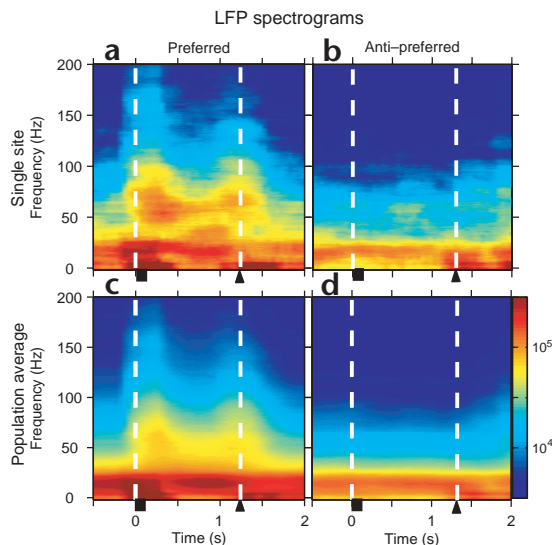


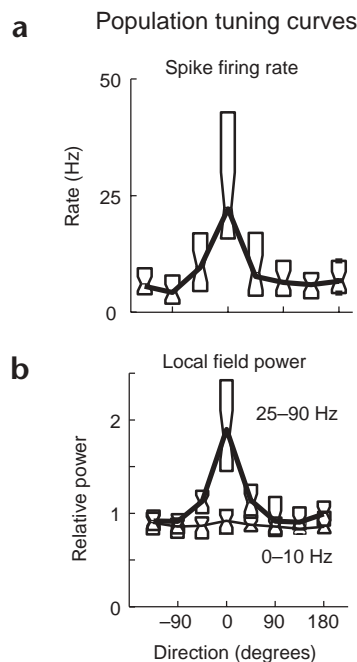
Fig. 3. Tuning of spiking and LFP activity. (a) Spike rate tuning curve with box plots showing mean firing rate across the population. (b) LFP tuning curve with box plots showing relative power in the gamma band (25–90 Hz; thick line) and at low frequency (0–10 Hz; thin line) across the population.

the spectrum was equal to its mean firing rate (dotted line). The spectrum of a Poisson process with this rate would fall on this line. During the memory period, the spectrum was significantly different from a Poisson process (Fig. 4b). There was a significant peak in the spectrum in the gamma band at 50 Hz and significant suppression at low frequency (0–10 Hz). This was representative across a population of cells. Some cells showed significant spectral structure during the baseline, but the population average was approximately Poisson (Fig. 4c). In contrast, the same average during memory contained a broad peak in the gamma frequency band (Fig. 4d). So in addition to an increase in firing rate, spiking activity during memory contained a broad gamma-band spectral peak.

We estimated the spectrogram of spike activity during trials to the preferred direction to see how this structure developed during the trial. The spectrogram of single-cell activity during trials to the preferred direction showed that the gamma-band peak is sustained during the memory period (Fig. 5a). However, because the spike spectrum goes to the rate at high frequency, it is dominated by changes in the rate. Before constructing the population average spike spectrogram, we rescaled the spectrum to normalize for variations in the rate. This rate-normalized spectrogram emphasizes changes in temporal structure that are independent of changes in the rate. The rate-normalized spectrogram showed that the gamma band temporal structure began after the offset of the spatial cue (Fig. 5b). Because the firing rate sharply increased during the cue, this feature indicates that the gamma-band activity is not just due to an increase in firing rate.

Coherency between spiking and LFP activity

The temporal structure we observed in both spiking and the LFP suggests spike timing may be correlated with the gamma-band oscillatory process in the LFP. Such structure has been reported in spiking and LFP activity in the visual cortex^{16,31}. We investigated this correlation by studying the spike-triggered average extracel-



lular potential (STA) and by calculating the coherency between spiking and LFP activity during baseline and memory periods.

We estimated the STA potential for all spikes during the baseline period (Fig. 6a) and during working memory (Fig. 6b) at a single site. The average during working memory showed a gamma-band oscillatory component that is absent from the baseline. This oscillatory activity remained just below the error bars across much of the window, although it is visible to the eye. We reasoned this was because error bars constructed in the time domain were localized in time and not suitable to detect signals localized in frequency.

To address this problem, we estimated the coherency between spiking and LFP activity on a moving window during saccades to the preferred direction. An increase in the gamma-band coherency with zero phase was present after the spatial cue and during the memory period (Fig. 7a). As indicated by the arrows, the phase of the coherency was approximately 0° when the coherency was significant ($P < 0.01$; t -test). This zero phase means the cell fired on or just before the peak of the LFP oscillation (according to our convention for the LFP; Methods). Moreover, the phase of the coherency was relatively constant across the gamma band. This constancy means spike and LFP activity showed some phase-locking during memory and did not lag each other in time. The population average coherency also showed a significant increase in gamma band between the LFP and spiking ($P < 0.01$; t -test; Fig. 7b). Additional analysis of the phase distribution between the spikes and gamma-band LFP activity during memory across the population of recordings found the distribution to be unimodal and peaked at 0 degrees. This population analysis showed spiking is coherent with gamma-band activity in the LFP during

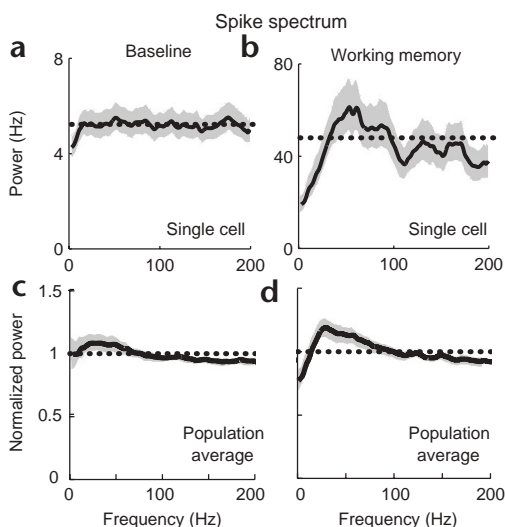


Fig. 4. Spike spectrum. (a) Line plot of the spectrum of a single cell during the baseline period (solid) with 95% error bars (shaded) estimated with a jackknife across tapers and trials. High-frequency limit (dotted). (b) Single cell during working memory. (c) Population average during the baseline. (d) Population average during working memory.

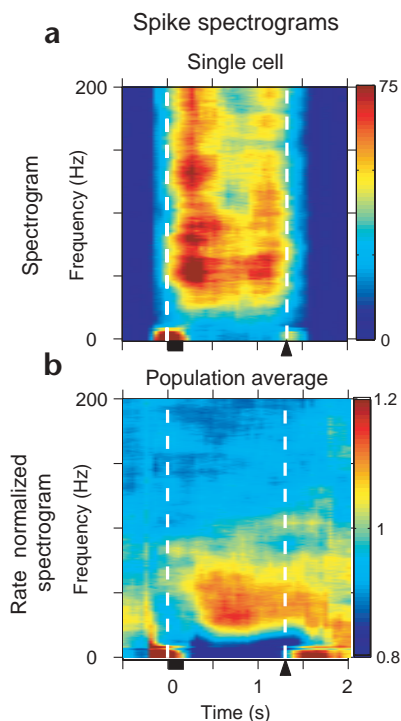


Fig. 5. Spike spectrograms. (a) Two-dimensional plot of the spectrogram for a single cell. (b) Two-dimensional plot of the population average spectrogram normalized by the rate. Time is on the x-axis; frequency is on the y-axis. Power is color-coded on a linear scale. Square, time of the cue. Triangle, mean time of the saccade.

working memory, with the same preferred phase across the sites in area LIP we sampled.

Single-trial decode of spiking and LFP activity

The presence of dynamic memory fields in LFP activity that are similar to memory fields defined by the spike rate suggests that the LFP spectrum could be decoded on a single-trial basis to control a neural prosthesis. Because our analysis of mean spectral responses during the task suggests that LFP activity encodes both the direction and time of the saccade, we sought to extract both features from the LFP spectrum on a single-trial basis. Using linear discriminant analysis, we decoded direction by comparing the LFP activity during the memory period for saccades in the preferred and anti-preferred direction to determine how well it predicts the direction of the impending saccade. We also decoded behavioral state by comparing LFP activity during the memory period and the perisaccadic period to determine whether the LFP activity could indicate the transition from planning to execution state, irrespective of saccade direction. Finally, we repeated both the direction and state decodes using the simultaneously recorded spike rate to examine whether it carried more or less information about saccade direction and timing, on a single-trial basis.

For the direction decode, we found that the spike rate and LFP spectrum could be decoded with approximately equal accuracy (Fig. 8a). Over the database, the average probability of a correct prediction for a saccade to the preferred direction was 87% using the spike rate and 87% using the LFP spectrum. The probability for the anti-preferred direction was 78% using spikes and 87% using the LFP. The frequency of the LFP spectrum that gave the best direction decode varied from site to site in the range

30–100 Hz. The state could not be decoded from the spike rate using this simple procedure, whereas it could be decoded from the LFP spectrum (Fig. 8b). Average probabilities for the plan state were 56% using spikes and 71% using the LFP and for the execution state were 57% using spikes and 71% using the LFP. The frequency of the LFP spectrum that gave the best state decode varied from site to site in the range 0–20 Hz. This is a different frequency band than was used for the direction decode.

These results show that the direction of a planned movement but not the execution time can be easily decoded from the spike rate of a single cell. In contrast, the LFP spectrum codes for the direction and time of a planned movement. Therefore, the LFP at a single site contains information about movement planning and execution multiplexed in different frequency bands that can be decoded from a single trial.

DISCUSSION

We investigated the relation of neuronal dynamics to working memory by examining temporal structure in spiking and LFP activity in area LIP during a memory-saccade task. We report four principal findings. LFP activity reveals spatially tuned dynamic memory fields in the gamma band (25–90 Hz) but not at lower frequencies. Spike activity contains temporal structure during working memory and not simple fixation. Spiking and LFP activity are coherent in the gamma band with the same phase across LIP during working memory. LFP activity can be decoded to give the direction and time of a planned movement. We first discuss an important methodological issue relating to this work.

Correlation functions in time and frequency domains

The presence of temporal structure in neural activity is of much interest³², and we have investigated it in these data using correlation function measures in the temporal domain³³. Correlation functions are often estimated in the time domain but suffer serious problems of estimation bias and variance³⁴. These problems often cannot be addressed by pooling observations across large

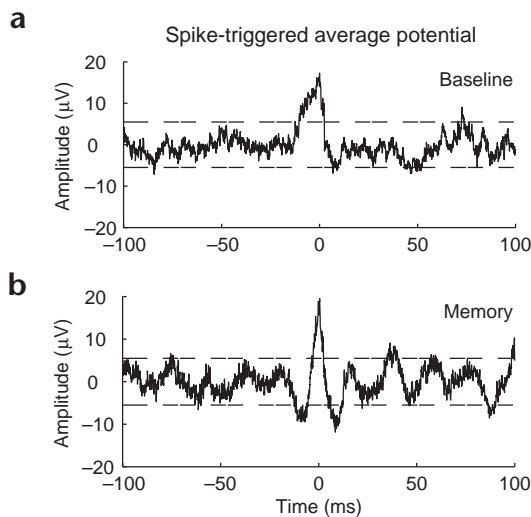


Fig. 6. Spike-triggered average potential from a single cell at a single site. (a) Baseline. (b) Memory.



Fig. 7. Coherency of spiking and LFP activity across time. **(a)** Coherency at a single site in the preferred direction. **(b)** Coherency for the population average in the preferred direction. Time is on the x-axis; frequency is on the y-axis. Square, time of the cue. Triangle, mean time of the saccade. The coherence is color-coded on a linear scale. Arrows denote the phase where the coherence is significant ($P < 0.01$; *t*-test).

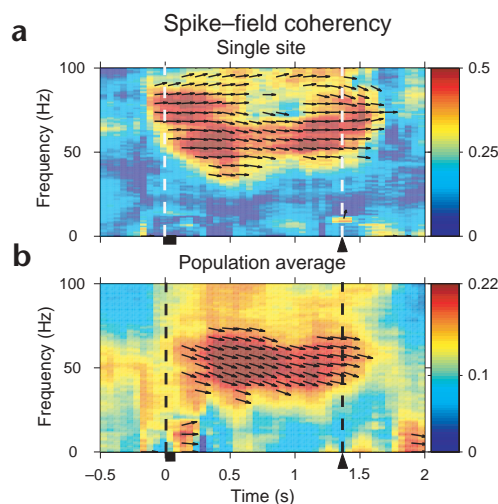
periods of time, as changing behavior leads to violation of the stationarity assumption and potential misinterpretation of the data.

Here we resolve tuned temporal correlations in neuronal activity in parietal cortex using spectral analysis. Spectral quantities are estimated in the frequency domain, and although the corresponding population quantities are mathematically equivalent to correlation functions, spectral estimates have superior statistical estimation properties. The problems of estimation bias and variance can be controlled by using multitaper methods^{35,36}. The advantages of multitaper methods for spectral analysis are most pronounced when studying short segments of possibly nonstationary data, which makes them well-suited for studying neuronal dynamics in behaving animals³⁷ (see **Supplementary Methods** online).

Temporal structure in the LFP

Our results show that gamma band (25–90 Hz) LFP activity is organized in dynamic memory fields whose spatial tuning suggests they reflect the columnar organization of area LIP³⁸. Additional temporal structure at lower frequencies (0–20 Hz) is related to movement preparation and execution. These results indicate that temporal structure in the LFP in parietal cortex contains a variety of specific information about the direction and time of planned movements. In future work, it will be important to determine how directional tuning of the LFP is organized in cortex. The relationship of coherent gamma and unit activity may only apply to a small volume of tissue, possibly localized to specific cortical layers.

Posterior parietal cortex is implicated in higher cognitive function. The spiking activity in many parietal areas codes for movement plans^{6,8,9}, and there is interest in decoding these planning signals for neural prosthetics (ref. 39 and Shenoy *et al.*, *Soc. Neurosci. Abstr.* 25, 383, 1999). However, acquiring spiking activity from many cells with chronically implanted electrodes over long periods of time is difficult. Because the LFP contains specific information about movement planning and is easier to acquire than spikes, our findings suggest that decoding the temporal structure in LFP activity could accelerate the development of this



application. Indeed, we found that the LFP temporal structure at a single site contained as much information about the planned direction of movement and even more information about the behavioral state than the activity of a single cell.

Understanding the temporal structure in spike activity

Spikes exhibited phase coherence to the LFP and preferentially fire at the peak of a broadband LFP oscillation throughout the memory period and not during simple fixation. As this temporal patterning was modulated by behavior, it may reflect neural processing. Such activity need not be specific to working memory. Area LIP is activated by other oculomotor tasks that require spatial cognition⁴⁰, and coherent gamma band activity may also be present during those tasks.

Our finding of correlations in a broad gamma frequency band modulated by behavior is very similar to previous work in the visual cortex¹⁸ and complements recent reports of gamma-band activity during attention in the monkey¹⁹ and working memory in human EEG²⁸. These and other studies suggest that investigating correlations between spikes and LFP activity may help bridge the gap between modulated temporal structure in EEG activity in humans and spike activity in non-human primates.

The finding of robust oscillatory activity in the gamma band of the LFP and the temporal coherence of the spikes to these oscillations suggests that reverberant circuits are a basis for working memory¹⁴. The

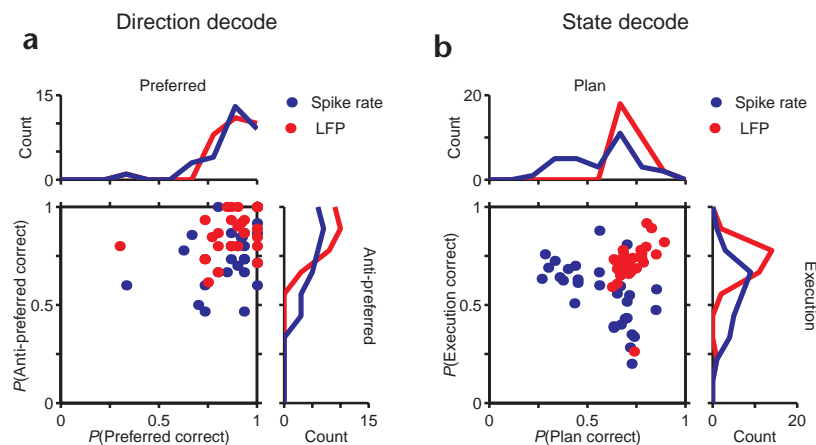


Fig. 8. Single-trial decoding of a movement plan. **(a)** Direction was decoded using spike rate (blue) and the LFP spectrum (red). Each dot represents a single cell or site. Horizontal axis is the probability that a saccade to the preferred direction is decoded correctly. Vertical axis is the probability that a saccade to the anti-preferred direction is decoded correctly. Line plots show the histograms of cell/site counts for each direction. **(b)** State decode. Horizontal axis is the probability that the activity from the plan state is decoded correctly. Vertical axis is the probability that the activity from the execution state is decoded correctly. Line plots show the histograms of cell/site counts for each state.



temporal structure we observed also may have implications for how neural information is processed in LIP. For instance, it has been proposed that synchronous activity may be involved in attention by increasing the synaptic efficiency of projections onto downstream cortical areas¹⁹. Synchronization of activity has also been proposed as a method for binding features processed in different cortical areas³² or as a method for parsing spike activity for classifying different odors⁴¹. Our current experiments cannot determine decisively whether temporal structure is a signature of reverberation for working memory or, for that matter, any of the other possible functional roles mentioned above. However, the finding of highly structured and robust temporal activity in the spikes and field potentials of area LIP does open an intriguing line of research for establishing its functional meaning, which may be far reaching in terms of its role in working memory and the processing of neural signals.

METHODS

Animal preparation, behavioral task and electrophysiology. Recordings were made from two adult male Rhesus monkeys (*Macaca mulatta*) during head restraint. All surgical procedures and animal care protocols were approved by the California Institute of Technology Institutional Animal Care and Use Committee and were in accordance with National Institutes of Health Guidelines. Recordings were made while animals performed a memory-saccade task (Fig. 1a). The baseline period was 1000 ms, the cue duration was 100 ms, and the memory period was 1000 ms. Cued locations were randomly interleaved to collect 10–15 successful trials for each location in blocked fashion. Stimuli were generated with a two-beam optical bench from incandescent sources. Electrical activity was recorded using single tetrodes adapted for use in awake monkeys⁴². Tetrodes were made from 12 μm tungsten wire (California Fine Wire, Grover Beach, California) and positioned using a standard hydraulic microdrive (Fred Haer Corp, Brunswick, Maine). Neural signals were amplified, filtered (low-pass, $f_c=10$ kHz, Tucker-Davis Technologies (TDT), Gainesville, Florida) and digitized ($f_s=20$ kHz, also TDT). Digital data were then streamed to disk and written to CD-ROM. The polarity of the signal was reversed to give positive-going spike activity. Continuous extracellular traces were processed off-line to extract and classify spike events and calculate the LFP. Response criteria and stereotaxic location were consistent with recording from LIP. In one animal, locations were histologically verified; in the other, *in-vivo* MRI techniques were used for verification. Recordings were made from a 2×3 mm region in the chamber.

Spike sorting. Spike activity was extracted from the digitized recordings by an automated procedure that identified and sorted spike waveforms into clusters, each of which was presumed to arise from a single cell⁴³. In brief, events were clustered by fitting spike vectors to a mixture of Gaussian distributions. The fit was performed using the relaxation expectation-maximization algorithm, and the number of Gaussian components was determined by cascading model selection. Events were assigned to cells according to a maximum *a posteriori* rule. To ensure that the clustering was robustly determined, the segment of data used to fit the model was varied, and only models that were consistent for all segments were included in the database.

Data analysis. The collected data are spiking activity (a point process) and the LFP (a continuous valued time series), making this a hybrid data set. Spikes were binned at 1-ms time resolution to give spike trains. LFP time series were calculated from the extracellular recording on one tetrode channel by low-pass filtering the signal at 250 Hz. To reduce artifacts due to spikes leaking through the filter into the LFP, we estimated the LFP from the channel of the tetrode that had the smallest spike amplitude and subtracted a mean spike waveform from the extracellular recording at each spike time before filtering.

The baseline period extended 750–250 ms before the onset of the spatial cue. The memory period extended 450–950 ms following the offset of the spatial cue. The perisaccadic period extended 250 ms either side of the saccade onset. Baseline activity was estimated by pooling activity

from all successful trials. Memory and perisaccadic activity were estimated by pooling activity from successful trials according to saccade direction. The preferred direction of spike activity was the direction that elicited the maximum firing rate during the memory period. The preferred direction of the LFP was defined to be the same as the preferred direction of spiking. Activity from different locations was aligned to the preferred direction before estimating population averages.

We used spectral analysis to characterize temporal structure in the data. The relationship between spiking and LFP activity was measured using the spike-field coherency. The coherency is a complex quantity whose magnitude, the coherence, is normalized between 0 and 1. It measures the degree of predictability of one process using a linear function of the other⁴⁴. The coherency is given by the cross-spectrum of each process normalized by the spectrum of each process. This means changes in coherence do not result from changes in firing rate and allows the coherency to be meaningfully averaged across different pairs of time series. The coherency is used in preference to the cross-correlation function, as the cross-correlation function is typically normalized by the zero lag estimates of the auto-correlation function. This normalization is *ad hoc*, as it depends on the bin size used to estimate the zero lag auto-correlation. We used multitaper methods of spectral estimation to construct spectral estimators^{35–37} (Supplementary Methods).

The LFP spectrum was estimated on a 500 ms window with 5 Hz resolution using five Slepian data tapers. The spectrum of the spike counts, spike spectrum, was estimated on a 500-ms window with 10 Hz resolution using nine Slepian data tapers. When averaging across cells, we normalized the spike spectrum. For each 500-ms window, we divided the spectrum by the mean firing rate to give the 'rate-normalized spectrum'. Spike-field coherency was also estimated on a 500-ms window with 10 Hz resolution using nine Slepian data tapers.

All spectral quantities were smoothed with a lag window for visualization purposes³⁶. Significance levels were computed from estimates of the variance using a jackknife over tapers and trials^{45,46} (Supplementary Methods). Time-frequency representations of the activity were calculated on a 500-ms window that was stepped by 50 ms between estimates through the trial with the time index aligned to the center of the analysis window. Trials were aligned in time to the spatial cue.

The STA was estimated during baseline and memory before a saccade in the preferred direction for that cell. Segments of the raw voltage trace 200 ms long and centered on the spike were extracted. Spike waveforms were suppressed by subtracting a 2-ms mean spike shape waveform. The traces were then averaged to give the STA potential. The 95% confidence intervals were calculated by estimating the standard error of the mean and were used to test the hypothesis that the STA potential was zero.

Spike rate and LFP activity were decoded on a single-trial basis using linear discriminant analysis⁴⁷. The direction, either preferred or anti-preferred, was decoded from spike rate and LFP spectrum in the memory period. The state, either plan or execution, was decoded from the spike rate and LFP spectrum in the memory period for the plan state or perisaccadic period for the execution state. Only trials to the preferred and anti-preferred directions were included in the direction decode, giving 20–30 trials for analysis. Trials to all directions were included in the state decode, giving 80–120 trials for analysis. For each decode, the LFP spectrum was estimated with 10-Hz resolution using nine Slepian data tapers. Increasing the bandwidth of the spectral estimate increased its statistical stability, which was useful when analyzing the data on a single-trial basis. The LFP spectrum was log transformed and sampled at the single frequency with greatest discriminability across trials. This frequency was different for each site. Trials used to train the discriminant function were not used in the test set.

Note: Supplementary information is available on the Nature Neuroscience website.

Acknowledgments

This work was supported by the DARPA grant MDA972-00-1-0029, NIH grants EY05522-21 and MH62528-01, ONR grant N00014-94-0412, the Keck Foundation, the McKnight Foundation, the Sloan-Swartz Foundation, the Center for Neuromorphic Systems Engineering at Caltech and the Workshop for the Analysis of Neural Data (<http://www.vis.caltech.edu/~WAND>).

Competing interests statement

The authors declare that they have no competing financial interests.

RECEIVED 25 JANUARY; ACCEPTED 26 JUNE 2002

- Fuster, J. *Memory in Cerebral Cortex: An Empirical Approach to Neural Networks in the Human and Nonhuman Brain* (MIT Press, Cambridge, Massachusetts, 1995).
- Bruce, C. J. & Goldberg, M. E. Primate frontal eye fields: I. Single neurons discharging before saccades. *J. Neurophysiol.* 53, 603–635 (1985).
- Gnadt, J. W. & Andersen, R. A. Memory related motor planning activity in posterior parietal cortex of macaque. *Exp. Brain Res.* 70, 216–220 (1988).
- Koch, K. W. & Fuster, J. M. Unit activity in monkey parietal cortex related to haptic perception and temporary memory. *Exp. Brain Res.* 76, 292–306 (1989).
- Miller, E. K., Erickson, C. A. & Desimone, R. Neural mechanisms of visual working memory in prefrontal cortex of the macaque. *J. Neurosci.* 16, 5154–5167 (1996).
- Andersen, R. A. Encoding of intention and spatial location in the posterior parietal cortex. *Cereb. Cortex* 5, 457–469 (1995).
- Funahashi, S., Bruce, C. J. & Goldman-Rakic, P. S. Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. *J. Neurophysiol.* 61, 331–349 (1989).
- Mazzoni, P., Bracewell, R. M., Barash, S. & Andersen, R. A. Motor intention activity in the macaque's lateral intraparietal area. I. Dissociation of motor plan from sensory memory. *J. Neurophysiol.* 76, 1439–1456 (1996).
- Snyder, L. H., Batista, A. P. & Andersen, R. A. Coding of intention in the posterior parietal cortex. *Nature* 386, 167–170 (1997).
- Goldman-Rakic, P. S. Cellular basis of working memory. *Neuron* 14, 477–485 (1995).
- Seung, H. S. How the brain keeps the eyes still. *Proc. Natl. Acad. Sci. USA* 93, 13339–13344 (1996).
- Wang, X.-J. Synaptic basis of cortical persistent activity: the importance of NMDA receptors to working memory. *J. Neurosci.* 19, 9587–9603 (1999).
- Hebb, D. *Organization of Behavior* (Wiley, New York, 1949).
- Amit, D. J. The hebbian paradigm reintegrated: local reverberations as internal representation. *Behav. Brain Sci.* 18, 617–626 (1995).
- Eckhorn, R., Frien, A., Bauer, R., Woelbern, T. & Kehr, H. High-frequency (60–90 Hz) oscillations in primary visual cortex of awake monkey. *Neuroreport* 4, 243–246 (1993).
- Gray, C. M. & Singer, W. Stimulus-specific neuronal oscillations in orientation columns of cat visual cortex. *Proc. Natl. Acad. Sci. USA* 86, 1698–1702 (1989).
- Gray, C. M., König, P., Engel, A. K. & Singer, W. Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Nature* 338, 334–337 (1989).
- Kreiter, A. K. & Singer, W. Stimulus-dependent synchronization of neuronal responses in the visual cortex of the awake macaque monkey. *J. Neurosci.* 16, 2381–2396 (1996).
- Fries, P., Reynolds, J. H., Rorie, A. E. & Desimone, R. Modulation of oscillatory neuronal synchronization by selective visual attention. *Science* 291, 1560–1563 (2001).
- Steinmetz, P. N. *et al.* Attention modulates synchronized neuronal firing in primate somatosensory cortex. *Nature* 404, 187–190 (2000).
- Rougeul, A., Bouyer, J. J., Dedet, L. & Debray, O. Fast somato-parietal rhythms during combined focal attention and immobility in baboon and squirrel monkey. *Electro. Clin. Neurophysiol.* 46, 310–319 (1979).
- Bressler, S. L., Coppola, R. & Nakamura, R. Episodic multiregional cortical coherence at multiple frequencies during visual task-performance. *Nature* 366, 153–156 (1993).
- Murthy, V. N. & Fetz, E. E. Oscillatory activity in sensorimotor cortex of awake monkeys: synchronization of local field potentials and relation to behavior. *J. Neurophysiol.* 76, 3949–3967 (1996).
- Roelfsema, P. R., Engel, A. K., König, P. & Singer, W. Visuomotor integration is associated with zero time-lag synchronization among cortical areas. *Nature* 385, 157–161 (1997).
- Donoghue, J. P., Sanes, J. N., Hatsopoulos, N. G. & Gaal, G. Neural discharge and local field potential oscillations in primate motor cortex during voluntary movements. *J. Neurophysiol.* 79, 159–173 (1998).
- Lebedev, M. A. & Wise, S. P. Oscillations in the premotor cortex: single-unit activity from awake, behaving monkeys. *Exp. Brain Res.* 130, 195–215 (2000).
- Nakamura, K., Mikami, A. & Kubota, K. Oscillatory neuronal-activity related to visual short-term memory in monkey temporal pole. *Neuroreport* 3, 117–120 (1992).
- Tallon-Baudry, C., Bertrand, O., Peronnet, F. & Pernier, J. Induced gamma-band activity during the delay of a visual short-term memory task in humans. *J. Neurosci.* 18, 4244–4254 (1998).
- Chaffee, M. V. & Goldman-Rakic, P. S. Matching patterns of activity in primate prefrontal area 8a and parietal area 7ip during a spatial working memory task. *J. Neurophysiol.* 79, 2919–2940 (1998).
- Mitzdorf, U. Current source-density method and application in cat cerebral cortex: investigation of evoked potentials and EEG phenomena. *Physiol. Rev.* 65, 37–100 (1985).
- Engel, A. K., König, P., Gray, C. M. & Singer, W. Stimulus-dependent neuronal oscillations in cat visual-cortex - intercolumnar interaction as determined by cross-correlation analysis. *Eur. J. Neurosci.* 2, 588–606 (1990).
- Singer, W. & Gray, C. M. Visual feature integration and the temporal correlation hypothesis. *Annu. Rev. Neurosci.* 18, 555–586 (1995).
- Pezaris, J. S., Sahani, M. & Andersen, R. A. Response-locked changes in auto- and cross-covariations in parietal cortex. *Neurocomputing* 26–27, 471–476 (1999).
- Jarvis, M. R. & Mitra, P. P. Sampling properties of the spectrum and coherency of sequences of action potentials. *Neural Comput.* 13, 717–749 (2001).
- Thomson, D. J. Spectrum estimation and harmonic analysis. *Proc. IEEE* 70, 1055–1096 (1982).
- Percival, D. B. & Walden, A. T. *Spectral Analysis for Physical Applications* (Cambridge University Press, Cambridge, 1993).
- Mitra, P. P. & Pesaran, B. Analysis of dynamic brain imaging data. *Biophys. J.* 76, 691–708 (1999).
- Blatt, G., Andersen, R. A. & Stoner, G. Visual receptive field organization and cortico-cortical connections of area LIP in the macaque. *J. Comp. Neurol.* 299, 421–445 (1990).
- Wessberg, J. *et al.* Real-time prediction of hand trajectory by ensembles of cortical neurons in primates. *Nature* 408, 361–365 (2000).
- Colby, C. L. & Goldberg, M. E. Space and attention in parietal cortex. *Annu. Rev. Neurosci.* 22, 319–349 (1999).
- Wehr, M. & Laurent, G. Odour encoding by temporal sequences of firing in oscillating neural assemblies. *Nature* 384, 162–166 (1996).
- Pezaris, J. S., Sahani, M. & Andersen, R. A. in *Computational Neuroscience: Trends in Research* (ed. Bower, J. M.) 937–942 (Plenum, New York, 1996).
- Sahani, M., Pezaris, J. S. & Andersen, R. A. in *Advances in Neural Information Processing Systems 10* (eds Jordan, M. I., Kearns, M. J. & Solla, S. A.) 222–228 (MIT Press, Cambridge, Massachusetts, 1998).
- Brillinger, D. R. *Time Series* (Holt, Rinehart and Winston, New York, 1974).
- Efron, B. & Tibshirani, R. J. *An Introduction to the Bootstrap* (Chapman and Hall, London, 1993).
- Thomson, D. J. & Chave, A. D. *Advances in Spectrum Analysis and Array Processing* (Prentice Hall, Englewood Cliffs, New Jersey, 1991).
- Ripley, B. D. *Pattern Recognition and Neural Networks* (Cambridge University Press, Cambridge, 1996).